

Cisco ASR 5000 Series Product Overview

Release 12.0

Last Updated January 31, 2013

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-24419-05

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco ASR 5000 Series Product Overview

© 2013 Cisco Systems, Inc. All rights reserved.

CONTENTS

About this Guide xxxi
Conventions Used xxxii
Contacting Customer Support xxxiv
Additional Information xxxv
The Cisco® ASR 5000 Platform 37
Characteristics of the System 38
Features and Benefits 40
Product, Service and Feature Licenses 45
Supported Product/License Quick Reference 46
Session Use and Feature Use Licenses 50
Session Use Licenses 50
Feature Use Licenses 51
Default Licenses 53
ASR 5000 Hardware Platform Overview 55
The ASR 5000 Platform 56
Chassis Configurations 57
Chassis Description 61
Slot Numbering 61
Rear Slot Numbering for Half-Height Line Cards 62
Rear Slot Numbering with Full-height Line Cards 62
Mounting Options 62
Midplane Architecture 63
320 Gbps Switch Fabric 64
32 Gbps Control Bus 64
System Management Bus 64
280 Gbps Redundancy Bus 64
OC-48 TDM Bus 65
SPIO Cross-Connect Bus 65
Power Filter Units 67
Fan Tray Assemblies 68
Lower Fan Tray 68
Air Filter Assembly 68
Upper Fan Tray 69
Chassis Airflow 69
Application Cards 70
System Management Card (SMC) 70
SMC RAID Support 72
Packet Processing Cards: PSC, PSCA, PSC2, PSC3 and PPC 73
Packet Services Card (PSC) 75
Packet Services Card Type A (PSCA) 75
Packet Services Card 2 (PSC2) 75
Packet Services Card 3 (PSC3) 75
Packet Processor Card (PPC) Description 76

Line Cards 77
Switch Processor I/O (SPIO) Card 77
Management LAN Interfaces 78
Console Port 78
BITS Timing 79
Central Office Alarm Interface 79
Redundancy Crossbar Card (RCC) 79
Fast Ethernet Line Card (FELC or FLC2) 80
Gigabit Ethernet Line Card (GELC or GLC2) 81
Quad Gigabit Ethernet Line Card (QGLC) 83
10 Gigabit Ethernet Line Card (XGLC) 85
Optical Line Card (OLC2) 88
Channelized Line Card (CLC2) 90
Card Interlock Switch 94
Card Identifiers 95
Software Architecture 97
Understanding the Distributed Software Architecture 99
Software Tasks 99
Subsystems 99
Redundancy and Availability Features 103
Service Availability Features 104
Hardware Redundancy Features 104
Hardware Redundancy Configuration 104
Software Assurance Features 106
Session Recovery Feature 106
Interchassis Session Recovery 107
Mean Time Between Failure 108
System Availability 110
Spare Component Recommendations 111
Management System Overview 113
Out-of-Band Management 115
Command Line Interface 116
CLI Overview 116
Web Element Manager Application 118
Application Detection and Control Overview 119
ADC Overview 120
Platform Requirements 121
License Requirements 121
P2P Voice Call Duration 121
Random Drop Charging Action 121
How ADC Works 122
Advantages of P2P Processing Before DPI 122
ADC Session Recovery 123
Recovery from Task Failure 123
Recovery from CPU or PSC/PSC2 Failure 123
Limitations 123
BitTorrent 123
eDonkey 124
FastTrack 124
Gadu-Gadu 124
Gnutella / Morpheus 124
Jabber 124

MSN 124
Skype 124
Winny 125
Yahoo 125
Other Limitations 125
ASN Gateway Overview 127
ASN Mobility Management 128
EAP User Authentication 130
ASN Gateway and AAA 131
Profile Management 132
Inter-ASN Handovers 133
Supported Features 134
Simple IPv4 Support 134
DHCP Proxy Server 134
DHCP Relay Support 135
ASN Gateway Micro-Mobility 135
Uncontrolled Handovers 135
Controlled Handovers 135
WiMAX R4 Inter-ASN Mobility Management 136
WiMAX R3 CSN Anchored Mobility Management 136
Proxy Mobile IPv4 (PMIPv4) 137
Client Mobile IPv4 (CMIPv4) 137
Authenticator 137
EAP Authentication Methods 137
Supported RADIUS Methods 138
Supported Diameter Methods 138
WiMAX Prepaid Accounting 138
Volume and Duration-based Prepaid Accounting 139
Supported Enhanced Features 139
Lawful Intercept 139
Intelligent Traffic Control 140
Hotlining/Dynamic RADIUS Attributes 140
Multi-flow QoS 141
802.1P QoS Support 141
ASN Gateway Intra-Chassis Session Recovery 142
Robust Header Compression (ROHC) 142
Supported Inline Services 142
Enhanced Charging Service 142
Multi-host Support 142
How it Works 143
ASN Gateway in a WiMAX Network 145
Access Service Network (ASN) 145
Connectivity Service Network (CSN) 147
WiMAX Reference Points and Interfaces 148
Message Relay in ASN 148
ASN Gateway Architecture and Deployment Profiles 149
WiMAX Network Deployment Configurations 151
Standalone ASN Gateway/FA and HA Deployments 151
Co-Located Deployments 151
ASN Call Procedure Flows 153
Functional Components for Handover 153
Anchor ASN Gateway 153
Anchor Session 153
Non-Anchor ASN Gateway 154

Non-Anchor Session 154
Initial Network Entry and Data Path Establishment without Authentication 155
Initial Network Entry and Data Path Establishment with Authentication (Single EAP) 157
Unexpected Network Re-entry 159
MS Triggered Network Exit 160
Network Triggered Network Exit 161
Intra-ASN Gateway Handover 163
Intra-anchor ASN Gateway Uncontrolled Handover 163
Intra-anchor ASN Gateway Controlled Handover 165
Inter-ASN Gateway Handover 171
ASN Gateway Function for Handovers 172
Controlled Anchor ASN Gateway to Non-Anchor ASN Gateway Handover 173
Uncontrolled Anchor ASN Gateway to Non-Anchor ASN Gateway Handover 177
RADIUS-based Prepaid Accounting for WiMax 179
Obtaining More Quota after the Quota is Reached 179
Applying HTTP Redirection Rule when Quota is Reached 181
Applying HTTP Redirection Rule When CoA is Received 183
Terminating the Call when Quota is Reached 185
DHCP Relay Support for ASNGW 187
DHCP Keys 187
Key Generation 187
Key Distribution 187
DHCP Relay in ASNGW for PMIPv4 Calls 188
DHCP Relay Requirements for Simple IP Calls 189
DHCP Relay Requirements for PMIPv4 Calls 189
RADIUS Based Procedures for Prepaid Accounting 191
Flow- based Prepaid Accounting 191
Configuring Rating-Group-ID 192
Prepaid Tariff Switching 192
Hotlining with Flow-based Prepaid Accounting 192
RADIUS-based Procedures for WiMAX Hotlining 193
Types of Hotlining 193
Hotlining for Prepaid Accounting Sessions 193
Active Session Hot-lining Call Flows 193
Active Session Hotlining for Prepaid Users 195
Hotlining During Initial Network Entry 195
Tariff Switching for Prepaid Accounting 195
CSN Procedure Flows 197
PMIP4 Connection Setup and Call Flow with DHCP Proxy 197
PMIP4 Session Release 198
WiMAX Deployment with Legacy Core Networks 200
ASN Gateway Interoperability with 3GPP Overlay 200
ASN Gateway Interoperability with 3GPP2 Overlay 200
Session Continuity Support for 3GPP2 and WiMAX Handovers 201
NSP-ID and NAP-ID Functionality 202
Manual Mode 202
Automatic Mode 202
ASN GW and NAP-ID/NSP-ID Process 203
Data Tunnel Endpoint Support 204
ASNGW with a Different Tunnel Endpoint 204
No Handoff (INE) 204
Inter-ASNGW Handoff 204
Intra-ASNGW Handoff 205
Supported Standards 206

WiMAX/IEEE References 206
IEEE Standards 206
IETF References 206
Object Management Group (OMG) Standards 207
ASN Paging Controller and Location Registry Overview 209
Introduction 210
Description of PC/LR Support 211
Licenses 211
Paging and Location Update Procedures 211
Paging Controller (PC) 211
Paging Agent (PA) 212
Paging Group (PG) 212
Location Register (LR) 212
Location Update Procedure 212
Location Update with Paging Controller Relocation 214
Paging Operation 216
MS Initiated Idle Mode Entry 218
MS Initiated Idle Mode Exit 221
Supported Platforms and Software 224
CDMA2000 Wireless Data Services 225
Product Description 226
Product Specifications 227
Hardware Requirements 227
Platforms 227
ASR 5x00 Series Platform System Hardware Components 227
Features and Functionality—Base Software 228
Gx and Gy Support 228
RADIUS Support 229
Description 230
Access Control List Support 230
IP Policy Forwarding 231
Description 231
AAA Server Groups 232
Description 232
Overlapping IP Address Pool Support 232
Routing Protocol Support 232
Description 232
Management System Overview 234
Description 234
Bulk Statistics Support 234
Description 235
Threshold Crossing Alerts (TCA) Support 235
Description 236
IP Header Compression - Van Jacobson 236
Description 236
DSCP Marking 237
Features and Functionality - Optional Enhanced Software Features 238
Session Recovery Support 238
Description 238
IPv6 Support 239
Description 239
L2TP LAC Support 240
Description 240

L2TP LNS Support....240
Description....240
Proxy Mobile IP....241
Description....241
IP Security (IPSec)....241
Description....241
Traffic Policing and Rate Limiting....242
Description....242
Intelligent Traffic Control....243
Dynamic RADIUS Extensions (Change of Authorization)....243
Description....243
Web Element Management System....244
Description....244
Features and Functionality - Inline Service Support....245
Content Filtering....245
Integrated Adult Content Filter....245
ICAP Interface....245
Network Address Translation (NAT)....246
Peer-to-Peer Detection....246
Personal Stateful Firewall....247
Traffic Performance Optimization (TPO)....247
Features and Functionality - External Application Support....249
Mobility Unified Reporting....249
CDMA2000 Data Network Deployment Configurations....250
Standalone PDSN/FA and HA Deployments....250
Interface Descriptions....250
Co-Located Deployments....251
Understanding Simple IP and Mobile IP....253
Simple IP....253
How Simple IP Works....253
Mobile IP....256
Mobile IP Tunneling Methods....256
How Mobile IP Works....258
Proxy Mobile IP....261
How Proxy Mobile IP Works....262
Supported Standards....267
Requests for Comments (RFCs)....267
TIA and Other Standards....270
Telecommunications Industry Association (TIA) Standards....270
Object Management Group (OMG) Standards....270
3GPP2 Standards....270
IEEE Standards....271

Content Filtering Support Overview....273

Introduction....274
Platform Requirements....275
Licenses Requirements....275
URL Blacklisting Support....276
URL Blacklisting Solution Components....277
Web Element Manager (WEM)....277
How URL Blacklisting Works....278
Blacklist Updates....278
URL Blacklisting Action....279
Category-based Content Filtering Support....280
Benefits of Category-based Content Filtering....280

Static-and-Dynamic Content Filtering 281
TCP Proxy Functionality 281
ECS and Content Filtering Application 282
Components of Category-based Content Filtering Solution 283
Category-based Content Filtering Subsystem 284
Static Rating Categorization Database (SRDB) 284
Dynamic Static Rating Categorization Database 285
Rater Package Model Files 285
Content Rating Rules Update Server 286
Master Content Rating Database Server (MCRDBS) 286
External Storage System 286
RADIUS Server and Policy Manager 287
Web Element Manager (WEM) 287
Mobility Unified Reporting System 288
How Category-based Content Filtering Works 288
How URL Blacklisting and Category-based Content Filtering Work Concurrently 292
Content Filtering Server Group Support 293
External Storage System 295
Bulk Statistics Support 296
Minimum System Requirements and Recommendations 297
MCRDBS System Requirements 297
Hardware Requirements 297
Additional Requirements on Chassis 298

Enhanced Charging Service Overview 299

Introduction 300
Platform Requirements 300
License Requirements 300
Basic Features and Functionality 301
Shallow Packet Inspection 301
Deep Packet Inspection 301
Charging Subsystem 301
Traffic Analyzers 301
How ECS Works 302
ECS Deployment and Architecture 311
Enhanced Features and Functionality 313
Session Control in ECS 313
Time and Flow-based Bearer Charging in ECS 313
Fair Usage 315
Content Filtering Support 315
Content Filtering Server Group Support 315
In-line Content Filtering Support 316
IP Readdressing 316
Next-hop Address Configuration 316
Post Processing 317
How the Post-processing Feature Works 317
Time-of-Day Activation/Deactivation of Rules 318
How the Time-of-Day Activation/Deactivation of Rules Feature Works 318
URL Filtering 318
TCP Proxy 319
TCP Proxy Behavior and Limitations 320
X-Header Insertion and Encryption 323
License Requirements 324
X-Header Insertion 324
X-Header Encryption 324

Limitations to the Header Insertion and Encryption Features....325
Accounting and Charging Interfaces....327
GTPP Accounting....327
RADIUS Accounting and Credit Control....327
Diameter Accounting and Credit Control....328
Gx Interface Support....328
Gy Interface Support....328
Event Detail Records (EDRs)....329
Usage Detail Records (UDRs)....330
Charging Methods and Interfaces....331
Prepaid Credit Control....331
Postpaid....331
Prepaid Billing in ECS....332
How ECS Prepaid Billing Works....332
Credit Control Application (CCA) in ECS....333
How Credit Control Application (CCA) Works for Prepaid Billing....333
Postpaid Billing in ECS....334
How ECS Postpaid Billing Works....335
External Storage System....339
System Resource Allocation....340
Redundancy Support in ECS....341
Intra-chassis Session Recovery Interoperability....341
Recovery from Task Failure....341
Recovery from CPU or Packet Processing Card Failure....341
Inter-chassis Session Recovery Interoperability....341
Inter-chassis Session Recovery Architecture....342
Impact on xDR File Naming....342
Impact on xDR File Content....343
External Storage System Overview....345
ESS Overview....346
ESS Features and Functions....347
System Requirements....349
ASR 5x00 System Requirements....349
ESS System Requirements....349
ESS System Recommendations for Cluster Deployment....349
Femto Network Gateway Overview....351
Product Description....352
Platform Requirements....352
Licenses....352
Summary of FNG Features and Functions....353
Network Deployment(s) and Interfaces....354
Network Elements....354
Femtocell Access Point....354
Femtocell Management System....355
Femto Network Gateway....355
Femtocell AAA Server....355
IMS Core Network Elements....355
PDSN/HA....356
Basic Operation....356
Network Interfaces....356
Features and Functionality....358
FNG Service....358
IKEv2 and IP Security (IPSec) Encryption....359

X.509 Certificate-based Peer Authentication 359
A12 Aggregation 360
RADIUS Support 360
AAA Server Group Selection 360
FAP ID-based Duplicate Session Detection 361
Tunnel Cleanup on FAP Reboot 361
Child SA Rekey Support 361
Multiple Child SAs 361
DoS Protection Cookie Challenge 362
IKEv2 Keep-Alive Messages (Dead Peer Detection) 362
DSCP Marking 362
Custom DNS Handling 363
Session Recovery Support 363
Congestion Control 363
Bulk Statistics 364
Threshold Crossing Alerts 365
How the FNG Works 366
IPSec Tunnel Establishment 366
IPSec Tunnel Establishment with EAP-AKA Authentication 367
X.509 Certificate-based Peer Authentication 369
Supported Standards 372
3GPP2 References 372
IETF References 372

GGSN Support in GPRS/UMTS Wireless Data Services 375

Product Description 376
Product Specification 377
Licenses 377
Platform Requirements 377
Operating System Requirements 377
Network Deployment and Interfaces 378
GGSN in the GPRS/UMTS Data Network 378
Supported Interfaces 380
Features and Functionality - Base Software 382
16,000 SGSN Support 383
AAA Server Groups 383
Access Control List Support 383
ANSI T1.276 Compliance 384
APN Support 384
Bulk Statistics Support 385
Direct Tunnel Support 386
DHCP Support 387
DSCP Marking 388
Framed-Route Attribute Support 388
Generic Corporate APN 388
GnGp Handoff Support 388
GTPP Support 389
Host Route Advertisement 390
IP Policy Forwarding 391
IP Header Compression - Van Jacobson 391
IPv6 Support 392
Management System Overview 393
MPLS Forwarding with LDP 394
Overlapping IP Address Pool Support 395
PDP Context Support 395

Per APN Configuration to Swap out Gn to Gi APN in CDRs....396
Port Insensitive Rule for Enhanced Charging Service....396
Quality of Service Support....397
RADIUS Support....397
RADIUS VLAN Support....398
Routing Protocol Support....399
Subscriber Session Trace Support....400
Support of Charging Characteristics Provided by AAA Server....402
Support of all GGSN generated causes for partial G-CDR closure....402
Threshold Crossing Alerts (TCA) Support....402
Features and Functionality - Optional Enhanced Feature Software....404
Common Gateway Access Support....404
Dynamic RADIUS Extensions (Change of Authorization)....405
GRE Protocol Interface Support....405
Gx Interface Support....407
Inter-Chassis Session Recovery....408
IP Security (IPSec)....409
L2TP LAC Support....411
L2TP LNS Support....412
Lawful Intercept....412
Mobile IP Home and Foreign Agents....413
Mobile IP NAT Traversal....414
Multimedia Broadcast Multicast Services Support....414
Overcharging Protection on Loss of Coverage....414
Proxy Mobile IP....415
Session Persistence....415
Session Recovery Support....416
Traffic Policing and Rate Limiting....417
Web Element Management System....418
How GGSN Works....420
PDP Context Processing....420
Dynamic IP Address Assignment....421
Subscriber Session Call Flows....422
Transparent Session IP Call Flow....422
Non-Transparent IP Session Call Flow....424
Network-Initiated Session Call Flow....427
PPP Direct Access Call Flow....428
Virtual Dialup Access Call Flow....430
Corporate IP VPN Connectivity Call Flow....432
Mobile IP Call Flow....434
Proxy Mobile IP Call Flows....437
IPv6 Stateless Address Auto Configuration Flows....440
Supported Standards....442
3GPP References....442
IETF References....443
Object Management Group (OMG) Standards....446
HA Overview....447
Product Specifications....448
Hardware Requirements....448
Platforms....448
Components....448
Operating System Requirements....448
MPLS Forwarding with LDP....449
Supported Standards....450

Requests for Comments (RFCs) 450
Network Deployment Configurations 454
Standalone PDSN/FA and HA Deployments 454
Interface Descriptions 454
Co-Located Deployments 455
Mobile IP Tunneling Methods 456
How Mobile IP Works 459
Understanding Mobile IP 462
Session Continuity Support for 3GPP2 and WiMAX Handoffs 462
HNB Gateway in Wireless Network 463
Product Description 464
HNB Access Network Elements 465
Home NodeB 465
Security Gateway (SeGW) 466
HNB Gateway (HNB-GW) 466
HNB Management System (HMS) 466
Licenses 466
Platform Requirements 467
Network Deployment and Interfaces 468
HNB Gateway in 3G UMTS Network 468
Supported Logical Interfaces 468
Features and Functionality - Base Software 471
AAA Server Group Support 471
AAL2 Establish and Release Support 472
Access Control List Support 472
ANSI T1.276 Compliance 473
ATM VC Management Support 473
Congestion Control and Management Support 473
Emergency Call Handling 474
GTP-U Tunnels Management Support 475
HNB-UE Access Control 475
HNB Management Function 476
Intra-Domain Multiple CN Support Through Iu-Flex 476
Iu Signalling Link Management Support 477
IuH User-Plane Transport Bearer Handling Support 477
Multiple MSC Selection without Iu-Flex 477
Network Access Control Functions through SeGW 478
Authentication and Key Agreement (AKA) 478
3GPP AAA Server Support 478
X.509 Certificate-based Authentication Support 478
Open Access Mode Support 479
QoS Management with DSCP Marking 480
RADIUS Support 480
UE Management Function for Pre-Rel-8 UEs 481
System Management Features 481
Management System Overview 482
Bulk Statistics Support 483
Threshold Crossing Alerts (TCA) Support 485
ANSI T1.276 Compliance 486
Features and Functionality - Optional Enhanced Feature Software 487
Dynamic RADIUS Extensions (Change of Authorization) 487
IP Security (IPSec) 488
Session Recovery 488
Web Element Management System 489

How HNB-GW Works....490
HNB Provisioning and Registration Procedure....490
UE Registration Procedure....492
UE Registration Procedure of Non-CSG UEs or Non-CSG HNBs....492
Iu Connection Procedures....494
Iu Connection Establishment Procedure....494
Network Initiated Iu Connection Release Procedure....496
Paging and Serving RNS Relocation Procedures....498
Paging Procedure....498
SRNS Relocation Procedure....498
RANAP Reset Procedures....499
HNB Initiated RANAP Reset Procedure....499
CN Initiated RANAP Reset Procedure....499
HNB-GW Initiated RANAP Reset Procedure....499
Supported Standards....501
3GPP References....501
IETF References....502
ITU-T Recommendations....504
Object Management Group (OMG) Standards....505
HRPD Serving Gateway Overview....507
Product Description....508
Basic Features....509
Authentication....509
IP Address Allocation....510
Quality of Service....510
AAA, Policy and Charging....510
Platform Requirements....511
Licenses....511
Network Deployment(s)....512
HRPD Serving Gateway in an eHRPD Network....512
Supported Logical Network Interfaces (Reference Points)....513
Features and Functionality - Base Software....517
A10/A11....517
AAA Server Groups....518
ANSI T1.276 Compliance....518
Bulk Statistics Support....518
Congestion Control....519
DSCP Marking....520
Dynamic Policy and Charging: Gxa Reference Interface....520
EAP Authentication (STa)....521
IP Access Control Lists....521
Management System....522
Mobile IP Registration Revocation....523
Multiple PDN Support....524
Non-Optimized Inter-HSGW Session Handover....524
P-GW Selection (Discovery)....524
PPP VSNCP....525
Proxy Mobile IPv6 (S2a)....525
Rf Diameter Accounting....526
Threshold Crossing Alerts (TCA) Support....526
UE Initiated Dedicated Bearer Resource Establishment....527
Features and Functionality - External Application Support....528
Web Element Management System....528
Features and Functionality - Optional Enhanced Feature Software....530

Intelligent Traffic Control 530
IP Header Compression (RoHCv1 for IPv4/IPv6) 530
IP Security (IPSec) 531
Layer 2 Traffic Management (VLANs) 531
Session Recovery Support 532
Traffic Policing and Shaping 532
Traffic Policing 532
Traffic Shaping 533
Call/Session Procedure Flows 534
Initial Attach with IPv6/IPv4 Access 534
PMIPv6 Lifetime Extension without Handover 536
PDN Connection Release Initiated by UE 537
PDN Connection Release Initiated by HSGW 538
PDN Connection Release Initiated by P-GW 540
Supported Standards 542
Release 8 3GPP References 542
3GPP2 References 542
IETF References 543
Object Management Group (OMG) Standards 543

InTracer Overview 545

Introduction 546
InTracer Architecture 547
InTracer Components 548
InTracer System Requirements 554
Supported Features 556

IP Services Gateway Overview 557

Introduction 558
Platform Requirements 558
License Requirements 558
How it Works 559
RADIUS Server Mode 559
RADIUS Proxy 559
RADIUS Snoop Mode 560
In-line Services 561
Application Detection and Control 561
Content Filtering 561
Enhanced Charging Service 561
Enhanced Feature Support 562
Dynamic RADIUS Extensions (Change of Authorization) 562
Gx Interface Support 562
Gy Interface Support 563
Content Service Steering 563
Multiple IPSG Services 564
Session Recovery 564

Mobility Unified Reporting System Overview 565

Introduction 566
Report Types 567
Exporting Reports to Other File Formats 573
License Requirements 574
MUR Architecture 575
Distributed Architecture of MUR 577
How RDP works with MUR 578

Scalable MUR....580
Basic Scalability Model....580
Scalability Setup for New Deployments of MUR....582
MUR Features....583
MUR Deployment....584
MUR System Requirements....585
Server Recommendations for Use in Solaris Environment....585
Server Recommendations for Use in RHEL Environment....586
Hardware Requirements for Scalable Model of MUR....587
Software Requirements for Scalable Model of MUR....587
MUR Ports....588
Firewall Settings....588
Using Apache Port....589
Using Apache in Solaris....589
Using Apache in RHEL....589
Mobile Video Gateway Overview....591
Product Description....592
Platform Requirements....592
Licenses....592
Summary of Mobile Video Gateway Features and Functions....592
Network Deployments and Interfaces....595
The Mobile Video Gateway in an E-UTRAN/EPC Network....595
The Mobile Video Gateway in a GPRS/UMTS Network....596
The Mobile Video Gateway in a CDMA2000 Network....597
Mobile Video Gateway Logical Network Interfaces....597
Features and Functionality....599
Deep Packet Inspection....599
Transparent Video Re-addressing....599
HTTP X-Header Use in Transparent Video Re-addressing....600
Mobile Video Gateway to the CAE....600
CAE to the OS....601
CAE Load Balancing....601
CAE Load Balancer Function....602
CAE Health-Check Monitoring Function....603
Video Optimization Policy Control....603
Functional Overview....604
Video Optimization Policy Control Call Flows....605
Video White-listing....608
Video Pacing....608
Video Pacing Operation....608
Video Pacing Functions....610
Video Pacing Call Flows....611
Interactions with Related Functions....612
Supported Video Container File Formats....613
TCP Link Monitoring....613
TCP Link Monitoring System Flow....613
Functional Overview....614
Dynamically-enabled TCP Proxy....615
Traffic Performance Optimization....615
N+1 Stateful Redundancy....615
Threshold Crossing Alerts....615
Mobile Video Statistics....616
Bulk Statistics for Mobile Video....617
How the Mobile Video Gateway Works....619

Mobile Video Gateway with the Content Adaptation Engine 619
DPI of HTTP GET Request Identifying a Non-Video Request (MVG with the CAE) 619
DPI of HTTP RESPONSE Identifying a Non-Video Request (MVG with the CAE) 620
DPI of HTTP GET Request Identifying a Video Request (MVG with the CAE) 621
DPI of HTTP RESPONSE Identifying a Video Request (MVG with the CAE) 623
Mobile Video Gateway without the Content Adaptation Engine 625
DPI of HTTP GET Request Identifying a Non-Video Request (MVG without the CAE) 625
DPI of HTTP RESPONSE Identifying a Non-Video Request (MVG without the CAE) 626
DPI of HTTP GET Request Identifying a Video Request (MVG without the CAE) 627
DPI of HTTP RESPONSE Identifying a Video Request (MVG without the CAE) 629

Mobility Management Entity Overview 631

Product Description 632
Platform Requirements 634
Licenses 634
Network Deployment and Interfaces 635
MME in the E-UTRAN/EPC Network 635
Supported Logical Network Interfaces (Reference Points) 637
Features and Functionality - Base Software 643
3GPP R8 Identity Support 644
ANSI T1.276 Compliance 645
APN Restriction Support 645
Authentication and Key Agreement (AKA) 645
Bulk Statistics Support 646
Congestion Control 646
EPS Bearer Context Support 647
EPS GTPv2 Support on S11 Interface 647
HSS Support Over S6a Interface 648
Inter-MME Handover Support 649
Interworking Support 649
Interworking with SGSNs 649
Handover Support for S4-SGSNs 650
IPv6 Support 651
MME Interfaces Supporting IPv6 Transport 651
Load Balancing 652
Load Re-balancing 652
Management System Overview 652
MME Pooling 653
MME Selection 654
Mobile Equipment Identity Check 654
Mobility Restriction 654
Handover Restriction 654
Multiple PDN Support 655
NAS Protocol Support 655
EPS Mobility Management (EMM) 655
EPS Session Management (ESM) 655
NAS Signalling Security 655
Operator Policy Support 656
Overload Management in MME 656
Packet Data Network Gateway (P-GW) Selection 657
Radio Resource Management Functions 657
Reachability Management 657
SCTP Multi-homing Support 657
SCTP Multi-homing for S6a 658
SCTP Multi-homing for S1-MME 658

Serving Gateway Pooling Support....658
Serving Gateway Selection....658
Session and Quality of Service Management....658
Subscriber Level Session Trace....659
Threshold Crossing Alerts (TCA) Support....660
Tracking Area List Management....661
Features and Functionality - External Application Support....662
Web Element Management System....662
Features and Functionality - Licensed Enhanced Feature Software....664
Circuit Switched Fall Back (CSFB) and SMS over SGs Interface....664
IP Security (IPSec)....666
Lawful Intercept....667
Location Services....668
Optimized Paging Support....668
Session Recovery Support....668
User Location Information Reporting....669
How the MME Works....671
EPS Bearer Context Processing....671
Purge Procedure....671
Paging Procedure....671
Subscriber-initiated Initial Attach Procedure....672
Subscriber-initiated Detach Procedure....674
Service Request Procedures....675
UE-initiated Service Request Procedure....675
Network-initiated Service Request Procedure....677
Supported Standards....679
3GPP References....679
IETF References....679
Object Management Group (OMG) Standards....681

Network Address Translation Overview....683

NAT Overview....684
Platform Requirements....685
License Requirements....685
NAT Realms....685
NAT IP Pool Groups....687
NAT IP Address Allocation and Deallocation....687
NAT IP Address Allocation....688
NAT IP Address Deallocation....688
NAT Port-chunk Allocation and Deallocation....688
NAT Port-chunk Allocation....688
NAT Port-chunk Deallocation....689
NAT IP Address/Port Allocation Failure....689
TCP 2MSL Timer....689
Flow Mapping Timer....690
NAT Binding Records....690
NAT Binding Updates....691
CoA NAT Query....691
Firewall-and-NAT Policy....692
Disabling NAT Policy....693
Updating Firewall-and-NAT Policy in Mid-session....693
Target-based NAT Configuration....694
NAT Application Level Gateway....694
Supported NAT ALGs....695
H323 ALG Support....695

NAT Aware H323 Clients 696
EDRs and UDRs 696
EDRs 696
UDRs 696
Bulk Statistics 697
Alarms 697
Session Recovery and ICSR 698
Supported Standards 699
How NAT Works 700
Packet Data Interworking Function Overview 705
Product Description 706
Platform Requirements 706
Licenses 706
Interfaces 707
Sample Deployments 708
Mobile Station using Mobile IP with PDIF/FA 708
Overview 708
Mobile IP / Native Simple IP Call Minimum Requirements 709
Mobile IP Session Setup over IPSec 709
Simple IP and Simple IP Fallback 712
Simple IP Fallback Minimum Requirements 715
Features and Functionality - Base Software 716
PSC2 Support 716
Duplicate Session Detection 717
Unsupported Critical Payload Handling 717
Registration Revocation 718
CHILD SA Rekey Support 718
Denial of Service (DoS) Protection: 718
Cookie Challenge Statistics 719
MAC Address Validation 720
RADIUS Accounting 720
Special RADIUS Attribute Handling 720
Mobile IP and Proxy Mobile IP Attributes 721
IPv6 Support 721
IPv6 Neighbor Discovery 722
IPv6 Static Routing 722
Port-Switch-On-L3-Fail for IPv6 722
IKEv2 Keep-Alive (Dead Peer Detection (DPD)) 722
Congestion Control and Overload Disconnect 722
SCTP (Stream Control Transmission Protocol) Support 723
X.509 Digital Trusted Certificate Support 723
2048-bit Certificate Key Functionality 723
Collocation of Certificate with 1024-bit Key and 2048-bit Key 724
Distinguishing Between 1024-bit and 2048-bit Key Certificates 724
Custom DNS Handling 725
Gx+ and Gy Interface Support via HA 725
Features and Functionality - Licensed Enhanced Feature Support 728
PDIF Service 728
Multiple PDIF Services 729
Lawful Intercept 729
Diameter Authentication Failure Handling 730
Online Upgrade 730
The Active-Standby Upgrade Model 731
Operation Over a Common IPv4 Network 732

Operation Over a Common IPv6 Network....733
Other Devices....734
Session Recovery Support....735
IPSec/IKEv2....735
Simple IP Fallback....736
Simple IP....736
Proxy Mobile IP....737
Multiple Authentication in a Proxy Mobile IP Network....737
AAA Group Selection....738
RADIUS Authentication....738
First-Phase Authentication....739
Second-Phase Authentication....739
Termination....739
Session Recovery....740
Intelligent Packet Monitoring System (IPMS)....740
Multiple Traffic Selectors....741
Selective Diameter Profile Update Request Control....741
Support of SHA2 Algorithms....742
Supported Standards and RFCs....743
3GPP2 References....743
IETF References....743
Object Management Group (OMG) Standards....744
PDG/TTG Overview....745
Product Description....746
Platform Requirements....746
Licenses....746
Summary of PDG/TTG Features and Functions....746
Network Deployment(s) and Interfaces....748
The PDG in a GPRS/UMTS Data Network....748
The TTG in a GPRS/UMTS Data Network....748
PDG/TTG Logical Network Interfaces (Reference Points)....749
Features and Functionality....751
PDG Service....751
PDG Mode....752
TTG Mode....752
IKEv2 and IP Security (IPSec) Encryption....752
Multiple Digital Certificate Selection Based on APN....753
Subscriber Traffic Policing for IPSec Access....753
DSCP Marking for IPSec Access....754
WLAN Access Control....755
RADIUS and Diameter Support....756
Additional Command Option for APN Configuration for Diameter AAA....756
EAP Fast Re-authentication Support....757
Pseudonym NAI Support....757
Multiple APN Support for IPSec Access....758
Multiple Authentication Support....758
Fast Re-authentication and Pseudonym Re-authentication....758
Re-authorization and RADIUS CoA Handling....759
Lawful Intercept....759
IMS Emergency Call Handling....759
Session Recovery Support....759
Congestion Control....760
Bulk Statistics....760
Threshold Crossing Alerts....761

AAA Mediation Accounting and Offline Charging 762
Accounting Session Management 762
Triggers for WLAN CDRs Charging Information 762
Features Not Supported in This Release 764
How the PDG/TTG Works 765
PDG Connection Establishment 765
TTG Connection Establishment 768
Multiple Authentication Using EAP and PAP 772
Multiple Authentication Using EAP and CHAP 773
Multiple Authentication Using EAP and EAP 774
Multiple Authentication with Request from UE for Change of Second Phase Protocol 775
Supported Standards 776
3GPP References 776
IETF References 776

PDN Gateway Overview 779

Product Description 780
Platform Requirements 782
Licenses 782
Network Deployment(s) 783
PDN Gateway in the E-UTRAN/EPC Network 783
Supported Logical Network Interfaces (Reference Points) 784
PDN Gateway Supporting eHRPD to E-UTRAN/EPC Connectivity 788
Supported Logical Network Interfaces (Reference Points) 790
Features and Functionality - Base Software 795
3GPP R9 Volume Charging Over Gx 796
AAA Server Groups 796
ANSI T1.276 Compliance 796
APN Support 797
Bulk Statistics Support 797
Congestion Control 799
Default and Dedicated EPC Bearers 799
DHCP Support 800
Direct Tunnel Support 800
DSCP Marking 801
Dynamic Policy Charging Control (Gx Reference Interface) 801
Enhanced Charging Service (ECS) 802
Content Analysis Support 804
Content Service Steering 805
Support for Multiple Detail Record Types 805
Diameter Credit Control Application 806
Accept TCP Connections from DCCA Server 806
Gy Interface Support 806
Gn/Gp Handoff Support 807
IP Access Control Lists 808
IP Address Hold Timers 808
IPv6 Capabilities 808
Local Break-Out 809
Management System Overview 809
Mobile IP Registration Revocation 811
Multiple PDN Support 812
Non-Optimized e-HRPD to Native LTE (E-UTRAN) Mobility Handover 812
Online/Offline Charging 813
Online Charging 813
Offline Charging 814

Proxy Mobile IPv6 (S2a) 815
QoS Bearer Management 815
RADIUS Support 816
Source IP Address Validation 817
Subscriber Level Trace 817
Threshold Crossing Alerts (TCA) Support 818
Virtual APN Support 819
Features and Functionality - Inline Service Support 820
Content Filtering 820
Integrated Adult Content Filter 820
ICAP Interface 821
Header Enrichment: Header Insertion and Encryption 821
Mobile Video Gateway 822
Network Address Translation (NAT) 823
Peer-to-Peer Detection 823
Personal Stateful Firewall 824
Traffic Performance Optimization (TPO) 824
Features and Functionality - External Application Support 826
Web Element Management System 826
Features and Functionality - Optional Enhanced Feature Software 828
Dynamic RADIUS Extensions (Change of Authorization) 828
GRE Protocol Interface Support 829
Inter-Chassis Session Recovery 829
IP Security (IPSec) Encryption 830
L2TP LAC Support 831
Lawful Intercept 831
Layer 2 Traffic Management (VLANs) 831
Local Policy Decision Engine 832
NEMO Service Supported 832
Session Recovery Support 832
Traffic Policing and Shaping 833
Traffic Policing 833
Traffic Shaping 833
User Location Information Reporting 834
How the PDN Gateway Works 836
PMIPv6 PDN Gateway Call/Session Procedures in an eHRPD Network 836
Initial Attach with IPv6/IPv4 Access 836
PMIPv6 Lifetime Extension without Handover 838
PDN Connection Release Initiated by UE 839
PDN Connection Release Initiated by HSGW 840
PDN Connection Release Initiated by P-GW 842
GTP PDN Gateway Call/Session Procedures in an LTE-SAE Network 843
Subscriber-initiated Attach (initial) 843
Subscriber-initiated Detach 846
Supported Standards 848
Release 8 3GPP References 848
3GPP2 References 849
IETF References 849
Object Management Group (OMG) Standards 850

Personal Stateful Firewall Overview 851

Firewall Overview 852
Platform Requirements 852
License Requirements 852
Supported Features 853

Protection against Denial-of-Service Attacks 853
Types of Denial-of-Service Attacks 853
Protection against Port Scanning 855
Application-level Gateway Support 855
PPTP ALG Support 856
TFTP ALG Support 856
Stateful Packet Inspection and Filtering Support 857
Stateless Packet Inspection and Filtering Support 857
Host Pool, IMSI Pool, and Port Map Support 857
Host Pool Support 857
IMSI Pool Support 857
Port Map Support 858
Flow Recovery Support 858
SNMP Thresholding Support 858
Logging Support 859
How Personal Stateful Firewall Works 860
Disabling Firewall Policy 860
Mid-session Firewall Policy Update 861
How it Works 861
Understanding Rules with Stateful Inspection 865
Connection State and State Table in Personal Stateful Firewall 865
Transport and Network Protocols and States 866
Application-Level Traffic and States 867

Serving Gateway Overview 869

Product Description 870
Platform Requirements 872
Licenses 872
Network Deployment(s) 873
Serving Gateway in the E-UTRAN/EPC Network 873
Supported Logical Network Interfaces (Reference Points) 874
Features and Functionality - Base Software 879
ANSI T1.276 Compliance 879
Bulk Statistics Support 880
Circuit Switched Fall Back (CSFB) Support 880
Congestion Control 881
IP Access Control Lists 881
IPv6 Capabilities 882
Management System Overview 882
Multiple PDN Support 884
Online/Offline Charging 884
Online: Gy Reference Interface 884
Offline: Gz Reference Interface 885
Operator Policy Support 885
QoS Bearer Management 885
Subscriber Level Trace 886
Threshold Crossing Alerts (TCA) Support 887
Features and Functionality - External Application Support 889
Web Element Manager 889
Features and Functionality - Optional Enhanced Feature Software 890
Direct Tunnel 890
Inter-Chassis Session Recovery 891
IP Security (IPSec) Encryption 892
Lawful Intercept 893
Layer 2 Traffic Management (VLANs) 893

Session Recovery Support....893
How the Serving Gateway Works....895
GTP Serving Gateway Call/Session Procedures in an LTE-SAE Network....895
Subscriber-initiated Attach (initial)....895
Subscriber-initiated Detach....898
Supported Standards....900
3GPP References....900
Release 8 Supported Standards....900
3GPP2 References....901
IETF References....901
Object Management Group (OMG) Standards....902
Session Control Manager Overview....903
Product Description....904
IMS Architecture....904
Proxy-CSCF....907
Interrogating-CSCF....908
Serving-CSCF....908
Emergency-CSCF....910
A-BG....910
Technical Specifications....911
Platform Requirements....912
Licenses....912
Network Deployments and Interfaces....913
SCM in a CDMA2000 Data Network Deployment....913
Integrated CSCF / A-BG / HA....913
Logical Network Interfaces (Reference Points)....913
SCM in a GSM/UMTS Data Network Deployment....914
CSCF / A-BG / GGSN Deployment....914
Logical Network Interfaces (Reference Points)....915
Features and Functionality - Base Software....917
Bulk Statistics Support....918
Call Abort Handling....919
Call Forking....919
Call Types Supported....919
Congestion Control....919
DSCP Marking....920
Early IMS Security....920
Emergency Call Support....921
Error Handling....921
Future-proof Solution....921
Intelligent Integration....921
Interworking Function....921
IPv6 Support....921
Management System Overview....923
MSRP Support....924
PCRF Policy Control....925
Presence Enabled....925
Redirection....925
Redundancy and Session Recovery....925
Registration Event Package....925
Signaling Compression (SigComp)....925
SIP Denial of Service (DoS) Attack Prevention....925
SIP Intelligence at the Core....926
SIP Large Message Support....926

SIP Routing Engine 926
Shared Initial Filter Criteria (SiFC) 927
Telephony Application Server (TAS) Basic Supported 927
Threshold Crossing Alerts (TCA) Support 928
Trust Domain 929
Features and Functionality - External Application Support 930
Web Element Management System 930
Features and Functionality - Licensed Enhanced Feature Support 932
Interchassis Session Recovery 932
IPSec Support 933
IPv4-IPv6 Interworking 933
Lawful Intercept 935
Session Recovery Support 935
TLS Support in P-CSCF 936
How the SCM Works 937
Admission and Routing 937
CSCF Access Control Lists 937
Translation Lists 937
Route Lists 937
Signaling Compression 938
Supported Standards 939
Release 9 3GPP References 939
Release 8 3GPP References 939
3GPP2 References 941
IETF References 942
Other 944

Serving GPRS Support Node (SGSN) Overview 945

Product Description 946
Platform Requirements 946
Licenses 946
Network Deployments and Interfaces 948
SGSN and Dual Access SGSN Deployments 948
SGSN/GGSN Deployments 949
SGSN Logical Network Interfaces 950
SGSN Core Functionality 954
All-IP Network (AIPN) 954
SS7 Support 954
PDP Context Support 955
Mobility Management 955
GPRS Attach 955
GPRS Detach 956
Paging 956
Service Request 956
Authentication 956
P-TMSI Reallocation 957
P-TMSI Signature Reallocation 957
Identity Request 957
Location Management 957
Session Management 958
PDP Context Activation 958
PDP Context Modification 958
PDP Context Deactivation 959
PDP Context Preservation 959
Charging 959

SGSN in GPRS/UMTS Network 959
Features and Functionality 961
APN Aliasing 962
Default APN 962
APN Resolution with SCHAR or RNC-ID 963
Automatic Protection Switching (APS) 963
Authentications and Reallocations -- Selective 964
Avoiding PDP Context Deactivations 965
Bulk Statistics Support 965
CAMEL Service Phase 3, Ge Interface 966
CAMEL Service 966
CAMEL Support 966
Ge Interface 967
CAMEL Configuration 967
Configurable RAB Asymmetry Indicator in RAB Assignment Request 967
Direct Tunnel 967
DSCP Templates for Control and Data Packets - Iu or Gb over IP 967
ECMP over ATM 968
Equivalent PLMN 968
First Vector Configurable Start for MS Authentication 968
GMM-SM Event Logging 968
Gn/Gp Delay Monitoring 969
GTP-C Path Failure Detection and Management 969
Handling Multiple MS Attaches All with the Same Random TLLI 969
HSPA Fallback 970
Ignore Context-ID During 4G/3G Handovers 970
Intra- or Inter-SGSN Serving Radio Network Subsystem (SRNS) Relocation (3G only) 970
Lawful Intercept 971
Link Aggregation - Horizontal 971
Local DNS 971
Local Mapping of MBR 971
Local QoS Capping 972
Management System Overview 972
Multiple PLMN Support 973
Network Sharing 974
Benefits of Network Sharing 974
GWCN Configuration 974
MOCN Configuration 975
Implementation 976
NPU FastPath 977
NRPCA - 3G 978
Operator Policy 978
Some Features Managed by Operator Policies 978
Overcharging Protection 979
QoS Traffic Policing per Subscriber 979
QoS Classes 979
QoS Negotiation 979
DSCP Marking 979
Traffic Policing 980
Reordering of SNDCP N-PDU Segments 980
Session Recovery 981
SGSN Pooling and Iu-Flex / Gb-Flex 981
Gb/Iu Flex Offloading 982
Short Message Service (SMS over Gd) 982

SMS Authentication Repetition Rate 982
SMSC Address Denial 983
Threshold Crossing Alerts (TCA) Support 983
Tracking Usage of GEA Encryption Algorithms 984
VLR Pooling via the Gs Interface 984
How the SGSN Works 985
First-Time GPRS Attach 985
PDP Context Activation Procedures 986
Network-Initiated PDP Context Activation Process 988
MS-Initiated Detach Procedure 989
Supported Standards 991
IETF Requests for Comments (RFCs) 991
3GPP Standards 991
ITU Standards 994
Object Management Group (OMG) Standards 995
Traffic Performance Optimization Overview 997
Overview 998
TPO Deployment 998
Feature Specifications 1001
Platform Requirements 1001
License Requirements 1001
Supported Standards 1001
TCP Optimization 1001
TCP Optimization Techniques 1001
HTTP Optimizations 1004
HTTP Optimization Techniques 1004
HTTP Compression 1004
URL Rewrite 1006
Advertisement Filter 1008
Session Recovery 1012
TPO Administration 1013
Disabling/Enabling TPO Optimizations 1013
MUR Reporting for TPO 1013
Switching TPO Policy 1014
How TPO Works 1015
Terms and Definitions 1015
TPO Processing 1016
Policy-based TPO Processing 1016
Charging Action Based TPO Processing 1016
Web Element Manager Overview 1019
Supported Features 1022
FCAPS Support 1022
Fault Management 1022
Configuration Management 1022
Accounting Management 1023
Performance Management 1024
Security Management 1024
Additional Features 1026
High Availability Redundant Server Clustering 1026
Management Integration Capabilities 1026
Database Management and Redundancy Support 1026
Multiple Language Support 1026
Context-Sensitive Help System 1026

Stand-alone Offline Help System....1026
Multiple OS Support....1027
Web Element Manager System Requirements....1028
Server Hardware Requirements....1028
Sun Solaris Server Hardware Requirements....1028
Red Hat Enterprise Linux Server Hardware Requirements....1028
Operating System Requirements....1028
Sun Solaris Operating System Requirements....1029
Red Hat Enterprise Linux Operating System Requirements....1029
Client Access Requirements....1030
WEM Architecture....1031
Host Filesystem....1031
Apache Web Server....1031
WEM Server FCAPS Support....1031
Fault Management....1031
Configuration Management....1032
Accounting Management....1033
Performance Management....1033
Security Management....1034
WEM Process Monitor....1035
Bulk Statistics Server....1035
Script Server....1035
PostgreSQL Database Server....1036
Northbound Server....1036
WEM Logger....1038
Technical Specifications....1039
Physical Dimensions....1040
Weights....1041
Power Specifications....1042
Estimating Power Requirements....1042
Mounting Requirements....1043
Interface Specifications....1045
SPIO Card Interfaces....1045
Console Port Interface....1045
Fiber SFP Interface....1047
10/100/1000 Mbps RJ-45 Interface....1047
Central Office Alarm Interface....1048
BITS Timing Interface....1051
Fast Ethernet Line Card (FELC/FLC2) Interfaces....1051
10/100 Mbps RJ-45 Interface....1051
Gigabit Ethernet Card (GELC/GLC2/QGLC) SFP Interfaces....1052
1000Base-SX....1052
1000Base-LX Interface....1053
1000Base-T....1053
10 Gigabit Ethernet Line Card (XGLC) SFP+....1054
10GBase-SR....1054
10 Base-LR Interface....1054
Fiber ATM/POS OC-3 (OLC2) Multi-Mode Interface....1055
Fiber ATM/POS OC-3 SM IR-1 Interface....1055
Channelized Line Cards....1056
Channelized Line Card (CLC2) with Single-Mode Interface....1056
Channelized Line Cards (CLC2) with Multi-Mode Interface....1057
Safety, Electrical and EMC Certifications....1059

Federal Communications Commission Warning 1060
ICS Notice 1060
Laser Notice 1060
Safety Certifications 1061
Electrical Certifications 1062
Environmental Certifications 1063
Acoustic Noise 1064
Electromagnetic Compatibility (EMC) Compliance 1065
Japan VCCI-A 1065
Korean EMC 1065
Environmental Specifications 1067
Operating and Storage Parameters 1068
Supported Environmental Standards 1069
Chassis Air Flow 1070
Glossary of Terms 1071

About this Guide

This document pertains to the features and functionality that run on and/or that are related to the Cisco® ASR 5000 Chassis.

This preface includes the following sections:

- Conventions Used
- Contacting Customer Support
- Additional Information

Conventions Used

The following tables describe the conventions used throughout this documentation.

Icon	Notice Type	Description
	Information Note	Provides information about important features or instructions.
	Caution	Alerts you of potential damage to a program, device, or system.
	Warning	Alerts you of potential personal injury or fatality. May also alert you of potential electrical hazards.
	Electrostatic Discharge (ESD)	Warns you to take proper grounding precautions before handling ESD sensitive components or devices.

Typeface Conventions	Description
Text represented as a `screen display`	This typeface represents text that appears on your terminal screen, for example: `Login:`
Text represented as **`commands`**	This typeface represents commands that you enter at the CLI, for example: **`show ip access-list`** This document always gives the full form of a command in lowercase letters. Commands are not case sensitive.
Text represented as a **`command`** *`variable`*	This typeface represents a variable that is part of a command, for example: **`show card`** *`slot_number`* *`slot_number`* is a variable representing the desired chassis slot number.
Text represented as menu or sub-menu names	This typeface represents menus and sub-menus that you access within a software application, for example: Click the **File** menu, then click **New**.

Command Syntax Conventions	Description
{ **`keyword`** or *`variable`* }	Required keywords and variables are surrounded by braces. They must be entered as part of the command syntax.
[**`keyword`** or *`variable`*]	Optional keywords or variables that may or may not be used are surrounded by brackets.

Command Syntax Conventions	Description
|	Some commands support alternative variables. These "options" are documented within braces or brackets by separating each variable with a vertical bar. These variables can be used in conjunction with required or optional keywords or variables. For example: **{ nonce | timestamp }** OR [**count** *number_of_packets* |**size** *number_of_bytes*]

Contacting Customer Support

Go to http://www.cisco.com/cisco/web/support/ to submit a service request. A valid Cisco account (username and password) is required to access this site. Please contact your Cisco account representative for additional information.

Additional Information

Refer to the following guides for supplemental information about the system:

- *Cisco ASR 5000 Installation Guide*
- *Cisco ASR 5000 System Administration Guide*
- *Cisco ASR 5x00 CDMA Command Line Interface Reference*
- *Cisco ASR 5x00 eHRPD / LTE Command Line Interface Reference*
- *Cisco ASR 5x00 GPRS / UMTS Command Line Interface Reference*
- *Cisco ASR 5x00 Thresholding Configuration Guide*
- *Cisco ASR 5x00 SNMP MIB Reference*
- *Web Element Manager Installation and Administration Guide*
- *Cisco ASR 5x00 AAA Interface Administration and Reference*
- *Cisco ASR 5x00 GTPP Interface Administration and Reference*
- *Cisco ASR 5x00 Release Change Reference*
- *Cisco ASR 5x00 Statistics and Counters Reference*
- *Cisco ASR 5x00 Gateway GPRS Support Node Administration Guide*
- *Cisco ASR 5x00 HRPD Serving Gateway Administration Guide*
- *Cisco ASR 5000 IP Services Gateway Administration Guide*
- *Cisco ASR 5x00 Mobility Management Entity Administration Guide*
- *Cisco ASR 5x00 Packet Data Network Gateway Administration Guide*
- *Cisco ASR 5x00 Packet Data Serving Node Administration Guide*
- *Cisco ASR 5x00 System Architecture Evolution Gateway Administration Guide*
- *Cisco ASR 5x00 Serving GPRS Support Node Administration Guide*
- *Cisco ASR 5x00 Serving Gateway Administration Guide*
- *Cisco ASR 5000 Session Control Manager Administration Guide*
- *Cisco ASR 5000 Packet Data Gateway/Tunnel Termination Gateway Administration Guide*
- Release notes that accompany updates and upgrades to the StarOS for your service and platform

Chapter 1
The Cisco® ASR 5000 Platform

Designed exclusively for the wireless industry, the Cisco® ASR 5000 Chassis provides an ultra-high density solution for deployment in wireless carrier and operator environments.

The ASR 5000 is a high-performance, carrier-grade platform that offers industry-leading wireless data capacity while enabling numerous integrated applications for additional revenue generation.

Large, high-demand multimedia applications require an ever increasing amount of processing power and memory. The ASR 5000 has been designed to address these needs and provide a scalable platform to meet the needs of future fourth generation (4G) networks.

Figure 1. The Cisco® ASR 5000



Characteristics of the System

This section provides an overview of some of the key characteristics of the system. Detailed information for these characteristics is provided in subsequent chapters of this guide.

- Carrier-grade Hardware Design
  - NEBS Level 3 Compliant components
  - UL certified
  - Five 9s availability
  - Local alarming and alarm cut-off capabilities
  - High availability design (less than 4.35 minutes of downtime per year)
- Redundancy
  - 1:1 Switch Processor Card (SPC)/System Management Card (SMC) redundancy
  - 1:n Packet Services Cards (PSC/PSC2) redundancy - allowing redundancy of multiple active to multiple redundant for up to 14 total packet processing cards

    Important: 1:1 redundancy is supported for these cards however some subscriber sessions and accounting information may be lost in the event of a hardware or software failure even though the system remains operational.
  - 1:1 card-level redundancy for Switch Processor Input/Output (SPIO), and all types of line cards
  - 1:1 port-level redundancy for SPIO and all types of line cards
  - Integrated hardware and software redundancy with automatic failover features
  - Optional session recovery support for the following call types:

    - WiMAX ASN GW services supporting simple IP, Mobile IP, and Proxy Mobile IP

    - PDSN services supporting simple IP, Mobile IP, and Proxy Mobile IP

    - HA services supporting Mobile IP and/or Proxy Mobile IP session types with or without per-user Layer 3 tunnels

    - GGSN services for IPv4 and PPP PDP contexts

    - Home-NodeB Gateway (HNB-GW) services for Femto-UMTS session types

    - MME services for LTE/SAE networks and 3G services

    - LNS session types
  - Optional Interchassis Session Recovery
  - Hot swappable cards, allowing dynamic card replacement while the system is operational
  - Load sharing, hot swappable - 48 VDC power filters with redundant power circuitry throughout
- High Capacity Design
  - Self-healing 320 Gbps packet-based Switch Fabric
  - System Management Bus

  - 32 Gbps Control Bus
  - 140 Gbps Redundancy Bus
- Operating System
  - Linux™-based
  - Application hosting capabilities
  - Modular, distributed processing
  - Robust development environment

Features and Benefits

Some of the benefits found in deploying the system include.

Table 1. Features and Benefits of the System

Feature	Benefit
Policy Control and Charging Rule Function Support through Intelligent Policy Control Function (IPCF)	Provides the policy control decision and flow based charging control functionalities for the subscriber's data traffic in the broadband network. It performs the decision based on inputs received from PCEFs/BBERFs, the subscription profile repository (SSC/SPR), application function (AF), if available as well as operator's local policy. IPCF collates the subscriber and application data to authorize QoS resources and instructs the transport plane on how to proceed with the underlying data traffic in terms of charging and policy. PCC solution provides following components: • Intelligent Policy Control Function (IPCF) • Subscriber Service Controller (SSC) • Policy Provisioning Tool (PPT)
Mobility Management Entity (MME) service support in LTE/SAE networks	Delivers unrivaled throughput, call transaction rates, and packet processing, along with significant memory resources.Provides UE state management (attach, detach, idle, RAN mobility), authentication, paging, mobility with 3GPP 2G/3G nodes (SGSN), roaming, and other bearer management functions. Also provides Integration of multiple core network functions.High transaction rates for attaches, activations, TAUs, handoffs, and paging along with congestion management, load sharing, and MME pooling.MME provides intelligent signaling heuristics to maximize performance and self Optimizing Network (SON) capabilities to the radio and packet core network. It also provides dynamic optimization of network topology based on usage patterns to reduce latency and backhaul costs.Circuit Switch (CS) Fallback for voice traffic
Mobile data service support for WiMAX networks	Provides WiMAX ASN GW, WiMAX Foreign Agent (FA), and WiMAX Home Agent (HA) services within a single chassis, or as distributed network functions supporting both Simple and Mobile IP.Provides WiMAX ASN Paging Controller and Location Registry (ASN PC/LR) services within a single chassis, or co-located as distributed network functions supporting paging procedures for idle mode entry and exit and location update.Provides multiple host support behind a WiMAX Customer Premise Equipment (CPE) through one primary airlink session Provides optional base station monitoring feature to monitor base stations attached to it.

Feature	Benefit
Wireless data service support for 3G CDMA2000 and GPRS/UMTS and for 2.5G/3G GPRS/UMTS networks	• Provides Packet Data Service Node (PDSN), Foreign Agent (FA), and Home Agent (HA) services within a single chassis, or as distributed network functions supporting both Simple and Mobile IP. • Provides Gateway GPRS Support Node (GGSN), Foreign Agent (FA), and Home Agent (HA) services within a single chassis, or as distributed network functions supporting basic data and Mobile IP functionality. • Provides Serving GPRS Support Node (SGSN) and Gateway GPRS Support Node (GGSN) services within a single chassis, or as distributed network functions supporting both the control and data planes.
Wireless data service support for Femto (UMTS/CDMA) subscriber in 3G UTMS networks	• Provides Home-NodeB Gateway (HNB-GW) service for Femto access network user to connect voice and IP data traffic with CS/PS core network. • It supports multiple services on a single chassis or as distributed network functions supporting enhanced voice and IP data functionality.
Proxy Mobile IP	Provides a mobility solution for subscriber's with Mobile Nodes (MNs) that do not implement the Mobile IP protocol stack.
Full Handover Support	Compliance with 3GPP procedures for Mobility Management, Location Management, and Session Management ensure high volume, load-balancing, and successful handover.
Direct Tunneling	Reduces latency by creating GTP-U tunnel for data transport between the RNC and the GGSN while optimizing SGSN usage for control plane processing and user plane functionality for cases such as roaming and lawful intercept.
L2TP Tunneling	Layer 2 Tunneling Protocol (L2TP) support encapsulates data packets between an L2TP Access Concentrator (LAC) and an L2TP Network Server (LNS) to create a Virtual Private Network (VPN).The system can be configured as either an LAC or LNS in support of L2TP.LAC is an optional licensed feature.
Lawful Intercept (optional licensed feature)	Provides Telecommunication Service Providers (TSPs) with a mechanism to assist Law Enforcement Agencies (LEAs) in the monitoring of suspicious individuals (referred to as targets) for potential criminal activity.
IPSEC	Secure VPN Connectivity for the enterprise.Secure L2TP and mobile IP tunneling.System architecture provides IPSEC implementation with no performance degradation.Encryption Daughter Card (EDC) availability for hardware-based encryption.
OSPF Routing	Provides optional OSPFv2 routing with NSSA support.
BGP-4 Routing	Provides optional BGP-4 routing.
Flow-based Traffic Policing (optional licensed feature)	Provides the traffic policing to control session flow on a flow classification basis.Provides the QoS to control session flow on a flow classification basis.
Traffic Policing and Shaping (optional licensed feature)	Provides the ability to limit network bandwidth on a per subscriber basis.Provides the ability to buffer packets which exceeds the allowed limit and transmit them once the traffic flow comes below the exceed limit.
Dynamic QoS Renegotiation (ECS support required)	Provides the ability to manage the risk of bandwidth mis-appropriation. This feature allows the Enhanced Charging Service (ECS) to analyze application traffic, and triggers QoS renegotiation with the AGW to optimize service performance.It provides Network Controlled QoS (NCQoS) and traffic class-based QoS renegotiation support.

Feature	Benefit
GTPP Server Group support	• Provides more than one list of GTPP servers through GTPP server group feature at context level for GTPP accounting functionality • Provides GTPP accounting functionality to individual subscriber through APN
Session Redirection (“hotlining”) (optional licensed feature)	Provides the ability to redirect subscriber traffic to an external server through the application of Access Control List (ACL) rules.Relies on the Change of Authorization (CoA) feature for the dynamic redirection of subscriber IP datagrams.
PDSN RAN Optimization	Provides session redirection based on sessions having a specific MSID or received from specific PCF zones.
Change of Authorization (CoA) and Packet of Disconnect RADIUS message support	Allows system contexts to listen for and act upon CoA and/or disconnect messages from a RADIUS server.CoA messages enable the dynamic changing of subscriber attributesDisconnect Messages (DMs) allow the termination of subscriber sessions from a particular RADIUS serverCoA is supported for use with PDSN.
RADIUS Server Group support (optional licensed feature)	Provides more than one list of AAA servers through RADIUS server group feature at context level for AAA functionalityProvides AAA functionality to individual subscriber through realm (domain) APN
Adjunct Compression Server	Reduces network complexity and capital expenditure.Application based compression that helps conserve radio bandwidth resources.
802.1Q VLAN Tagging (optional licensed feature)	Provides layer 2 VPN connectivity.Simplified network configuration.Allows overlapping IP addresses within the same context.
Prepaid (optional licensed feature)	Provides subscriber billing based on data volume or session time.Mid-session account balance updates.
Robust Header Compression	Provides Robust Header Compression (ROHC) support for IP packets.
HA Proxy DNS Intercept	Provides a solution for unreachable (fire-walled) DNS servers in visited networks.
“In-Line” data services capability (optional licensed feature)	Allows for deep packet inspection to support enhanced/advanced billing techniquesImproved subscriber awareness to more quickly identify usage trends and tailor content to subscriber's patternsIncreased revenue opportunities through application of new services with no or minimal processing degradation
Carrier-grade design	Ensures maximum level of reliability and service availabilityAllows for installation and/or co-location in central office facilities
Multiple context support	Allows operator to support multiple enterprise and home networks from a single systemAllows operators to assign duplicate/overlapping IP address ranges in different contexts
Multimedia Broadcast and Multicast Service (MBMS)	Provides a solution for transferring light video and audio clips and also a suitable method for mass communications to operator.It eliminates unnecessary replication of data on UMTS wireless networks by transmitting a single stream of data to multiple users.
Integrated “control node” function	Eliminates processing bottlenecksIntelligently distributes processing across multiple system processors for increased throughput

Feature	Benefit
Session Recovery (optional licensed feature)	Recovers all fully established sessions upon single hardware or software failure for the following call types: • ASN GW services supporting simple IP, Mobile IP, and Proxy Mobile IP • PDSN services supporting simple IP, Mobile IP, and Proxy Mobile IP • HA services supporting Mobile IP and/or Proxy Mobile IP session types with or without per-user Layer 3 tunnels • GGSN services for IPv4 and PPP PDP contextsLNS session types • Restores data and control packet state information, subscriber data statistics, subscriber idle time and other timer-related data • Provides an in-service recovery mechanism to increase system availability and overall fault tolerance without significant interruption of subscriber services and without loss of accounting information.
MIP NAT Traversal (optional licensed feature)	This feature allows the HA to set up a UDP tunnel for an MN that is behind a NAT device.
IMS Authorization Service and Gx interface support (optional licensed feature)	Provides Gx and Gy interface support to implement IMS authorization in GPRS/UMTS network. described in 3GPP Release 6 and 7.Provide sufficient, uninterrupted, consistent, and seamless user experience to a roaming IMS subscriber for an application along with dynamic charging functionality for the particular IMS application used.
IMS Authorization Service and Ty interface support (optional licensed feature)	Provides Ty interface support for roaming IMS subscriber to implement IMS authorization in CDMA2000 network as described in 3GPP2 standards.Provide sufficient, uninterrupted, consistent, and seamless user experience to a roaming IMS subscriber for an application along with dynamic charging functionality for the particular IMS application used.
IP Services Gateway (optional licensed product)	Provides legacy access gateways (GGSNs, PDSNs, HAs, etc.) that are not service capable, to provide managed services such as enhanced charging, stateful firewalls, traffic performance optimization, and others.
Integrated Session Control Manager (SCM) Functionality	The SCM provides an easy on-ramp to deploying Session Initiation Protocol (SIP)-based services and a future-proof migration path to the IP Multimedia Subsystem/Multimedia Domain (IMS/MMD) architectures.The SCM consists of an IETF-compliant SIP Proxy/Registrar, a 3GPP/3GPP2-compliant Proxy Call Session Control Function (P-CSCF), and a Policy Agent (PA).
PCF/BS Monitoring	Provides PDSN service to monitor PCFs attached to it.Provides ASN BS monitoring facility to AS NGW service attached to it.
Linux-based operating system	Ensures compatibility with leading applicationsAllows for integration of third-party applications into system to host “in-line” services
Integrated data aggregation features	Delivers wire speed transport throughout systemEliminates need to add external routing devices to move from high-speed to low-speed links
Future-proof design	Robust hardware platform allows for easy migration to next-generation data services using the same chassisScalable hardware and software components allow you to cost effectively add capacity as your subscriber-base increases

Feature	Benefit
Web-based element management	Reduces operational complexityImproves overall system management accuracy and securityAllows for remote monitoring and configuration, using SNMPv1 and CORBAProvides security for management data using Secure Sockets Layer (SSL) encryptionAllows for seamless integration with external network, service, and business layer management applications through CORBA interface
Application Programming Interface for management	Allows for internal development of custom management applicationsAllows integration with new or existing service management applicationsUses industry standard Interface Definition Language (IDL) as API for integration
Intelligent Packet Monitoring System (IPMS)	Provides more detailed network performance information on control events and measures call success and protocolsVerifies accuracy of accounting records and analyzes set up failure causesIdentifies network faults and counts number of affected users when manager/line card/port fails and debugs connection issuesComprehensive query tool to simplify searches across multiple access gateways and ability to diagnose the calls based on disconnect reasons
Command Line Interface	Designed for intuitive use by experienced network administratorsCLI commands are designed to be conducive to scripting, allowing operators to easily issue commands using EXPECT scripts and interactive applications written in Tcl/TkHelps operators securely configure, upgrade, monitor, and set system triggers from remote locations, supporting Telnet and Secure Shell (SSH) protocolsRemote management features help manage and deploy large scale, carrier-class, highly available and very manageable, easily monitored networkContext-sensitive Help for all commands, keywords, and variables

Chapter 2
Product, Service and Feature Licenses

This chapter provides information regarding Cisco Systems' licensed products, services, and features. The following sections are included:

- Supported Product_License Quick Reference
- Session Use and Feature Use Licenses
- Default Licenses

Supported Product/License Quick Reference

The following table provides a quick reference list of supported products and the license name of services and features for the ASR 5000 platforms.

Table 2. Supported Products and Licenses

Inline Service/Feature Name	Product	License Name(s)
Always On Licensing	P-GW S-GW	Always On Licensing
Application Detection and Control	GGSN IPSG PDSN P-GW	Application Detection and Control
APN Remapping, Blank APN, Default APN, Wildcard APN	SGSN	APN Aliasing
BCMCS	PDSN	Broadcast & Multicast Services
Circuit Switched Fallback	MME	Circuit Switched Fallback
CoA, RADIUS DM, and Session Redirection (Hotlining)	ASN GW GGSN HA IPSG PDSN P-GW	Dynamic RADIUS extensions (CoA and PoD)
Content Filtering ICAP Interface Support	GGSN P-GW	Content Filtering ICAP Interface
Direct Tunnel, selective	SGSN	Direct Tunnel
Dynamic QoS Renegotiation (Traffic Class-based QoS and Network Controlled QoS)	GGSN	GGSN Dynamic QoS Renegotiation
Dynamic Mobile IP Key Update (DMU)	PDSN	Dynamic Mobile IP Key Update
Enhanced Charging Service	ASN GW GGSN HA IPSG PDSN P-GW	Enhanced Charging Bundle 1
Enhanced Charging Service with Diameter and DCCA functionality	GGSN HA IPSG PDSN P-GW	Enhanced Charging Bundle 2
Generic Routing Encapsulation (GRE) Interface Tunneling	GGSN P-GW	GRE Interface Tunneling

Inline Service/Feature Name	Product	License Name(s)
Gx Interface Support	GGSN HA HSGW IPSG PDSN P-GW SGW	Dynamic Policy Interface
HA DNS Intercept Proxy	HA	HA DNS Intercept Proxy
Header Enrichment	GGSN IPSG P-GW	Header Enrichment
HTTP Header Encryption	GGSN IPSG P-GW	HTTP Header Encryption
HTTP Header Enrichment and Encryption	GGSN IPSG P-GW	HTTP Header Enrichment and Encryption
Heuristic or Idle Mode Paging	MME	Optimized Paging
Integrated Content Filtering	GGSN HA IPSG PDSN P-GW	Integrated Content Filtering Service
Intelligent Traffic Control (ITC)	ASN GW HA HSGW PDSN	Intelligent Traffic control
Interchassis Session Recovery	GGSN HA P-GW SCM	Inter-Chassis Session Recovery
IP Header Compression	HSGW PDSN	Robust Header Compression
IP Security (IPSec)	ASN GW GGSN HA HSGW IPSG PDIF PDSN P-GW SCM S-GW	IPSec
IP Security (IPSec) for LTE	MME	Network Domain Security

Inline Service/Feature Name	Product	License Name(s)
Lawful Intercept	ASN GW GGSN HA LNS PDIF PDSN P-GW SCM SGSN	Lawful Intercept/Enhanced Lawful Intercept
Local Policy Decision Engine	P-GW	Local Policy Decision Engine
L2TP Access Concentrator	ASN GW GGSN IPSG PDSN P-GW	L2TP LAC
L2TP Network Server	GGSN/LNS PDSN/LNS	L2TP LNS
MIP NAT Traversal	HA P-GW	MIP NAT Traversal
Multi Protocol Label Switching (MPLS)	GGSN	MPLS
Multimedia Broadcast and Multicast Service	GGSN	MBMS
MSID and PCF Zone Based Call Redirection	HA	PDSN RAN Optimization, Bundle 1
Network Address Translation	GGSN HA PDSN P-GW	NAT/PAT With DPI
Network Mobility (NEMO) Service	P-GW	NEMO
PCF Monitoring	PDSN	PCF/BS Monitoring
Per Subscriber Stateful Firewall	GGSN HA IPSG PDSN P-GW	Per Subscriber Stateful Firewall
RADIUS Pre-paid Billing	HA PDSN	Prepaid Accounting/IS-835C Prepaid Bundle
Proxy-Mobile IP	ASN GW/FA GGSN/FA HSGW IPSG PDIF PDSN/FA P-GW	Proxy MIP

Inline Service/Feature Name	Product	License Name(s)
Remote Address-based RADIUS Accounting	ASN GW GGSN HA IPSG PDIF PDSN	Destination Based Accounting
Secure Sockets Layer (SSL) for Transport Layer Security (TLS)	SCM	SSL
Session Recovery	ASN GW GGSN HA HSGW IPSG MME PDIF PDSN P-GW SCM SGSN	Session Recovery
Single Radio Voice Call Continuity	MME	SRVCC Support
Traffic Policing and Shaping	ASN GW GGSN HA HSGW PDSN P-GW SCM S-GW	Per Subscriber Traffic Policing/Shaping
Ty Interface Support	HA PDSN	Dynamic Policy Interface
User Location Information Reporting	MME P-GW	ULI Reporting
VLANs	ASN GW GGSN HA HSGW IPSG PDIF PDSN P-GW SCM S-GW	Layer 2 Traffic Management
WiMAX Paging Controller	ASN GW	WiMAX Paging Controller/Location Register

Session Use and Feature Use Licenses

Session use and feature use licenses are software mechanisms used to provide session limit controls and enable special features within the system. These electronic licenses are stored in the system's configuration file that is loaded as part of the system software each time the system is powered on or restarted.

Session Use Licenses

Session use licenses limit the number of concurrent sessions that a system is capable of supporting per service type and are acquired on an as needed basis. This allows carriers to pay only for what they are using and easily increase capacity as their subscriber base grows. Session use licenses are available for the following services:

- Access Service Network Gateway (ASN GW) (Includes FA, DHCP, Proxy MIP and RADIUS AAA Server Groups)
- Enhanced Charging Service (ECS):
  - Enhanced Charging Bundle 1 1k Sessions
  - Enhanced Charging Bundle 2 1k Sessions (Includes Diameter and DCCA functionality with ECS)
- EV-DO Rev A PDSN (Includes FA, RADIUS AAA Server Groups, and PDSN RAN Optimization, Bundle 1)
- EV-DO Rev A / PDSN [UPGRADE] 1k Sessions or 10k Sessions (UPGRADE: Will convert PDSN into EV-DO Rev A / PDSN)
- Gateway GPRS Support Node (GGSN) (Includes RADIUS AAA Server Groups)
- HA license for GGSN (Includes RADIUS AAA Server Groups and MIP NAT Traversal)
- Home Agent (HA) (Includes RADIUS AAA Server Groups)
- HRPD Serving Gateway (HSGW) (Includes Dynamic Policy Interface, Session Recovery, IPv6, Intelligent Traffic Control, and Enhanced Charging Bundle 2)
- IP Services Gateway
- L2TP Network Server (LNS)
- Mobility Management Entity (MME) (Includes Session Recovery, and Enhanced Lawful Intercept)
- Packet Data Service Node (PDSN) (Includes RADIUS AAA Server Groups)
- PDN Gateway (P-GW) (Includes Dynamic Policy Interface, Lawful Intercept, Session Recovery, RADIUS AAA Server Groups, IPv6, Intelligent Traffic Control, and Enhanced Charging Bundle 2)
- PDIF-Service (Includes IPSec, FA, and RADIUS AAA Server Groups)
- SCM (Proxy/Registrar IETF RFC3261, Proxy P-CSCF, Serving S-CSCF, Interrogating I-CSCF, and A-SBC SW License)
- Serving Gateway (S-GW) (Includes Dynamic Policy Interface, Lawful Intercept, Session Recovery, Proxy MIP, and IPv6)

Feature Use Licenses

Feature use licenses enable certain features/functionality within the system and are distributed based on the total number of sessions supported by the system. Licenses are available for each of the following features:

- Always On Licensing
- Application Detection and Control
- Content Filtering ICAP Interface (Requires Enhanced Charging Bundle 1 or Enhanced Charging Bundle 2)
- Circuit Switched Fallback
- Destination Based Accounting
- DHCP
- Diameter Closed-Loop Charging Interface
- Diameter EAP
- Dynamic Mobile IP Key Update
- Dynamic Policy Interface (Includes DIAMETER Closed-Loop Charging Interface)
- Dynamic QoS Traffic Policing
- Dynamic Radius extensions (CoA and PoD)
- Enhanced Lawful Intercept
- External Service Steering
- Foreign Agent (FA)
- GGSN Dynamic QoS Renegotiation
- HA DNS Intercept Proxy
- IKEv2 including Multi-Authentication
- In-line services usage
- Intelligent Traffic Control
- Inter-chassis Session Recovery
- IP Security (IPSec)
- IPv4 Routing Protocols
- IPv6 (this is enabled by default)
- IS-835C Prepaid Bundle (Includes Change of Authorization, Destination Based Accounting, and Prepaid Accounting)
- Lawful Intercept
- Layer 2 Traffic Management (VLAN)
- Local Policy Decision Engine
- L2TP Access Concentrator (LAC) PDSN/HA/GGSN/ASN GW
- MAC Address Authorization and Sh Interface
- MIP NAT Traversal
- Mobile Enterprise Security Bundle (Includes IPSec, L2TP LAC PDSN, L2TP LAC HA)
- MPLS

- Network Domain Security
- Optimized Paging
- PCF/BS Monitoring
- PDSN RAN Optimization Bundle (Includes PCF Monitoring)
- Per Subscriber Traffic Policing/Shaping
- PPP Fast Setup
- Prepaid Accounting (Requires PDSN, HA, EV-DO Rev A / PDSN, and/or ASN GW)
- Proxy Mobile IP
- RADIUS AAA Server Groups
- RADIUS AAA Server Groups (Always On)
- Robust Header Compression (ROHC)
- RP Flow Control
- Secure Combo Phone Bundle (Includes IPSEC, IPSEC NAT Traversal)
- Session Recovery
- Simple IP Fallback
- SIP Application Serve
- Smartphone Tethering Detection
- SRVCC Support
- 3GPP2 Always-On Rp Extensions
- ULI Reporting
- User Layer 3 Tunneling
- Voice Media Gateway

Each license is associated with both SPC/SMC cards in a redundant SPC configuration ensuring correct support for the system in case of an SPC/SMC failover. This license is unique to each system and its respective SPC/SMC-based CompactFlash cards. Session use licenses can be upgraded remotely to increase system session capacity as new PAC/PSC cards are added.

Important: In the event that an SPC/SMC requires replacement, you will need to remove the CompactFlash card from the SPC/SMC being replaced and install it onto the replacement SPC/SMC. Failure to exchange the CompactFlash card on the SPC/SMC will cause the session license to not match both SPCs/SMCs. The system will recognize that one of the SPCs match, and the session use license for the system would still be valid. However, unmatched keys would result in a loss of redundancy for all license-enabled session use and features should the remaining SPC/SMC that possesses the correctly matched license fail.

Default Licenses

If a system boots with no license key installed, or an invalid license key is specified in the configuration file, a set of default limited session use and feature licenses is installed. The following Exec Mode command lists the license information;

```
show license information
```

The following shows the license information for a system with no license key installed. Notice that the session use licenses for PDSN, HA, GGSN, and L2TP LNS are limited to 10,000 sessions.

```
No license key installed. Using defaults.

Enabled Features:

  Feature                                  Applicable Part Numbers

  ---------------------------------------  ---------------------------

  PDSN:                                    [ 600-00-7501 / 600-00-7504 ]

   + FA                                    [ None ]

   + RADIUS AAA Server Groups              [ None ]

  HA:                                      [ 600-00-7502 / 600-00-7505

                                             600-00-7592 / 600-00-7593 ]

   + RADIUS AAA Server Groups              [ None ]

  GGSN:                                    [ 600-00-7544 / 600-00-7545 ]

   + DHCP                                  [ 600-00-7520 ]

   + RADIUS AAA Server Groups              [ None ]

  IPv4 Routing Protocols                   [ none ]

  IPSec                                    [ 600-00-7507 / 600-00-8537

  L2TP LAC                                  [ 600-00-7508 / 600-00-753

                                              600-00-8536 ]

  L2TP LNS                                 [ 600-00-7503 / 600-00-7506 ]

  Session Recovery                         [ 600-00-7513 / 600-00-7546

                                             600-00-7552 / 600-00-7554

                                             600-00-7566 / 600-00-7594

                                             600-00-9100 / 600-00-9101
```

```
                                                  600-00-7638 / 600-00-7640
                                                  600-00-7634 / 600-00-7595
                                                  600-00-7669 / None
                                                  600-00-7897 / None
                                                  600-20-0145 ]
  Proxy MIP:                                    [ 600-00-7512 / 600-00-7549
                                                  600-00-8538 ]
   + FA                                         [ None ]
  IPv6                                          [ 600-00-7521 / 600-00-7576 ]
  Lawful Intercept                              [ 600-00-7522 ]
  PCF/BS Monitoring                             [ 600-00-7514 / 600-00-8506 ]
  Dynamic Radius extensions (CoA and PoD)       [ 600-00-7518
  User Layer3 Tunneling                         [ none ]
  RADIUS AAA Server Groups                      [ None ]
Session Limits:
               Sessions  Session Type
               --------  ----------------------
                  10000  PDSN
                  10000  HA
                  10000  GGSN
                  10000  L2TP LNS
```

Chapter 3
ASR 5000 Hardware Platform Overview

This chapter describes the hardware components that comprise the ASR 5000.

It includes the following sections:

- The ASR 5000 Platform
- Chassis Configurations
- Chassis Description
- Power Filter Units
- Fan Tray Assemblies
- Application Cards
- Line Cards
- Card Interlock Switch
- Card Identifiers

The ASR 5000 Platform

The ASR 5000 multimedia core platform is designed for deployment in multimedia-enabled core networks. It features a distributed architecture that allows all tasks and services to be allocated across the entire platform. This platform allows operators to deploy more efficient mobile networks that support a greater number of concurrent calls, optimize resource usage, and deliver enhanced services, while providing scalability.

ASR 5000 hardware components support the following features:

- 1:1 redundancy for all hardware elements
- Hot-swappable sub-components
- Inter-chassis session recovery (ICSR)
- CLI support for telnet, SSH, and local login through a console port
- SNMP support for event notification

Figure 2. ASR 5000 Chassis



Chassis Configurations

The system is designed to scale from a minimum configuration, as shown in the table below, to a fully-loaded redundant configuration containing a maximum of 48 cards.

If session recovery is enabled, the minimum number of packet processing cards per chassis increases from one to four cards. Three packet processing cards are active and one packet processing card is standby (redundant). This minimal configuration is designed to protect against software failures only.

Table 3. Minimum Card Configurations

Component	Supported ASR 5000 Product	Redundant HW Configuration	Redundant HW + SW Configuration	Maximum per Chassis
Application Cards				
System Management Card (SMC)	All	2	2	2
Packet Services Card (PSC)	All	3 (2 active +1 standby) See Note 3 below.	4 (3 active +1 standby) See Note 3 and Note 4 below.	14
Packet Services Card Type A (PSCA)				
Packet Services Card 2 (PSC2)				
Packet Services Card 3 (PSC3)				
Packet Processing Card (PPC)				
Chassis Subcomponents				
Power Filter Unit (PFU)	All	2	2	2
Upper Fan Tray Assembly	All	1	1	1
Lower Fan Tray Assembly	All	1	1	1
Line Cards				
Switch Processor I/O (SPIO) Card	All	2	2	2
Redundancy Crossbar Card (RCC)	All	2	2	2
Fast Ethernet Line Card (FELC)	All	2	See Note 5 below.	28*
Fast Ethernet Line Card 2 (FLC2)	All	2		28*
Gigabit Ethernet Line Card (GELC)	All	2		28*
Gigabit Ethernet Line Card 2 (GLC2)	All	2		28*

Component	Supported ASR 5000 Product	Redundant HW Configuration	Redundant HW + SW Configuration	Maximum per Chassis
Quad Gigabit Ethernet Line Card (QGLC)	All	2		28*
10 Gigabit Ethernet Line Card (XGLC)	All	2		14**
Optical Line Card 2 (OLC2)	SGSN only	2		28*
Channelized Line Card 2 (CLC2)	SGSN only	2		28*

Notes:

1. These numbers represent the minimum number of components for hardware redundancy. Additional components are required if Session Recovery is to be supported.

2. These numbers represent the minimum number of components for: a) hardware and software redundancy; b) platforms with combined services. Additional components are required if Session Recovery is to be supported.

3. PSCs and PSCAs can be mixed in the same chassis. However, in a mixed card environment, system capacity reverts to the lower limit of a PSC and incompatible components prevent the running of crypto functions. PSCs, PSC2s, PSC3s and PPCs must not be mixed with other packet processing card types.

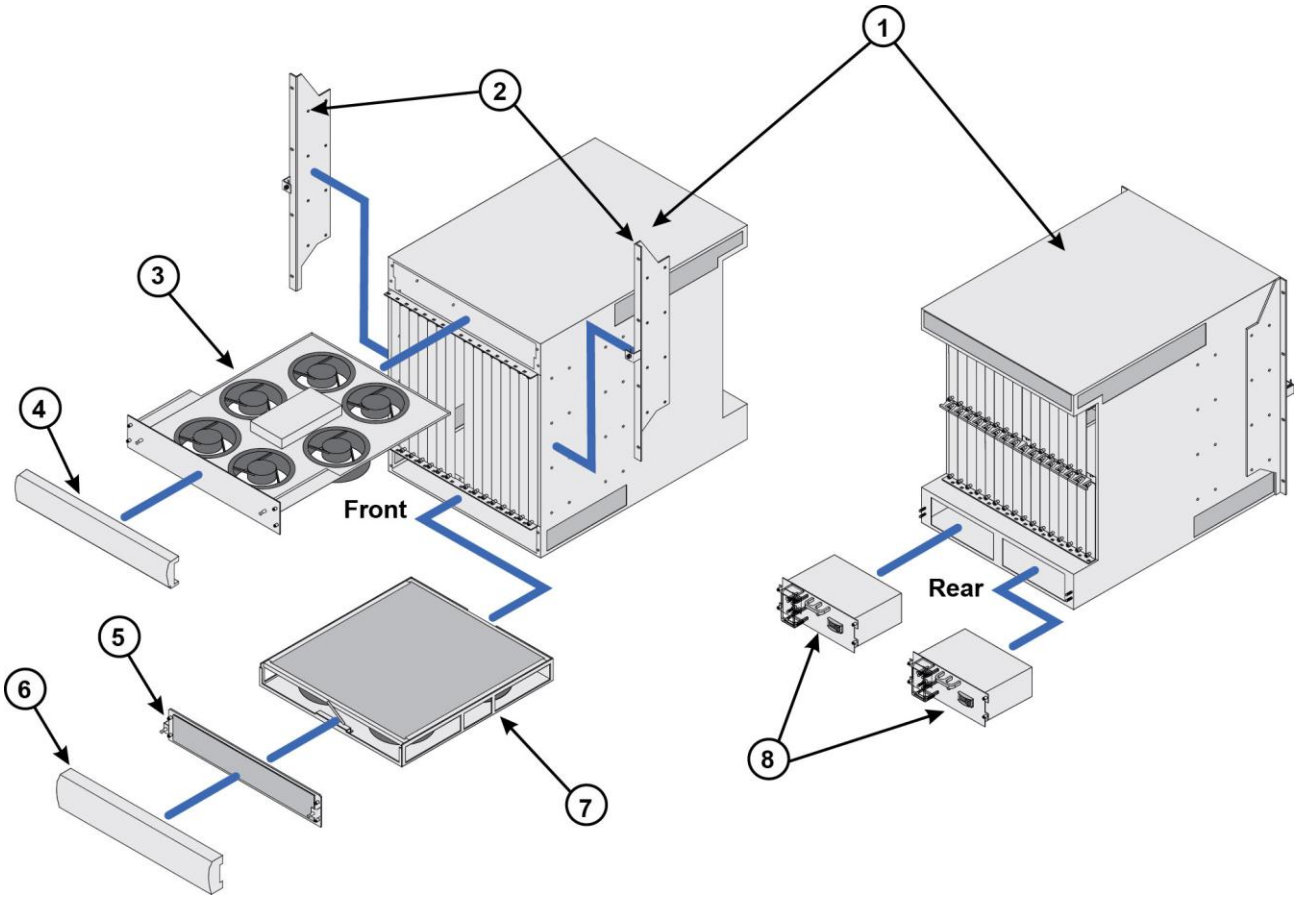
4. This is the minimum configuration for redundant SGSN service and MME service.

5. This number varies based on network deployment requirements.

*The maximum number of half-height line cards you can install is 28. However, redundant configurations may use fewer than the physical maximum number of line cards since they are inactive behind standby packet service cards.

**The 10 Gigabit Ethernet Line Card (XGLX) is a full-height line card that takes up the upper and lower slots in the back of the chassis. When referring to an installed XGLC, use the upper slot number only. Slot numbering for other installed half-height cards is maintained: 17 to 32 and 33 to 48, regardless of the number of installed XGLCs.

Figure 3. Chassis Components (front and rear views)



This diagram shows exploded views of the front and rear chassis components. They are described below.

Table 4. Chassis and Sub-component Identification Key

Item	Description
1	**Chassis:** Supports 16 front-loading slots for application cards and 32 rear-loading slots for line cards. To support the XGLC, a full-height line card, remove the half-height guide from the rear slots. The chassis ships with blanking panels over every slot except the following: 1, 8, 17, and 24. These are intentionally left uncovered for initial installation of application and line cards.
2	**Mounting brackets:** Support installation in a standard 19-inch rack or telecommunications cabinet. Flush and mid-mount options are supported. In addition, each bracket contains an electrostatic discharge jack for use when handling equipment.
3	**Upper fan tray:** Draws air through the chassis for cooling and ventilation. It then exhausts warmed air through the vents at the upper-rear of the chassis.
4	**Upper bezel:** Covers the upper fan tray bay.
5	**Lower fan tray cover/EMI shield:** Secures the lower fan tray assembly in place and serves as an EMI shield. The cover also provides an air baffle allowing air to enter into the chassis.
6	**Lower bezel:** Covers the lower fan tray bay.

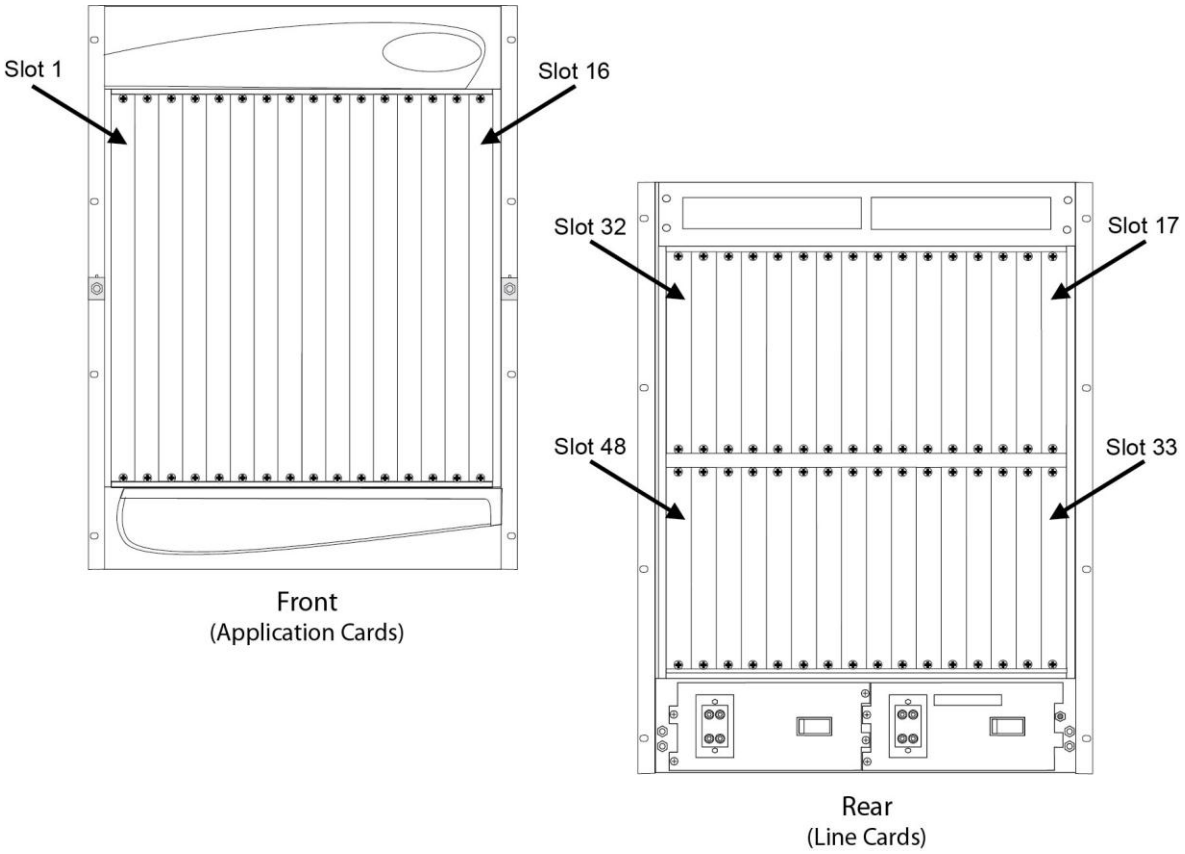
Item	Description
7	**Lower fan tray assembly:** Draws ambient air through the chassis' front and sides for cooling and ventilation. It is equipped with a particulate air filter to prevent dust and debris from entering the system.
8	**Power Filter Units (PFUs):** Each of the system's two PFUs provides -48 VDC power to the chassis and its associated cards. Each load-sharing PFU operates independently of the other to ensure maximum power feed redundancy.

Chassis Description

Slot Numbering

The ASR 5000 chassis features a 48-slot design with 16 front-loading slots for application cards and 32 rear-loading slots (16 upper and 16 lower) for line cards. ASR 5000.

Figure 4. Chassis Slot Numbers



The following table shows the front slot numbers and their corresponding rear slot numbers.

Table 5. Front and Rear Slot Numbering Relationship

Position	Slot Number															
Front	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
Rear Top Slots	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17

Position	Slot Number															
Rear Bottom Slots	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33

Rear Slot Numbering for Half-Height Line Cards

Rear-installed line cards must be installed directly behind their respective front-loaded application card. For example, an application card in Slot 1 must have a corresponding line card in Slot 17. The redundant line card for this configuration would be placed in Slot 33. This establishes a directly mapped communication path through the chassis midplane between the application and line cards.

To help identify which rear slot corresponds with the front-loaded application card, the upper rear slot numbers are equal to the slot number of the front-loaded card plus 16. For example, to insert a line card to support an application card installed in slot 1, add 16 to the slot number of the front-loaded application card (Slot 1 + 16 slots = Slot 17). Slot 17 is the upper right-most slot on the rear of the chassis, directly behind Slot 1.

For lower rear slot numbers, add 32. Again, a redundant line card for an application card in Slot 1 would be (Slot 1 + 32 = Slot 33). Slot 33 is the lower right-most slot on the rear of the chassis, also behind Slot 1.

Rear Slot Numbering with Full-height Line Cards

ASR 5000 systems may be configured with 10 Gigabit Ethernet Line Cards (XGLCs). These are full-height line cards that require the removal of the half-height card guides to accommodate the cards. In this case, only the upper slot number is used to refer to the XGLC. For half-height cards installed with the XGLCs, the half-height slot numbering scheme is maintained.

For example, XGLCs installed in slots 17 and 32 also take up slots 33 and 48, but are referred to as cards in slots 17 and 32 only. The slots in which the SPIOs and RCCs are installed in the same configuration, are slots 24 and 25, and 40 and 41, respectively.

Mounting Options

The chassis is designed for installation in a standard (EIA-310-D, IEC 60297) 19-inch wide (482.6 mm) equipment rack or telco cabinet. Additional rack hardware (such as extension brackets) may be used to install the chassis in a 23-inch (584.2 mm) rack. Each chassis is 24.50 inches (62.23 cm) high. This equates to 14 Rack Units (1 RU = 1.75 in. [44.5 mm]).

You can mount a maximum of three ASR 5000 chassis in a 2- or 4-post equipment rack, or telco cabinet, provided that all system cooling and ventilation requirements are met. Three stacked chassis will occupy a minimum of 42 RUs.

There are two options for mounting the chassis in an equipment rack or telco cabinet:

- **Flush mount**: In this configuration, the flanges of the mounting brackets are flush with the front of the chassis. This method is typically used with 4-post racks and telco equipment cabinets. This is the default configuration as shipped.
- **Mid-mount**: In this configuration, the flanges of the mounting brackets are recessed from the front of the chassis. This method is typically used with 2-post racks. You must remove and re-install the mounting brackets in the middle of the chassis on both sides.

Caution: The equipment rack or cabinet hardware must not hinder air flow at any of the intake or exhaust vents. The ambient environment (conditioned space) must allow the system to function within its specified operating limits.

Midplane Architecture

The midplane separates the front and rear chassis slots. The connectors on the midplane provide intra-chassis communications, power connections, and data transport paths between the various installed cards.

The midplane also contains two independent -48 VDC busses (not shown) that distribute redundant power to each card within the chassis.

Figure 5. Midplane/Switch Fabric Architecture

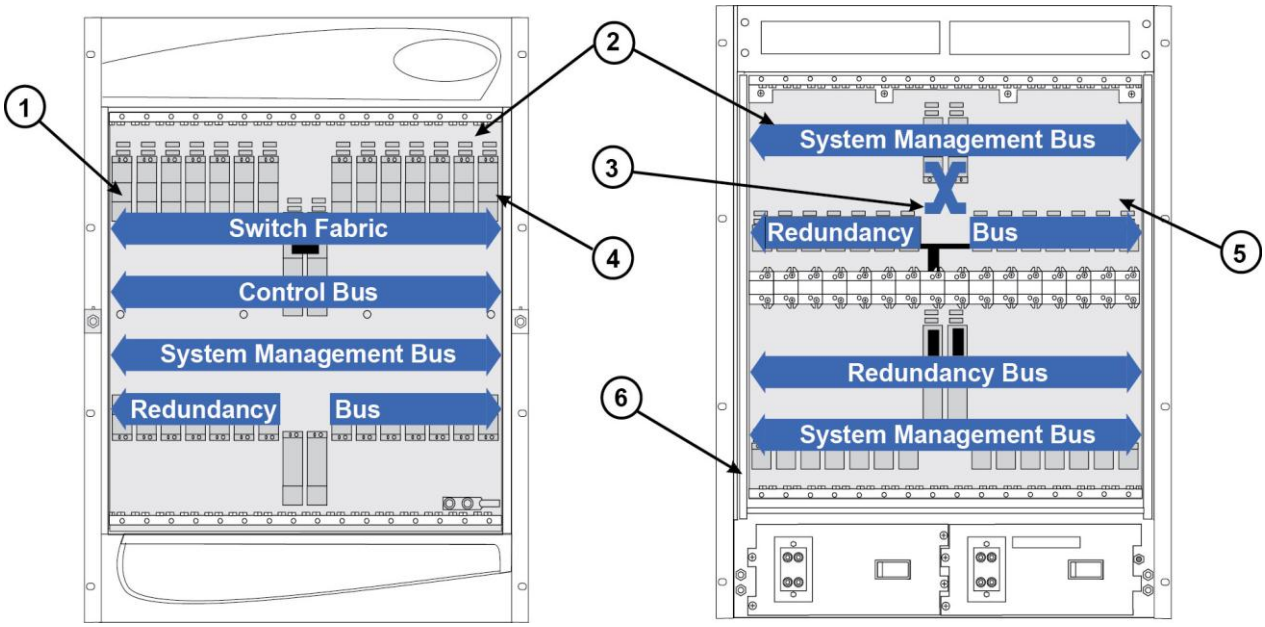


Table 6. Midplane and Bus Descriptions

Item	Description
1	Slot number 1 (left-most application card slot)
2	Chassis midplane: provides intra-chassis communications and data transport paths between the various installed cards
3	SPIO cross-connect bus
4	Chassis slot number 16: right-most application card slot
5	Chassis slot number 17: upper right-most line card slot. The 10 Gigabit Ethernet Line Card (XGLC) is a full-height line card that takes up the upper and lower slots in the back of the chassis. Use the upper slot number only when referring to installed XGLCs. Slot numbering for other half-height lines cards is maintained: 17 to 32 and 33 to 48, regardless of the number of installed XGLCs.

Item	Description
6	Chassis slot number 48: lower left-most line card slot

The following subsections describe each bus.

320 Gbps Switch Fabric

The System Management Card (SMC) is an IP-based (packetized) switch fabric that provides a transport path for user data throughout the system. Its 320 Gbps switch fabric establishes inter-card communication between the SMCs and other application cards within the chassis along with their associated line cards.

32 Gbps Control Bus

The Control Bus features redundant 32 Gbps Ethernet paths that interconnect all control and management processors within the system. The bus uses a full-duplex Gigabit Ethernet (GigE) switching hierarchy from both SMCs to each of the 14 application card slots in the chassis. Each application card is provisioned with a GigE switch to meet its specific needs. This bus also interconnects the two SMC modules.

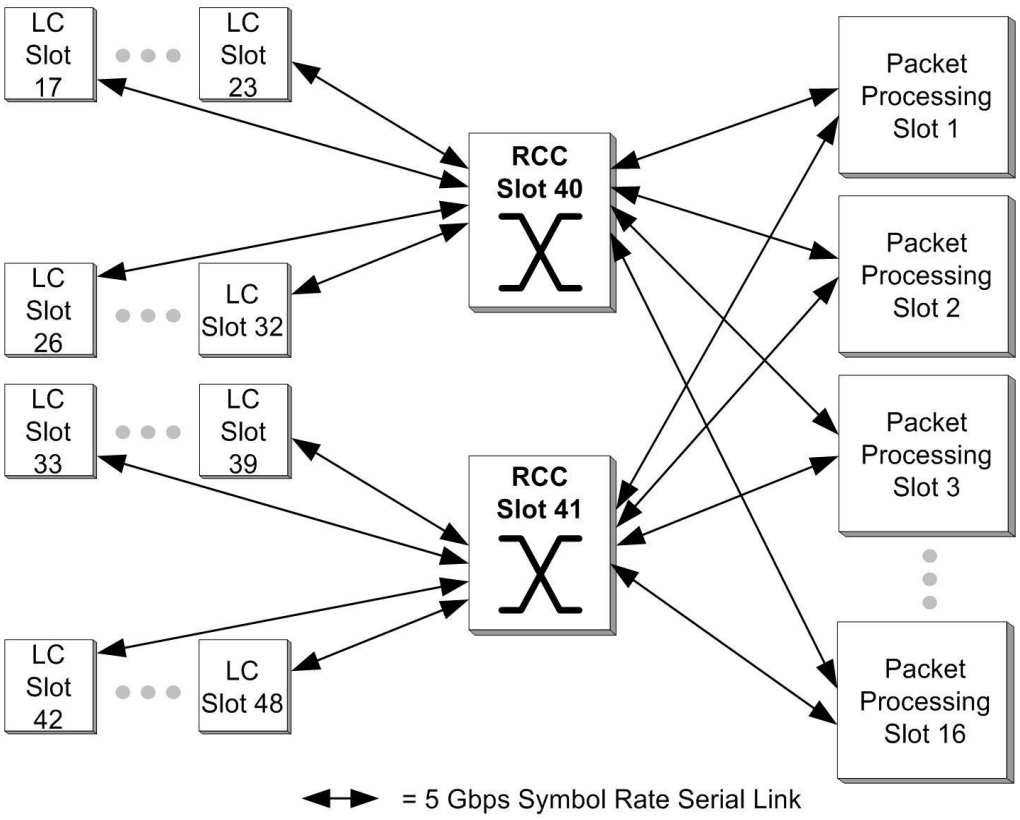
System Management Bus

The System Management Bus supports management access to each component within the chassis. It provides a communication path from each SMC to every card in the system with a 1 Mbps transfer rate to each card. This allows the SMCs to manage several low-level system functions (such as, supplying power, monitoring temperature, board status, pending card removals, data path errors, redundant/secondary path switchovers, card resets and failovers). Additionally, the System Management Bus monitors and controls the fan trays, PFUs, and alarming functions.

280 Gbps Redundancy Bus

The Redundancy Bus consists of multiple, full-duplex serial links providing packet processing card-to-line card redundancy through the chassis' Redundancy Crossbar Cards (RCCs) as shown below.

Figure 6. RCC Logical View



Each RCC facilitates 28 links:

- One link with each of the 14 packet processing card slots
- One link with each of the line card slots
  - The RCC in slot 40 supports line card slots 17-23 and 26-32 (upper-rear slots)
  - The RCC in slot 41 supports line card slots 33-39 and 42-48 (lower-rear slots)

Each serial link facilitates up to 5 Gbps symbol rate, equivalent to 4 Gbps of user data traffic, in each direction. Therefore, the Redundancy Bus provides 140 Gbps symbol rate (112 Gbps user data) of throughput per RCC, 280 Gbps symbol rate (224 Gbps user data) total for both.

OC-48 TDM Bus

The system also hosts a dual OC-48 TDM bus consisting of 128 independent TDM paths each consisting of 512 DS0 channels. This bus supports voice services on the system. Higher speed TDM traffic requirements are addressed using the system's data fabric.

SPIO Cross-Connect Bus

To provide redundancy between Switch Processor I/O (SPIO) cards, the system possesses a physical interconnect between the ports on the SPIOs. This cross-connect allows management traffic or alarm outputs to be migrated from an active SPIO experiencing a failure to the redundant SPIO.

While an SPIO should be installed directly behind its corresponding SMC, this bus allows either SMC to utilize either SPIO.

Power Filter Units

Located at the bottom rear of the chassis are slots for two 165-amp Power Filter Unit (PFU) assemblies. Each PFU provides DC power from the site's power distribution frame (PDF) to the chassis and its associated cards. Each load-sharing PFU operates independently of the other to ensure maximum power feed redundancy. The maximum input operating voltage range of the PFU is -40 VDC to -60 VDC; the nominal range is -48 VDC to -60 VDC.

Important: The ASR 5000 does not offer an AC power supply option. If only AC power is available at the installation site, an adequately sized AC-to-DC converter will be required to supply -48 VDC power to the chassis.

The following drawing shows the PFU and its connectors.

Figure 7. PFU Components

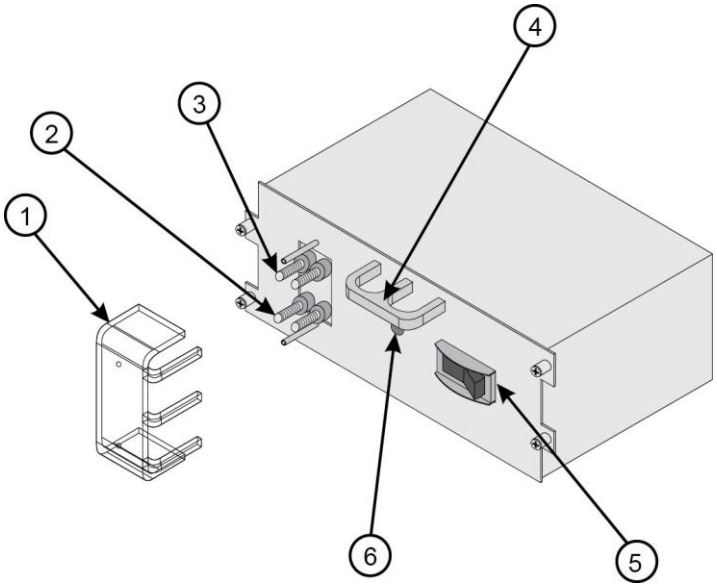


Table 7. Power Filter Unit Component Descriptions

Item	Description
1	Plastic terminal cover
2	VDC (-48 VDC input terminals)
3	RTN (voltage return terminals)
4	PFU handle
5	Circuit breaker (On/Off) rated at 165A
6	Power LED

Fan Tray Assemblies

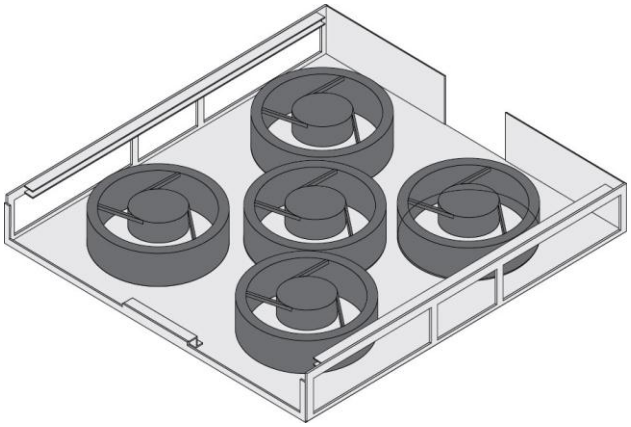
There are two fan tray assemblies within the chassis. A lower fan tray intakes ambient air and an upper fan tray exhausts warmed air from the chassis. Each fan tray is connected to both PFUs to ensure power feed redundancy. Both fan tray assemblies are variable speed units that automatically adjust fan speed based on temperature or failover situations.

Thermal sensors monitor temperatures within the chassis. In the event of a fan failure or other temperature-related condition, the SMC notifies all operable fans in the system to switch to high speed, and generates an alarm.

Lower Fan Tray

The lower fan tray assembly contains multiple fans and pulls ambient air into the chassis from the lower front and sides of the chassis. The air is then pushed upward across the cards and midplane to support vertical convection cooling.

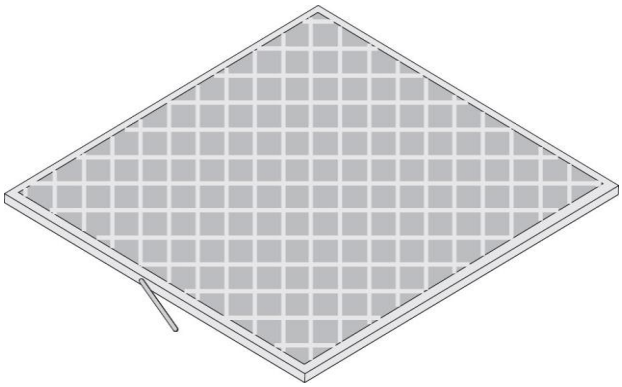
Figure 8. Lower Fan Tray



Air Filter Assembly

The chassis supports a replaceable particulate air filter that meets UL 94-HF-1 standards for NEBS-compliant electronics filtering applications. This filter mounts above the lower fan tray assembly and removes contaminants before they enter the system. Temperature sensors measure the temperature at various points throughout the chassis. The system monitors this information, and generates a maintenance alarm, if necessary.

Figure 9. Particulate Air Filter

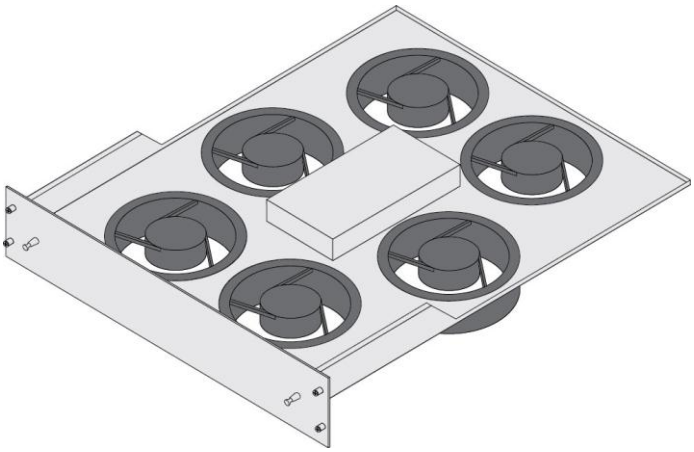


Important: A replacement air filter is shipped with each chassis. A minimum of one replacement air filter for each deployed chassis should be kept on site. This allows qualified service personnel to quickly replace the filter when necessary.

Upper Fan Tray

The upper fan tray unit contains multiple fans that exhaust air from the upper rear and sides of the chassis.

Figure 10. Upper Fan Tray



Chassis Airflow

Airflow within the chassis complies with Telcordia recommendations to ensure proper vertical convection cooling of the system.

Application Cards

The following application cards are supported by the system.

System Management Card (SMC)

The SMC serves as the primary system controller, initializing the entire system and loading the software's configuration image into other cards in the chassis as applicable.

SMCs are installed in the chassis slots 8 and 9. During normal operation, the SMC in slot 8 serves as the primary card and the SMC in slot 9 serves as the secondary. Each SMC has a dual-core central processing unit (CPU) and 4 GB of random access memory (RAM).

There is a single PC-card slot on the front panel of the SMC that supports removable ATA Type I or Type II PCMCIA cards. Use these cards to load and store configuration data, software updates, buffer accounting information, and store diagnostic or troubleshooting information.

There is also a Type II CompactFlash™ slot on the SMC that hosts configuration files, software images, and the session limiting/feature use license keys for the system.

The SMC provides the following major functions:

- Non-blocking low latency inter-card communication
- 1:1 or 1:N redundancy for hardware and software resources
- System management control
- Persistent storage via CompactFlash and PCMCIA cards (for field serviceability), and a hard disk drive for greater storage capabilities
- Internal gigabit Ethernet switch fabrics for management and control plane communication

The front panel of the SMC with its major components is shown below:

Figure 11. System Management Card (SMC)

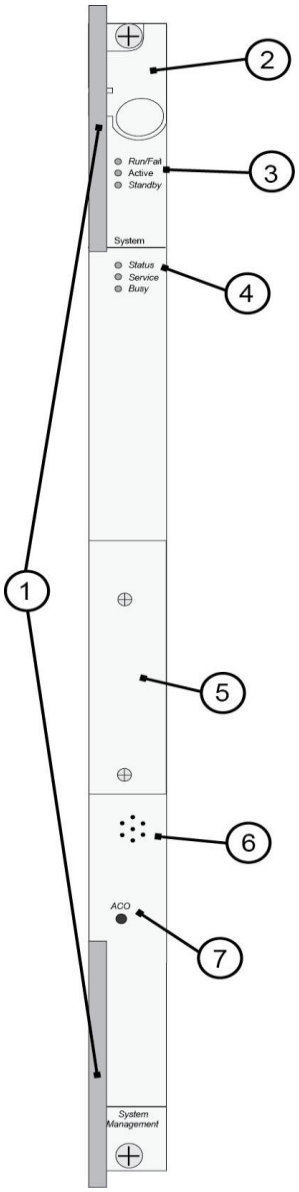


Table 8. SMC Callouts

Item	Description
1	**Card Ejector Levers** —Use to insert/remove card to/from chassis.
2	**Interlock Switch** —Sliding this switch downward on an active SMC initiates an immediate switchover to the standby SMC.
3	**Card Level Status LEDs**—Show the status of the card.
4	**System Level Status LEDs**—Show the status of overall system health and/or maintenance requirements.

Item	Description
5	**PC-Card/PCMCIA Slot**—Stores or moves software, diagnostics, and other information.
6	**System Alarm Speaker**—Sounds an audible alarm when specific system failures occur.
7	**Alarm Cut-Off (ACO)** —Press and release this recessed toggle switch to reset the system alarm speaker and other audible or visual alarm indicators connected to the CO Alarm interface on the SPIO.

SMC RAID Support

Each SMC is equipped with a hard disk, commonly referred to as a Small Form Factor (SFF) disk.

Important: The hard disk is not physically accessible. Disk failure constitutes an SMC failure.

If there is a redundant SMC in the chassis, the standby disk mirror the disk in the active SMC, forming an active Redundant Array of Inexpensive Disks (RAID).

HD RAID ids configurable via CLI commands. RAID control mechanisms allow xDR charging data to be written to the hard disks on both the active and standby SMCs for later upload to a suitable local or remote storage server.

Event logs related to disk and RAID include disk name, serial number and RAID UUID for reference. They are generated at the Critical, Error, Warning, and Informational levels.

Event logs at the Critical level are generated for service-affecting events such as:

- RAID failure, including failures during runtime and various cases of initial RAID discovery and disk partition failures
- File system failure when the system fails to initialize or mount file systems
- Network failure for NFS server-related errors

Event logs at the Error level are generated for important failures:

- RAID disk failure, including failures during runtime
- Internal errors, including forking process failures

Event logs at Warning level are generated for important abnormal cases:

- Overwriting a valid or invalid disk partition, RAID image, and file system
- RAID construction in progress and possible failure
- Low disk space
- Files deleted to free up disk space

Event logs at the Informational level are generated for normal situations:

- Disk partition completion
- RAID discovery results without overwriting
- RAID construction completion
- RAID disk added or removed
- File system initialization
- NFS service start

- Files copied/removed from CDR module to RAID disk

The hard disk supports SNMP notifications which are described in the *SNMP MIB Reference*.

Packet Processing Cards: PSC, PSCA, PSC2, PSC3 and PPC

The packet processing cards provide packet processing and forwarding capabilities within a system. Each card type supports multiple contexts, which allows an operator to overlap or assign duplicate IP address ranges in different contexts.

Caution: With the exception of the PSC and PSCA, you cannot mix packet processing card types (PSCs, PSCAs, PSC2s, PSC3s or PPCs) in the same chassis.

Specialized hardware engines support parallel distributed processing for compression, classification, traffic scheduling, forwarding, packet filtering, and statistics.

The packet processing cards use control processors to perform packet-processing operations, and a dedicated high-speed network processing unit (NPU). The NPU does the following:

- Provides "Fast-path" processing of frames using hardware classifiers to determine each packet's processing requirements
- Receives and transmits user data frames to and from various physical interfaces
- Performs IP forwarding decisions (both unicast and multicast)
- Provides per interface packet filtering, flow insertion, deletion, and modification
- Manages traffic and traffic engineering
- Modifies, adds, or strips datalink/network layer headers
- Recalculates checksums
- Maintains statistics
- Manages both external line card ports and the internal connections to the data and control fabrics

To take advantage of the distributed processing capabilities of the system, you can add packet processing cards to the chassis without their supporting line cards, if desired. This results in increased packet handling and transaction processing capabilities. Another advantage is a decrease in CPU utilization when the system performs processor-intensive tasks such as encryption or data compression.

Packet processing cards can be installed in chassis slots 1 through 7 and 10 through 16. Each card can either Active (available to the system for session processing) or redundant (a standby component available in the event of a failure).

The front panel of a packet processing card with its major components is shown below.

Figure 12. Packet Processing Card (Generic)

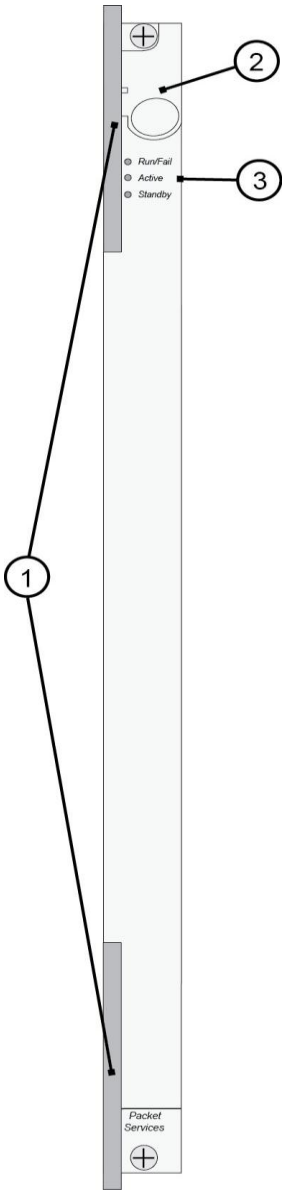


Table 9. Packet Processing Card (Generic) Callouts

Item	Description
1	**Card Ejector Levers** —Use to insert/remove card to/from chassis.
2	**Interlock Switch** —In its Down position the interlock switch notifies the system to safely power down the card prior to its removal.
3	**Card Level Status LEDs** —Show the current status of the card.
4	**Card Identification Label** —Indicates the type of packet processing card. See the table at the end of this chapter.

Packet Services Card (PSC)

Each PSC has two x86-based control processor (CP) subsystems that perform the bulk of the packet-based user service processing. The main x86 CP contains four cores split across two chips. It is equipped with 16 GB of RAM. Therefore, a fully-loaded system consisting of 14 PSCs, provides 224 GB of RAM dedicated to packet processing tasks. The second CP on the PSC is in the NPU. This CP contains 1.5 GB of memory, but only 512 MB is available to the OS for use in session processing. Hardware encryption devices are standard on the PSC.

Packet Services Card Type A (PSCA)

The PSCA is a direct replacement for the PSC and can be mixed with PSCs in the same chassis. It is an economical alternative for customers who wish to add more processing power to their ST40 or ASR 5000 chassis. Although the PSCA has a higher throughput capacity than the PSC, when mixed with PSCs, the PSCA will not overrun the PSCs. The PSCA can also be used to populate a new chassis for higher throughput than a PSC but for less cost and throughput than a PSC2.

Important: When PSCs and PSCAs are mixed in a chassis, incompatible components prevent the running of crypto functions. The chassis must also be configured to run in hybrid mode before PSCAs will be allowed to boot.

Packet Services Card 2 (PSC2)

The PSC2 provides increased aggregate throughput and performance, and a higher number of subscriber sessions in comparison with the PSC or PSCA.

The PSC2 uses a faster network processor unit, featuring two quad-core x86 CPUs and 32 GB of RAM. These processors run a single copy of the operating system. The operating system running on the PSC2 treats the two dual-core processors as a 4-way multi-processor.

The PSC2 has a dedicated security processor that provides the highest performance for cryptographic acceleration of next-generation IP Security (IPSec), Secure Sockets Layer (SSL) and wireless LAN/WAN security applications with the latest security algorithms.

PSC2s should not be mixed with PSCs, PSCAs, PSC3s. or PPCs. Due to the different processor speeds and memory configurations, the PSC2 cannot be combined in a chassis with other packet processing card types.

The PSC2 can dynamically adjust the line card connection mode to support switching between XGLCs and non-XGLCs with minimal service interruption.

PSC2 is fully redundant with a spare PSC2.

Important: ICSR is not supported between a chassis using PSC2s and a chassis using PSCs, PSC3s or PPCs due to the different capabilities of the two chassis.

Packet Services Card 3 (PSC3)

The PSC3 provides increased aggregate throughput and performance and a higher number of subscriber sessions than other ASR 5000 packet processing cards. Specialized hardware engines support parallel distributed processing for compression, classification, traffic scheduling, forwarding, packet filtering, and statistics.

The PSC3 features two 6-core CPUs and 64 GB of RAM. These processors run a single copy of the operating system. The operating system running on the PSC3 treats the two core processors as a 6-way multi-processor.

To optimize network efficiency and minimize down time, the system supports 1:n redundancy for PSC3s. If session recovery is enabled, the minimum number of PSC3s per chassis increases from one to four cards. Three PSC3s are active and one PSC3 is standby (redundant). This minimum configuration protects against software failures only. In addition to increased hardware requirements, Session Recovery may reduce subscriber session capacity, performance, and data throughput.

In the event of PSC3 failure, tasks are migrated from the active PSC3 to the standby card. The line card installed behind the PSC3 that was formerly active maintains the interfaces to the external network equipment. Redundancy Crossbar Cards (RCCs) provide a path for signaling and data traffic between the line card and the now active packet processing card.

PSC3s must not be mixed with PSCs, PSCAs, PSC2s. or PPCs.

The PSC3 is fully redundant with a spare PSC3.

Packet Processor Card (PPC) Description

The PPC is an economical alternative for smaller operators who do not require the throughput of PSC2s or PSC3s. It features a quad-core x86 CPU and 16GB of RAM. The processor runs a single copy of the operating system. The operating system running on the PPC treats the dual-core processor as a two-way multi-processor. The PPC supports CDMA, HA, and GGSN functionality.

A second-generation data transport fixed programmable gate array (FPGA) connects the PPC NPU bus to the switch fabric interface. The FPGA also provides a bypass path between the line card or Redundancy Crossbar Card (RCC) and the switch fabric for ATM traffic. Traffic from the line cards or the RCC is received over the FPGA serial links and is sent to the NPU on its switch fabric interface. The traffic destined for the line cards or RCC is diverted from the NPU interface and sent over the serial links.

A PCI-E interface allows the control processors to perform register accesses to the FPGA and some components attached to it, and also allows DMA operations between the NPU and the control processors' memory. A statistics engine is provided in the FPGA. Two reduced latency DRAM (RLDRAM) chips attached to the FPGA provide 64 MB of storage for counters.

The PPC has a dedicated security processor that provides the highest performance for cryptographic acceleration of next-generation IP Security (IPsec), Secure Sockets Layer (SSL) and wireless LAN/WAN security applications with the latest security algorithms.

PPCs must not be mixed with PSCs, PSCAs, PSC2s, or PSC3s.

The PPC is fully redundant with a spare PPC.

Line Cards

The following rear-loaded cards are currently supported by the system.

Switch Processor I/O (SPIO) Card

The SPIO card provides connectivity for local and remote management, CO alarming, and Building Integrated Timing Supply (BITS) timing input. SPIOs are installed in chassis slots 24 and 25, behind SMCs. During normal operation, the SPIO in slot 24 works with the active SMC in slot 8. The SPIO in slot 25 serves as a redundant component. In the event that the SMC in slot 8 fails, the redundant SMC in slot 9 becomes active and works with the SPIO in slot 24. If the SPIO in slot 24 should fail, the redundant SPIO in slot 25 takes over.

The following shows the front panel of the SPIO card, its interfaces, and other major components.

Figure 13. Switch Processor I/O (SPIO) Card

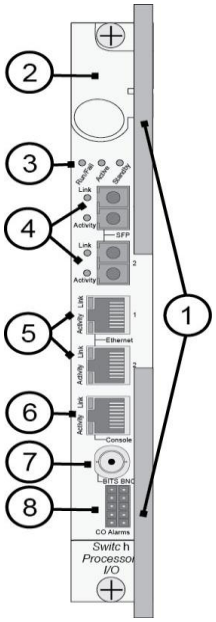


Table 10. SPIO Callouts

Item	Description
1	**Card Ejector Levers**—Use to insert/remove card to or from the chassis.
2	**Interlock Switch**—In its Down position the interlock switch notifies the system to safely power down the card prior to its removal.
3	**Card Level Status LEDs**—Show the status of the card.
4	**Optical Gigabit Ethernet Management LAN Interfaces**—Two Small Form-factor Pluggable (SFP) optical Gigabit Ethernet interfaces to connect optical transceivers.

Item	Description
5	**10/100/1000 Mbps Ethernet Management LAN Interfaces**—Two RJ-45 interfaces, supporting 10/100 Mbps or 1 Gbps Ethernet.
6	**Console Port**—RJ-45 interface used for local connectivity to the command line interface (CLI).
7	**BITS Timing Interface** (Optional)—Either a BNC interface or 3-pin wire wrap connector. Used for application services that use either the optical or channelized line cards.
8	**CO Alarm Interface**—Dry contact relay switches, allowing connectivity to central office, rack, or cabinet alarms.

Management LAN Interfaces

SPIO management LAN interfaces connect the system to the carrier's management network and applications, normally located remotely in a Network Operations Center (NOC). You can use the RJ-45 copper 10/100/1000 Mbps Ethernet interfaces or optical SFP Gigabit Ethernet interfaces to connect to the management network.

When using the RJ-45 interfaces, use CAT5 shielded twisted pair (STP) cabling.

Important: Use shielded cabling whenever possible to further protect the chassis and its installed components from ESD or other transient voltage damage.

Table 11. SFP Interface Supported Cable Types

Module Type	Card Identification	Interface Type	Cable Specifications
1000Base-SX	Ethernet 1000 SX	Fiber, LC duplex female connector	**Fiber Type:** Multi-mode fiber (MMF), 850 nm wavelength **Core Size (microns)/Range:** • 62.5/902.23 feet (275 meters) • 50/1640.42 feet (500 meters) **Minimum Tx Power:** -9.5 dBm **Rx Sensitivity:** -17 dBm

Console Port

The console uses an RS-232 serial communications port to provide local management access to the command line interface (CLI). A 9-pin-to-RJ-45 console cable is supplied with each SPIO card. The console cable must provide carrier-detect when attached in a null modem configuration.

Should connection to a terminal server or other device requiring a 25-pin D-subminiature connector be required, a specialized cable can be constructed to support DB-25 to RJ-45 connectivity. The baud rate for this interface is configurable between 9600 bps and 115,200 bps (default is 9600 bps).

BITS Timing

The Building Integrated Timing Supply (BITS) timing interface is optional and required only when the system is used in support of non-data applications. A BITS module is available on two versions of the SPIO: one supports a BNC interface and the other a 3-pin interface. If your system uses optical and/or channelized line cards for SDH/SONET, you can configure it to have the BITS module provide the transmit timing source, compliant with Stratum 3 requirements, for all the line cards in the chassis.

Central Office Alarm Interface

The CO alarm interface is a 10-pin connector for up to three dry-contact relay switches for connection to a CO alarm monitoring panel. The three Normally Closed alarm relays can be wired to support Normally Open or Normally Closed devices, indicating minor, major, and critical alarms.

A CO alarm cable is shipped with the product so you can connect the CO Alarm interfaces on the SPIO card to your alarming devices. The "Y" cable design ensures CO alarm redundancy by connecting to both primary and secondary SPIO cards.

Redundancy Crossbar Card (RCC)

The RCC uses 5 Gbps serial links to ensure connectivity between rear-mounted line cards and every non-SMC front-loaded application card slot in the system. This creates a high availability architecture that minimizes data loss and ensures session integrity. If a packet processing card were to experience a failure, IP traffic would be redirected to and from the LC to the redundant packet processing card in another slot. Each RCC connects up to 14 line cards and 14 packet processing cards for a total of 28 bidirectional links or 56 serial 2.5 Gbps bidirectional serial paths.

The RCC provides each packet processing card with a full-duplex 5 Gbps link to 14 (of the maximum 28) line cards placed in the chassis. This means that each RCC is effectively a 70 Gbps full-duplex crossbar fabric, giving the two RCC configuration (for maximum failover protection) a 140 Gbps full-duplex redundancy capability.

The RCC located in slot 40 supports line cards in slots 17 through 23 and 26 through 32 (upper rear slots). The RCC in slot 41 supports line cards in slots 33 through 39 and 42 through 48 (lower rear slots):

Figure 14. Redundancy Crossbar Card (RCC)

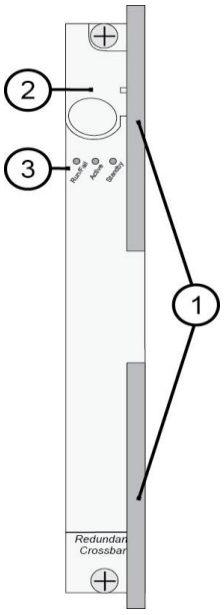


Table 12. RCC Callouts

Item	Description
1	**Card Ejector Levers**—Use to insert/remove a card to and from the chassis.
2	**Interlock Switch**—In its Down position the interlock switch notifies the system to safely power down the card prior to its removal.
3	**Card Level Status LEDs**—Show the status of the card.

Fast Ethernet Line Card (FELC or FLC2)

The FELC/FLC2 installs directly behind its respective packet processing card, providing network connectivity to the RAN interface and the packet data network. Each FELC/FLC2 (Ethernet 10/100) has eight RJ-45 interfaces. Each of these IEEE 802.3-compliant interfaces supports auto-sensing 10/100 Mbps Ethernet. Allowable cabling includes:

- 100Base-Tx—full or half duplex Ethernet on CAT 5 shielded twisted pair (STP) or unshielded twisted pair (UTP) cable
- 10Base-T—full or half duplex Ethernet on CAT 3, 4, or 5 STP or UTP cable

Important: Use shielded cabling whenever possible to further protect the chassis and its installed components from ESD or other transient voltage damage.

The FLC2 supports the Star Channel (1 Gbps) for faster FPGA upgrades and is Restriction of Hazardous Substances (RoHS) 6/6 compliant.

The FELC/FLC2 can be installed in chassis slots 17 through 23, 26 through 39, and 42 through 48. These cards are always installed directly behind their respective packet processing cards, but are not required to be placed behind any redundant packet processing cards (those operating in Standby mode).

The following shows the panel of the FELC/FLC2 with its interfaces and major components.

Figure 15. Fast Ethernet Line Card (FELC/FLC2)

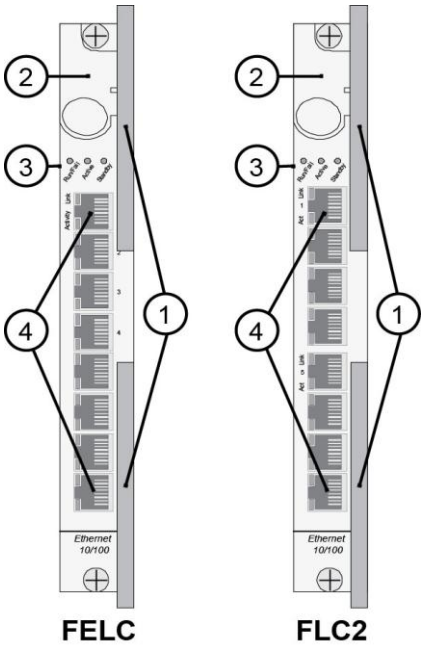


Table 13. FELC/FLC2 Callouts

Item	Description
1	**Card Ejector Levers**—Use to insert/remove card to/from chassis.
2	**Interlock Switch**—In its Down position the interlock switch notifies the system to safely power down the card prior to its removal.
3	**Card Level Status LEDs—**Show the status of the card.
4	**RJ-45 10/100 Ethernet Interfaces**—Eight auto-sensing RJ-45 interfaces for R-P interface connectivity carrying user data. Ports are numbered 1 through 8 from top to bottom.

Gigabit Ethernet Line Card (GELC or GLC2)

The GELC/GLC2 installs directly behind its respective packet processing card, providing network connectivity to the packet data network. The GELC/GLC2 (Ethernet 1000) supports a variety of 1000 Mbps optical and copper interfaces based on the type of Small Form-factor Pluggable (SFP) modules installed on the card.

Table 14. SFP Modules Supported by the GELC/GLC2

Module Type	Card Identification	Interface Type	Cable Specifications
1000Base-SX	Ethernet 1000 SX	Fiber, LC duplex female connector	**Fiber Type:** Multi-mode fiber (MMF), 850 nm wavelength **Core Size (microns)/Range:** • 62.5/902 feet (275 meters) • 50/1640 feet (500 meters) **Minimum Tx Power:** -9.5 dBm **Rx Sensitivity:** -17 dBm
1000Base-LX	Ethernet 1000 LX	Fiber, LC duplex female connector	**Fiber Type:** Single-mode fiber (SMF), 1310 nm wavelength **Core Size (microns)/Range:** 9/32808 feet (10 kilometers) **Minimum Tx Power:** -9.5 dBm **Rx Sensitivity:** -19 dBm
1000Base-T	Ethernet 1000 Copper	RJ-45	Operates in full-duplex up to 100 meters of CAT-5 Shielded Twisted Pair (STP) cable with BER less than 10e-10.

Important: This product has been tested and found to comply with the limits for Class 1 laser devices for IEC825, EN60825, and 21CFR1040 specifications.

The GLC2 supports the Star Channel (1 Gbps) for faster FPGA upgrades and is Restriction of Hazardous Substances (RoHS) 6/6 compliant.

The GELC/GLC2s can be installed in chassis slots 17 through 23, 26 through 39, and 42 through 48. These cards are always installed directly behind their respective or packet processing cards, but they are not required behind any redundant packet processing cards (those operating in Standby mode).

The following diagram shows the front panel of the GELC/GLC2 with an optical connector, identifying its interfaces and major components.

Figure 16. Gigabit Ethernet Line Card (GELC/GLC2)

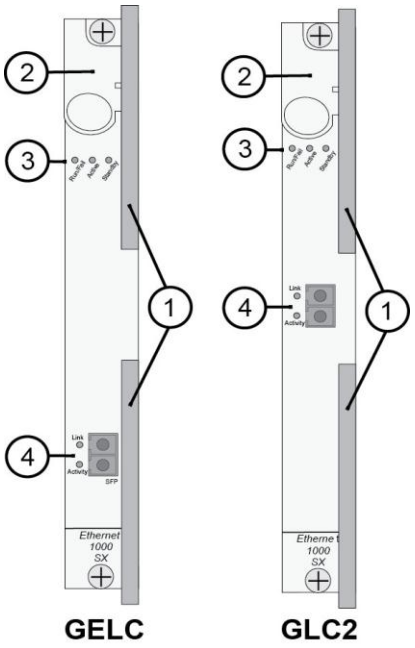


Table 15. GELC/GLC2 Callouts

Item	Description
1	**Card Ejector Levers**—Use to insert/remove card to/from chassis.
2	**Interlock Switch**—In its Down position, the interlock switch notifies system to safely power down card prior to removal.
3	**Card Level Status LEDs**—Show the status of the card. See *Applying Power and Verifying Installation* for definitions.
4	**Gigabit Ethernet Interface**—Gigabit Ethernet (GE) SFP modules. 1000Base-SX, 1000Base-LX, and 1000Base-T interfaces are supported depending on the SFP module installed.

Quad Gigabit Ethernet Line Card (QGLC)

The QGLC is a 4-port Gigabit Ethernet line card that installs directly behind its associated packet processing card to provide network connectivity to the packet data network. There are several different versions of Small Form-factor Pluggable (SFP) modules available for the QGLC.

Table 16. SFP Modules Supported by the QGLC

Module Type	Card Identification	Interface Type	Cable Specifications

Module Type	Card Identification	Interface Type	Cable Specifications
1000Base-SX	Ethernet 1000 SX	Fiber, LC duplex female connector	**Fiber Type:** Multi-mode fiber (MMF), 850 nm wavelength **Core Size (microns)/Range:** • 62.5/902 feet (275 meters) • 50/1640 feet (500 meters) **Minimum Tx Power:** -9.5 dBm **Rx Sensitivity:** -17 dBm
1000Base-LX	Ethernet 1000 LX	Fiber, LC duplex female connector	**Fiber Type:** Single-mode fiber (SMF), 1310 nm wavelength **Core Size (microns)/Range:** 9/32808 feet (10 kilometers) **Minimum Tx Power:** -9.5 dBm **Rx Sensitivity:** -19 dBm
1000Base-T	Ethernet 1000 Copper	RJ-45	Operates in full-duplex up to 100 meters of CAT-5 Shielded Twisted Pair (STP) cable with BER less than 10e-10.

Important: This product has been tested and found to comply with the limits for Class 1 laser devices for IEC825, EN60825, and 21CFR1040 specifications.

The QGLC supports the Star Channel (1 Gbps) for faster FPGA upgrades and is Restriction of Hazardous Substances (RoHS) 6/6 compliant.

Install QGLCs in chassis slots 17 through 23, 26 through 39, and 42 through 48. Always install these cards directly behind their respective packet processing cards. They are not required behind any redundant packet processing cards (those operating in Standby mode).

The following shows the front panel of the QGLC, identifying its interfaces and major components:

Figure 17. Quad Gigabit Ethernet Line Card (QGLC)

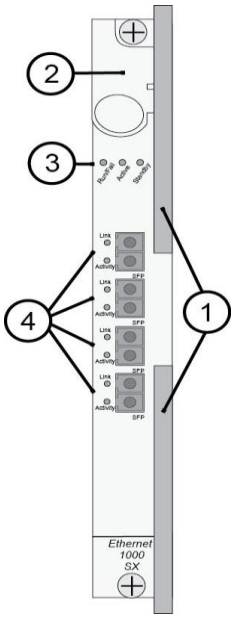


Table 17. QGLC Callouts

Item	Description
1	**Card Ejector Levers**—Use to insert/remove card to/from chassis.
2	**Interlock Switch**—In its Down position the interlock switch notifies system to safely power down the card prior to its removal.
3	**Card Level Status LEDs**—Show the status of the card. See *Applying Power and Verifying Installation* for definitions.
4	**Gigabit Ethernet Interface(s)** —Gigabit Ethernet (GE) SFP modules. 1000Base-SX, 1000Base-LX, and 1000Base-T interfaces are supported depending on the SFP module installed.

10 Gigabit Ethernet Line Card (XGLC)

The XGLC supports higher speed connections to packet core equipment, increases effective throughput between the ASR 5000 and the packet core network, and reduces the number of physical ports needed on the ASR 5000.

The XGLC (10G Ethernet) is a full-height line card, unlike the other line cards, which are half height. To install an XGLC, you must remove the half-height card guide to create a full-height slot.

The single-port XGLC supports the IEEE 802.3-2005 revision which defines full duplex operation of 10 Gigabit Ethernet. PSC2s or PSC3s are required to achieve maximum sustained rates with the XGLC.

The XGLC use a Small Form Factor Pluggable Plus (SFP+) module. The modules support one of two media types: 10GBASE-SR (Short Reach) 850nm, 300m over multimode fiber (MMF), or 10GBASE-LR (Long Reach) 1310nm, 10km over single mode fiber (SMF).

The XGLC is configured and monitored by the SMC via the system's control bus. If the firmware needs to be upgraded. the XGLC uses the Star Channel for a faster download.

Install XGLCs in chassis slots 17 through 23 and 26 through 32. These cards should always be installed directly behind their respective packet processing cards, but they are not required behind any redundant packet processing cards (those operating in Standby mode).

The supported redundancy schemes for XGLC are L3, Equal Cost Multi Path (ECMP) and 1:1 side-by-side redundancy.

Side-by-side redundancy allows two XGLC cards installed in neighboring slots to act as a redundant pair. Side-by-side pair slots are 17-18, 19-20, 21-22, 23-26, 27-28, 29-30, and 31-32.

Side-by-side redundancy only works with XGLC cards. When configured for non-XGLC cards, the cards are brought offline. If the XGLCs are not configured for side-by-side redundancy, they run independently without redundancy.

When you first configure side-by-side redundancy, the higher-numbered slot's configuration is erased and then duplicated from the lower-numbered slot. The lower-numbered top slot retains all other configuration settings. While side by side redundancy is configured, all other configuration commands work as if the side by side slots were top-bottom slots. Configuration commands directed at the bottom slots either fail with errors or are disallowed.

When you unconfigure side-by-side redundancy, the configuration for the higher-numbered top and bottom slots are initialized to the defaults. The configuration for the lower-numbered slot retains all other configuration settings. If you install non-XGLC cards in the slots, you may bring them back online.

Table 18. SFP Modules Supported by the XGLC

Module Type	Card Identification	Interface Type	Cable Specifications
10GBase-SR	Ethernet 10G SR	Fiber, LC duplex female connector	**Fiber Type:** Multi-mode fiber (MMF), 850 nm wavelength **Core Size (microns)/Range:** • 62.5/902.23 feet (275 meters) • 50/1640.42 feet (500 meters) • 62.5um/33m (OM1) • 50um 500MHz-km/82m (OM2) • 50um 2000MHz-km/300m (OM3) **Minimum Tx Power:** -7.3 dBm **Rx Sensitivity:** -11.1 dBm
10GBase-LR	Ethernet 10G LR	Fiber, LC duplex female connector	**Fiber Type:** Single-mode fiber (SMF), 1310 nm wavelength **Core Size (microns)/Range:** 9/32808.4 feet (10 Kilometers) **Minimum Tx Power:** -11.0 dBm **Rx Sensitivity:** -19 dBm

Important: This product has been tested and found to comply with the limits for Class 1 laser devices for IEC825, EN60825, and 21CFR1040 specifications.

The following shows the front panel of the XGLC, identifying its interfaces and major components:

Figure 18. 10 Gigabit Ethernet line Card (XGLC)

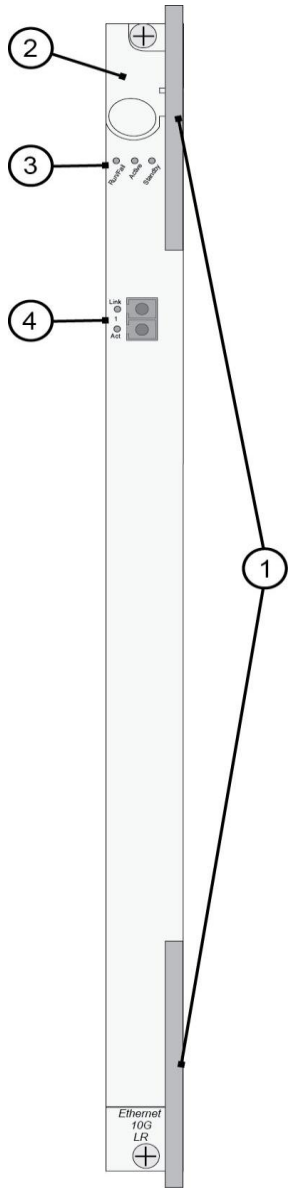


Table 19. XGLC Callouts

Item	Description
1	**Card Ejector Levers**—Use to insert/remove card to/from chassis.
2	**Interlock Switch**—In its Down position the interlock switch notifies system to safely power down the card prior to its removal.
3	**Card Level Status LEDs**—Show the status of the card.

Item	Description
4	**Gigabit Ethernet Interface(s)—10** Gigabit Ethernet (GE) SFP+ modules. 10Base-SR and 10Base-LR interfaces are supported, depending on the SFP+ module installed.

Optical Line Card (OLC2)

The OLC2 is labeled OLC2 OC-3/STM-1 Multi Mode (or Single Mode depending on SFP type). The OLC2 supports either OC-3 or STM-1 signaling and ATM.

The OLC2 support both SDH and SONET. The basic unit of framing in SDH is STM-1 (Synchronous Transport Module level - 1), which operates at 155.52 Mbps. SONET refers to this basic unit as STS-3c (Synchronous Transport Signal - 3, concatenated), but its high-level functionality, frame size, and bit-rate are the same as STM-1.

SONET offers an additional basic unit of transmission, STS-1 (Synchronous Transport Signal - 1), operating at 51.84 Mbps—exactly one third of an STM-1/STS-3c. The OLC2 concatenates three STS-1 (OC-1) frames to provide transmission speeds up to 155.52 Mbps with payload rates of 149.76 Mbps and overhead rates of 5 .76 Mbps.

The OLC2 optical fiber line card supports network connectivity through Iu or IuPS interfaces to the UMTS Terrestrial Radio Access Network (UTRAN). These interfaces are commonly used with our SGSN products to provide either non - IP 3G traffic or all IP 3G traffic (for all-IP packet-based networking) over ATM (Asynchronous Transfer Mode).

Each OLC2 provides four physical interfaces (ports) that are populated by Small Form-factor Pluggable (SFP) modules which include LC-type connectors. The Optical (ATM) line card supports two types of SFP modules (ports) and applicable cabling, but each card supports only one type at-a-time, as indicated in the following table.

Table 20. SFP Modules Supported by the OLC2

Module Type	Card Identification	Interface Type	Cable Specifications
Single-mode Optical Fiber	ATM/POS OC-3 SM IR-1	Single-mode Fiber, LC duplex female connector	**Fiber Types:** Single-mode optical fiber **Wavelength:** 1310 nm **Core Size:** 9 micrometers **Cladding Diameter:** 125 micrometers **Range:** Intermediate/21 kilometers **Attenuation:** 0.25 dB/KM **Min/Max Tx Power:** -15 dBm/-8 dBm **Rx Sensitivity:** -28 dBm

Module Type	Card Identification	Interface Type	Cable Specifications
Multi-mode Optical Fiber	ATM/POS OC-3 Multi-Mode	Multi-mode Fiber, LC duplex female connector	**Fiber Types:** Multi-mode optical fiber **Wavelength:** 1310 nm **Core Size:** 62.5 micrometers **Cladding Diameter:** 125 micrometers **Range:** Short/2 kilometers **Min/Max Tx Power:** -19 dBm/-14 dBm **Rx Sensitivity:** -30 dBm

The OLC2 supports the Star Channel (1 Gbps) for faster FPGA upgrades and is Restriction of Hazardous Substances (RoHS) 6/6 compliant.

Install the OLC2 directly behind its respective (active) packet processing card. As with other line cards, install the Optical (ATM) Line Card in slots 17 through 23, 26 through 39, and 42 through 48. The following figures show the panel of the OLC2, indicating its ports and major components.

Figure 19. Optical Line Card (OLC2)

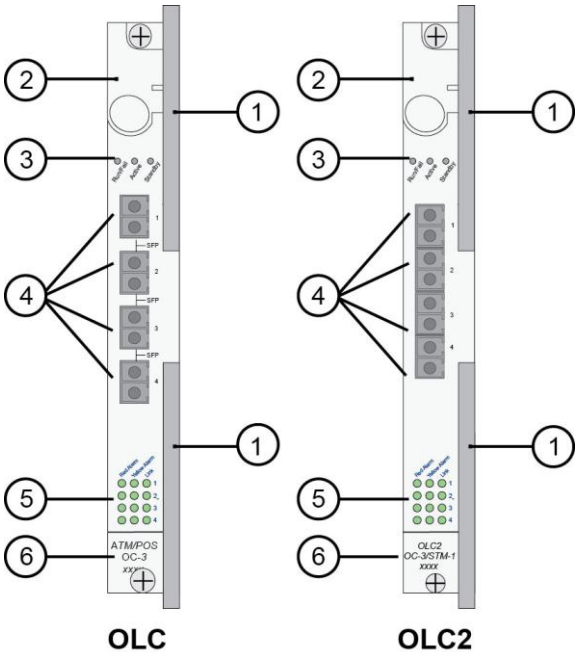


Table 21. OLC2 (ATM) Callouts

Item	Description
1	**Card Ejector Levers**—Use to insert/remove card to/from chassis.
2	**Interlock Switch**—In its Down position, the interlock switch notifies the system to safely power down the card prior to its removal.

Item	Description
3	**Card Level Status LEDs**—Show the status of the card.
4	**Port connectors**—Fiber LC duplex female connector.
5	**Port Level Status LEDs**—Show the status of a port.
6	**Line Card Label**—Identifies the type of SFP modules and cabling supported: • OLC2, OC-3/STM-1, Single Mode • OLC2, OC-3/STM-1, Multi-Mode

Channelized Line Card (CLC2)

The CLC2 is also referred to as the Frame Relay line card. It provides frame relay over SONET or SDH. The CLC2 supports network connectivity through a gigabit interface to connect to the Packet Control Unit (PCU) of the base station subsystem (BSS) in a mobile network. These interfaces are commonly used with the SGSN product to support frame relay.

In North America, the card supplies ANSI SONET STS-3 (optical OC-3) signaling. In Europe, the card supplies SDH STM-1 (optical OC-3). The transmission rate for the card is 155.52 Mbps with 336 SONET channels for T1 and 252 SDH channels for E1.

Each CLC2 provides four optical fiber physical interfaces (ports). The ports are populated by a Small Form-factor Pluggable (SFP) modules which include an LC-type connector. The ports of the CLC2 supports two types of SFP modules and cabling, as shown in the following table.

Table 22. SFP Modules supported by the CLC2

Module Type	Card Identification	Interface Type	Cable Specifications
Single-mode Optical Fiber	Channelized (STM-1/OC-3) SM IR-1	Single-mode Fiber, LC duplex female connector	**Fiber Types:** Single-mode optical fiber **Wavelength:** 1310 nm **Core Size:** 9 micrometers **Cladding Diameter:** 125 micrometers **Range:** Intermediate/21 kilometers **Attenuation:** 0.25 dB/KM **Min/Max Tx Power:** -15 dBm/-8 dBm **Rx Sensitivity:** -28 dBm

Module Type	Card Identification	Interface Type	Cable Specifications
Multi-mode Optical Fiber	Channelized (STM-1/OC-3) Multi-Mode	Multi-mode Fiber, LC duplex female connector	**Fiber Types:** Multi-mode optical fiber **Wavelength:** 1310 nm **Core Size:** 62.5 micrometers **Cladding Diameter:** 125 micrometers **Range:** Short/2 kilometers **Min/Max Tx Power:** -19 dBm/-14 dBm **Rx Sensitivity:** -30 dBm

The CLC2 supports the Star Channel (1 Gbps) for faster FPGA upgrades and is Restriction of Hazardous Substances (RoHS) 6/6 compliant.

Install the CLC2 directly behind its respective (Active) packet processing card. You may optionally install CLC2s behind a redundant (Standby) packet processing card. As with other line cards, install the Channelized Line Cards in slots 17 through 23, 26 through 39, and 42 through 48.

The following figures show the panel of the CLC2 Channelized Line Cards, identifying their interfaces and major components.

Figure 20. Channelized Line Card (CLC/CLC2)

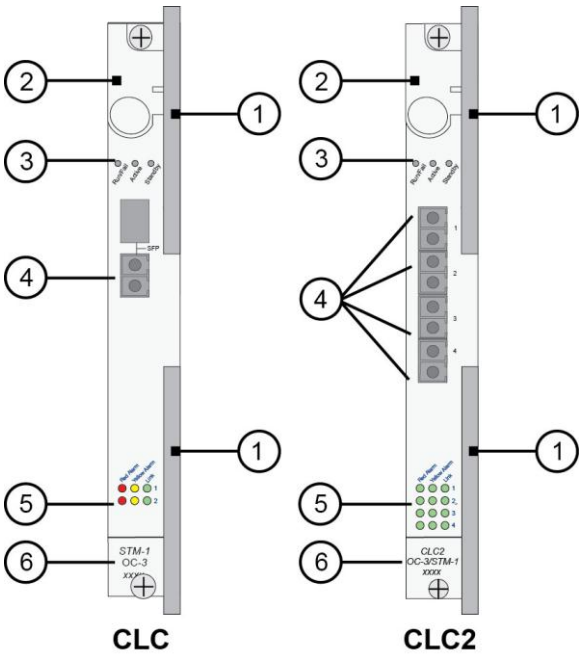


Table 23. CLC2 (Frame Relay) Callouts

Item	Description
1	**Card Ejector Levers**—Use to insert/remove card to/from chassis.

Item	Description
2	**Interlock Switch**—In its Down position the interlock switch notifies the system to safely power down the card prior to its removal.
3	**Card Level Status LEDs**—Show the status of the card.
4	**Port connectors**—Fiber LC duplex female connector.
5	**Port Level Status LEDs**—Show the status of a port.
6	**Line Card Label—**Identifies the type of SFP modules and cabling supported: • CLC2, OC-3/STM-1, Single Mode • CLC2, OC-3/STM-1, Multi-Mode

The Channelized Line Card (CLC2) was developed in compliance with the following standards:

- ITU-T - Recommendation G.704 - Synchronous Frame Structures Used at 1544, 6312, 2048, 8448 and 44736 kbps Hierarchical Levels, October, 1998.
- ITU-T - Recommendation G.706 - Frame Alignment and Cyclic Redundancy Check (CRC) Procedures Relating to Basic Frame Structures Defined in Recommendation G.704, April 1991.
- ITU-T - Recommendation G.707 Network Node Interface for the Synchronous Digital Hierarchy (SDH), December 2003.
- ITU-T - Recommendation G.747 Second Order Digital Multiplex Equipment Operating at 6312 kbps and Multiplexing Three Tributaries at 2048 kbps, 1993.
- ITU-T - Recommendation G.751 Digital Multiplex Equipment Operating at the Third Order Bit Rate of 34 368 kbps and the Fourth Order Bit Rate of 139 264 kbps and Using Positive Justification, 1993.
- ITU-T - Recommendation G.775, - Loss of Signal (LOS) and Alarm Indication Signal (AIS) Defect Detection and Clearance Criteria, November 1994.
- ITU-T - Recommendation G.783 Characteristics of Synchronous Digital Hierarchy (SDH) Equipment Functional Blocks, February 2004.
- ITU-T - Recommendation G.823, -The Control of Jitter and Wander within Digital Networks which are based on the 2048 kbps Hierarchy, March 2000.
- ITU-T - Recommendation G.824 The Control of Jitter and Wander within Digital Networks which are based on the 1544 kbps Hierarchy, March 2000.
- ITU-T - Recommendation G.825 Control of Jitter and Wander within Digital Networks Which are Based on the Synchronous Digital Hierarchy (SDH) Series G: Transmission Systems and Media, Digital Systems and Networks Digital Networks - Quality and Availability Targets, March 2000.
- ITU-T - Recommendation G.832 Transport of SDH elements on PDH networks Frame and multiplexing structures, October 1998.
- ITU-T - Recommendation G.957 Optical interfaces for equipment and systems relating to the Synchronous Digital Hierarch, March 2006.
- ITU-T - Recommendation I.431 - Primary Rate User-Network Interface Layer 1 Specification, March 1993.
- ITU-T - Recommendation O.150 - General Requirements for Instrumentation Performance Measurements on Digital Transmission Equipment, May 1996.
- ITU-T - Recommendation O.151 - Error Performance Measuring Equipment Operating at the Primary Rate and Above, October 1992.

- ITU-T - Recommendation O.152 - Error Performance Measuring Equipment for Bit Rates of 64 kbps and N x 64 kbps, October 1992.
- ITU-T - Recommendation O.153 - Basic Parameters for the Measurement of Error Performance at Bit Rates below the Primary Rate, October 1992.
- ITU-T - Recommendation Q.921 - ISDN User-Network Interface - Data Link Layer Specification, September 1997.
- ITU-T - Recommendation Q.922 - ISDN data link layer specification for frame mode bearer services.
- ITU-T - Recommendation Q.933 Annex E.
- Frame Relay Forum - FRF 1.2 - User-to-Network Interface (UNI).
- Frame Relay Forum - FRF 2.1 - Frame Relay Network-to-Network Interface (NNI).
- Frame Relay Forum - FRF 5.0 - Network Interworking.
- Frame Relay Forum - FRF 8.1 - Service Interworking.
- Frame Relay Forum - FRF 12.0 - Frame Relay Fragmentation.

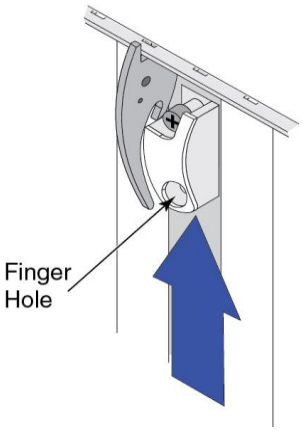
Card Interlock Switch

Each card has a switch interlock mechanism that is integrated with the upper card ejector lever. The interlock ensures proper notification to the system before a card is removed. You cannot configure or place a card into service until you slide the card interlock switch upward. This locks the upper ejector lever in place and signals the system that the card is ready for use.

Sliding the interlock downward to the unlocked position allows you to operate the upper ejector lever. This sliding lock mechanism notifies the system to migrate various processes on the card prior to its removal.

The following figure shows how the card interlock switch works in conjunction with the ejector lever.

Figure 21. Card Interlock Switch in the Lever-locked Position



Card Identifiers

The table below cross-references ASR 5000 application and line cards by acronym, label, variant, and Cisco part identifier (PID).

Table 24. ASR 5000 Component References

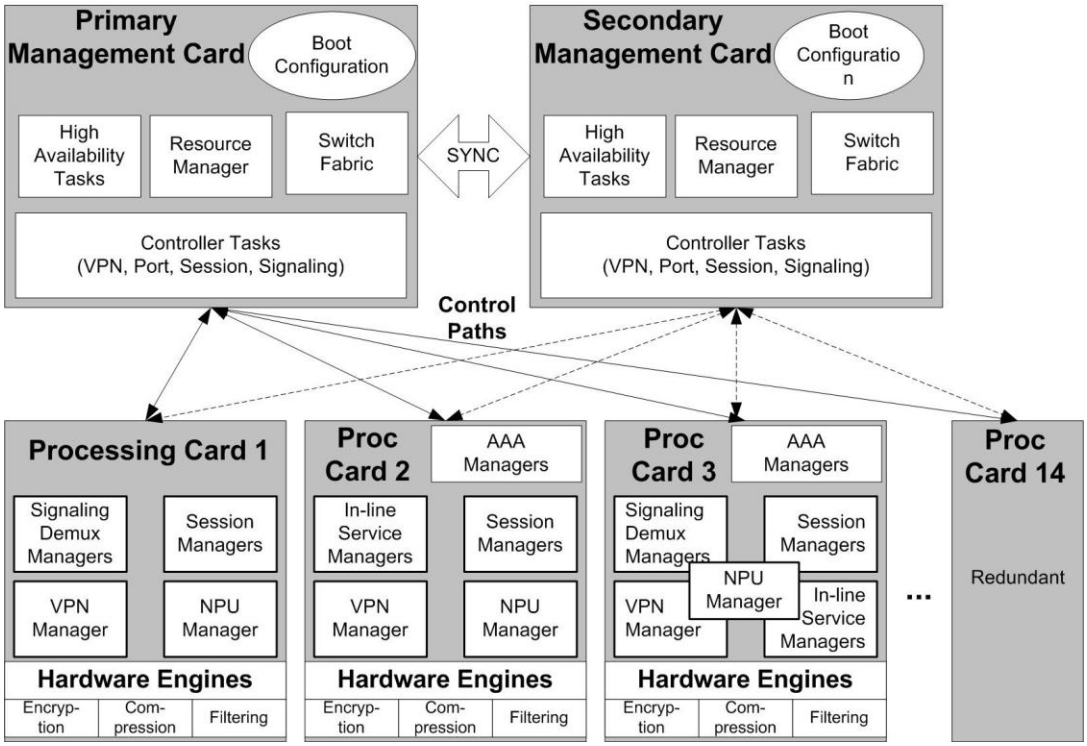
Descriptor	Acronym	Label	Variant	Cisco PID
Application Cards				
System Management Card	SMC	System Management	None	ASR5K-SMC-K9
Packet Services Card 16GB	PSC-A	PSC	None	ASR5K-PSC-16G-K9
Packet Services Card 32GB	PSC2	Packet Services 2 32GB	None	ASR5K-PSC-32G-K9
Packet Services Card 64GB	PSC3	Packet Services 3 64GB	None	ASR5K-PSC-64G-K9
Packet Processing Card 16GB	PPC	PPC	None	ASR5K-PPC-K9
Switch Processor Input/Output	SPIO	Switch Processor I/O	Switch Processor I/O, BNC BITS	ASR5K-SPIO-BNC-K9
			Switch Processor I/O, 3-Pin BITS	ASR5K-SPIO-3PN-K9
			Switch Processor I/O, BNC BITS with Stratum 3 module	ASR5K-SPS3-BNC-K9
			Switch Processor I/O, 3-Pin BITS with Stratum 3 module	ASR5K-SPS3-3PN-K9
Line Cards				
Redundancy Crossbar	RCC	Redundancy Crossbar	None	ASR5K-RCC-K9
FELC Ethernet 10/100 Line Card	FELC	Ethernet 10/100	None	ASR5K-01100E-K9
FELC Ethernet 10/100 Line Card 2	FLC2	Ethernet 10/100	None	ASR5K-08100E-K9
GELC Ethernet 1000 Line Card	GELC	Ethernet 1000 SX	with SX MM Short Haul SFP	ASR5K-011GE-SX-K9
		Ethernet 1000 LX	with LX SM SFP	ASR5K-011GE-LX-K9
		Ethernet 1000 T	with copper SFP	ASR5K-011GE-T-K9

Descriptor	Acronym	Label	Variant	Cisco PID
GLC2 Ethernet 1000 Line Card	GLC2	Ethernet 1000 SX	with SX MM Short Haul SFP	ASR5K-011G2-SX-K9
		Ethernet 1000 LX	with LX SM SFP	ASR5K-011G2-LX-K9
		Ethernet 1000 T	with copper SFP	ASR5K-011GE-T-K9
QGLC 4-Port Ethernet 1000 Line Card	QGLC	Ethernet 1000 SX	with SX MM Short Haul SFP	ASR5K-041GE-SX-K9
		Ethernet 1000 LX	with LX SM SFP	ASR5K-041GE-LX-K9
		Ethernet 1000 T	with copper SFP	ASR5K-011G2-T-K9
QGLC Rev2 4-Port Ethernet 1000 Line Card	QGLC	Ethernet 1000 SX	with SX MM Short Haul SFP	ASR5K-042GE-SX-K9
		Ethernet 1000 LX	with LX SM SFP	ASR5K-042GE-LX-K9
		Ethernet 1000 T	with copper SFP	ASR5K-042GE-T-K9
XGLC 1-Port 10 Gigabit Ethernet Line Card	XGLC	Ethernet 10G SR	with MM SFP+	ASR5K-0110G-MM-K9
		Ethernet 10G LR	with SM SFP+	ASR5K-0110G-SM-K9
Channelized 4-port Line Card	CLC2	CLC2 OC-3/STM-1	with MM SFP	ASR5K-C4OC3-MM-K9
			with SM SFP	ASR5K-C4OC3-SM-K9
Optical 4-port (ATM) Line Card	OLC2	OLC2 OC-3/STM-1	with MM SFP	ASR5K-4OC3C-MM-K9
			with SM SFP	ASR5K-4OC3C-SM-K9

Chapter 4
Software Architecture

The operating system software is based on a Linux software kernel and runs specific applications in the system such as monitoring tasks, various protocol stacks, and other items. The following figure shows an example block diagram of the operating system's software architecture.

Figure 22. Software Architecture Block Diagram



The software architecture is designed for high availability, flexibility, and performance. The system achieves these goals by implementing the following key software features:

- **Scalable control and data operations:**

  System resources can be allocated separately for control and data paths. For example, certain processing cards could be dedicated to performing routing or security control functions while other cards are dedicated to processing user session traffic. As network requirements grow and call models change, hardware resources can be added to accommodate processes, such as encryption, packet filtering, etc., that require more processing power. Additionally, certain software task sizes are dynamically sized based on hardware and installed licenses thus conserving system memory.

- **Fault containment:**

The system isolates faults at the lowest possible levels through its High Availability Task (HAT) function that monitors all system entities for faults and performs automatic recovery and failover procedures using its Recovery Control Task (RCT).

Processing tasks are distributed into multiple instances running in parallel so if an unrecoverable software fault occurs, the entire processing capabilities for that task are not lost. User session processes can be sub-grouped into collections of sessions so that if a problem is encountered in one sub-group users in another sub-group will not be affected by that problem. The architecture also allows check-pointing of processes, which is a mechanism to protect the system against any critical software processes that may fail.

The self-healing attributes of the software architecture protects the system by anticipating failures and instantly spawning mirror processes locally or across card boundaries to continue the operation with little or no disruption of service. This unique architecture allows the system to perform at the highest level of resiliency and protects the user's data sessions while ensuring complete accounting data integrity.

- **Promotes internal location transparency:**

  Processes can be distributed across the system to fit the needs of the network model and specific process requirements. For example, most tasks can be configured to execute on an SMC or a processing card, while some processor intensive tasks can also be performed across multiple processing cards to utilize multiple CPU resources. Distribution of these tasks is invisible to the user.

- **Leverages third party software components:**

  The use of the Linux operating system kernel enables reuse of many well-tested, stable, core software elements such as protocol stacks, management services, and application programs.

- **Supports dynamic hardware removal/additions:**

  By migrating tasks from one card to another via software controls, application cards can be “hot swapped” to dynamically add capacity and perform maintenance operations without service interruption**.**

- **Multiple context support:**

  The system can be fully virtualized to support multiple logical instances of each service. This eliminates the possibility of any one domain disrupting operations for all users in the event of a failure.

  Further, multiple context support allows operators to assign duplicate/overlapping IP address ranges in different contexts.

Understanding the Distributed Software Architecture

To better understand the advantages of the system's distributed software architecture, this section presents an overview of the various components used in processing a subscriber session. Numerous benefits are derived from the system's ability to distribute and manage sessions across the entire system. The following information is intended to familiarize you with some of the components and terminology used in this architecture.

Software Tasks

To provide unprecedented levels of software redundancy, scalability, and robust call processing, the system's software is divided into a series of tasks that perform specific functions. These tasks communicate with each other as needed to share control and data information throughout the system.

A task is a software process that performs a specific function related to system control or session processing. There are three types of tasks that operate within the system:

- Critical tasks

  These tasks control essential functions to ensure the system's ability to process calls. Examples of these would be system initialization and automatic error detection and recovery tasks.

- Controller tasks

  These tasks, often referred to as "Controllers", serve several different purposes. These include:

- Monitoring the state of their subordinate managers and allowing for intra-manager communication within the same subsystem.
- Enabling inter-subsystem communication by communicating with controllers belonging to other subsystems

  Controller tasks mask the distributed nature of the software from the user - allowing ease of management.

- Manager tasks

  Often referred to as "Managers", these tasks control system resources and maintain logical mappings between system resources. Some managers are also directly responsible for call processing.

  System-level processes can be distributed across multiple processors, thus reducing the overall workload on any given processor—thereby improving system performance. Additionally, this distributed design provides fault containment that greatly minimizes the impact to the number of processes or PPP sessions due to a failure.

  The SMC has a single Control Processor (CP) that is responsible for running tasks related to system management and control.

  Each PSC contains two CPs (CPU 0 and CPU 1) The CPs on the processing cards are responsible for PPP and call processing, and for running the various tasks and processes required to handle the mobile data call. In addition to the CPs, the processing cards also have a high-speed Network Processor Unit (NPU) used for enhanced IP forwarding.

Subsystems

Individual tasks that run on CPs can be divided into subsystems. A subsystem is a software element that either performs a specific task or is a culmination of multiple other tasks. A single subsystem can consist of critical tasks, controller tasks, and manager tasks.

Following is a list of the primary software subsystems:

- **System Initiation Task (SIT) Subsystem:** This subsystem is responsible for starting a set of initial tasks at system startup and individual tasks as needed.
- **High Availability Task (HAT) Subsystem:** Working in conjunction with the Recovery Control Task (RCT) subsystem, HAT is responsible for maintaining the operational state of the system. HAT maintains the system by monitoring the various software and hardware aspects of the system. On finding any unusual activities, such as the unexpected termination of a task, the HAT would take a suitable action like triggering an event prompting the RCT to take some corrective action or report the status.

  The benefit of having this subsystem running on every processor is that should an error occur, there is minimal or no impact to the service.
- **Recovery Control Task (RCT) Subsystem:** Responsible for executing a defined recovery action for any failure that occurs in the system. The RCT subsystem receives recovery actions from the HAT subsystem.

  The RCT subsystem only runs on the active SMC and synchronizes the information it contains with the mirrored RCT subsystem on the standby management card.
- **Shared Configuration Task (SCT) Subsystem:** Provides the system with a facility to set, retrieve, and be notified of system configuration parameter changes. This subsystem is primarily responsible for storing configuration data for the applications running within the system.

  The SCT subsystem runs only on the active SMC and synchronizes the information it contains with the mirrored SCT subsystem on the standby management card.
- **Resource Management (RM) Subsystem:** The RM subsystem is responsible for assigning resources to every system task upon their start-up. Resources are items such as CPU loading and memory. RM also monitors these items to verify the allocations are being followed. This subsystem is also responsible for monitoring all sessions and communicating with the Session Controller, a subordinate task of the Session subsystem, to enforce capacity licensing limits.
- **Virtual Private Network (VPN) Subsystem:** Manages the administrative and operational aspects of all VPN-related entities in the system. The types of entities managed by the VPN subsystem include:
  - Creating separate VPN contexts
  - Starting the IP services within a VPN context
  - Managing IP pools and subscriber IP addresses
  - Distributing the IP flow information within a VPN context

  All IP operations within the system are done within specific VPN contexts. In general, packets are not forwarded across different VPN contexts. The only exception to this rule is the Session subsystem.
- **Network Processing Unit (NPU) Subsystem:** The NPU subsystem is responsible for the following:
  - "Fast-path" processing of frames using hardware classifiers to determine each packet's processing requirements
  - Receiving and transmitting user data frames to/from various physical interfaces
  - IP forwarding decisions (both unicast and multicast)
  - Per interface packet filtering, flow insertion, deletion, and modification
  - Traffic management and traffic engineering
  - Passing user data frames to/from processing card CPUs
  - Modifying/adding/stripping datalink/network layer headers

- Recalculating checksums
- Maintaining statistics
- Managing both external line card ports and the internal connections to the data and control fabrics

- **Card/Slot/Port (CSP) Subsystem:** Responsible for coordinating the events that occur when any card is inserted, locked, unlocked, removed, shut down, or migrated, the CSP subsystem is responsible for all card activity for each of the 48 slots in the chassis. It is also responsible for performing auto-discovery and configuration of ports on a newly inserted line card, and determining how line cards map to processing cards (including through an RCC in failover situations).

  The CSP subsystem runs only on the active SMC and synchronizes the information it contains with the mirrored SCT subsystem on the standby management card. It is started by the SIT subsystem, and monitored by the HAT subsystem for failures.

- **Session Subsystem:** The Session subsystem is responsible for performing and monitoring the processing of a mobile subscriber's data flows. Session processing tasks for mobile data calls include: A10/A11 termination for CDMA2000 networks, GSM Tunneling Protocol (GTP) termination for GPRS and/or UMTS networks, asynchronous PPP processing, packet filtering, packet scheduling, Diffserv codepoint marking, statistics gathering, IP forwarding, and AAA services. Responsibility for each of these items is distributed across subordinate tasks (called Managers) to provide for more efficient processing and greater redundancy. A separate Session Controller task serves as an integrated control node to regulate and monitor each of the Managers and to communicate with the other active subsystems.

  This subsystem also manages all specialized user data processing, such as for payload transformation, filtering, statistics collection, policing, and scheduling.

Chapter 5
Redundancy and Availability Features

Every minute of downtime and every dropped session represents lost revenue to the wireless operator resulting in potential customer loss and reduced profitability. With this understanding, we have developed a system that exceeds the availability features found in the majority of today's wireless and wireline access devices.

This chapter includes the following topics:

- Service Availability Features
- Mean Time Between Failure
- System Availability
- Spare Component Recommendations

Service Availability Features

In its recommended redundant configuration, the system provides the highest level of service assurance. Following is detailed information describing the service availability features found in the system.

Hardware Redundancy Features

In addition to providing the highest transaction rates and session capacity, the system is designed to provide robust hardware reliability and service assurance features.

Features of the hardware design include:

- 1:1 System Management Card (SMC) redundancy
- 1:n packet processing card redundancy, allowing redundancy of multiple active to multiple redundant cards for up to 14 total packet processing cards.

  1:1 redundancy is supported for these cards however some subscriber sessions and accounting information may be lost in the event of a hardware or software failure even though the system remains operational.
- 1:1 Optical (ATM) line card (LC) redundancy (OLC and OLC2)
- 1:1 Channelized (STM-1/OC-3) Line Card redundancy (CLC and CLC2)
- 1:1 Quad Gigabit Ethernet Line Card (QGLC)
- 1:1 10 Gigabit Ethernet Line Card (XGLC)
- 1:1 Switch Processor I/O (SPIO) card redundancy
- 1:1 Fast Ethernet Line Card (FELC)
- 1:1 Gigabit Ethernet Line Cards (GELC, GLC2)
- Configurable port redundancy (Ethernet, ATM (SGSN), and SPIO line cards)
- Redundancy Crossbar Card (RCC) for processor-card-to-line card failover using the 280 Gbps Redundancy Bus
- Self-healing redundant 320 Gbps switching fabric
- Redundant 32 Gbps Control Bus
- Redundant Power Filter Units (PFUs)
- Hot-swappable cards, allowing dynamic replacement while the system is operational

Hardware Redundancy Configuration

The maximum redundant configuration for a fully loaded system supporting data services consists of the following:

- Two SMCs: 1 active and 1 standby (redundant)
- 14 processing cards: 13 active and 1 standby
- Two SPIOs: 1 active and 1 standby
- 26 half-height Ethernet line cards: 13 active and 13 standby (10/100 Ethernet Line Card (FELC), 1000 Gigabit Line Card (GELC, GLC2), and Quad Gigabit Ethernet Line Card (QGLC). these cards support top-bottom (vertical) redundancy.

- Two full-height 10 Gigabit Ethernet Line Cards (XGLC): 1 active, 1 standby. The XGLC supports side-by-side (horizontal) redundancy.
- 26 Optical (ATM) line cards: 13 active and 13 standby (OLC and OLC2)
- 26 Channelized line cards: 13 active and 13 standby (CLC and CLC2)
- Two RCCs: 2 standby

This configuration allows for the highest session capacity while still providing redundancy. The following figures depict this recommended maximum redundant configuration.

Figure 23. Recommended Redundant Configuration for Data Services - Front View

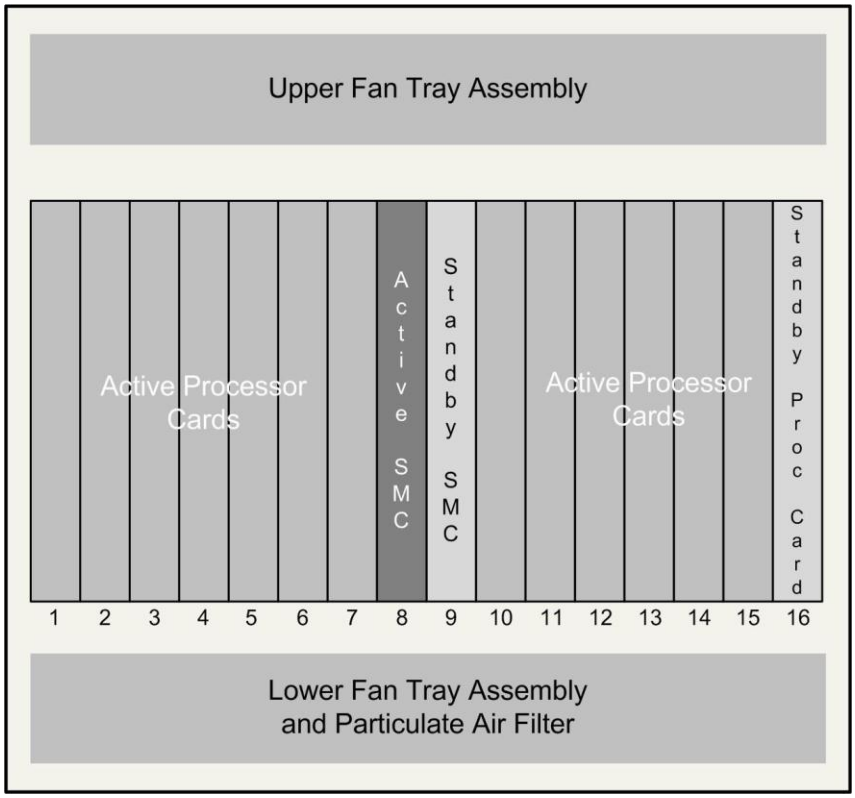
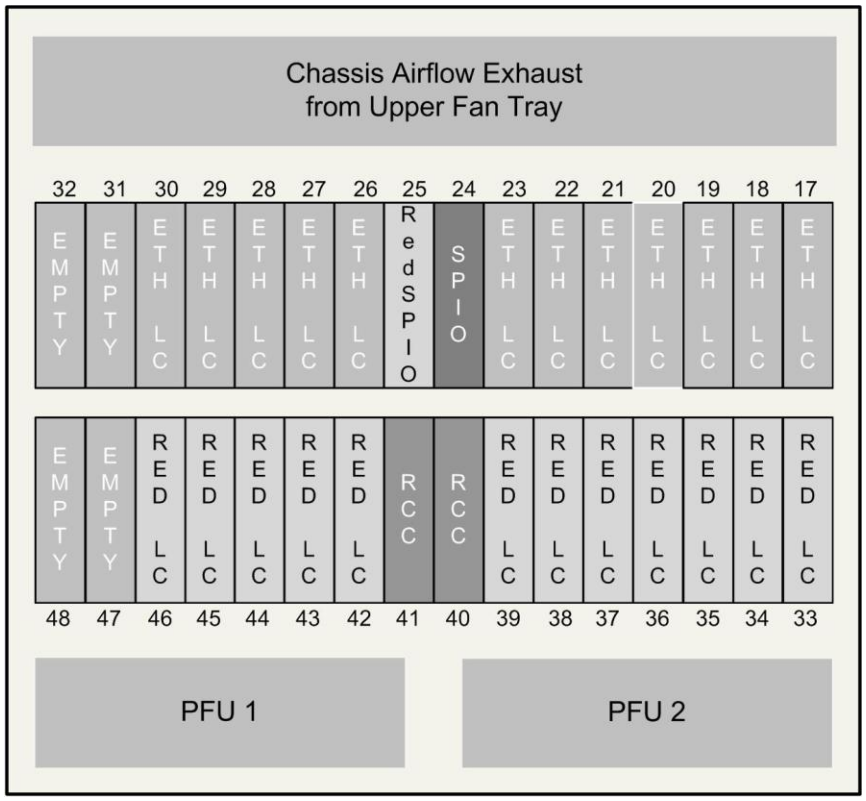


Figure 24. Recommended Redundant Configuration for Data Services - Rear View



Software Assurance Features

Numerous features are built into the system software to ensure the continuation of service in the case of software process failures. SMC software controls the management contexts and overall system control, while processing card software controls the PPP sessions, AAA, and VPN processes.

Session Recovery Feature

This licensed software feature performs an automatic recovery of all fully established subscriber sessions should a session manager task failure occur.

With this feature enabled, there is no loss of session information as described in table above. Session recovery consists of the migration and recreation of control and data packet state information, subscriber session statistics, or session time parameters such as idle timer and others.

This feature is enabled/disabled on a chassis-wide basis and requires additional processing card hardware to ensure that enough reserve resources (memory, processing, etc.) are available to fully recover session in the event of a software or hardware failure.

Interchassis Session Recovery

The Interchassis Session Recovery (ICSR) feature provides the highest possible availability for continuous call processing without interrupting subscriber services. This is accomplished through the use of redundant chassis. The chassis are configured as primary and backup with one being active and one inactive. Both chassis are connected to the AAA server. When calls pass the checkpoint duration timer, checkpoint data is sent from the active chassis to the inactive chassis. If the active chassis handling the call traffic goes out of service, the inactive chassis transitions to the active state and continues processing the call traffic without interrupting the subscriber session.

The chassis determine which is active through a propriety TCP-based connection called a redundancy link. This link is used to exchange status messages between the primary and backup chassis and must be maintained for proper system operation. In the event the redundancy link goes out of service, interchassis session recovery (ICSR) is maintained through the use of authentication probes and BGP peer monitoring. BGP routing must be enabled.

Mean Time Between Failure

Mean Time Between Failure (MTBF) data is used to provide statistical information as to the length of time that should expire before a particular card or system fails. This information is calculated using the following method:

Calculated MTBF - Expected elapsed time before failure occurs using the method defined in Telcordia TR-NWT-000332-CORE. This is based on reliability of components and design factors.

Failure per million hours (Fpmh) identifies the predicted failure rate per one million hours (for every 1,000,000 hours of operation, "FITS number" of failures would be expected to occur) for a component of the system.

The following table shows the MTBF characteristics of each major component of the system.

Table 25. Mean Time Between Failure Statistics

Cisco PID	Description	MTBF (Hours)	MTBF (Years)	Fpmh (Failure per million hours)
ASR5000-CHSSYS-K9=	Chassis with Midplane	16,386,995	1869.38	0.061
ASR5K-SMC-K9=	System Management Card (SMC)	104,372	11.91	9.58
ASR5K-PSC-K9=	Packet Services Card (PSC)	102,294	11.68	9.78
ASR5K-PSC-32G-K9= ASR5K-PSC-64G-K9= ASR5K-PSC-16G-K9= ASR5K-PPC-K9=	Packet Services Card (PSC2, PSC3, PSCA, PPC)	93,950	10.75	10.64
ASR5K-0110G-MM-K9= ASR5K-0110G-SM-K9=	10 Gigabit Ethernet Line Card (XGLC)	247,720	28.28	4.04
ASR5K-041GE-SX-K9= ASR5K-041GE-T-K9= ASR5K-041GE-LX-K9=	Quad Gig-E Card (QGLC)	258,606	29.52	3.867
ASR5K-01OC3-SM-K9=	ATM/POS OC-3 SM IR-1 Card optical daughter card	214,492 1,419,581	48.6 73.4	4.66 0.70

Cisco PID	Description	MTBF (Hours)	MTBF (Years)	Fpmh (Failure per million hours)
ASR5K-SPIO-BNC-K9= ASR5K-SPIO-3PN-K9= ASR5K-SPS3-BNC-K9= ASR5K-SPS3-3PN-K9=	Switch Processor I/O Card (SPIO)	333,999	38.13	2.99
ASR5K-RCC-K9=	Redundancy Crossbar Card (RCC)	555,862	63.46	1.79
ASR5K-01100E-K9=	Fast Ethernet Card (FELC)	495,886	56.61	2.01
ASR5K-011GE-SX-K9= ASR5K-011GE-LX-K9= ASR5K-011GE-T-K9=	Gigabit Ethernet Card (GELC)	396,715	45.29	2.52
ASR5K-011G2-SX-K9 ASR5K-011G2-LX-K9 ASR5K-011G2-T-K9	Gigabit Ethernet Card v2 (GLC2)	396,715	45.29	2.52
ASR5K-PFU=	Power Filter Unit (PFU)	967,118	110.40	1.03
ASR5K-FANT-LW=	Fan Tray Unit - Lower	70,517	8.05	19.51
ASR5K-FANT-UP=	Fan Blower Unit - Upper	120,178	13.72	18.72

System Availability

System-level Mean Time To Failure (MTTF), is the average interval of time that a component will operate before failing. Reliability information is based on the number of overall anticipated failures of the individual components, in conjunction with any redundancy schemes employed to minimize the impact of such failures.

The following table provides service availability calculations (based on reliability modeling) for the ASR 5000 platform.

Table 26. Hardware Platform Availability Information

Platform	Operational Uptime	Yearly Downtime	MTTF	
	(%)	(minutes)	Hours	Years
ASR 5000	99.999978	0.12	14,077,473	1605.91

One suggestion to help improve overall system availability is to institute an on-site spares program, wherein key components are housed locally with the deployed equipment. The following section defines a recommended spares program and quantities for the system.

Mean Time To Repair (MTTR) is the amount of time needed to repair a component, recover the system, or otherwise restore service after a failure. System availability calculations are based on the industry standard of four hours.

Spare Component Recommendations

This section provides a recommended quantity of spare parts to be used as part of a spare components program for the system. The information contained is for informational purposes only, and should only be used as a guideline for designing a spares program that meets your company's design, deployment, and availability goals.

It is recommended that your company either has fully-trained personnel available to effect the exchange of Field Replaceable Units (FRUs) within your network, or requests on-site or field engineering resources to perform such duties.

Based on industry-leading redundancy and failover features found in the system, the following minimum spare parts levels for any planned deployment are recommended.

Table 27. Recommended FRU Parts Sparing Quantities

Component Name	Minimum number of spares	For every "n" number of deployed components
ASR 5000 Chassis with Midplane	1	20
System Management Card (SMC)	1	10
Packet Services Cards (PSCx, PPC)	1	12
Quad Gig-E Line Card (QGLC)	1	20
10 Gigabit Ethernet Line Card (XGLC)	1	20
Optical Line Card (OLC or OLC2)	1	20
Channelized Line Card (CLC or CLC2)	1	20
Switch Processor I/O Card (SPIO)	1	18
Redundancy Crossbar Card (RCC)	1	30
Fast Ethernet Line Card (FELC)	1	25
Gigabit Ethernet Line Card (GELC or GLC2)	1	25
Power Filter Unit (PFU)	1	30
Upper Fan Tray Unit	1	8
Lower Fan Tray Unit	1	5
Particulate Air Filter	1	1

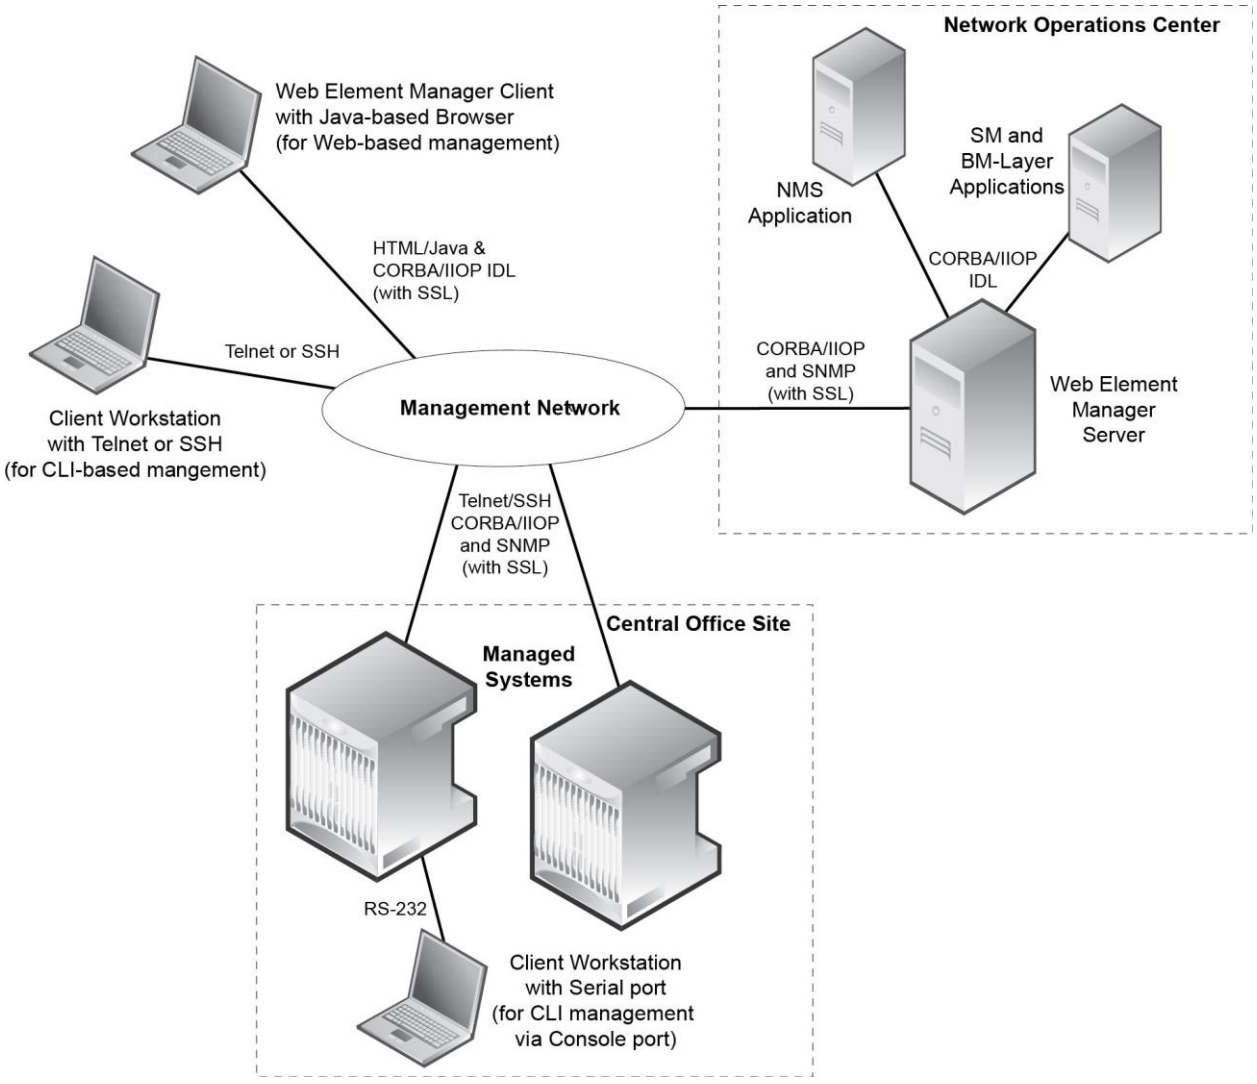
Chapter 6
Management System Overview

This chapter outlines the various methods of managing the system. There are multiple ways to locally or remotely manage the system using its out-of-band management interfaces. These include:

- Using the Command Line Interface (CLI)
    - Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card Ethernet management interfaces
    - Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager application
    - Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000 Base-SX management interfaces on the SPIO
    - Client-Server model supports any browser (i.e. Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
    - Supports Common Object Request Broker Architecture (CORBA) protocol, Secure Sockets Layer (SSL) for encryption of management data, and Simple Network Management Protocol version 1 (SNMPv1) for fault management
    - Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
    - Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network. Element Management Methods

Figure 25. Element Management Methods



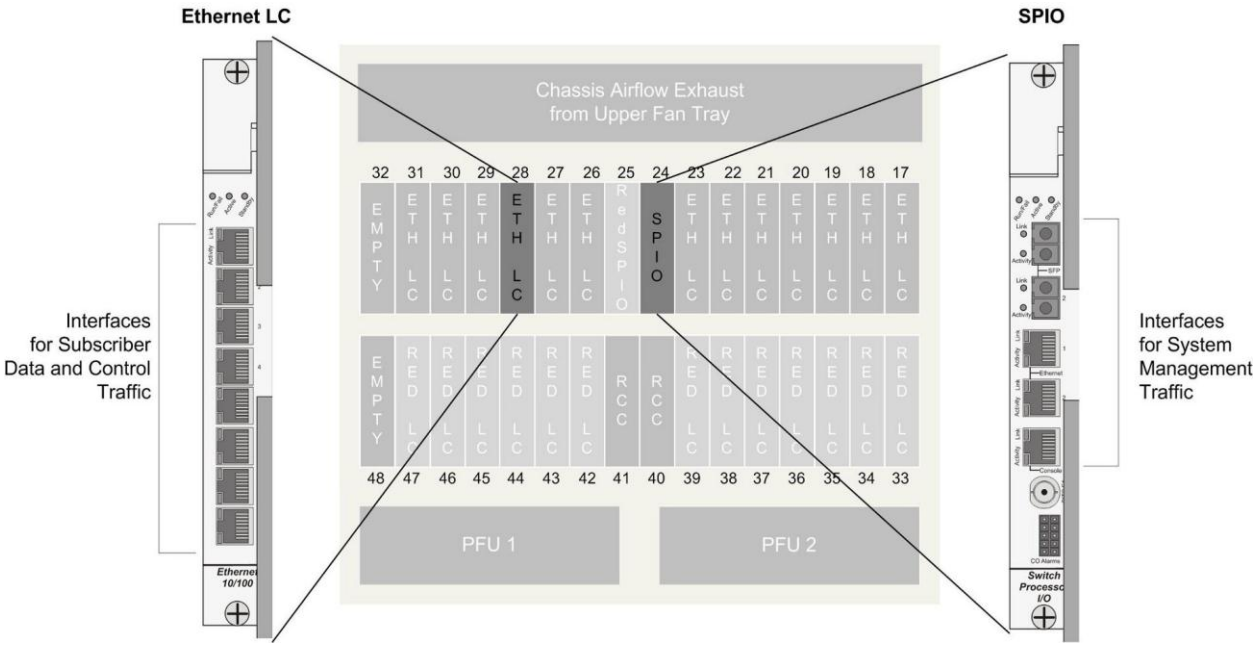
The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management - focusing on providing superior quality Network Element (NE) and element management system (Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1), into higher-level management systems - giving wireless operators the ability to integrate the system into their overall network, service, and business management systems.

Overview information about each of these methods follows. For detailed information, please see the System Administration and Configuration Reference, the Web Element Manager Getting Started Guide, or the Web Element Manager's robust Help system.

Out-of-Band Management

Management of the system is performed using Out-Of-Band (OOB) transmission methods through either the Console port or one of the Ethernet management ports on the SPIO. OOB management ensures that no management traffic can be accessed or viewed by any subscriber. Management data is separated on different physical interfaces from those used to transport user data. The following figure shows this separation.

Figure 26. Separation of Management Data From User Data



Additionally, the system uses the **local** context solely for system management purposes. Contexts are described in this document's Glossary, but basically they provide a way to host multiple virtual service or configuration parameter groups in a single physical device. To ensure OOB management, users are required to create other service-specific contexts for user data.

By using the **local** context as the separate management context, network operations personnel are able to utilize their own RADIUS services for management authentication and accounting, further maintaining the separation of user and management data.

Command Line Interface

CLI Overview

The CLI is a multi-threaded man machine interface that allows users to manipulate, configure, control, and query the various components that make up the system and the services hosted within the system. The CLI contains numerous command sets that perform various pre-defined functions when entered by a user. The CLI communicates with other controls and software tasks that make up the operating system.

The CLI provides numerous features, including:

- Simultaneous multiple CLI user support, providing a CLI instance for every context.

  The maximum number of multiple CLI session support is based on the amount of available memory. The Resource Manager, however, reserves enough resources so that the following minimum number of CLI sessions are assured:

  - For ASR 5000s: 15

    In both cases, one of the assured sessions is reserved for use exclusively by a CLI session on an SPIO console interface.

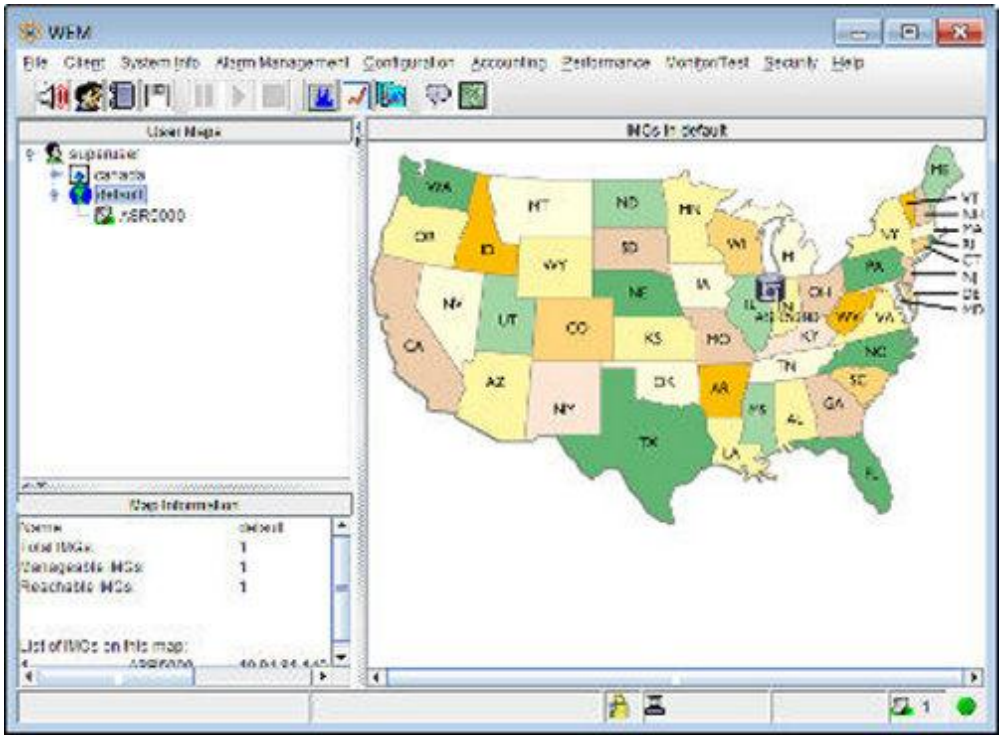
- Local or remote management login support
- Hierarchical structure supporting two command modes
  - Exec (execute) Mode, supporting basic commands that allow users to maneuver around system and perform monitoring functions
  - Config (configuration) Mode, providing global system configuration and context and service-specific configuration functions
- Differentiated administrative user privileges
  - Inspector users have minimal read-only privileges
  - Operator users have read-only privileges. They can maneuver across multiple contexts, but cannot perform configuration operations
  - Administrator users have read-write privileges and full access to all contexts and command modes (except for a few security functions)
  - Security Administrator users have read-write privileges and full access to all contexts and command modes
- Intuitive CLI command prompt displaying user's exact location within the CLI, command mode, and user privilege level
- CLI command auto-completion feature that allows users to enter only enough characters to make a command unique, prompting the system to complete the rest of the command or keyword by pressing the <Tab> key
- CLI auto-pagination, improving the readability of command output displays
- Complete command history features, allowing users to review all commands previously entered during current session, and EMACS-style command line manipulation features increasing CLI usability
- Interactive, context-sensitive Help, providing two levels of help for CLI commands, keywords, and variables

For more detailed information, reference *Command Line Interface Overview* chapter in the *System Administration and Configuration Reference*.

Web Element Manager Application

The Web Element Manager is a client-server application providing complete element management of the system. The UNIX-based server application works with clients using virtually any Java-enabled web browser to remotely manage the network elements within the system using the Common Object Request Broker Architecture (CORBA) standard. The Secure Sockets Layer (SSL) protocol can be used to encrypt management data traffic between the client and the server. The following figure shows the Web Element Manager topology window:

Figure 27. WEM Topology



In addition to its element management capabilities, the Web Element Manager can be integrated with higher-layer network, service, and business management applications using its northbound CORBA interface.

For more information on Web Element Manager application, refer *Web Element Manager Overview* section.

Chapter 7
Application Detection and Control Overview

This chapter provides an overview of the Application Detection and Control (ADC) in-line service, formerly known as Peer-to-Peer Detection.

The System Administration Guide provides basic system configuration information, and the product administration guides provide procedures to configure basic functionality of core network service. It is recommended that you select the configuration example that best meets your service model, and configure the required elements for that model, as described in the respective product Administration Guide, before using the procedures in this chapter.

This chapter covers the following topics:

- ADC Overview
- How ADC Works

ADC Overview

P2P is a term used in two slightly different contexts. At a functional level, it means protocols that interact in a peering manner, in contrast to client-server manner. There is no clear differentiation between the function of one node or another. Any node can function as a client, a server, or both—a protocol may not clearly differentiate between the two. For example, peering exchanges may simultaneously include client and server functionality, sending and receiving information. ADC utilizes the Enhanced Charging Service (ECS) functionality. For information about ECS, refer to the *Enhanced Charging Services Administration Guide*.

The ADC in-line service works in conjuction with the following products:

- GGSN
- IPSG
- PDSN
- P-GW

Detecting P2P protocols requires recognizing, in real time, some uniquely identifying characteristic of the protocols. Typical packet classification only requires information uniquely typed in the packet header of packets of the stream(s) running the particular protocol to be identified. In fact, many P2P protocols can be detected by simple packet header inspection. However, some P2P protocols are different, preventing detection in the traditional manner. This is designed into some P2P protocols to purposely avoid detection. The creators of these protocols purposely do not publish specifications. A small class of P2P protocols is stealthier and more challenging to detect. For some protocols, no set of fixed markers can be identified with confidence as unique to the protocol.

Operators care about P2P traffic because of the behavior of some P2P applications (for example, Bittorrent, Skype, and eDonkey). Most P2P applications can hog the network bandwidth such that 20% ADC users can generate as much traffic as generated by the rest 80% non-ADC users. This can result into a situation where non-ADC users may not get enough network bandwidth for their legitimate use because of excess usage of bandwidth by the ADC users. Network operators need to have dynamic network bandwidth / traffic management functions in place to ensure fair distributions of the network bandwidth among all the users. And this would include identifying P2P traffic in the network and applying appropriate controlling functions to the same (for example, content-based premium billing, QoS modifications, and other similar treatments).

The Application Detection and Control technology makes use of innovative and highly accurate protocol behavioral detection techniques. In order to ensure its effectiveness, the ADC continually supports detection of newer protocols and versions. For more information on the supported protocols and applications, refer to the *Peer-to-Peer Protocol and Application Detection Support* appendix in this guide.

ADC supports statistics reporting and postpaid charging policies. Per-protocol statistics via bulkstats and via report records including:

- UDR types: Summarizing data usage for a given content type
- EDR types: Specific to a particular event
- e-GCDRs: Specific to 3GPP

Upon detection of a P2P protocol for a particular flow, one of the following actions can be applied:

- Blocking P2P traffic—blocking protocol(s) and discarding traffic
- Bandwidth policing—limiting the bandwidth, applied per PDP context per P2P application type
- Flow policing—limiting the number of simultaneous P2P flows
- QoS support—including policing

- TOS marking—applied per P2P protocol type
- Prepaid and postpaid ADC content-based billing
- Statistics reporting—analyzing per-protocol statistics using bulkstats

Platform Requirements

The ADC in-line service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the Installation Guide for the chassis and/or contact your Cisco account representative.

License Requirements

The Application Detection and Control is a licensed Cisco feature. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

P2P Voice Call Duration

The ADC product has the capability to detect network traffic created by P2P VoIP clients such as Skype, Yahoo, MSN, Gtalk, Oscar. The VoIP call duration is a direct indication to the revenue impact of the network operator. The ADC product is well poised to process the network traffic online to detect and control the VoIP presence, and generate records that can be used to calculate the VoIP call durations.

Random Drop Charging Action

The random drop charging action is added as an option to degrade P2P voice calls. This is achieved by randomly dropping packets of the voice calls over the voice call period.

Voice data is encoded in multiple packets by the codec. Since there is a possibility of packets being dropped in a network, the codec replicates the same information across multiple packets. This provides resilience to random packet drops in the network. For a considerable degradable voice quality, a chunk of packets need to be dropped. By this way, the codec will be unable to decode the required voice information. The chunk size for achieving degradation of voice call varies from one protocol to another.

The Random Drop decision has to be made once for a chunk of packets. By choosing the random drop time from a configured range, the drop is achieved at random seconds within a configured range. The packets will drop within a known period of time. For example, if a voice call happens for 2 minutes and if we configure a drop interval of 12–15 seconds, then a packet will be dropped within the first 15 seconds of the voice call.

Important: This feature is applicable only for VOIP calls.

How ADC Works

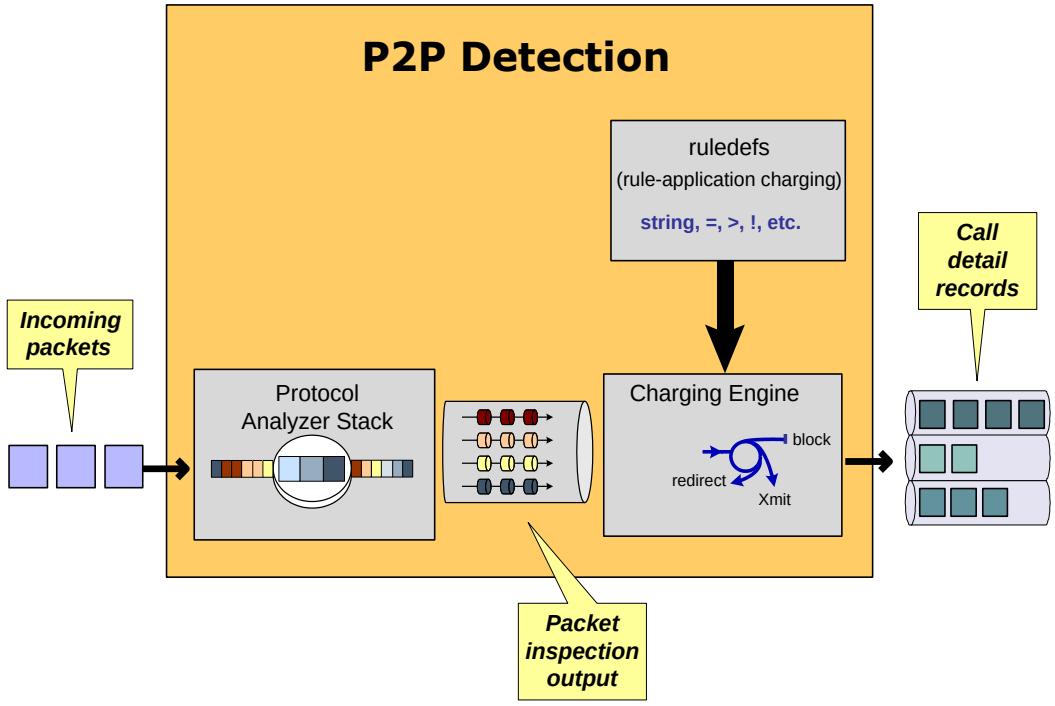
ADC interfaces to a PCRF Diameter Gx interface to accept policy ACLs and rulebases from a PDF. ADC supports real-time dynamic policy updates during a subscriber session. This includes modifying the subscriber's policy rules during an active session by means of ACL name and Rulebase name.

In Rel. 7 Gx interface, a Charging Rulebase will be treated as a group of ruledefs. A group of ruledefs enables grouping rules into categories, so that charging systems can base the charging policy on the category. When a request contains names of several Charging Rulebases, groups of ruledefs of the corresponding names are activated. For P2P rules to work in the group of ruledefs, P2P detection has to be enabled in the rulebase statically.

Static policy is supported initially. A default subscriber profile is assumed and can be overwritten on the gateway. Per-subscriber static policy is pulled by the gateway from the AAA service at subscriber authentication.

The following figure illustrates how packets travel through the system. The packets are investigated and then handled appropriately using ruledefs for charging.

Figure 28. Overview of Packet Processing in ECSv2



Advantages of P2P Processing Before DPI

The advantages of P2P processing before DPI:

- Some protocols like BitTorrent and Orb use HTTP traffic for initial setup. If P2P analysis is done after HTTP, it is possible that these protocols may go undetected.

- Protocols like Skype use well known ports (like 80 & 443). In these scenarios, the HTTP engine reports these as invalid packets. For protocol detection, it is desirable to have P2P detection before Deep Packet Inspection (DPI).
- Stateless detection of protocols based on signature will be easier when the P2P analysis is done before DPI.

ADC Session Recovery

Intra-chassis session recovery is coupled with SessMgr recovery procedures.

Intra-chassis session recovery support is achieved by mirroring the SessMgr and AAAMgr processes. The SessMgrs are paired one-to-one with the AAAMgrs. The SessMgr sends checkpointed session information to the AAAMgr. ACS recovery is accomplished using this checkpointed information.

Important: In order for session recovery to work there should be at least four packet processing cards (PSCs/PSC2s), one standby and three active. Per active CPU with active SessMgrs, there is one standby SessMgr, and on the standby CPU, the same number of standby SessMgrs as the active SessMgrs in the active CPU.

There are two modes of session recovery, one from task failure and another on failure of CPU or PSC/PSC2.

Recovery from Task Failure

When a SessMgr failure occurs, recovery is performed using the mirrored "standby-mode" SessMgr task running on the active packet processing card. The "standby-mode" task is renamed, made active, and is then populated using checkpointed session information from the AAAMgr task. A new "standby-mode" SessMgr is created.

Recovery from CPU or PSC/PSC2 Failure

When a packet processing card hardware failure occurs, or when a planned packet processing card migration fails, the standby packet processing card is made active and the "standby-mode" SessMgr and AAAMgr tasks on the newly activated packet processing card perform session recovery.

Limitations

This section lists the limitations of ADC in this release.

BitTorrent

- Some clients (like Azureus 3.0) provide an advanced user interface which can include an embedded web browser. These are not detected as BitTorrent. Also other features like chat or instant messaging are not detected as BitTorrent. These features are client specific and not related to the BitTorrent protocol.
- Certain clients also display advertisements. These images are downloaded through plain HTTP and are not detected as BitTorrent.

eDonkey

- The eDonkey client eMule supports a protocol named Kademlia. This protocol is an implementation of a DHT (Distributed Hash Table). Kademlia is only used for searching new peers which have the file the user wants to download. The download itself uses the eDonkey protocol. However, the Kademlia protocol is not detected as eDonkey.
- The eDonkey client eMule supports a text chat that is not detected as eDonkey.

FastTrack

SSL packets and HTTP packets from the Kazaa client is not detected. Only data transfer is detected.

Gadu-Gadu

Radio traffic passes through HTTP and is not detected.

Gnutella / Morpheus

- Some of the clients that use Gnutella protocol for file sharing can also use other file sharing protocols. The part of traffic that follows Gnutella Protocol will only be detected as Gnutella.
- Client specific patterns which are not part of the Gnutella Protocol will not be detected as Gnutella. UDP contributes to about 20-30 % of most Gnutella clients. Detection is based on some strange patterns in the first packet of these UDP flows. Untested Gnutella clients may have more strange patterns, causing drop in the detection %.
- The Morpheus Client creates a lot of TCP flows, without any string pattern in the application header. These flows are not currently detected.

Jabber

- Most clients that use Jabber for IM offer other services like Voice Call or File Transfer. These services are not detected as Jabber.
- Jabber with SSL encryption cannot be detected, because it uses SSL.

MSN

MSN HTTP downloads such as MSN Games and MSN Applications are not detected. Traffic from these MSN applications use a different protocol for their traffic.

Skype

- The Skype detection cannot detect traffic of most of the third-party plug-ins. The plug-ins use Skype only for marketing and presentation purposes such as opening a window within a Skype window or modifying the main Skype window with buttons or sounds. These plug-ins do NOT use the Skype protocol to transfer data over the network.
- Other than Skype Voice, all detected Skype traffic is marked as Skype. Distinctions between different data types within Skype (i.e. text chat, file transfer, and so on) cannot be made.

- Skype voice detection may not be accurate if it happens with other traffic (file transfer, video, etc.) on the same flow.

Winny

The Winny client also supports bbs. This is currently not detected.

Yahoo

Yahoo! HTTP downloads for yahoo games, images and ads that come during yahoo messenger startup are not detected as Yahoo!. If configured, these can be passed on to the HTTP analyzer for HTTP Deep Packet Inspection.

Other Limitations

- Most of the heuristic analysis for a subscriber is stateful and depends on building an internal state based on certain patterns seen by the analyzer. Patterns occur over multiple packets in a single flow and over multiple flows for a subscriber. If the system loses the state (due to a task failure for example), then the detection can fail for the affected subscribers/flows after recovery.

  Most P2P protocols emit these patterns regularly (sometimes as early as the next flow created by the application). When the system sees the pattern again, it re-learns the subscriber state and starts detecting the protocol.
- In this release, P2P rules cannot be combined with UDP and TCP rules in one ruledef.

Chapter 8
ASN Gateway Overview

Access Service Network Gateway (ASN Gateway) is the subscriber-aware mobility access gateway for IEEE 802.16 mobile WiMAX radio access networks. These carrier- and enterprise-class platforms provide exceptional reliability and performance characteristics for mobile WiMAX operators.

The ASN Gateway provides inter-technology mobility for 3GPP, 3GPP2, DSL, and WiFi access technologies. This assures common billing and seamless inter-technology handover.

ASN Gateway is available for all chassis running StarOS Release 7.1 or later.

Important: The ASN Gateway is a licensed product that requires an Access Service Network Gateway support license.

ASN Gateway provides the following functionality, all of which is integrated with the chassis:

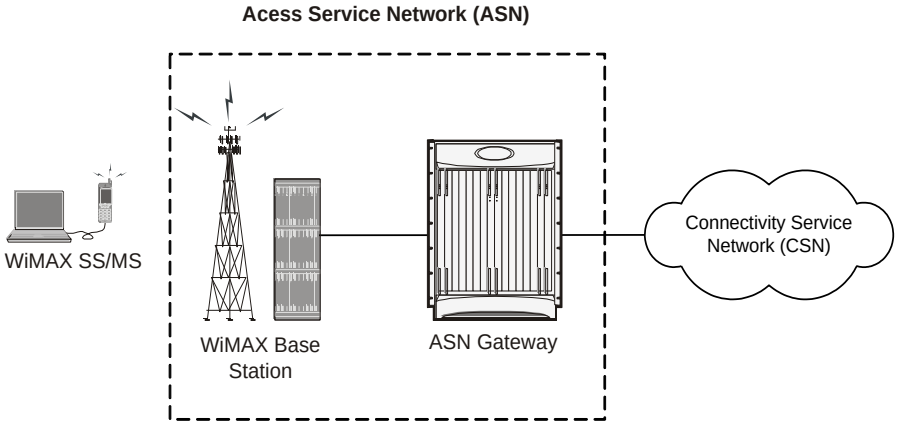
- ASN Mobility Managment
- DHCP proxy server
- DHCP Relay
- Connectivity Service Network (CSN) mobility
- Intra-ASN and inter-ASN handover
- Paging controller/location register
- Radio resource controller relay function
- Service Flow Authenticator (SFA)
- Proxy-Mobile Internet Protocol (P-MIP) client
- Mobile IP Foreign Agent (MIP FA) protocol
- Data path function
- Context server function
- Handover relay function
- WiMax NSP-ID functionality
- 802.1P QoS support
- RADIUS-based prepaid accounting and hotlining for WiMAX
- DHCP relay support for ASNGW
- Creation, modification, and deletion of pre-provisioned/dynamic service flows

ASN Mobility Management

The Access Service Network Gateway (ASN Gateway) processes subscriber control and bearer data traffic. It also supports connection and mobility management across cell sites and inter-service provider network boundaries. An ASN Gateway is a logical entity in the Access Service Network (ASN) of a WiMAX radio access network and interfaces directly with base transceiver station or base station via an R6 GRE reference interface. An ASN Gateway performs control plane functions, bearer plane routing or bridging functions, resident functions in the connectivity service network, or a function in another ASN.

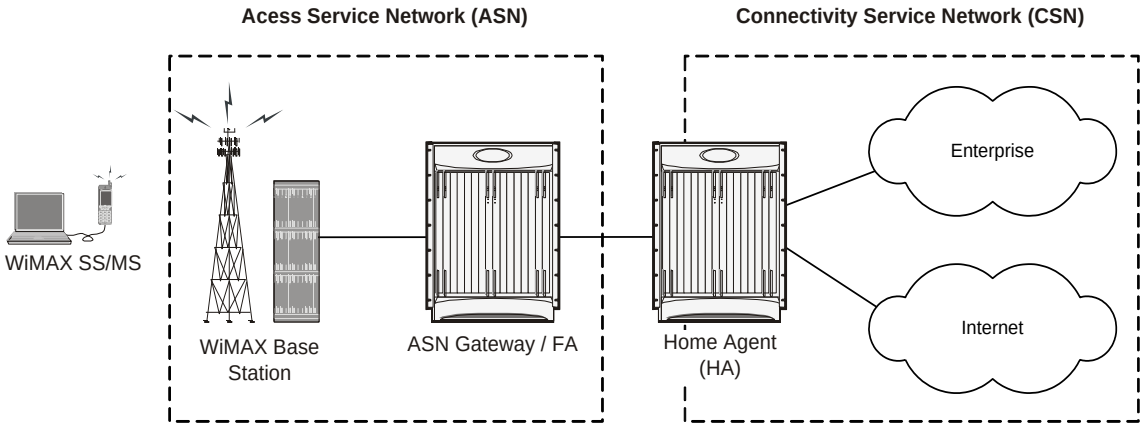
The ASN Gateway is placed at the edge of an ASN and is the link to the CSN. Each ASN Gateway can concentrate traffic from multiple radio base stations. This reduces the number of devices to manage and minimizes connection set-up latency by decreasing the number of call handovers in the network.

Figure 29. Basic ASN Gateway Network



To support Mobile IP and/or Proxy Mobile IP data applications, you can configure the system to perform the role of the ASN Gateway/foreign agent and/or the home agent within the connectivity service network (CSN) of your WiMAX data network. When functioning as a home agent, the system can be located within your WiMAX network or in the CSN of an external enterprise or ISP network. In either case, the ASN Gateway/foreign agent terminates the mobile subscriber's call session and then routes the subscriber's data to and from the appropriate home agent.

Figure 30. Basic ASN Gateway Mobile IP Network



EAP User Authentication

The ASN Gateway serves as the Extensible Authentication Protocol (EAP) authenticator and mobility key holder for subscriber connections and RADIUS clients to attached Authorization, Authentication, and Accounting (AAA) servers.

ASN Gateway and AAA

ASN control is handled by the ASN Gateway and the base station. The ASN Gateway control plane handles the feature set, including AAA functions, context management, profile management, service flow authorization, paging, radio resource management, and handover. The data plane feature set includes mapping radio bearer to the IP network, packet inspection, tunneling, admission control, policing, QoS, and data forwarding.

The ASN Gateway acts as an authenticator. It operates in pass-through mode for EAP authentication between the EAP client (the mobile station) and the EAP (AAA) server. After successful EAP authentication, the AAA server sends the master session key (MSK) to the ASN Gateway. The ASN Gateway, as authenticator, performs authorization key (AK) context management. It derives the AK from the MSK and sends it to the base station. As part of the AK context, other information, such as the AkID and CMAC are sent to the base station to secure the R1 interface.

An AAA module in the ASN Gateway provides flow information for accounting. Every detail about a flow, such as the transferred or received number of bits, the duration of the connection, and the applied policy, is retrievable from the data plane.

Profile Management

The ASN Gateway provides profile management and a policy function that resides in the connectivity network. Profile management identifies a subscriber's feature set, such as the allowed QoS rate, number of flows, and type of flows.

In addition, the ASN Gateway maintains a context for the mobile subscriber and the base station. Each subscriber's context contains the subscriber's profile and security context, and the characteristics of the subscriber's mobile device. The subscriber's context is retrieved and exchanged between the serving base station and a target base station during handover.

The ASN Gateway authorizes service flows according to the subscriber's profile. Allowed service flows and active service flows can change over time, so the ASN Gateway provides admission control for downlink traffic. The ASN Gateway creates a GRE tunnel per service flow.

Inter-ASN Handovers

During a handover, the ASN Gateway provides the subscriber's context to a target base station and when requested, changes the data path. To minimize latency and packet loss, the ASN Gateway implements data integrity through bi-casting or multi-casting. For paging, buffering is also supported. A foreign agent maintains the IP connectivity if the mobile subscriber initiates an inter-ASN handover. The ASN Gateway supports either Proxy-Mobile IP (PMIP) or Client-Mobile IP (CMIP) in order to communicate with home agents.

The ASN Gateway maintains location information to provide the paging service that tracks subscribers when they are operating in idle mode. If there is any download traffic, ASN Gateway requests the PC to trigger paging. During active operation, location information is also updated as the mobile subscriber moves to a new base station.

Supported Features

The Access Service Network Gateway (ASN Gateway) provides ASN Gateway control and bearer plane routing functions:

- BS Interface: R6 IP/GRE bearer plane
- Inter-ASN handovers to other ASN Gateways: R4 IP/GRE bearer plane
- Interactions with AAA management or policy servers: R3 RADIUS interface
- Mobile IP Interface to HA in Connectivity Service Network: R3 IP-in-IP tunneling

A Profile C ASN Gateway is one of three alternative designs for radio resource management proposed by the WiMAX Forum. In Profile C architecture, the handover control component resides in the base stations. The ASN Gateway represents a transparent message relay point between neighboring base stations. The Radio Resource Controller (RRC) component in every BTS periodically polls its neighbors to build a resource availability database that it checks prior to triggering call handovers.

Profile C provides a high performance ASN Gateway platform with the following supported features in the current software version.

Important: Not all features are supported on all platforms.

Simple IPv4 Support

A Simple IP model supports non-mobile IP terminals and provides ASN-anchored mobility for fixed, nomadic, or portable mobility applications. A Simple IP architecture removes dependencies for separate foreign agent and home agent functions. ASN Gateway handles simultaneous combinations of Simple IP, Mobile IP, or Proxy Mobile IP calls. A Simple IP model permits the ASN to be combined or split from the CSN, depending upon the need for roaming. The Simple IP implementation includes a DHCP Proxy Server function for local or AAA-provided IP address assignment.

Simple IP provides a solution for stationary wireless DSL-like applications. It enables mobility on intra-ASN handovers between neighboring base stations and permits inter-ASN mobility via an R4 interface between ASN Gateways.

DHCP Proxy Server

Compared to 3G wireless technologies such as EV-DO (Evolution-Data Optimized) or PDP (Packet Data Protocol) Type PPP (Point-to-Point Protocol) contexts in General Packet Radio Service/Wideband Code division Multiple Access (GPRS/W-CDMA) networks, WiMAX networks do not use a PPP data link layer between access devices and the ASN Gateway. An alternative approach to IP address allocation is needed in Simple IP and Proxy Mobile IP usage models.

The ASN-GW includes a DHCP proxy/server/relay that interacts with the DHCP client function on the access device. In a Simple IP usage model, the DHCP server allocates dynamic addresses from a local address pool or fetches static addresses from subscriber profiles during authentication from a AAA server. Alternatively, the ASN-GW uses a DHCP relay process to forward the DHCP request to an external DHCP server.

In a Proxy Mobile IP use case, the ASN-GW uses a DHCP proxy to trigger a local foreign agent function to initiate a Mobile IP Request via the R3 interface to a home agent. The home agent returns the address via the Mobile IP Response. The DHCP Proxy component on the ASN Gateway conveys the address in a DHCP Response message to the DHCP client running on the user's access device.

This solution enables mobility on intra-ASN handovers between neighboring base stations. It also permits inter-ASN mobility via an R4 interface between ASN Gateways.

DHCP Relay Support

Following are the enhancements and changes to the existing functionality of the DHCP relay. Refer to the DHCP Relay section for details.

- DHCP Module enhancements, including the addition of DHCP Relay Agent option in the relayed messages to the DHCP server, DHCP Auth sub-option in the relayed messages to the DHCP server. and decoding of the DHCP Auth sub-option from the DHCP server messages
- DHCP Relay for SIP and PMIPv4 calls
- DHCP Relay support for Init reboot scenario
- DHCP NAK is also handled through the DHCP Relay
- AAA related enhancements, including receipt of the DHCP server address, DHCP-RK, DHCP-Key-ID attributes from the AAA, and decryption of the DHCP-RK
- Session recovery support for SIP and PMIPv4 call lines with DHCP Relay functionality
- CLI related enhancements, such as changes to subscriber configuration to enable the DHCP relay option, and configuration of DHCP Server Address, and DHCP-RK and DHCP-Key-ID in the DHCP service

ASN Gateway Micro-Mobility

ASN Gateway micro-mobility provides ASN Gateway-anchored L2 handovers. This low-latency procedure assures the seamless mobility of mobile access devices within a WiMAX network. The ASN Gateway supports both uncontrolled and controlled handovers for micro-mobility.

Uncontrolled Handovers

In an uncontrolled handover scenario, a mobile subscriber attempts to re-enter the WiMAX network at a target base station without the handover preparation procedures with the serving base station.

In order to authenticate the roaming user, the target base station obtains the subscriber and security context information from the serving ASN. The anchor authenticator ASN Gateway conveys the context response message and assists in the establishment of a new R6 GRE bearer connection to the target base station. It is referred to as an L2 operation because the previously assigned IP address for the binding remains the same on the anchor authenticator/data path ASN Gateway while the L2 BSID (Ethernet MAC address) is updated for the target base station.

Uncontrolled handovers are supported for both Simple IP or Mobile IP use cases. With uncontrolled L2 handover procedures, interactive and non-real-time applications incur minimal performance degradation and packet loss during subscriber movement between cell sites.

Controlled Handovers

A controlled handover occurs when a subscriber access device explicitly requests handover assistance from the serving base station to a new target base station. This process minimizes packet loss to the WiMAX access device. During the handover request, the serving base station provides the subscriber's context information to the anchor authenticator ASN Gateway and a list of target base stations that are preferred by the mobile device. Upon a successful response from potential target base stations, the anchor authenticator ASN Gateway initiates a data path for the mobile subscriber to

the target base station. It also transfers all contextual information for the session to the target base station. The downlink traffic for the mobile subscriber is simultaneously broadcast and subsequently buffered by each of the target base stations.

Controlled handovers may be triggered by the mobile access device or the serving base station as a congestion overload control mechanism.

Controlled handovers and associated data path pre-registrations minimize the impact on performance to a greater extent than uncontrolled handovers and significantly reduce datapath outages.

WiMAX R4 Inter-ASN Mobility Management

R4 inter-ASN mobility management procedures enable low latency call handovers between neighboring ASN Gateways located in different geographical regions or different operator networks. During mobility operations, the call is anchored on the anchor authenticator ASN Gateway. When a mobile subscriber roams to a destination cell site, the target base station connects to the anchor gateway over the serving ASN Gateway's R4 interface. The R4 interface provides control functions such as security context transfers and IP/GRE bearer level connections. The data conveyed to the subscriber by the remote hosts is subsequently tunneled over R4 by the anchor authenticator gateway to the serving gateway. The current ASN Gateway implementation supports the co-existence of anchor authenticator and anchor datapath functions in the same ASN Gateway.

Supported R4 functionality includes:

- R4 over Simple IP connections
- R4 over Mobile IP connections
- Anchor Gateway bi-casting over simultaneous R6 and R4 sessions
- Co-location of DHCPv4 Proxy and PMIPv4 FA on anchor authenticator gateway
- Support for multiple QoS service flows per-session via R4 tunnels

Important: Both the anchor gateway session and non-anchor gateway sessions are counted towards the session license separately. Licensed session limits are enforced based on the total number of anchor and non-anchor sessions.

WiMAX R3 CSN Anchored Mobility Management

The R3 reference point defines a set of control plane protocols between the Access Service Network (ASN) and Connectivity Service Network (CSN) to support AAA, policy enforcement, and mobility management functions. The R3 reference interface is used in a mobile IP application with the home agent acting as the call anchor point. In contrast to L2-based ASN anchored mobility procedures, CSN anchored mobility is L3-based and supports both proxy mobile IP and mobile IP calls. The R3 interface uses mobile IP signaling and IP-in-IP tunneling or GRE tunneling and includes standard features such as dynamic Home of Address (HoA) address allocation. Mobility signaling messages are authenticated by the home agent based on a dynamic user identity called a pseudo-NAI which changes after each authentication.

Mobile IP applications are well suited for inter-provider roaming applications and inter-technology handovers such as WiMAX-HRPD Rev A, WiMAX-WiFi, and WiMAX-W-CDMA. Mobile IP also provides an attractive solution for operators with a heterogeneous radio access network who want to support seamless mobility across base transfer stations from multiple RAN suppliers.

Important: Support for this function requires the HA feature license key.

Proxy Mobile IPv4 (PMIPv4)

The P-MIP procedure is designed only for Simple IP-capable access devices for which mobility procedures are performed entirely in the network. Certain events on the access device require relocation of the L3 anchor point (for example, CoA). One case is for the initial connection establishment in which the home agent or H-AAA server assigns an IP address and generates the mobility binding. Another is when the mobile subscriber roams across cell sites or ASNs and attaches to a target ASN Gateway.

Client Mobile IPv4 (CMIPv4)

CMIPv4 provides mobility procedures for mobile IP-capable access devices. In contrast to PMIPv4, where stateful DHCP proxy signaling triggers R3 signaling between the ASN Gateway and the home agent, CMIPv4 uses agent advertisement between the foreign agent component in the ASN Gateway and mobile IP client on subscriber access device. Mobile IP signaling occurs directly between the access device and the anchor foreign agent component in the ASN Gateway.

Authenticator

The authenticator function in the ASN Gateway acts as an anchored authenticator for a subscriber for the duration of the session. For example, as a subscriber moves between base stations served by the ASN Gateway, the authenticator anchor remains stationary. If a subscriber moves to a base station served by a different ASN Gateway, the anchor authenticator is hosted at that ASN Gateway. If the R4 interface is not supported between both gateways, only the subscriber needs to be re-authenticated.

The RADIUS client for authentication and accounting is collocated with the authenticator function. The ASN Gateway acts as an EAP relay and is agnostic to the EAP method. EAP transport between the ASN Gateway and the base station is performed as a control exchange. The base station functions as an EAP relay, converting Pair-wise Master Key version 2 (PKMv2) to the EAP messages for the ASN Gateway. The ASN Gateway works in pass-through mode and any EAP method that generates keys, such as MSK or EMSK, is supported in the system.

PKMv2 performs over-the-air user authentication. PKMv2 transfers EAP over the IEEE 802.16 air interface between the MS and the base station. The base station relays the EAP messages to the authenticator in the ASN Gateway. The AAA client on the authenticator encapsulates the EAP message in AAA protocol packets, and forwards them through one or more AAA proxies to the AAA server in the CSN of the home NSP. In roaming scenarios, one or more AAA brokers with AAA proxies may exist between the authenticator and the AAA server. AAA sessions always exist between the Authenticator and AAA server, with optional AAA brokers providing a conduit for NAI realm-based routing.

EAP Authentication Methods

WiMAX networks use Ethernet as the L2 protocol for network access authentication. The Extensible Authentication Protocol (EAP) provides the network authorization function. The ASN Gateway represents the EAP authenticator and supports a transparent relay point between the EAP client on the subscriber access device and EAP server on the AAA. The ASN Gateway triggers an EAP-identity request to the subscriber device. The subscriber device responds with an EAP-identity response. It subsequently unpacks EAP messages over the R6 interface and transfers them via RADIUS or Diameter signaling to the AAA server.

EAP authentication provide multiple authentication methods that can be tailored to the operator's preference toward user-level, device-level, or user- and device-level network authorization. At the H-AAA server in Home Network Service Provider (H-NSP), device-level authentication in a roaming application guards against unauthorized network access by users with stolen access devices.

Supported RADIUS Methods

ASN Gateway supports following EAP authentication and authorization methods using RADIUS:

- EAP-Pre-shared Key (EAP-PSK)
- EAP-Transport Layer Security (EAP-TLS)
- EAP-Tunneled Transport Layer Security (EAP-TTLS)
- EAP-Authentication and Key Agreement (EAP-AKA)

EAP-Pre-shared Key (EAP-PSK)

EAP-PSK is a symmetric mutual authentication method that uses manually provisioned pre-shared keys between an EAP client on an access device and an EAP server component on AAA. The size of the pre-shared key can be up to 256 bytes.

EAP-Transport Layer Security (EAP-TLS)

EAP-TLS is an asymmetric authentication method that uses X.509 digital certificates, for example public/private key pairs, and enables device-based authentication.

EAP-Tunneled Transport Layer Security (EAP-TTLS)

EAP-TTLS is a multi-level authentication scheme to enable device and user-based authentication. The first level handshake provides device-level authentication and uses the same encryption and ciphering algorithms as EAP-TLS. The device can, but is not required to, be authenticated via a CA-signed PKI certificate to the server. The secure connection established through the first level handshake is then extended with MS-CHAP-V2 authentication to verify user credentials. As with other EAP methods, successful EAP transactions at AAA result in a Master Session Key (MSK) that is returned over an encrypted connection. The ASN Gateway uses the key to generate a derivative key for securing the air interface between ASN and user access device.

EAP-Authentication and Key Agreement (EAP-AKA)

EAP-AKA uses symmetric cryptography based on pre-shared private client/server keys and challenge-response mechanisms similar to other EAP methods. It verifies credentials for users of Removable User Identity Modules (R-UIMs).

Supported Diameter Methods

ASN Gateway supports the following Diameter methods for EAP authentication and authorization:

EAP-Authentication and Key Agreement (EAP-AKA)

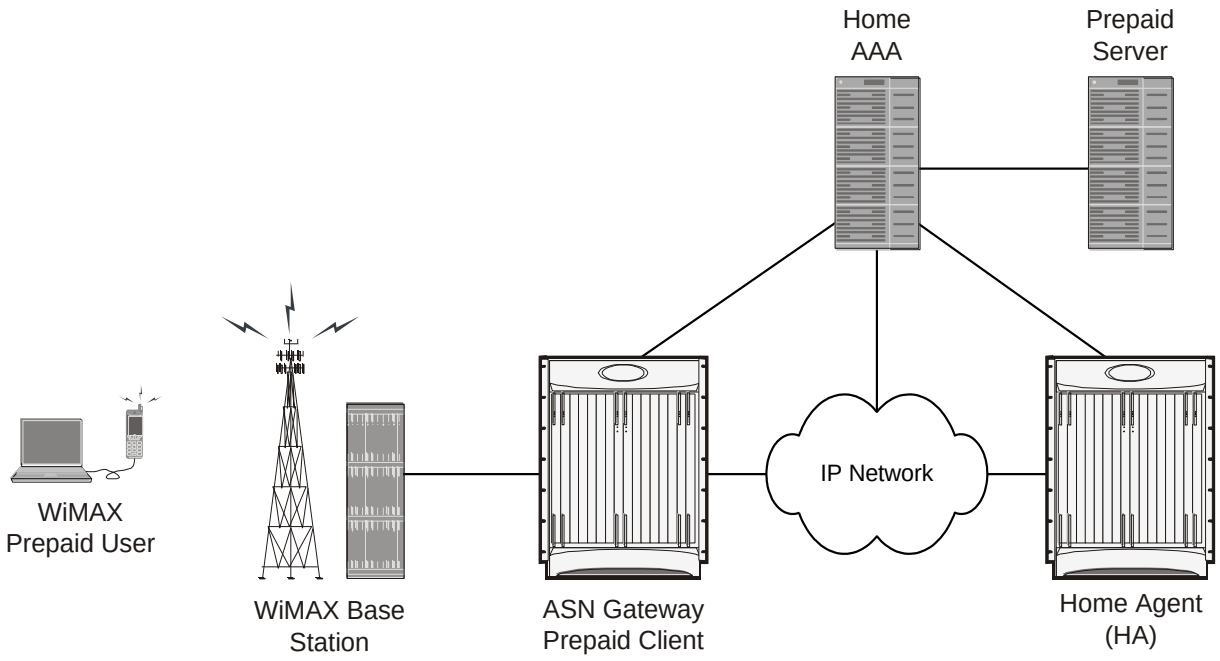
EAP-AKA uses symmetric cryptography based on pre-shared private client/server keys and challenge-response mechanisms similar to other EAP methods. It verifies credentials for users of Removable User Identity Modules (R-UIMs).

WiMAX Prepaid Accounting

The system supports prepaid accounting for clients on the ASN Gateway.

Clients can communicate directly to a home AAA server or be proxied through a visited network's AAA server. The following figure shows a typical prepaid network topology.

Figure 31. Prepaid Network Topology



Volume and Duration-based Prepaid Accounting

Prepaid accounting is a licensed-enabled feature. The ASN Gateway supports both volume threshold and duration threshold based prepaid accounting. Even though session-level accounting is performed for both volume and duration, the number of bytes in a multi-flow session is applied to a duration-based configuration.

RADIUS attributes identify thresholds and quotas for both volume (number of bytes) and duration (length of session).

Supported Enhanced Features

All enhanced features described in this section require the appropriate feature license keys.

Lawful Intercept

The Cisco Lawful Intercept feature is supported on the ASN Gateway. Lawful Intercept is a licensed-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your local Cisco account representative.

Intelligent Traffic Control

Intelligent Traffic Control (ITC) supports customizable policy definitions. The policies enforce and manage service level agreements for a subscriber profile, thus enabling differentiated levels of services for native and roaming subscribers.

ITC includes features such as traffic prioritization, for example, marking DiffServ codepoints to enable unique treatments for the five WiMAX classes of service, queue redirection, and per-subscriber/per-flow traffic bandwidth control. Traffic policing enables maximum rate-based services and tiered bandwidth charging models. ITC includes a local policy engine that runs on an ASN Gateway in a Simple IP usage model, or as a home agent in a Mobile IP application. You can configure ITC policies statically with Class-Maps to identify applications flows that use L3/L4 5-tuple identifiers. You can then apply the resulting policy actions through policy maps and policy groups. The detection and programming of the local policy engine can alternatively be triggered on network access at the ASN Gateway as it retrieves QoS profiles for each authenticated user.

This feature provides a policy mechanism so you can enable user entitlements and provision treatments for native users and applications relative to roaming subscribers, Mobile Virtual Network Operators (MVNOs), and offnet P2P traffic.

Hotlining/Dynamic RADIUS Attributes

WiMAX is an all IP-based networking technology in which mobile operators seek a more profitable business model. One way to do this is to avoid traditional device subsidization that accompanies the sale of locked devices that restrict access to provisioned subscribers of an operator's network. The WiMAX Forum has proposed remote Over-the-Air (OTA) activation protocols such as Open Mobile Alliance Device Management (OMA DM) to enable self-provisioned, self-configured, retail subscription models.

The ASN GW supports hotlining on a session basis. This capability is enabled by default. The rule-based hotlines use an IP redirection rule with the standard attribute Filter-ID. The server sends the ACL names in the Filter-ID attribute. This in turn, locates the rules.

Upon receiving a RADIUS Access-Accept message containing the Filter-ID attribute, the ASN GW locates the rule list, using the name contained in Filter-ID, and applies them to the session.

Configure the rules locally on the ASN GW under ACL groups.

In this scenario:

- A user with an unprovisioned access device registers with a special decorated NAI that represents him/her as a non-subscriber to the AAA.
- The AAA grants limited network access by returning a hotlining filter rule to the ASN Gateway. ASN GW hotlining support uses the standard attribute Filter-ID, along with the session identification parameters User-Name, Calling-Station-ID, and AAA-Session-ID.
- An IP address is assigned during initial network entry. The ASN Gateway uses the redirect address associated with the filter rule to hotline the call to a web activation portal.
- The user profile and subscription activation process is completed. The call is forwarded to the OMA DM server.
- The OMA DM server triggers a network-initiated bootstrapping session with the OMA DM client on the user access device.
- The OMA DM uses XML messaging over a secure OTA connection to remotely configure the access device.
- If a session and an ACL list are located, the rules are applied to the session and a COA-ACK is returned. The AAA server transmits a RADIUS message to the ASN Gateway instructing it to "unhotline" the session.
- At this point, the user is a known subscriber to the back-end subscription database and is granted unrestricted access to the network.

This feature facilitates a non-subsidized retail activation model through over-the-air user-driven subscription and remote device configuration. It also prevents unprovisioned users unrestricted access to the wireless operator's network. This is a complementary technique you can use with operator fraud prevention systems by quarantining fraudulent user sessions or redirecting them to a billing/web portal.

Multi-flow QoS

Within a WiMAX ASN, QoS enforcement is administered by the Service Flow Authorization (SFA) component in the ASN Gateway (also referred to as Anchor Policy Charging Enforcement Function, or A-PCEF). SFA provides traffic management and QoS policy management for subscriber service flows.

Multi-flow QoS enables the establishment of static traffic policies for various subscriber application level service flows. It can be used in Simple IP or Mobile IP usage scenarios. The policies are stored in a Subscriber Policy Repository (SPR) database and retrieved as authenticated QoS profiles by the ASN Gateway. The A-PCEF negotiates via R6 with the Service Flow Manager (SFM) function on the base station. If the authorized QoS profile matches the available base station resources, the request is granted. The A-PCEF provides the following:

- Traffic classification
- Admission control
- Prioritization (DSCP marking)
- Per-session/per-flow bandwidth control
- Flow mapping across application-specific R6/R4 GRE tunnels

In conjunction with multiflow QoS, the ASN Gateway offers configurable accounting on a per-session, per-R6, or per-service flow basis. Multi-flow QoS enables the OFDM radio access connection to be separated into multiple logical Connection ID's (CIDs) with each pair of forward and reverse sub-channels transporting one or more application flows.

ASN Gateway supports pre-provisioned as well as dynamic service flows. A total of up to 12 uni-directional service flows per subscriber R6 or R4 session are possible.

Multi-flow QoS provides enhanced an user experience via end-to-end differentiated QoS connection-oriented services and stringent treatment for isochronous voice and delay-sensitive multimedia applications over broadband WiMAX networks. This feature also enables service convergence and is the foundation for delivery of IMS service control.

802.1P QoS Support

Quality of Service (QoS) is a method to better handle the challenges of reliability and quality despite increasing bandwidth demands. 802.1p which is a signaling technique for prioritizing network traffic at the data-link/MAC sublayer (OSI Reference Model Layer 2).

Network administrators have two major types of Quality of Service (QoS) techniques available. They can negotiate, reserve, and hard-set capacity for certain types of service, or simply prioritize data without reserving any capacity setting.

The 802.1p header includes a three-bit field for prioritization, which allows packets to be grouped into various traffic classes. The packets contain a 32-bit tag header located after a destination and source address header. IEEE 802.1p-compliant switches pick up this tag, read it, and put the packet in the appropriate priority queue. No bandwidth is reserved nor requested with this technique.

There are eight priority levels, numbered 0 throuh 7, and consequently, eight possible queues that can be created. Level 7 represents the highest priority and is assigned to mission-critical applications. Levels 6 and 5 are designed for delay-sensitive applications such as interactive video and voice. Levels 4 to 1 are suitable for regular enterprise data transfer, as well as streaming video. Level 0 is assigned to traffic that can tolerate a best-effort protocol.

For more information and configuration examples, see "ASN Gateway QoS and Service Flow Configuration" in this guide.

ASN Gateway Intra-Chassis Session Recovery

This feature enables the system to recover from single software or hardware faults without interrupting subscriber sessions or losing accounting information. Intra-chassis session recovery uses regular task check-pointing of active call states to insure that the fail-over task has the identical configuration and state as the failed process.

Session recovery is supported for the following major features:

- Simple IP, Proxy Mobile IP or Client Mobile IP calls
- R6 or R4 control signaling and bearer level subscriber traffic
- Paging Controller/Location Register (PC/LR) idle mode sessions. PC/LR is a licensed-based feature.
- L2TP LAC & LNS tunnels and sessions

Important: Minimum hardware requirements consist of four processing cards (3 Active, 1 Standby). When session recovery is enabled, overall system capacity may be reduced, depending upon configuration.

Intra-chassis session recovery provides hitless in-service recovery that increases system availability. This eliminates the need for the Radio Access Network to re-register large blocks of simultaneous users. It also minimizes the likelihood of revenue leakage due to the failure of network elements.

This feature requires a feature license key for ASN Gateway session recovery.

Robust Header Compression (ROHC)

Header compression is applied to ASNGW service flows when ROHC is enabled on the ASNGW and the MS and the AAA server authorize ROHC for the ASNGW call. ROHC parameter values are negotiated over R6. Unidirectional and bi-directional ROHC service flows are supported.

Supported Inline Services

All inline services described in this section require the appropriate feature license keys.

Enhanced Charging Service

The Enhanced Charging Service (ECS) is an in-line service feature integrated with the system. ECS provides flexible, differentiated, and detailed billing to subscribers by using Layer 3 through Layer 7 packet inspection. ECS can integrate with a back-end billing system. ECS functionality is supported at the point where sessions are anchored—for example, on the ASN Gateway for Simple IP sessions and on the home agent for Mobile IP sessions.

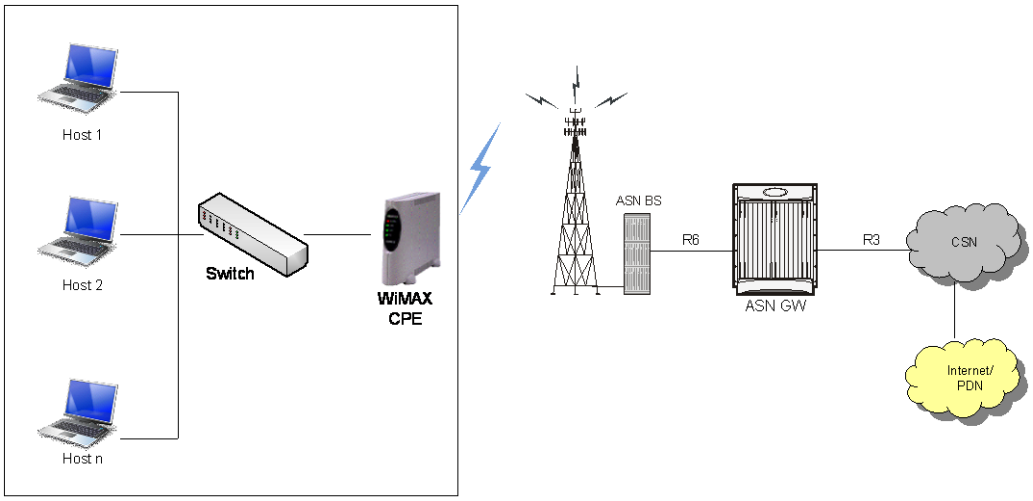
For more information about ECS, refer to the Enhanced Charging Services Administration Guide.

Multi-host Support

ASN Gateway's multi-host feature provides multiple host connectivity.

A WiMAX CPE modem supports multiple IP hosts in fixed/nomadic applications. The modem shares a single WiMAX airlink to connect to the WiMAX IP network. This feature is an effective solution for small or home office users to provide multiple station connectivity through one airlink.

Figure 32. Multi Host Support in WiMAX Network



The WiMAX ASN Gateway allows each WiMAX MS (identified by its 6-byte MSID) to be assigned a single IP address. IP accounting is maintained for the IP address.

How it Works

The DHCP proxy server and the IP pool hosted locally on the ASN Gateway provide the primary IP address from a primary IP pool to the WiMAX customer premise equipment (CPE). The CPE is identified by its WiMAX R6 MSID (6-byte MAC address).

Important: Multiple IP hosts feature is not supported for Proxy-MIP session.

Once a primary IP address is assigned dynamically to the WiMAX CPE, additional IP addresses are assigned dynamically to other IP hosts. Each of the IP hosts is identified by its unique 6-byte MAC address. The DHCP proxy on the ASN Gateway manages the IP addresses by mapping them to the unique MAC addresses supplied by the client in the **chaddr** option field in DHCP DISCOVER or REQUEST messages.

The primary IP address is assigned to the CPE first via DHCP. It is followed by requests for additional IP addresses by individual IP hosts behind the CPE. The ASN Gateway allocates secondary hosts on-demand, up to the configured limit of 4.

Primary IP addresses assigned to WiMAX CPE and secondary IP addresses assigned to the IP hosts, are configured in separate IP pools or the same IP pool. Accounting is based on the primary IP address assigned to CPE and UDR accounting is enabled only for the primary session (flow/session based). No accounting is performed for secondary sub-sessions.

Using the device credentials of the WiMAX CPE, authentication is performed with the EAP-TLS method. There is no authentication for each assigned IP address, and no validation of MAC addresses contained in DHCP requests, except to make sure that they are unique across all subscribers connected to the DHCP proxy server.

IP Address Allocation through DHCP

The dynamic IP address allocation procedure for primary node and secondary hosts is described below:

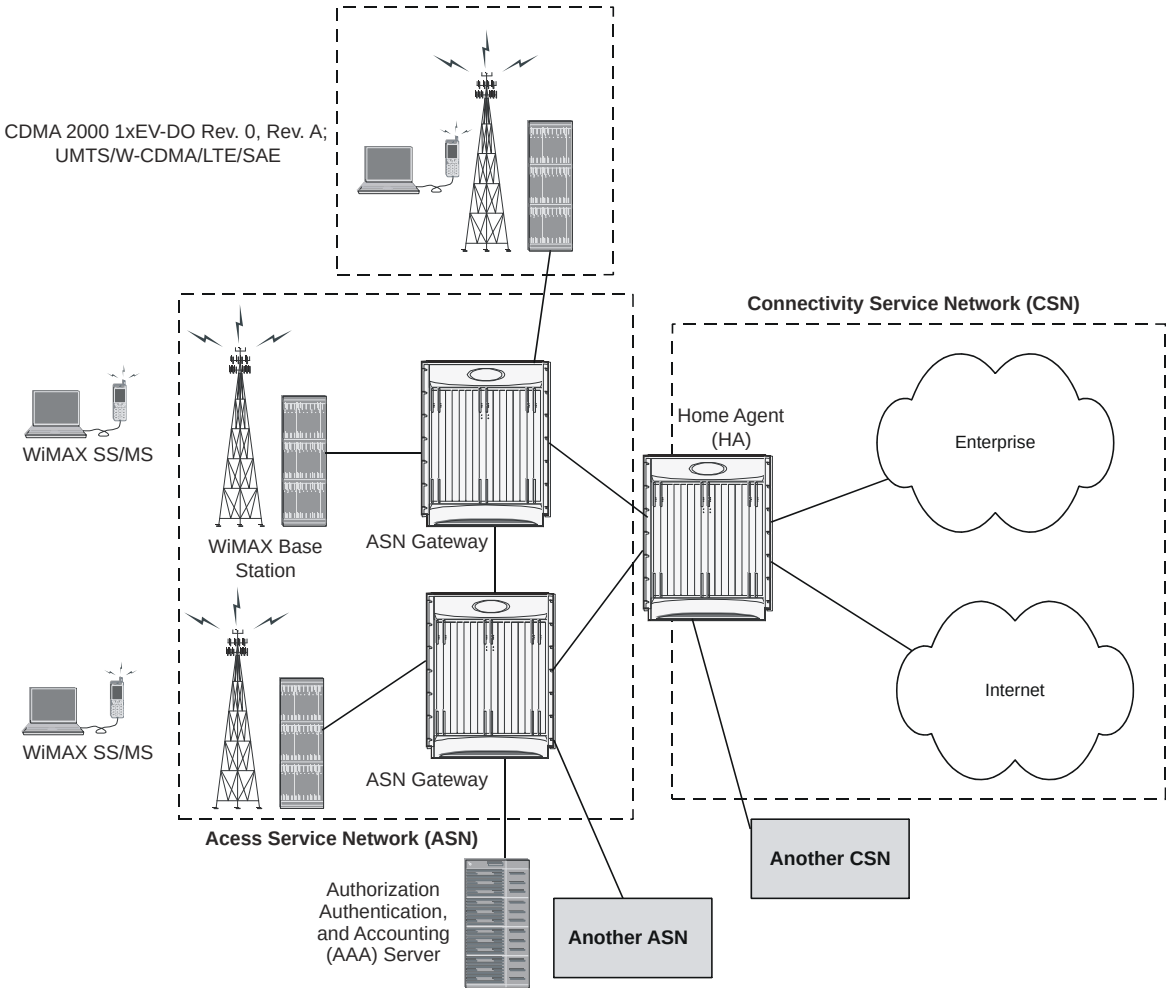
- After the initial network entry for WiMAX CPE is completed, the WiMAX CPE acts as a primary node and starts the DHCP process with the WiMAX ASN Gateway.
- The DHCP proxy server hosted on the ASN Gateway allocates the Primary IP address to the WiMAX CPE as a primary node from the configured primary IP Pool.
- The primary IP address is the first IP address assigned to the WiMAX CPE. The DHCP DISCOVER and REQUEST messages for this must contain the WiMAX R6 MSID as the **chaddr** field. After this IP address is assigned, the session goes into Connected state and is ready to accept DHCP requests for additional IP addresses for other IP hosts.
- Once the primary IP address is assigned to the primary node (WiMAX CPE), hosts behind the CPE start the DHCP process with the WiMAX ASN Gateway for each host mapping to its 6-byte MAC address.
- The DHCP proxy server hosted in the ASN Gateway allocates the secondary IP addresses to the hosts behind the CPE as an auxiliary node from the configured secondary IP Pool.
- When session termination is requested, the primary IP address is the last IP address to be released by the clients and ASN Gateway. This means the primary IP address must be in use and in lease for the session to continue in Connected state. When the Primary IP address is released, the ASN Gateway session is terminated and all IP addresses are freed.
- The auxiliary IP addresses can be assigned and freed any time during the call via DHCP messages.

ASN Gateway in a WiMAX Network

In a WiMAX network architecture, each of the entities, Subscriber Station (SS)/Mobile Station (MS), Access Service Network (ASN) and Connectivity Service Network (CSN) represent a grouping of functional entities.

Each of these functions may be in a single physical device or distributed over multiple physical devices to meet functional and interoperability requirements. The following figure shows a high-level example of WiMAX network architecture

Figure 33. WiMAX Network Architecture



Access Service Network (ASN)

The ASN is an aggregation of functional entities and corresponding message flows associated with the access services. The ASN represents a boundary for functional interoperability with WiMAX clients, WiMAX connectivity service functions, and other vendor-specific functions.

An ASN is defined as a complete set of network functions that provide radio access to a WiMAX subscriber. The ASN provides the following functions:

- WiMAX Layer-2 (L2) connectivity with WiMAX SS/MS
- The transfer of AAA messages to WiMAX subscribers' Home Network Service Provider (H-NSP) for authentication, authorization, and session accounting for subscriber sessions
- Network discovery and the selection of an appropriate NSP from which WiMAX subscribers accesses WiMAX service(s)
- Relay functionality for establishing Layer-3 (L3) connectivity with a WiMAX SS/MS (IP address allocation)
- Radio resource management
- ASN-CSN tunneling

In addition to the above mandatory functions, for a portable and mobile environment the ASN supports the following functions:

- ASN anchor mobility
- CSN anchor mobility
- Paging and location management

The ASN has the following network elements:

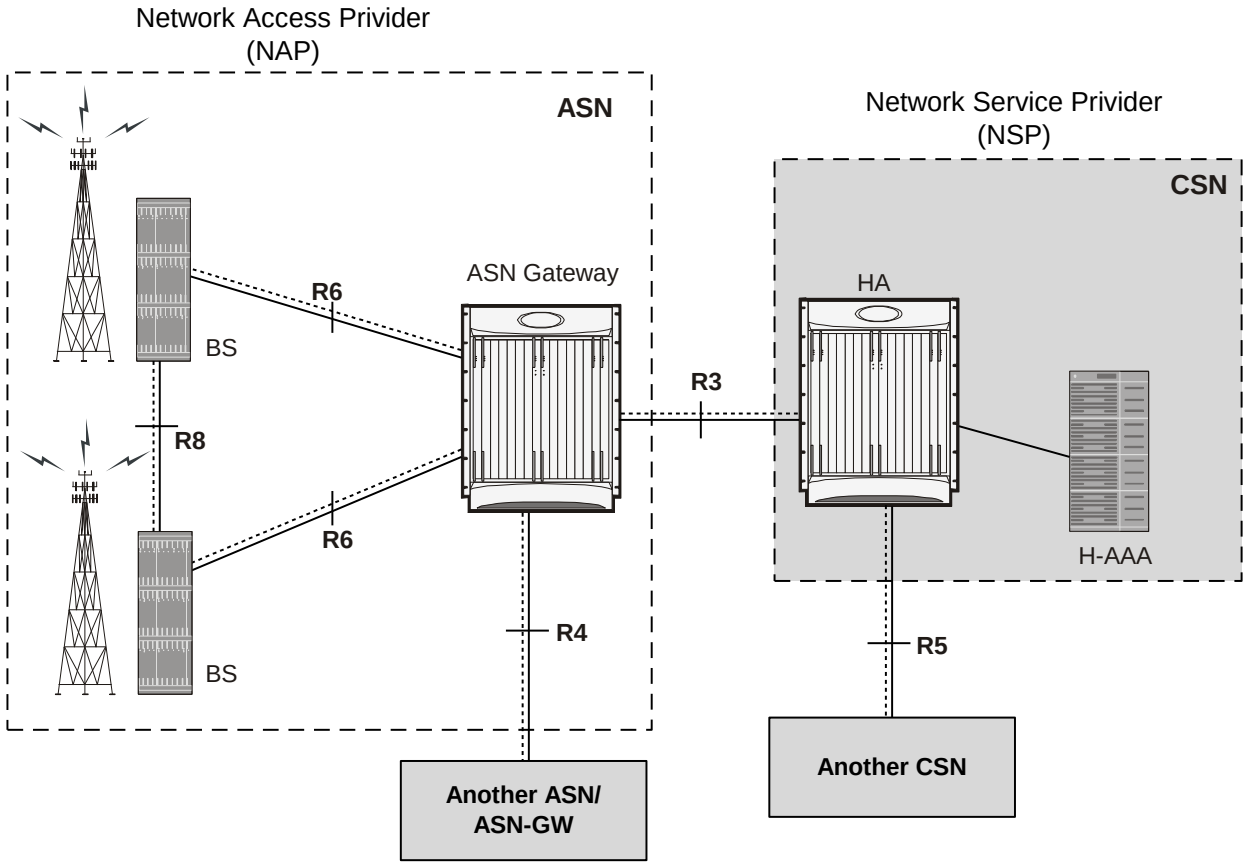
- The WiMAX base station, which is a logical entity that embodies a full instance of the WiMAX Medium Access Control (MAC) layer and physical layer in compliance with the IEEE 802.16 suite of applicable standards. The base station may host one or more access functions and is logically connected to one or more ASN Gateways.
- The ASN Gateway (ASN Gateway), which is a logical entity that represents an aggregation of control plane functional entities. These entities are paired with a corresponding function in the ASN, for example a base station instance, a resident function in the CSN, or a function in another ASN.

The ASN Gateway may also perform bearer plane routing or bridging functions.

The ASN consists of at least one instance of a base station and at least one instance of an ASN Gateway (ASN Gateway). An ASN may be shared by more than one Connectivity Service Networks (CSN).

The ASN decomposition with Network Reference Model (NRM) is shown in the following figure.

Figure 34. ASN Network Reference Model with ASN Gateway



Connectivity Service Network (CSN)

The Connectivity Service Network (CSN) is a set of network functions that provide IP connectivity services to the WiMAX subscriber. A CSN provides the following functions:

- SS/MS IP address and endpoint parameter allocation for user sessions
- Internet access
- AAA proxy or server
- Policy and admission control based on user subscription profiles
- ASN-CSN tunneling support,
- WiMAX subscriber billing and inter-operator settlement
- Inter-CSN tunneling for roaming
- Inter-ASN mobility
- Home agent

The CSN also provides location-based services, connectivity for peer-to-peer services, provisioning, authorization and/or connectivity to IP multimedia services, and support for lawful intercept services in the WiMAX radio access network.

Important: CSN is out of the scope of this document.

WiMAX Reference Points and Interfaces

A reference point (RP) in a WiMAX network is a conceptual link. An RP connects two groups of functions that reside in different functional entities of an ASN, CSN, or mobile station (MS). It is not necessarily a physical interface; an RP becomes a physical interface only when the functional entities on either side of it are contained in different physical devices.

Following are the reference points implemented with the ASN Gateway for WiMAX mobility functions:

- **R3 Reference Point**—Consists of the set of control plane protocols between the ASN and the CSN to support AAA, policy enforcement, and mobility management capabilities. It also encompasses the bearer plane methods (for example, tunneling) to transfer user data between the ASN and the CSN. R3 supports three types of clients: CMIPv4 (for MS/client capable of mobile IP), and PMIPv4 and PMIPv6 (proxy mobile IPv4 and IPv6, where ASNGW/FA proxy mobile IP supports MS/client).
- **R4 Reference Point**—Consists of the set of control and bearer plane protocols originating and terminating in various functional entities of an ASN that coordinate MS mobility between ASNs and ASN Gateways. R4 is the only interoperable RP between similar or heterogeneous ASNs.
- **R5 Reference Point**—Consists of the set of control plane and bearer plane protocols for internetworking between the CSN operated by the home NSP and that operated by a visited NSP.
- **R6 Reference Point**—Consists of the set of control and bearer plane protocols for communication between the base station and the ASN Gateway. The bearer plane is an intra-ASN datapath between the base station and ASN gateway. The control plane includes protocols for datapath establishment, modification, and release control, in accordance with the MS mobility events. R6, in combination with R4, may serve as a conduit for exchange of MAC state information between base stations that cannot interoperate over R8.
- **R7 Reference Point**—Consists of an optional set of control plane protocols, for example, AAA and policy coordination in the ASN gateway as well as other protocols for coordination between the two groups of functions identified in R6. The decomposition of the ASN functions using the R7 protocols is optional.

Important: To provide high throughput and high density call processing, the ASN Gateway integrates both the Decision Point and Enforcement Point functions. Therefore, the R7 reference point is not exposed.

Message Relay in ASN

The ASN Gateway provides relay procedures to send or distribute received messages with responses from a base station or another ASN Gateway. Supported types of relay functions are:

- **Passive Relay**: In this type of message relay, when the ASN Gateway receives a message on an R4 or R6 interface, it retrieves the destination ID and forwards the same request message to the given destination.
- **Active Relay**: In this type of message relay, upon receiving the message on R4/R6 interface, the ASN Gateway creates a similar R4/R6 message on the basis of original message and relays it to the destination. For example, if during the inter-ASN Gateway handover a non-anchor ASN Gateway receives the data path registration request from the target base station, it creates a new data path registration request and sends it to the anchor ASN Gateway. After receiving the duplicate message, the anchor ASN Gateway sends the data path

registration response to the non-anchor ASN Gateway. When it receives that message, the non-anchor ASN Gateway creates a new response message and sends the new data path registration response to the target base station.

ASN Gateway Architecture and Deployment Profiles

The ASN Gateway is part of the Access Service Network (ASN) within the WiMAX network. The ASN Gateway comprises logical and functional elements that provide different functionality in an ASN.

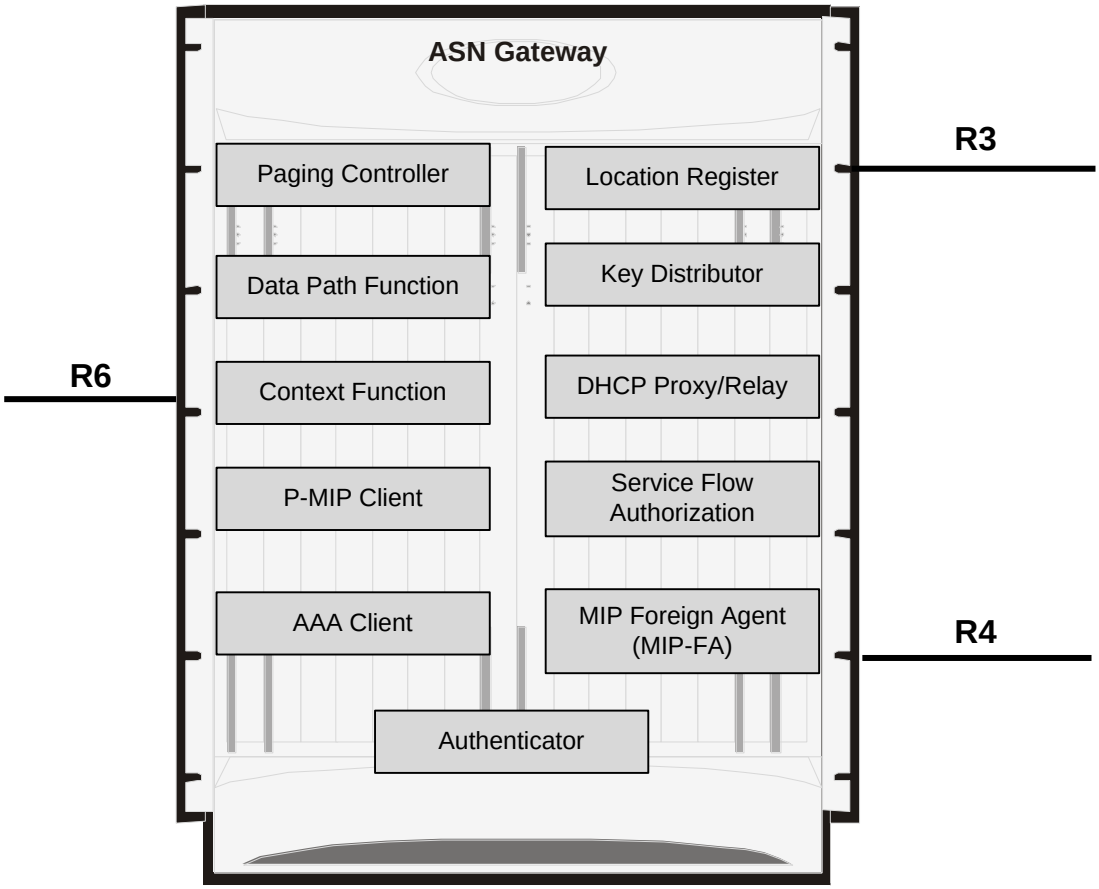
ASN profiles provide a framework for interoperability among entities within an ASN. At a high level, the WiMAX forum has defined groups of functionality for an ASN. These are called Profile Mappings A, B, and C. The key attributes of the profile mappings are:

- **ASN Profile-A**
  - Handover control and Radio Resource control (RRC) in the ASN Gateway
  - ASN anchored mobility among base stations using R6 and R4 reference points
  - CSN anchored mobility among ASNs using PMIP/CMIP (R3)
  - Paging Controller and Location Register in the ASN Gateway
- **Profile-B**: ASN Profile-B removes the ASN Gateway altogether and pushes all its functionality into the base station. This functionality includes the following:
  - Radio Resource control (RRC) handling within the base station
  - R3 reference point
  - R4 reference point
- **Profile-C**: ASN Profile-C functionality is a subset of Profile-A with following functionality in Base Station:
  - HO control
  - Radio Resource Controller (RRC)

The ASN Gateway supports ASN Profile-C functionality. Form more information on supported features and functionality, refer to the Supported Feature section.

The following figure shows the mapping of functional entities in an ASN Gateway for Profile-C.

Figure 35. Functional view of ASN Gateway Profile-C



WiMAX Network Deployment Configurations

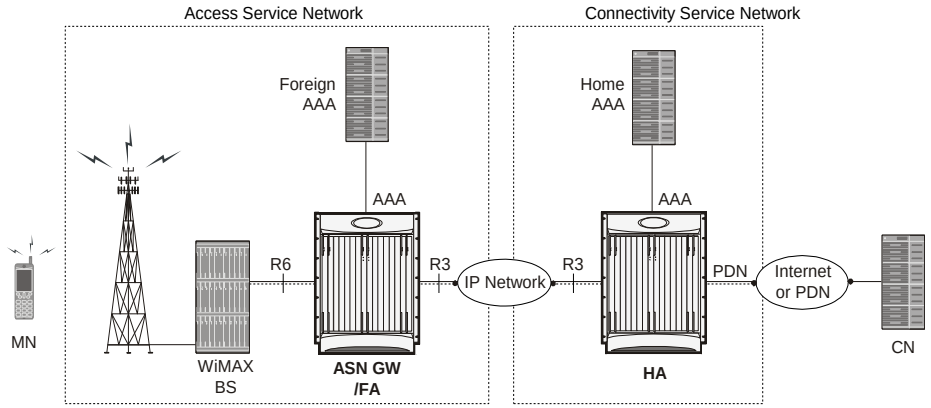
This section provides examples of how the system can be deployed within a WiMAX carrier's network. As noted previously, the system can be deployed in standalone configurations, serving as an Access Service Network Gateway/Foreign Agent (ASN Gateway/FA), a Home Agent (HA), or in a combined ASN Gateway/FA/HA configuration which provides all services from a single chassis.

Standalone ASN Gateway/FA and HA Deployments

The ASN Gateway/foreign agent (FA) serves as an integral part of a WiMAX network by providing packet processing and re-direction to a mobile user's home network through communications with the home agent (HA). No redirection is required when mobile users connect to an ASN Gateway that serves their home network.

The following figure shows an example of a network configuration in which the ASN Gateway/FA and HA are separate systems.

Figure 36. ASN Gateway/FA and HA Network Deployment Configuration Example

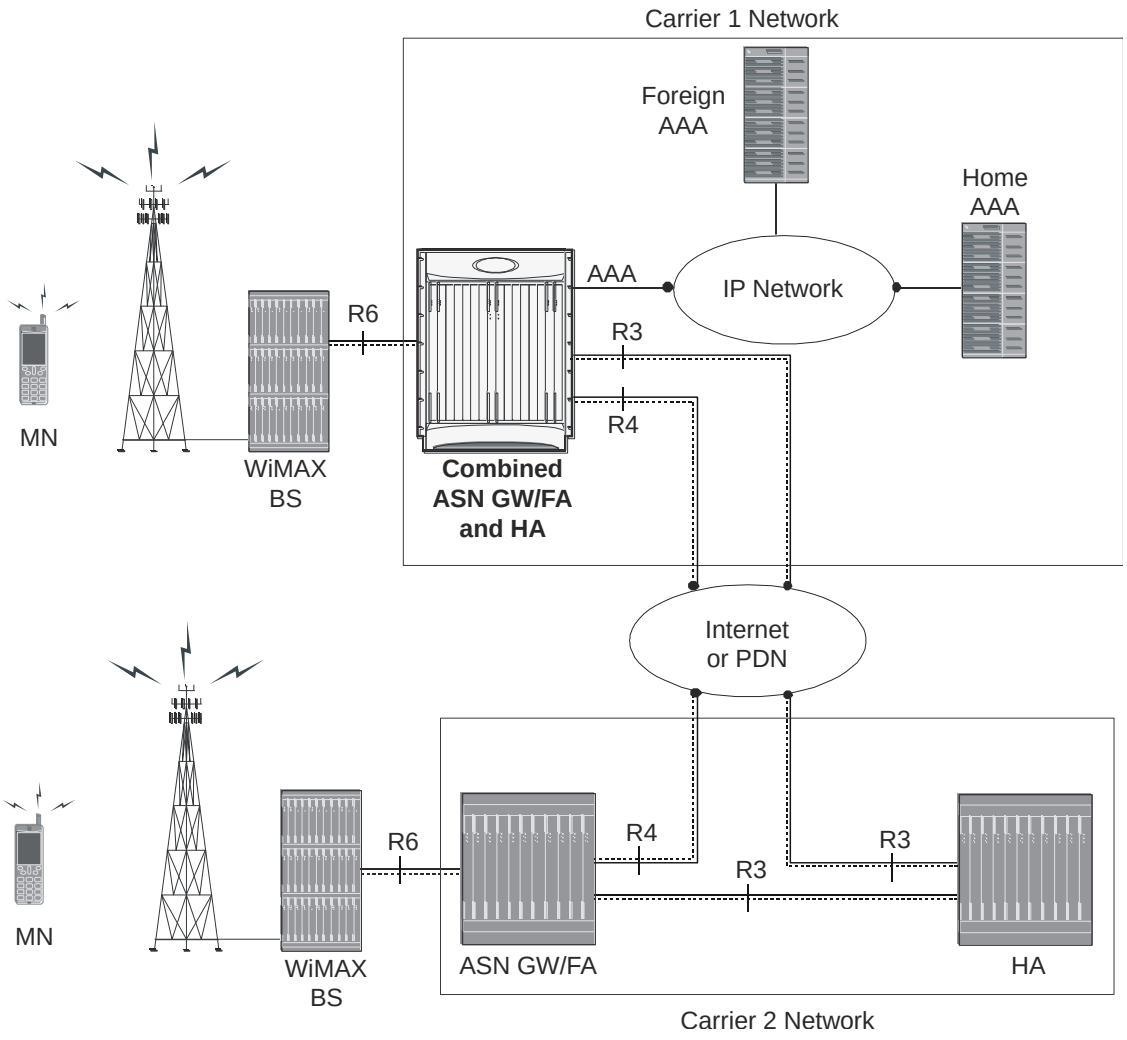


Co-Located Deployments

An advantage of the system is its ability to support both high-density ASN Gateway/FA and HA configurations within the same chassis. The economies of scale presented in this configuration example provide both improved session handling and reduced cost in deploying a WiMAX data network.

The following figure shows an example of a co-located deployment.

Figure 37. Co-located ASN Gateway/FA and HA Network Deployment Configuration Example



ASN Call Procedure Flows

This section provides information on the function of the ASN Gateway in a WiMAX network and presents call procedure flows for different stages of session setup.

Functional Components for Handover

This section describes the functional components used during handover between ASN Gateways on R4 and R6 interfaces.

Anchor ASN Gateway

The anchor ASN Gateway is the ASN Gateway that holds the anchor data path functions for a given MS. As shown in the following figure, the anchor ASN Gateway hosts the following functions:

- Authenticator (includes Accounting Client)
- Anchor DP function
- DHCP proxy
- DHCP Relay
- PMIP client
- MIP FA
- Anchor SFA
- DHCP proxy function

The ASN Gateway service IP address is the R6 and R4 tunnel endpoint and handles both R6 and R4 traffic.

Anchor Session

The following identifiers identify the anchor ASN Gateway session:

- MSID
- MS NAI
- MS IP address
- DHCP MAC address

The ASN Gateway session consists of an access R6 session and a MIP FA network session. The R6 session has a GRE data path to a base station for an active session. In this session the ASN Gateway service IP address is the R6 and R4 tunnel endpoint and handles both R6 and R4 traffic.

Upon initial network entry, when the DPF is in the anchor ASN Gateway, there is no R4 session. After a MS does a handover to a target BS, it connects to the anchor GW over R4 via a different serving ASN Gateway. At this point, the anchor GW session has an access R4 session and a MIP FA network session. The anchor GW can maintain the R6 session and a R4 session simultaneously.

Note that R6 and R4 tunnels are handled uniformly by the anchor GW as both are access-side tunnels. The anchor GW can check the IP address of the non-anchor GW peer against the configured list of peer ASN Gateway's, so that it can control which R4 connections are accepted.

The anchor GW handles all the Layer 3 processing for the subscriber without including any other rule and policy.

When an anchor GW receives a request message, it reads the source ID in this request and sends the response to this source ID as destination ID. The anchor ASN Gateway remembers the source IP address of the peer from where the message was received, if it is different from the source ID of the message. The response message is sent to this peer IP address, which is the immediate peer.

Non-Anchor ASN Gateway

The non-anchor ASN Gateway hosts the following functions:

- **Serving DP Function**: The subscriber data is not processed in the non-anchor GW. It relays the subscriber data to anchor ASN Gateway over R4. When the inner IP packet emerges from the R6 tunnel at the non-anchor ASN Gateway, the packet is sent over R4 data path tunnel to the Anchor ASN Gateway.
- **Serving SFA Function**: No packet classification is performed in this function. It provides only tunnel switching between R4 to R6 or vice versa.
- **DHCP Proxy relay Function**: DHCP messages are not processed in the non-anchor GW and relayed to the DHCP proxy in the anchor ASN Gateway over R4. When the inner IP packet emerges from the R6 tunnel at the non-anchor ASN Gateway, a check is made to see whether the DHCP proxy is co-located in the ASN Gateway and whether or not to process DHCP packet locally. If the session is not anchored locally, that is, the DHCP proxy is not co-located, the non-anchor ASN Gateway sends the DHCP packet over an R4 data path tunnel to the anchor ASN Gateway.
- **Relay Function**: The non-anchor ASN Gateway provides relay functions to distribute received messages and subscriber information. The message relay is supported for following functions:
  - Context transfer
  - Paging
  - Accounting
  - Authentication
  - Handover (HO)
  - Radio Resource Controller (RRC)

Non-Anchor Session

A non-anchor session is created upon receiving an R6 Data Path Registration Request from the target base station. Note that the non-anchor ASN Gateway session is identified by MSID only. This non-anchor ASN Gateway does NOT know the MS NAI and MS IP address of the subscriber, since the authenticator, DHCP and PMIP functions are not exposed here and the MSID is used as the username in session manager. The non-anchor session has the following attributes:

- The Registration Type in the request is set to HO.
- The Destination ID in the message does not match the destination IP address of the message. It needs to match the anchor ASN Gateway ID in the message if an R6 and R4 Data Path setup is intended.
- The anchor ASN Gateway is one of the peer ASN Gateway configured in the ASN Gateway service.

Initial Network Entry and Data Path Establishment without Authentication

This section describes the procedure of initial entry and data session establishment for a WiMAX subscriber station (SS) or MS without authentication by ASN Gateway.

Figure 38. Initial Network Entry and Data Session Establishment without Authentication Call Flow

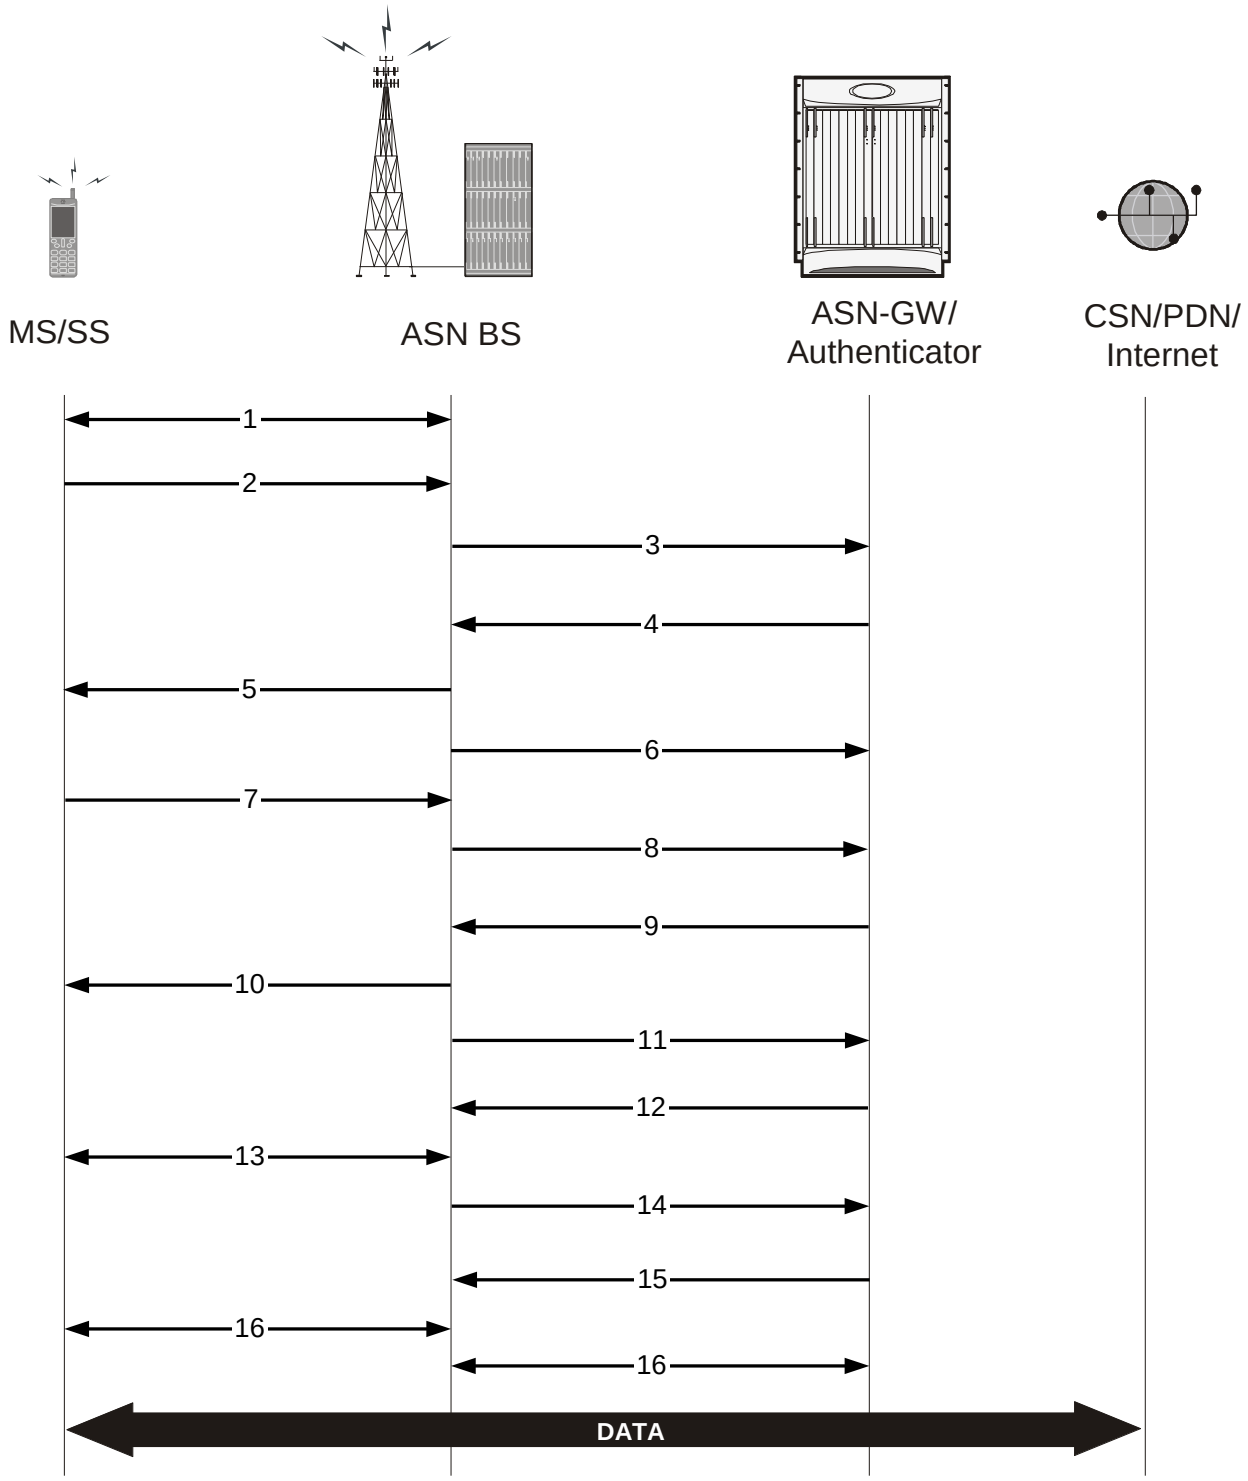


Table 28. Initial Network Entry and Data Session Establishment without Authentication Call Flow Description

Step	Description
1	MS performs initial ranging with the ASN BS. Ranging is a process by which an MS becomes time-aligned with the ASN BS. The MS is synchronized with the BS at the successful completion of ranging and is ready to set up a connection.
2	MS sends basic capability exchange request (SBC-REQ) to ASN BS.
3	ASN BS sends MS-Pre-Attachment Request (authorization policy request) to ASN Gateway.
4	ASN Gateway sends MS-Pre-Attachment Response on the basis of authorization policy to ASN BS for MS.
5	ASN BS sends basic capability exchange response (SBC-RSP) to MS.
6	If authorization policy allows, ASN BS sends MS Pre-Attachment Acknowledgement to ASN Gateway.
7	MS sends Registration-Request (REG-REQ) to ASN BS.
8	ASN BS sends MS-Attachment-Request to ASN Gateway.
9	ASN Gateway sends MS-Attachment-Response to ASN BS and reserves the resource.
10	ASN BS sends Registration-Response to MS.
11	ASN BS sends MS-Attachment-Acknowledgement to ASN Gateway.
12	ASN Gateway sends Path Registration Request to ASN BS.
13	ASN BS creates 802.16 connection and establishes path with MS.
14	ASN BS sends Path Registration Response to ASN Gateway and ASN Gateway creates service flow with CSN over which PDUs can be sent and received.
15	ASN Gateway sends Path Registration Acknowledgment to ASN BS.
16	GRE tunnel mapped to 802.16 connection between MS and ASN BS.
17	R6 GRE data path established between ASN BS and ASN Gateway and data flow starts.

Initial Network Entry and Data Path Establishment with Authentication (Single EAP)

This section describes the procedure of initial entry and data session establishment for a WiMAX Subscriber Station (SS) or MS with single EAP authentication.

The following figure provides a high-level view of the steps involved for initial network entry of an SS/MS with EAP authentication and data link establishment. The following table explains each step in detail.

Figure 39. Initial Network Entry and Data Session Establishment with Authentication Call Flow

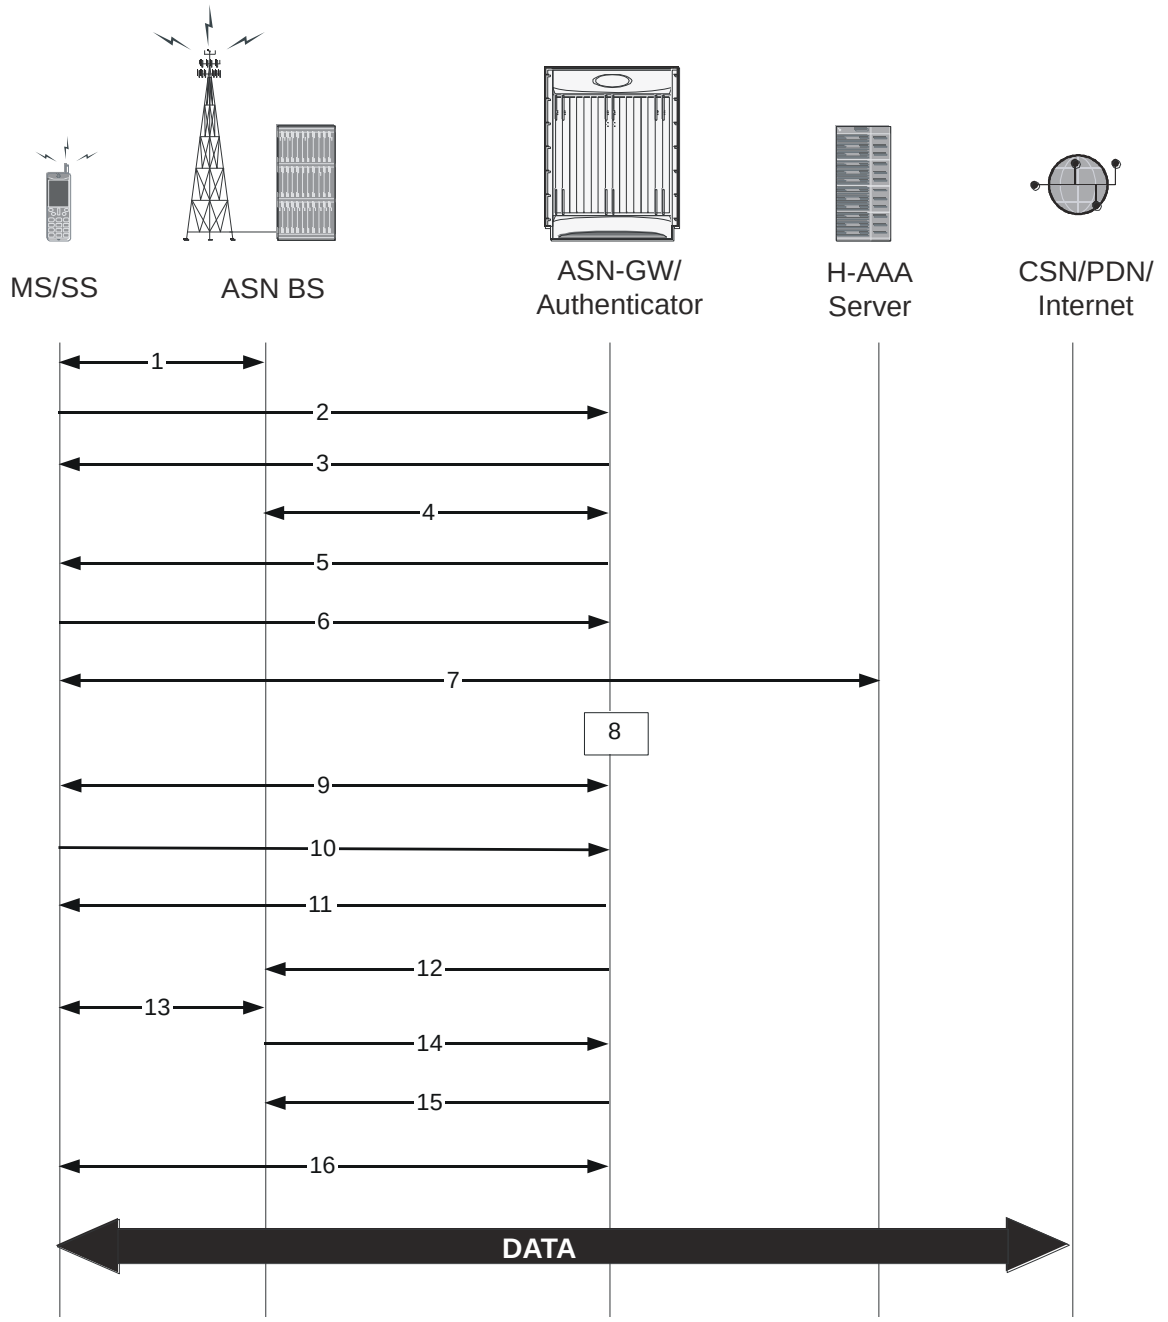


Table 29. Initial Network Entry and Data Session Establishment with Authentication Call Flow Description

Step	Description
1	MS performs initial ranging with the BS. Ranging is a process by which an MS becomes time aligned with the BS. The MS is synchronized with the BS at the successful completion of ranging and is ready to set up a connection.

Step	Description
2	SS Basic capability exchange (SBC-REQ) between MS and BS starts and MS-Info-Request for authorization policy sent to AAA client/authenticator in ASN Gateway.
3	AAA client/authenticator (ASN Gateway) sends MS-Info-Report to BS and BS sends SS Basic Capability Response (SBC-RSP) to MS.
4	BS acknowledges the MS-Info-Report to AAA client/authenticator.
5	AAA client/authenticator (ASN Gateway) starts EAP transfer request to BS and MS.
6	MS and BS sends EAP transfer response to AAA client/authenticator.
7	The MS progresses to an authentication phase with home AAA Server. Authentication is based on PKMv2 as defined in the IEEE standard 802.16 specification. EAP authentication process starts
8	EAP authentication successful and AAA client/authenticator starts security context transfer.
9	PKMv.2-RSP/EAP-Transfer/SA-TEK-Challenge-Request-Response/Key-Request-Response exchange between MS and BS.
10	MS sends 802.16 Registration Request (REG-REQ) to ASN BS and ASN BS sends MS-Info-Request to AAA client/authenticator.
11	AAA client/authenticator sends MS-Info-Report to BS and BS sends Registration Response (REG-RESP) to MS and MS-Info-Report Acknowledge to AAA client/authenticator.
12	ASN Gateway sends Path Registration Request to ASN BS.
13	ASN BS creates 802.16e connection and establishes path with MS.
14	ASN BS sends Path Registration Response to ASN Gateway and ASN Gateway creates service flow with CSN over which PDUs can be sent and received.
15	ASN Gateway sends Path Registration Acknowledgment to ASN BS.
16	GRE tunnel mapped to 802.16 connection between MS and ASN BS.
17	R6 GRE data path established between ASN BS and ASN Gateway and data flow starts.

Unexpected Network Re-entry

An unexpected network re-entry is when a mobile station starts the process of initial network entry to the ASN Gateway via the same or new base station while an existing call for the MS is still in progress or being set up. When this occurs, the ASN Gateway's default behavior is to:

- Accept the new call regardless of the existing call state if the pre-attachment request of the new call comes from a different BS.
- Accept the new call if the original call is in any state past the pre-attachment phase and the pre-attachment request of the new call comes from the same BS.
- Drop the original call in favor of new call.

To disable this default behavior use the **`policy ms-unexpected-network-reentry`** command in the ASN Gateway Service Configuration Mode. For more information regarding this command, refer to the *Command Line Interface Reference*.

MS Triggered Network Exit

This section describes the procedure of MS Triggered network exit for a WiMAX Subscriber Station (SS) or MS in normal mode.

The following figure provides a high-level view of the steps involved for network exit of an SS/MS in normal mode. The following table explains each step in detail.

Figure 40. MS Triggered Network Exit Call Flow

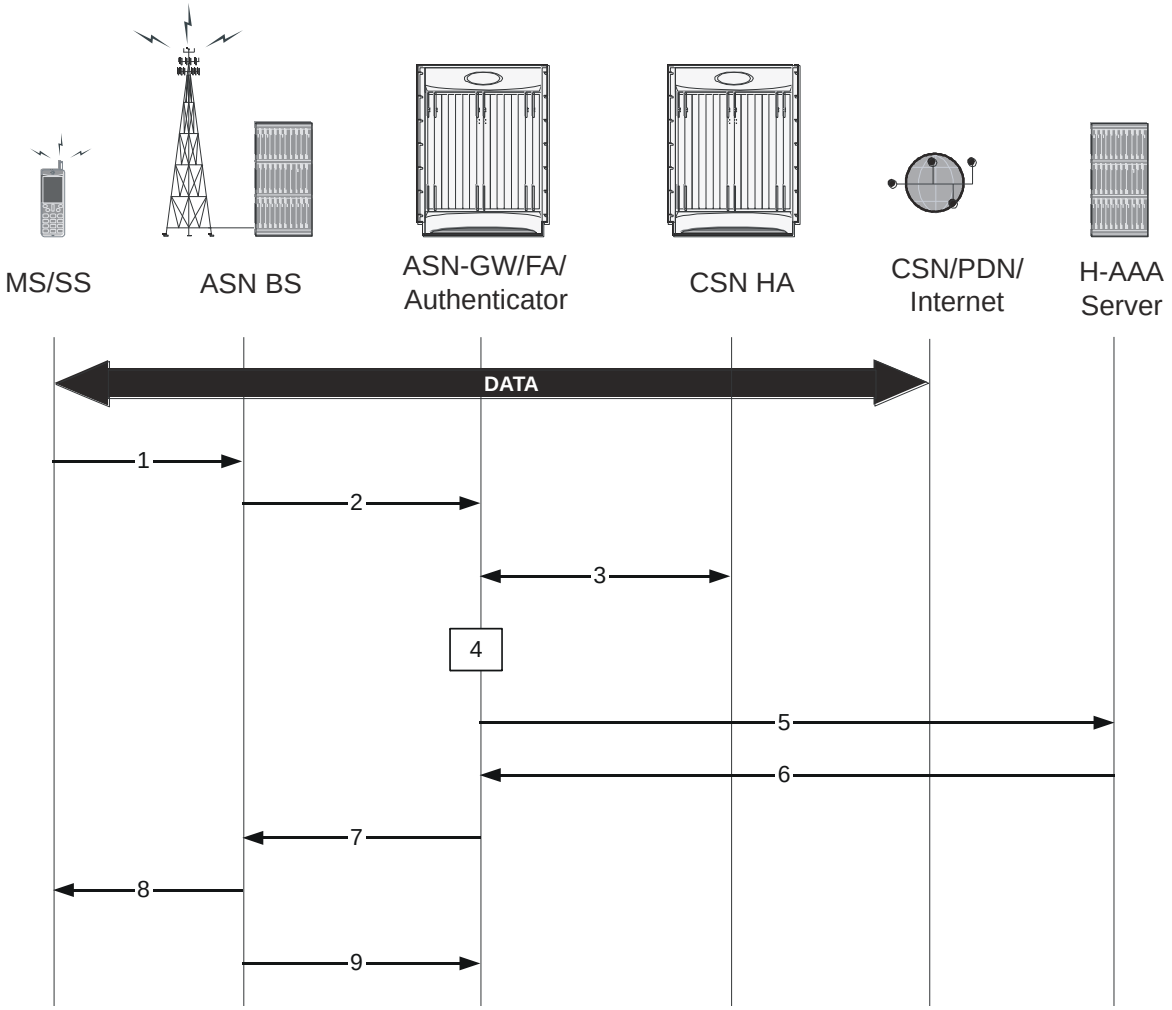


Table 30. MS Triggered Network Exit Call Flow Description

Step	Description
1	MS sends DREG_REQ message to ASN BS in serving ASN, including De-Registration_Request Code=0x00.
2	ASN BS sends R6 Path_Dereg_Req message to ASN Gateway.
3	ASN Gateway/FA and HA starts MIP release procedure.

Step	Description
4	ASN Gateway/FA starts MS context delete procedure.
5	ASN Gateway sends Accounting-Stop-Request (Release Indication) message to AAA.
6	AAA replies with Accounting-Stop-Response message to ASN Gateway.
7	ASN Gateway/FA replies with Path_Dereg_Response message to ASN BS.
8	ASN BS sends DREG_CMD message to MS, including Action Code=0x04.
9	ASN BS sends R6 Path_Dereg_Ack to the ASN Gateway and related entities releases the retained MS context and the assigned data path resource for the MS.

Network Triggered Network Exit

This section describes the procedure of a network triggered network exit for a WiMAX Subscriber Station (SS) or MS in normal mode.

The following figure provides a high-level view of the steps involved for a network-triggered network exit of an SS/MS in normal mode. The following table explains each step in detail.

Figure 41. Network Triggered Network Exit Call Flow

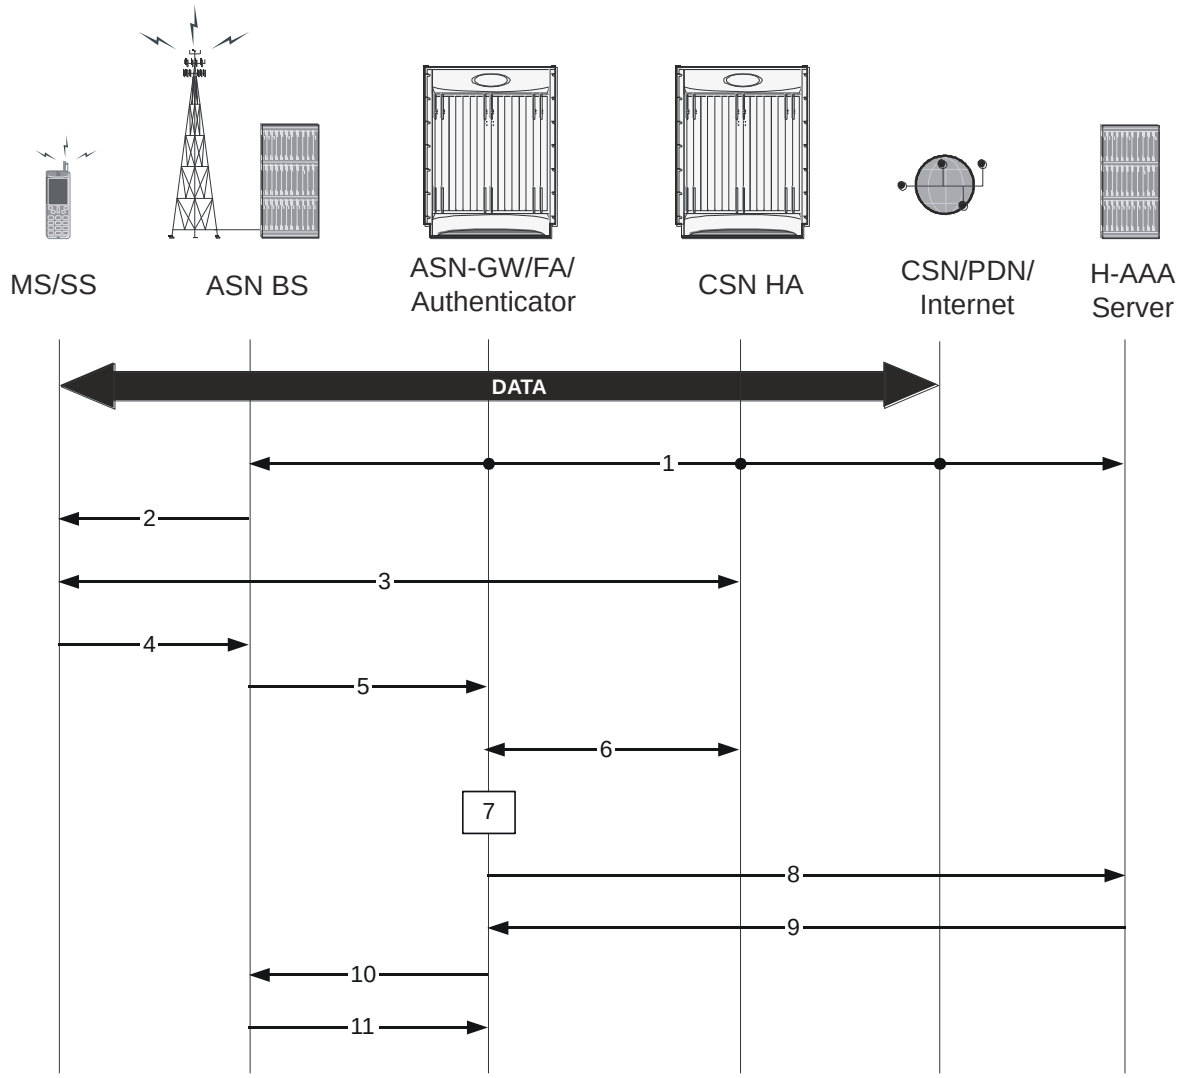


Table 31. Network Triggered Network Exit Call Flow Description

Step	Description
1	Network entities, such as AAA Server, ASN Gateway FA/HA, trigger Session Release Trigger to ASN BS. This can be from H-AAA ServerAnchor ASN Gateway/FA/HAServing ASN BS, etc.
2	ASN BS sends DREG_CMD message to MS, including Action Code=0x00 to indicate MS existing network.
3	IP session for DHCP/MIP release starts between MS and network entities.
4	MS sends DREG_REQ to ASN BS with De-Registration_Request_Code=0x02.
5	ASN BS sends Path_Dereg_Req message to ASN Gateway.
6	Auth Context Request and Report Exchange between target BS and ASNGW.

Step	Description
7	ASN Gateway/FA and HA starts MIP release procedure.
8	ASN Gateway/FA exchanges NetExit_MS_State_Change_Req and NetExit_MS_State_Change_Rsp messages with the anchor accounting client, anchor authenticator, and MIP client to delete MS contexts.
9	ASN Gateway sends Accounting-Stop-Request (Release Indication) message to H-AAA.
10	AAA replies with Accounting-Stop-Response message to ASN Gateway.
11	ASN Gateway/FA replies with Path_Dereg_Response message to ASN BS.
12	HO Complete message sent from target BS to serving BA via ASNGW.
13	ASN BS sends R6 Path_Dereg_Ack to the ASN Gateway and related entities releases the retained MS context and the assigned data path resource for the MS.

Intra-ASN Gateway Handover

This section describes the handover procedure between two ASN BSs connected to one ASN Gateway. The ASN Gateway supports following types of handover:

- Intra-anchor ASN Gateway Uncontrolled Handover
- Intra Non-anchor ASN Gateway Uncontrolled Handover
- Intra-anchor ASN Gateway Controlled Handover
- Intra Non-anchor ASN Gateway Controlled Handover

Details regarding controlled and uncontrolled handovers for the anchor ASN gateways are provided below.

Intra-anchor ASN Gateway Uncontrolled Handover

This section describes the procedure for an uncontrolled intra-anchor ASN Gateway handover for a WiMAX Subscriber MS.

The following figure provides a high-level view of the steps involved in an intra-anchor ASN Gateway uncontrolled handover of an SS/MS. The following table explains each step in detail.

Figure 42. Intra-ASN Gateway Uncontrolled Handover Call Flow

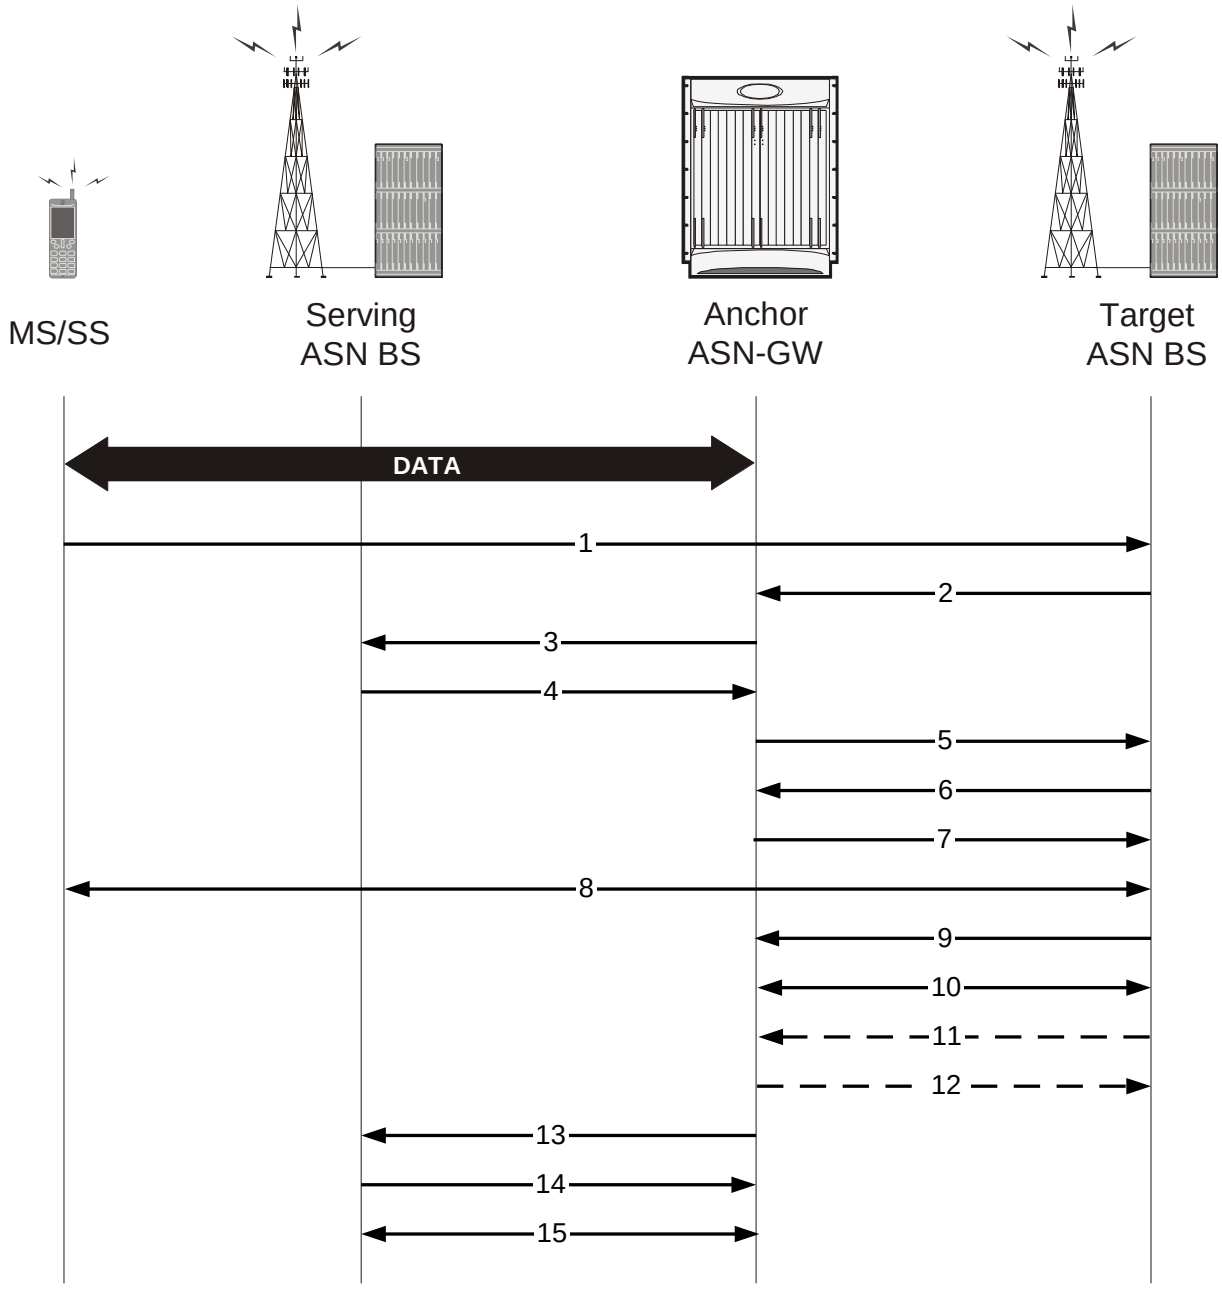


Table 32. Intra-ASN Gateway Uncontrolled Handover Call Flow Description

Step	Description
1	MS sends RNG-REQ message to target ASN BS.
2	Target ASN BS sends Context-Request message to anchor ASN Gateway for this MS.
3	Anchor ASN Gateway forwards Context-Request message to serving ASN BS.

Step	Description
4	Serving ASN BS sends Context-Report message with MS context information to anchor ASN Gateway.
5	Anchor ASN Gateway forwards Context-Report message with MS context information to target ASN BS.
6	Target ASN BS sends Path Registration Request to anchor ASN Gateway.
7	Anchor ASN Gateway replies with Path Registration Response to target ANS BS.
8	Target ANS BS sends ranging response with RNG_RSP message to MS.
9	Target ASN BS sends Path Registration Acknowledge to anchor ASN Gateway.
10	R6 GRE data path established between target ASN BS and anchor ASN Gateway and data flow starts.
11	Target ASN BS sends CMAC Key Count Update message to anchor ASN Gateway.
12	Anchor ASN Gateway replies with CMAC Key Count Update ACK message to target ASN BS.
13	Anchor ASN Gateway sends Path_De-Reg_Req message to release data path to serving BS.
14	Serving ASN BS sends Path_De-Reg_Rsp message to anchor ASN Gateway.
15	R6 GRE data path terminated between serving ASN BS and anchor ASN Gateway.

Intra-anchor ASN Gateway Controlled Handover

An intra-anchor ASN Gateway controlled handover consists of the following types and phases.

MS Initiated Intra-anchor ASN Gateway Controlled Handover

This section describes the intra-anchor ASN Gateway controlled handover between two base stations initiated by a mobile station.

HO Preparation Phase

This is the initial phase for a controlled handover between two BSs.

The following figure and table describe the call flow for the steps involved in a controlled intra-ASN Gateway handover preparation phase between two BSs.

Figure 43. MS initiated Controlled Intra-ASN Gateway Handover Preparation Phase

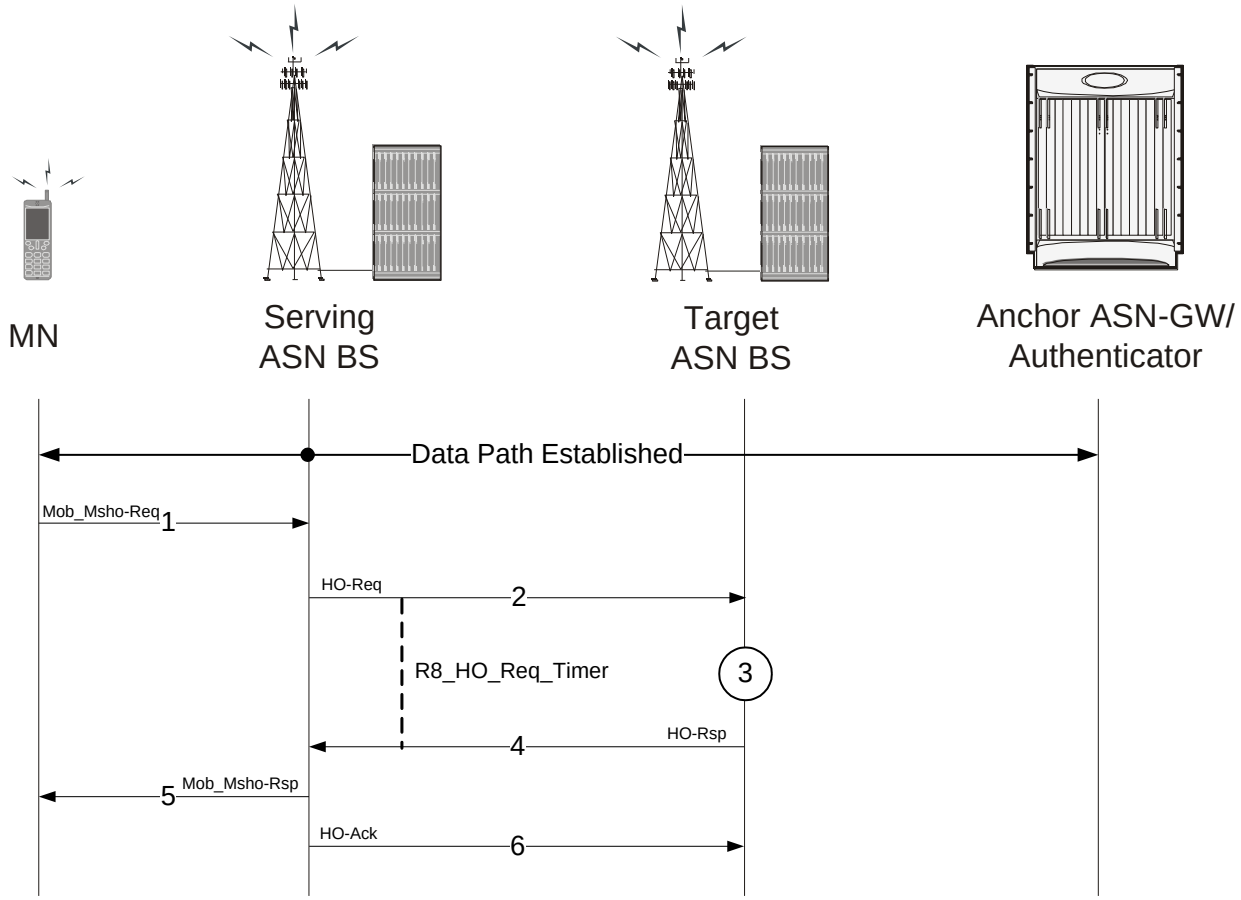


Table 33. MS initiated Controlled Intra-ASN Gateway Handover Preparation Phase Description

Step	Description
1	MS sends MOB_MSHO_REQ messages to serving BS
2	Upon receiving MS initiated handover request (MOB_MSHO_REQ), the serving BS sends HO_Req messages to target BS selected by MS and starts R8_HO_Req timer
3	Targeted BS tests the acceptability of the requested HO by comparing the amount of available resources and required bandwidth/QoS parameters in the HO request received from serving BS
4	Once a target BS accepts the request it sends the HO_Rsp message to the serving BS
5	Serving BS sends MOB_MSHO_RSP response t o MS
6	Serving BS sends HO_Ack message to the target BS and HO preparation phase is completed

HO Action Phase

The following figure and table describe the call flow for the steps involved in uncontrolled intra-ASN Gateway handover action phase between two BSs.

Figure 44. MS initiated Controlled Intra-ASN Gateway Handover Action Phase

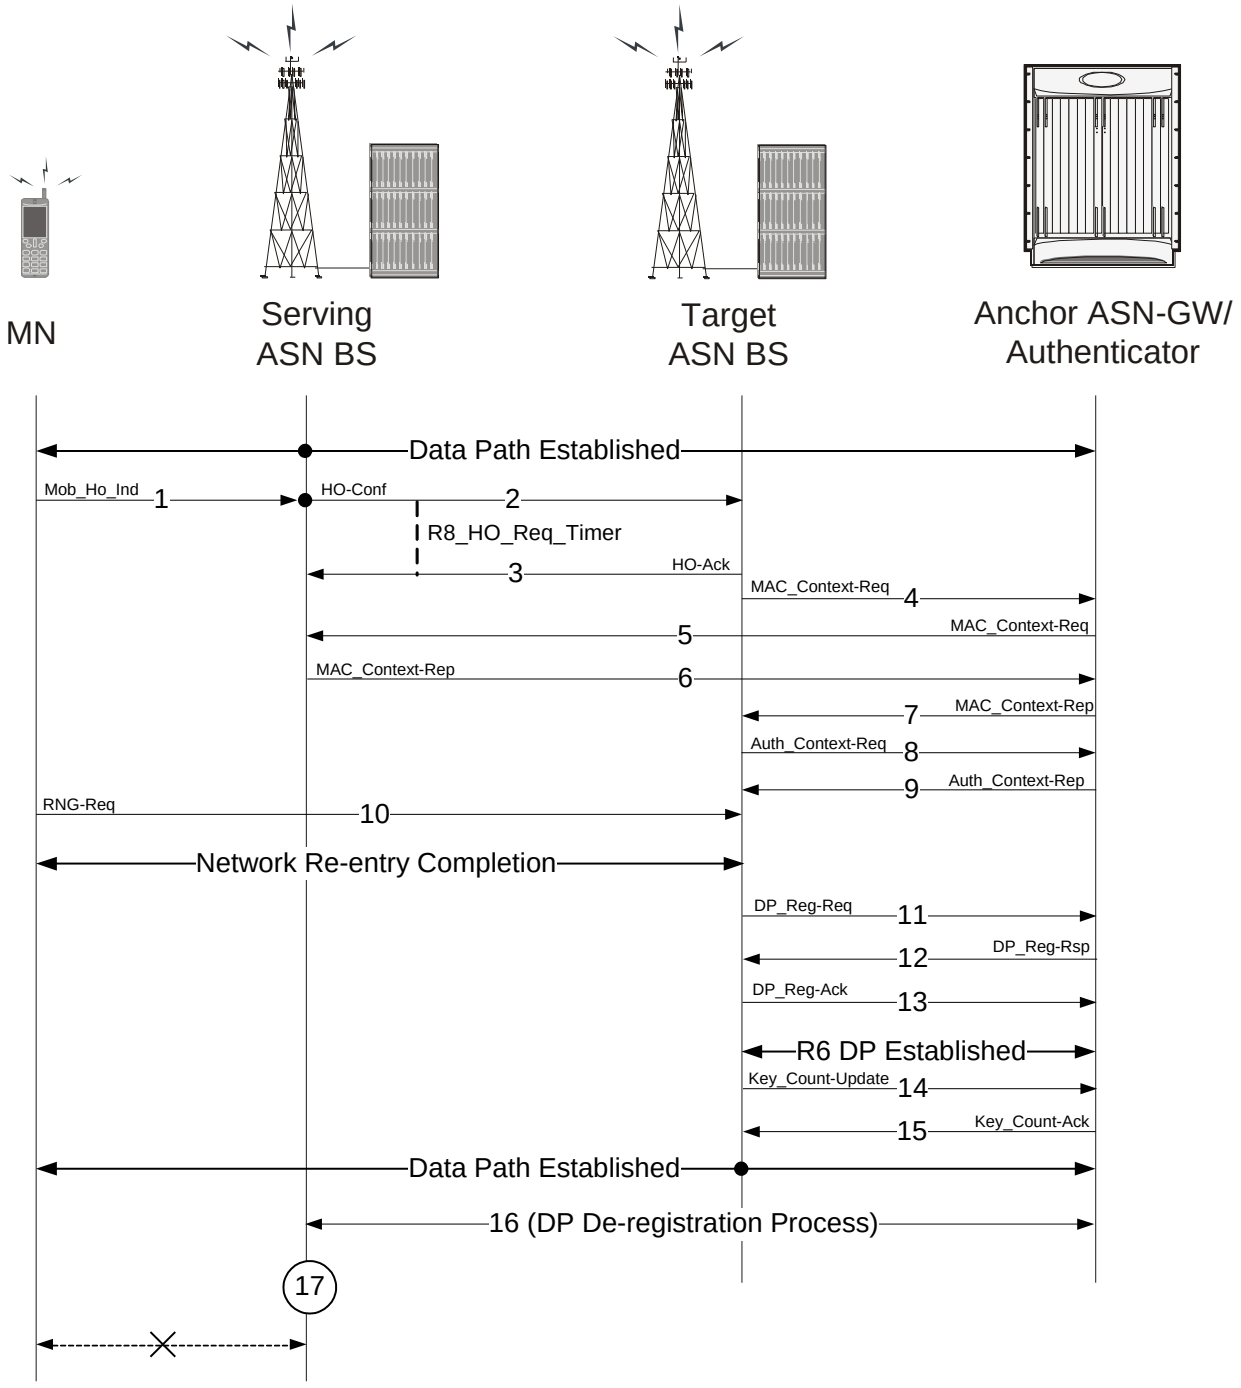


Table 34. MS initiated Controlled Intra-ASN GW Handover Phase

Step	Description

Step	Description
1	Once HO preparation phase is completed and target BS receives HO-Ack message, the MS sends MOB_HO-IND messages to the serving BS.
2	The serving BS sends HO_Conf messages to the selected target BS with other context information and starts R8_HO_Confirm Timer.
3	The target BS accepts the request and sends the HO_Ack message to serving BS and serving BS stops R8_HO_Confirm Timer.
4	Target BS sends MAC Context Request message to the anchor ASN Gateway.
5	The anchor ASN Gateway forwards the MAC Context Request to the serving BS.
6	Serving BS sends MAC Context Report information to anchor ASN Gateway.
7	Anchor ASN Gateway forwards MAC Context Report information to the target BS.
8	Target BS sends Authentication Context Request to anchor ASN Gateway.
9	Anchor ASN Gateway transfers Authentication Context information to target BS.
10	MS starts ranging with target BS and sends RNG-REQ to the target BS and network reentry completed.
11	Target BS sends Data Path Registration Request to anchor ASN Gateway.
12	Anchor ASN Gateway sends Data Path Registration Response to target BS.
13	Target BS sends Data Path Registration Ack message to Anchor ASN Gateway and R6 data path is established.
14	Target BS sends CMAC Key count Update message to anchor ASN Gateway.
15	Anchor ASN Gateway sends CMAC Key Count Update Ack message to target BS and handover completed.
16	Anchor AS NGW starts Data Path De-registration process with serving BS.
17	Serving BS releases all resources and terminates data path with MS.

BS Initiated Intra Anchor ASN Gateway Controlled Handover

This section describes the intra-anchor ASN Gateway controlled handover between two base stations initiated by serving base station.

HO Preparation Phase

This is the initial phase for a controlled handover between two BSs.

The following figure and table describe the call flow for the steps involved in uncontrolled intra-ASN Gateway handover preparation phase between two BSs.

Figure 45. BS initiated Controlled Intra-ASN Gateway Handover Preparation Phase

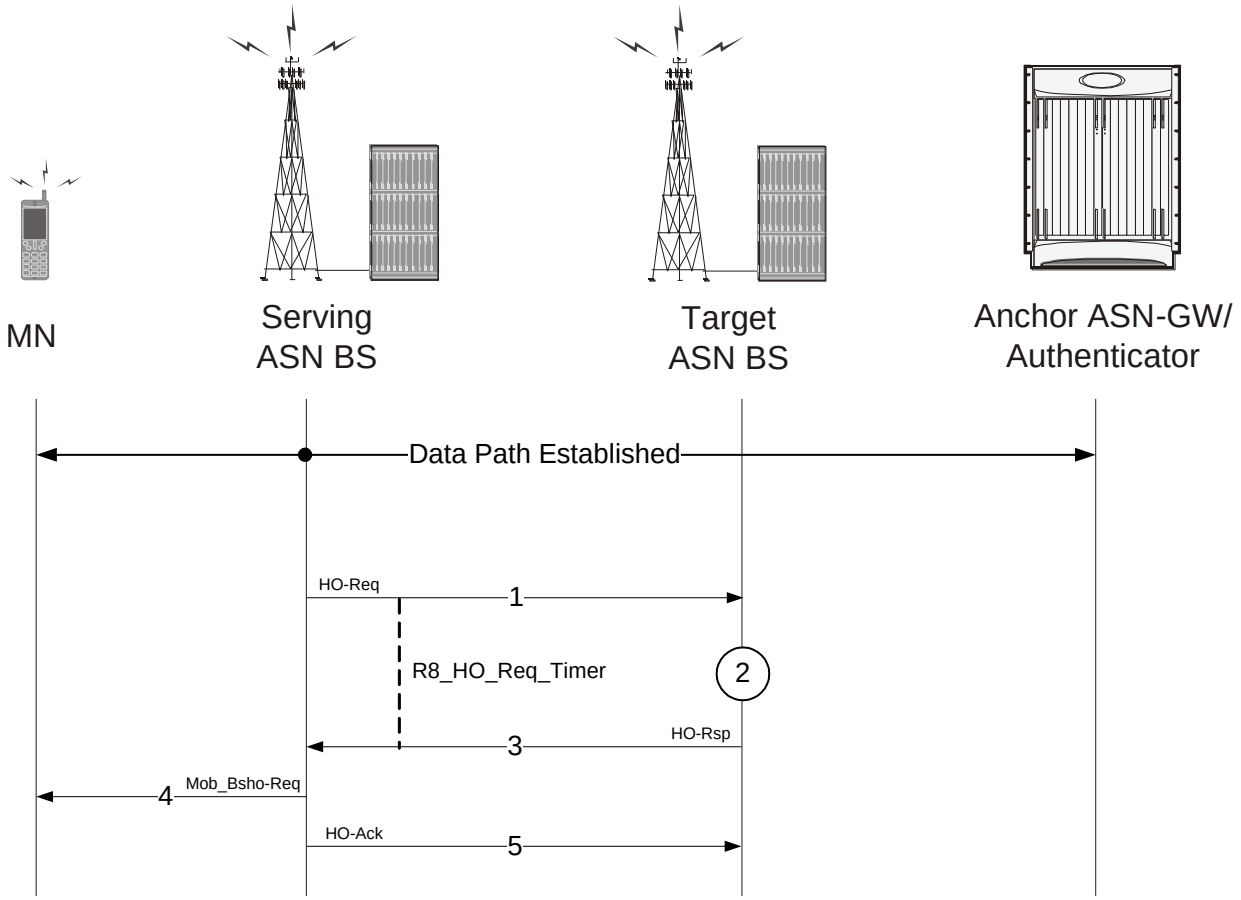


Table 35. BS initiated Controlled Intra-ASN Gateway Handover Preparation Phase Description

Step	Description
1	In BS initiated HO scenario, the serving BS sends HO_Req messages to target BS from its peer list and starts R8_HO_Req timer.
2	Targeted BS tests the acceptability of the requested HO by comparing the amount of available resources and required bandwidth/QoS parameters in the HO request received from serving BS.
3	Once a target BS accepts the request it sends the HO_Rsp message to the serving BS.
4	Serving BS sends MOB_MSHO_RSP response t o MS.
5	Serving BS sends HO_Ack message to the target BS and HO preparation phase is completed.

HO Action Phase

The following figure and table describe the call flow for the steps involved in an uncontrolled intra-ASN Gateway handover action phase between two BSs.

Figure 46. BS initiated Controlled Intra-ASN Gateway Handover Action Phase

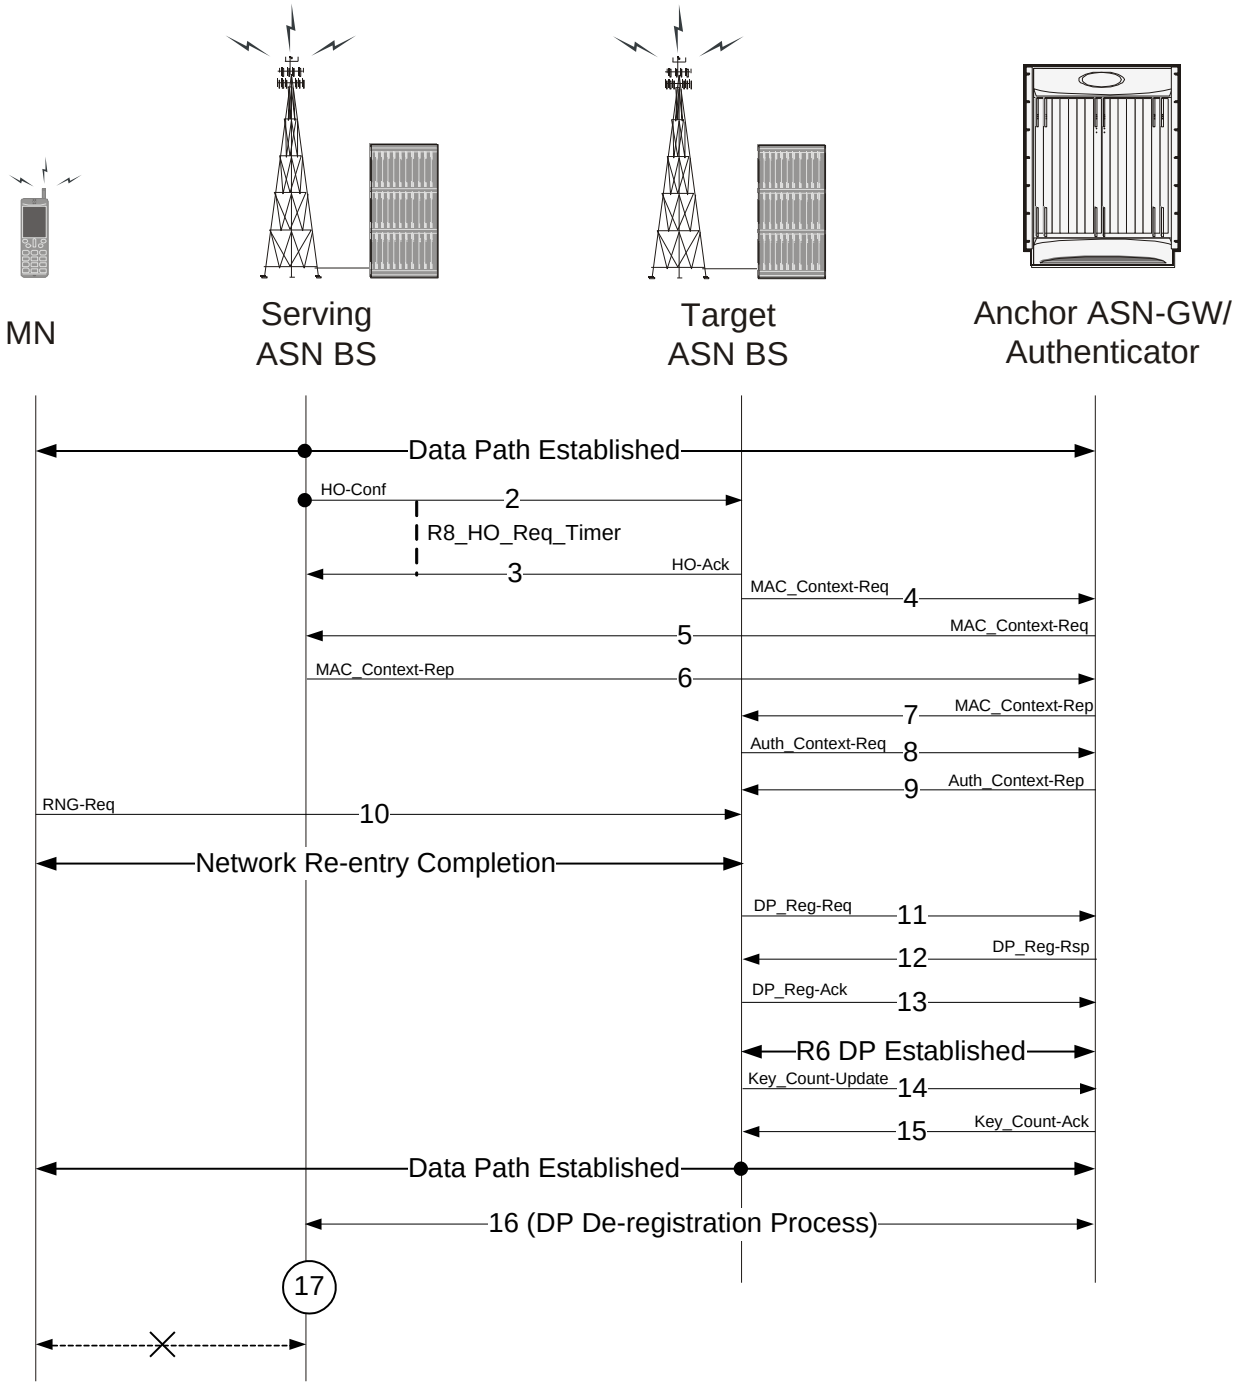


Table 36. BS initiated Controlled Intra-ASN Gateway Handover Action Phase Description

Step	Description
1	Handover preparation phase is completed and data path is established.

Step	Description
2	The serving BS sends HO_Conf messages to the selected target BS with other context information and starts R8_HO_Confirm Timer.
3	The target BS accepts the request and sends the HO_Ack message to serving BS and serving BS stops R8_HO_Confirm Timer.
4	Target BS sends MAC Context Request message to the anchor ASN Gateway.
5	The Anchor ASN Gateway forwards the MAC Context Request to the serving BS.
6	Serving BS sends MAC Context Report information to anchor ASN Gateway.
7	Anchor ASN Gateway forwards MAC Context Report information to the target BS.
8	Target BS sends Authentication Context Request to anchor ASN Gateway.
9	Anchor ASN Gateway transfers Authentication Context information to target BS.
10	MS starts ranging with target BS and sends RNG-REQ to the target BS and network reentry completed.
11	Target BS sends Data Path Registration Request to anchor ASN Gateway.
12	Anchor ASN Gateway sends Data Path Registration Response to target BS.
13	Target BS sends Data Path Registration Ack message to anchor ASN Gateway and R6 data path established.
14	Target BS sends CMAC Key count Update message to anchor ASN Gateway.
15	Anchor ASN Gateway sends CMAC Key Count Update Ack message to target BS and handover completed.
16	Anchor AS NGW starts Data Path De-registration process with serving BS.
17	Serving BS releases all resources and terminates data path with MS.

Inter-ASN Gateway Handover

This section describes the procedure of inter-ASN Gateway handovers through an R4 interface for a WiMAX Subscriber Station (SS). The R4 reference is the interface over which ASN control and data messages are exchanged between two ASN Gateways, either within the same ASN or across separate ASNs.

For a given subscriber, a WiMAX session may be handled by ASN Gateway functions located in different physical nodes in the network. For example, the authenticator and FA may be located in ASN Gateway *x* and the R6 Data Path Function in ASN Gateway *y*. The various ASN Gateway functions communicate over the R4 interface.

The following inter-ASN Gateway handover scenarios are supported on the ASN Gateway over the R4 interface:

Important: Not all features are supported on all platforms.

- Controlled Anchor ASN Gateway to Non-Anchor ASN Gateway Handover
- Controlled Non-Anchor ASN Gateway to Anchor ASN Gateway Handover
- Controlled Non-Anchor ASN Gateway to Non-Anchor ASN Gateway Handover
- Uncontrolled Anchor ASN Gateway to Non-Anchor ASN Gateway Handover
- Uncontrolled Non-Anchor ASN Gateway to Anchor ASN Gateway Handover

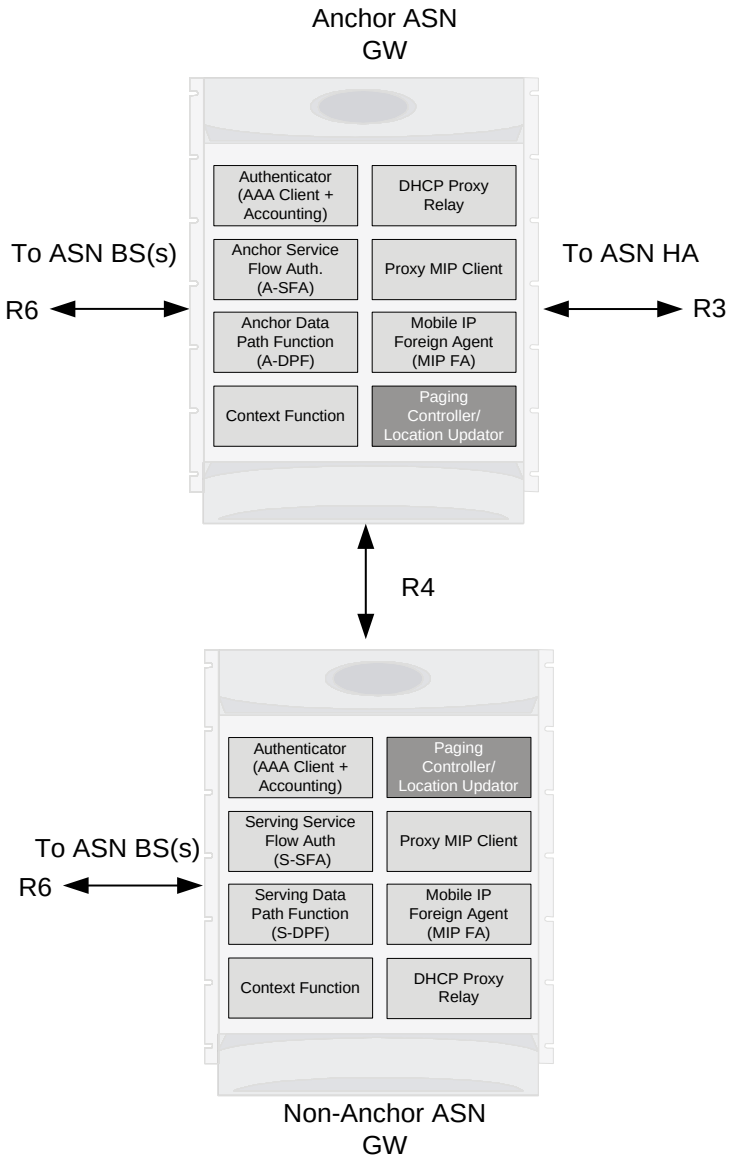
- Uncontrolled Non-Anchor ASN Gateway to Non-Anchor ASN Gateway Handover

ASN Gateway Function for Handovers

An ASN Gateway configured for inter-ASN Gateway handovers requires the following functionality to support the handover via an R4 interface.

The following figure provides a high-level view of the components and functions distribution in ASN Gateway.

Figure 47. Distribution of Components and Function in ASN Gateway for Handover



Controlled Anchor ASN Gateway to Non-Anchor ASN Gateway Handover

For Controlled handovers, the ASN Gateway provides and/or supports the following functions:

- **Message Relay**: The ASN Gateway provides the passive relay function for HO Request, HO Response, HO Ack, HO Confirm, and HO Complete messages in a stateless fashion. The gateway keeps the statistics of the different types of messages it has relayed. Retransmission of these messages is handled by the BS.

  The serving BS generates these messages. The serving BS generates a different HO Request transaction for each target BS. In other words, the gateway does not generate multiple HO Request messages after receiving a single HO Request message with multiple target BSs. Generally, the HO transaction is initiated by the serving BS which also chooses the selected target BS to which the handover will take place.
- **Security Context Retrieval**: The ASN Gateway supports the retrieval of the security context using Context Request and Context Report messages. This retrieval is also stateless. The context retrieval operation can be performed at any time during the lifetime of a call.
- **Data Path Registration**: After Pre-Registration, the target BS performs Data Path Registration. Data Path Registration is performed using a 3-way handshake. If Pre-Registration has occurred, the Data Path Registration messages do not contain any service flow information.
  - If Pre-Registration has not occurred, the Data Path Registration messages carry the service flow information.
  - Data Path Pre-Registration and Data Path Registration is initiated by the BS.

Preparation Phase

The following figure and table provides a high-level view of the steps involved during the preparation phase of a controlled inter-ASN Gateway handover of an SS/MS from an anchored gateway to a non-anchored gateway.

Figure 48. Controlled Inter-ASN Gateway Handover Procedure - Preparation Phase

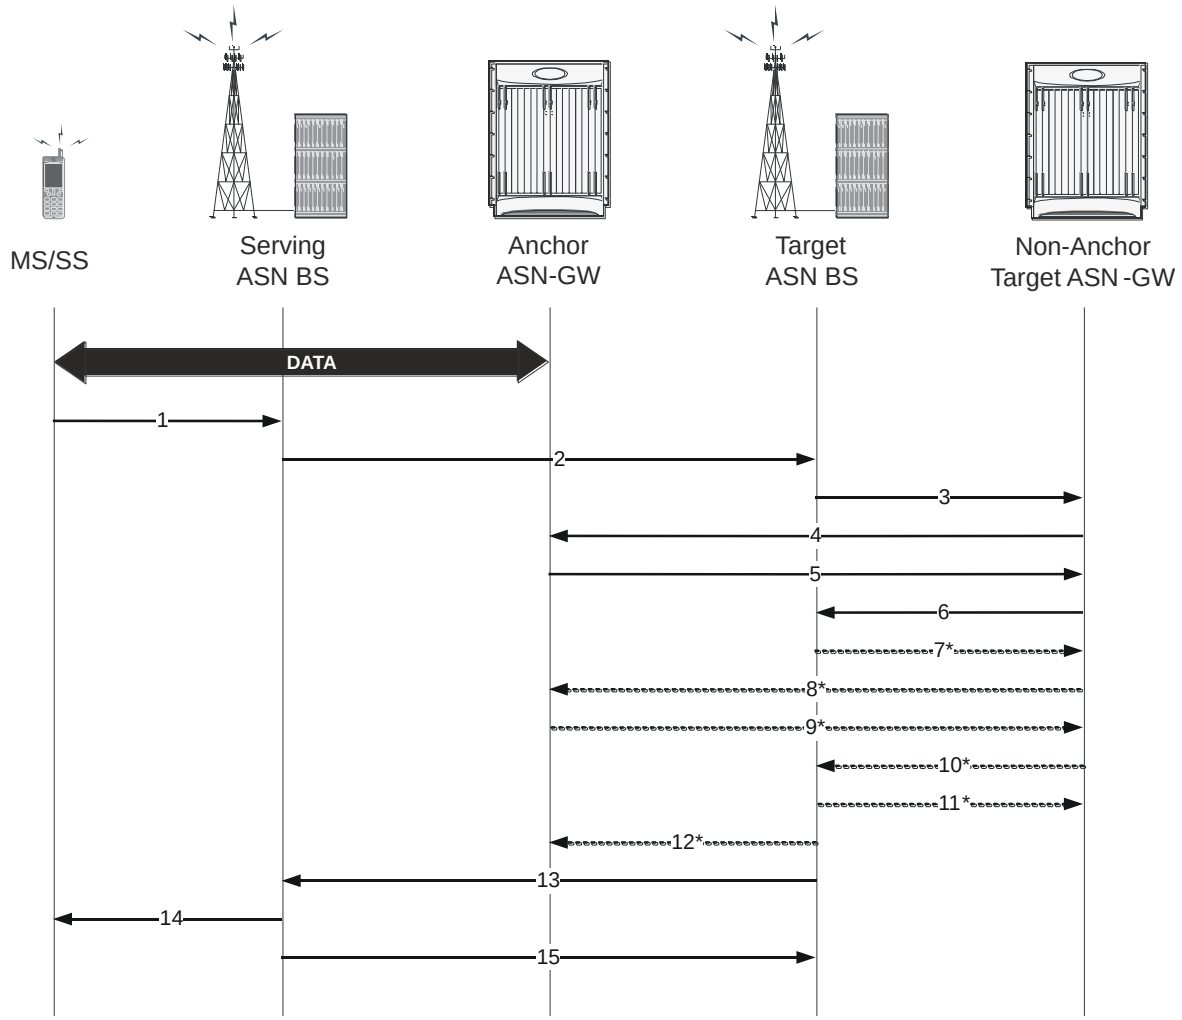


Table 37. Controlled Inter-ASN Gateway Handover Procedure - Preparation Phase Description

Step	Description
1	MS sends a MOB_MSHO-REQ message to the serving ASN BS.
2	Serving ASN BS sends a Handover Request message to the target ASN BS.
3	Target ASN BS sends a Context-Request message to the target non-anchor ASN Gateway for this MS.
4	Target non-anchor ASN Gateway forwards the Context-Request message to the anchor ASN Gateway.
5	Anchor ASN Gateway sends a Context-Report message to the target non-anchor ASN Gateway.
6	Target non-anchor ASN Gateway forwards the Context-Report message to the target ASN BS.
7	Target ASN BS sends a Path Pre-Registration Request message to the target non-anchor ASN Gateway. Pre-registration is optional.

Step	Description
8	Target non-anchor ASN Gateway forwards the Path Pre-Registration Request message to the anchor ASN Gateway. Pre-registration is optional.
9	Anchor ASN Gateway sends a Path Pre-Registration Response message to the target non-anchor ANS GW. Pre-registration is optional.
10	Target non-anchor ASN Gateway forwards the Path Pre-Registration Response message to the target ASN BS. Pre-registration is optional.
11	Target ASN BS sends a Path Pre-Registration Acknowledge message to the target non-anchor ASN Gateway. Pre-registration is optional.
12	Target non-anchor ASN Gateway forwards the Path Pre-Registration Acknowledge message to the anchor ASN Gateway. Pre-registration is optional.
13	Target BS sends a Handover Response message to the serving BS.
14	Serving BS sends a MOB_BSHO-RSP message to the MS.
15	Serving BS sends a Handover Acknowledge message to the target BS.

Action Phase

The following figure and table provides a high-level view of the steps involved during the action phase of a controlled inter-ASN Gateway handover of an SS/MS from an anchored gateway to a non-anchored gateway.

Figure 49. Controlled Inter-ASN Gateway Handover Procedure - Action Phase

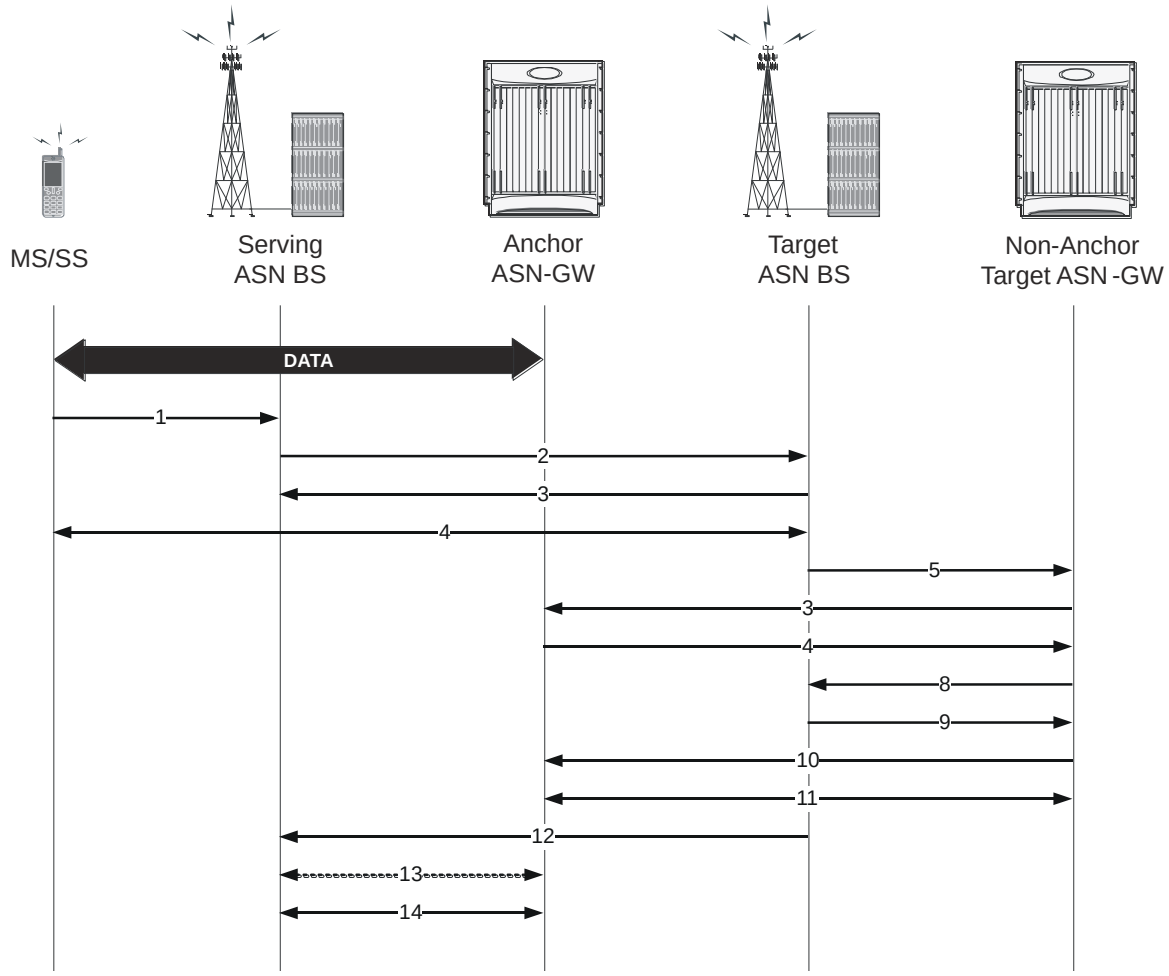


Table 38. Controlled Inter-ASN Gateway Handover Procedure - Action Phase Description

Step	Description
1	MS sends a MOB_MSHO-IND message to the serving ASN BS.
2	Serving ASN BS sends a Handover Confirm message to the target ASN BS.
3	Target ASN BS sends a Handover Acknowledge message to the serving ASN BS.
4	MS moves off of the serving ASN Gateway and re-enters the network through target ASN BS.
5	Target ASN BS sends a Path Registration Request message to the target non-anchor ASN Gateway.
6	Target non-anchor ASN Gateway forwards the Path Registration Request message to the anchor ASN Gateway.
7	Anchor ASN Gateway sends a Path Registration Response message to the target non-anchor ANS GW.
8	Target non-anchor ASN Gateway forwards the Path Registration Response message to the target ASN BS.
9	Target ASN BS sends a Path Registration Acknowledge message to the target non-anchor ASN Gateway.

Step	Description
10	Target non-anchor ASN Gateway forwards the Path Registration Acknowledge message to the anchor ASN Gateway.
11	Target non-anchor ASN Gateway sends/receives CMAC Key Count Update and Acknowledge messages to/from anchor ASN Gateway.
12	Target ASN BS sends a Handover Complete message to the serving ASN BS.
13	Anchor ASN Gateway sends/receives Path De-Reg Req/Rsp/Ack messages (to release the data path) to/from Serving BS.
14	R6 GRE data path terminated between Serving ASN BS and Anchor ASN Gateway.

Uncontrolled Anchor ASN Gateway to Non-Anchor ASN Gateway Handover

The following figure and table provides a high-level view of the steps involved in an uncontrolled inter-ASN Gateway handover of an SS/MS from an anchored gateway to a non-anchored gateway.

Figure 50. Uncontrolled Inter-ASN Gateway Handover Procedure

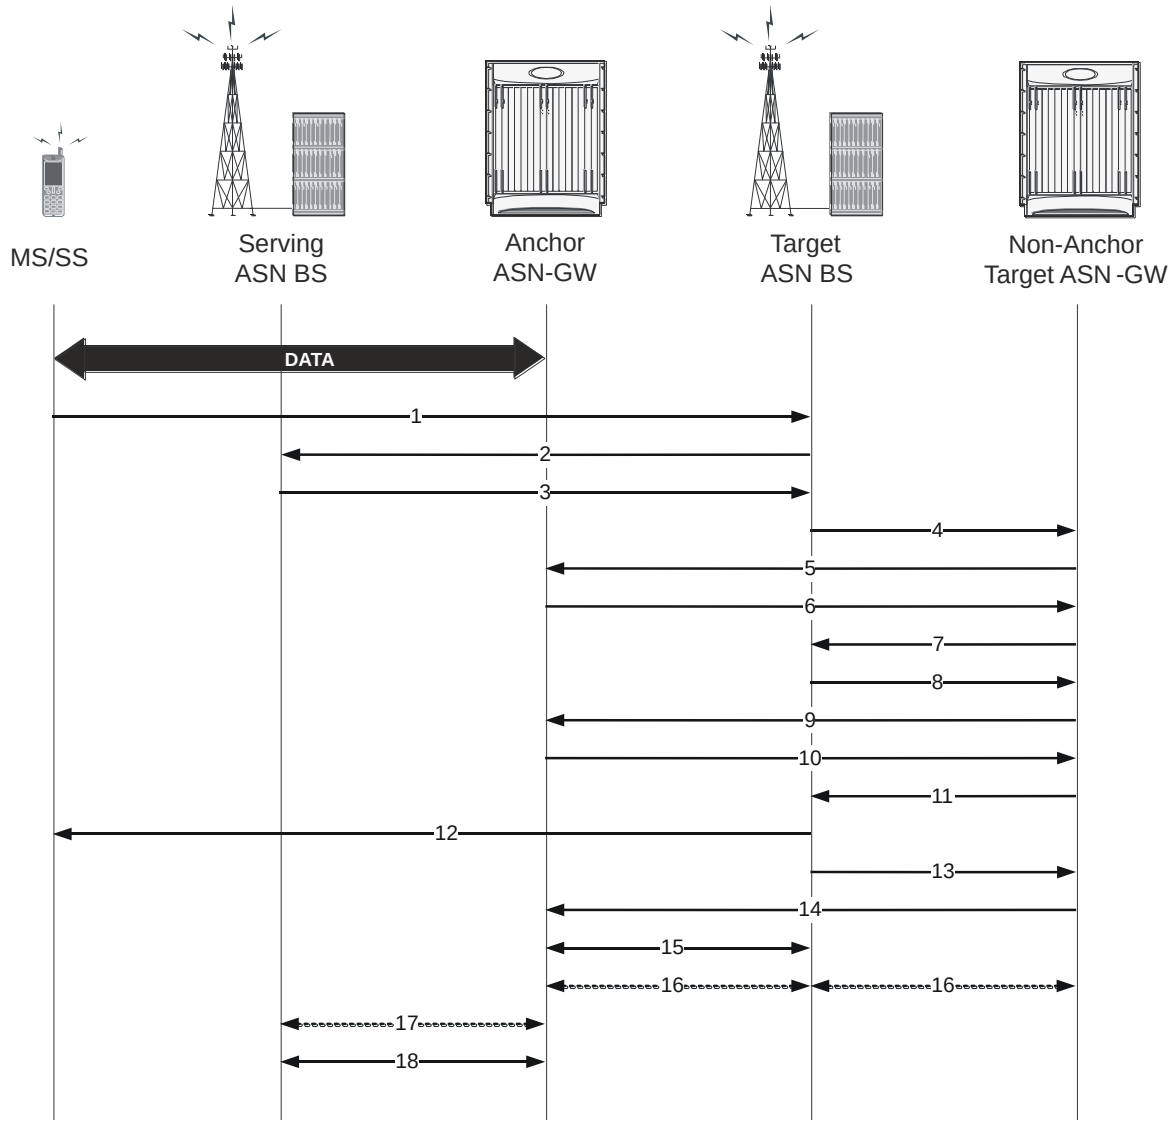


Table 39. Uncontrolled Inter-ASN Gateway Handover Procedure Description

Step	Description
1	MS sends RNG-REQ message to target ASN BS.
2	Target ASN BS sends Context-Request message to serving ASN BS.
3	Serving ASN BS sends Context-Report message with MS context information to target ASN BS.
4	Target ASN BS sends Context-Request message to target non-anchor ASN Gateway.
5	Target non-anchor ASN Gateway forwards Context-Request message to anchor ASN Gateway.
6	Anchor ASN Gateway sends Context-Report message with MS context information to target non-anchor ASN Gateway.

Step	Description
7	Target non-anchor ASN Gateway forwards Context-Report message to target ASN BS.
8	Target ASN BS sends Path Registration Request to target non-anchor ASN Gateway.
9	Target non-anchor ASN Gateway forwards Path Registration Request to anchor ASN Gateway.
10	Anchor ASN Gateway replies with Path Registration Response to target non-anchor ANS GW.
11	Target non-anchor ASN Gateway forwards Path Registration Response to target ASN BS.
12	Target ANS BS sends ranging response with RNG_RSP message to MS.
13	Target ASN BS sends Path Registration Acknowledge to target non-anchor ASN Gateway.
14	Target non-anchor ASN Gateway forwards Path Registration Acknowledge to anchor ASN Gateway.
15	R6 GRE data path established between Target ASN BS and anchor ASN Gateway. Data flow starts.
16	Target ASN BS sends/receives CMAC Key Count Update and Acknowledge messages to/from anchor ASN Gateway via target non-anchor ASN Gateway.
17	Anchor ASN Gateway sends/receives Path De-Reg Req/Rsp/Ack messages to release data path to/from serving BS.
18	R6 GRE data path terminated between Serving ASN BS and anchor ASN Gateway.

RADIUS-based Prepaid Accounting for WiMax

Online accounting is set up by the exchange of RADIUS Access-Request and Access-Accept packets. The initial Access-Request packet from the ASN GW and/or the home agent includes a prepaid accounting capability (PPAC) vendor specific attribute too the prepaid server (PPS). This indicates support for online accounting at the ASN and/or the home agent. If the subscriber's session requires online charging, the PPS assigns a prepaid accounting quota (PPAQ) to the PPC with RADIUS Access-Accept packets. As the session continues, the PPC and the PPS replenish the quotas by exchanging RADIUS packets.

Note the following:

- ASN GW operates as the prepaid client (PPC).
- In the case of a mobile IP call, both the ASN GW and the home agent work independently as the prepaid client. Both the ASN GW and the home agent send online access requests to the configured RADIUS servers independently.
- Only session-based online accounting is supported.

Obtaining More Quota after the Quota is Reached

The following figure and table provide a high-level view of the steps involved in allocating additional quotas for prepaid calls once the original quota is reached.

Figure 51. Call Flow Showing How Additional Quota is Obtained

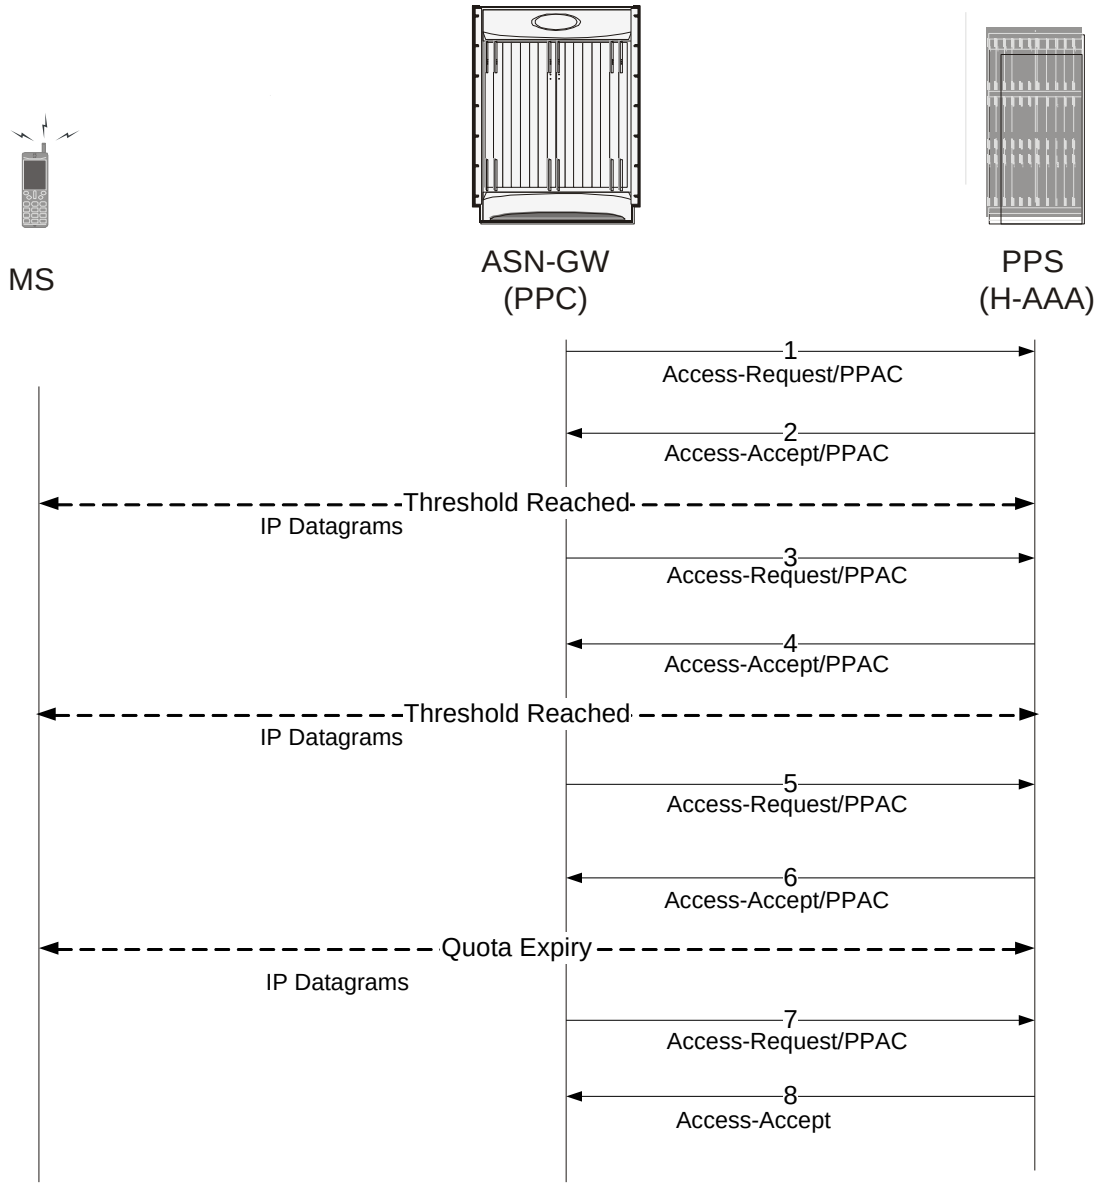


Table 40. Call Flow Showing How Additional Quota is Obtained

Step	Description
1	During network entry, a NAS sends an Access-Request packet to the HCSN. If the NAS supports a PPC, the NAS includes the PPAC attributes, indicating it prepaid capabilities.
2	If the subscriber session is a prepaid session, the PPS (HAAA) assigns the initial prepaid quota(s) by including one or more PPAQ attributes in the Access-Accept packet.
3	Once the threshold for the quota(s) is reached, the PPC sends an Authorize-Only Access-Request to request additional quota. The request contains one or more PPAQs that indicate which quota(s) need to be replenished to the PPS.

Step	Description
4	The PPS responds with an Access-Accept packet that contains one or more replenished quotas.
5	Once again, a threshold is reached for one or more of the quotas. The PPC sends an Authorize-Only Access-Request to the PPS to request more quota.
6	The PPS responds with the final quota in an Access-Accept. The final quota is indicated by the presence of the Terminate-Action subtype. The Terminate-Action subtype includes the action for the PPC to take once the quota is reached.
7	The quota expires. The PPC sends an Authorize-Only Access-Request packet to indicate that the quota has expired.
8	The PPS responds with an Access-Accept. If there are additional resources, the PPS allocates additional quotas and the service continues.

Applying HTTP Redirection Rule when Quota is Reached

The following figure and table provide a high-level view of the steps showing how the HTTP Redirection Rule is applied once a quota is reached.

Figure 52. Call Flow for Applying HTTP Redirection Rule on Quota-Reach

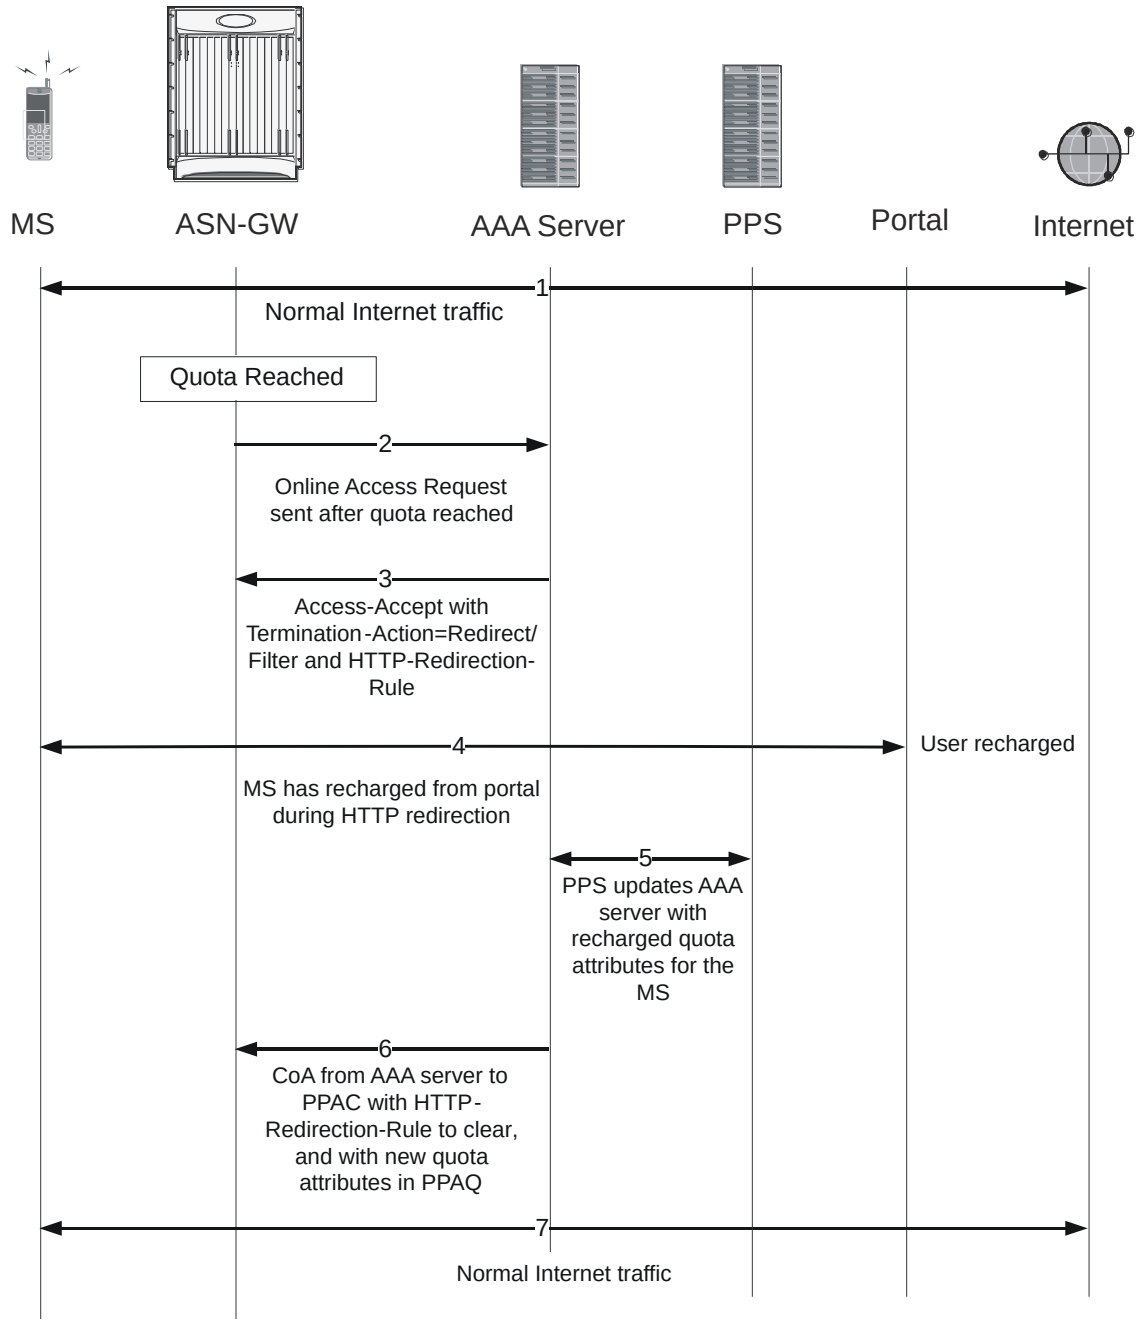


Table 41. Call Flow for Applying HTTP Redirection Rule on Quota-Reach

Step	Description
1	The Volume or Duration quota is reached. The Termination-Action is Request More Quota.
2	The PPC sends an Online Access Request to the AAA server and waits for Access-Accept.

Step	Description
3	The Access-Accept is received. It contains no additional quota attributes. The Termination-Action is Redirect/Filter. There is an HTTP Redirection Rule with redirect rule present in the Access-Accept.
4	The PPC (home agent) applies the HTTP Redirection Rule for the HTTP traffic. All other traffic is dropped. During this period, the MS recharges from the portal.
5	The PPC sends updated quota attributes to the AAA server based on the MS recharge from the portal.
6	The AAA server sends a CoA message to the PPC (home agent) with the new quota attributes in PPAQ and also sends the HTTP Redirection Rule to clear the HTTP Redirection rule at the PPC.
7	Normal traffic, including HTTP traffic, is allowed, per the new quota attributes.

Applying HTTP Redirection Rule When CoA is Received

The following figure and table show the steps involved in applying the HTTP Redirection Rule when the PPAC receives a change of authorization (CoA) from a AAA server.

Figure 53. Call Flow for Applying HTTP-Redirection Rule when CoA is Received

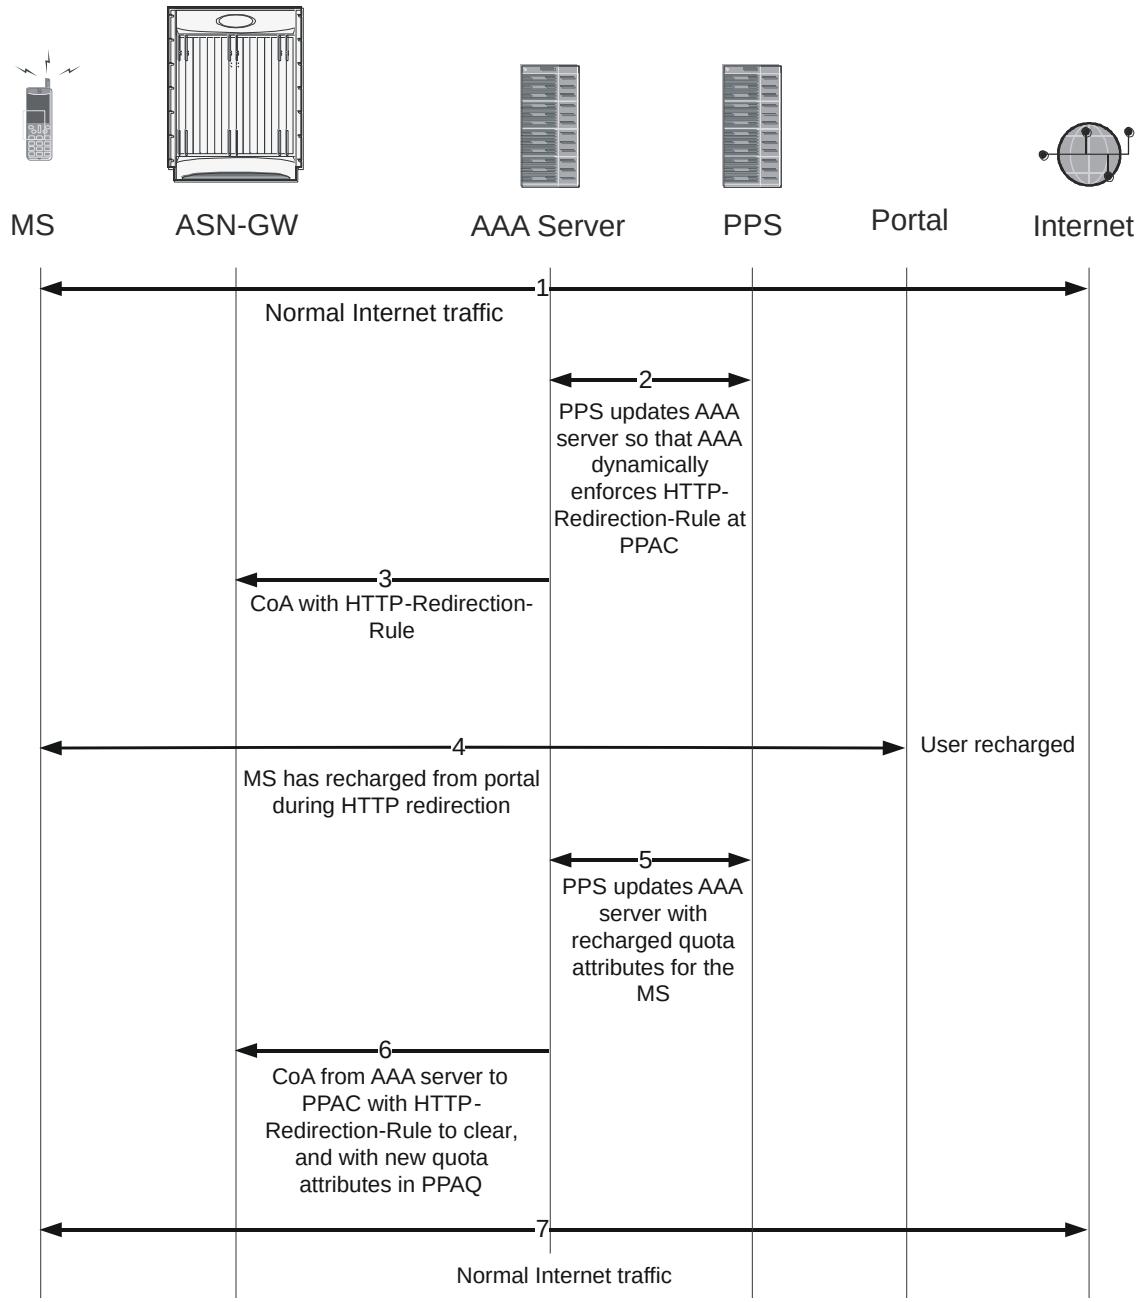


Table 42. Call Flow for Applying HTTP-Redirection Rule Received by CoA

Step	Description
1	The PPS updates the AAA server so that the AAA server dynamically enforces HTTP Redirection Rule at the PPC.
2	The AAA server sends a CoA message to the PPC (home agent) with the HTTP Redirection Rule.

Step	Description
3	The PPC (home agent) applies the HTTP Redirection Rule for the HTTP traffic. All other traffic is dropped. During this period, the MS is recharged from the portal.
4	The PPC sends updated quota attributes to the AAA server based on the MS recharge from the portal.
5	The AAA server sends a CoA message to the PPC (home agent) with the new quota attributes in PPAQ and also sends the HTTP Redirection Rule to clear the HTTP Redirection rule at the PPC.
6	Normal traffic, including HTTP traffic, is allowed, per the new quota attributes.

Terminating the Call when Quota is Reached

The following figure and table provide a high-level view of the steps involved in allocating additional quotas for prepaid calls once the original quota is reached.

Figure 54. Call Flow for Terminating the Call on Quota-Reach

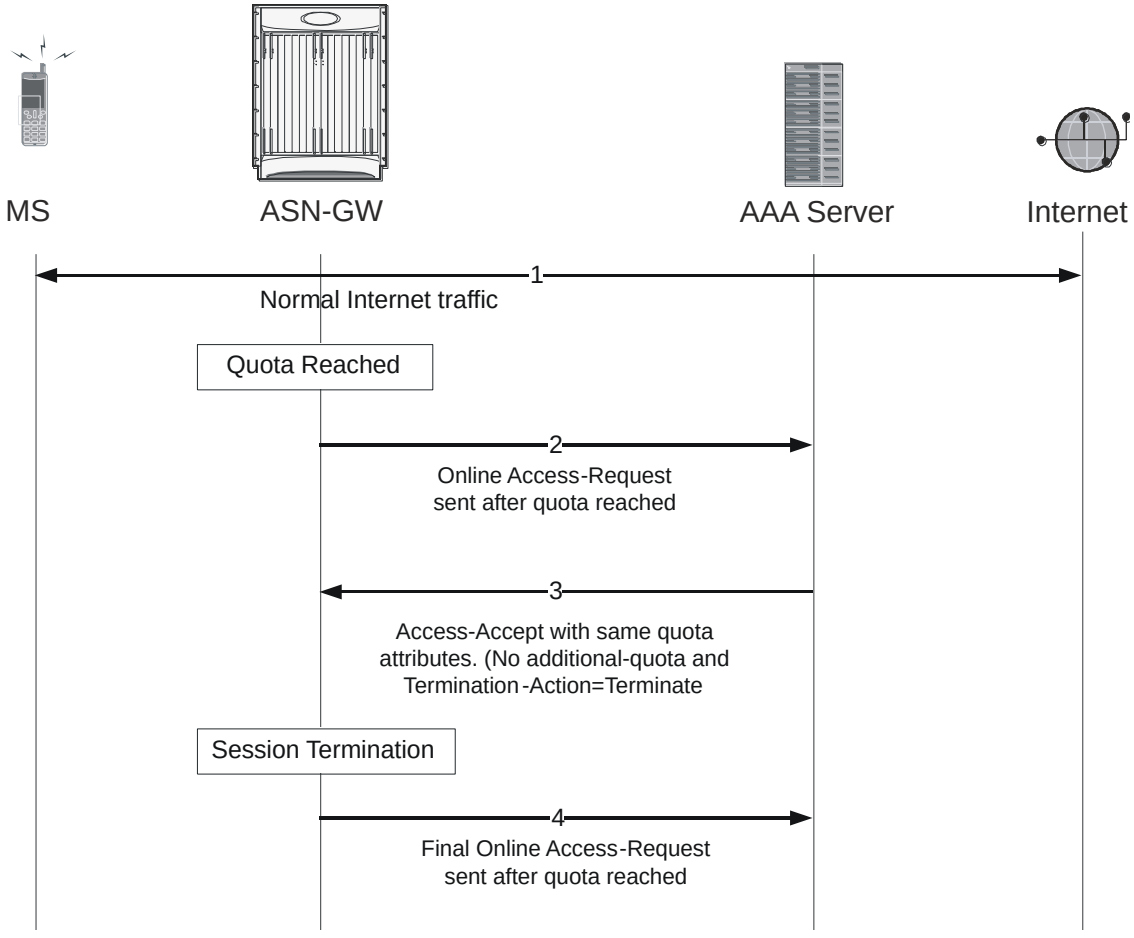


Table 43. Call Flow for Terminating the Call on Quota-Reach

Step	Description
1	Volume or Duration quota is reached. If the termination-action is Request-More-Quota, step 2 occurs next. If termination-action is Terminate, step 4 occurs next.
2	If the termination-action is Request-More-Quota, the PPC sends an Online-Access-Request to the AAA server and waits for Access-Accept.
3	The PPC receives the Access-Accept, which contains no additional quota attributes.
4	Session is terminated at the PPC (home agent) and at the ASN GW.
5	The PPC sends the final Online-Access-Request.

DHCP Relay Support for ASNGW

Following is a description of DHCP functionality to support ASNGW service. The functionality assumes the AAA to be RADIUS. To support Diameter AAA, the Access-Accept messages are correlated to DEA messages of the Diameter Protocol.

DHCP Keys

DHCP messages between the DHCP relay and DHCP server are authenticated by the DHCP Authentication Suboption. This requires that the DHCP relay and the DHCP server have a shared secret we call the DHCP-key.

The DHCP-key is specific between each DHCP Relay and DHCP server. The DHCP keys are derived from the DHCP-RK. The DHCP-RK key generation is internal to the AAA server and is transported as necessary to the authenticator and DHCP server using AAA protocol. The DHCP Key is derived from the DHCP-RK at the authenticator and at the DHCP server.

DHCP-RK and derived keys are not bound to individual user or authentication sessions, but to a specific DHCP server and DHCP relay/DHCP server) pairs. DHCP-RK is generated only on demand, but not for each EAP (re-)authentication taking place. Nevertheless, a DHCP-RK key along with the key identifier and lifetime values, are delivered to the authenticator during network access authentication of a MS (that is, it is carried by but otherwise unrelated to this specific MS). The lifetime and key identifier of DHCP-RK is managed by the AAA server. It is the responsibility of the AAA server to deliver a new DHCP-RK to the authenticator and DHCP server prior to the expiration of the DHCP-RK.

Key Generation

The DHCP-RK is created by the AAA server assigning the DHCP server to an authenticating subscriber. A different 160-bit random DHCP-RK is generated for every DHCP server.

From the DHCP-RK an authenticator generates DHCP-key for a specific DHCP relay/ DHCP server pair if requested by this DHCP relay. The DHCP-key is also generated by the DHCP server when a DHCP message arrives from a DHCP relay for which the DHCP server has no key yet.

From the DHCP-RK an authenticator generates a DHCP-key for a specific DHCP relay/ DHCP server pair if requested by this DHCP relay. The DHCP-key is also generated by the DHCP server when a DHCP message arrives from a DHCP relay for which the DHCP server has no key yet.

Key Distribution

The DHCP-RK keys are generated by the AAA server and are transported to the DHCP server and the Authenticator with the AAA protocol. The DHCP key generated by the authenticator is transported to the DHCP relay via WiMAX specific R4 signaling (Context Req/Rsp msg). The DHCP keys generated by the DHCP server are never transported outside of the DHCP server.

- **DHCP-RK Key**: Generated by AAA, transported to authenticator and DHCP server through AAA protocol.
- **DHCP Key**: Generated by authenticator and DHCP server, transported to DHCP relay and DHCP server via WiMAX-specific R4 signaling. Never transported outside of the DHCP server.

DHCP key distribution covers two scenarios. In the first scenario the authenticator and DHCP relay are co-located in the same entity. In the second scenario, no re-authentication takes place when the MS moves to a different anchor ASN

hosting a new DHCP relay, so the anchor authenticator is continued to be used, and provisions the new DHCP relay with the required keys.

DHCP Relay in ASNGW for PMIPv4 Calls

The following steps are based on R3 already being secured. If R3 is not secured, the DHCP relay adds the authentication sub-option (as explained in step 12) to have data integrity and replay protection for relayed DHCP messages.

1. The MS sends a DHCP Discover as a broadcast message. The DHCP message is sent on the MS's Initial service flow setup over R1 interface to the BS.
2. The DHCP Discover message is forwarded from the BS to the DHCP Relay present in the ASN through the data path established for the ISF (Initial Service Flow) traffic.
3. The DHCP Relay in ASN intercepts and changes the destination IP address from broadcast to unicast and configures the giaddr field in the DHCP payload. It then sends the DHCP Discover message to the DHCP server of the MS based on configuration information. The configuration information in the most generic case is downloaded via the AAA but it may also be statically provisioned. If the Datapath is per MS or SF, the MS context is found based on the Datapath and not on the MAC address. If the Datapath is per BS the MS context is found based on the MAC address or MS NAI
4. DHCP servers receiving the DHCP Discover request reply by sending a DHCP Offer message including an offered IP address.
5. The DHCP Relay in the ASN forwards the DHCP replies to the MS. The DHCP Offer message is sent from ASN GW to BS through the Data Path. The destination IP address of the DHCP Offer message sent to MS is a unicast one. Normally DHCP servers or relay agents attempt to deliver the DHCP Offer to a MS directly using unicast delivery. However, some MS implementations are unable to receive such unicast IP datagram until they know their own IP addresses. To work around this, an MS's broadcast address may be used in the DHCP Offer message. The ASN checks the BROADCAST (B) flag in the DHCP Offer message. If this flag is set, the ASN uses a broadcast address to send DHCP Offer message. Otherwise the unicast address is sent, but the delivery is over a unicast CID
6. BS sends DHCP Offer message to the MS on the MS's Initial Service Flow.
7. The MS receives a DHCP Offer message and sends a DHCP Request to the selected DHCP server as a broadcast message confirming its choice of the DHCP Server.
8. DHCP Request message is sent from BS to DHCP relay in ASN through the established Data Path.
9. The DHCP Relay in the ASN prompts the PMIP client to initiate the Mobile IP Registration procedure. The PMIP client uses the HoA information to construct a Mobile IP Registration Request message. This message contains HoA and CoA for this MS. The source address for this R3 message is CoA, and the destination address is HA address.
10. The HA responds with the Mobile IP Registration Response message in which the source address for this R3 message is the HA address, and the destination address is CoA.
11. After the establishment of the MIP tunnel the PMIP client informs the DHCP Relay with the MIP registration result. The DHCP Relay in ASN relays the DHCP Request with the optional MIP registration result encapsulated in the WiMAX vendor specific relay agent suboption to the DHCP server.
12. The selected DHCP server receives the DHCP Request and replies with a DHCP Ack containing the configuration information requested by the MS.
13. The DHCP Relay relays the DHCP Ack to the BS.
14. BS sends DHCP Ack message to the MS on the MS's provisioned Initial Service Flow.If MS doesn't receive a DHCP Ack, or DHCP NAK message when timeout, it retransmits the DHCP Request. If neither DHCP Ack nor DHCP NAK is received when the maximum retransmission reached, MS restarst the IP initialization process.

DHCP Relay Requirements for Simple IP Calls

The DHCP relay handles all DHCP messages sent by the MS to the broadcast IP address.

- The DHCP relay is configured with the DHCP server address during the MS authentication. The AAA server sends the address of the DHCP server in the Access Accept message. The DHCP relay uses this address to relay the DHCP messages from the MS to the DHCP server.
- Upon receiving a DHCP DISCOVER message from the MS, the DHCP relay verifies the "chaddr" field in the DHCP header and matches the MS MAC address associated with the R6/R4 over which the DHCP message is received. This is feasible without any additional option in the DHCP message since the DHCP relay is collocated with the Anchor ASN (ASN-GW for profiles A and C). This is done to prevent MAC address spoofing by a rogue MS.
- If the DHCP relay determines that the MS has included a MAC address in the chaddr field of the DHCP DISCOVER message that does not match the known MAC address of the R6/R4 over which the DHCP message was received, the DHCP relay consider the following action: A rogue MS is trying to spoof MAC address. In this case, the DHCP relay informs the DPF to initiate R6 teardown.
- After determining the NAI to be used for the request, the DHCP relay adds the relay agent option to the original DHCP message and sets the Subscriber-ID suboption to the NAI used for MIP associated with MS. If there is a secure communication channel between the DHCP relay and the DHCP server, the relay and server may choose to omit the authentication suboption.
- If a DHCP Decline message is received, the DHCP Relay forwards the message to the DHCP Server. The messaging between the DHCP relay and DHCP server is transported over the R3 interface.
- When DHCP relay receives the DHCP OFFER message from the DHCP server, it relays it to the MS. If the DHCP server included the authentication suboption is in the relay agent option, the DHCP relay validates it before relaying the DHCP OFFER to the MS.
- The DHCP relay behavior for handling DHCPREQUEST or DHCPDECLINE from the MS is same as in the case of DHCP DISCOVER.
- The DHCP relay intercepts the DHCP renewal and release messages, and verifies the content of the message. If R3 is not secured (e.g., by IPSec), the DHCP relay adds the relay agent authentication suboption to the message before relaying it to the DHCP server.

DHCP Relay Requirements for PMIPv4 Calls

The DHCP relay handles all DHCP messages sent by the MS to the broadcast IP address. The DHCP relay is configured with the DHCP server address during the MS authentication. The AAA server sends the address of the DHCP server in the RADIUS Access Accept message or Diameter WDEA command. The DHCP relay uses this address to relay the DHCP messages from the MS to the DHCP server.

- Upon receiving a DHCP DISCOVER message from the MS, the DHCP relay verifies the chaddr field in the DHCP header matches the MS MAC address associated with the R6/R4 over which the DHCP message is received. This is feasible without any additional option in the DHCP message since the DHCP relay is collocated with the Anchor ASN (ASN-GW for profiles A and C). This is done to prevent MAC address spoofing by a rogue MS.
- If the DHCP relay determines that the MS has included a MAC address in the chaddr field of the DHCP DISCOVER message does not match with the known MAC address associated with the R6/R4 over which the DHCP message was received, the DHCP relay considers that a rogue MS is trying to spoof the MAC address. In this case, the DHCP relay may inform the DPF to initiate R6 teardown.

- After determining the NAI to be used for the request, the DHCP relay adds the relay agent option to the original DHCP message and sets the Subscriber-ID suboption to the NAI used for MIP associated with MS. If there is a secure communication channel between the DHCP relay and the DHCP server, the relay and server may choose to omit the authentication suboption. T
- If a DHCP Decline message is received, the DHCP Relay forwards the message sto the DHCP Server.
- When DHCP relay receives the DHCP OFFER message from the DHCP server, it relays it to the MS. If the DHCP server included the authentication suboption in the relay agent option, the DHCP relay validates it before relaying the DHCP OFFER to the MS.
- The DHCP relay behavior for handling DHCP REQUEST or DHCP DECLINE from the MS is same as in the case of DHCP DISCOVER
- When DHCP relay receives the DHCP REQUEST message from the MS, it prompts the PMIP4 client to initiate MIPv4 registration procedures. It passes the requested IPv4 address and the HA information to the PMIP4 client. The PMIP4 client performs the registration with the FA and HA on behalf of the MS. The PMIP4 client informs the DHCP relay with the MIPv4 registration result.
- Upon receipt of such indication, the DHCP relay relays the DHCP REQUEST message with the MIP registration result encapsulated in the vendor specific relay agent suboption code 1 to the DHCP Server. If this suboption is not sent to the DHCP server and the MIP registration indicates a failure, the DHCP relay does not forward the DHCP REQUEST message to the DHCP server and the network performs an exit for the corresponding MS. When the DHCP relay receives the DHCPACK message from the DHCP Server, it relays the DHCPACK message to the MS.
- Since AAA can assign different HAs (e.g., when dynamically assigning HA from a pool) and each HA handles different MS subnets, the assigned HA needs to be passed to DHCP server to allow choosing the matching MS address pool. DHCPDISCOVER and DHCPREQUEST from MS SHOULD include HA IP address in same vendor specific relay agent suboption code 2 to the DHCP Server

RADIUS Based Procedures for Prepaid Accounting

In prepaid accounting, the ASNGW works as the prepaid client (PPC). When there is a mobile IP call, both the ASNGW and HA can work as the PPC independently. However, either the HA or ASNGW, not both, is responsible for prepaid accounting. In the case of Simple IP calls, the ASNGW can act as the prepaid client for prepaid subscriber calls.

Online accounting is set up by the exchange of RADIUS Access-Request and Access-Accept packets. The initial Access-Request packet from the ASNGW and or the HA includes a prepaid accounting capability (PPAC) VSA to the PPS indicating support for On-line accounting at the ASN and or the HA.

If the subscription session requires online charging, in the case of session-based prepaid accounting, the AAA server (or PPS) assigns a prepaid accounting quota (PPAQ) to the PPC (HA or ASNGW) by encoding the PPAQ values in the RADIUS Access-Accept packets. As the session continues, the PPC and the PPS replenish the quotas by exchanging RADIUS packets. In the case of IP 5-tuple flow-based prepaid accounting, when traffic is received for a configured prepaid IP flow, the PPAC sends an authorize-only access-request to the AAA server to request and receive quota attributes for that flow.

In the case of session-based prepaid accounting, quotas are allocated to each session. The Service-ID in the PPAQ is sent to the IP-Address corresponding to the IP-Session.

Flow- based Prepaid Accounting

As per NWG WiMAX architecture, flow-based prepaid accounting can be either IP 5-tuple flow-based, PDF ID based. or SDF ID based. Currently, only the IP 5-tuple flow-based prepaid accounting is supported.

The ASNGW receives PPAQ attributes separately for each flow that requires flow-based accounting. The ASNGW performs pre-paid accounting for each of the flows for which prepaid accounting is configured. However, there can be flows for which pre-paid accounting is not requested. For those flows, session-based off-line accounting procedures are applied.

PPAQ attribute contains the Rating-Group-ID attribute which corresponds to one or more IP 5-tuple flows, depending on the configuration. The Rating-Group-ID identifies the flow for which pre-paid accounting is applicable. The format of Service-ID for flow-based accounting is as follows:

- When the first data packet which belongs to a Rating-Group-ID for which prepaid is configured is received, the PPC (ASNGW or HA) requests quota attributes for this Rating-Group-ID. The PPC sends an Authorize-Only access-request and encoding PPAQ with Rating-Group-ID in it. The AAA server sends back the available/allocated quota attributes in an Access-Accept by encoding the PPAQ with the same Rating-Group-ID along with quota parameters.
- In the case of IP 5-tuple flow-based prepaid accounting, the quota is requested based on Rating-Group-ID only. Configuration is required to map IP flows to Rating-Group-IDs.
- On quota expiry for a Rating-Group-ID, the PPC (ASNGW/HA) sends an on-line access request to renew the quota. The ASNGW/HA includes the PPAQ for the corresponding Rating-Group-ID only in the on-line access request.
- In the case of quota expiry for a flow, the ASNGW/HA renews the quota for that particular Rating-Group-ID. If there is no quota allocated for the rating-group, the traffic for that rating-group is dropped for some period (based on configuration). If there is any more traffic matching that rating-group after that period (based on configuration), the quota is requested again for that rating group.

Configuring Rating-Group-ID

IP 5-tuples are configured by using current active charging service ruledef configurations. Each ruledef is associated with a rating group ID by configuration. Several ruledefs (i.e., several IP 5-tuple rules) may be mapped to a single rating-group ID. Prepaid quota attributes are requested only on a ratinggroup ID basis.

Prepaid Tariff Switching

Prepaid Tariff Switching will be supported for IP 5-.tuple flow-based prepaid accounting.

Hotlining with Flow-based Prepaid Accounting

Hotlining is applied on a session basis. That is, when hotling has to be applied because of either quota expiry for a rating-group or because a CoA is received with hotlining parameters for a session, the entire session is hotlined even though there may be some rating-groups with quota remaining.

RADIUS-based Procedures for WiMAX Hotlining

The hotlining feature provides a WiMAX operator with the capability to address issues with users that would otherwise be unauthorized to access packet data services. The hotlining device (HLD) can be located at the ASN or CSN.

Types of Hotlining

There are two methods defined by which the HAAA indicates that a user is to be hot-lined:

- **Profile based hotlining**: Hot-line profile(s) with all hotlining rules are pre-provisioned at the HAAA. The HAAA sends a hot-line profile identifier in the RADIUS message (Access-Accept and Change of Authorization) when the hotlining is activated.
- **Rule based hotlining**: Hotlining rules (filter rules, IP or HTTP redirection rules) are sent in the RADIUS message (Access-Accept and Change of Authorization) by the HAAA when the hotlining is activated. Cisco Systems supports profile based hotlining and rule-based hotlining with HTTP redirection rules only. Rule-based hotlining with IP-Redirection-Rules and NAS-Filter-Rules are currently not supported.

Based on the status of the user's session, there are two ways users can be hot-lined:

- **Active session hotlining**: The user starts a normal packet data session. At some point in the session, the HAAA receives a trigger for hotlining from the hotlining application (HLA).
- **New session hotlining**: The trigger from the HLA arrives prior to the user access authentication.

Once the hotlining is resolved, the packet data session returns to normal. Both these approaches are discussed in the following sub-sections.

Hotlining for Prepaid Accounting Sessions

In the case of mid-session hotlining, if an access-request is sent by the system to replenish the quota, the server may not have enough quota to allocate for the session. The AAA server may send hotlining attributes (profile-id or hotlining rules) and the termination-action in the PPAQ set to "Redirect/Filter" in the access-accept. In this case, the HA or ASNWG installs the hotlining rules when the quota is exhausted and the session becomes a hotlined session. If a CoA is received with hotline attributes, the hotlining rules are installed and the session becomes a hotlined session.

In the case of hotlining during initial session establishment, the AAA server may send hotlining attributes and 0 quota in the PPAQ. The termination-action in the PPAQ is set to "Redirect/Filter". In this case, the session would start as a hotlined session and the hotline rules are nstalled.

Hotline rules are either derived from the hotline rule attributes in access-accept in the case of rule based hotlining, or from the locally configured hotline profile-id (Filter/Redirection rules which are configured under the locally configured active-charging rulebase). The hotline profile-id is received from the AAA server in access-accept or in a CoA

Active Session Hot-lining Call Flows

The active IP session hot-lining is invoked when the user is currently engaged in a packet data session and the HAAA receives hot lining trigger from the HLA. Below figure depicts the call flow between the HLD, HAAA and HLA.

1. User is in an active IP session which is not hotlined.
2. The HLA detects that the user needs to be hot-lined. This is indicated to the HAAA by sending a Hotlining Active Trigger.

3. When the home AAA server receives notification from the hotline applicaiton, it records the hotlining state against the user record in the database. The home AAA server determines whether the user has an ongoing packet data session. If the user has an ongoing packet data session, the home AAA server initiates the Active-Session Hotlining procedure. The home AAA server uses the contents of the Hotline Capability VSA and other local policies to determine which access device will be the Hotlining Device for the session. The home AAA server sends a RADIUS CoA-Request to the HLD with either Profile based hotlining or Rule based hotlining
4. Upon receipt of the RADIUS COA, if the HLD can honor the request, it responds with a RADIUS COA Ack to the HAAA. If the HLD cannot honor the request, it responds with a COA NAK message. Based on the local policy, HAAA may either retry sending the hotlining request to the HLD or it may send a RADIUS Disconnect Message (DM) to the HLD to terminate the session. ASNGW/HA always sends COA ACK if the session is present.
5. The HLD sends a RADIUS Accounting Request (Stop) indication for the active data session, with Session Continue set to true.
6. The HLD sends a RADIUS Accounting Request (Start) for the hotlined session with Beginning of Session set to False. If the Session-Timeout attribute was included in step 3, the HLD initiates session teardown (i.e., tear down of the service flows associated with the IP session) when the duration specified in the Session-Timeout attribute has elapsed and the user's session is still hotlined. After tearing down the service flow(s), the HLD sends an Accounting Request (Stop) to the HAAA to inform that the user's IP session has ended. If the Session-Timeout times out while a session is hot-lined, the session is terminated.
7. Since the user's data session is hot-lined in mid session, the user's data traffic is affected. Based on the hotlining rules set at the HAAA and indicated by it in the RADIUS COA earlier, the uplink and/or downlink data traffic of the user is either dropped, disconnected, or blocked, and redirected to the HLA by the HLD.
8. Once the hotline status is applied to the user status, the HLA notifies the user of his/her hotlined status and resolves the issue. If the condition that triggered the hot-lining session is not cleared, the HLA may terminate the session. In this case, the HAAA is notified by the HLA. Upon receipt of this notification, the HAAA sends a RADIUS Disconnect Message to the HLD. The accounting records are stopped and session termination is initiated. This may happen automatically at the HLD if the user's hotlined status does not change during the Session-Timeout value. If the condition that triggered the hotlining session is cleared, the HLA sends the Hotlinig Inactive Trigger to the HAAA to clear the hotlined status of the user.
9. Upon receipt of the Hotlining Inactive Trigger, the HAAA sends a RADIUS COA message to the HLD with appropriate attributes. Note that this may not be the same HLD that initially handled the activation of the hotlining. This may occur due to events such as handoff.
10. Upon receipt of the RADIUS COA, if the HLD can honor the request it will respond with a RADIUS COA Ack to the HAAA. At that point the Hotline Session-Timeout timer is turned off. If the HLD cannot honor the request, it responds with a COA NAK message. Based on the local policy, the HAAA may retry sending the hotlining signal to the HLD or it may send a RADIUS Disconnect Message to the HLD to terminate the session. In this case, the HLD sends a RADIUS Accounting Request (Stop) message to the HAAA to indicate the end of the IP session for the user. This occurs after the Disconnect Message is processed and the service flow(s) associated with the IP session are torn down.
11. The HLD generates a RADIUS Accounting Request (Stop) with Session Continue set to True message for the hotlined packet data session.
12. The RADIUS Accounting Request (Stop) message with Session Continue set to True is followed by a RADIUS Accounting Request (Start) message with Beginning of Session set to False. This indicates the start of the normal packet data session.
13. The user continues the packet data session and the traffic is routed normally. During the hotlined active status in the HLD, the byte, packet, and duration counts for user's hotlined IP session is counted towards the overall byte and packet counts.

Active Session Hotlining for Prepaid Users

Active IP session hotlining is invoked when the prepaid user is engaged in a packet data session and the HAAA/PPS does not grant additional quota to the user. The figure below shows the call flow between HLD/PPC, HAAA/PPS and HLA.

1. The prepaid user is in an active IP session that is not hotlined.
2. The threshold for the prepaid quota(s) is reached.
3. PPC requests additional quota by sending an Authorize-Only Access-Request to the PPS containing one or more PPAQ(s) indicating which quota(s) need to be replenished.
4. PPS responds with an Access-Accept packet. The balance on the user account is too low for additional quota to be allocated. Hotlining is triggered for the user to replenish his/her account. Access-Accept is sent to the HLD with either Profile-based Hot-lining or Rule-based Hot-lining.
5. From this point on all steps are identical to "Active Session Hotlining for Prepaid Users."

Hotlining During Initial Network Entry

During initial network entry, hotlining is invoked. (Triggers for invoking hotlining are not within the scope of this section.) Examples include limited access to emergency services, empty prepaid accounts, or mobility restriction applied to a fixed or nomadic subscription. This occurs when H-AAA detects that initial network entry is being performed at a BS that does not belong to the network entry zone of the MS. I

1. The MS performs EAP authentication of initial network entry.
2. The Authenticator sends Access-Request as part of the authentication procedure. The H-AAA server acquires the ASN hot-lining capabilities.
3. If H-AAA activates hotlining, it sends an Access-Accept to the Authenticator/HLD with the appropriate attributes. The H-AAA may activate hotlining, depending on application-specific conditions such as emergency network entry (indicated by ES specific NAI), mobility restrictions applying to fixed or nomadic subscribers, or an empty prepaid account.
4. The anchor SFA located with the authenticator establishes the initial service flow (ISF) for the MS.
5. The MS gets an IP address from network side if an IP address is required for hotlining.
6. The authenticator/HLD sends a RADIUS Accounting Request (Start) for the hotlined session to indicate the activation of hot-lining.
7. Based on the hotlining rules received from the H-AAA server, the uplink and/or downlink data traffic of the user is either dropped or blocked, and redirected to the HLA by the HLD.
8. If the HAAA detects that the condition that triggered the hot-lining of the session is cleared, the HAAA sends a Radius COA message to the Authenticator/HLD with the appropriate attributes.
9. Upon receipt of the Radius COA, the authenticator/HLD responds with a Radius COA Ack to the HAAA.
10. The authenticator/HLD sends a Radius Accounting Request (Stop) to the HAAA to indicate the inactivation of the Hotlining.

Tariff Switching for Prepaid Accounting

The PPC and the PPS may support the tariff switching mechanism described in this section. This mechanism is useful if, for example, traffic before 18:00 is rated at rate r1 and traffic after 18:00 is rated at rate r2. The mechanism requires the PPC to report usage before and after the switch occurrs.

The PPC indicates support for tariff switching by setting the appropriate bit in the PPAC. If the PPS needs to signal a tariff switch time, it sends a PTS attribute that indicates the point in time when the switch will occur. This indication represents the number of seconds from current time (TariffSwitchInterval TSI).

At some point after the tariff switch, the PPC sends another Access-Request, as a result of the user having logged off or the volume threshold being reached. The PPC reports how much volume was used in total (in a PPAQ attribute) and how much volume was used after the tariff switch (in a PTS VUATS subtype attribute). In situations with multiple tariff switches, the PPS must specify the length of the tariff switch period using the TimeIntervalAfterTariffSwitchUpdate (TITSU) in the PTS attribute.

When a TITSU is specified in the PTS, the PPC generates an Access-Request within the time after TSI and before TITSU expires. Note that, typically, the PPC is triggered by the Volume Threshold. However, it is possible that during period r2, resources may not be entirely consumed and the threshold not reached. The TITSU attribute ensures that, even in this case, the PPC will generate the new Access-Request in good time.

Note that it is not efficient to use the tariff switching mechanism for services that are metered based on time and uninterrupted consumption. Also note that separate services flows may have individual tariff periods.

CSN Procedure Flows

This section provides an overview of CSN procedure and working of ASN Gateway in CSN procedure.

PMIP4 Connection Setup and Call Flow with DHCP Proxy

This section describes the CSN procedure of simple IP with DHCP proxy triggering PMIPv4 for a WiMAX subscriber.

The following figure and table provide a high-level view of the steps involved in PMIP4 connection and call flow of an SS/MS.

Figure 55. PMIP4 Connection Setup Call Flow

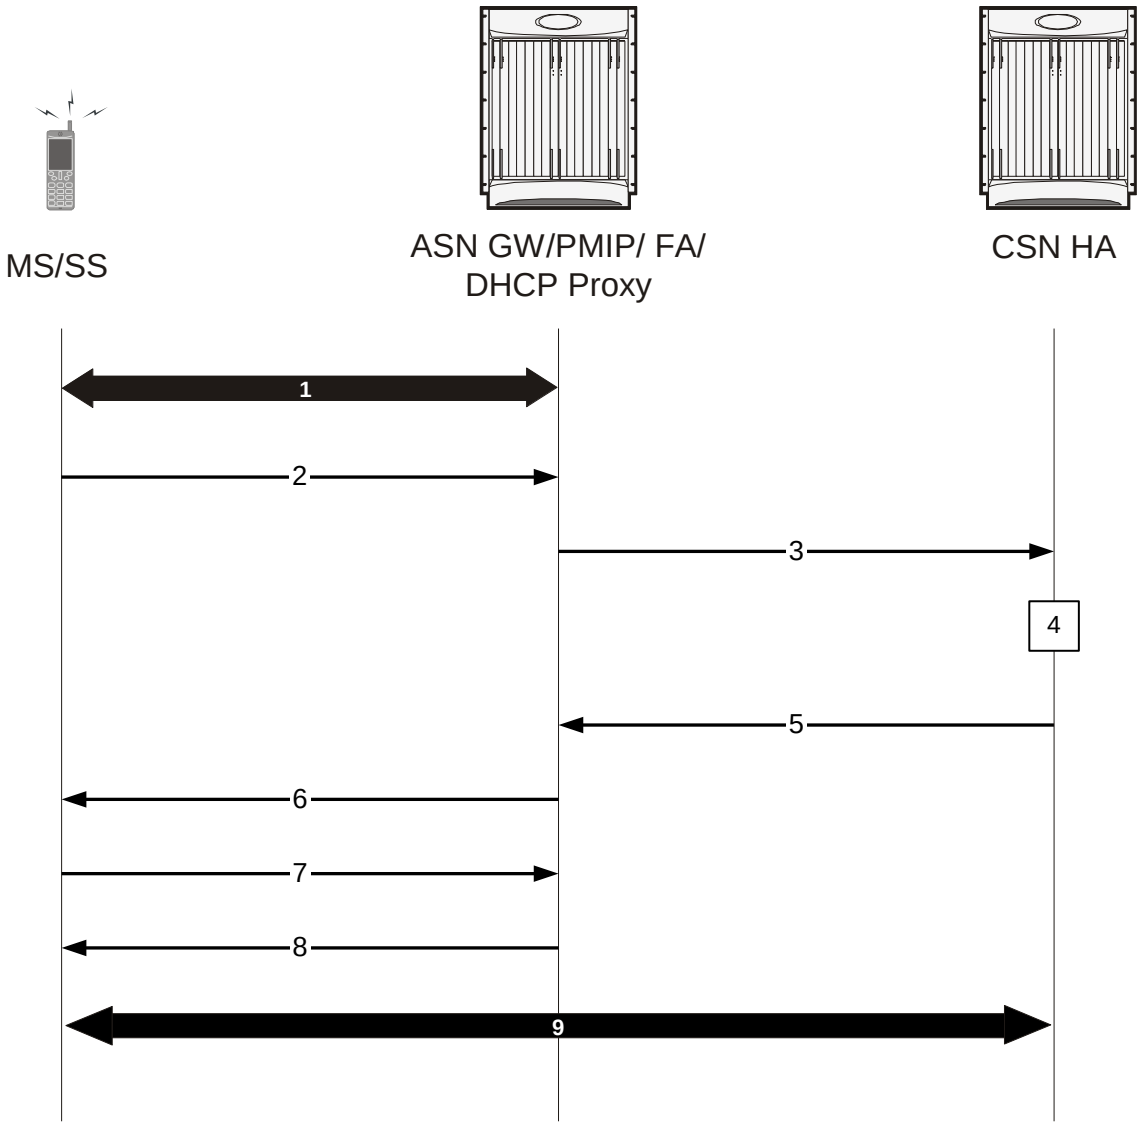


Table 44. PMIP4 Connection Setup Call Flow Description

Step	Description
1	Initial network entry completed as described in ASN Procedures.
2	MS sends DHCP DISCOVER message to DHCP Proxy (co-located with ASN Gateway) to discover a DHCP server for IP host configuration.
3	Upon receiving the DHCP DISCOVER message, the DHCP Proxy in the NAS triggers the PMIP4 client to initiate 8 the Mobile IPv4 Registration procedure.The PMIP4 client uses the HoA information and constructs a Mobile IPv4 Registration Request message and sends the Mobile IPv4 Registration Request to the FA address. The FA forwards the registration request to the CSN HA.
4	CSN HA processes the MIPv4 Registration Request.If a HoA is 0.0.0.0 in the Mobile IP Registration Request message, the HA assigns a HoA. Otherwise, the HoA in the Mobile IP Registration Request message is used.
5	The HA responds with the Mobile IP Registration Response message.The source address for this Mobile IPv4 message over R3 is HA, and the destination address is FA-CoA.The FA forwards the message to the PMIP4 client. The PMIP4 client passes this information to the DHCP proxy.
6	The DHCP proxy sends the DHCP OFFER message to the MS.
7	MS sends a DHCP REQUEST to the DHCP Proxy with the information received in the DHCP OFFER.
8	The DHCP Proxy acknowledges the use of this IP address and other configuration parameters by sending the DHCP ACK message.
9	WiMAX session established between MS and CSN HA.

PMIP4 Session Release

This section describes the CSN procedure of PMIPv4 session release during a WiMAX subscriber session.

The following figure and table provide a high-level view of the steps involved in PMIPv4 session release and termination of connection an SS/MS.

Figure 56. PMIP4 Session Release Call Flow

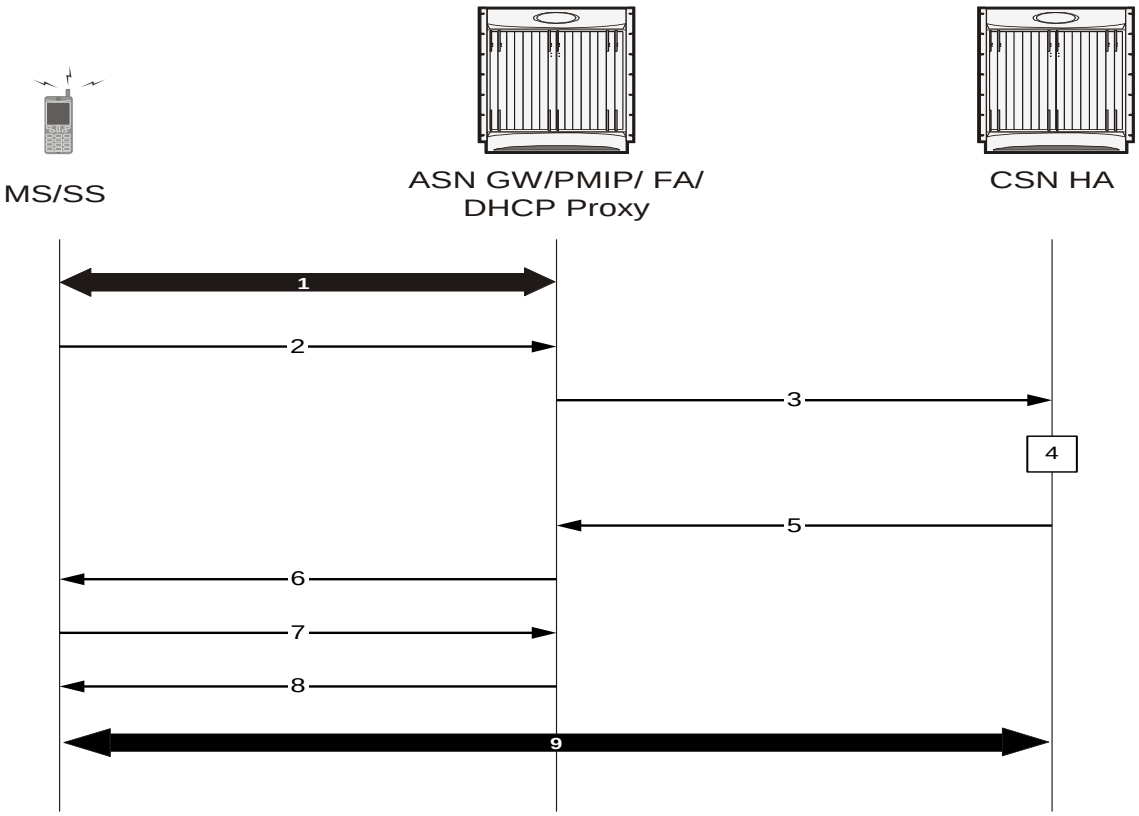


Table 45. PMIP4 Session Release Call Flow Description

Step	Description
1	The session release trigger send by MS sending DHCP-Release message to the ASN GS or DHCP proxy has expired on lease time or FA initiates session release.
2	ASN Gateway initiates the session release with PMIPv4 client by sending FA_Revoke_Req and sends PMIP De-Reg RRQ (Registration Revocation) message to CSN HA.
3	CSN HA starts release of MIP binding.
4	CSN HA sends PMIP De-Reg RRQ (Registration Revocation) message to ASN Gateway and PMIP client sends GA_Revoke_Rsp message to ASN Gateway.
9	WiMAX session terminated between MS and CSN HA.

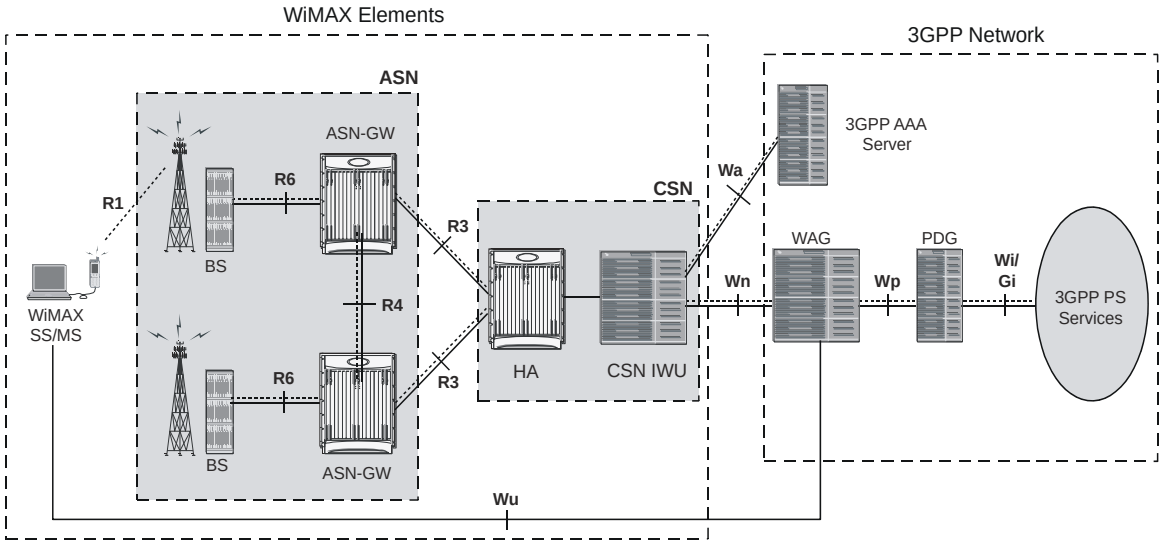
WiMAX Deployment with Legacy Core Networks

ASN Gateway can inter-operate with a 3GPP overlay and 3GPP2 overlay and provide continuity support for 3GPP2 and WiMAX handovers.

ASN Gateway Interoperability with 3GPP Overlay

The following figure shows a typical interoperability scenario between WiMAX and 3GPP legacy networks with reference points and interfaces.

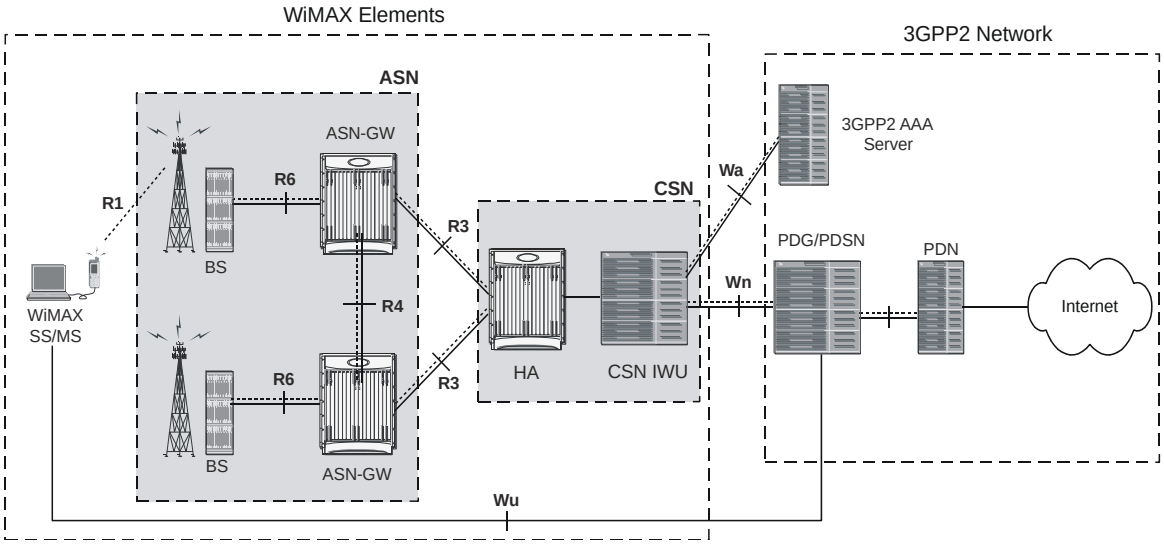
Figure 57. ASN Gateway with 3GPP Overlay



ASN Gateway Interoperability with 3GPP2 Overlay

The following figure shows a typical interoperability scenario between WiMAX and 3GPP2 legacy networks with reference points and interfaces.

Figure 58. ASN Gateway with 3GPP2 Overlay



Session Continuity Support for 3GPP2 and WiMAX Handovers

This feature provides seamless 3GPP2 session mobility for WiMAX subscribers and other access technology subscribers. With the implementation of this feature, the HA can be configured for:

- 3GPP2 HA service
- 3GPP HA service
- WiMAX HA service
- A combination of 3GPP2 and WiMAX HA services

The above configurations provide the session continuity capability that enables a dual-mode device (a multi-radio device) to continue its active data session as it changes its active network attachment from 3GPP2 to Wimax and vice versa, with no perceived impact from a user perspective. This capability brings the following benefits:

- Common billing and customer care
- Accessing home 3GPP2 service through Wimax network and vice versa
- Better user experience with seamless session continuity

For more information on this support, refer to the HA Administration Guide.

NSP-ID and NAP-ID Functionality

NSP-ID and NAP-ID functionality enable the MS to discover all accessible network service providers (NSPs) in a WiMAX coverage area, and to indicate the NSP selection during connectivity to the ASN. The actual NSP selection by the MS may be based on various preference criteria, depending on the configuration information.

Configuration information includes:

- Information useful in the MS discovery of the network access provider (NAP), including channel center frequency and PHY profile,
- Information useful in the MS decision mechanism to prioritize NSPs for automatic service selection, including a list of authorized NAPs and NSPs,
- A list of authorized share or roaming relationships between authorized NAPs and NSPs and partner NAPs and NSPs,
- Identity credentials provided by the NSPs to which the MS has a business relationship, and
- The mapping relation table between 24-bit NSP identities and corresponding realms of the NSPs.

Configuration information may be provided on a pre-provisioned basis or at the time of MS dynamic service subscription.

Manual Mode

The NSP Enumeration List is presented to the user for selection. Each entry presents only the verbose NSP name to the user. If more than one NAP can be used to establish a direct connection with a NSP, the MS may indicate each of the candidate NAPs along with the NSP or verbose NSP name to the user in the following order:

- Home NSP.
- If the "User Controlled RAPL" (Roaming Contractual Agreements Preference List) is available, NSPs or their corresponding verbose NSP names in the "User Controlled RAPL" in the MS (in priority order),
- If the "Operator Controlled RAPL" is available, NSPs or their corresponding verbose NSP names in the "Operator Controlled RAPL" in the MS (in priority order),
- Any other NSP or their corresponding verbose NSP names in random order.

Upon selection and successful authentication to the selected NSP, the MS indicates the selected NSP. If no NSP is found, the MS behavior is implementation dependent.

Automatic Mode

For Automatic Mode, without user intervention, the MS selects a NAP that has a direct connection to the home NSP. If more than one NAP can be used to establish a direct connection with an NSP, the MS selects a NAP by using "User Controlled CAPL" (Contractual Agreements Preference List) or "Operator Controlled CAPL".

If a NAP that has direct connection to the home NSP is not found, the MS attempts to select a NAP that has connection to one of the NSPs in the Preferred NSPs list. The order that the MS follows for selection from the NSP Enumeration List is determined by the "User Controlled CAPL" and "Operator Controlled CAPL", if available in the configuration information. The MS selects and tries to authenticate with an available and allowable NSP, in the following precedence:

- Home NSP,

- If the "User Controlled RAPL" is available, NSPs in the "User Controlled RAPL" in the MS (in priority order),
- If the "Operator Controlled RAPL" is available, NSPs in the "Operator Controlled RAPL" in the MS (in priority order),
- Any other NSP in random order.

Upon selection and successful authentication to the selected NSP, the MS indicates the selected NSP. If no NSP is found, the MS behavior is implementation dependent.

ASN GW and NAP-ID/NSP-ID Process

Following is an overview of NAP-ID and NSP-ID process from the ASN GW's perspective:

1. NAP/NSP advertisement: BS advertises the available NAP/NSP to MS. The MS chooses one of the preferred NAP/NSPs and performs INE with that NAP/NSP. The BS/MS supports this function; the ASN GW does not play a role.
2. Once the MS selects an NSP, the MS performs INE with the selected NSP to authenticate the MS at the selected NSP. The ASN GW enables the MS to get service from that NSP. In short, the ASN GW routes the Authentication request to the right AAA server, based on the NAI.
3. The NSP-ID is included in the access request from the ASN GW.
4. The ASN GW and HA sends the NSPID in authentication and accounting procedures to AAA server. The ASNGW does not send NAPID in authentication and accounting procedures to the AAA server, since the ASNGW sends the BSID to the AAA server.

Data Tunnel Endpoint Support

ASNGW supports forwarding downlink data to different endpoint tunnels other than the control address on the BS. In addition, the ASNGW can process uplink data and control traffic on different paths (VLANs) on the ASNGW if the data and control traffic are hosted on a different IP address.

The control and data tunnel endpoint can be different for the BS or the ASNGW. If the data tunnel endpoint is different from the control endpoint, it applies for both downlink and uplink traffic destined to, or received from, the peer (BS or ASNGW). This means that when downlink traffic is sent to the peer, the data tunnel endpoint is the destination IP address; when the uplink traffic is received from this peer, the source IP address is the data tunnel endpoint. The GRE key is unique for downlink and uplink data paths.

ASNGW with a Different Tunnel Endpoint

The ASNGW supports different data tunnel points on the BS for downlink traffic. The BS conveys its data tunnel endpoint through the existing R6 TLVs within MS Info.

If the uplink has a different data tunnel point, the ASNGW adds this information in the message that carries MS information, including the TLV or data path TLV that describes uplink service flow. There is a unique GRE key assigned to the uplink data path.

The ASNGW includes the tunnel endpoint TLV in the data path messages to BS or from the non-anchor GW to the anchor GW and vice-versa, to support the handoff functionality. After receiving the tunnel endpoint TLV within the data path messages, the BS forwards all the uplink data traffic to the same address.

No Handoff (INE)

For the control and data path setup for the INE, the BS/ASN specifies the different IP addresses for control and data traffic. For example, DT1 and BS1 are the data and control endpoints on the BS side. AT1 and AS1 are the data and control endpoints respectively on the ASNGW side. If the ASNGW requires different data tunnel endpoints instead of control addresses, the tunnel endpoint IP address has to be populated in the MS information TLV if it is per BS for DP Reg Request/Response message.

The ASNGW creates an NPU flow using the tunnel IP address if received in the DP Reg Response message for the uplink packet. From now on, all the data packets received from the BS have the source address of DT1. For all the downlink traffic, the ASNGW forms the data packet using the ASNGW Data IP address as the source address; the AS1 and destination address of the tunnel are the tunnel endpoint (that is, DT1). If no tunnel end point attribute is received from BS/ASNGW, by default, the control IP address is used for data traffic.

If the ASNGW requires a different data tunnel endpoint instead of a control address, the tunnel endpoint IP address is populated in the MS information TLV if it is per BS for DP Reg Request/Response message.

Inter-ASNGW Handoff

For control and data path setup for an inter-ASNGW handoff, the serving base station (SBS) is connected to the anchor GW and the target base station (TBS) is connected to the non-anchor GW. During INE, the anchor GW specifies a different IP address for data traffic. For example, if AT1 and DT1 are the data endpoints on the ASNGW and SBS side, the ASNGW informs the SBS about a different address for data by sending out a tunnel endpoint IP address in the MS information TLV in the DP-Reg Request during INE.

Similarly, during inter-ASNGW handoff, the non-anchor GW specifies the different data tunnel endpoint for uplink traffic in the DP-Reg Rsp message. The same address on the non-anchor GW receives the downlink data packets from the anchor GW. AT2 and DT2 are the data tunnel endpoint for the non-anchor GW and TBS, respectively. The tunnel endpoint IP address is populated in the MS information TLV in DP Reg Rsp message from non-anchor GW to the TBS.

Intra-ASNGW Handoff

For intra-ASNGW handoff, during INE, the ASNGW specifies the different data tunnel endpoint for receiving uplink traffic in the DP-Reg Req message to the serving base station (SBS). After handoff, the ASNGW specifies the different data tunnel endpoint to receive the Uplink traffic in the DP-Reg Rsp message. For example, AT1 is the tunnel endpoint on ASNGW for data traffic. DT1 and DT2 are the data tunnel endpoints on the SBS and TBS, respectively. The ASNGW provides the tunnel endpoint in the MS information TLV within the data path messages

AS1 is the control address on the ASNGW to negotiate R6 control traffic. SB1 and TB1 are the control addresses on SBS and TBS, respectively. The ANCHOR GW specifies the tunnel endpoint to receive the uplink traffic in the DP-Reg Rsp message. AT1 and AT2 are the data tunnel endpoints on the anchor and non-anchor GWs, respectively to negotiate R6 control traffic. SB1 and TB1 is the control address on SBS and TBS, respectively.

Supported Standards

WiMAX/IEEE References

- "WiMAX ASN Profiles", WiMAX Forum
- "Initial Network Entry Stage 3", Draft T33-001-R016v01-G Specification WiMAX Forum
- "Procedures and Messages for ASN Anchored Mobility with Profile C: Stage 3" draft, WiMAX Forum
- "Procedures for CSN Anchored Mobility Stage 3" draft, WiMAX Forum
- "WiMAX End-to-End Network Systems Architecture: Stage 2 Draft Specification", Release 1.0.0 Draft, March 28, 2007, WiMAX Forum
- "WiMAX End-to-End Network Systems Architecture: Stage 3: Detailed Protocols and Procedures", Release 1.0.0 Draft, March 28, 2007, WiMAX Forum
- "WiMAX Forum Network Architecture - Stage 3-Detailed Protocols and Procedures" - DRAFT-T33-001-R015v01-J

IEEE Standards

- IEEE 802.16e/D12 September 2005, Local and Metropolitan Area Networks – Part 16: Air Interface for Fixed Broadband Wireless Access Systems, Feb 2006.
- 802.1P QoS at MAC Level
- 802.1Q VLAN Standard
- IEEE 802.16e/D12 September 2005, Local and Metropolitan Area Networks - Part 16: Air Interface for Fixed Broadband Wireless Access Systems, March 2004.

IETF References

- RFC-1701, Generic Routing Encapsulation (GRE), October 1994
- RFC-2131, Dynamic Host Configuration Protocol (DHCP), March 1997
- RFC-2794, Mobile NAI Extension
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-3012, Mobile Ipv4 Challenge/Response Extensions, November 2000
- RFC-3024, Reverse Tunneling for Mobile IP, revised, January 2001
- RFC-3046, DHCP Relay Agent Information Option, January 2001
- RFC-3344, Mobile IP support for Ipv4, August 2002
- RFC-3579, RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP), September 2003
- RFC-3588, Diameter Base Protocol, September 2003

- RFC-3748, Extensible Authentication Protocol, June 2004
- RFC 1918, NWG, Stage 2 Architecture, 121505
- RFC 3115, Mobile IP Vendor/Organization-specific Extensions

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

Chapter 9
ASN Paging Controller and Location Registry Overview

The ASN Paging Controller and Location Registry (PC/LR) provides paging and location updates to WiMAX subscribers in IEEE 802.16 Mobile WiMAX radio access networks. This service can be used as a standalone product or in combination with ASN Gateway as co-located services on the same chassis.

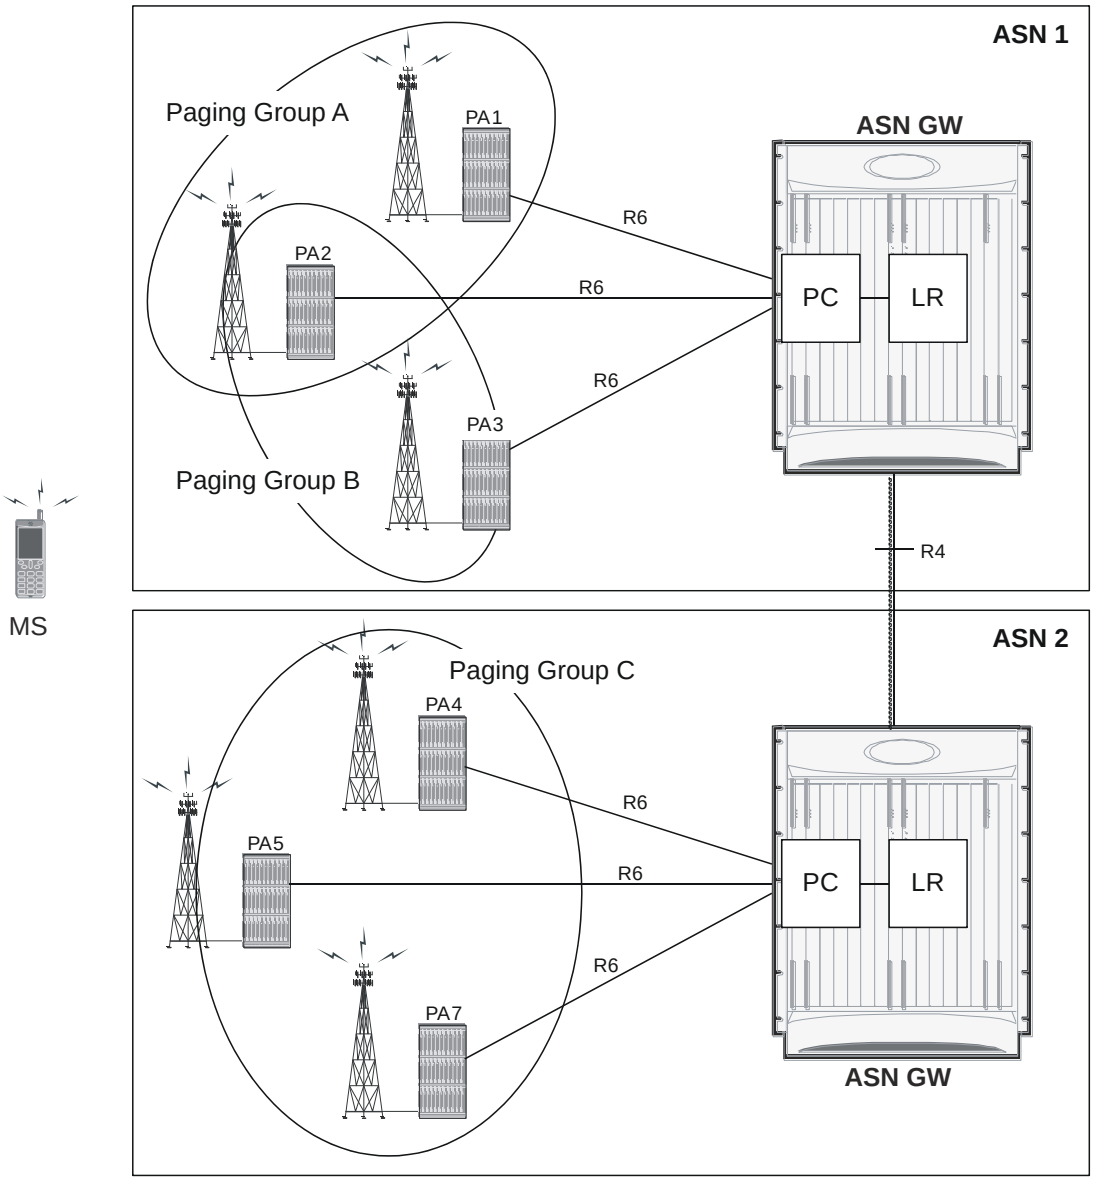
Introduction

ASN Paging Controller and Location Registry (PC/LR) supports connection management and mobility across cell sites and inter-service provider network boundaries by processing subscriber control and bearer data traffic.

Each ASN Gateway can concentrate traffic from many radio base stations. This reduces the required number of devices under management and minimizes connection set-up latency by decreasing the number of call hand-offs in the network.

Paging and Idle Mode Operation maintains a track and alert for MSs when they are in idle mode to save battery power. Paging is executed to alert MSs when there is an incoming message. Figure 8 illustrates the paging operation and paging and idle mode elements in the WiMAX network system.

Figure 59. ASN Paging Controller and Location Registry in WiMAX Networks



In WiMAX networks, a mobile station is tracked when it is in idle mode. The information is stored to a location register (LR). The tracking area is larger than the cell size because a paging group (PG) comprises multiple cells. When a mobile station moves across paging groups, its location is updated via R6 and/or R4. The paging controller (PG) in ASN-GW retrieves the location from the LR and alerts the paging agent in (PA) in the base station to signal to the mobile station.

Location information for idle mode subscribers is maintained in a location register central database that is co-located on an anchor paging controller. Idle mode can be initiated by the mobile device or the network. The paging controller retains subscriber session context information in addition to supervising paging activities. It also represents an authentication liaison between the user device and the AAA server. As the subscriber roams across cell sites, it is associated with a group of base stations known as a paging group. Location updates to the LR database are conveyed over R6 and R4 messages between the relay paging controller serving ASN and the A-PC/LR. When a remote host attempts to reach an idle mode subscriber device, the anchor paging controller alerts the paging group members when it receives downlink traffic by requesting the paging agent in the base station to signal the idle mode subscriber.

Description of PC/LR Support

The PC/LR runs as a stand-alone function in a separate chassis or as an integrated service on same chassis as the Anchor Authenticator (A-PC)/Anchor Datapath (A-DP) ASN Gateway. The idle mode LR database uses distributed software architecture and provides an LR manager task that partitions smaller database volumes across separately running session manager tasks in the system. The implementation is based on a topologically unaware paging scheme in which the A-PC does not have global awareness of all member base stations in a paging group. The A-PC uses a single-step paging operation where paging notifications are sent to the last-reported serving paging controller or directly attached base station.

Idle mode operation is very important in order for any cellular system to keep the mobile device reachable when it is inactive. It enables mobility in addition to conserving battery life. Idle mode paging also eliminates the requirements of independent VLRs/HLRs, when it is supported as an integrated function in the ASN Gateway system.

Licenses

The ASN PC/LR service is a separate product from the ASN Gateway. You must purchase the WiMAX Paging Controller/Location Register product license separately to enable this service.

Paging and Location Update Procedures

This section provides an overview of the ASN Gateway's paging and location update procedures.

The system provides following components for the paging controller, paging group and location registry functionality.

Paging Controller (PC)

The paging controller is a functional entity that administers the activity of idle mode mobile stations in the network. It is identified by PC ID, which maps to the address of a functional entity in a WiMAX network. In this implementation, the PC is co-located with ASN Gateway. There are two types of PCs:

- **Anchor PC**: For each idle mode MS, there is a single anchor PC that contains the updated location information of the MS.
- **Relay PC**: There are one or more other PCs in the network, called relay PCs, that participate in relaying paging and location management messages between the paging agent and the anchor PC.

Paging Agent (PA)

The paging agent is a functional entity, implemented in an ASN base station, that handles the interaction between PC- and paging-related functionality.

Paging Group (PG)

A paging group is a logical entity comprising one or more paging agents. A paging group resides entirely within a NAP boundary. Paging groups are managed by the network management system and provisioned per the access network operator's provisioning requirements.

Location Register (LR)

A location register is a distributed database, with each instance corresponding to an anchor PC. Location registers contain information about idle mode MSs. The information for each MS includes:

- MS paging information about each MS that has registered in the past in the network but is currently in idle mode
- Current paging group ID (PGID)
- PAGING_CYCLE
- PAGING_OFFSET
- Last reported BSID
- Last reported relay PCID
- MS service flow information
    - Idle mode retention information for each MS in idle mode
    - Information about the service flows associated with the MS

An instance of a location register is associated with every anchor PC.

Paging Controller and Location Update functionality supports following operation and procedures in ASN Gateway:

Location Update Procedure

Location Update with Paging Controller Relocation

Paging Operation

MS Initiated Idle Mode Entry

MS Initiated Idle Mode Exit

Location Update Procedure

This section describes the secure location update procedure for a WiMAX MS.

The following figure and table provides a high-level view of the steps involved in a secure location update.

Figure 60. Location Update Flow

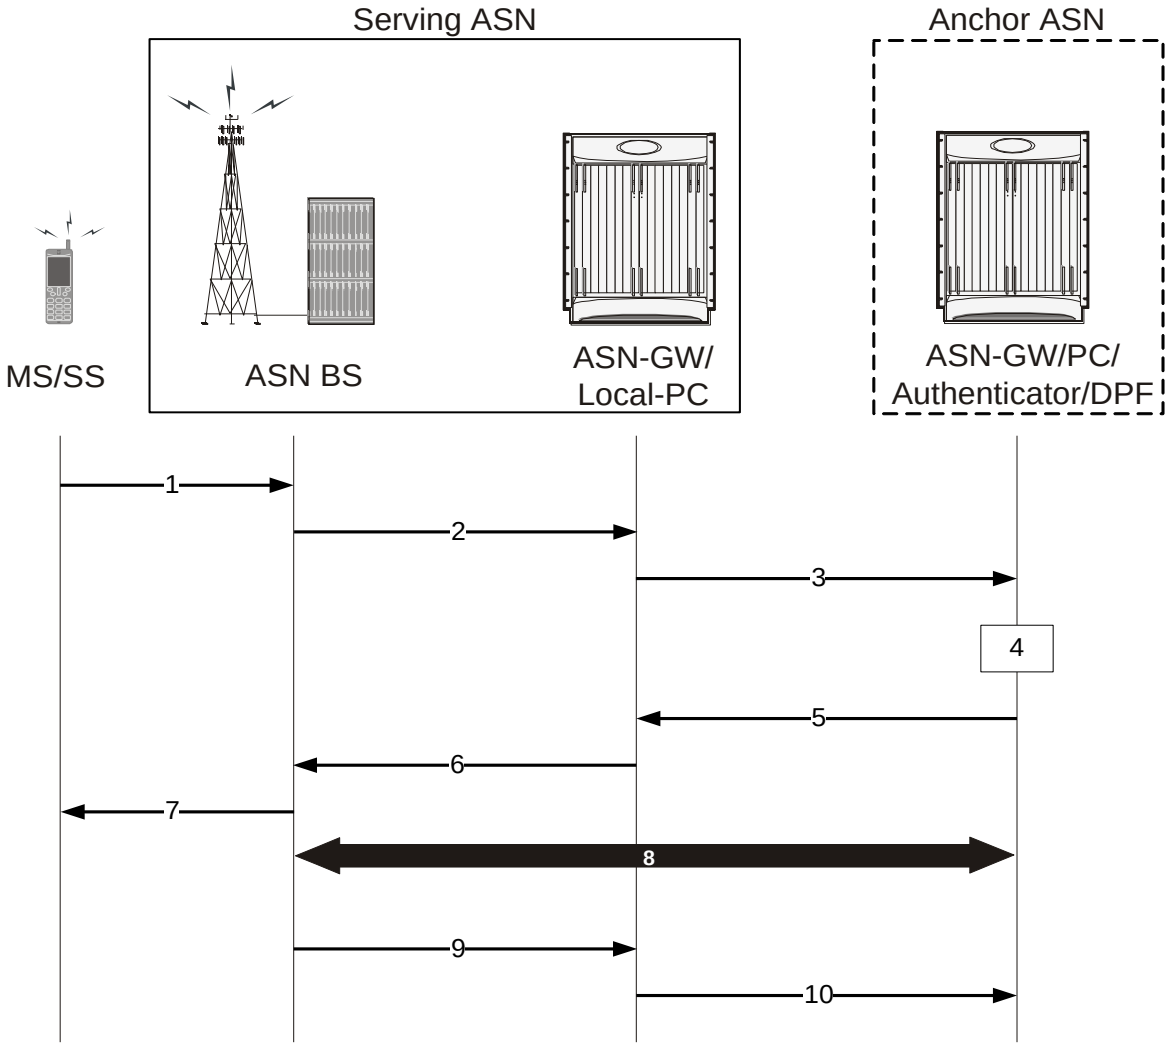


Table 46. Location Update Procedure Flow Description

Step	Description
1	The MS initiates a secure Location Update procedure by sending a RNG-REQ message to serving ASN BS, which includes the Ranging Purpose Indication TLV set to indicate Idle Mode Location Update, the PC ID TLV which points to the anchor ASN Gateway acting as the anchor PC function for the MS, and the HMAC/CMAC tuple.
2	The serving ASN BS sends an R6 LU_Req message to the serving ASN Gateway and starts timer TR6_LU_Req. The message may include the PG ID, Paging Offset, and Paging Cycle TLVs if the serving ASN BS proposes an update to these parameters.
3	The serving ASN Gateway (associated with the local Paging Controller) sends an R4 LU_Req message to the anchor PC (associated with Anchor ASN Gateway) and starts timer TR6_LU_Req. The message may include the PG ID, Paging Offset, and Paging Cycle TLVs if the serving ASN Gateway proposes an update to these parameters. Note: This message may be relayed by several intermittent ASNs before reaching the anchor PC (Anchor ASN Gateway).

Step	Description
4	If the anchor PC retains context information for the MS including its Authenticator ID, the anchor PC initiates a Context Request procedure with the anchor authenticator/ASN Gateway. If the anchor authenticator/ASN Gateway has valid key material for the MS, it returns AK context for the MS to the Anchor PC.
5	Upon successful retrieval of the AK context, the anchor PC sends an R4 LU_Rsp message back to the serving ASN Gateway and starts timer TR4_LU_Conf. The message includes the MSID, BSID, Authenticator ID, assigned PGID, Paging Offset, Paging Cycle, Anchor PC ID TLVs, and Location Update Status TLV set to Accept. Upon receipt of the R4 LU_Rsp message, the serving ASN Gateway stops timer TR4_LU_Req.
6	Upon receipt of the R4 LU_Rsp message, the serving ASN Gateway stops timer TR4_LU_Req, sends an R6 LU_Rsp message to the serving ASN BS, and starts timer TR6_LU_Conf. The message includes the Location Update Status TLV set to Accept, AK Context TLVs, as well as the assigned Paging Information TLV if they were included in the corresponding R4 message.
7	Based on the AK and AK context received from the anchor PC, the serving BS (associated with Local PC/Relay PC in serving ASN Gateway) successfully authenticates the RNG_REQ message received from the MS and sends a RNG_RSP message with HMAC/CMAC and Successful LU_Rsp indication to the MS.
8	The serving ASN BS initiates an R6 CMAC Key Count Update procedure with the ASN Gateway. The serving ASN Gateway initiates an R4 CMAC Key Count Update procedure with the authenticator ASN to update it with the latest CMAC Key Count.
9	The serving ASN BS sends an R6 LU_Cnf message to the serving ASN Gateway with Location Update TLV indicating success. Upon receipt of the message, the serving ASN Gateway stops timer TR6_LU_Conf.
10	The serving ASN Gateway sends an R4 LU_Cnf message with a successful LU indication to the anchor PC and stops timer TR6_LU_Req. Upon receipt of the message, the anchor PC updates the LR with MS Idle Mode information and stops timer TR4_LU_Conf.

Location Update with Paging Controller Relocation

This section describes the secure location update with PC relocation procedure for a WiMAX MS.

The following figure and table provides a high-level view of the steps involved in a secure location update with PC relocation.

Table 47. Location Update with PC Relocation - Procedure Flow

Step	Description
1	The MS initiates a secure Location Update procedure by sending a RNG-REQ message to the serving ASN BS, which includes the Ranging Purpose Indication TLV set to indicate Idle Mode Location Update, the PC ID TLV which points to the anchor ASN Gateway acting as the anchor PC function for the MS, and the HMAC/CMAC tuple.
2	The serving BS sends an R6 LU_Req message to the serving ASN Gateway and starts timer TR6_LU_Req. The message may include the PG ID, Paging Offset, and Paging Cycle TLVs if the serving BS proposes an update to these parameters.
3	The serving ASN Gateway (associated with the serving BS and local PC) sends an R4 LU_Req message to the anchor PC ASN associated and starts timer TR4_LU_Req. The message may include the PG ID, Paging Offset, and Paging Cycle TLVs if the serving ASN proposes an update to these parameters. Note that this message may be relayed by several intermittent ASNs before reaching the current anchor PC ASN. The serving ASN or any intermittent ASN along the path may request PC relocation.

Step	Description
4	Upon receipt of the R4 LU_Req message, a relay PC ASN adds the anchor PC Relocation Destination TLV to initiate PC relocation. The message is forwarded to the anchor PC ASN. The new anchor PC ASN starts timer TR4_LU_Request.
5	If the current anchor PC ASN retains context information for the MS, including its authenticator ID, the current anchor PC ASN initiates a Context Request procedure with the anchor authenticator ASN. If the anchor authenticator ASN has valid key material for the MS, it returns AK context for the MS to the anchor PC ASN.
6	The current anchor PC ASN sends an R4 LU_Rsp message back to the new anchor PC ASN and starts timer TR4_LU_Conf. The message includes the MSID, BSID, authenticator ID, assigned PGID, Paging Offset, Paging Cycle, Anchor PC ID TLVs, and Location Update Status TLV set to Accept. The anchor PC Relocation Request Response TLV is set to Accept to indicate that the current anchor PC ASN accepted the PC_Relocation_Req and the anchor PC ID TLV is set to the identifier of the new anchor PC ASN ID which was received in the anchor PC Relocation Destination TLV in the R4 LU_Req message. The R4 LU_Rsp message also includes MS Info TLV containing the MS context for transfer to the new anchor PC ASN. If the new anchor PC ASN does not request PC Relocation, the current anchor PC MAY still request to perform the procedure by including the PC Relocation Indication TLV. If the new anchor PC does not accept the relocation, it reports a failure in step 17.
7	Upon receipt of the R4 LU_Rsp message from current anchor PC ASN, the new anchor PC ASN stops timer TR4_LU_Req, stores the MS context received from current anchor PC ASN, updates the paging information (Paging Group ID, Paging Cycle, Paging Offset), forwards the R4 LU_Rsp message on to the serving ASN, and starts timer TR4_LU_Conf.
8	Upon receipt of the R4 LU_Rsp message, the serving ASN-GW stops timer TR4_LU_Req, sends an R6 LU_Rsp message to the S-BS, and starts timer TR6_LU_Conf. The message includes the Location Update Status TLV set to Accept, MS Info, AK Context, anchor PC ID, and the old anchor PC ID TLV. The message may include the paging information TLVs if they were included in the corresponding R4 message.
9	Based on the AK and AK context received from the current anchor PC, the serving BS (associated with local PC/relay PC) successfully authenticates the RNG_REQ message received from the MS. The serving BS sends a RNG_RSP message with HMAC/CMAC and Successful Location Update Response indication to the MS.
10	The serving BS sends an R6 LU_Cnf message to the serving ASN-GW with Location Update TLV indicating success. Upon receipt of the message, the serving ASN-GW stops timer TR6_LU_Conf.
11	The serving ASN sends an R4 LU_Cnf message with a successful LU indication to new anchor PC ASN (as indicated by the anchor PC ID received from the BS) and stops timer TR6_LU_Req. Alternatively, the relay PC ASN forwards LU_Cnf to the ASN associated with new anchor PC with the result indication reassigned by the relay PC. Upon receipt of the message, new anchor PC ASN stops timer TR4_LU_Conf.
12	Upon receipt of the LU_Cnf message, the new anchor PC ASN sends an R4 PC_Relocation_Ind to the anchor DP/FA ASN, and starts timer TR4_PC_Reloc_Upd_ADP.
13	The anchor DP/FA ASN updates the anchor PC for the MS with the new anchor PC ASN ID and responds with an R4 PC_Relocation_Ack message confirming the anchor PC update. Upon receipt of the message, the new anchor PC ASN stops timer TR4_PC_Reloc_Upd_ADP. The new anchor PC ASN hosts the anchor PC function and becomes the new current anchor PC ASN for the MS. The anchor PC is de-allocated from the old current anchor PC ASN.
14	Simultaneous with sending PC_Relocation_Ind to the anchor DP/FA, the new anchor PC sends an R4 PC Relocation Indication to the anchor authenticator ASN to inform the change of the anchor PC, and starts timer TR4 - PC_Reloc_Upd_AA.
15	The anchor authenticator ASN updates the anchor PC for the MS with the new anchor PC ASN ID and responds with an R4 PC_Relocation_Ack message confirming the anchor PC update. Upon receipt of the message, the new anchor PC ASN stops timer TR4-PC_Reloc_Upd_AA. At this point, new anchor PC ASN hosts the anchor PC function and becomes the new current Anchor PC ASN for the MS. The anchor PC is de-allocated from the old current anchor PC ASN.

Step	Description
16	The new anchor PC ASN sends an R4 LU_Cnf message with a successful LU indication to the current anchor PC ASN and stops timer TR4_LU_Conf. The old current anchor PC ASN clears its LR context for the MS.
17	This step is optional. If the anchor PC ASN receives CMAC Key Count TLV update in LU_Cnf message, it should perform an R4 CMAC Key Count Update procedure with the authenticator ASN to update it with the latest CMAC Key Count. Refer to section 4.13 for the call flow.

Paging Operation

This section describes the paging operation for a WiMAX MS.

The following figure and table provides a high-level view of the steps involved in the paging operation call flow of an MS.

Figure 61. Paging Operation Procedure Flow

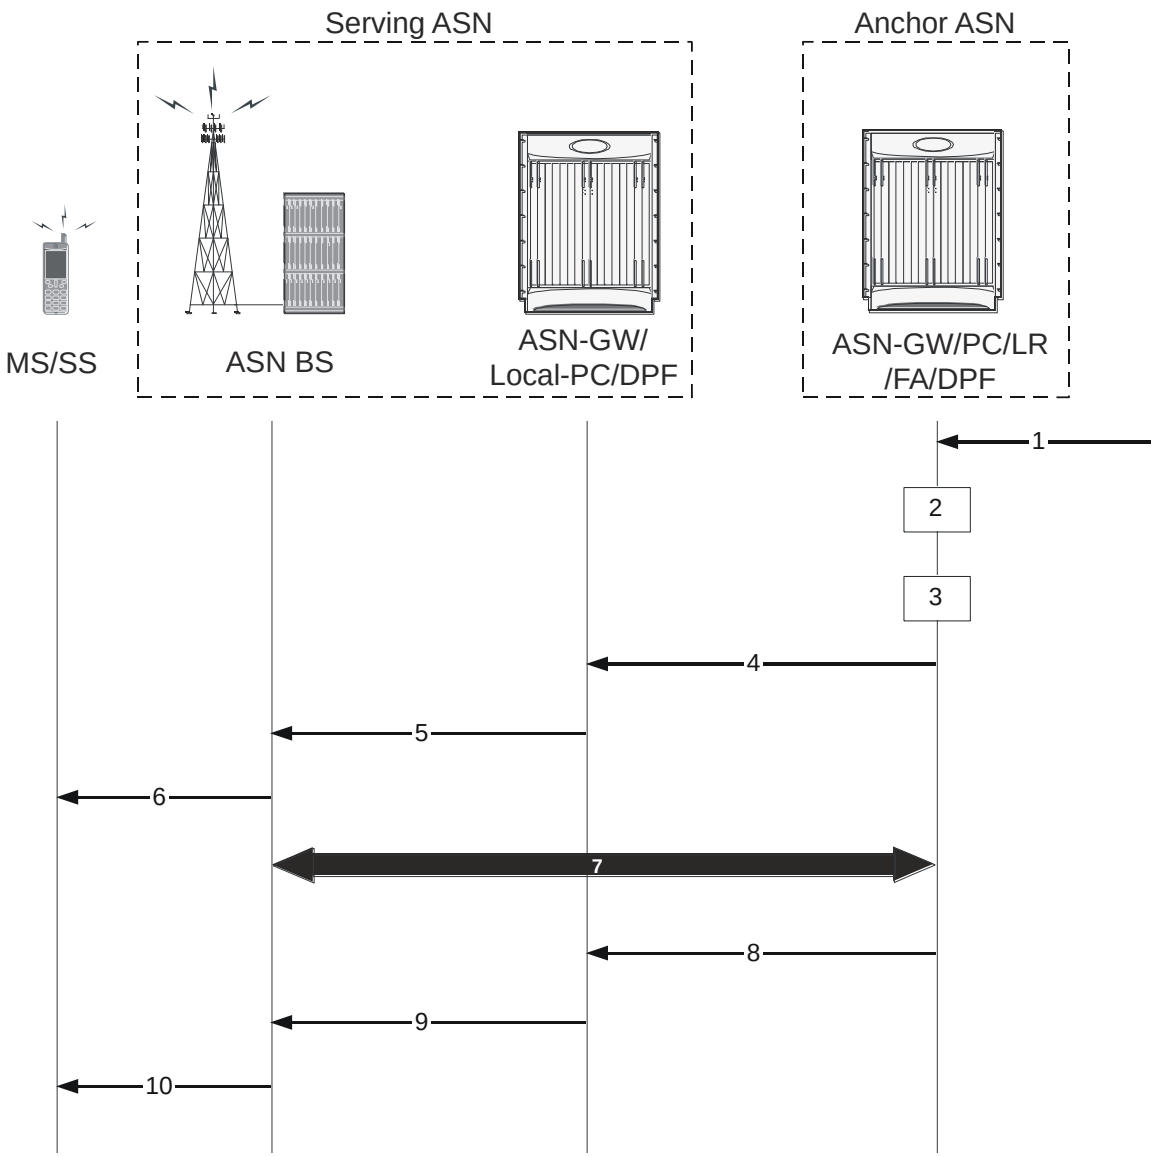


Table 48. Paging Operation Procedure Flow Description

Step	Description
1	Data from HA arrives through the tunnel at the FA and its associated DPF. The anchor DPF buffers the data.
2	Anchor Data Path Function (DPF) sends an R4 Initiate_Paging_Req message to the anchor PC/LR to request paging. Optionally the R4 Initiate_Paging_Req message contains the QoS parameters of the flow for which the data arrived at the anchor DPF. This helps set priority treatment of the paging operation based on the QoS parameters and flow types. The anchor DPF may have policies for triggering paging based on the QoS parameters for the data received. The anchor DP Function starts timer TInit_Page_Req.Note: When MS is in Idle Mode, if data not belonging to any saved Service Flow (SF) of the MS arrives, the decision to initiate paging or not is on the basis of operator's setting.

Step	Description
3	The anchor PC/LR retrieves the information related to the MS and sends an R4 Initiate_Paging_Rsp to Anchor Data Path function. This message indicates whether the MS context as contained in the PC/LR is correct and the requested paging action is authorized. Exclusion of the Response Code TLV indicates intent to page the MS. Upon receipt of this message the anchor DP Function starts timer TInit_Page_Req if running.
4	If paging action is authorized, the anchor PC retrieves the MS paging information and constructs Paging_Announce message. The anchor PC issues one or more Paging_Announce messages based on its knowledge of the paging region topology. The anchor PC starts a timer TR4_Paging_Announce when it sends out the first Paging_Announce message and waits for the paging response. The anchor PC sets a paging re-transmission counter *N*. If the anchor PC does not receive a paging response, it retransmits the Paging_Announce message prior to the expiration of the timer TR4_Paging_Announce. If the anchor PC is topologically aware of the defined Paging Group (PG), including the last BS from which the MS performed location update, the anchor PC directly issues Paging_Announce messages to all or some subset of the paging group members. The members consist of BSs and/or relay PCs in the region. If the anchor PC is topologically unaware of the paging region or the BSs defined in the paging group, the Paging_Announce messages are sent to the known relay PC(s). The relay PC(s) forwards the announce message to one or more BSs in the paging region.
5	The ASN Gateway that contains the local/relay PC function for the MS initiates the paging operation and sends the R6 Paging_Announce message to the BS(s) associated with the Paging Group ID (PGID) received in R4 Paging_Announce. The ASN Gateway performs single- or multi-step paging based on whether the BS ID TLV or the L-BSID TLV is present. Associated with each R4 Paging_Announce message, the ASN Gateway starts timer TR6_Paging_Announce.
6	Once the Paging Agent (PA) at the BS receives the Paging_Announce message with the requested action set to Start, it extracts the relevant paging parameters for the MS (Paging Cycle, Paging Offset). It then initiates the paging action requested by sending out MOB-PAG_ADV message over the airlink as per the indicated paging cycle and the paging offset. The optional SF Flow info in the message helps the BS implement a paging priority scheme for faster call setup when bandwidth is constrained or for resource allocation. The PA continues to page the MS for the duration specified by the Paging Announce Timer TLV, until the appropriate response is received from the MS, or a stop page indication is received from the Local PC.
7	Upon being successfully paged the MS performs a Idle Mode Exit or a location update procedure. If any Paging Agent (PA) receives a successful reply from the paged MS, the Paging Agent notifies the Local PC by sending an R6 LU_Req message, or an R6 IM_Exit_State_Change_Req message, in the case of data delivery to MS in idle mode. Upon receipt of a such a message the Local PC stops timer TR6_Paging_Announce if running, and sends the appropriate R4 LU_Req or R4 IM_Exit_State_Change_Req message to the anchor PC. Upon receipt of such a message, the anchor PC stops timer TR4_Paging_Announce, if running. The anchor PC also initiate stop paging procedures as described at step 8 and onward.
8	Upon receipt of a response from the MS as mentioned at step 7, and anchor PC wants to initiate stop paging procedure, the anchor PC sends a R4 Paging_Announce message to all BSs in the PG. The R4 Paging_Announce message has the Paging Start/Stop TLV set to 0.
9	The local PC sends a R6 Paging_Announce message to the BSs. The R6 Paging_Announce message has the Paging Start/Stop TLV set to 0.
10	Upon receipt of the R6 Paging_Announce message with Paging Start/Stop = 0, the BS terminate/cease a MOB_PAG-ADV messages over the air.

MS Initiated Idle Mode Entry

This section describes the MS-initiated idle mode entry procedure for a WiMAX subscriber.

The following figure and table provides a high-level view of the steps involved in MS-initiated idle mode entry call flow of an SS/MS.

Figure 62. MS Initiated Idle Mode Entry Procedure Flow

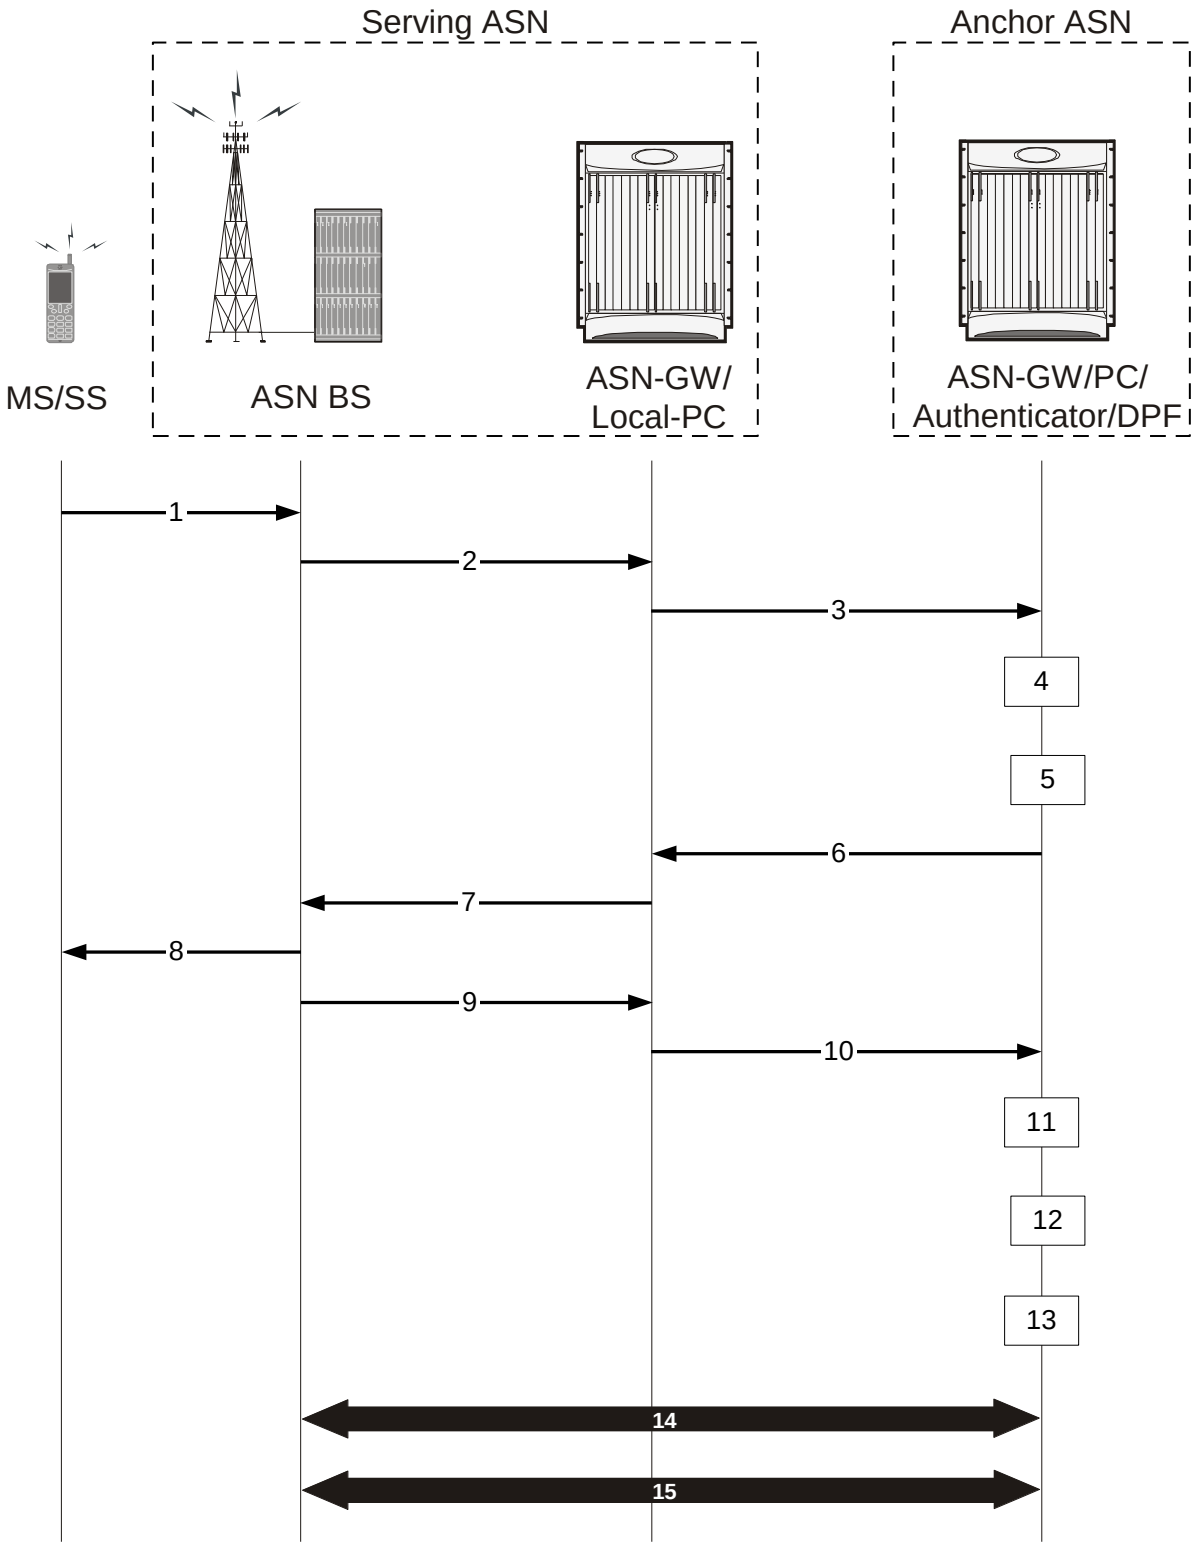


Table 49. MS Initiated Idle Mode Entry Procedure Flow Description

Step	Description
1	MS decides to enter Idle Mode and sends DREG_REQ formatted as described in IEEE 802.16e.The De-Registration Request code is set to 0x01 indicating that the MS intends to enter Idle Mode.
2	Based on the MS's request, the serving ASN BS (Paging Agent) in serving ASN sends an R6 IM_Entry_State_Change_Req message to its ASN Gateway. Timer TR4_IM_Entry_Req is started to monitor R6 IM_Entry_State_Change_Rsp at the serving ASN BS(PA).
3	The local Relay PC in Serving ASN Gateway chooses an anchor PC for the MS and sends inter-ASN R4 IM_Entry_State_Change_Req message to the Anchor ASN associated with the chosen anchor PC. Timer TR4_IM_Entry_Req_ASN is started to monitor the R4 IM_Entry_State_Change_Rsp.
4	The anchor PC/LR, sends R4 IM_Entry_State_Change_Req to the anchor authenticator to verify whether the MS is allowed to go into Idle mode. Timer TR4_IM_Entry_Req_APC is started at this time to monitor the R4 IM_Entry_State_Change_Rsp from the anchor authenticator. This step is optional if the anchor authenticator and anchor PC/LR are collocated in the same ASN Gateway.
5	The anchor authenticator checks if the MS is allowed to enter Idle Mode and saves necessary information if allowed, then sends back R4 IM_Entry_State_Change_Rsp to the anchor PC/LR including MSID, IDLE mode authorization indication. If the anchor authenticator rejects the Idle mode entry request, the Idle Mode Authorization TLV contains the rejection code. When R4 IM_Entry_State_Change_Rsp for MS entering Idle Mode is send successfully, The anchor authenticator stores the anchor PC ID for this MS. Upon receipt of this message at the anchor PC, the TR4_IM_Entry Req_APC is stopped. This step is optional if the anchor authenticator and anchor PC/LR are collocated in the same ASN Gateway.
6	According to the reported information in R4 IM_Entry_State_Change_Rsp, based on the content of Idle mode authorization indication IE, the anchor PC updates the LR with current MS location information (PGID) and other parameters, and sends back R4 IM_Entry_State_Change_Rsp message to the serving ASN Gateway. When this message is received at the serving ASN Gateway, timer TR4_IM_Entry_Req_ASN is stopped.
7	Serving ASN Gateway forwards the R6 IM_Entry_State_Change_Rsp to serving BS (PA) including IDLE Mode authorization indication and accepted Paging parameters. Upon receipt of this message at the BS, timer TR6_IM_Entry_Req is stopped.
8	Serving ASN BS sends DREG_CMD to the MS. The DREG_CMD conveys PC ID field pointing to the anchor PC for the MS and allocated Idle mode parameters.
9	After sending the DREG_CMD to the MS, the serving ASN BS (PA) acknowledges the successful delivery of DREG_CMD to the local relay PC in the serving ASN Gateway by sending R6 IM_Entry_State_Change_Ack.
10, 11	The local relay PC in the serving ASN Gateway forwards the successful entry of MS to Idle mode to the anchor PC in the anchor ASN Gateway by sending an R4 IM_Entry_State_Change_Ack. Upon receipt of this message at the anchor PC, timer TR4_IM_Entry_Rsp is stopped.
12	The anchor ASN Gateway associated with the anchor PC/LR updates the information of MS into LR database and sends an Anchor PC Indication message to the anchor DPF/FA to reflect the success of MS entering Idle Mode. Timer TR4_APC_Ind is started at this time when the anchor PC Indication is sent, to monitor the response.
13	The anchor DPF/FA finally updates the information of MS, including the anchor PC ID of this MS, and acknowledges to the anchor PC/LR by an Anchor PC Ack message. When the Anchor PC Ack is received at the anchor ASN Gateway, timer TR4_APC_Ind is stopped.
14	After the expiration of the Management Resource Holding Timer (an 802.16e parameter), the serving BS initiates the related R6 data Path Dereg procedure by sending R6 Path_Dereg_Req to the anchor ASN Gateway.

Step	Description
15	The serving ASN Gateway completes the data path de-registration from its side and sends an R4 Path_Dereg_Ack to the anchor DPF/FA. Upon receipt of this message, the anchor ASN Gateway stops timer TPath_Dereg_Rsp_ADPFt and the serving BS (PA) updates the anchor authenticator with the CMAC Key count for the MS via the serving ASN Gateway as per the CMAC Key count update procedure. The anchor authenticator acknowledges the CMAC update for the MS. Optionally this procedure may be invoked anytime after step 11.

MS Initiated Idle Mode Exit

This section describes the MS-initiated idle mode exit procedure for a WiMAX subscriber.

The following figure and table provides a high-level view of the steps involved in MS- initiated idle mode exit call flow of an SS/MS.

Figure 63. MS Initiated Idle Mode Exit Procedure Flow

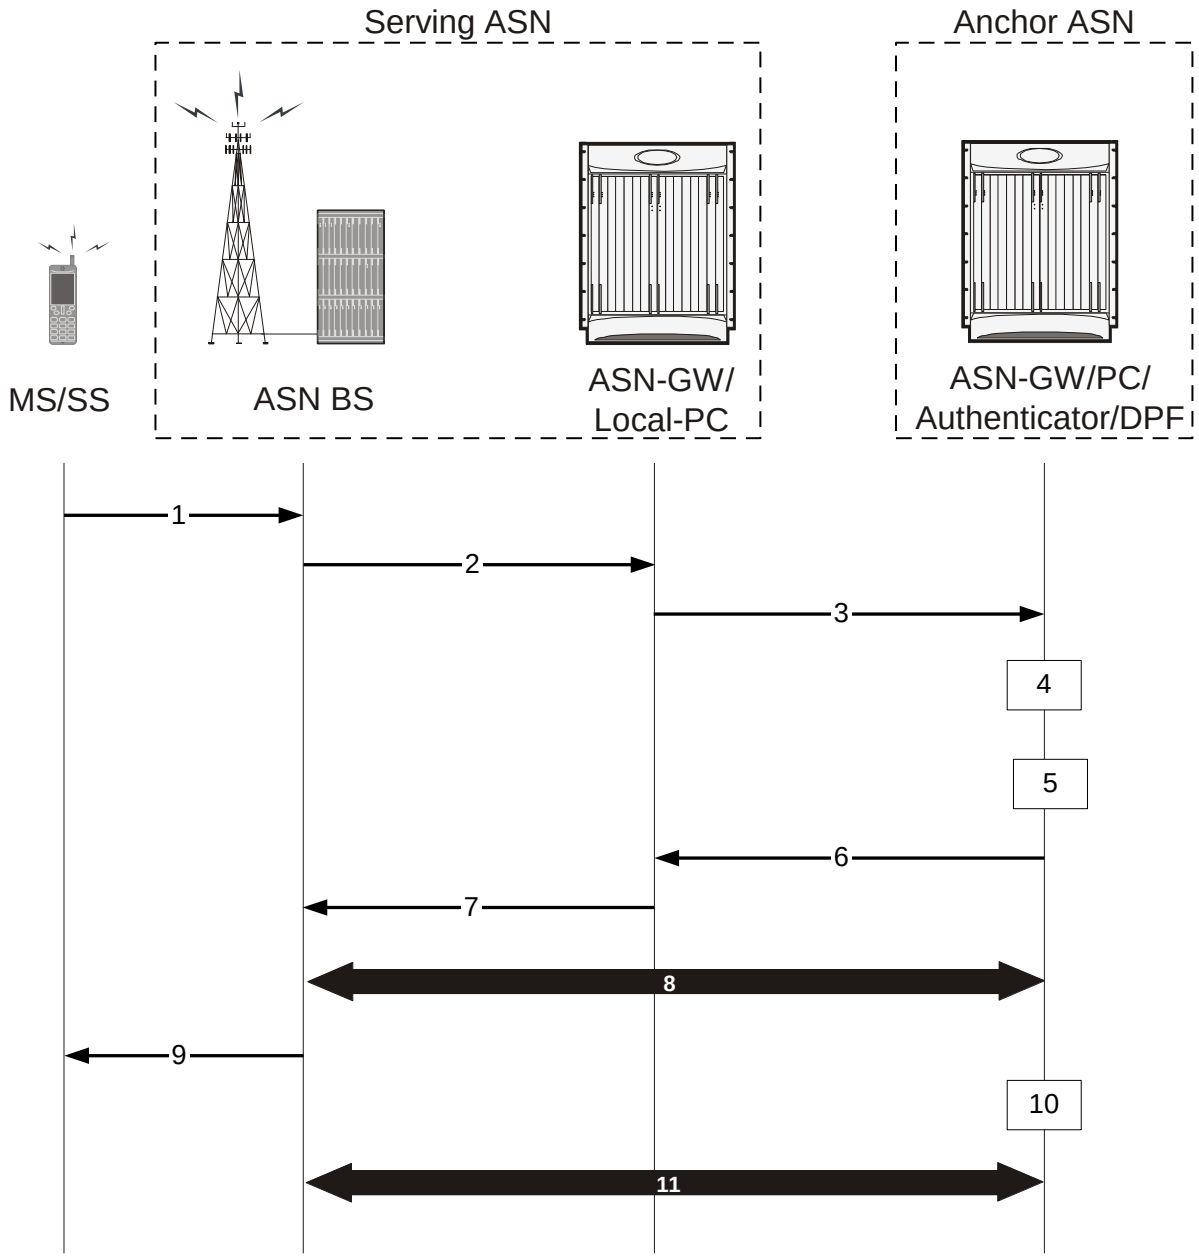


Table 50. MS Initiated Idle Mode Exit Procedure Flow Description

Step	Description
1	MS initiates exit procedure from IDLE mode and sends RNG_REQ to the serving ASN BS. The Ranging Purpose Indication TLV is set to 1 and the PC ID TLV is included, thus indicating that the MS intends to Re-Entry from Idle Mode.

Step	Description
2	The ASN BS receives the RNG_REQ message from MS indicating Idle mode exit and sends R6 IM_Exit_State_Change_Req to the relay PC in the ASN Gateway, indicating that the MS wants to become active. Timer TR6_IM_Exit_Ctx_Req is started at this point by the BS to monitor the response for this message.
3	The relay PC in the serving ASN Gateway receives the R6 IM_Exit_State_Change_Req from the BS indicating Idle mode exit and sends R4 IM_Exit_State_Change_Req to the anchor PC/LR in the anchor ASN Gateway, indicating that the MS wants to become active. Timer TR4_IM_Exit_Ctx_Req is started at this point by the anchor ASN Gateway to monitor the response for this message. In the event that the relay PC is the anchor PC, this step is not required.
4	On receiving the R4 IM_Exit_State_Change_Req, the anchor PC/LR proceeds to request the security context from the anchor authenticator in the anchor ASN Gateway using the R4 IM_Exit_State_Change_Req. Timer TR4_lMexit_ctx_req_PC is started at this point by the anchor PC to monitor the response for this message. This step is optional if the anchor authenticator and anchor PC/LR are co-located in the same ASN Gateway.
5	The anchor authenticator responds with the security context back to the anchor PC/LR with R4 IM_Exit_State_Change_Rsp message. Once the anchor PC receives this message, Timer TIM_Exit_Ctx_Req_PC is stopped. This step is optional if the anchor authenticator and the anchor PC/LR are collocated in the same ASN Gateway.
6	The anchor PC/LR, sends R4 IM_Exit_State_Change_Rsp to the relay PC. Once the relay PC receives this message, Timer TR4_IM_Exit_Ctx_Req is stopped.R4 IM_Exit_State_Change_Rsp contains the stored information for the MS at the anchor PC.
7	The serving ASN Gateway retrieves the MS context from anchor PC ASN and forwards the MS context to the serving BS on the R6 interface. Once the BS receives this message, timer TR6_IM_Exit_Ctx_Req is stopped. The AK fetched from the authenticator is used to verify the RNG-REQ.
8	After successful authentication, the BS starts data path establishment across the serving BS, serving ASN Gateway, relay PC, anchor PC, authenticator, and DPF.
9	The serving BS uses MS service and operational information indicated by IDLE Mode Retain Info obtained by step 7 to construct HO Process Optimization TLV settings in the RNG-RSP based on local policy, then sends RNG_RSP message to the MS formatted according to IEEE 802.16e specification. This message delivers all the required information to resume service in accordance with Idle Mode Retain information.
10	When R4 Path_Reg_Ack is received at the anchor DPF, the Data Path function associated with FA sends a Delete_MS_Entry_Req message to PC/LR in order to delete the Idle mode entry associated with the MS. If the MS is exiting Idle mode due to a network initiated Idle mode exit, the PC/LR will cease all Paging Announce operations.
11	The serving BS updates the anchor authenticator with the CMAC Key count for the MS via the serving ASN Gateway. The anchor authenticator acknowledges the CMAC update for the MS.

Supported Platforms and Software

ASN PC-LR is available for all chassis running StarOS Release 8.0 or later.

Chapter 10
CDMA2000 Wireless Data Services

The ASR 5x00 provides wireless carriers with a flexible solution that functions as a Packet Data Support Node (PDSN) in CDMA 2000 wireless data networks.

This overview provides general information about the PDSN including:

- Product Description
- Product Specifications
- Features and FunctionalityBase Software
- Features and Functionality - Optional Enhanced Software Features
- CDMA2000 Data Network Deployment Configurations
- Understanding Simple IP and Mobile IP
- Supported Standards

Product Description

The system provides wireless carriers with a flexible solution that can support both Simple IP and Mobile IP applications (independently or simultaneously) within a single scalable platform.

When supporting Simple IP data applications, the system is configured to perform the role of a Packet Data Serving Node (PDSN) within the carrier's 3G CDMA2000 data network. The PDSN terminates the mobile subscriber's Point-to-Point Protocol (PPP) session and then routes data to and from the Packet Data Network (PDN) on behalf of the subscriber. The PDN could consist of Wireless Application Protocol (WAP) servers or it could be the Internet.

When supporting Mobile IP and/or Proxy Mobile IP data applications, the system can be configured to perform the role of the PDSN/Foreign Agent (FA) and/or the Home Agent (HA) within the carrier's 3G CDMA2000 data network. When functioning as an HA, the system can either be located within the carrier's 3G network or in an external enterprise or ISP network. Regardless, the PDSN/FA terminates the mobile subscriber's PPP session, and then routes data to and from the appropriate HA on behalf of the subscriber.

Product Specifications

This section describes the hardware and software requirements for a PDSN service.

Hardware Requirements

This section describes the hardware required to enable the PDSN service.

Platforms

The PDSN service service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

ASR 5x00 Series Platform System Hardware Components

The following application and line cards are required to support CDMA2000 wireless data services on the system:

- **System Management Cards (SMCs):** Provides full system control and management of all cards within the ASR 5x00 platform. Up to two SMC can be installed; one active, one redundant.
- **Packet Services Cards (PSCs):** Within the ASR 5x00 platform, PSCs provide high-speed, multi-threaded PPP processing capabilities to support either PDSN/FA or HA services. Up to 14 PSCs can be installed, allowing for multiple active and/or redundant cards.
- **Switch Processor Input/Outputs (SPIO):** Installed in the upper-rear chassis slots directly behind the SPCs/SMCs, SPIOs provide connectivity for local and remote management, Central Office (CO) alarms. Up to two SPIOs can be installed; one active, one redundant.
- **Ethernet 10/100 and/or Ethernet 1000/Quad Gig-E Line Cards (QGLC):** Installed directly behind PSCs, these cards provide the RP, AAA, PDN, and Pi interfaces to elements in the data network. Up to 26 line cards should be installed for a fully loaded system with 13 active PSCs, 13 in the upper-rear slots and 13 in the lower-rear slots for redundancy. Redundant PSCs do no not require line cards.
- **Redundancy Crossbar Cards (RCCs):** Installed in the lower-rear chassis slots directly behind the SMCs, RCCs utilize 5 Gbps serial links to ensure connectivity between Ethernet 10/100 or Ethernet 1000 line cards/QGLCs and every PSC in the system for redundancy. Two RCCs can be installed to provide redundancy for all line cards and PSCs.

Important: Additional information pertaining to each of the application and line cards required to support CDMA2000 wireless data services is located in the *Product Overview Guide*.

Features and Functionality—Base Software

This section describes the features and functions supported by default in base software on PDSN service and do not require any additional licenses.

Important: To configure the basic service and functionality on the system for PDSN service, refer to the configuration examples provided in the PDSN Administration Guide.

This section describes following features:

- Gx and Gy Support
- RADIUS Support
- Access Control List Support
- IP Policy Forwarding
- AAA Server Groups
- Overlapping IP Address Pool Support
- Routing Protocol Support
- Management System Overview
- Bulk Statistics Support
- Threshold Crossing Alerts (TCA) Support
- IP Header Compression - Van Jacobson
- DSCP Marking

Gx and Gy Support

The PDSN supports 3GPP Release 8 standards based policy interface with the Policy and Charging Rules Function (PCRF). The policy interface is based on a subset 3GPP 29.212. based Gx interface specification. The PDSN policy interface fully supports installation/modification of dynamic and predefined rules from the PCRF.

The enforcement of dynamic and predefined PCC rules installed from the PCRF is done using Enhanced Charging Services (ECS).The full ECS functionality including the DPI and P2P detection can b e enabled via predefined rules using the Gx interface.

The PDSN supports a subset of event triggers as defined in 29.212. Currently the event trigger support is limited to the following:

- RAT Change
- User location change (BSID)
- AN GW change (during inter PCF handoff)

The PDSN also supports triggering of online charging via the policy interface. 3GPP Release 8 Gy interface as defined in 32.299 is used for online charging.

The PDSN supports connectivity to multiple PCRF's . The PCRF's may be referred to by an FQDN. Load balancing of sessions across multiple servers are achieved by using a round robin algorithm. Redundancy between servers can be achieved by configuring multiple weighted sets of servers.

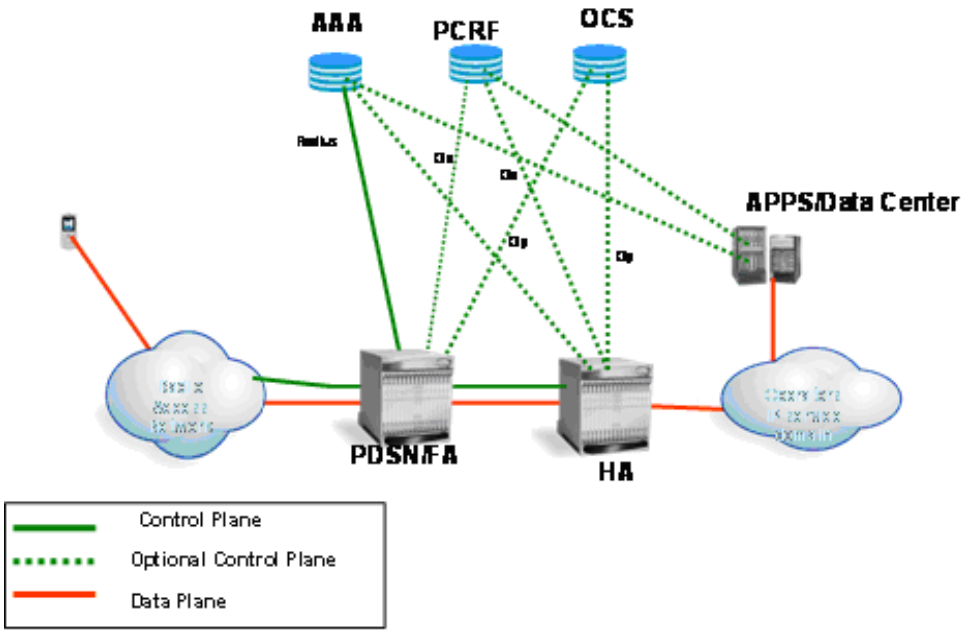
The configuration allows Policy support to be enabled on a per subscriber/APN basis.

The policy features supported on PDSN and GGSN will be quite similar. On PDSN the Gx will only be supported for Simple IP calls.

On PDSN additional event triggers rat type change and location change will be supported.On PDSN Gy , standard DCCA based credit control is supported , 3GPP related trigger functionality is not supported on PDSN Gy.

The following figure shows the Gx support for Simple IP.

Figure 64. Gx for Simple IP



RADIUS Support

Provides a mechanism for performing authorization, authentication, and accounting (AAA) for subscriber PDP contexts based on the following standards:

- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000
- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000

Description

The Remote Authentication Dial-In User Service (RADIUS) protocol is used to provide AAA functionality for subscriber PDP contexts.

Within context contexts configured on the system, there are AAA and RADIUS protocol-specific parameters that can be configured. The RADIUS protocol-specific parameters are further differentiated between RADIUS Authentication server RADIUS Accounting server interaction.

Among the RADIUS parameters that can be configured are:

- **Priority**: Dictates the order in which the servers are used allowing for multiple servers to be configured in a single context.
- **Routing Algorithm**: Dictate the method for selecting among configured servers. The specified algorithm dictates how the system distributes AAA messages across the configured AAA servers for new sessions. Once a session is established and an AAA server has been selected, all subsequent AAA messages for the session will be delivered to the same server.

In the event that a single server becomes unreachable, the system attempts to communicate with the other servers that are configured. The system also provides configurable parameters that specify how it should behave should all of the RADIUS AAA servers become unreachable.

The system provides an additional level of flexibility by supporting the configuration RADIUS server groups. This functionality allows operators to differentiate AAA services based on the subscriber template used to facilitate their PDP context.

In general, 128 AAA Server IP address/port per context can be configured on the system and it selects servers from this list depending on the server selection algorithm (round robin, first server). Instead of having a single list of servers per context, this feature provides the ability to configure multiple server groups. Each server group, in turn, consists of a list of servers.

This feature works in following way:

- All RADIUS authentication/accounting servers configured at the context-level are treated as part of a server group named “default”. This default server group is available to all subscribers in that context through the realm (domain) without any configuration.
- It provides a facility to create “user defined” RADIUS server groups, as many as 399 (excluding “default” server group), within a context. Any of the user defined RADIUS server groups are available for assignment to a subscriber through the subscriber configuration within that context.

Since the configuration of the subscriber can specify the RADIUS server group to use as well as IP address pools from which to assign addresses, the system implements a mechanism to support some in-band RADIUS server implementations (i.e. RADIUS servers which are located in the corporate network, and not in the operator's network) where the NAS-IP address is part of the subscriber pool. In these scenarios, the PDSN supports the configuration of the first IP address of the subscriber pool for use as the RADIUS NAS-IP address.

Important: In 12.3 and earlier releases, refer to the *AAA and GTPP Interface Administration and Reference* for more information on RADIUS AAA configuration.

Access Control List Support

Access Control Lists provide a mechanism for controlling (i.e permitting, denying, redirecting, etc.) packets in and out of the system.

IP access lists, or Access Control Lists (ACLs) as they are commonly referred to, are used to control the flow of packets into and out of the system. They are configured on a per-context basis and consist of "rules" (ACL rules) or filters that control the action taken on packets that match the filter criteria. Once configured, an ACL can be applied to any of the following:

- An individual interface
- All traffic facilitated by a context (known as a policy ACL)
- An individual subscriber
- All subscriber sessions facilitated by a specific context

There are two primary components of an ACL:

- Rule: A single ACL consists of one or more ACL rules. As discussed earlier, the rule is a filter configured to take a specific action on packets matching specific criteria. Up to 128 rules can be configured per ACL.

  Each rule specifies the action to take when a packet matches the specifies criteria. This section discusses the rule actions and criteria supported by the system.
- Rule Order: A single ACL can consist of multiple rules. Each packet is compared against each of the ACL rules, in the order in which they were entered, until a match is found. Once a match is identified, all subsequent rules are ignored.

Important: For more information on Access Control List configuration, refer to the IP Access Control List chapter in System Administration Guide.

IP Policy Forwarding

IP Policy Forwarding enables the routing of subscriber data traffic to specific destinations based on configuration. This functionality can be implemented in support of enterprise-specific applications (i.e. routing traffic to specific enterprise domains) or for routing traffic to back-end servers for additional processing.

Description

The system can be configured to automatically forward data packets to a predetermined network destination. This can be done in one of three ways:

- **IP Pool-based Next Hop Forwarding** - Forwards data packets based on the IP pool from which a subscriber obtains an IP address.
- **ACL-based Policy Forwarding** - Forwards data packets based on policies defined in Access Control Lists (ACLs) and applied to contexts or interfaces.
- **Subscriber specific Next Hop Forwarding** - Forwards all packets for a specific subscriber.

The simplest way to forward subscriber data is to use IP Pool-based Next Hop Forwarding. An IP pool is configured with the address of a next hop gateway and data packets from all subscribers using the IP pool are forward to that gateway.

Subscriber Next Hop forwarding is also very simple. In the subscriber configuration a nexthop forwarding address is specified and all data packets for that subscriber are forwarded to the specified nexthop destination.

ACL-based Policy Forwarding gives you more control on redirecting data packets. By configuring an Access Control List (ACL) you can forward data packets from a context or an interface by different criteria, such as; source or destination IP address, ICMP type, or TCP/UDP port numbers.

ACLs are applied first. If ACL-based Policy Forwarding and Pool-based Next Hop Forwarding or Subscriber are configured, data packets are first redirected as defined in the ACL, then all remaining data packets are redirected to the next hop gateway defined by the IP pool or subscriber profile.

AAA Server Groups

Value-added feature to enable VPN service provisioning for enterprise or MVNO customers. Enables each corporate customer to maintain its own AAA servers with its own unique configurable parameters and custom dictionaries.

Description

This feature provides support for up to 800 AAA (RADIUS and Diameter) server groups and 800 NAS IP addresses that can be provisioned within a single context or across the entire chassis. A total of 128 servers can be assigned to an individual server group. Up to 1,600 accounting, authentication and/or mediation servers are supported per chassis and may be distributed across a maximum of 1,000 subscribers. This feature also enables the AAA servers to be distributed across multiple subscribers within the same context.

Important: Due to additional memory requirements, this service can only be used with 8GB Packet Accelerator Cards (PACs) or Packet Service Cards (PSCs)

Important: In 12.3 and earlier releases, refer to the *AAA and GTPP Interface Administration and Reference* for more information on AAA Server Group configuration.

Overlapping IP Address Pool Support

Overlapping IP Address Pools provides a mechanism for allowing operators to more flexibly support multiple corporate VPN customers with the same private IP address space without the expensive investments in physically separate routers, or expensive configurations using virtual routers.

Important: For more information on IP pool overlapping configuration, refer to the VLANs chapter in the *System Administration Guide*.

Routing Protocol Support

The system's support for various routing protocols and routing mechanism provides an efficient mechanism for ensuring the delivery of subscriber data packets.

Description

The following routing mechanisms and protocols are supported by the system:

- **Static Routes:** The system supports the configuration of static network routes on a per context basis. Network routes are defined by specifying an IP address and mask for the route, the name of the interface in the currant context that the route must use, and a next hop IP address.

- **Open Shortest Path First (OSPF) Protocol version 2:** A link-state routing protocol, OSPF is an Interior Gateway Protocol (IGP) that routes IP packets based solely on the destination IP address found in the IP packet header using the shortest path first. IP packets are routed "as is", meaning they are not encapsulated in any further protocol headers as they transit the network.

  Variable length subnetting, areas, and redistribution into and out of OSPF are supported.

  OSPF routing is supported in accordance with the following standards:

  - RFC-1850, OSPF Version 2 Management Information Base, November 1995
  - RFC-2328, OSPF Version 2, April 1998
  - RFC-3101 OSPF-NSSA Option, January 2003

- **Border Gateway Protocol version 4 (BGP-4):** The system supports a subset of BGP (RFC-1771, A Border Gateway Protocol 4 (BGP-4)), suitable for eBGP support of multi-homing typically used to support geographically redundant mobile gateways, is supported.

  EBGP is supported with multi-hop, route filtering, redistribution, and route maps. The network command is support for manual route advertisement or redistribution.

  BGP route policy and path selection is supported by the following means:

  - Prefix match based on route access list
  - AS path access-list
  - Modification of AS path through path prepend
  - Origin type
  - MED
  - Weight

- **Route Policy:** Routing policies modify and redirect routes to and from the system to satisfy specific routing needs. The following methods are used with or without active routing protocols (i.e. static or dynamic routing) to prescribe routing policy:

  - **Route Access Lists:** The basic building block of a routing policy, route access lists filter routes based upon a specified range of IP addresses.
  - **IP Prefix Lists:** A more advanced element of a routing policy. An IP Prefix list filters routes based upon IP prefixes.
  - **AS Path Access Lists:** A basic building block used for Border Gateway Protocol (BGP) routing, these lists filter Autonomous System (AS) paths.

- **Route Maps:** Route-maps are used for detailed control over the manipulation of routes during route selection or route advertisement by a routing protocol and in route redistribution between routing protocols. This detailed control is achieved using IP Prefix Lists, Route Access Lists and AS Path Access Lists to specify IP addresses, address ranges, and Autonomous System Paths.

- **Equal Cost Multiple Path (ECMP):** ECMP allows distribution of traffic across multiple routes that have the same cost to the destination. In this manner, throughput load is distributed across multiple path, typically to lessen the burden on any one route and provide redundancy. The mobile gateway supports from four to ten equal-cost paths.

Important: For more information on IP Routing configuration, refer to the Routing chapter in the *System Administration Guide*.

Management System Overview

The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management -- focusing on providing superior quality Network Element (NE) and element management system (Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems -- giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

Description

Cisco's O&M module offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces.

These include:

- Using the Command Line Interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager application
- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000
- Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (i.e. Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

Important: For more information on command line interface based management, refer to the CDMA Command Line Interface Reference and PDSN Administration Guide.

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

Description

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. The following schemas are supported:

- System: Provides system-level statistics
- **Card**: Provides card-level statistics
- **Port**: Provides port-level statistics
- **BCMCS**: Provides BCMCS service statistics
- **FA**: Provides FA service statistics
- **HA**: Provides HA service statistics
- **IP Pool**: Provides IP pool statistics
- **MIPv6HA**: Provides MIPv6HA service statistics
- **PPP**: Provides Point-to-Point Protocol statistics
- **RADIUS**: Provides per-RADIUS server statistics
- **ECS**: Provides Enhanced Charging Service Statistics

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the IMG or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, IMG host name, IMG uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

Description

The following thresholding models are supported by the system:

- **Alert:** A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm:** Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps:** SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values.

  Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.
- **Logs:** The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

  Logs are supported in both the Alert and the Alarm models.
- **Alarm System:** High threshold alarms generated within the specified polling interval are considered "outstanding" until a the condition no longer exists or a condition clear alarm is generated. "Outstanding" alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

  The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

IP Header Compression - Van Jacobson

Implementing IP header compression provides the following benefits:

- Improves interactive response time
- Allows the use of small packets for bulk data with good line efficiency
- Allows the use of small packets for delay sensitive low data-rate traffic
- Decreases header overhead
- Reduces packet loss rate over lossy links

Description

The system supports the Van Jacobson (VJ) IP header compression algorithms by default for subscriber traffic.

The VJ header compression is supported as per RFC 1144 (CTCP) header compression standard developed by V. Jacobson in 1990. It is commonly known as VJ compression. It describes a basic method for compressing the headers of IPv4/TCP packets to improve performance over low speed serial links.

By default IP header compression using the VJ algorithm is enabled for subscribers. You can also turn off IP header compression for a subscriber.

Important: For more information on IP header compression support, refer to the IP Header Compression chapter.

DSCP Marking

Provides support for more granular configuration of DSCP marking.

For different Traffic class, the PDSN supports per-service and per-subscriber configurable DSCP marking for Uplink and Downlink direction based on Allocation/Retention Priority in addition to the current priorities.

Features and Functionality - Optional Enhanced Software Features

This section describes the optional enhanced features and functions for PDSN service.

Each of the following features require the purchase of an additional license to implement the functionality with the PDSN service.

This section describes following features:

- Session Recovery Support
- IPv6 Support
- L2TP LAC Support
- L2TP LNS Support
- Proxy Mobile IP
- IP Security (IPSec)
- Traffic Policing and Rate Limiting
- Dynamic RADIUS Extensions (Change of Authorization)
- Web Element Management System

Session Recovery Support

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

Description

Session recovery is performed by mirroring key software processes (e.g. session manager and AAA manager) within the system. These mirrored processes remain in an idle state (in standby-mode), wherein they perform no processing, until they may be needed in the case of a software failure (e.g. a session manager task aborts). The system spawns new instances of "standby mode" session and AAA managers for each active Control Processor (CP) being used.

Additionally, other key system-level software tasks, such as VPN manager, are performed on a physically separate Packet Services Card (PSC) to ensure that a double software fault (e.g. session manager and VPN manager fails at same time on same card) cannot occur. The PSC used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

The additional hardware resources required for session recovery include a standby System Processor Card (SPC) and a standby PSC.

There are two modes for Session Recovery.

- **Task recovery mode:** Wherein one or more session manager failures occur and are recovered without the need to use resources on a standby PSC. In this mode, recovery is performed by using the mirrored "standby-mode" session manager task(s) running on active PACs. The "standby-mode" task is renamed, made active, and is then populated using information from other tasks such as AAA manager.

- **Full recovery mode:** Used when a PSC hardware failure occurs, or when a PSC migration failure happens. In this mode, the standby PSC is made active and the "standby-mode" session manager and AAA manager tasks on the newly activated PSC perform session recovery.

Session/Call state information is saved in the peer AAA manager task because each AAA manager and session manager task is paired together. These pairs are started on physically different PACs to ensure task recovery.

Important: For more information on session recovery support, refer to the Session Recovery chapter in the *System Administration Guide*.

IPv6 Support

This feature allows IPv6 subscribers to connect via the CDMA 2000 infrastructure in accordance with the following standards:

- RFC 2460: Internet Protocol, Version 6 (IPv6) Specification
- RFC 2461: Neighbor Discovery for IPv6
- RFC 2462: IPv6 Stateless Address Autoconfiguration
- RFC 3314: Recommendations for IPv6 in 3GPP Standards
- RFC 3316: Internet Protocol Version 6 (IPv6) for Some Second and Third Generation Cellular Hosts
- RFC 3056: Connection of IPv6 domains via IPv4 clouds
- 3GPP TS 23.060: General Packet Radio Service (GPRS) Service description
- 3GPP TS 27.060: Mobile Station Supporting Packet Switched Services
- 3GPP TS 29.061: Interworking between the Public Land Mobile Network (PLMN) supporting Packet Based Services and Packet Data Networks (PDN)

Description

The PDSN allows a subscriber to be configured for IPv6 PDP contexts. Also, a subscriber may be configured to simultaneously allow IPv4 PDP contexts.

The PDSN supports IPv6 stateless dynamic auto-configuration. The mobile station may select any value for the interface identifier portion of the address. The link-local address is assigned by the PDSN to avoid any conflict between the mobile station link-local address and the PDSN address. The mobile station uses the interface identifier assigned by the PDSN during the stateless address auto-configuration procedure. Once this has completed, the mobile can select any interface identifier for further communication as long as it does not conflict with the PDSN's interface identifier that the mobile learned through router advertisement messages from the PDSN.

Control and configuration of the above is specified as part of the subscriber configuration on the PDSN, e.g., IPv6 address prefix and parameters for the IPv6 router advertisements. RADIUS VSAs may be used to override the subscriber configuration.

Following IPv6 PDP context establishment, the PDSN can perform either manual or automatic 6to4 tunneling, according to RFC 3056, Connection of IPv6 Domains Via IPv4 Clouds.

L2TP LAC Support

The system configured as a Layer 2 Tunneling Protocol Access Concentrator (LAC) enables communication with L2TP Network Servers (LNSs) for the establishment of secure Virtual Private Network (VPN) tunnels between the operator and a subscriber's corporate or home network.

Description

The use of L2TP in VPN networks is often used as it allows the corporation to have more control over authentication and IP address assignment. An operator may do a first level of authentication, however use PPP to exchange user name and password, and use IPCP to request an address. To support PPP negotiation between the PDSN and the corporation, an L2TP tunnel must be setup in the PDSN running a LAC service.

L2TP establishes L2TP control tunnels between LAC and LNS before tunneling the subscriber PPP connections as L2TP sessions. The LAC service is based on the same architecture as the PDSN and benefits from dynamic resource allocation and distributed message and data processing. This design allows the LAC service to support over 4000 setups per second or a maximum of over 3G of throughput. There can be a maximum up to 65535 sessions in a single tunnel and as many as 500,000 L2TP sessions using 32,000 tunnels per system.

The LAC sessions can also be configured to be redundant, thereby mitigating any impact of hardware of software issues. Tunnel state is preserved by copying the information across processor cards.

Important: For more information on L2TP Access Concentrator support, refer to the L2TP Access Concentrator chapter.

L2TP LNS Support

The system configured as a Layer 2 Tunneling Protocol Network Server (LNS) supports the termination secure Virtual Private Network (VPN) tunnels between from L2TP Access Concentrators (LACs).

Description

The LNS service takes advantage of the high performance PPP processing already supported in the system design and is a natural evolution from the LAC. The LNS can be used as a standalone, or running alongside a PDSN service in the same platform, terminating L2TP services in a cost effective and seamless manner.

L2TP establishes L2TP control tunnels between LAC and LNS before tunneling the subscriber PPP connections as L2TP sessions. There can be a maximum of up to 65535 sessions in a single tunnel and up to 500,000 sessions per LNS.

The LNS architecture is similar to the PDSN and utilizes the concept of a de-multiplexer to intelligently assign new L2TP sessions across the available software and hardware resources on the platform without operator intervention..

Important: For more information on L2TP LNS support support, refer to the L2TP Access Concentrator chapter.

Proxy Mobile IP

Mobility for subscriber sessions is provided through the Mobile IP protocol as defined in RFCs 2002-2005. However, some older Mobile Nodes (MNs) do not support the Mobile IP protocol. The Proxy Mobile IP feature provides a mobility solution for these MNs.

Description

For IP PDP contexts using Proxy Mobile IP, the MN establishes a session with the PDSN as it normally would. However, the PDSN/FA performs Mobile IP operations with an HA (identified by information stored in the subscriber's profile) on behalf of the MN (i.e. the MN is only responsible for maintaining the IP PDP context with the PDSN, no Agent Advertisement messages are communicated with the MN).

The MN is assigned an IP address by either the HA, an AAA server, or on a static-basis. The address is stored in a Mobile Binding Record (MBR) stored on the HA. Therefore, as the MN roams through the service provider's network, each time a hand-off occurs, the MN will continue to use the same IP address stored in the MBR on the HA.

Proxy Mobile IP can be performed on a per-subscriber basis based on information contained in their user profile, or for all subscribers facilitated by a specific subscriber. In the case of non-transparent IP PDP contexts, attributes returned from the subscriber's profile take precedence over the configuration of the subscriber.

Important: For more information on Proxy Mobile IP configuration, refer to the Proxy Mobile IP chapter.

IP Security (IPSec)

IP Security provides a mechanism for establishing secure tunnels from mobile subscribers to pre-defined endpoints (i.e. enterprise or home networks) in accordance with the following standards:

- RFC 2401, Security Architecture for the Internet Protocol
- RFC 2402, IP Authentication Header (AH)
- RFC 2406, IP Encapsulating Security Payload (ESP)
- RFC 2409, The Internet Key Exchange (IKE)
- RFC-3193, Securing L2TP using IPSEC, November 2001

Description

IP Security (IPSec) is a suite of protocols that interact with one another to provide secure private communications across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways. IPSec can be implemented on the system for the following applications:

- PDN Access: Subscriber IP traffic is routed over an IPSec tunnel from the system to a secure gateway on the Packet Data Network (PDN) as determined by Access Control List (ACL) criteria.
- Mobile IP: Mobile IP control signals and subscriber data is encapsulated in IPSec tunnels that are established between Foreign Agents (FAs) and Home Agents (HAs) over the Pi interfaces.

  Once an IPSec tunnel is established between an FA and HA for a particular subscriber, all new Mobile IP sessions using the same FA and HA are passed over the tunnel regardless of whether or not IPSec is supported for the new subscriber sessions. Data for existing Mobile IP sessions is unaffected.
- L2TP: L2TP-encapsulated packets are routed from the system to an LNS/secure gateway over an IPSec tunnel.

Important: For more information on IPSec support, refer to the IP Security chapter.

Traffic Policing and Rate Limiting

Allows the operator to proportion the network and support Service-level Agreements (SLAs) for customers

Description

The Traffic-Policing/Shaping feature enables configuring and enforcing bandwidth limitations on individual PDP contexts of a particular 3GPP traffic class. Values for traffic classes are defined in 3GPP TS 23.107 and are negotiated with the SGSN during PDP context activation using the values configured for the subscriber on the PDSN. Configuration and enforcement is done independently on the downlink and the uplink directions for each of the 3GPP traffic classes. Configuration is on a per-subscriber basis, but may be overridden for individual subscribers or subscriber tiers during RADIUS authentication/authorization.

A Token Bucket Algorithm (a modified trTCM, as specified in RFC2698) is used to implement the Traffic-Policing feature. The algorithm measures the following criteria when determining how to mark a packet.

- **Committed Data Rate (CDR):** The guaranteed rate (in bits per second) at which packets may be transmitted/received for the subscriber during the sampling interval.
- **Peak Data Rate (PDR):** The maximum rate (in bits per second) that packets may be transmitted/received for the subscriber during the sampling interval.
- **Burst-size:** The maximum number of bytes that may be transmitted/received for the subscriber during the sampling interval for both committed (CBS) and peak (PBS) rate conditions. This represents the maximum number of tokens that can be placed in the subscriber's "bucket". Note that the committed burst size (CBS) equals the peak burst size (PBS) for each subscriber.

  Tokens are removed from the subscriber's bucket based on the size of the packets being transmitted/received. Every time a packet arrives, the system determines how many tokens need to be added (returned) to a subscriber's CBS (and PBS) bucket. This value is derived by computing the product of the time difference between incoming packets and the CDR (or PDR). The computed value is then added to the tokens remaining in the subscriber's CBS (or PBS) bucket. The total number of tokens can not be greater than the configured burst-size. If the total number of tokens is greater than the burst-size, the number is set to equal the burst-size. After passing through the Token Bucket Algorithm, the packet is internally classified with a color, as follows:

  - There are not enough tokens in the PBS bucket to allow a packet to pass, then the packet is considered to be in violation and is marked "red" and the violation counter is incremented by one.
  - There are enough tokens in the PBS bucket to allow a packet to pass, but not in the CBS "bucket", then the packet is considered to be in excess and is marked "yellow", the PBS bucket is decremented by the packet size, and the exceed counter is incremented by one.
  - There are more tokens present in the CBS bucket than the size of the packet, then the packet is considered as conforming and is marked "green" and the CBS and PBS buckets are decremented by the packet size.

  The subscriber on the PDSN can be configured with actions to take for red and yellow packets. Any of the following actions may be specified:

  - **Drop**: The offending packet is discarded.
  - **Transmit**: The offending packet is passed.
  - **Lower the IP Precedence**: The packet's ToS octet is set to "0", thus downgrading it to Best Effort, prior to passing the packet.

- **Buffer the Packet**: The packet stored in buffer memory and transmitted to subscriber once traffic flow comes in allowed bandwidth.

Different actions may be specified for red and yellow, as well as for uplink and downlink directions and different 3GPP traffic classes.

Refer to the Intelligent Traffic Control section for additional policing and shaping capabilities of the PDSN.

Important: For more information on per subscriber traffic policing and shaping, refer to the Traffic Policing and Shaping section.

Intelligent Traffic Control

Enables operators to provide differentiated tiered service provisioning for native and non-native subscribers.

Description

Mobile carriers are looking for creative methods for maximizing network resources while, at the same time, enhancing their end users overall experience. These same mobile operators are beginning to examine solutions for providing preferential treatment for their native subscribers and services as compared to, for example, roaming subscribers, Mobile Virtual Network Operators (MVNOs) and/or Peer-to-Peer (P2P) applications. The overall end goal is to provide superior levels of performance for their customers/services, while ensuring that non-native users/applications do not overwhelm network resources.

ITC provides the ability to examine each subscriber session and respective flow(s) such that selective, configurable limits on a per-subscriber/per-flow basis can be applied. Initially, QoS in this context is defined as traffic policing on a per-subscriber/per-flow basis with the potential to manipulate Differentiated Services Code Points (DSCPs), queue redirection (i.e. move traffic to a Best Effort (BE) classification) and/or simply dropping out of profile traffic. ITC enables 5 tuple packet filters for individual application flows to be either manually configured via CLI or dynamically established via RSVP TFT information elements in 1xEV-DO Rev A or as a consequence of PDP context establishments in CDMA networks. Policy rules may be locally assigned or obtained from an external PCRF via push/pull policy signaling interactions. Policies may be applied on a per-subscriber, per-context and/or chassis-wide basis.

Important: For more information on intelligent traffic control support, refer to the Intelligent Traffic Control chapter.

Dynamic RADIUS Extensions (Change of Authorization)

Dynamic RADIUS extension support provide operators with greater control over subscriber PDP contexts by providing the ability to dynamically redirect data traffic, and or disconnect the PDP context.

This functionality is based on the RFC 3576, Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS), July 2003 standard.

Description

The system supports the configuration and use of the following dynamic RADIUS extensions:

- **Change of Authorization**: The system supports CoA messages from the AAA server to change data filters associated with a subscriber session. The CoA request message from the AAA server must contain attributes to identify NAS and the subscriber session and a data filter ID for the data filter to apply to the subscriber session.
- **Disconnect Message:** The DM message is used to disconnect subscriber sessions in the system from a RADIUS server. The DM request message should contain necessary attributes to identify the subscriber session.

The above extensions can be used to dynamically re-direct subscriber PDP contexts to an alternate address for performing functions such as provisioning and/or account set up. This functionality is referred to as Session Redirection, or Hotlining.

Session redirection provides a means to redirect subscriber traffic to an external server by applying ACL rules to the traffic of an existing or a new subscriber session. The destination address and optionally the destination port of TCP/IP or UDP/IP packets from the subscriber are rewritten so the packet is forwarded to the designated redirected address.

Return traffic to the subscriber has the source address and port rewritten to the original values. The redirect ACL may be applied dynamically by means of the Radius Change of Authorization (CoA) extension.

Important: For more information on dynamic RADIUS extensions support, refer to the CoA, RADIUS, And Session Redirection (Hotlining) chapter.

Web Element Management System

Provides a Graphical User Interface (GUI) for performing Fault, Configuration, Accounting, Performance, and Security (FCAPS) management of the ASR 5x00 system.

Description

The Web Element Manager is a Common Object Request Broker Architecture (CORBA)-based application that provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) management capability for the system.

For maximum flexibility and scalability, the Web Element Manager application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the fault-tolerant Sun® Solaris® operating system.

Important: For more information on WEM support, refer to the *WEM Installation and Administration Guide*.

Features and Functionality - Inline Service Support

This section describes the features and functions of inline services supported on the PDSN. These services require additional licenses to implement the functionality.

Content Filtering

The Cisco PDSN offers two variants of network-controlled content filtering / parental control services. Each approach leverages the native DPI capabilities of the platform to detect and filter events of interest from mobile subscribers based on HTTP URL or WAP/MMS URI requests:

- Integrated Content Filtering: A turnkey solution featuring a policy enforcement point and category based rating database on the Cisco PDSN. An offboard AAA or PCRF provides the per-subscriber content filtering information as subscriber sessions are established. The content filtering service uses DPI to extract URL's or URI's in HTTP request messages and compares them against a static rating database to determine the category match. The provisioned policy determines whether individual subscribers are entitled to view the content.
- Content Filtering ICAP Interface: This solution is appropriate for mobile operators with existing installations of Active Content Filtering external servers. The service continues to harness the DPI functions of the ASR 5000 platform to extract events of interest. However in this case, the extracted requests are transferred via the Integrated Content Adaptation Protocol (ICAP) with subscriber identification information to the external ACF server which provides the category rating database and content decision functions.

Integrated Adult Content Filter

Provides a value-added service to prevent unintended viewing of objectionable content that exploits underage children. Content Filtering offers mobile operators a way to increase data ARPU and subscriber retention through a network-based solution for parental controls and content filtering. The integrated solution enables a single policy decision and enforcement point thereby streamlining the number of signaling interactions with external AAA/Policy Manager servers. When used in parallel with other services such as Enhanced Content Charging (ECS) it increases billing accuracy of charging records by insuring that mobile subscribers are only charged for visited sites they are allowed to access.

The Integrated Adult Content Filter is a subscriber-aware inline service provisioned on an ASR 5000 running PDSN services. Integrated Content Filtering utilizes the local DPI engine and harnesses a distributed software architecture that scales with the number of active PDSN sessions on the system.

Content Filtering policy enforcement is the process of deciding if a subscriber should be able to receive some content. Typical options are to allow, block, or replace/redirect the content based on the rating of the content and the policy defined for that content and subscriber. The policy definition is transferred in an authentication response from a AAA server or Diameter policy message via the Gx reference interface from an adjunct PCRF. The policy is applied to subscribers through rulebase or APN/Subscriber configuration. The policy determines the action to be taken on the content request on the basis of its category. A maximum of one policy can be associated with a rulebase.

ICAP Interface

Provides a value-added service to prevent unintended viewing of objectionable content that exploits underage children. Content Filtering offers mobile operators a way to increase data ARPU and subscriber retention through a network-based solution for parental controls and content filtering. The Content Filtering ICAP solution is appropriate for operators with existing installations of Active Content Filtering servers in their networks.

The Enhanced Charging Service (ECS) provides a streamlined Internet Content Adaptation Protocol (ICAP) interface to leverage the Deep Packet Inspection (DPI) to enable external Application Servers to provide their services without performing the DPI functionality and without being inserted in the data flow. The ICAP interface may be attractive to mobile operators that prefer to use an external Active Content Filtering (ACF) Platform. If a subscriber initiates a WAP (WAP1.x or WAP2.0) or Web session, the subsequent GET/POST request is detected by the deep packet inspection function. The URL of the GET/POST request is extracted by the local DPI engine on the ASR 5000 platform and passed, along with subscriber identification information and the subscriber request, in an ICAP message to the Application Server (AS). The AS checks the URL on the basis of its category and other classifications like, type, access level, content category and decides if the request should be authorized, blocked or redirected by answering the GET/POST message. Depending upon the response received from the ACF server, the PDSN either passes the request unmodified or discards the message and responds to the subscriber with the appropriate redirection or block message.

Network Address Translation (NAT)

NAT translates non-routable private IP address(es) to routable public IP address(es) from a pool of public IP addresses that have been designated for NAT. This enables to conserve on the number of public IP addresses required to communicate with external networks, and ensures security as the IP address scheme for the internal network is masked from external hosts, and each outgoing and incoming packet goes through the translation process.

NAT works by inspecting both incoming and outgoing IP datagrams and, as needed, modifying the source IP address and port number in the IP header to reflect the configured NAT address mapping for outgoing datagrams. The reverse NAT translation is applied to incoming datagrams.

NAT can be used to perform address translation for simple IP and mobile IP. NAT can be selectively applied/denied to different flows (5-tuple connections) originating from subscribers based on the flows' L3/L4 characteristics—Source-IP, Source-Port, Destination-IP, Destination-Port, and Protocol.

NAT supports the following mappings:

- One-to-One
- Many-to-One

Important: For more information on NAT, refer to the *Cisco ASR 5000 Series Network Address Translation Administration Guide*.

Peer-to-Peer Detection

Allows operators to identify P2P traffic in the network and applying appropriate controlling functions to ensure fair distribution of bandwidth to all subscribers.

Peer-to-Peer (P2P) is a term used in two slightly different contexts. At a functional level, it means protocols that interact in a peering manner, in contrast to client-server manner. There is no clear differentiation between the function of one node or another. Any node can function as a client, a server, or both—a protocol may not clearly differentiate between the two. For example, peering exchanges may simultaneously include client and server functionality, sending and receiving information.

Detecting peer-to-peer protocols requires recognizing, in real time, some uniquely identifying characteristic of the protocols. Typical packet classification only requires information uniquely typed in the packet header of packets of the stream(s) running the particular protocol to be identified. In fact, many peer-to-peer protocols can be detected by simple packet header inspection. However, some P2P protocols are different, preventing detection in the traditional manner. This is designed into some P2P protocols to purposely avoid detection. The creators of these protocols purposely do not

publish specifications. A small class of P2P protocols is stealthier and more challenging to detect. For some protocols no set of fixed markers can be identified with confidence as unique to the protocol.

Operators care about P2P traffic because of the behavior of some P2P applications (for example, Bittorrent, Skype, and eDonkey). Most P2P applications can hog the network bandwidth such that 20% P2P users can generate as much as traffic generated by the rest 80% non-P2P users. This can result into a situation where non-P2P users may not get enough network bandwidth for their legitimate use because of excess usage of bandwidth by the P2P users. Network operators need to have dynamic network bandwidth / traffic management functions in place to ensure fair distributions of the network bandwidth among all the users. And this would include identifying P2P traffic in the network and applying appropriate controlling functions to the same (for example, content-based premium billing, QoS modifications, and other similar treatments).

Cisco's P2P detection technology makes use of innovative and highly accurate protocol behavioral detection techniques.

Important: For more information on peer-to-peer detection, refer to the *Cisco ASR 5000 Series Application Detection and Control Administration Guide*.

Personal Stateful Firewall

The Personal Stateful Firewall is an in-line service feature that inspects subscriber traffic and performs IP session-based access control of individual subscriber sessions to protect the subscribers from malicious security attacks.

The Personal Stateful Firewall supports stateless and stateful inspection and filtering based on the configuration.

In stateless inspection, the firewall inspects a packet to determine the 5-tuple—source and destination IP addresses and ports, and protocol—information contained in the packet. This static information is then compared against configurable rules to determine whether to allow or drop the packet. In stateless inspection the firewall examines each packet individually, it is unaware of the packets that have passed through before it, and has no way of knowing if any given packet is part of an existing connection, is trying to establish a new connection, or is a rogue packet.

In stateful inspection, the firewall not only inspects packets up through the application layer / layer 7 determining a packet's header information and data content, but also monitors and keeps track of the connection's state. For all active connections traversing the firewall, the state information, which may include IP addresses and ports involved, the sequence numbers and acknowledgement numbers of the packets traversing the connection, TCP packet flags, etc. is maintained in a state table. Filtering decisions are based not only on rules but also on the connection state established by prior packets on that connection. This enables to prevent a variety of DoS, DDoS, and other security violations. Once a connection is torn down, or is timed out, its entry in the state table is discarded.

The Enhanced Charging Service (ECS) / Active Charging Service (ACS) in-line service is the primary vehicle that performs packet inspection and charging. For more information on ECS, see the *Cisco ASR 5000 Series Enhanced Charging Service Administration Guide*.

Important: For more information on Personal Stateful Firewall, refer to the *Cisco ASR 5000 Series Personal Stateful Firewall Administration Guide*.

Traffic Performance Optimization (TPO)

Though TCP is a widely accepted protocol in use today, it is optimized only for wired networks. Due to inherent reliability of wired networks, TCP implicitly assumes that any packet loss is due to network congestion and consequently invokes congestion control measures. However, wireless links are known to experience sporadic and

usually temporary losses due to several reasons, including the following, which also trigger TCP congestion control measures resulting in poor TCP performance.

Reasons for delay variability over wireless links include:

- Channel fading effect, subscriber mobility, and other transient conditions
- Link-layer retransmissions
- Handoffs between neighboring cells
- Intermediate nodes, such as SGSN and e-NodeB, implementing scheduling polices tuned to deliver better QoS for select services; resulting is variable delay in packet delivery for other services

The TPO inline service uses a combination of TCP and HTTP optimization techniques to improve TCP performance over wireless links.

Important: For more information on TPO, refer to the *Cisco ASR 5000 Series Traffic Performance Optimization Administration Guide*.

Features and Functionality - External Application Support

This section describes the features and functions of external applications supported on the PDSN. These services require additional licenses to implement the functionality.

Mobility Unified Reporting

The Cisco Mobility Unified Reporting (MUR) system is a Web-based application providing a unified reporting interface for diverse data from Cisco Systems In-line service and storage applications.

The MUR application provides comprehensive and consistent set of statistics and customized reports, report scheduling and distribution from ASR chassis / in-line service product. For example, a subscriber's Quality of Experience, top 10 sites visited, top 10 users, and so on.

The MUR application provides reporting capability for Content Filtering (CF) data, bulk statistics, Key Performance Indicators (KPIs), EDRs data from in-line service and storage applications. The MUR application facilitates and enhances the operators' ability to simply and easily determine the health and usage of the network.

Important: For more information on MUR support, refer to the *MUR Installation and Administration Guide*.

CDMA2000 Data Network Deployment Configurations

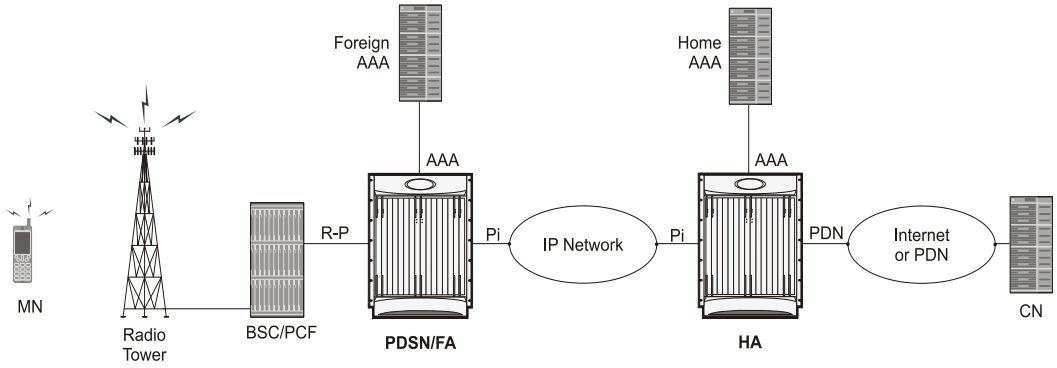
This section provides examples of how the system can be deployed within a wireless carrier's network. As noted previously in this chapter, the system can be deployed in standalone configurations, serving as a Packet Data Serving Node/Foreign Agent (PDSN/FA), a Home Agent (HA), or in a combined PDSN/FA/HA configuration providing all services from a single chassis. Although XT-2 systems are highly flexible, but XT-2 systems are pre-loaded with purchased services and operator can not add additional services through license. Operator needs to predefine the services required on a system.

Standalone PDSN/FA and HA Deployments

The PDSN/FA serves as an integral part of a CDMA2000 network by providing the packet processing and re-direction to the mobile user's home network through communications with the HA. In cases where the mobile user connects to a PDSN that serves their home network, no re-direction is required.

The following figure depicts a sample network configuration wherein the PDSN/FA and HA are separate systems.

Figure 65. PDSN/FA and HA Network Deployment Configuration Example



The HA allows mobile nodes to be reached, or served, by their home network through its home address even when the mobile node is not attached to its home network. The HA performs this function through interaction with an FA that the mobile node is communicating with using the Mobile IP protocol. Such transactions are performed through the use of virtual private networks that create Mobile IP tunnels between the HA and FA.

Interface Descriptions

This section describes the primary interfaces used in a CDMA2000 wireless data network deployment.

R-P Interface

This interface exists between the Packet Control Function (PCF) and the PDSN/FA and implements the A10 and A11 (data and bearer signaling respectively) protocols defined in 3GPP2 specifications.

The PCF can be co-located with the Base Station Controller (BSC) as part of the Radio Access Node (RAN). The PDSN/FA is connected to the RAN via Ethernet line cards installed in the rear of the chassis. The system supports either 8-port Fast Ethernet line cards (Ethernet 10/100) or single-port small form-factor pluggable (SFP) optical gigabit Ethernet line cards (Ethernet 1000) or four-port Quad Gig-E line cards (QGLC). These line cards also support outbound

IP traffic that carries user data to the HA for Mobile IP services, or to the Internet or Wireless Access Protocol (WAP) gateway for Simple IP services.

Pi Interfaces

The Pi interface provides connectivity between the HA and its corresponding FA. The Pi interface is used to establish a Mobile IP tunnels between the PDSN/FA and HA.

PDN Interfaces

PDN interface provide connectivity between the PDSN and/or HA to packet data networks such as the Internet or a corporate intranet.

AAA Interfaces

Using the LAN ports located on the Switch Processor I/O (SPIO) and Ethernet line cards, these interfaces carry AAA messages to and from RADIUS accounting and authentication servers. The SPIO supports RADIUS-capable management interfaces using either copper or fiber Ethernet connectivity through two auto-sensing 10/100/1000 Mbps Ethernet interfaces or two SFP optical gigabit Ethernet interfaces. User-based RADIUS messaging is transported using the Ethernet line cards.

While most carriers will configure separate AAA interfaces to allow for out-of-band RADIUS messaging for system administrative users and other operations personnel, it is possible to use a single AAA interface hosted on the Ethernet line cards to support a single RADIUS server that supports both management users and network users.

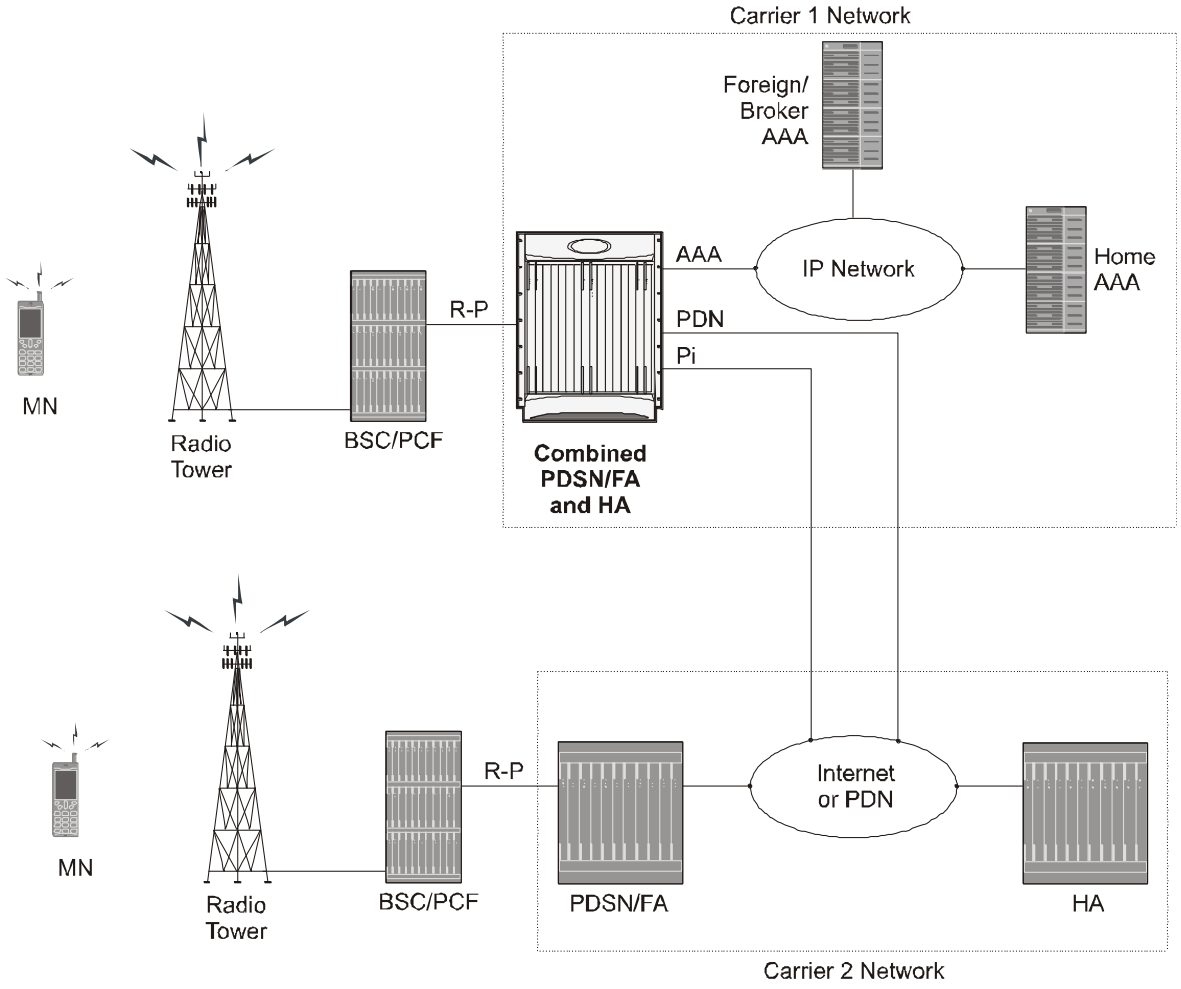
Important: Subscriber AAA interfaces should always be configured using Ethernet line card interfaces for the highest performance. The out-of-band local context should not be used for service subscriber AAA functions.

Co-Located Deployments

An advantage of the system is its ability to support both high-density PDSN/FA and HA configurations within the same chassis. The economies of scale presented in this configuration example provide for both improved session handling and reduced cost in deploying a CDMA2000 data network.

The following figure depicts a sample co-located deployment.

Figure 66. Co-located PDSN/FA and HA Configuration Example



It should be noted that all interfaces defined within the 3GPP2 standards for 1x deployments exist in this configuration as they are described in the two previous sections. This configuration can support communications to external, or standalone, PDSNs/FAs and/or HAs using all prescribed standards.

Understanding Simple IP and Mobile IP

From a mobile subscriber's perspective, packet data services are delivered from the service provider network using two access methods:

- Local and public network access
- Private network access

Within the packet data network, access is similar to accessing the public Internet through any other access device. In a private network access scenario, the user must be tunneled into the private network after initial authentication has been performed.

These two methods are provided using one of the following access applications:

- **Simple IP:** The mobile user is dynamically assigned an IP address from the service provider. The user can maintain this address within a defined geographical area, but when the user moves outside of this area, their IP address will be lost. This means that whenever a mobile user moves to a new location, they will need to re-register with the service provider to obtain a new IP address.
- **Mobile IP:** The mobile subscriber uses either a static or dynamically assigned IP address that belongs to their home network. As the subscriber roams through the network, the IP address is maintained providing the subscriber with the opportunity to use IP applications that require seamless mobility such as performing file transfers.
- **Proxy Mobile IP:** Provides a mobility solution for subscribers whose Mobile Nodes (MNs) do not support the Mobile IP protocol. The PDSN/FA proxy the Mobile IP tunnel with the HA on behalf of the MS. The subscriber receives an IP address from either the service provider or from their home network. As the subscriber roams through the network, the IP address is maintained providing the subscriber with the opportunity to use IP applications that require seamless mobility such as transferring files.

The following sections outline both Simple IP, Mobile IP, and Proxy Mobile IP and how they work in a 3G network.

Simple IP

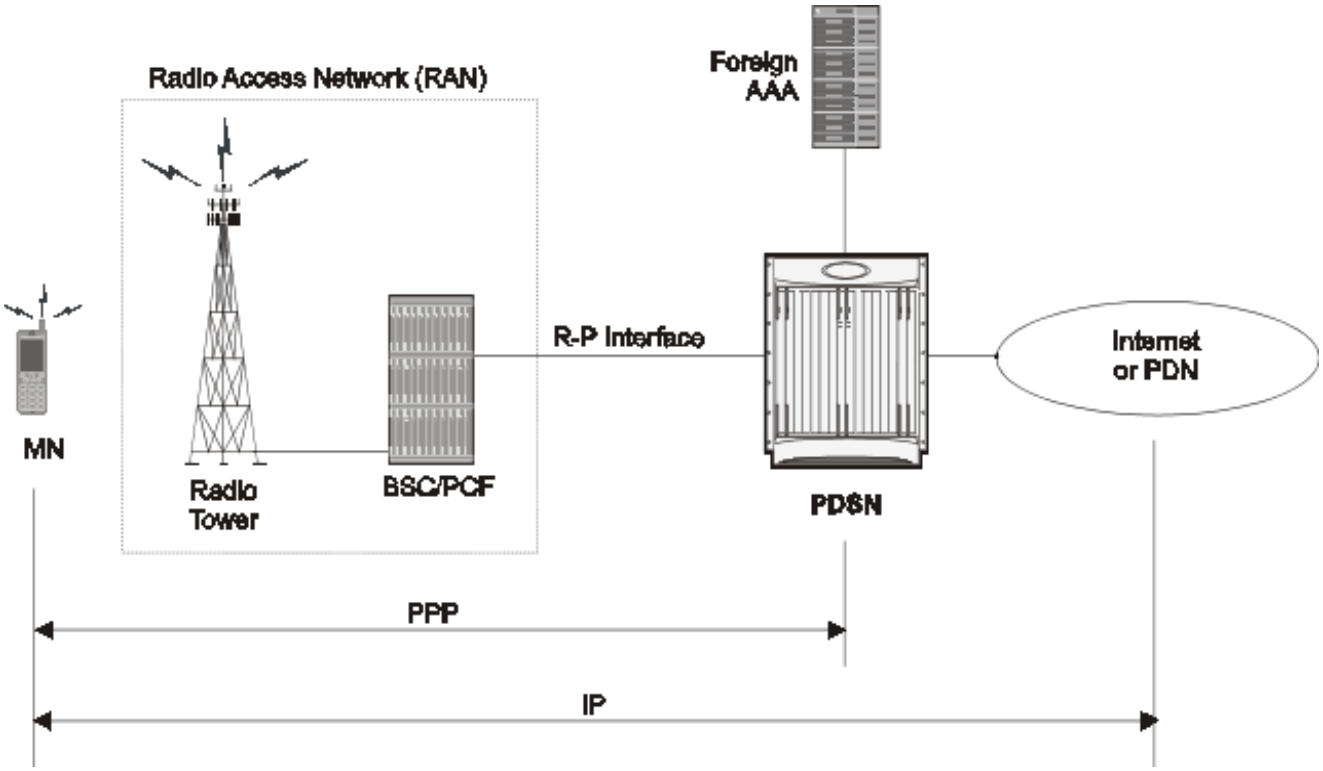
From a packet data perspective, Simple IP is similar to how a dial-up user would connect to the Internet using the Point-to-Point Protocol (PPP) and the Internet Protocol (IP) through an Internet Service Provider (ISP). With Simple IP, the mobile user is assigned a dynamic IP address from a PDSN or AAA server that is serving them locally (a specific geographic area). Once the mobile user is connected to the particular radio network that the assigning PDSN belongs to, an IP address is assigned to the mobile node. The PDSN provides IP routing services to the registered mobile user through the wireless service provider's network.

There is no mobility beyond the PDSN that assigns the dynamic IP address to the mobile user, which means that should the mobile user leave the geographic area where service was established (moves to a new radio network service area), they will need to obtain a new IP address with a new PDSN that is serving the new area. This new connection may or may not be provided by the same service provider.

How Simple IP Works

As described earlier, Simple IP uses two basic communications protocols, PPP and IP. The following figure depicts where each of these protocols are used in a Simple IP call.

Figure 67. Simple IP Protocol Usage



As depicted in the figure above, PPP is used to establish a communications session between the MN and the PDSN. Once a PPP session is established, the Mobile Node (MN) and end host communicate using IP packets.

The following figure and table provides a high-level view of the steps required to make a Simple IP call that is initiated by the MN to an end host. Users should keep in mind that steps 2, 3, 11, and 12 in the call flow are related to the Radio Access Node (RAN) functions and are intended to show a high-level overview of radio communications iterations, and as such are outside the scope of packet-based communications presented here.

Figure 68. Simple IP Call Flow

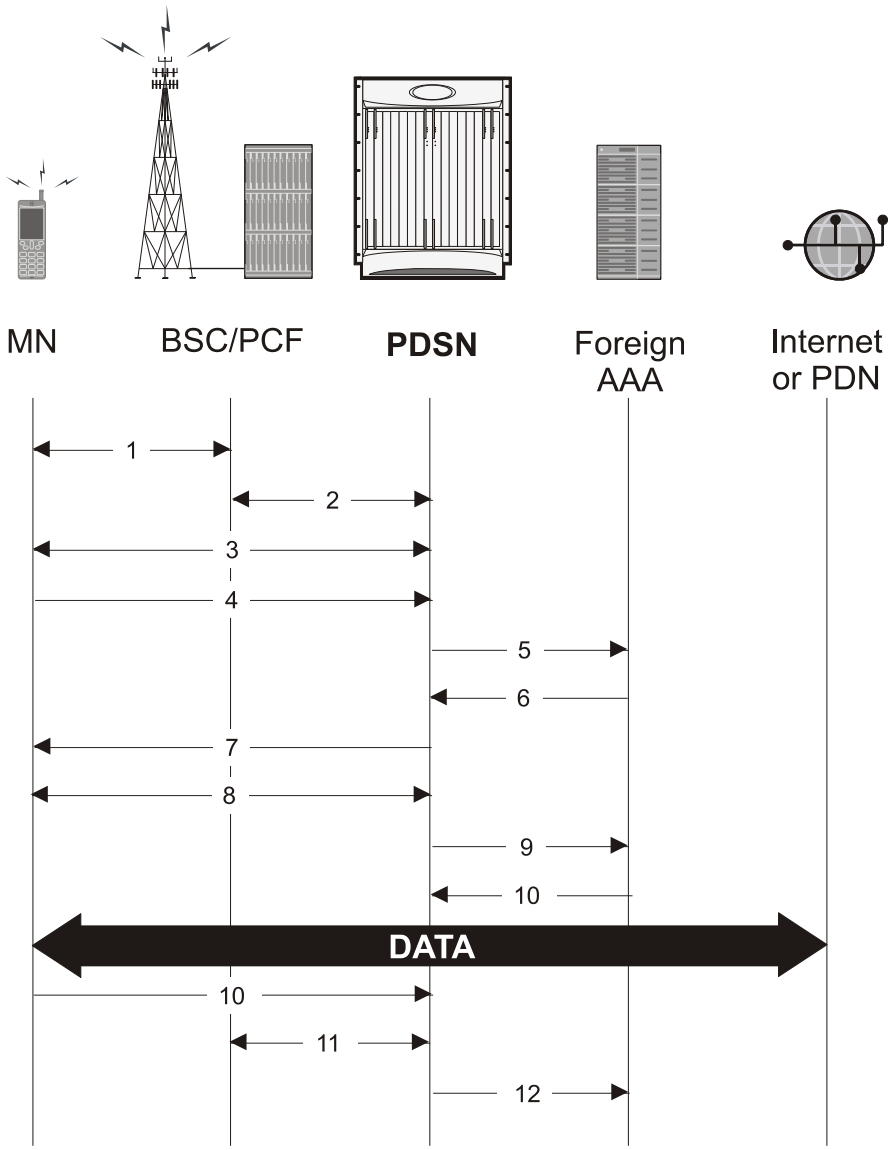


Table 51. Simple IP Call Flow Description

Step	Description
1	Mobile Node (MN) secures a traffic channel over the airlink with the RAN through the BSC/PCF.
2	The PCF and PDSN establish the R-P interface for the session.
3	The PDSN and MN negotiate Link Control Protocol (LCP).
4	Upon successful LCP negotiation, the MN sends a PPP Authentication Request message to the PDSN.
5	The PDSN sends an Access Request message to the RADIUS AAA server.

Step	Description
6	The RADIUS AAA server successfully authenticates the subscriber and returns an Access Accept message to the PDSN. The Accept message may contain various attributes to be assigned to the MN.
7	The PDSN sends a PPP Authentication Response message to the MN.
8	The MN and the PDSN negotiate the Internet Protocol Control Protocol (IPCP) that results in the MN receiving an IP address.
9	The PDSN forwards a RADIUS Accounting Start message to the AAA server fully establishing the session allowing the MN to send/receive data to/from the PDN.
10	Upon completion of the session, the MN sends an LCP Terminate Request message to the PDSN to end the PPP session.
11	The BSC closes the radio link while the PCF closes the R-P session between it and the PDSN. All PDSN resources used to facilitate the session are reclaimed (IP address, memory, etc.).
12	The PDSN sends accounting stop record to the AAA server, ending the session.

Mobile IP

Mobile IP provides a network-layer solution that allows mobile nodes (MNs, i.e. mobile phones, wireless PDAs, and other mobile devices) to receive routed IP packets from their home network while they are connected to any visitor network using their permanent or home IP address. Mobile IP allows mobility in a dynamic method that allows nodes to maintain ongoing communications while changing links as the user traverses the global Internet from various locations outside their home network.

In Mobile IP, the Mobile Node (MN) receives an IP address, either static or dynamic, called the “home address” assigned by its Home Agent (HA). A distinct advantage with Mobile IP is that MNs can hand off between different radio networks that are served by different PDSNs.

In this scenario, the PDSN in the visitor network performs as a Foreign Agent (FA), establishing a virtual session with the MN's HA. Each time the MN registers with a different PDSN/FA, the FA assigns the MN a care-of-address. Packets are then encapsulated into IP tunnels and transported between FA, HA, and the MN.

Mobile IP Tunneling Methods

Tunneling by itself is a technology that enables one network to send its data via another network's connections. Tunneling works by encapsulating a network protocol within a packet, carried by the second network. Tunneling is also called encapsulation. Service providers typically use tunneling for two purposes; first, to transport otherwise un-routable packets across the IP network and second, to provide data separation for Virtual Private Networking (VPN) services. In Mobile IP, tunnels are used to transport data packets between the FA and HA.

The system supports the following tunneling protocols, as defined in the IS-835-A specification and the relevant Request For Comments (RFCs) for Mobile IP:

IP in IP tunnels

IP in IP tunnels basically encapsulate one IP packet within another using a simple encapsulation technique. To encapsulate an IP datagram using IP in IP encapsulation, an outer IP header is inserted before the datagram's existing IP header. Between them are other headers for the path, such as security headers specific to the tunnel configuration. Each header chains to the next using IP Protocol values. The outer IP header Source and Destination identify the “endpoints” of the tunnel. The inner IP header Source and Destination identify the original sender and recipient of the datagram,

while the inner IP header is not changed by the encapsulator, except to decrement the TTL, and remains unchanged during its delivery to the tunnel exit point. No change to IP options in the inner header occurs during delivery of the encapsulated datagram through the tunnel. If needed, other protocol headers such as the IP Authentication header may be inserted between the outer IP header and the inner IP header.

The Mobile IP working group has specified the use of encapsulation as a way to deliver datagrams from an MN's HA to an FA, and conversely from an FA to an HA, that can deliver the data locally to the MN at its current location.

GRE tunnels

The Generic Routing Encapsulation (GRE) protocol performs encapsulation of IP packets for transport across disparate networks. One advantage of GRE over earlier tunneling protocols is that any transport protocol can be encapsulated in GRE. GRE is a simple, low overhead approach—the GRE protocol itself can be expressed in as few as eight octets as there is no authentication or tunnel configuration parameter negotiation. GRE is also known as IP Protocol 47.

Important: The chassis simultaneously supports GRE protocols with key in accordance with RFC-1701/RFC-2784 and "Legacy" GRE protocols without key in accordance to RFC-2002.

Another advantage of GRE tunneling over IP-in-IP tunneling is that GRE tunneling can be used even when conflicting addresses are in use across multiple contexts (for the tunneled data).

Communications between the FA and HA can be done in either the forward or reverse direction using the above protocols. Additionally, another method of routing information between the FA and various content servers used by the HA exists. This method is called Triangular Routing. Each of these methods is explained below.

Forward Tunneling

In the wireless IP world, forward tunneling is a tunnel that transports packets from the packet data network towards the MN. It starts at the HA and ends at the MN's care-of address. Tunnels can be as simple as IP-in-IP tunnels, GRE tunnels, or even IP Security (IPSec) tunnels with encryption. These tunnels can be started automatically, and are selected based on the subscriber's user profile.

Reverse Tunneling

A reverse tunnel starts at the MN's care-of address, which is the FA, and terminates at the HA.

When an MN arrives at a foreign network, it listens for agent advertisements and selects an FA that supports reverse tunnels. The MN requests this service when it registers through the selected FA. At this time, the MN may also specify a delivery technique such as Direct or the Encapsulating Delivery Style.

Using the Direct Delivery Style, which is the default mode for the system, the MN designates the FA as its default router and sends packets directly to the FA without encapsulation. The FA intercepts them, and tunnels them to the HA.

Using the Encapsulating Delivery Style, the MN encapsulates all its outgoing packets to the FA. The FA then de-encapsulates and re-tunnels them to the HA, using the FA's care-of address as the entry-point for this new tunnel.

Following are some of the advantages of reverse tunneling:

- All datagrams from the mobile node seem to originate from its home network
- The FA can keep track of the HA that the mobile node is registered to and tunnel all datagrams from the mobile node to its HA

Triangular Routing

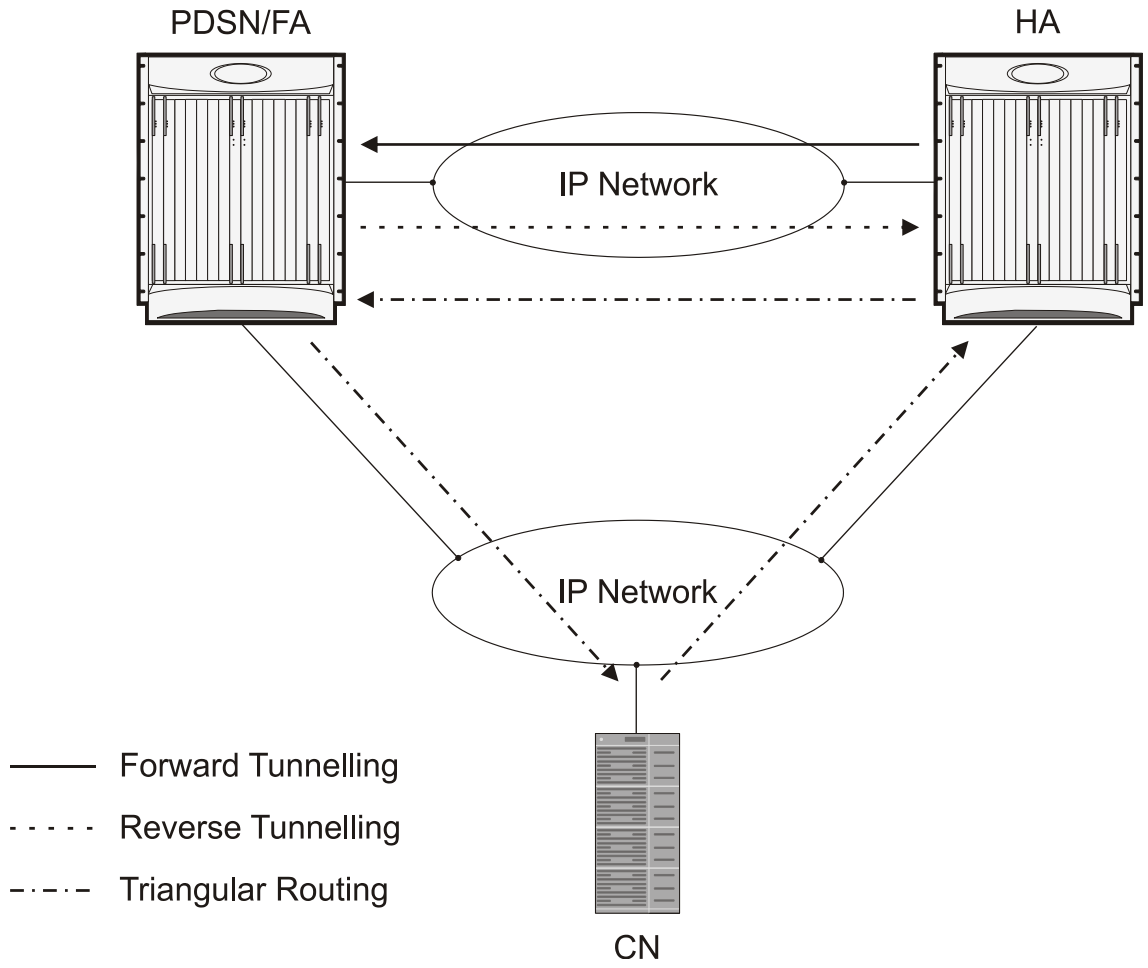
Triangular routing is the path followed by a packet from the MN to the Correspondent Node (CN) via the FA. In this routing scenario, the HA receives all the packets destined to the MN from the CN and redirects them to the MN's care-of-address by forward tunneling. In this case, the MN sends packets to the FA, which are transported using conventional IP routing methods.

A key advantage of triangular routing is that reverse tunneling is not required, eliminating the need to encapsulate and de-capsulate packets a second time during a Mobile IP session since only a forward tunnel exists between the HA and PDSN/FA.

A disadvantage of using triangular routing is that the HA is unaware of all user traffic for billing purposes. Also, both the HA and FA are required to be connected to a private network. This can be especially troublesome in large networks, serving numerous enterprise customers, as each FA would have to be connected to each private network.

The following figure shows an example of how triangular routing is performed.

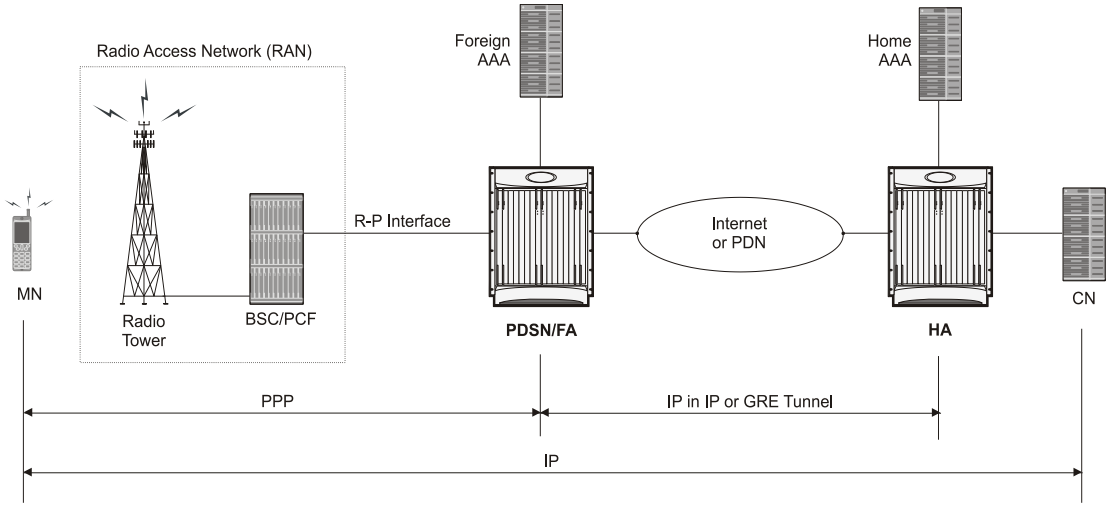
Figure 69. Mobile IP, FA and HA Tunneling/Transport Methods



How Mobile IP Works

As described earlier, Mobile IP uses three basic communications protocols; PPP, IP, and Tunneled IP in the form of IP-in-IP or GRE tunnels. The following figure depicts where each of these protocols are used in a basic Mobile IP call.

Figure 70. Mobile IP Protocol Usage



As depicted in the figure above, PPP is used to establish a communications session between the MN and the FA. Once a PPP session is established, the MN can communicate with the HA, using the FA as a mediator or broker. Data transport between the FA and HA use tunneled IP, either IP-in-IP or GRE tunneling. Communication between the HA and End Host can be achieved using the Internet or a private IP network and can use any IP protocol.

The following figure provides a high-level view of the steps required to make a Mobile IP call that is initiated by the MN to a HA and table that follows, explains each step in detail. Users should keep in mind that steps in the call flow related to the Radio Access Node (RAN) functions are intended to show a high-level overview of radio communications iterations, and as such are outside the scope of packet-based communications presented here.

Figure 71. Mobile IP Call Flow

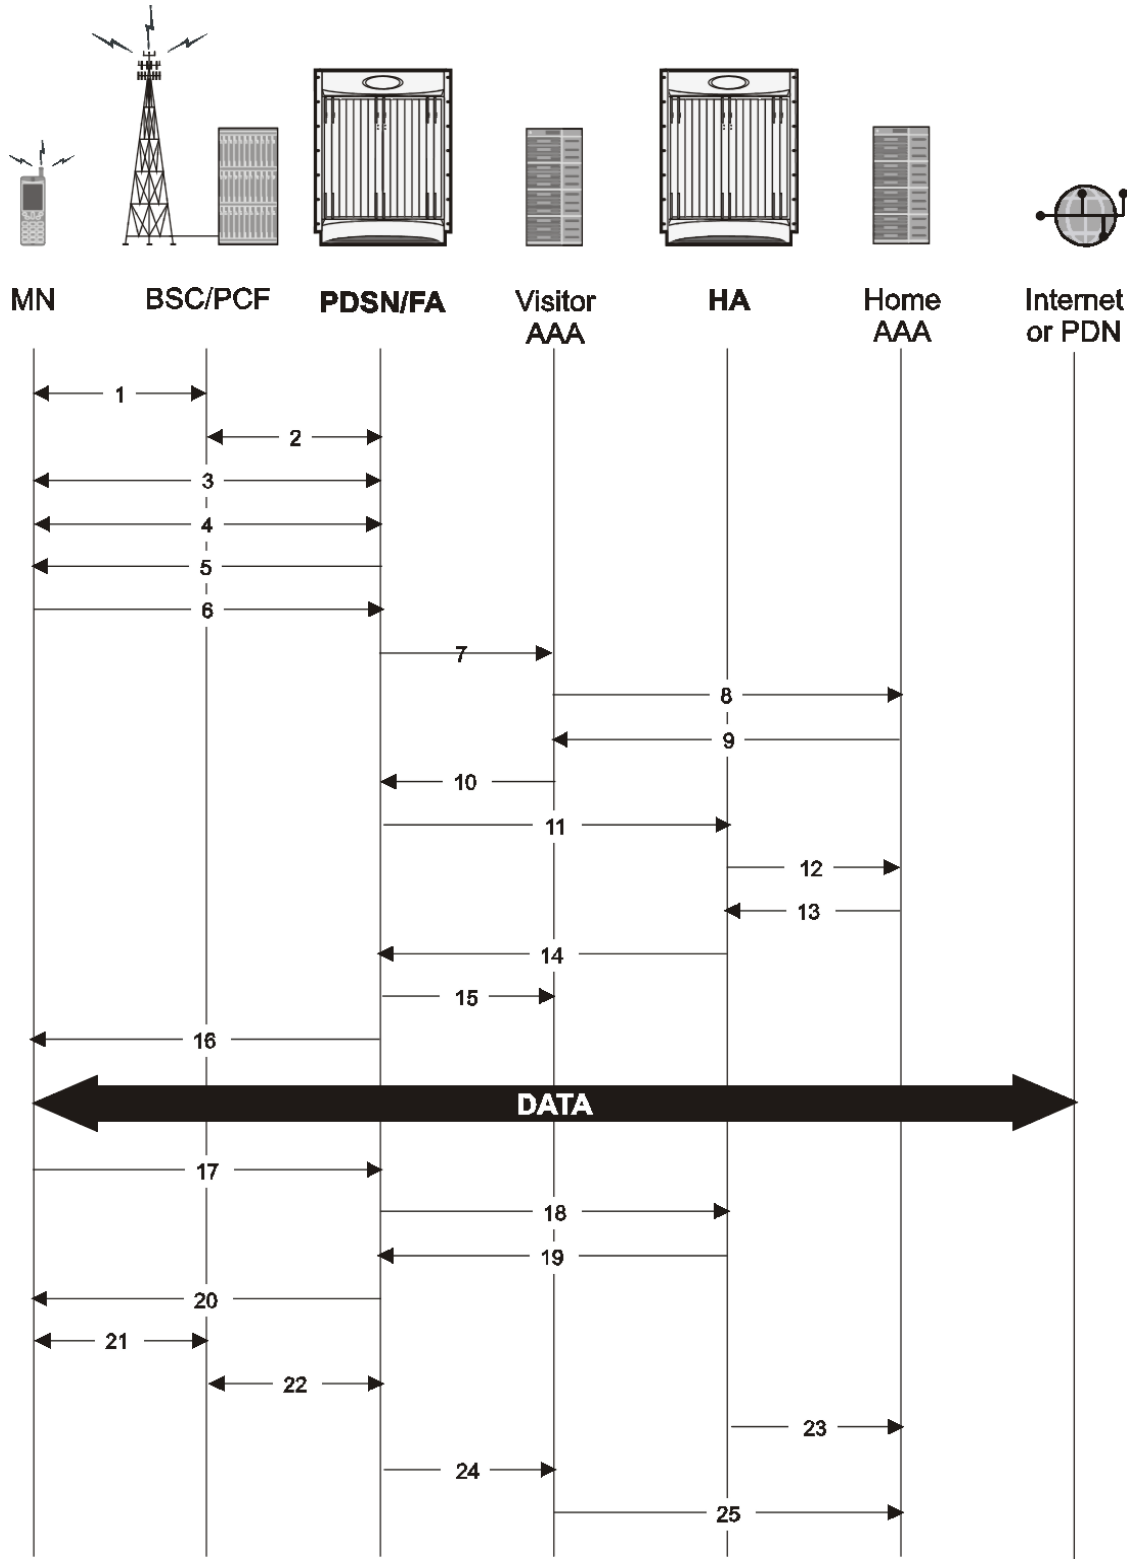


Table 52. Mobile IP Call Flow Description

Step	Description
1	Mobile Node (MN) secures a traffic channel over the airlink with the RAN through the BSC/PCF.
2	The PCF and PDSN establish the R-P interface for the session.
3	The PDSN and MN negotiate Link Control Protocol (LCP).
4	The PDSN and MN negotiate the Internet Protocol Control Protocol (IPCP).
5	The PDSN/FA sends an Agent Advertisement to the MN.
6	The MN sends a Mobile IP Registration Request to the PDSN/FA.
7	The PDSN/FA sends an Access Request message to the visitor AAA server.
8	The visitor AAA server proxies the request to the appropriate home AAA server.
9	The home AAA server sends an Access Accept message to the visitor AAA server.
10	The visitor AAA server forwards the response to the PDSN/FA.
11	Upon receipt of the response, the PDSN/FA forwards a Mobile IP Registration Request to the appropriate HA.
12	The HA sends an Access Request message to the home AAA server to authenticate the MN/subscriber.
13	The home AAA server returns an Access Accept message to the HA.
14	Upon receiving response from home AAA, the HA sends a reply to the PDSN/FA establishing a forward tunnel. Note that the reply includes a Home Address (an IP address) for the MN.
15	The PDSN/FA sends an Accounting Start message to the visitor AAA server. The visitor AAA server proxies messages to the home AAA server as needed.
16	The PDSN return a Mobile IP Registration Reply to the MN establishing the session allowing the MN to send/receive data to/from the PDN.
17	Upon session completion, the MN sends a Registration Request message to the PDSN/FA with a requested lifetime of 0.
18	The PDSN/FA forwards the request to the HA.
19	The HA sends a Registration Reply to the PDSN/FA accepting the request.
20	The PDSN/FA forwards the response to the MN.
21	The MN and PDSN/FA negotiate the termination of LCP effectively ending the PPP session.
22	The PCF and PDSN/FA close terminate the R-P session.
23	The HA sends an Accounting Stop message to the home AAA server.
24	The PDSN/FA sends an Accounting Stop message to the visitor AAA server.
25	The visitor AAA server proxies the accounting data to the home AAA server.

Proxy Mobile IP

Proxy Mobile IP provides mobility for subscribers with MNs that do not support the Mobile IP protocol stack.

For subscriber sessions using Proxy Mobile IP, R-P and PPP sessions get established as they would for a Simple IP session. However, the PDSN/FA performs Mobile IP operations with an HA (identified by information stored in the subscriber's profile) on behalf of the MN while the MN performs only Simple IP processes. The protocol details are similar to those displayed in figure earlier for Mobile IP.

The MN is assigned an IP address by either the PDSN/FA or the HA. Regardless of its source, the address is stored in a Mobile Binding Record (MBR) stored on the HA. Therefore, as the MN roams through the service provider's network, each time a hand-off occurs, the MN will receive the same IP address stored in the MBR on the HA.

Note that unlike Mobile IP-capable MNs that can perform multiple sessions over a single PPP link, Proxy Mobile IP allows only a single session over the PPP link. In addition, simultaneous Mobile and Simple IP sessions will not be supported for an MN by an FA currently facilitating a Proxy Mobile IP session for the MN.

How Proxy Mobile IP Works

This section contains call flows displaying successful Proxy Mobile IP session setup scenarios. Two scenarios are described based on how the MN receives an IP address:

- **Scenario 1:** The AAA server specifies an IP address that the PDSN allocates to the MN from one of its locally configured static pools.
- **Scenario 2:** The HA assigns an IP address to the MN from one of its locally configured dynamic pools.

Scenario 1: AAA server and PDSN/FA Allocate IP Address

The following figure and table display and describe a call flow in which the MN receives its IP address from the AAA server and PDSN/FA.

Figure 72. AAA/PDSN Assigned IP Address Proxy Mobile IP Call Flow

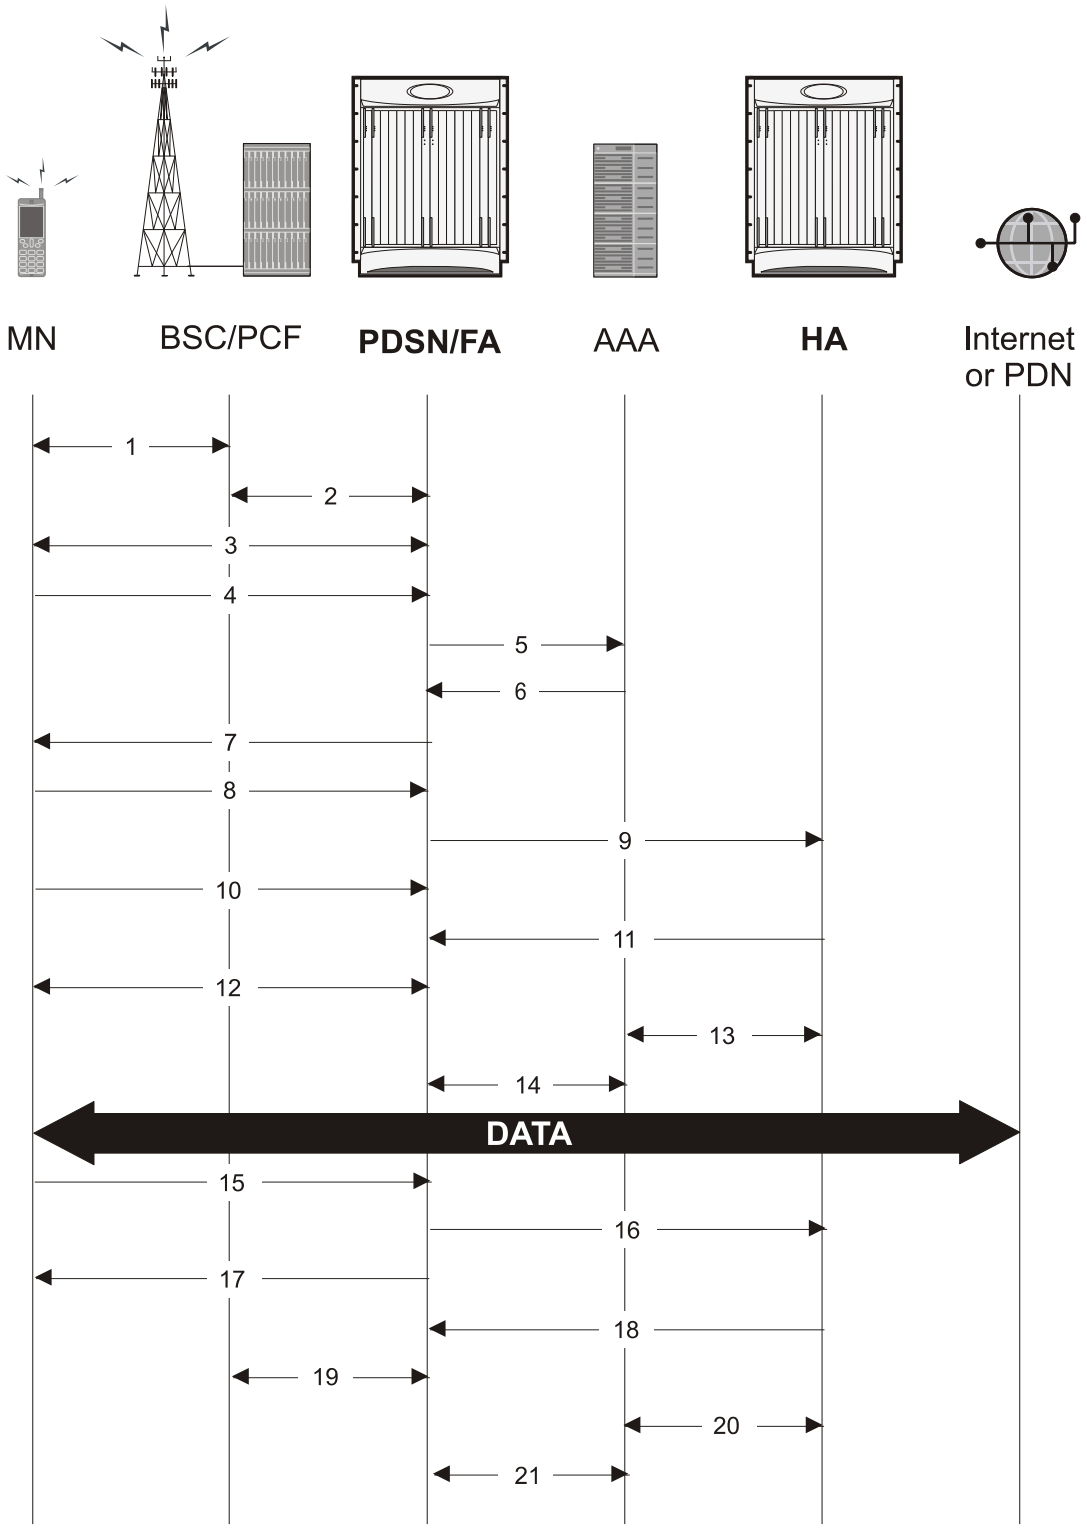


Table 53. AAA/PDSN Assigned IP Address Proxy Mobile IP Call Flow Description

Step	Description
1	Mobile Node (MN) secures a traffic channel over the airlink with the RAN through the BSC/PCF.
2	The PCF and PDSN/FA establish the R-P interface for the session.
3	The PDSN/FA and MN negotiate Link Control Protocol (LCP).
4	Upon successful LCP negotiation, the MN sends a PPP Authentication Request message to the PDSN/FA.
5	The PDSN/FA sends an Access Request message to the RADIUS AAA server.
6	The RADIUS AAA server successfully authenticates the subscriber and returns an Access Accept message to the PDSN/FA. The Accept message may contain various attributes to be assigned to the MN including the MN's Home Address (IP address) and the IP address of the HA to use.
7	The PDSN/FA sends a PPP Authentication Response message to the MN.
8	The MN sends an Internet Protocol Control Protocol (IPCP) Configuration Request message to the PDSN/FA with an MN address of 0.0.0.0.
9	The PDSN/FA forwards a Proxy Mobile IP Registration Request message to the HA. The message includes such things as the MN's home address, the IP address of the FA (the care-of-address), and the FA-HA extension (security parameter index (SPI)).
10	While the FA is communicating with the HA, the MN may send additional IPCP Configuration Request messages.
11	The HA responds with a Proxy Mobile IP Registration Response after validating the home address against it's pool(s). The HA also creates a Mobile Binding Record (MBR) for the subscriber session.
12	The MN and the PDSN/FA negotiate IPCP. The result is that the MN is assigned the home address originally specified by the AAA server.
13	While the MN and PDSN/FA are negotiating IPCP, the HA and AAA server initiate accounting.
14	Upon completion of the IPCP negotiation, the PDSN/FA and AAA server initiate accounting fully establishing the session allowing the MN to send/receive data to/from the PDN.
15	Upon completion of the session, the MN sends an LCP Terminate Request message to the PDSN to end the PPP session.
16	The PDSN/FA sends a Proxy Mobile IP De-registration Request message to the HA.
17	The PDSN/FA send an LCP Terminate Acknowledge message to the MN ending the PPP session.
18	The HA sends a Proxy Mobile IP De-Registration Response message to the FA terminating the Pi interface
19	The PDSN/FA and the PCF terminate the R-P session.
20	The HA and the AAA server stop accounting for the session.
21	The PDSN and the AAA server stop accounting for the session.

Scenario 2: HA Assigns IP Address to MN from Locally Configured Dynamic Pools

The following figure and table display and describe a call flow in which the MN receives its IP address from the AAA server and PDSN/FA.

Figure 73. HA Assigned IP Address Proxy Mobile IP Call Flow

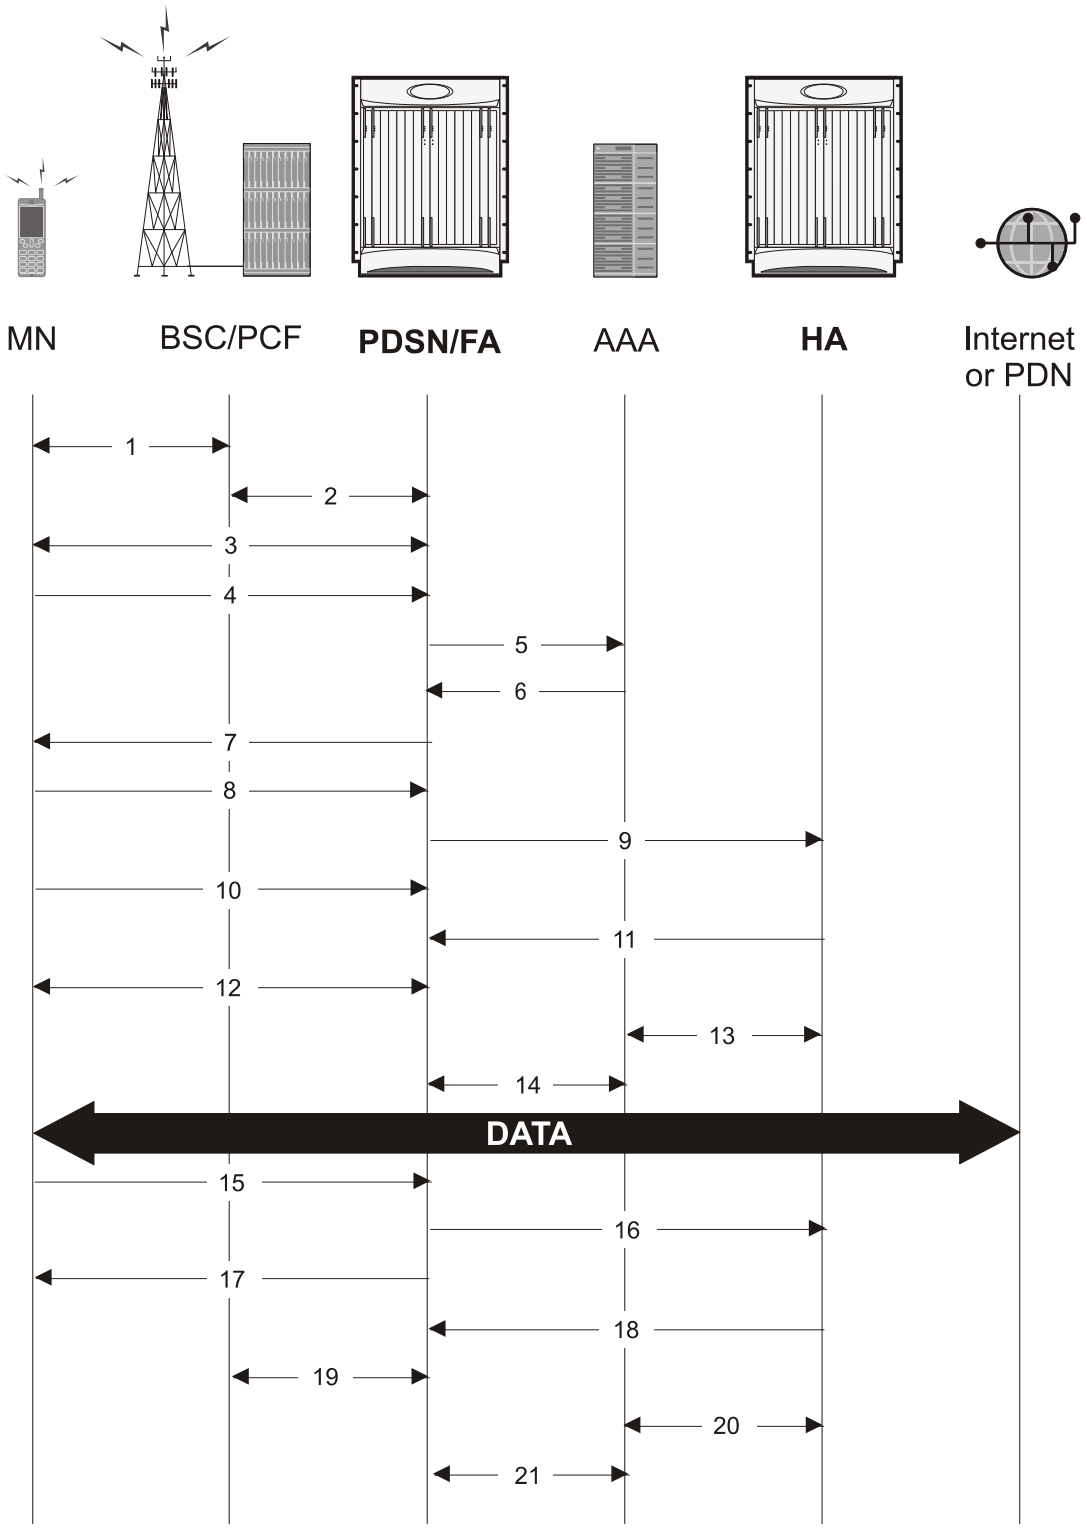


Table 54. HA Assigned IP Address Proxy Mobile IP Call Flow Description

Step	Description
1	Mobile Node (MN) secures a traffic channel over the airlink with the RAN through the BSC/PCF.
2	The PCF and PDSN/FA establish the R-P interface for the session.
3	The PDSN/FA and MN negotiate Link Control Protocol (LCP).
4	Upon successful LCP negotiation, the MN sends a PPP Authentication Request message to the PDSN/FA.
5	The PDSN/FA sends an Access Request message to the RADIUS AAA server.
6	The RADIUS AAA server successfully authenticates the subscriber and returns an Access Accept message to the PDSN/FA. The Accept message may contain various attributes to be assigned to the MN including the IP address of the HA to use.
7	The PDSN/FA sends a PPP Authentication Response message to the MN.
8	The MN sends an Internet Protocol Control Protocol (IPCP) Configuration Request message to the PDSN/FA with an MN address of 0.0.0.0.
9	The PDSN/FA forwards a Proxy Mobile IP Registration Request message to the HA. The message includes such things as a Home Address indicator of 0.0.0.0, the IP address of the FA (the care-of-address), the IP address of the FA (the care-of-address), and the FA-HA extension (Security Parameter Index (SPI)).
10	While the FA is communicating with the HA, the MN may send additional IPCP Configuration Request messages.
11	The HA responds with a Proxy Mobile IP Registration Response. The response includes an IP address from one of its locally configured pools to assign to the MN (its Home Address). The HA also creates a Mobile Binding Record (MBR) for the subscriber session.
12	The MN and the PDSN/FA negotiate IPCP. The result is that the MN is assigned the home address originally specified by the AAA server.
13	While the MN and PDSN/FA are negotiating IPCP, the HA and AAA server initiate accounting.
14	Upon completion of the IPCP negotiation, the PDSN/FA and AAA server initiate accounting fully establishing the session allowing the MN to send/receive data to/from the PDN.
15	Upon completion of the session, the MN sends an LCP Terminate Request message to the PDSN to end the PPP session.
16	The PDSN/FA sends a Proxy Mobile IP De-registration Request message to the HA.
17	The PDSN/FA send an LCP Terminate Acknowledge message to the MN ending the PPP session.
18	The HA sends a Proxy Mobile IP De-Registration Response message to the FA terminating the Pi interface
19	The PDSN/FA and the PCF terminate the R-P session.
20	The HA and the AAA server stop accounting for the session.
21	The PDSN and the AAA server stop accounting for the session.

Supported Standards

The system supports the following industry standards for 1x/CDMA2000/EV-DO devices.

Requests for Comments (RFCs)

- RFC-768, User Datagram Protocol (UPD), August 1980
- RFC-791, Internet Protocol (IP), September 1982
- RFC-793, Transmission Control Protocol (TCP), September 1981
- RFC-894, A Standard for the Transmission of IP Datagrams over Ethernet Networks, April 1984
- RFC-1089, SNMP over Ethernet, February 1989
- RFC-1144, Compressing TCP/IP headers for low-speed serial links, February 1990
- RFC-1155, Structure and Identification of Management Information for TCP/IP-based Internets, May 1990
- RFC-1157, Simple Network Management Protocol (SNMP) Version 1, May 1990
- RFC-1212, Concise MIB Definitions, March 1991
- RFC-1213, Management Information Base for Network Management of TCP/IP-based Internets: MIB-II, March 1991
- RFC-1215, A Convention for Defining Traps for use with the SNMP, March 1991
- RFC-1224, Techniques for Managing Asynchronously Generated Alerts, May 1991
- RFC-1256, ICMP Router Discovery Messages, September 1991
- RFC-1305, Network Time Protocol (Version 3) Specification, Implementation and Analysis, March 1992
- RFC-1332, The PPP Internet Protocol Control Protocol (IPCP), May 1992
- RFC-1398, Definitions of Managed Objects for the Ethernet-Like Interface Types, January 1993
- RFC-1418, SNMP over OSI, March 1993
- RFC-1570, PPP LCP Extensions, January 1994
- RFC-1643, Definitions of Managed Objects for the Ethernet-like Interface Types, July 1994
- RFC-1661, The Point to Point Protocol (PPP), July 1994
- RFC-1662, PPP in HDLC-like Framing, July 1994
- RFC-1701, Generic Routing Encapsulation (GRE), October 1994
- RFC-1771, A Border Gateway Protocol 4 (BGP-4)
- RFC-1850, OSPF Version 2 Management Information Base, November 1995
- RFC-1901, Introduction to Community-based SNMPv2, January 1996
- RFC-1902, Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1903, Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996

- RFC-1904, Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1905, Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1906, Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1907, Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1908, Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework, January 1996
- RFC-1918, Address Allocation for Private Internets, February 1996
- RFC-1919, Classical versus Transparent IP Proxies, March 1996
- RFC-1962, The PPP Compression Control Protocol (CCP), June 1996
- RFC-1974, PPP STAC LZS Compression Protocol, August 1996
- RFC-2002, IP Mobility Support, May 1995
- RFC-2003, IP Encapsulation within IP, October 1996
- RFC-2004, Minimal Encapsulation within IP, October 1996
- RFC-2005, Applicability Statement for IP Mobility Support, October 1996
- RFC-2118, Microsoft Point-to-Point Compression (MPPC) Protocol, March 1997
- RFC-2136, Dynamic Updates in the Domain Name System (DNS UPDATE)
- RFC-2211, Specification of the Controlled-Load Network Element Service
- RFC-2246, The Transport Layer Security (TLS) Protocol Version 1.0, January 1999
- RFC-2290, Mobile IPv4 Configuration Option for PPP IPCP, February 1998
- RFC-2328, OSPF Version 2, April 1998
- RFC-2344, Reverse Tunneling for Mobile IP, May 1998
- RFC-2394, IP Payload Compression Using DEFLATE, December 1998
- RFC-2401, Security Architecture for the Internet Protocol, November 1998
- RFC-2402, IP Authentication Header (AH), November 1998
- RFC-2406, IP Encapsulating Security Payload (ESP), November 1998
- RFC-2408, Internet Security Association and Key Management Protocol (ISAKMP), November 1998
- RFC-2409, The Internet Key Exchange (IKE), November 1998
- RFC-2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, December 1998
- RFC-2475, An Architecture for Differentiated Services, December 1998
- RFC-2484, PPP LCP Internationalization Configuration Option, January 1999
- RFC-2486, The Network Access Identifier (NAI), January 1999
- RFC-2571, An Architecture for Describing SNMP Management Frameworks, April 1999
- RFC-2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP), April 1999

- RFC-2573, SNMP Applications, April 1999
- RFC-2574, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3), April 1999
- RFC-2597, Assured Forwarding PHB Group, June 1999
- RFC2598 - Expedited Forwarding PHB, June 1999
- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2661, Layer Two Tunneling Protocol "L2TP", August 1999
- RFC-2697, A Single Rate Three Color Marker, September 1999
- RFC-2698, A Two Rate Three Color Marker, September 1999
- RFC-2784, Generic Routing Encapsulation (GRE) - March 2000, IETF
- RFC-2794, Mobile IP Network Access Identifier Extension for IPv4, March 2000
- RFC-2809, Implementation of L2TP Compulsory Tunneling via RADIUS, April 2000
- RFC-2845, Secret Key Transaction Authentication for DNS (TSIG), May 2000
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000
- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000
- RFC-3007, Secure Domain Name System (DNS) Dynamic Update, November 2000
- RFC-3012, Mobile IPv4 Challenge/Response Extensions, November 2000
- RFC-3095, Robust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP and uncompressed, July 2001
- RFC-3101, OSPF NSSA Option, January 2003.
- RFC-3141, CDMA2000 Wireless Data Requirements for AAA, June 2001
- RFC-3143, Known HTTP Proxy/Caching Problems, June 2001
- RFC-3193, Securing L2TP using IPSEC, November 2001
- RFC-3241 Robust Header Compression (ROHC) over PPP, April 2002
- RFC-3409, Lower Layer Guidelines for Robust (RTP/UDP/IP) Header Compression, December 2002
- RFC-3519, NAT Traversal for Mobile IP, April 2003
- RFC-3543, Registration Revocation in Mobile IPv4, August 2003
- RFC 3576 - Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS), July 2003
- RFC-3706, A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers, February 2004
- RFC-3759, Robust Header Compression (ROHC): Terminology and Channel Mapping Examples, April 2004
- RFC-3588, Diameter Based Protocol, September 2003
- RFC-4005, Diameter Network Access Server Application, August 2005
- RFC-4006, Diameter Credit-Control Application, August 2005

- Draft, Generalized Key Distribution Extensions for Mobile IP
- Draft, AAA Keys for Mobile IP

TIA and Other Standards

Telecommunications Industry Association (TIA) Standards

- TIA/EIA/IS-835-A, CDMA2000 Wireless IP Network Standard, April 2001
- TIA/EIA/IS-835-B, CDMA2000 Wireless IP Network Standard, October 2002
- TIA/EIA/IS-835-C, CDMA2000 Wireless IP Network Standard, August 2003
- TIA/EIA/IS-707-A-1, Data Service Options for Wideband Spread Spectrum Systems
- TIA/EIA/IS-707-A.5 Packet Data Services
- TIA/EIA/IS-707-A.9 High Speed Packet Data Services
- TIA/EIA/IS-2000.5, Upper Layer (Layer 3) Signaling for CDMA2000 Spread Spectrum Systems
- TIA/EIA/IS-2001, Interoperability Specifications (IOS) for CDMA2000 Access Network Interfaces
- TIA/EIA/TSB100, Wireless Network Reference Model
- TIA/EIA/TSB115, CDMA2000 Wireless IP Architecture Based on IETF Protocols
- TIA/EIA J-STD-025 PN4465, TR-45 Lawfully Authorized Electronic Surveillance

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

3GPP2 Standards

- 3GPP2 A.S0001-A v2: 3GPP2 Access Network Interfaces Interoperability Specification (also known as 3G-IOS v4.1.1)
- 3GPP2 P.S0001-A-3: Wireless IP Network Standard
- 3GPP2 P.S0001-B: Wireless IP Network Standard
- 3GPP2 S.R0068: Link Layer Assisted Robust Header Compression
- [9] 3GPP2 C.S0047-0: Link Layer Assisted Service Options for Voice-over-IP: Header Removal (SO60) and Robust Header Compression (SO61)
- 3GPP2 A.S0008 v3.0 Interoperability Specification (IOS) for High Rate Packet Data (HRPD) Access Network Interfaces
- 3GPP2 A.S0015-0 v2: Interoperability Specification (IOS) for CDMA2000 1 1 Access Network Interfaces — Part 5 (A3 and A7 1 2 Interfaces) (Partial Support) (also know as 3G-IOSv4.2)
- 3GPP2 P.S0001-B V1.0.0 Wireless IP Network Standard October 25, 2002 (relating to MIP interactions with IPSEC)
- 3GPP2 P.S0001 (TIA/EIA/IS-835-1) Version 1.0, Wireless IP Network Standard - December 10, 1999
- 3GPP2 P.R0001 (TSB115) Version 1.0.0, Wireless IP: Architecture Based on IETF Protocols - July 14, 2000

- 3GPP2 3GPP2 X.S0011-005-C Version: 1.0.0, CDMA2000 Wireless IP Network Standard: Accounting Services and 3GPP2 RADIUS VSAs - August 2003
- 3GPP2 X.S0011-006-C Version: 1.0.0, CDMA2000 Wireless IP Network Standard: PrePaid Packet Data Service - Date: August 2003
- 3GPP2 TSGA A.S0013-c v0.4 Interoperability Specification (IOS) for CDMA2000 June 2004
- 3GPP2 TSG-A A.S.0017-C baseline Interoperability Specification (IOS) for CDMA2000 Access Network Interfaces - Part 7(A10 and A11 Interfaces) (IOS v5.0 baseline) June 2004
- 3GPP2 A.S0012-D Segmentation for GRE January, 2005
- Inter-operability Specification (IOS) for CDMA2000 Access Network Interfaces
- 3GPP2 X.S0011-005-D Accounting Services and 3GPP2 RADIUS VSAs, February 2006
- 3GPP2 TSG-X (PSN) X.P0013-014-0, Service Based Bearer Control – Ty Interface Stage-3

IEEE Standards

- 802.1Q VLAN Standard

Chapter 11
Content Filtering Support Overview

This chapter provides an overview of the Content Filtering In-line Service feature.

This chapter covers the following topics:

- Introduction
- Platform Requirements
- Licenses Requirements
- URL Blacklisting Support
- Category-based Content Filtering Support
- Content Filtering Server Group Support
- External Storage System
- Bulk Statistics Support
- Minimum System Requirements and Recommendations

Introduction

Content Filtering is an in-line service available for 3GPP and 3GPP2 networks to filter HTTP and WAP requests from mobile subscribers based on the URLs in the requests. This enables operators to filter and control the content that an individual subscriber can access, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences.

The CF in-line service works in conjunction with the following products:

- GGSN
- HA
- IPSG
- PDSN
- P-GW

The Content Filtering service offers the following solutions:

- URL Blacklisting:

  In the URL Blacklisting solution, all HTTP/WAP URLs in subscriber requests are matched against a database of "blacklisted" URLs. If there is a match, the flow is discarded, redirected, or terminated as configured. If there is no match, subscribers view the content as they would normally.

  URL Blacklisting may/may not be a subscriber opt-in service, operators can enable URL Blacklisting either for all subscribers or for a subset of subscribers. Typical cases include applying a blacklisted database of child porn URLs to all subscribers so that they are inadvertently not exposed to such universally unacceptable content.

- Category-based Static Content Filtering:

  In Category-based Static Content Filtering, all HTTP/WAP URLs in subscriber requests are matched against a static URL categorization database. Action is taken based on a URL's category, and the action configured for that category in the subscriber's content filtering policy. Possible actions include permitting, blocking, redirecting, and inserting content.

This release supports the following types of Category-based Content Filtering:

- Category-based Static Content Filtering:

  In Category-based Static Content Filtering, all HTTP/WAP URLs in subscriber requests are matched against a static URL categorization database. Action is taken based on a URL's category, and the action configured for that category in the subscriber's content filtering policy. Possible actions include permitting, blocking, redirecting, and inserting content.

- Category-based Static-and-Dynamic Content Filtering:

  In Category-based Static-and-Dynamic Content Filtering, if static rating categorizes a URL as either "dynamic" or "unknown", the "requested content" is sent for dynamic rating. Wherein the "requested content" is analyzed and categorized. Action is taken based on the category determined by dynamic rating, and the action configured for that category in the subscriber's content filtering policy. Possible actions include permitting, blocking, redirecting, and inserting content.

Typically Category-based Content Filtering is an opt-in service, subscribers self-choose a content-filtering policy or plan, such as Teen, Child, Adult, etc., and are subjected to content filtering as per their chosen plan. Also, the content-filtering policies of different subscribers may be different, enabling differential access of content to them. This solution provides maximum flexibility, and is also referred to as the Policy-based Content Filtering.

Both URL Blacklisting and Category-based Content Filtering support can be concurrently enabled on a system.

Content Filtering uses Deep Packet Inspection (DPI) feature of Enhanced Charging Service (ECS) / Active Charging Service (ACS) to discern HTTP and WAP requests.

Platform Requirements

The Content Filtering in-line service runs on a Cisco® ASR 5000 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

Licenses Requirements

The URL Blacklisting, Category-based Content Filtering and External Content Filtering Server support through Internet Content Adaptation Protocol (ICAP) interface are licensed Cisco features. Separate feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

For more information on ICAP feature, see the *ICAP Interface Support* appendix in the administration guide for the product that you are deploying.

URL Blacklisting Support

In the URL Blacklisting solution, a blacklist is a list of known URLs/URIs, which for some reason are being denied recognition. The blacklist can be obtained from a known source such as the National Center for Missing & Exploited Children (NCMEC, http://www.missingkids.com), or any other IP source. The blacklist file is obtained from various sources in known formats and converted into a non human-readable optimized format (OPTBLDB) and then made available in the system. NCMEC database is provided in plain text whereas IWF database is provided in OPTBLDB format. For more information on the blacklist file, please contact your Cisco account service representative.

Unlike the Category-based Content Filtering solution, which categorizes URLs as per a static database and takes different actions based on the different policies associated with subscribers, URL Blacklisting is applicable to all subscribers associated with a blacklisting-enabled rulebase. The same blacklist database is used for all subscribers, and for a specific URL, the same action is taken for all subscribers.

The blacklist file is downloaded and converted into a non human-readable optimized format (OPTBLDB) and then made available in the system. Once in place, all HTTP and WAP requests from subscribers are inspected in order to determine the requested destination URL/URI. If the URL/URI is not present in the blacklist then the request is passed on as usual. If the URL/URI is present in the blacklist, the request is dropped, or the flow is redirected or terminated as configured. There is no indication/messaging sent to the requesting subscribers that the requested HTTP/WAP URL/URI was rejected due to a blacklist match.

The blacklisting file can contain up to 32K URLs and the expected average size of blacklisting database (DB) is 2.5 to 5K.

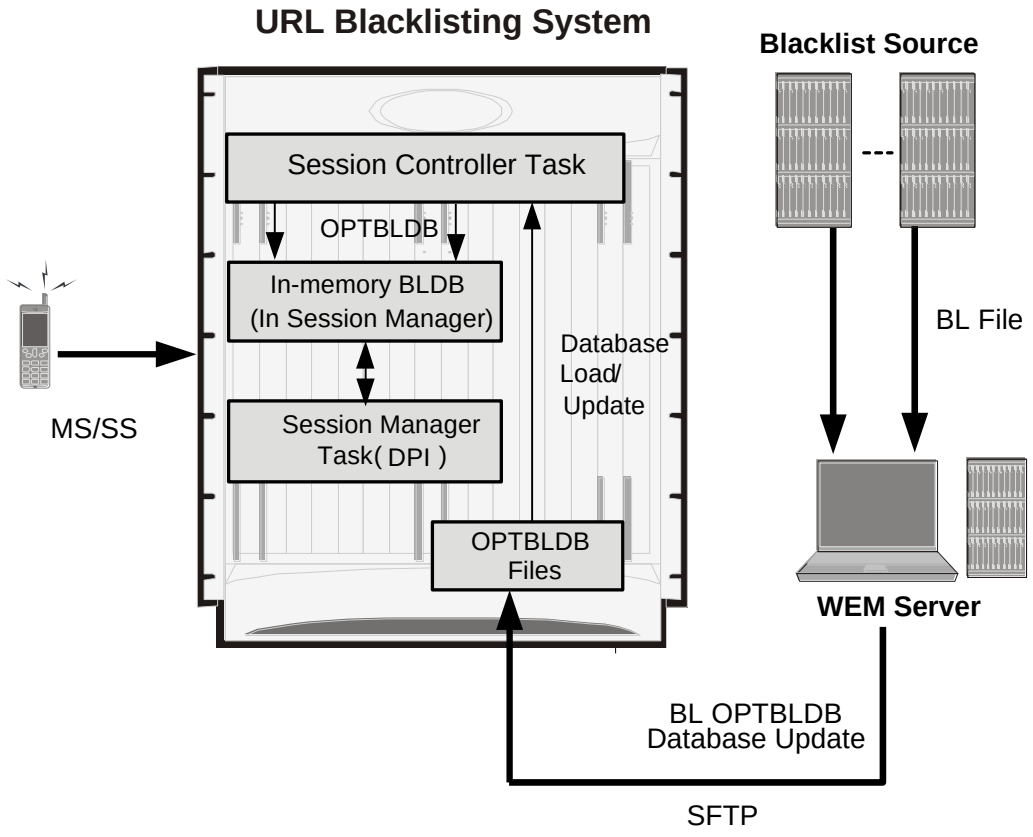
The URL Blacklisting match-method can be configured to either be generic or to look for any URL/URI in its exact, literal form.

The system generates usage/event data that can be utilized as the basis for blacklist reporting. The offline reports consist of, at a minimum, a running total of the number of times a match was made against the blacklist without any information regarding the specifics of the request.

The default/configured number of versions of the Blacklist database are maintained on the chassis (both the SPCs). This enables reverting to a particular version if required.

The following figure shows the high-level URL Blacklisting architecture with ECS, and other components in a deployment scenario.

Figure 74. High-Level Architecture of URL Blacklisting with ECS



URL Blacklisting Solution Components

The URL Blacklisting solution uses the deep-packet inspection capabilities of ECS for URL/URI extraction.

ECS functionality is managed by the following components:

- **Session Controller (SessCtrl)**: The SessCtrl runs on the primary SPC/SMC and is responsible for managing ECS and URL Blacklisting services.
- **Session Manager (SessMgr)**: A single SessMgr treats ECS charging and URL Blacklisting that is applicable to common subscriber sessions.

Apart from ECS, the URL Blacklisting solution uses the following components:

- Content Filtering Subsystem in ECS
- Web Element Manager (WEM)

Web Element Manager (WEM)

The WEM is a server-based application enabling complete element management of the system. The UNIX-based server application works with the network elements within the system using the Common Object Request Broker Architecture (CORBA) standard.

The WEM server must be set up with access to the following networks:

- Internet—to communicate with the source of the blacklist file (NCMEC/other)

The WEM application includes the following features:

- Single point of management for a large operator deployment
  - Service configuration and monitoring
  - Alarm/trap management for the WEM server
- URL Blacklisting database management functions:
  - Downloads the URL Blacklist database in *cumulative.csv* format from NCMEC, at configured schedule.
  - Converts the URL Blacklist database file to Starent Format Master Database (SFMDB) file.
  - Computes OPTBLDB suitable for updating the system.
- Distributes OPTBLDB/OPTBLDB-INC files to the chassis automatically at configured interval.

Important: For more information on WEM, refer to the *Cisco Web Element Manager Installation and Administration Guide*.

How URL Blacklisting Works

This section describes how URL Blacklisting works.

Blacklist Updates

The following steps describe how the blacklist is updated in the system:

Step 1 The WEM downloads the blacklist file from the specified source (NCMEC/other). NCMEC provides the file in clear text format which is then converted into a non-human readable optimized format (OPTBLDB), and IWF database provides the file in OPTBLDB format. The OPTBLDB file is then pushed to the chassis.

Step 2 The WEM pushes the optblk.bin file to the chassis (to the *flash/pcmcia* device) at pre-determined intervals. The optblk.bin file contains the full blacklist. If this file is verified to be correct it replaces the optblk.bin file on the chassis, and the last optblk.bin is rolled over.

Step 3 The blacklist file is auto-detected by the Session Controller (SessCtrl), which verifies the integrity of the Blacklist database using checksums, and then loads it.

The new blacklist is loaded only if it has been received properly. If the full Blacklist database is not found, corrupted, or if the loading fails, traps are generated. Correspondingly clear traps are also generated on a valid Blacklist database being available, and after a successful load.

Step 4 The SessMgrs read the file and load the blacklisted URLs in a local in-memory database.

Important: The URL Blacklisting feature is enabled only if the url-blacklisting action is set in any of the rulebases. Thus, the automatic detection of the Blacklist database, storing it in memory, and loading onto the SessMgrs will happen only if the url-blacklisting action is set in any of the rulebases.

Step 5 The Blacklist database is loaded on each SessMgr as and when they come up (if URL Blacklisting is set in any rulebase) or when URL Blacklisting gets set in any of the rulebases.

When the SessMgrs start for the first time or after recovery, if URL Blacklisting is set in any of the rulebases, the stored Blacklist database at SessCtrl is loaded onto the SessMgrs. This holds true for standby managers as well i.e., when standby managers come up, the Blacklist database is loaded onto them.

Whenever a SessMgr is killed, standby manager which already has the Blacklist database loaded takes its place, and a new standby manager is created which loads the Blacklist database as part of SessMgr getting started for the first time.

If SessCtrl is killed, while recovering it checks if URL Blacklisting is set in any of the rulebases, if set it will store the Blacklist database onto itself and load all the SessMgrs as well.

Step 6 When a new Blacklist database is loaded on to the SessMgrs, the new database (and any stored versions that have rolled over) are synced to the other SPC so that after switchover, the proper Blacklist database can be accessed.

URL Blacklisting Action

The following steps describe how the URL Blacklisting feature works:

Step 1 When an initial HTTP/WAP request comes for ECS processing and is processed by the ECS subsystem, a check is made to see if the URL Blacklisting support is enabled.

Step 2 If enabled, the URL is extracted from the incoming request and is matched with the local in-memory Blacklist database.

If a match is found for the URL in the Blacklist database, the packets are treated as per the blacklisting action configured—Discard, Redirect, or Terminate flow.

In case of multiple HTTP requests in the same TCP packet, if any of the URLs match the packet is treated as per the blacklisting action configured.

If a match is not found, the request is allowed to pass through.

Category-based Content Filtering Support

The Category-based Content Filtering application is a fully integrated, subscriber-aware in-line service provisioned on chassis running HA services. This application is transparently integrated within the ECS, and utilizes a distributed software architecture that scales with the number of active HA sessions on the system.

Content Filtering policy enforcement is the process of deciding if a subscriber should be able to receive some content. Typical options are to allow, block, or replace/redirect the content based on the rating of the content and the policy defined for that content. For the list of content categories, refer to the *Category List* appendix in this guide.

Benefits of Category-based Content Filtering

The Category-based Content Filtering solution enables operators to ensure a simplified end-to-end traffic flow with a simple network topology. In-line deployment of Content Filtering provides a more attractive solution in contrast to out-of-line solutions where the filtering and policy enforcement is provided at some offload point that is decoupled from the bearer-processing layer.

The out-of-line model forces a session to make multiple hops through a redundant array of equipment which has a negative impact on traffic latency and limits subscriber and network visibility. In addition, the out-of-line model requires all subscriber sessions to be steered to the adjunct Content Filtering platform for policy enforcement regardless of whether this additional processing is needed. This leads to increased bandwidth provisioning requirements on gateway routers.

To facilitate network simplicity, it makes sense to leverage the benefits of deep packet inspection at a single policy enforcement point that is tied to the bearer processing layer. The advantages of this approach implemented in include the following benefits:

- **Reduced processing latency**: In-line service processing eliminates unnecessary hand-offs and forwarding to external network elements.
- **Simplified policy provisioning**: Enables all policies like Content Filtering, ECS and QoS to be retrieved from same AAA/Policy Manager signaling interface thus reducing total volume of control transactions and associated delay.
- **Simplified provisioning and complete service integration**: Provisioning of separate resources like packet processing cards for processing subscriber data sessions and discrete services are eliminated. The same CPU can contain active Session Manager tasks for running Content Filtering and ECS charging.
- **Integration with Content Service Steering (CSS) architecture**: Enables applicable sessions to be forwarded to the in-line content filtering subsystem while delay and time sensitive voice/multimedia services immediately forwarded to Internet.
- **Service control**: Precise control over the interaction and service order handling of bearer flows with required applications like Content Filtering, ECS, Subscriber-aware Stateful Firewall, integrated Policy Charging and Rules Function (PCRF) for Service Based Bearer Control.

Apart from the advantages described previously, Category-based Content Filtering service reduces the requirement of over-provisioning of capacity at neighboring gateway routers. It also eliminates requirements of external Server Load Balancers and enhances the accuracy in subscriber charging records.

The Category-based Content Filtering solution has the following logical functions:

- Deep Packet Inspection (DPI) for Content Rating (event detection and content extraction)
- Content Rating Function with Static Rating of URLs and Dynamic Rating of content

- Content Rating Policy Enforcement; for example, permit, discard, deny, redirect
- Content-ware accounting CF-EDR generation for events of interest

Static-and-Dynamic Content Filtering

With Static Category-based Content Filtering, the filtering is only as good as the collection of URLs in the database. Even the largest URL database covers only a fraction of the Surface Web and virtually none of the Deep Web. It is quite impossible to find, review, and categorize enough of the available Web sites to keep the database current.

Also, many mobile sites are classified as dynamic sites. A dynamic site may return either acceptable or inappropriate content from the same URL. For example, search engines, news portals, or auction sites that return variable results depending upon subscriber requests.

When the Content Filtering subsystem receives a request for dynamic content it becomes necessary to categorize pages in real-time to determine how to classify content the provider is delivering at that moment. The "Static Rating" solution that relies exclusively on previously categorized rating for sites may fail to categorize dynamic sites appropriately.

Dynamic Content Filtering enables on-the-fly content analysis of Web traffic using different content analysis techniques. When a Web page is received, it is analyzed and then categorized according to the content found in the page. Whether a Web site has existed for five months or for five minutes does not matter since determination of the category to which the Web page belongs is made just at the time of request. Therefore, dynamic filters have no problem keeping up with the growth and changing content of the Internet. A combination of static filtering and dynamic inspection provides real accuracy and scalability as the Web weaves an increasingly sophisticated network of sites.

Important: Category-based Content Filtering can only work in static-only or in static-and-dynamic modes. Dynamic-only Content Filtering mode is not supported.

In Static-and-Dynamic Content Filtering, every URL will first undergo static rating, if the URL cannot be rated by the static database, or if the URL's statistic rating is categorized as DYNAM or UNKNOW, then it will go for dynamic rating. After the content has been analyzed, as with static content filtering, dynamic rating actions include acceptance, blocking, redirection, and/or replacement of content.

Static-and-Dynamic Content Filtering must be enabled at the global and rulebase levels. Before enabling static-and-dynamic rating in the rulebase, it must be enabled at the global level as the resources required for dynamic rating are allocated at the global level. When enabled in a rulebase, it is applied for subscribers using that rulebase.

Limitations of Dynamic Content Filtering

- Only text-based dynamic rating is supported, image-based rating is not supported.
- Only one category "PORN" is supported in all languages, while 14 other categories are supported in English.
- Content in zipped/encoded form will not be supported for dynamic rating.
- Three packet processing cards are required to support Dynamic Content Filtering.

TCP Proxy Functionality

The CF solution utilizes the services of TCP proxy to receive all the packets of a response and takes appropriate actions after rating the response. This functionality can be implemented for HTTP1.0 and HTTP1.1 protocols. For more information on the TCP Proxy feature and its implementation, refer to the *Enhanced Charging Services Administration Guide*.

Important: For the dynamic CF to be functional, the TCP Proxy feature is required.

When TCP proxy is enabled, the CF behaves as a proxy between the MS and the origin server, and two sockets are created on the respective interfaces (Gn and Gi). The Gn side socket interfaces with MS and the Gi side socket interfaces with the origin server. The TCP proxy application uses the TCP specific API to send and receive data to/from the endpoints.

When TCP proxy is configured to work with dynamic CF, the CF solution starts to buffer the packets in a temporary memory until the complete HTTP response page is received. When the entire response cannot be buffered, apply the action specified in default policy. If no default policy is specified, then allow the content to pass through. The complete response will be reassembled at CF and sent for dynamic rating only if the HTTP response code is in 2xx range. Otherwise, the CF will stream the response back to MS with no further CF action.

TCP proxy must be enabled at rulebase level. When enabled in a rulebase, it is applied for subscribers using that rulebase. For information on how to configure TCP proxy, refer to the *Configuring TCP Proxy for CF* section in the *Content Filtering Service Configuration* chapter.

Dynamic CF implementation is accurate since the rating is performed based on the complete response data and not just the first packet. The maximum limit of the response size is 256KB.

Important: In this release, the static CF rating works with both TCP proxy enabled and disabled, and the dynamic CF rating works only if TCP proxy is enabled.

Important: Dynamic CF is performed only on those responses which are either rated DYNAM or UNKNOW during static rating.

ECS and Content Filtering Application

The Category-based Content Filtering subsystem is integrated within the Enhanced Charging Service (ECS) subsystem. Although it is not necessary to provision content-based charging in conjunction with content filtering, it is highly desirable as it enables a single point of deep-packet inspection for both services. It also enables a single policy decision and enforcement point for both services thereby streamlining the required number of signaling interactions with external AAA/Policy Manager servers. Utilizing both services also increases billing accuracy of charging records by insuring that mobile subscribers are only charged for visited sites content.

The Category-based Content Filtering solution uses Content Filtering Policy to analyze the content requested by subscribers. Content Filtering Policy provides a decision point for analyzed content on the basis of its category and priority.

The Category-based Content Filtering solution also utilizes ECS rulebases in order to determine the correct policy decision and enforcement action such as accept, block, redirect, or replace. Rulebase names are retrieved during initial authentication from the AAA/Policy Manager. Some possible examples of rulebase names include Consumer, Enterprise, Child, Teen, Adult, and Sport. Rulebase names are used by the ECS subsystem to instantiate the particular rule definition that applies for a particular session. Rulebase work in conjunction with a content filtering policy and only one content filtering policy can be associated with a rulebase.

Important: For more information on rulebases and rule definitions, refer to the *Enhanced Charging Services Administration Guide*.

The ECS subsystem includes L3–L7 deep packet inspection capabilities. It correlates all L3 packets with higher layer criteria such as URL detection within an HTTP header, it also provides stateful packet inspection for complex protocols like FTP, RTSP, and SIP that dynamically open ports for the data path.

The Content Filtering subsystem uses the deep-packet inspection capabilities of ECS for URL/URI extraction.

ECS functionality is managed by the following components:

- **Session Controller (SessCtrl)**: The SessCtrl runs on the primary SPC/SMC and is responsible for managing ECS and Content Filtering services.
- **Session Manager (SessMgr)**: A single SessMgr treats ECS charging and Content Filtering that is applicable to common subscriber sessions.

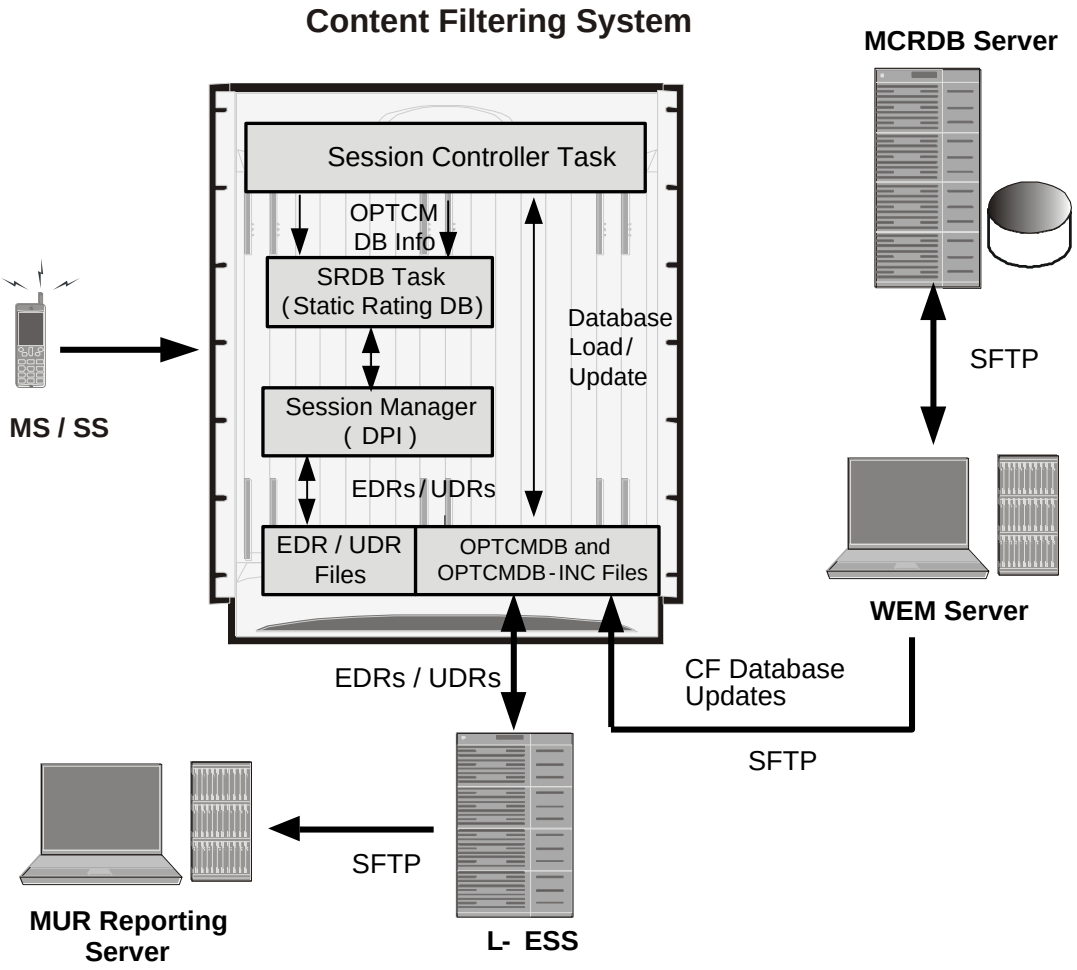
Components of Category-based Content Filtering Solution

The Category-based Content Filtering solution uses the following components:

- Content Filtering Subsystem in ECS
- Content Rating Rules Update Server
- Master Content Rating Database Server (MCRDBS)
- ECS Storage System (ESS)
- RADIUS Server/Policy Manager
- Web Element Manager (WEM)
- Mobility Unified Reporting (MUR) System

The following figure shows a high-level view of the Category-based Content Filtering architecture with ECS, and other components in a deployment scenario.

Figure 75. High-Level Architecture of Category-based Content Filtering



Category-based Content Filtering Subsystem

The Content Filtering solution comprises the following content rating and category databases:

- Static Rating Categorization Database
- Dynamic Static Rating Categorization Database

Static Rating Categorization Database (SRDB)

This is an internal categorization database (periodically synchronized with an external server) that provides ratings for publicly accessible traditional and mobile Web sites. When the SessMgr passes a URL/URI to internal list server, the list server returns a list of matching category ratings.

The list server is used to determine whether a Web site has already been classified. When the list server passes back a category rating to the filtering application, the rating is compared against the Category Policy ID applied for the subscriber to determine the appropriate action like accept, block, redirect, or replace. If the list server returns a clean rating, there is no need to perform a real-time analysis of any content delivered by the site.

When a blocked or rejected content rating is returned, the SessMgr can insert data such as a redirect server address into the bearer data stream. If no rating is returned this means the site is capable of returning either clean or unacceptable content. In this case, the Content Filtering application uses the real-time dynamic analysis engine to examine additional content served by the site.

Each SRDB contains a replication object consisting of hash tables that map known Web sites and their subdirectories to their respective category ratings. The SessCtrl reads the index of SRDB tables with a data structure that associates keys with URL rating values and loads it onto the SRDB managers.

To boost performance and provide high availability, SRDB Manager provides functionality to load the Optimized Content Rating Master Database (OPTCMDB) volumes from its peer SRDB task. If the peer SRDB task is not in loading state then the OPTCMDB loading is done through SessCtrl to the recovered SRDB task.

Dynamic Static Rating Categorization Database

When Static-and-Dynamic Content Filtering is enabled, Dynamic SRDB tasks are spawned based on available packet processing cards.

Important: To support dynamic rating, a minimum of three active packet processing cards are required, that will have one Dynamic SRDB. The number of Dynamic SRDBs may increase with an increase in the number of packet processing cards. The load for rating dynamic responses is distributed equally across all the Dynamic SRDBs created.

First one Dynamic SRDB task is created with two Standby SRDBs on the same CPU, then eight Static SRDBs are created, and then more Dynamic SRDB tasks are created if memory is available. A Dynamic SRDB is always created with two Standby SRDBs with it on the same CPU.

The Dynamic Rater Package, which contains the model files (used for language detection and category recognition) and feature counters (used to decide whether or not to evaluate the Web page against the respective model file) are loaded on the SRDB Managers. The rater package is loaded only on the active SRDB and not on the standbys.

The Rater package containing the model files (used for language detection and category recognition) and feature counters (used to decide whether or not to evaluate the Web page against the respective model file) is stored at pcmcia1/cf. After loading the static database, ECS will read the Rater package and load it onto the Dynamic SRDB Managers.

The Rater package will also be loaded on SRDBs on recovery/reconciliation if static-and-dynamic Content Filtering is enabled.

The rater package loaded onto SRDBs can be upgraded using an upgrade CLI, that will look for the upgrade file in the form “rater_f.pkg” at a specific location and if found load the new package onto the SRDBs. On successful loading, the “rater_f.pkg” is replaced with “rater.pkg” and versioned. In case of loading or upgrade failures, appropriate traps are generated.

Rater Package Model Files

The real-time analyzer requires a model file that defines the features which are necessary to classify a Web page as belonging to a specific category and language. A model file per category is created by analyzing the traits of th ousands of pages of that category and thousands of pages that does not belong to that category. For some categories, a feature counter file is used to decide whether or not to evaluate the Web page against the respective model file.

When URL Blacklisting solution is the only content filtering enabled on a system, no SRDB tasks are spawned at startup. Only when either Category-based Content Filtering is enabled in isolation, or with URL Blacklisting, the SRDB tasks are spawned.

Content Rating Rules Update Server

This is a third-party content rating solution for exporting content filtering rules database information to the Category-based Content Filtering system. In addition, while exporting database updates, it collects reports of URLs processed by ECS and Content Filtering services that are reported as unknown in the deployed static rating database. This server analyzes these URLs and provides the rating in future updates for static rating database.

This server provides the following support to Master Content Rating Database Server (MCRDBS) for the content rating function:

- Provides full Vendor Format Master Database files (VFMDB) to MCRDB server on request from MCRDBS.
- Provides incremental Vendor Format Master Static URL Database file (VFMDB-INC) to MCRDBS when any incremented VFMDB is available and requested from MCRDBS.
- Receives the Unknown URLs file (Vendor Format Unknown Database File (VFUNKDB)) from MCRDBS.

Master Content Rating Database Server (MCRDBS)

The Category-based Content Filtering solution provides a Master Content Rating Database Server to convert the VFMDB to SFMDB. It handles both full and incremental updates and processes them on a configured schedule.

This server is also responsible for distribution of SFMDB data files to WEM servers in the customer support infrastructure on a configured interval.

The server is responsible for following functionality as the MCRDBS solution:

- Database fetching: Pulls VFMDB files from third-party Content Rating Server to MCRDBS.
- Database conversion: Converts VFMDB files to SFMDB files. It also handles the incremented and unknown database files.
- Database poller: Provides the converted SFMDB database files for WEM in a preconfigured path.
- E-mail notification: Provides alerts and notification to the administrator for alarms.

External Storage System

The local external storage server is a part of ECS Storage System in the ECS solution architecture.

The L-ESS is a storage application running on redundant highly available servers that collect and process EDRs and UDRs from which billing events and reports are generated. Either the system pushes the EDR/UDR files to the L-ESS, or the L-ESS fetches them from the system and processes them into formats suitable for billing mediation servers and MUR server. The L-ESS consolidates the processed EDR/UDR files into a database for report generation through MUR. The database generated on an ESS by processing EDR/UDR records is a superset of the database required by MUR.

Important: For more information on External Storage System, refer to the *ESS Installation and Administration Guide*.

RADIUS Server and Policy Manager

The function of the RADIUS Server/Policy Manager in the Content Filtering solution is to provide per-subscriber Content Filtering provisioning information when a subscriber's session is established. It can also issue a Change-of-Authorization (CoA) to update an in-progress session to modify the Content Filtering policy for a subscriber.

The following are the basic functions provided by a RADIUS Server/Policy Manager in the Content Filtering solution:

- Support for the in/out ACL attributes to direct traffic through ECS for processing of subscriber traffic
- Support for ECS rulebase VSA to select the ECS rulebase to be applied to filtered traffic
- Support for Content Filtering Policy identifier VSA to select the content filtering policy within the selected rulebase for a subscriber
- Support exporting a subscriber provisioning record based on MSID to the customer service interface (Customer Care Interface) so that operator's customer care executive can see the provisioned content filtering policy for a subscriber

Web Element Manager (WEM)

The WEM is a server-based application providing complete element management of the system. The UNIX-based server application works with the network elements within the system using the Common Object Request Broker Architecture (CORBA) standard.

WEM server must be set up with access to the following networks:

- Internet: To communicate with the Master Content Rating Database Server (MCRDBS) which provides update files.

For Category-based Content Filtering, the WEM application includes the following features:

- Single point of management for a large Content Filtering Service operator deployment:
  - Content Filtering service configuration and monitoring
  - Alarm/trap management
- Configures and manages the operator-defined White/Black static rating database (WBLIST) for the network (WBLIST is maintained in SFMDB format)
- Content filtering database management functions:
  - Performs database processing in the background
  - Imports full and incremental SFMDB and SFMDB-INC files from the MCRDBS on a configured schedule
  - Processes incremental SFMDB-INC updates from MCRDBS maintaining an updated SFMDB file
  - Merge the operator's WBLIST database with the most recent SFMDB creating a SFCMDB
  - Computes an incremental update to the OPTCMDB-INC suitable for updating the Content Filtering subsystem that contains a previous version OPTCMDB
- Distributes OPTCMDB/OPTCMDB-INC files to the chassis automatically at configured interval

Important: For information on WEM, refer to the *Cisco Web Element Manager Installation and Administration Guide*.

Mobility Unified Reporting System

The Mobility Unified Reporting (MUR) application is a Web-based application providing a unified reporting interface for diverse data from the in-line service and storage applications. The MUR application provides comprehensive and consistent set of statistics and customized reports, statistical trending, report scheduling and distribution from chassis / in-line service product. For example, a subscriber's Quality of Experience, top 10 sites visited, top 10 users, and so on. The MUR application facilitates and enhances the operators' ability to simply and easily determine the health and usage of the network.

The MUR application supports the generation of various reports including CF-EDR reports in PDF and XML formats. The CF-EDR reports provide the summary of traffic over CF categories, CF actions, and CF ratings. It also provides the list of top N subscribers and URLs based on their unique subscriber's hit count and total usage.

- Summary Reports:
  - Category summary (volume/hits)
  - Action summary (volume/hits)
  - Rating summary (volume/hits)
- Top N Reports:
  - Top N Subscribers by volume/hits
  - Top N URLs by volume/hits

The CF-EDR files are pushed from L-ESS to MUR at a configured time interval and stored in a specified data directory on the MUR server. It can also create the files from CF-EDRs for unrated URLs which can be pulled by WEM.

For more information on the reports, refer to the *Mobility Unified Reporting System Online Help* documentation.

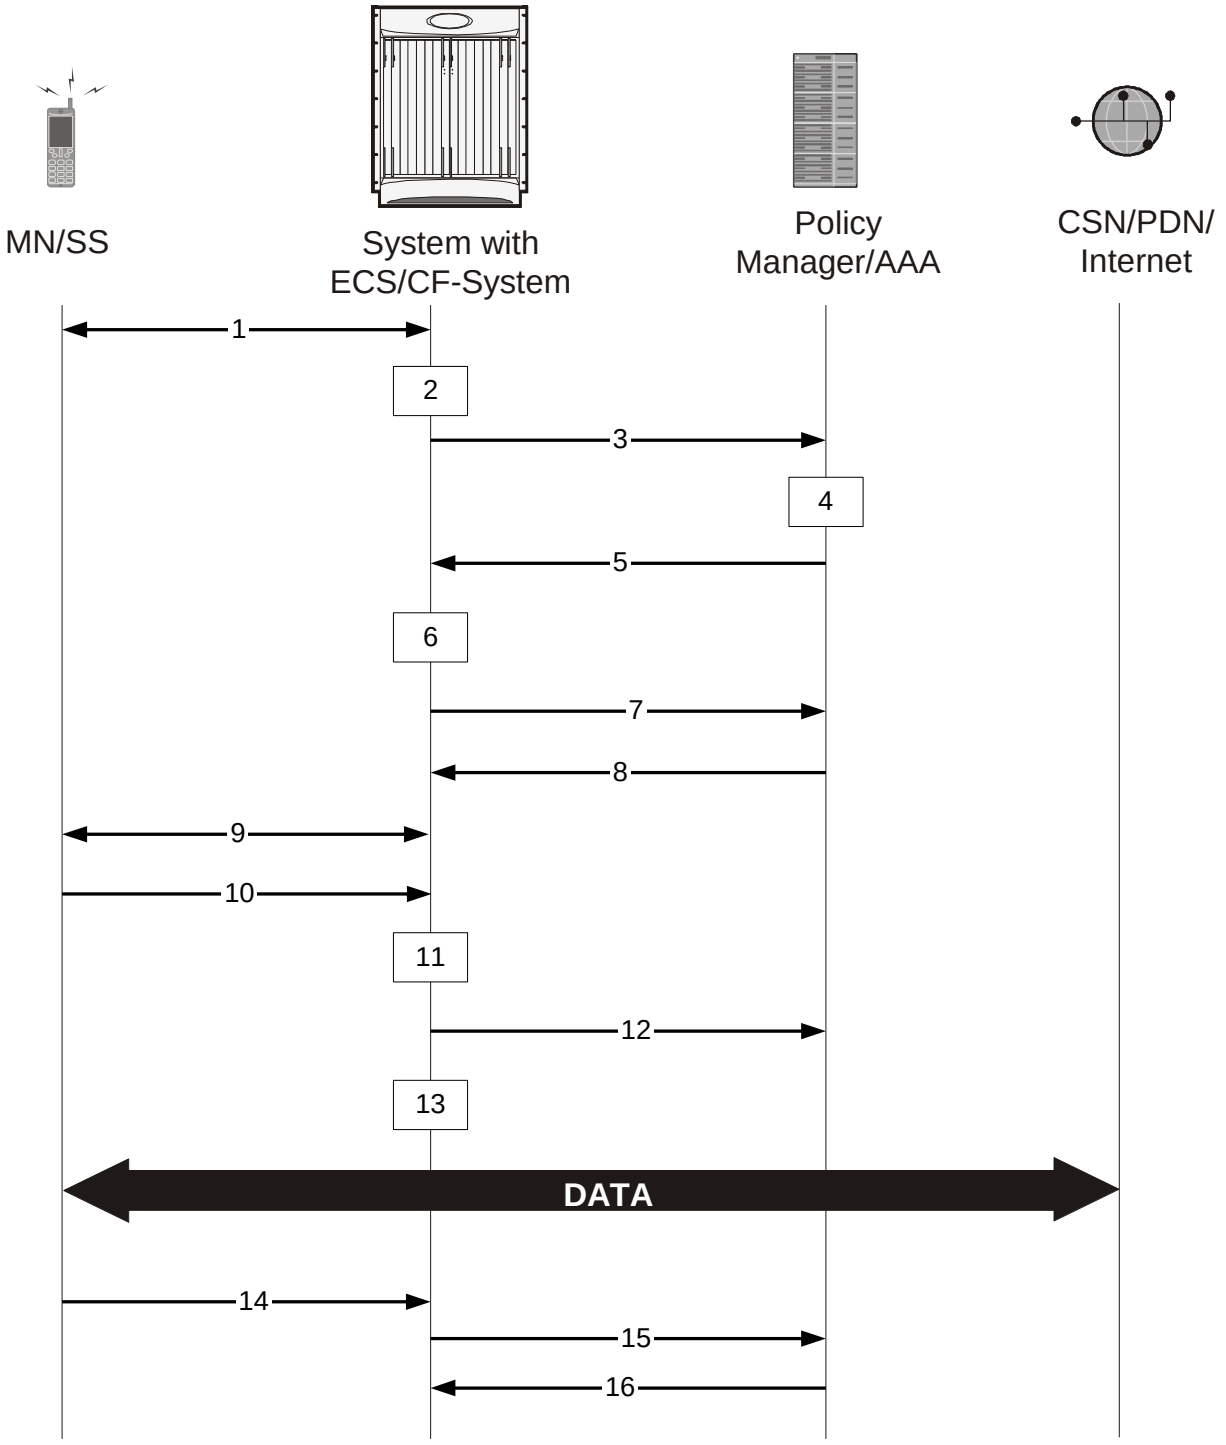
How Category-based Content Filtering Works

The Content Filtering Subsystem which is integrated into the ECS subsystem consists of an onboard static categorization database and a dynamic rating engine. The filtering service uses the Deep Packet Inspection (DPI) capabilities of the ECS subsystem to classify and partition application or protocol specific flows into virtual sessions.

Content analyzers are used to identify various types of flows such as HTTP, MMS/WAP, and POP3 E-mail. A typical HTTP request for a Web page, for example, invokes TCP and HTTP traffic analyzers. Any HTTP field including URLs or URIs can be identified. When a subscriber session is bound by CSS to an ECS running content filtering service, the URL/URI is extracted and compared against the static categorization database.

The following figure and the steps describe how Category-based Content Filtering works during a subscriber call:

Figure 76. Content Filtering Call Flow



Step 1 MS requests for registration to the system.

Step 2 System processes MS-related information with Content Filtering subsystem.

Step 3 System sends the AAA Access Request to AAA server for MS.

Step 4 AAA server processes the AAA Access Request from the Content Filtering subsystem to create the session, and the Policy Manager in AAA server uses subscriber identification parameters including NAI (*username@domain*), Calling Station ID (IMSI, MSID) and Framed IP Address (HoA) as the basis for subscriber lookup.

Step 5 The Policy Manager and AAA generate and send an Access Accept message including all policy and other attributes to establish the session to the Content Filtering subsystem.

The Policy Manager and/or AAA include the following attributes in the Access Accept message:

- **Filter-Id or Access Control List Name**: Applied to subscriber session. It typically contains the name of the Content Service Steering (CSS) ACL. The CSS ACL establishes the particular service treatments such as Content Filtering, ECS, Traffic Performance Optimization, Stateful Firewall, VPN, etc. to apply to a subscriber session and the service order sequence to use in the inbound (SN-IP-In-ACL AVP) or outbound (SN-IP-Out-ACL AVP) directions. Real-time or delay sensitive flows are directly transmitted to the Internet with no further processing required. In this case, no CSS ACL or Filter ID is included in the Access Response.
- **SN-CFPolicy-ID**: Applied to the subscriber content flow. Policy ID included in this attribute overrides the policy identifier applied to subscriber through rulebase or APN/Subscriber configuration. This content filtering policy determines the action to be taken on a content request from subscriber on the basis of its category. At anytime only one content filtering policy can be associated with a rulebase.
- **SN1-Rulebase**: This custom attribute contains information such as consumer, business name, child/adult/teen, etc.). The rulebase name identifies the particular rule definitions to apply. Rulebase definitions are used in ECS as the basis for deriving charging actions such as prepaid/postpaid volume/duration/destination billing and charging data files (EDRs/UDRs). Rulebase definitions are also used in content filtering to determine whether a type of user class such as teenagers should be permitted to receive requested content belonging to a particular type of category such as adult entertainment, gambling or hate sites. Rulebase definitions are generated in the Active Charging Configuration Mode and can be applied to individual subscribers, to domains or on per-context basis.

Step 6 Content Filtering subsystem creates a new session for MS.

Step 7 Content Filtering subsystem sends Accounting-Start messages to AAA server.

Step 8 AAA server sends Accounting-Start response message to Content Filtering subsystem.

Step 9 Content Filtering subsystem establishes data flow with MS.

Step 10 MS requests for data with URL name.

Step 11 Within the system access control list (ACL) processes the request and directs the request to ECS/Content Filtering subsystem based on the subscriber configuration.

Step 12 System performs ECS action on the content and then applies content filtering if required.

Within the system, if the bearer flow is treated by Content Filtering or other in-line services, the SessMgr feeds it to the Content Service Steering (CSS) API. If Content Filtering is the first service touch point, TCP and HTTP traffic analyzers within a given SessMgr utilize deep-packet inspection to extract the requested URL.

Step 13 The Content Filtering subsystem processes the URL access request.

When only Static Content Filtering is enabled, first the URL is looked-up in the cache maintained at SessMgr for static URL requests, if there is a hit, the category is returned, if its a miss, a URL look-up is performed by an onboard SRDB for static rating.

- If a category is returned, action is taken as configured for that category in the subscriber's Content Filtering policy:
  - allow: If the category is permitted by the subscriber's content filtering policy, the request is sent to the server, and the response transmitted to the subscriber's mobile.
  - content-insert: The system notifies the subscriber's mobile of the blocked content by inserting a specified message within the IP data stream, and prevents access to the requested content. The insert string is as specified in the subscriber's content filtering policy.
  - discard: The system silently discards the request packet(s).
  - redirect-url: The system inserts a specified redirect server address in the bearer data stream and returns an HTTP error message to the subscriber's mobile. The redirect address is as specified in the subscriber's content filtering policy.

    The redirect server may prompt the subscriber to send additional security credentials in order to access the requested content.
  - terminate-flow: The system gracefully terminates the TCP connection between the subscriber and server, and sends a TCP FIN to the subscriber and a TCP RST to the server.
  - www-reply-code-and-terminate-flow: The system terminates the flow with a specified reply code to the subscriber's mobile. The reply code is as specified in the subscriber's content filtering policy.
- If a category is not returned / the URL is not present in the database, the system takes the action as configured for the UNKNOW category in the subscriber's Content Filtering policy.
- If for the category returned there is no action configured in the subscriber's content filtering policy, the default action is taken.

When TCP proxy is enabled to work with static CF / dynamic CF, the CF implementation remains the same unless the packet content is modified. When the packet content is modified, the CF solution uses the TCP specific APIs to send the packets to the client.

Handling for concatenated and pipelined responses is the same as in Static Content Filtering. The action taken is based on the highest priority category among the pipelined responses.

- Content Filtering EDRs are generated for action taken after dynamic rating.

  Content Filtering EDRs are the same as for static rating. However if static rating fails and the request goes for dynamic rating, then Content Filtering EDRs will be generated only after dynamic rating has been completed and not when static rating failed.

If the SRDB task is timed out or some other failure happens, the action configured for failure is taken.

Step 14 MS requests for session termination.

Step 15 System sends Accounting-Stop Request to the AAA server.

Step 16 AAA server stops the accounting for the MS for content filtering session and sends Accounting-Stop-Response to the system.

How URL Blacklisting and Category-based Content Filtering Work Concurrently

Both URL Blacklisting and Category-based Content Filtering can be concurrently enabled in a system. The following describes how URL blacklisting and content filtering are performed on HTTP/WAP traffic when concurrently enabled on a system:

Step 1 If both URL Blacklisting and Category-based Content Filtering are enabled, first URL blacklist matching is performed, and then, if required, content filtering is performed.

When an HTTP/WAP request comes for ECS processing, a check is made to see if the URL Blacklisting feature is enabled. If enabled, the URL is extracted from the incoming request and is matched with the local Blacklist database.

- If a match is found for the URL in the Blacklist database, the packets are subjected to the blacklisting action configured in the rulebase—Discard, Redirect, or Terminate flow. In case of multiple HTTP requests in the same TCP packet, if any of the URLs is blacklisted, then action is taken on the packet.
- If a match is not found in the Blacklist database, then Category-based Content Filtering is performed.
  - If Category-based Static Content Filtering is enabled, static rating is performed and action taken as configured for the category returned in the subscriber's content filtering policy.

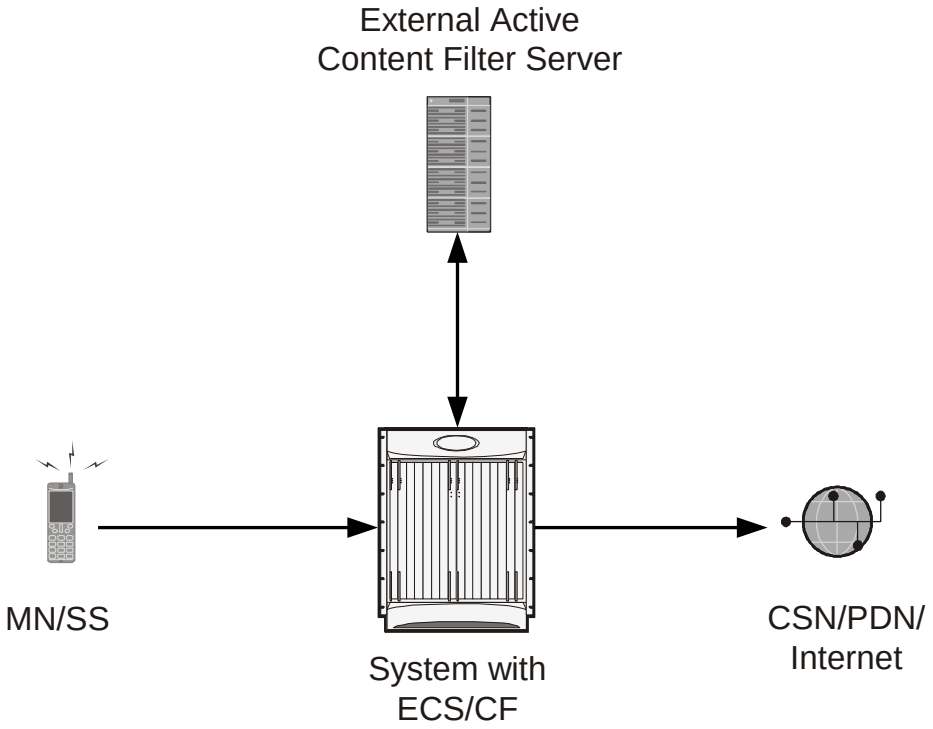
Step 2 If URL Blacklisting is enabled and Category-based Content Filtering is disabled, and a match is not found for the URL in the Blacklist database, the request is allowed to pass through, and no Content Filtering EDRs are generated for those flows.

Content Filtering Server Group Support

ECS supports the streamlined ICAP interface to leverage Deep Packet Inspection to enable external application servers to provide their services without performing DPI, and without being inserted in the data flow. For example, with an external Active Content Filtering (ACF) platform.

A high-level view of the streamlined ICAP interface support for external ACF is shown in the following figure.

Figure 77. High-Level View of Streamlined ICAP Interface with External ACF



The system with ECS is configured to support DPI and the system uses this capability for content charging as well.

If a subscriber initiates a WAP (WAP1.x or WAP2.0) or Web session, the subsequent GET/POST request is detected by the DPI function. The URL of the GET/POST request is extracted and passed, along with subscriber identification information and the subscriber request, in an ICAP message to the application server.

In the case of Category-based Content Filtering solution, the application server checks the URL on the basis of its category and other classifications like type, access level and content category and decides if the request should be authorized, blocked, or redirected by answering to the GET/POST with:

- A 200 OK message if the request is accepted.
- A 302 Redirect message in case of redirection. This redirect message includes the URL to which the subscriber should be redirected.
- A 403 Denied message is the request should be blocked.

Depending on the response received, the system with ECS will either pass the request unmodified, or discard the message, and respond to the subscriber with the appropriate redirection or block message.

Content Charging is performed by the ECS only after the request has been controlled by the application server. This guarantees the appropriate interworking between the external application and content-based billing. In particular, this guarantees that charging will be applied to the appropriate request in case of redirection, and that potential charging based redirections (i.e. Advice of Charge, Top Up page, etc.) will not interfere with the decisions taken by the application server.

The ACF performs the following functions:

- Retrieval of subscriber policies based on the subscriber identity passed in the ICAP message.
- Determining the appropriate action (permit, deny, redirect) to take for this type of content based on subscriber profile.
- Communication of the action (permit, deny, or redirect) decision for the URL back to the ECS subsystem.

For information on configuring the ICAP interface functionality for external ACF servers, see the *ICAP Interface Support* chapter of the *System Administration Guide*.

External Storage System

ESS supports generation of EDR/UDR/FDR (xDR) files from the chassis. To store generated xDR files, on the Cisco chassis, the system allocates 512 MB of memory on the packet processing card's RAM. The generated xDRs are stored in CSV format in the */records* directory on the packet processing card RAM. These generated xDRs can be used for billing as well as for generation of reports to analyze network usage and subscriber trends. As this temporary storage space (size configurable) reaches its limit, the system deletes older xDRs to make room for new xDRs. Setting gzip file compression extends the storage capacity by approximately 10:1.

Because of the volatile nature of the memory, xDRs can be lost due to overwriting, deletion, or unforeseen events such as power or network failure or unplanned chassis switchover. To avoid loosing charging and network analysis information, configure the CDR subsystem in conjunction with the External Storage System (ESS) to offload the xDRs for storage and analysis.

For more information on the ESS, refer to the *ESS Installation and Administration Guide*.

Bulk Statistics Support

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Content Filtering bulk statistics support only System schema.

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the chassis or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files. The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, host name, chassis uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

For more information, see the *Configuring and Maintaining Bulk Statistics* chapter of the *System Administration Guide*.

Minimum System Requirements and Recommendations

This section identifies the minimum system requirements for components of the URL Blacklisting / Category-based Content Filtering solutions.

Important: The hardware required for these components may vary, depending on the number of clients that require access, components managed, and other variables like EDR generation rate or CDR storage and processing requirements.

Certain basic server requirements are recommended for WEM and MUR system to exploit the CF solution. For information on these system requirements, refer to *Cisco Web Element Manager Installation and Administration Guide* and *Cisco Mobility Unified Reporting System Installation and Administration Guide*.

MCRDBS System Requirements

This section provides information on the system requirements for MCRDBS.

Important: You must ensure that the minimum system requirements are met before proceeding with the MCRDBS installation.

Hardware Requirements

- Dell PowerEdge 1950 server
  - 1.86 GHz Dual quad-core Intel Xeon CPU
  - 8 GB RAM
  - 2 * 146 GB RAID hard disk drive. The hard disk can be expanded up to 300 GB.
  - Gigabit Ethernet interfaces
  - CD-ROM Drive
- Operating Environment:
  - Red Hat Enterprise Linux 5.4

- or -

- Sun Microsystems Netra™ X4270 server
  - Quad-Core two socket Intel Xeon L5518 processor (1 * 4GB memory kit, 1333 MHz)
  - 32GB RAM
  - 2 * 300GB 10K RPM SAS disks
  - SATA DVD drive
  - 8-port internal SAS HBA
  - Choice of AC or DC power supplies
- Operating Environment:

    - Red Hat Enterprise Linux 5.4
- ZFS is the recommended file system with two ZFS pools.

Important: For the MCRDBS 10.0 and earlier releases, it is recommended to use the hardware configurations of Dell PowerEdge 1950 server.

Important: For the MCRDBS 11.0 and later versions, please use the hardware recommendations of X4270 server.

Additional Requirements on Chassis

The chassis requires the following additional hardware and memory to handle the Content Rating Master Databases; for example, for Category-based Content Filtering OPTCMDB. The memory required may vary with the size of rating databases used for content rating service.

- Minimum of two active packet processing cards s are required
- Minimum 4 GB memory:
    - in Cisco chassis on Flash memory

Chapter 12
Enhanced Charging Service Overview

This chapter provides an overview of the Enhanced Charging Service (ECS) in-line service, also known as Active Charging Service (ACS).

This chapter covers the following topics:

- Introduction
- Basic Features and Functionality
- ECS Deployment and Architecture
- Enhanced Features and Functionality

Introduction

The Enhanced Charging Service (ECS) is an in-line service feature that enables operators to reduce billing-related costs and gives the ability to offer tiered, detailed, and itemized billing to their subscribers. Using shallow and deep packet inspection (DPI), ECS allows operators to charge subscribers based on actual usage, number of bytes, premium services, location, and so on. ECS also generates charging records for postpaid and prepaid billing systems.

The ECS is an enhanced or extended premium service. The *System Administration Guide* provides basic system configuration information, and the product administration guides provide information to configure the core network service functionality. It is recommended that you select the configuration example that best meets your service model, and configure the required elements for that model before using the procedures in this document.

Platform Requirements

The ECS in-line service runs on Cisco ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

License Requirements

The ECS in-line service is a licensed Cisco feature. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements.

For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Basic Features and Functionality

This section describes basic features of the ECS in-line service.

Shallow Packet Inspection

Shallow packet inspection is the examination of the layer 3 (IP header) and layer 4 (for example, UDP or TCP header) information in the user plane packet flow. Shallow packet analyzers typically determine the destination IP address or port number of a terminating proxy.

Deep Packet Inspection

Deep-packet inspection is the examination of layer 7, which contains Uniform Resource Identifier (URI) information. In some cases, layer 3 and 4 analyzers that identify a trigger condition are insufficient for billing purposes, so layer 7 examination is used. Whereas, deep-packet analyzers typically identify the destination of a terminating proxy.

For example, if the Web site "www.companyname.com" corresponds to the IP address 1.1.1.1, and the stock quote page (www.companyname.com/quotes) and the company page (www.companyname.com/business) are chargeable services, while all other pages on this site are free. Because all parts of this Web site correspond to the destination address of 1.1.1.1 and port number 80 (http), determination of chargeable user traffic is possible only through the actual URL (layer 7).

DPI performs packet inspection beyond layer 4 inspection and is typically deployed for:

- Detection of URI information at level 7 (for example, HTTP, WTP, RTSP URLs)
- Identification of true destination in the case of terminating proxies, where shallow packet inspection would only reveal the destination IP address/port number of a terminating proxy such as the OpCo's WAP gateway
- De-encapsulation of nested traffic encapsulation, for example MMS-over-WTP/WSP-over-UDP/IP
- Verification that traffic actually conforms to the protocol the layer 4 port number suggests

Charging Subsystem

ECS has protocol analyzers that examine uplink and downlink traffic. Incoming traffic goes into a protocol analyzer for packet inspection. Routing rules definitions (ruledefs) are applied to determine which packets to inspect. This traffic is then sent to the charging engine where charging rules definitions are applied to perform actions such as block, redirect, or transmit. These analyzers also generate usage records for the billing system.

Traffic Analyzers

Traffic analyzers in ECS are based on configured ruledefs. Ruledefs used for traffic analysis analyze packet flows and create usage records. The usage records are created per content type and forwarded to a prepaid server or to a billing system.

The Traffic Analyzer function can perform shallow (layer 3 and layer 4) and deep (above layer 4) packet inspection of IP packet flows. It is able to correlate all layer 3 packets (and bytes) with higher layer trigger criteria (for example, URL detected in an HTTP header). It also performs stateful packet inspection for complex protocols like FTP, RTSP, and SIP that dynamically open ports for the data path and this way, user plane payload is differentiated into "categories". Traffic

analyzers can also detect video streaming over RTSP, and image downloads and MMS over HTTP and differential treatment can be given to the Vcast traffic.

Traffic analyzers work at the application level as well, and perform event-based charging without the interference of the service platforms.

The ECS content analyzers can inspect and maintain state across various protocols at all layers of the OSI stack. The ECS supports inspecting and analyzing the following protocols:

- Domain Name System (DNS)
- File Transfer Protocol (FTP)
- Hyper Text Transfer Protocol (HTTP)
- Internet Control Message Protocol (ICMP)
- Internet Control Message Protocol version 6 (ICMPv6)
- Internet Message Access Protocol (IMAP)
- Internet Protocol version 4 (IPv4)
- Internet Protocol version 6 (IPv6)
- Multimedia Messaging Service (MMS)
- Post Office Protocol version 3 (POP3)
- RTP Control Protocol/Real-time Transport Control Protocol (RTCP)
- Real-time Transport Protocol (RTP)
- Real Time Streaming Protocol (RTSP)
- Session Description Protocol (SDP)
- Secure-HTTP (S-HTTP)
- Session Initiation Protocol (SIP)
- Simple Mail Transfer Protocol (SMTP)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Wireless Session Protocol (WSP)
- Wireless Transaction Protocol (WTP)

Apart from the above protocols, ECS also supports analysis of downloaded file characteristics (for example, file size, chunks transferred, and so on) from file transfer protocols such as HTTP and FTP.

How ECS Works

This section describes the base components of the ECS solution, and the roles they play.

Content Service Steering

Content Service Steering (CSS) enables directing selective subscriber traffic into the ECS subsystem (in-line services internal to the system) based on the content of the data presented by mobile subscribers.

CSS uses Access Control Lists (ACLs) to redirect selective subscriber traffic flows. ACLs control the flow of packets into and out of the system. ACLs consist of "rules" (ACL rules) or filters that control the action taken on packets matching the filter criteria.

ACLs are configurable on a per-context basis and applies to a subscriber through either a subscriber profile (for PDSN) or an APN profile (for GGSN) in the destination context.

Important: For more information on CSS, refer to the *Content Service Steering* chapter of the *System Administration Guide*. For more information on ACLs, refer to the *IP Access Control Lists* chapter of the *System Administration Guide*.

Protocol Analyzer

The Protocol Analyzer is the software stack responsible for analyzing the individual protocol fields and states during packet inspection.

The Protocol Analyzer performs two types of packet inspection:

- **Shallow Packet Inspection**—Inspection of the layer 3 (IP header) and layer 4 (for example, UDP or TCP header) information.
- **Deep Packet Inspection**—Inspection of layer 7 and 7+ information. DPI functionality includes:
  - Detection of Uniform Resource Identifier (URI) information at level 7 (for example, HTTP, WTP, and RTSP URLs)
  - Identification of true destination in the case of terminating proxies, where shallow packet inspection would only reveal the destination IP address/port number of a terminating proxy
  - De-encapsulation of upper layer protocol headers, such as MMS-over-WTP, WSP-over-UDP, and IP-over GPRS
  - Verification that traffic actually conforms to the protocol the layer 4 port number suggests

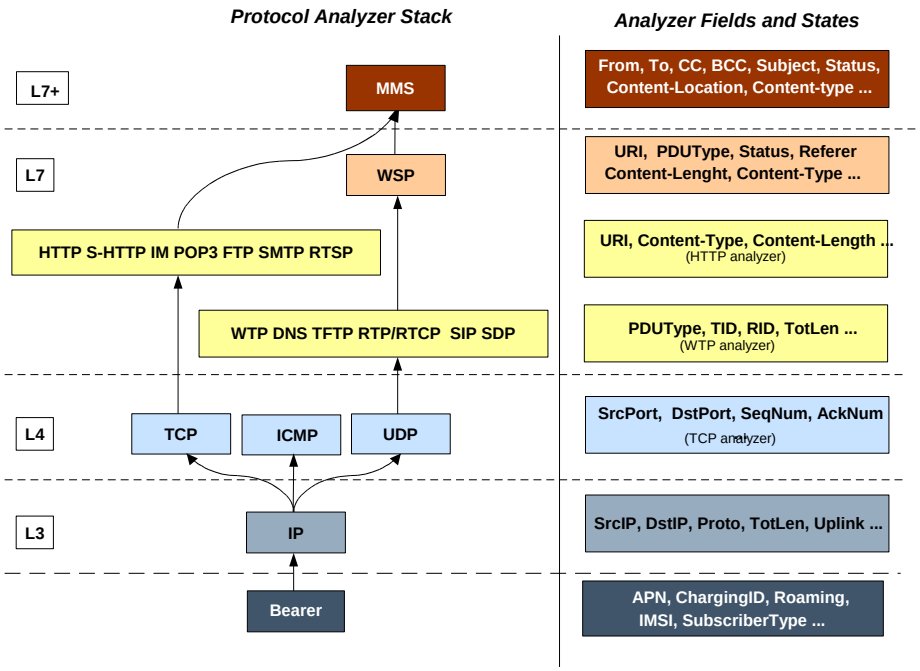
The Protocol Analyzer performs a stateful packet inspection of complex protocols, such as FTP, RTSP, and SIP, which dynamically open ports for the data path, so the payload can be classified according to content.

The Protocol Analyzer is also capable of determining which layer 3 packets belong (either directly or indirectly) to a trigger condition (for example, URL). In cases where the trigger condition cannot be uniquely defined at layers 3 and 4, then the trigger condition must be defined at layer 7 (that is, a specific URL must be matched).

Protocol Analyzer Software Stack

Every packet that enters the ECS subsystem must first go through the Protocol Analyzer software stack, which comprises of individual protocol analyzers for each of the supported protocols.

Figure 78. ECS Protocol Analyzer Stack



Note that protocol names are used to represent the individual protocol analyzers.

Each analyzer consists of fields and states that are compared to the protocol-fields and protocol-states in the incoming packets to determine packet content.

Rule Definitions

Rule definitions (ruledefs) are user-defined expressions based on protocol fields and protocol states, which define what actions to take on packets when specified field values match.

Rule expressions may contain a number of operator types (string, =, >, and so on) based on the data type of the operand. For example, "string" type expressions like URLs and host names can be used with comparison operators like "contains", "!contains", "=", "!=", "starts-with", "ends-with", "!starts-with" and "!ends-with". In 14.0 and later releases, ECS also supports regular expression based rule matching. For more information, refer to the *Regular Expression Support for Rule Matching* section.

Integer type expressions like "packet size" and "sequence number" can be used with comparison operators like "=", "!=", ">=", "<=". Each ruledef configuration consists of multiple expressions applicable to any of the fields or states supported by the respective analyzers.

Ruledefs are of the following types:

- **Routing Ruledefs**—Routing ruledefs are used to route packets to content analyzers. Routing ruledefs determine which content analyzer to route the packet to when the protocol fields and/or protocol-states in ruledef expression are true. Up to 256 ruledefs can be configured for routing.
- **Charging Ruledefs**—Charging ruledefs are used to specify what action to take based on the analysis done by the content analyzers. Actions can include redirection, charge value, and billing record emission. Up to 2048 charging ruledefs can be configured in the system.

- **Post-processing Ruledefs**—Used for post-processing purposes. Enables processing of packets even if the rule matching for them has been disabled.

Important: When a ruledef is created, if the rule-application is not specified for the ruledef, by default the system considers the ruledef as a charging ruledef.

Ruledefs support a priority configuration to specify the order in which the ruledefs are examined and applied to packets. The names of the ruledefs must be unique across the service or globally. A ruledef can be used across multiple rulebases.

Important: Ruledef priorities control the flow of the packets through the analyzers and control the order in which the charging actions are applied. The ruledef with the lowest priority number invokes first. For routing ruledefs, it is important that lower level analyzers (such as the TCP analyzer) be invoked prior to the related analyzers in the next level (such as HTTP analyzer and S-HTTP analyzers), as the next level of analyzers may require access to resources or information from the lower level. Priorities are also important for charging ruledefs as the action defined in the first matched charging rule apply to the packet and ECS subsystem disregards the rest of the charging ruledefs.

Each ruledef can be used across multiple rulebases, and up to 2048 ruledefs can be defined in a charging service.

Ruledefs have an expression part, which matches specific packets based upon analyzer field variables. This is a boolean (analyzer_field operator value) expression that tests for analyzer field values.

The following is an example of a ruledef to match packets:

```
http url contains cnn.com
```

–or–

```
http any-match = TRUE
```

In the following example the ruledef named “rule-for-http” routes packets to the HTTP analyzer:

```
route priority 50 ruledef rule-for-http analyzer http
```

Where, **`rule-for-http`** has been defined with the expressions: **`tcp either-port = 80`**
The following example applies actions where:

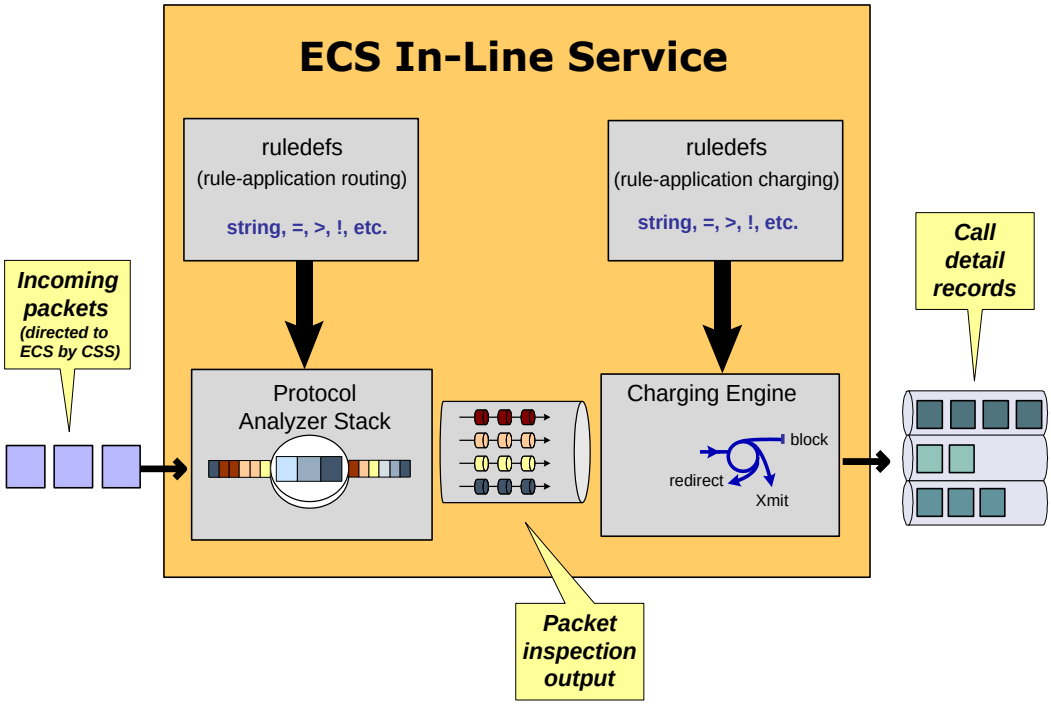
- Subscribers whose packets contain the expression “bbc-news” are not charged for the service.
- All other subscribers are charged according to the duration of use of the service.

```
ruledef port-80
   tcp either-port = 80
   rule-application routing
   exit
ruledef bbc-news
   http url starts-with http://news.bbc.co.uk
   rule-application charging
   exit
```

```
ruledef catch-all
   ip any-match = TRUE
   rule-application charging
   exit
charging-action free-site
   content-id 100
   [ ... ]
   exit
charging-action charge-by-duration
   content-id 101
   [ ... ]
   exit
rulebase standard
   [ ... ]
   route priority 1 ruledef port-80 analyzer http
   action priority 101 ruledef bbc-news charging-action free-site
   action priority 1000 ruledef catch-all charging-action charge-by-
duration
   [ ... ]
   exit
```

The following figure illustrates how ruledefs interact with the Protocol Analyzer Stack and Action Engine to produce charging records.

Figure 79. ECS In-line Service Processing

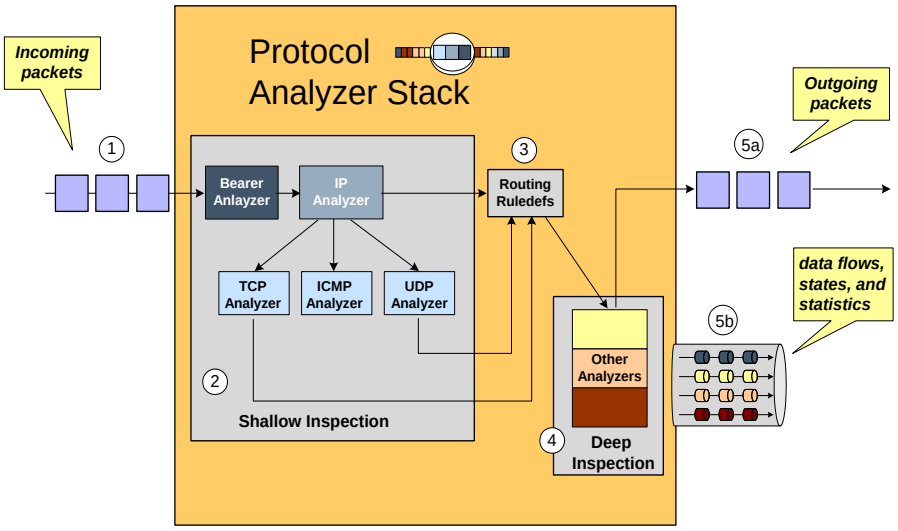


Packets entering the ECS subsystem must first pass through the Protocol Analyzer Stack where routing ruledefs apply to determine which packets to inspect. Then output from this inspection is passed to the charging engine, where charging ruledefs apply to perform actions on the output.

Routing Ruledefs and Packet Inspection

The following figure and the steps describe the details of routing ruledef application during packet inspection.

Figure 80. Routing Ruledefs and Packet Inspection



Step 1 The packet is redirected to ECS based on the ACLs in the subscriber's template/APN and packets enter ECS through the Protocol Analyzer Stack.

Step 2 Packets entering Protocol Analyzer Stack first go through a shallow inspection by passing through the following analyzers in the listed order:

Step a Bearer Analyzer

Step b IP Analyzer

Step c ICMP, TCP, or UDP Analyzer as appropriate

Important: In the current release traffic routes to the ICMP, TCP, and UDP analyzers by default. Therefore, defining routing ruledefs for these analyzers is not required.

Step 3 The fields and states found in the shallow inspection are compared to the fields and states defined in the routing ruledefs in the subscriber's rulebase.

The ruledefs' priority determines the order in which the ruledefs are compared against packets.

Step 4 When the protocol fields and states found during the shallow inspection match those defined in a routing ruledef, the packet is routed to the appropriate layer 7 or 7+ analyzer for deep-packet inspection.

Step 5 After the packet has been inspected and analyzed by the Protocol Analyzer Stack:

Step a The packet resumes normal flow and through the rest of the ECS subsystem.

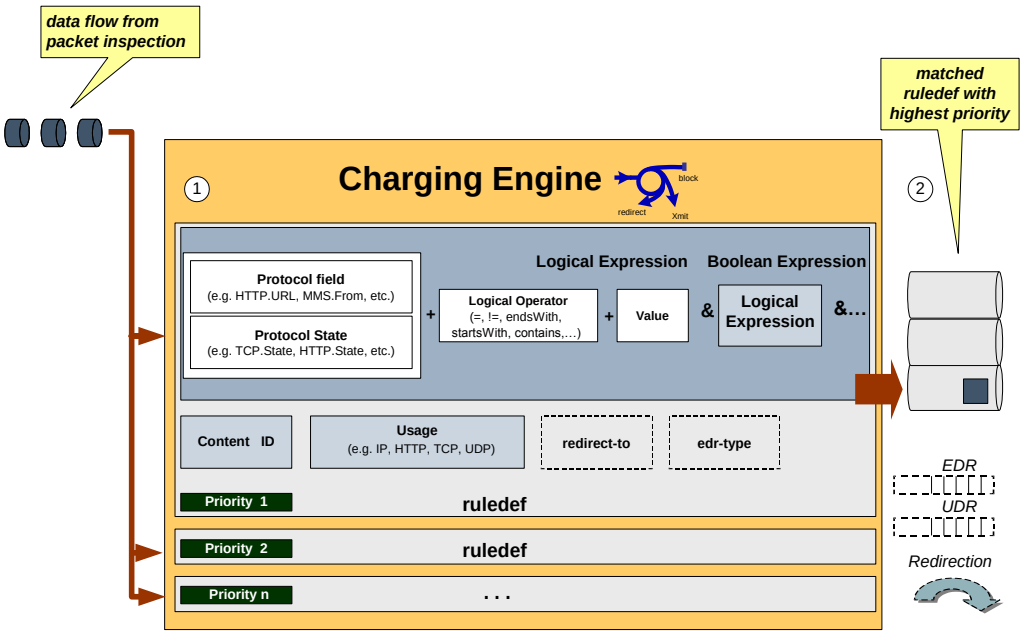
Step b The output of that analysis flows into the charging engine, where an action can be applied. Applied actions include redirection, charge value, and billing record emission.

Charging Ruledefs and the Charging Engine

This section describes details of how charging ruledefs are applied to the output from the Protocol Analyzer Stack.

The following figure and the steps that follow describe the process of charging ruledefs and charging engines.

Figure 81. Charging Ruledefs and Charging Engine



(rule-application charging)

Step 1 In the Classification Engine, the output from the deep-packet inspection is compared to the charging ruledefs. The priority configured in each charging ruledef specifies the order in which the ruledefs are compared against the packet inspection output.

Step 2 When a field or state from the output of the deep-packet inspection matches a field or state defined in a charging ruledef, the ruledef action is applied to the packet. Actions can include redirection, charge value, or billing record emission. It is also possible that a match does not occur and no action will be applied to the packet at all.

Group-of-Ruledefs

Group-of-Ruledefs enable grouping ruledefs into categories. When a group-of-ruledefs is configured in a rulebase, if any of the ruledefs within the group matches, the specified charging-action is applied, any more action instances are not processed.

A group-of-ruledefs may contain optimizable ruledefs. Whether a group is optimized or not is decided on whether all the ruledefs in the group-of-ruledefs can be optimized, and if the group is included in a rulebase that has optimization turned on, then the group will be optimized.

When a new ruledef is added, it is checked if it is included in any group-of-ruledefs, and whether it needs to be optimized, and so on.

The group-of-ruledefs configuration enables setting the application for the group (group-of-ruledefs-application parameter). When set to gx-alias the group-of-ruledefs is expanded only to extract the rule names out of it (with their original priority and charging actions) ignoring the field priority set within the group. This is just an optimization over the PCRF to PCEF interface where there exists a need to install/remove a large set of predefined rules at the same time. Though this is possible over the Gx interface (with a limit of 256) it requires a lot of PCRF resources to encode each name. This also increases the message size.

This aliasing function enables to group a set of ruledef names and provides a simple one name alias that when passed over Gx, as a Charging-Rule-Base-Name AVP, is expanded to the list of names with each rule being handled

individually. From the PCEF point of view, it is transparent, as if the PCRF had activated (or deactivated) those rules by naming each one.

Rulebase

A rulebase allows grouping one or more rule definitions together to define the billing policies for individual subscribers or groups of subscribers.

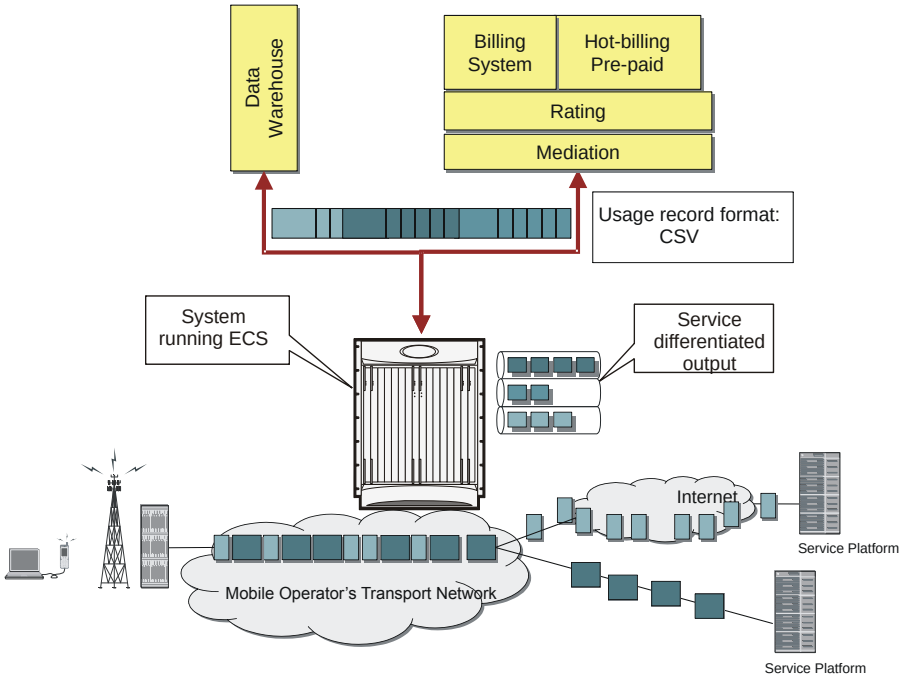
A rulebase is a collection of ruledefs and their associated billing policy. The rulebase determines the action to be taken when a rule is matched. A maximum of 512 rulebases can be specified in the ECS service.

It is possible to define a ruledef with different actions. For example, a Web site might be free for postpaid users and charge based on volume for prepaid users. Rulebases can also be used to apply the same ruledefs for several subscribers, which eliminate the need to have unique ruledefs for each subscriber.

ECS Deployment and Architecture

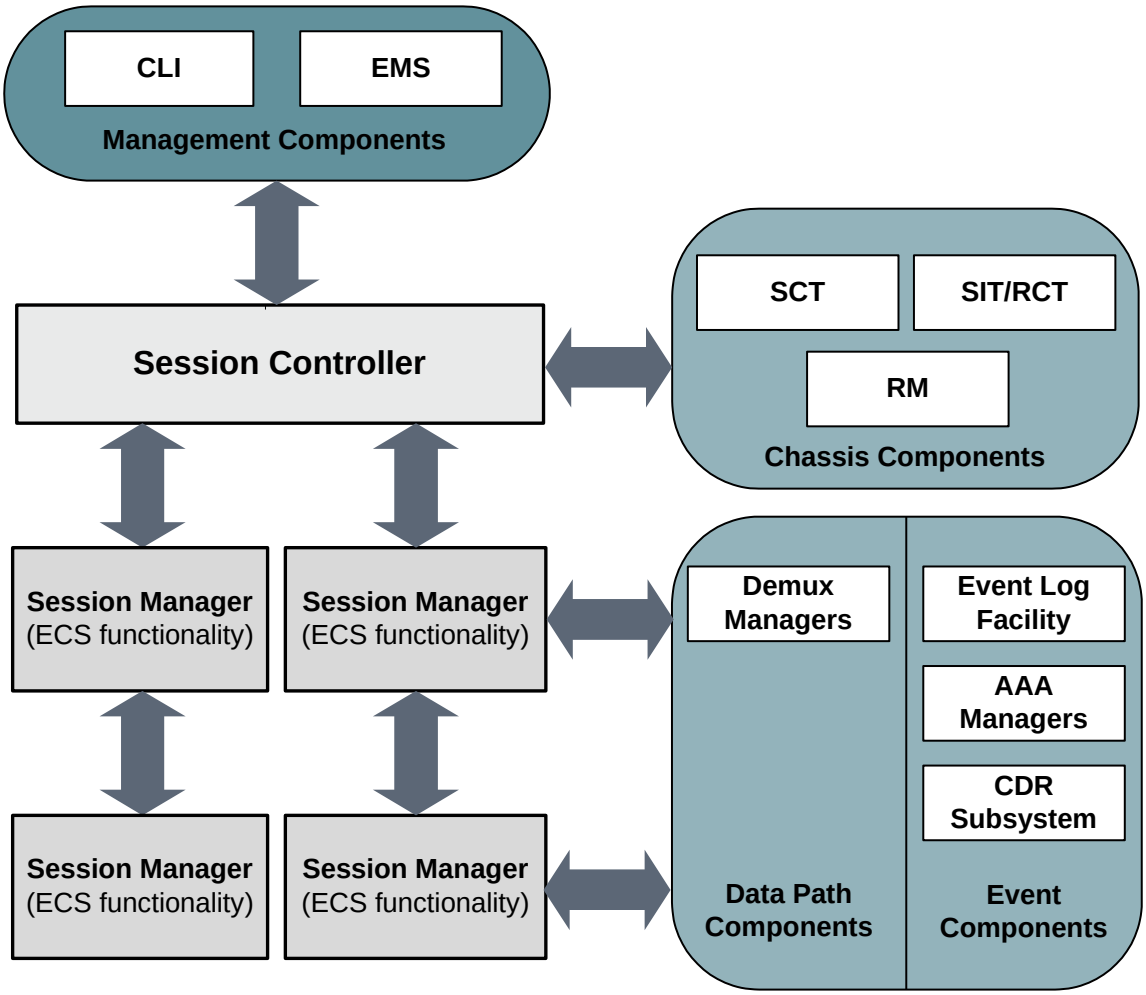
The following figure shows a typical example of ECS deployment in a mobile data environment.

Figure 82. Deployment of ECS in a Mobile Data Network



The following figure depicts the ECS architecture managed by the Session Controller (SessCtrl) and Session Manager (SessMgr) subsystems.

Figure 83. ECS Architecture



Enhanced Features and Functionality

This section describes enhanced features supported in ECS.

Important: The features described in this section may be licensed Cisco features. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Session Control in ECS

In conjunction with the Cisco ASR 5x00 chassis, the ECS provides a high-level network flow and bandwidth control mechanism in conjunction with the Session Control subsystem. ECS Session Control feature uses the interaction between SessMgr subsystem and Static Traffic Policy Infrastructure support of the chassis to provide an effective method to maximize network resource usage and enhancement of overall user experience.

This feature provides the following functionality:

- **Flow Control Functionality**—Provides the ability to define and manage the number of simultaneous IP-based sessions and/or the number of simultaneous instances of a particular application permitted for the subscriber.

  If a subscriber begins a packet data session and system is either pre-configured or receives a subscriber profile from the AAA server indicating the maximum amount of simultaneous flow for a subscriber or an application is allowed to initiate. If subscriber exceeds the limit of allowed number of flows for subscriber or type of application system blocks/redirect/discard/terminate the traffic.

  The following type of flow quotas are available for Flow Control Functionality:

  - **Subscriber-Level Session Quota**—Configurable on a per-rulebase basis
  - **Application-Level Session Quota**—Configurable on a per-charging-action basis
- **Bandwidth Control Functionality**—Allows the operator to apply rate limit to potentially bandwidth intensive and service disruptive applications.

  Using this feature the operator can police and prioritize subscribers' traffic to ensure that no single or group of subscribers' traffic negatively impacts another subscribers' traffic.

  For example, if a subscriber is running a peer-to-peer (P2P) file sharing program and the system is pre-configured to detect and limit the amount of bandwidth to the subscriber for P2P application. The system gets the quota limit for bandwidth from PDP context parameter or individual subscriber. If the subscriber's P2P traffic usage exceeds the pre-configured limit, the Session Control discards the traffic for this subscriber session.

  Session Control feature in ECS also provides the controls to police any traffic to/from a subscriber/application with the chassis.

Time and Flow-based Bearer Charging in ECS

ECS supports Time-based Charging (TBC) to charge customers on either actual consumed time or total session time usage during a subscriber session. TBC generates charging records based on the actual time difference between receiving the two packets, or by adding idle time when no packet flow occurs.

ECS also supports Flow-based Charging (FBC) based on flow category and type.

PDP context charging allows the system to collect charging information related to data volumes sent to and received by the MS. This collected information is categorized by the QoS applied to the PDP context. FBC integrates a Tariff Plane Function (TPF) to the charging capabilities that categorize the PDP context data volume for specific service data flows.

Service data flows are defined by charging rules. The charging rules use protocol characteristics such as:

- IP address
- TCP port
- Direction of flow
- Number of flows across system
- Number of flows of a particular type

FBC provides multiple service data flow counts, one each per defined service data flow. When FBC is configured in the ECS, PDP context online charging is achieved by FBC online charging using only the wildcard service data flow.

When further service data flows are specified, traffic is categorized, and counted, according to the service data flow specification. You can apply wildcard to service data flow that do not match any of the specific service data flows.

The following are the chargeable events for FBC:

- **Start of PDP context**—Upon encountering this event, a Credit Control Request (CCR) starts, indicating the start of the PDP context, is sent towards the Online Charging Service. The data volume is captured per service data flow for the PDP context.
- **Start of service data flow**—An interim CCR is generated for the PDP context, indicating the start of a new service data flow, and a new volume count for this service data flow is started.
- **Termination of service data flow**—The service data flow volume counter is closed, and an interim CCR is generated towards the Online Charging Service, indicating the end of the service data flow and the final volume count for this service data flow.
- **End of PDP context**—Upon encountering this event, a CCR stop, indicating the end of the PDP context, is sent towards the Online Charging Service together with the final volume counts for the PDP context and all service data flows.
- **Expiration of an operator configured time limit per PDP context**—This event triggers the emission of an interim CCR, indicating the elapsed time and the accrued data volume for the PDP context since the last report.
- **Expiration of an operator configured time limit per service data flow**—The service data flow volume counter is closed and an interim CCR is sent to the Online Charging Service, indicating the elapsed time and the accrued data volume since the last report for that service data flow. A new service data flow container is opened if the service data flow is still active.
- **Expiration of an operator configured data volume limit per PDP context**—This event triggers the emission of an interim CCR, indicating the elapsed time and the accrued data volume for the PDP context since the last report.
- **Expiration of an operator configured data volume limit per service data flow**—The service data flow volume counter is closed and an interim CCR is sent to the Online Charging Service, indicating the elapsed time and the accrued data volume since the last report for that service data flow. A new service data flow container is opened if the service data flow is still active.
- **Change of charging condition**—When QoS change, tariff time change are encountered, all current volume counts are captured and sent towards the Online Charging Service with an interim CCR. New volume counts for all active service data flows are started.
- **Administrative intervention** by user/service also force trigger a chargeable event.

The file naming convention for created xDRs (EDR/UDR/FDRs) are described in the Impact on xDR File Naming section.

Fair Usage

The Fair Usage feature enables resource management at two levels:

- **Instance-Level Load Balancing**—Enables load balancing of calls based on resource usage for in-line service memory allocations. If an in-line service is configured on the chassis, all resource allocation and release (memory) would involve maintaining instance-level credit usage. Sessions would be more equally distributed based on the in-line memory credits rather than the number of sessions running on individual instances.
- **Subscriber Session Resource Monitoring**—Enables monitoring individual subscriber resource usage and restricting unacceptable usage of resources by subscriber sessions. Every subscriber session will have a free ride until the operator configured threshold is reached. After that, any new resource requirement may be allowed based on the entitled services and the currently available memory in the system. Once resource allocation is failed for a subscriber session, the in-line service application requesting resource manages the failure handling.

  Operators can configure when the monitor action is initiated as a percentage of memory usage. By default, if the feature is enabled, the monitor action would be initiated at 50% threshold. Monitor action includes allowing or failing a particular resource allocation request. Any new resource allocation, once after the threshold is hit, is subject to an evaluation before allowing allocation. The requested resource is compared against the average available memory per session with a configurable per session waiver (default set to 10%). If the instance credit usage goes below a certain configurable percentage of the threshold, monitor action is disabled (default set to 5%) to avoid possible ping pong effect of enabling and disabling monitor action.

  Monitoring memory usage of individual subscriber sessions and providing preferential treatment is based on the services that the subscriber is entitled to. The subscriber session will be entitled to services as configured in the rulebase. Every subscriber session would have free ride until the operator configured threshold is reached. After that, any new resource requirement may be allowed based on the following:

  - Rulebase configured for the subscriber
  - Current available memory in the system

It is recommended that the parameters be configured only after continuous evaluation and fine tuning of the system.

Content Filtering Support

ECS provides off-line content filtering support and in-line static and dynamic content filtering support to control static and dynamic data flow and content requests.

Content Filtering Server Group Support

ECS supports external Content Filtering servers through Internet Content Adaptation Protocol (ICAP) implementation between ICAP client and Active Content Filter (ACF) server (ICAP server).

ICAP is a protocol designed to support dynamic content filtering and/or content insertion and/or modification of Web pages. Designed for flexibility, ICAP allows bearer plane nodes such as firewalls, routers, or systems running ECS to interface with external content servers such as parental control (content filtering) servers to provide content filtering service support.

In-line Content Filtering Support

Content Filtering is a fully integrated, subscriber-aware in-line service available for 3GPP and 3GPP2 networks to filter HTTP and WAP requests from mobile subscribers based on the URLs in the requests. This enables operators to filter and control the content that an individual subscriber can access, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences. Content Filtering uses Deep Packet Inspection (DPI) capabilities of ECS to discern HTTP and WAP requests.

Important: For more information on Content Filtering support, refer to the *Content Filtering Services Administration Guide*.

IP Readdressing

The IP Readdressing feature enables redirecting unknown gateway traffic based on the destination IP address of the packets to known/trusted gateways.

IP Readdressing is configured in the flow action defined in a charging action. IP readdressing works for traffic that matches particular ruledef, and hence the charging action. IP readdressing is applicable to both uplink and downlink traffic. In the Enhanced Charging Subsystem, uplink packets are modified after packet inspection, rule matching, and so on, where the destination IP/port is determined, and replaced with the readdress IP/port just before they are sent out. Downlink packets (containing the readdressed IP/port) are modified as soon as they are received, before the packet inspection, where the source IP/port is replaced with the original server IP/port number.

For one flow from an MS, if one packet is re-addressed, then all the packets in that flow will be re-addressed to the same server. Features like DPI and rule-matching remain unaffected. Each IP address + port combination will be defined as a ruledef.

In case of IP fragmentation, packets with successful IP re-assembly will be re-addressed. However, IP fragmentation failure packets will not be re-addressed.

Next-hop Address Configuration

ECS supports the ability to set the next-hop default gateway IP address as a charging action associated with any ruledef in a rulebase. This functionality provides more flexibility for service based routing allowing the next-hop default gateway to be set after initial ACL processing. This removes need for AAA to send the next-hop default gateway IP address for CC opted in subscribers.

How it works:

Step 1 The next-hop address is configured in the charging action.

Step 2 Uplink packet sent to ECS is sent for analysis.

Step 3 When the packet matches a rule and the appropriate charging action is applied, the next-hop address is picked from the charging action is copied to the packet before sending the packet to Session Manager.

Step 4 Session Manager receives the packet with the next-hop address, and uses it accordingly.

Post Processing

The Post Processing feature enables processing of packets even if the rule matching for them has been disabled. This enables all the IP/TCP packets including TCP handshaking to be accounted and charged for in the same bucket as the application flow. For example, delay-charged packets for IP Readdressing and Next-hop features.

- Readdressing of delay-charged initial hand-shaking packets.
- Sending the delay-charged initial packets to the correct next-hop address.
- DCCA—Taking appropriate action on retransmitted packets in case the quota was exhausted for the previous packet and a redirect request was sent.
  - DCCA with buffering enabled—Match CCA rules, charging-action will decide action—terminate flow/redirect
  - DCCA with buffering disabled—Match post-processing rules, and take action
- Content ID based ruledefs—On rule match, if content ID based ruledef and charging action are present, the rule is matched, and the new charging action will decide the action

A ruledef can be configured as a post-processing rule in the ruledef itself using rule-application of the ruledef. A rule can be charging, routing, or a post-processing rule. If the same ruledef is required to be a charging rule in one rulebase and a post-processing rule in another one, then two separate identical ruledefs must be defined.

How the Post-processing Feature Works

The following steps describe how the Post-processing feature works:

Step 1 Charging rule-matching is done on packets and the associated charging-action is obtained.

Step 2 Using this charging-action the disposition-action is obtained.

Step 3 If the disposition action is to either buffer or discard the packets, or if it is set by the ACF, or if there are no post-processing rules, the packets are not post processed. The disposition action is applied directly on the packets. Only if none of the above conditions is true, post processing is initiated.

Step 4 Post-processing rules are matched and the associated charging-action and then the disposition-action obtained through control-charge.

Step 5 If both match-rule and control-charge for post processing succeed, the disposition-action obtained from post-processing is applied. Otherwise, the disposition-action obtained from charging rule-matching is used.

If no disposition action is obtained by matching post-processing rules, the one obtained by matching charging-rules will be applied.

Irrespective of whether post processing is required or not, even if a single post-processing rule is configured in the rulebase, post processing will be done.

The following points should be considered while configuring post-processing rules for next-hop/readdressing.

- The rules will be L3/L4 based.
- They should be configured in post-processing rules' charging actions.

For x-header insertion, there should either be a post-processing rule whose charging-action gives no disposition-action or the packet should not match any of the post-processing rules so that the disposition action obtained from charging-rule matching is applied.

Time-of-Day Activation/Deactivation of Rules

Within a rulebase, ruledefs/groups-of-ruledefs are assigned priorities. When packets start arriving, as per the priority order, every ruledef/group-of-ruledefs in the rulebase is eligible for matching regardless of the packet arrival time. By default, the ruledefs/groups-of-ruledefs are active all the time.

The Time-of-Day Activation/Deactivation of Rules feature uses time definitions (timedefs) to activate/deactivate static ruledefs/groups-of-ruledefs such that they are available for rule matching only when they are active.

Important: The time considered for timedef matching is the system's local time.

How the Time-of-Day Activation/Deactivation of Rules Feature Works

The following steps describe how the Time-of-Day Activation/Deactivation of Rules feature enables charging according to the time of the day/time:

Step 1 Timedefs are created/deleted in the ACS Configuration Mode.

A maximum of 10 timedefs can be created in an ECS service.

Step 2 Timedefs are configured in the ACS Timedef Configuration Mode. Within a timedef, timeslots specifying the day/time for activation/deactivation of rules are configured.

A maximum of 24 timeslots can be configured in a timedef.

Step 3 In the ACS Rulebase Configuration Mode, timedefs are associated with ruledefs /groups-of-ruledefs along with the charging action.

One timedef can be used with several ruledefs/group-of-ruledefs. If a ruledef/group-of-ruledefs does not have a timedef associated with it, it will always be considered as active.

Step 4 When a packet is received, and a ruledef/group-of-ruledefs is eligible for rule matching, if a timedef is associated with the ruledef/group-of-ruledefs, before rule matching, the packet-arrival time is compared with the timeslots configured in the timedef. If the packet arrived in any of the timeslots configured in the associated timedef, rule matching is undertaken, else the next ruledef/group-of-ruledefs is considered.

This release does not support configuring a timeslot for a specific date.

If, in a timeslot, only the time is specified, that timeslot will be applicable for all days.

If for a timeslot, "start time" > "end time", that rule will span the midnight. That is, that rule is considered to be active from the current day until the next day.

If for a timeslot, "start day" > "end day", that rule will span over the current week till the end day in the next week.

In the following cases a rule will be active all the time:

- A timedef is not configured in an action priority
- A timedef is configured in an action priority, but the named timedef is not defined
- A timedef is defined but with no timeslots

URL Filtering

The URL Filtering feature simplifies using rule definitions for URL detection.

The following configuration is currently used for hundreds of URLs:

```
ruledef HTTP://AB-WAP.YZ

   www url starts-with HTTP://CDAB-SUBS.OPERA-MINI.NET/HTTP://AB-WAP.YZ

   www url starts-with HTTP://AB-WAP.YZ

   multi-line-or all-lines

   exit
```

In the above ruledef:

- The HTTP request for the URL "http://ab-wap.yz" is first sent to a proxy "http://cdab-subs.opera-mini.net/".
- The URL "http://cdab-subs.opera-mini.net/" will be configured as a prefixed URL.

Prefixed URLs are URLs of the proxies. A packet can have a URL of the proxy and the actual URL contiguously. First a packet is searched for the presence of proxy URL. If the proxy URL is found, it is truncated from the parsed information and only the actual URL (that immediately follows it) is used for rule matching and EDR generation.

The group-of-ruledefs can have rules for URLs that need to be actually searched (URLs that immediately follow the proxy URLs). That is, the group-of-prefixed-URLs will have URLs that need to be truncated from the packet information for further ECS processing, whereas, the group-of-ruledefs will have rules that need to be actually searched for in the packet.

URLs that you expect to be prefixed to the actual URL can be grouped together in a group-of-prefixed-URLs. A maximum of 64 such groups can be configured. In each such group, URLs that need to be truncated from the URL contained in the packet are specified. Each group can have a maximum of 10 such prefixed URLs. By default, all group-of-prefixed-URLs are disabled.

In the ECS rulebase, you can enable/disable the group-of-prefixed-URLs to filter for prefixed URLs.

Important: A prefixed URL can be detected and stripped if it is of the type "http://www.xyz.com/http://www.abc.com". Here, "http://www.xyz.com" will be stripped off. But in "http://www.xyz.com/www.abc.com", it cannot detect and strip off "http://www.xyz.com" as it looks for occurrence of "http" or "https" within the URL.

TCP Proxy

The TCP Proxy feature enables the ASR 5x00 to function as a TCP proxy. TCP Proxy is intended to improve ECS subsystem's functionality in case of Content Filtering, ICAP, RADIUS Prepaid, Redirection, Header Enrichment, Stateful Firewall, Application Detection and Control, DCCA, and Partial Application Headers features.

TCP Proxy along with other capabilities enables the ASR 5x00 to transparently split every TCP connection passing through it between sender and receiver hosts into two separate TCP connections, and relay data packets from the sender host to the receiver host via the split connections. This results in smaller bandwidth delay and improves TCP performance.

The TCP Proxy solution comprises of two main components:

- User-level TCP/IP Stacks — The TCP Proxy implementation uses two instances of the User Level TCP/IP stack. The stack is integrated with ECS and acts as packet receiving and sending entity. These stacks modify the behavior in which the connection is handled.

- Proxy Application — The Proxy application binds ECS, stack, and all the applications. It is the only communicating entity between the two stacks and the various applications requiring the stack. The TCP Proxy application manages the complete connection. It detects connection request, connection establishment, and connection tear-down, and propagates the same to the applications. Whenever the buffers are full, the Proxy application also buffers data to be sent later.

On an ASR 5x00 chassis, the TCP Proxy functionality can be enabled or disabled and configured from the CLI, enabling the ASR 5x00 to perform either in proxy or non-proxy mode. TCP Proxy can either be enabled for all connections regardless of the IP address, port, or application protocol involved, or for specific flows based on the configuration, for example, TCP Proxy can be enabled for some specific ports. TCP Proxy must be enabled at rulebase level. When enabled in a rulebase, it is applied on subscribers' flows using that rulebase.

TCP Proxy can be enabled in static or dynamic modes. In static mode TCP proxy is enabled for all server ports/flows for a rulebase. In the dynamic mode/Socket Migration TCP Proxy is enabled dynamically based on specified conditions. In case TCP proxy is started dynamically on a flow, the original client (MS) first starts the TCP connection with the final server. ECS keeps on monitoring the connection. Based on any rule-match/charging-action, it may happen that the connection will be proxied automatically. This activity is transparent to original client and original server. After dynamically enabling the proxy, ECS acts as TCP endpoint exactly in the same way it is when connection is statically proxied.

The functional/charging behavior of ECS for that particular connection before the dynamic proxy is started is exactly same as when there is no proxy. After the dynamic proxy is started on the connection, the functional/charging behavior of the ECS for that particular connection will be exactly similar to the ECS static proxy behavior. When the socket migration is underway, the functional/charging behavior for that particular connection is exactly the same as when there is no proxy for that flow.

TCP Proxy impacts post-recovery behavior and the charging model. With TCP Proxy, whatever packets are received from either side is charged completely. The packets that are sent out from the ECS are not considered for charging. This approach is similar to the behavior of ECS without proxy.

The following packets will be charged at ECS:

- Uplink packets received at Gn interface
- Downlink packets received at Gi interface

The following packets will not be considered for charging:

- Uplink packets forwarded/sent out by ECS/Stack on the Gi interface.
- Downlink packets forwarded/sent out by ECS/Stack on the Gn interface.

Important: After TCP Proxy is enabled for a connection, the connection will remain proxied for it's lifetime. TCP Proxy cannot be disabled for the flow.

TCP Proxy Behavior and Limitations

The following are behavioral changes applicable to various ECS features and on other applications after enabling TCP Proxy.

- TCP Proxy Model: Without TCP Proxy, for a particular flow, there is only a single TCP connection between subscriber and server. ECS is a passive entity with respect to flows and the packets received on ingress were sent out on egress side (except in case where some specific actions like drop are configured through CLI) transparently.

  With TCP Proxy, a flow is split into two TCP connections — one between subscriber and proxy and another between chassis and server.

- Ingress Data Flow to Proxy: For all uplink packets, ingress flow involves completing the following steps and then enters the Gn side TCP IP Stack of proxy:
  1. IP Analysis (support for IP reassembly)
  2. Shallow/Deep Packet TCP Analysis (support for TCP OOO)
  3. Stateful Firewall Processing
  4. Application Detection and Control Processing
  5. DPI Analysis
  6. Charging Function (including rule-matching, generation of various records, and applying various configured actions)

  For all downlink packets, ingress flow would involve completing the following steps, and then enters the Gi side TCP IP Stack of proxy:
  1. IP Analysis (support for IP reassembly)
  2. Network Address Translation Processing
  3. Shallow/Deep Packet TCP Analysis (support for TCP OOO)
  4. Stateful Firewall Processing
  5. Application Detection and Control Processing
  6. DPI Analysis
  7. Charging Function (including rule-matching, generation of various records, and applying various configured actions)
- Egress Data Flow from Proxy: All egress data flow is generated at proxy stack. For uplink packets, egress data flow would involve the following and then are sent out of the chassis:
  1. IP Analysis
  2. Shallow/Deep Packet TCP Analysis
  3. Stateful Firewall processing
  4. Network Address Translation processing

  For downlink packets, egress data flow would involve the following and then are sent out of the chassis:
  1. IP Analysis
  2. Shallow/Deep Packet TCP Analysis
  3. Stateful Firewall processing

On enabling TCP Proxy the behavior of some ECS features will get affected. For flows on which TCP Proxy is enabled it is not necessary that all the packets going out of the Gn (or Gi) interface are the same (in terms of number, size, and order) as were on Gi (or Gn).

- IP Reassembly: If the fragments are successfully reassembled then DPI analysis is done on the reassembled packet.

  Without TCP Proxy, fragmented packets will go out on the other side. With TCP proxy, normal (non-fragmented) IP packets will go out on the other side (which will not be similar to the incoming fragmented packets).

  With or without TCP Proxy, if fragment reassembly was not successful, then all the fragments will be dropped except under the case where received fragments were sufficient enough to identify the 5-tuple TCP flow and the flow had TCP Proxy disabled on it.

- TCP OOO Processing: Without TCP Proxy if it is configured to send the TCP OOO packets out (as they come), without TCP proxy such packets were sent out. With TCP Proxy, OOO packets coming from one side will go in-order on the other side. For proxied flows TCP OOO expiry timer will not be started and hence there will be no specific handling based on any such timeouts. Also, TCP OOO packets will not be sent to other side unless the packets are re-ordered.
- TCP Checksum Validation: Without TCP Proxy TCP Checksum validation is optional (configurable through "transport-layer-checksum verify-during-packet-inspection tcp" CLI command). With TCP Proxy TCP checksum is automatically done irrespective of whether the CLI command is configured or not. If the checksum validation fails, the packet is not processed further and so it does not go for application layer analysis.
- TCP Reset Packet Validation: Without TCP Proxy TCP reset packet is not validated for Seq and ACK number present in the segment and the flow is cleared immediately.

  With TCP Proxy TCP Reset packet validation is done. The flow will be cleared only if a valid TCP Reset segment is arrived. This validation is not configurable.
- TCP Timestamp (PAWS) Validation: Without TCP Proxy timestamp verification is not performed and even if there is any timestamp error, the packet is processed normally and goes for further analysis and rule-matching.

  With TCP Proxy if the connection is in established state, timestamp validation for packets is performed. If TCP timestamp is less than the previous timestamp, the packet is marked TCP error packet and is dropped. The packet is not analyzed further and not forwarded ahead. This packet should match TCP error rule (if configured). This validation is not configurable.
- TCP Error Packets: Without TCP Proxy ECS being a passive entity, most of the errors (unless configured otherwise) were ignored while parsing packets at TCP analyzer and were allowed to pass through. With TCP Proxy TCP error packets are dropped by Gi and Gn side TCP IP stack. However, since the ECS processing is already done before giving the packet to the stack, these packets are charged but not sent out by proxy on the other end.
- Policy Server Interaction (Gx): With TCP Proxy, application of policy function occurs on two separate TCP connections (non-proxy processed packets on Gn/Gi). Only external packets (the ones received from Radio and Internet) will be subject to policy enforcement at the box. This does not have any functional impact.
- Credit Control Interaction (Gy): With TCP Proxy, application of Credit Control function occur on two separate TCP connections (non-proxy processed packets on Gn/Gi). Only external packets (the ones received from Radio and Internet) will be subject to credit control at the box. This does not have any functional impact.
- DPI Analyzer: With TCP Proxy, application of DPI analyzer occurs on two separate TCP connections (non-proxy processed packets on Gn/Gi). Only external packets (the ones received from Radio and Internet) will be subject to DPI analyzer at the chassis. Any passive analyzer in the path would be buffering packet using the existing ECS infrastructure.
- ITC/BW Control: With TCP Proxy, only incoming traffic is dropped based on bandwidth calculation on ingress side packets. The BW calculation and dropping of packet is be done before sending packet to ingress TCP IP Stack. ToS and DSCP marking will be on flow level. The ToS and DSCP marking can be done only once for whole flow and once the ToS is marked for any packet either due to "ip tos" CLI command configured in the charging action or due to ITC/BW control, it will remain same for the whole flow.
- Next Hop and VLAN-ID: Without TCP Proxy nexthop feature is supported per packet, that is nexthop address can be changed for each and every packet of the flow depending on the configuration in the charging action. With TCP Proxy only flow-level next-hop will be supported. So, once the nexthop address is changed for any packet of the flow, it will remain same for the complete flow. The same is the case for VLAN-ID.
- TCP state based rules: Without TCP Proxy there is only one TCP connection for a flow and the TCP state based rules match to state of subscriber stack. With TCP Proxy there are two separate connections when TCP proxy is enabled. TCP state ("tcp state" and "tcp previous-state") based rules will match to MS state on egress side.

Two new rules (tcp proxy-state and tcp proxy-prev-state) have been added to support the existing cases (of TCP state based rules). "tcp proxy-state" and "tcp proxy-prev-state" are the state of the embedded proxy server, that is the proxy ingress-side. These rules will not be applicable if proxy is not enabled.

Using both "tcp state" and "tcp proxy-state" in the same ruledef is allowed. If proxy is enabled, they would map to Gi-side and Gn-side, respectively. If TCP Proxy is not enabled, the "tcp proxy-state" and "tcp proxy-prev-state" rules will not be matched because proxy-state will not be applicable.

Since TCP state and previous-state rules are now matched based on state on Gi side connection, ECS will not be able to support all the existing use-cases with the existing configuration. New ruledefs based on the new rules (tcp proxy-state and tcp proxy-prev-state) need to be configured to support existing use cases. Note that even by configuring using new rules; all use-cases may not be supported. For example, detection of transition from TIME-WAIT to CLOSED state is not possible now.

- TCP MSS: TCP IP Stack always inserts MSS Field in the header. This causes difference in MSS insertion behavior with and without TCP Proxy.
  - TCP CFG MSS limit-if-present: If incoming SYN has MSS option present, in outgoing SYN and SYN-ACK MSS value is limited to configured MSS value (CFG MSS)
  - TCP CFG MSS add-if-not-present: If incoming SYN does not have MSS option present, in outgoing SYN and SYN-ACK MSS configured MSS value is inserted (CFG MSS)
  - TCP CFG MSS limit-if-present add-if-not-present: If incoming SYN has MSS option present, in outgoing SYN and SYN-ACK MSS value is limited to configured MSS value (CFG MSS), OR if incoming SYN does not have MSS option present, in outgoing SYN and SYN-ACK MSS configured MSS value is inserted (CFG MSS).
- Flow Discard: Flow discard occurring on ingress/egress path of TCP Proxy would be relying on TCP-based retransmissions. Any discard by payload domain applications would result in data integrity issues as this might be charged already and it may not be possible to exclude packet. So it is recommended that applications in payload domain (like dynamic CF, CAE readdressing) should not be configured to drop packets. For example, dynamic content filtering should not be configured with drop action. If drop is absolutely necessary, it is better to use terminate action.
- DSCP/IP TOS Marking: Without TCP Proxy DSCP/IP TOS marking is supported per packet, that is IP TOS can be changed for each and every packet of the flow separately based on the configuration. With TCP Proxy flow-level DSCP/IP TOS marking is supported. So, once the IP TOS value is changed for any packet of the flow, it will remain same for the complete flow.
- Redundancy Support (Session Recovery and ICSR): Without TCP Proxy after recovery, non-syn flows are not reset. With TCP Proxy session recovery checkpointing is bypassing any proxied flows (currently on NAT flows support recovery of flows). If any flow is proxied for a subscriber, after recovery (session recovery or ICSR), if any non-syn packet is received for that subscriber, ECS sends a RESET to the sender. So, all the old flows will be RESET after recovery.
- Charging Function: Application of charging function would occur on two separate TCP connections (non proxy processed packets on Gn/Gi). Only external packets (the ones received from Radio and Internet) shall be subject to Policy enforcement at the box. Offline charging records generated at charging function would pertain to different connections hence.

X-Header Insertion and Encryption

This section describes the X-Header Insertion and Encryption features, collectively known as Header Enrichment, which enable to append headers to HTTP/WSP GET and POST request packets for use by end applications, such as mobile advertisement insertion (MSISDN, IMSI, IP address, user-customizable, and so on).

Important: In this release, the X-Header Insertion and Encryption features are supported only on the GGSN, IPSG, and P-GW.

License Requirements

X-Header Insertion and Encryption are both licensed Cisco features. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

X-Header Insertion

This section provides an overview of the X-Header Insertion feature.

Extension header (x-header) fields are the fields not defined in RFCs or standards but can be added to headers of protocol for specific purposes. The x-header mechanism allows additional entity-header fields to be defined without changing the protocol, but these fields cannot be assumed to be recognizable by the recipient. Unrecognized header fields should be ignored by the recipient and must be forwarded by transparent proxies.

The X-Header Insertion feature enables inserting x-headers in HTTP/WSP GET and POST request packets. Operators wanting to insert x-headers in HTTP/WSP GET and POST request packets, can configure rules for it. The charging-action associated with the rules will contain the list of x-headers to be inserted in the packets.

For example, if you want to insert the field *x-rat-type* in the HTTP header with a value of *rat-type*, the header inserted should be:

x-rat-type: geran

where, *rat-type* is *geran* for the current packet.

Configuring the X-Header Insertion feature involves:

Step 1 Creating/configuring a ruledef to identify the HTTP/WSP packets in which the x-headers must be inserted.

Step 2 Creating/configuring a rulebase and configuring the charging-action, which will insert the x-header fields into the HTTP/WSP packets.

Step 3 Creating/configuring the x-header format.

Step 4 Configuring insertion of the x-header fields in the charging action.

X-Header Encryption

This section provides an overview of the X-Header Encryption feature.

X-Header Encryption enhances the X-header Insertion feature to increase the number of fields that can be inserted, and also enables encrypting the fields before inserting them.

If x-header insertion has already happened for an IP flow (because of any x-header format), and if the current charging-action has the first-request-only flag set, x-header insertion will not happen for that format. If the first-request-only flag is not set in a charging-action, then for that x-header format, insertion will continue happening in any further suitable packets in that IP flow.

Changes to x-header format configuration will not trigger re-encryption for existing calls. The changed configuration will however, be applicable for new calls. The changed configuration will also apply at the next re-encryption time to

those existing calls for which re-encryption timeout is specified. If encryption is enabled for a parameter while data is flowing, since its encrypted value will not be available, insertion of that parameter will stop.

Important: Recovery of flows is not supported for this feature.

The following steps describe how X-Header Encryption works:

Step 1 X-header insertion, encryption, and the encryption certificate is configured in the CLI.

Step 2 When the call gets connected, and after each regeneration time, the encryption certificate is used to encrypt the strings.

Step 3 When a packet hits a ruledef that has x-header format configured in its charging-action, x-header insertion into that packet is done using the given x-header-format.

Step 4 If x-header-insertion is to be done for fields which are marked as encrypt, the previously encrypted value is populated for that field accordingly.

Limitations to the Header Insertion and Encryption Features

This section lists known limitations to insertion and encryption of x-header fields in HTTP and WSP headers.

The following are limitations to insertion and encryption of x-header fields in HTTP headers.

Limitations in StarOS 12.3 and earlier releases:

- The packet size is assumed to be less than "Internal MED MTU size, the size of header fields inserted". If the total length of packet exceeds the internal MTU size, header insertion will not occur after the addition of fields.
- Header insertion occurs for both HTTP GET and POST requests. However, for POST requests, the resulting packet size will likely be larger than for GET requests due to the message body contained in the request. If the previous limitation applies, then POST request will suffer a bigger limit due to this.
- Header insertion does not occur for retransmitted packets.
- Header insertion does not occur for packets with incomplete HTTP headers.
- Header insertion does not occur for TCP OOO and IP fragmented packets.
- If a flow has x-header insertion and later some IP fragments are received for which reassembly fails, sequence space of that segment will be mismatched.
- ECS does not support applying more than one modifying action on an inbound packet before sending it on the outbound interface. For example, if header insertion is applied on a packet, then the same packet is not allowed to be modified for NAT/ALG and MSS insertion.
- If a packet is buffered by DCCA/ICAP, header insertion will not occur for that packet.
- Receive window will not be considered during header enrichment. That is, after header enrichment if packet exceeds receive window, ECS will not truncate the packet.
- Packet size limit is 2400 bytes, if due to header insertion packet size exceeds this limit, behavior is unpredictable.
- Only those x-header fields in header portion of application protocol that begin with "-x" are parsed at HTTP analyzer. In URL and data portion of HTTP any field can be parsed.

The following are limitations to insertion and encryption of x-header fields in WSP headers:

- x-header fields are not inserted in IP fragmented packets.

- In case of concatenated request, x-header fields are only inserted in first GET or POST request (if rule matches for the same). X-header fields are not inserted in the second or later GET/POST requests in the concatenated requests. For example, if there is ACK+GET in packet, x-header is inserted in the GET packet. However, if GET1+GET2 is present in the packet and rule matches for GET2 and not GET1 x-header is still inserted in GET2. In case of GET+POST also, x-header is not inserted in POST.
- In case of CO, x-header fields are not inserted if the WTP packets are received out of order (even after proper re-ordering).
- If route to MMS is present, x-headers are not inserted.
- x-headers are not inserted in WSP POST packet when header is segmented. This is because POST contains header length field which needs to be modified after addition of x-headers. In segmented WSP headers, header length field may be present in one packet and header may complete in another packet.
- x-headers are not inserted in case of packets buffered at DCCA.

Accounting and Charging Interfaces

ECS supports different accounting and charging interfaces for prepaid and postpaid charging and record generation.

Important: Some feature described in this section are licensed Cisco features. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Accounting Interfaces for Postpaid Service: ECS supports the following accounting interfaces for postpaid subscribers:

- Remote Authentication Dial-In User Service (RADIUS) Interface
- GTPP Accounting Interface

Accounting and Charging Interface for Prepaid Service: ECS supports the following Credit Control Interfaces for prepaid subscribers:

- RADIUS Prepaid Credit Control interface
- Diameter Prepaid Credit Control Application (DCCA) Gy Interface
- Diameter Gx interface

Charging Records in ECS: ECS provides the following charging records for postpaid and prepaid charging:

- GGSN-Call Detail Records (G-CDRs)
- Enhanced GGSN-Call Detail Records (eG-CDRs)
- Event Detail Records (EDRs)
- Usage Detail Records (UDRs)

GTPP Accounting

ECS enables the collection of counters for different types of data traffic, and including that data in CDRs that is sent to a Charging Gateway Function (CGF).

For more information on GTPP accounting, refer to the *AAA and GTPP Interface Administration and Reference*.

RADIUS Accounting and Credit Control

The Remote Authentication Dial-In User Service (RADIUS) interface in ECS is used for the following purposes:

- **Subscriber Category Request**—ECS obtains the subscriber category from the AAA server (either prepaid or postpaid) when a new data session is detected. The AAA server used for the subscriber category request can be different from the AAA server used for service authorization and accounting.
- **Service Access Authorization**—ECS requests access authorization for a specific subscriber and a newly detected data session. The AAA server is the access Policy Decision Point and the ECS the Policy Enforcement Point.
- **On-line Service Accounting (Prepaid)**—ECS reports service usage to the AAA server. The AAA server acts as a prepaid control point and the ECS as the client. Accounting can be applied to a full prepaid implementation

or just to keep ECS updated of the balance level and trigger a redirection if the subscriber balance reaches a low level.

Diameter Accounting and Credit Control

The Diameter Credit Control Application (DCCA) is used to implement real-time online or offline charging and credit control for a variety of services, such as network access, messaging services, and download services.

In addition to being a general solution for real-time cost and credit control, DCCA includes these features:

- **Real-time Rate Service Information**—DCCA can verify when end subscribers' accounts are exhausted or expired; or deny additional chargeable events.
- **Support for Multiple Services:** DCCA supports the usage of multiple services within one subscriber session. Multiple service support includes:
  - The ability to identify and process the service or group of services that are subject to different cost structures.
  - Independent credit control of multiple services in a single credit control sub-session.

Gx Interface Support

The Gx interface is used in IMS deployment in GPRS/UMTS networks. Gx interface support on the system enables wireless operators to intelligently charge the services accessed depending on the service type and parameters with rules. It also provides support for IP Multimedia Subsystem (IMS) authorization in a GGSN service. The goal of the Gx interface is to provide network-based QoS control as well as dynamic charging rules on a per bearer basis for an individual subscriber. The Gx interface is in particular needed to control and charge multimedia applications.


Important: For more information on Gx interface support, see the *Gx Interface Support* appendix in the administration guide for the product that you are deploying.

Gy Interface Support

The Gy interface provides a standardized Diameter interface for real-time content-based charging of data services. It is based on the 3GPP standards and relies on quota allocation.

It provides an online charging interface that works with the ECS Deep Packet Inspection feature. With Gy, customer traffic can be gated and billed in an "online" or "prepaid" style. Both time- and volume-based charging models are supported. In all these models, differentiated rates can be applied to different services based on shallow or deep-packet inspection.

Gy is a Diameter interface. As such, it is implemented atop, and inherits features from, the Diameter Base Protocol. The system supports the applicable base network and application features, including directly connected, relayed or proxied DCCA servers using TLS or plain text TCP.

In the simplest possible installation, the system will exchange Gy Diameter messages over Diameter TCP links between itself and one "prepay" server. For a more robust installation, multiple servers would be used. These servers may optionally share or mirror a single quota database so as to support Gy session failover from one server to the other. For a more scalable installation, a layer of proxies or other Diameter agents can be introduced to provide features such as multi-path message routing or message and session redirection features.

The Diameter Credit Control Application (DCCA) which resides as part of the ECS manages the credit and quota for a subscriber.

Important: For more information on Gy interface support, see the *Gy Interface Support* appendix in the administration guide for the product that you are deploying.

Event Detail Records (EDRs)

Event Detail Records (EDRs) are usage records with support to configure content information, format, and generation triggers by the system administrative user.

EDRs are generated according to explicit action statements in rule commands. Several different EDR schema types, each composed of a series of analyzer parameter names, are specified in EDR. EDRs are written at the time of each event in CSV format. EDRs are stored in timestamped files that can be downloaded via SFTP from the configured context.

EDRs are generated on per flow basis, and as such they catch whatever bytes get transmitted over that flow including retransmitted.

EDR format

The EDRs can be generated in comma separated values (CSV) format as defined in the traffic analysis rules.

Important: In EDRs, the maximum field length for normal and escaped strings is 127 characters. If a field's value is greater than 127 characters, in the EDR it is truncated to 127 characters.

Flow-overflow EDR

Flow-overflow EDR or Summary FDR is a feature to count the data bytes from the subscriber that are missed due to various reasons in ECS.

In case any condition that affects the callline (FLOW end-condition like hagr, handoff) occurs, flow-overflow EDR generation is enabled, an extra EDR is generated. Based on how many bytes/packets were transferred from/to the subscriber for which ECS did not allocate data session. This byte/packet count is reflected in that extra EDR. This extra EDR is nothing but "flow-overflow" EDR or Summary FDR.

The extra EDR is generated if all of the following is true:

- Subscriber affecting condition occurs (session-end, hand-off, hagr)
- Flow-overflow EDR generation is enabled
- EDR generation on session-end, hand-off or hagr is enabled
- Number of bytes/packets for flow-overflow EDR is non-zero.

The bytes/packet count will be printed as a part of "sn-volume-amt" attribute in the EDR. Hence, this attribute must be configured in the EDR format.

EDR Generation in Flow-end and Transaction Complete Scenarios with sn-volume Fields

"sn-volume-amt" counters will be re-initialized only when the fields are populated in EDRs. For example, consider the following two EDR formats:

```
edr-format edr1

   rule-variable http url priority 10
```

```
        attribute sn-volume-amt ip bytes uplink priority 500
        attribute sn-volume-amt ip bytes downlink priority 510
        attribute sn-volume-amt ip pkts uplink priority 520
        attribute sn-volume-amt ip pkts downlink priority 530
        attribute sn-app-protocol priority 1000
        exit
    edr-format edr2
        rule-variable http url priority 10
        attribute sn-app-protocol priority 1000
        exit
```

"sn-volume-amt counters" will be re-initialized only if these fields are populated in the EDRs. Now if edr2 is generated, these counters will not be re-initialized. These will be re-initialized only when edr1 is generated. Also, note that only those counters will be re-initialized which are populated in EDR. For example, in the following EDR format:

```
    edr-format edr3
        rule-variable http url priority 10
        attribute sn-volume-amt ip bytes uplink priority 500
        attribute sn-volume-amt ip bytes downlink priority 510
        attribute sn-app-protocol priority 1000
        exit
```

If edr3 is generated, only uplink bytes and downlink bytes counter will be re-initialized and uplink packets and downlink packets will contain the previous values till these fields are populated (say when edr1 is generated).

For the voice call duration for SIP reporting requirements, ECS SIP analyzer keeps timestamp of the first INVITE that it sees. It also keeps a timestamp when it sees a 200 OK for a BYE. When this 200 OK for a BYE is seen, SIP analyzer triggers creation of an EDR of type ACS_EDR_VOIP_CALL_END_EVENT. This will also be triggered at the time of SIP flow termination if no 200 OK for BYE is seen. In that case, the last packet time will be used in place of the 200 OK BYE timestamp. The EDR generation logic calculates the call duration based on the INVITE and end timestamps, it also accesses the child RTP/RTCP flows to calculate the combined uplink/downlink bytes/packets counts and sets them in the appropriate fields.

Usage Detail Records (UDRs)

Usage Detail Records (UDRs) contain accounting information based on usage of service by a specific mobile subscriber. UDRs are generated based on the content-id for the subscriber, which is part of charging action. The fields required as part of usage data records are configurable and stored in the System Configuration Task (SCT).

UDRs are generated on any trigger of time threshold, volume threshold, handoffs, and call termination. If any of the events occur then the UDR subsystem generates UDRs for each content ID and sends to the CDR module for storage.

UDR format

The UDRs are generated in Comma Separated Values (CSV) format as defined in the traffic analysis rules.

Charging Methods and Interfaces

Prepaid Credit Control

Prepaid billing operates on a per content-type basis. Individual content-types are marked for prepaid treatment. A match on a traffic analysis rule that has a prepaid-type content triggers prepaid charging management.

In prepaid charging, ECS performs the metering function. Credits are deducted in real time from an account balance or quota. A fixed quota is reserved from the account balance and given to the system by a prepaid rating and charging server, which interfaces with an external billing system platform. The system deducts volume from the quota according to the traffic analysis rules. When the subscriber's quota gets to the threshold level specified by the prepaid rating and charging server, system sends a new access request message to the server and server updates the subscriber's quota. The charging server is also updated at the end of the call.

ECS supports the following credit control applications for prepaid charging:

- **RADIUS Credit Control Application**—RADIUS is used as the interface between ECS and the prepaid charging server. The RADIUS Prepaid feature of ECS is separate to the system-level Prepaid Billing Support and that is covered under a different license key.
- **Diameter Credit Control Application**—The Diameter Credit Control Application (DCCA) is used to implement real-time credit control for a variety of services, such as networks access, messaging services, and download services.

In addition to being a general solution for real-time cost and credit control, DCCA includes the following features:

- **Real-time Rate Service Information**—DCCA can verify when end subscribers' accounts are exhausted or expired; or deny additional chargeable events.
- **Support for Multiple Services**—DCCA supports the usage of multiple services within one subscriber session. Multiple service support includes:
  - The ability to identify and process the service or group of services that are subject to different cost structures.
  - Independent credit control of multiple services in a single credit control sub-session.

Postpaid

In a postpaid environment, the subscribers pay after use of the service. AAA/RADIUS server is responsible for authorizing network nodes to grant access to the user, and the CDR system generates G-CDRs/eG-CDRs/EDRs/UDRs for billing information on pre-defined intervals of volume or per time.

Important: G-CDRs and eG-CDRs are only available in UMTS networks.

ECS also supports FBC and TBC methods for postpaid billing. For more information on FBC and TBC in ECS, see the Time and Flow-based Bearer Charging in ECS section.

Prepaid Billing in ECS

In a prepaid environment, the subscribers pay for service prior to use. While the subscriber is using the service, credit is deducted from subscriber's account until it is exhausted or call ends. The prepaid charging server is responsible for authorizing network nodes to grant access to the user, as well as grant quotas for either time connected or volume used. It is up to the network node to track the quota use, and when these use quotas run low, the network node sends a request to the prepaid server for more quota.

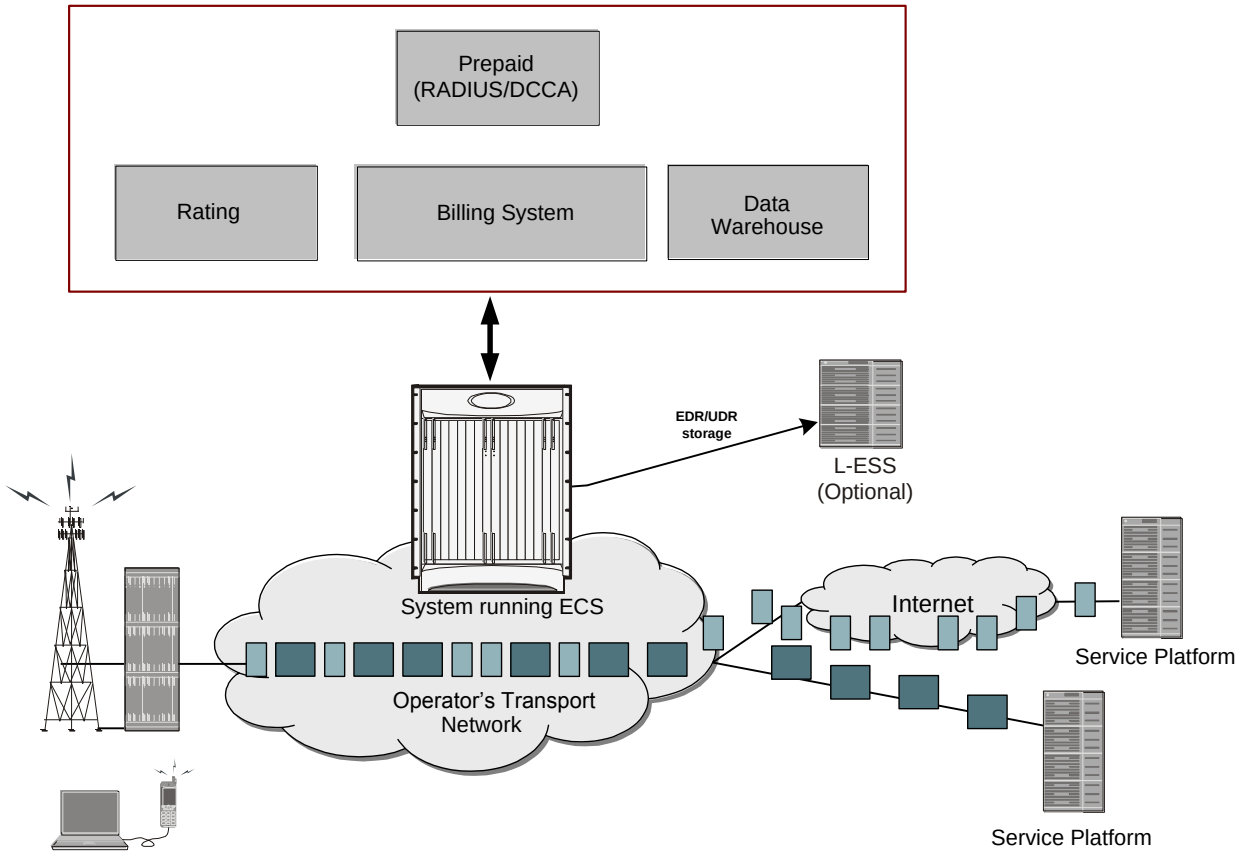
If the user has not used up the purchased credit, the server grants quota and if no credit is available to the subscriber the call will be disconnected. ECS and DCCA manage this functionality by providing the ability to set up quotas for different services.

Prepaid quota in ECS is implemented using RADIUS and DCCA as shown in the following figure.

How ECS Prepaid Billing Works

The following figure illustrates a typical prepaid billing environment with system running ECS.

Figure 84. Prepaid Billing Scenario with ECS



Credit Control Application (CCA) in ECS

This section describes the credit control application that is used to implement real-time credit-control for a variety of end user services such as network access, Session Initiation Protocol (SIP) services, messaging services, download services, and so on. It provides a general solution to the real-time cost and credit control.

CCA with RADIUS or Diameter interface uses a mechanism to allow the user to be informed of the charges to be levied for a requested service. In addition, there are services such as gaming and advertising that may debit from a user account.

How Credit Control Application (CCA) Works for Prepaid Billing

The following figure and steps describe how CCA works with in a GPRS/UMTS or CDMA-2000 network for prepaid billing.

Figure 85. Prepaid Charging in GPRS/UMTS/CDMA-2000 Networks

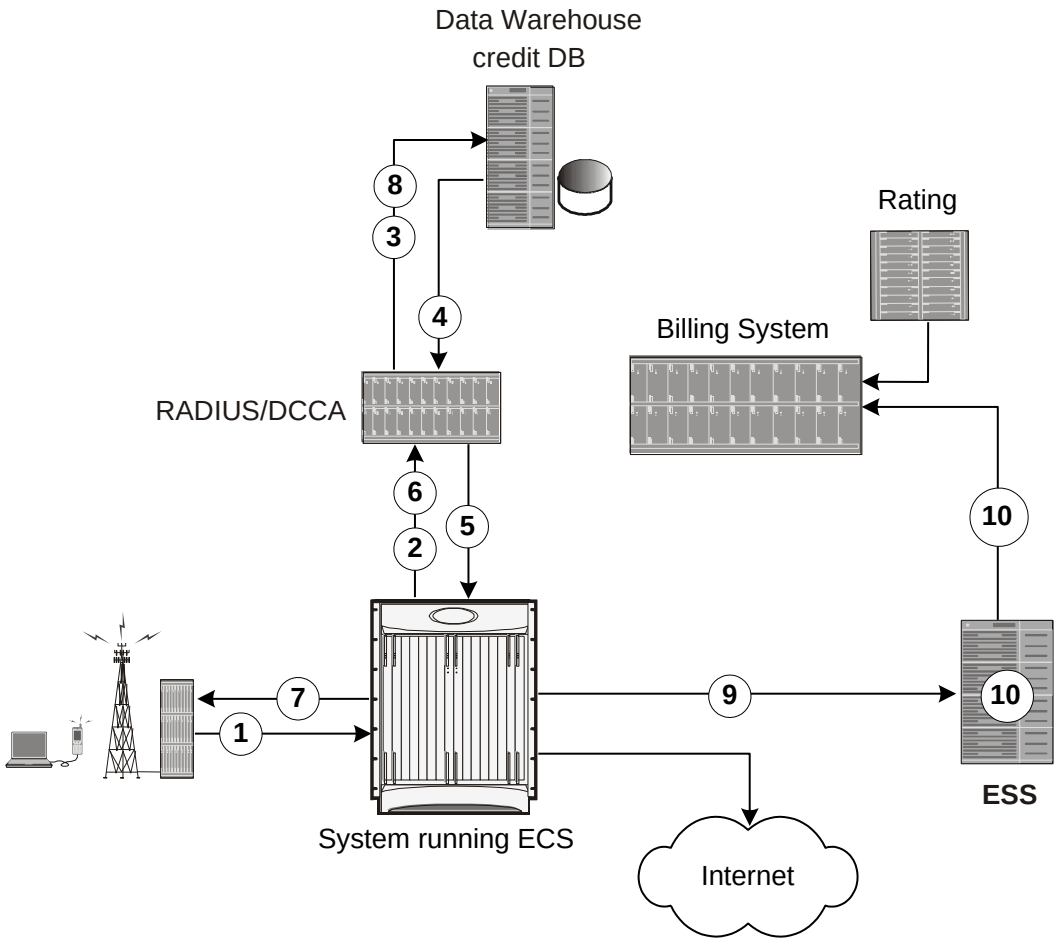


Table 55. Prepaid Charging in GPRS/UMTS/CDMA-2000 Networks

Step No.	Description
1	Subscriber session starts.
2	System sends request to CCA for subscriber's quota.
3	CCA sends request to Data Warehouse (DW) credit quota for subscriber.
4	Credit Database in DW sends pre-configured amount of usage limit from subscriber's quota to CCA. To reduce the need for multiple requests during subscriber's session configured amount of usage limit a major part of available credit quota for subscriber is set.
5	CCA sends the amount of quota required to fulfill the subscriber's initial requirement to the system.
6	When the initial amount of quota runs out, system sends another request to the CCA and the CCA sends another portion of available credit quota.
7	Subscriber session ends after either quota exhausts for subscriber or subscriber terminates the session.
8	CCA returns unused quota to DW for update to subscribers Credit DB.
9	EDRs and UDRs are periodically SFTPd from system memory to the ESS/external storage, if deployed or to billing system directly as they are generated. Or, if configured, pushed to the ESS/external storage at user-configurable intervals.
10	The ESS/external storage periodically sends records to the billing system or charging reporting and analysis system.

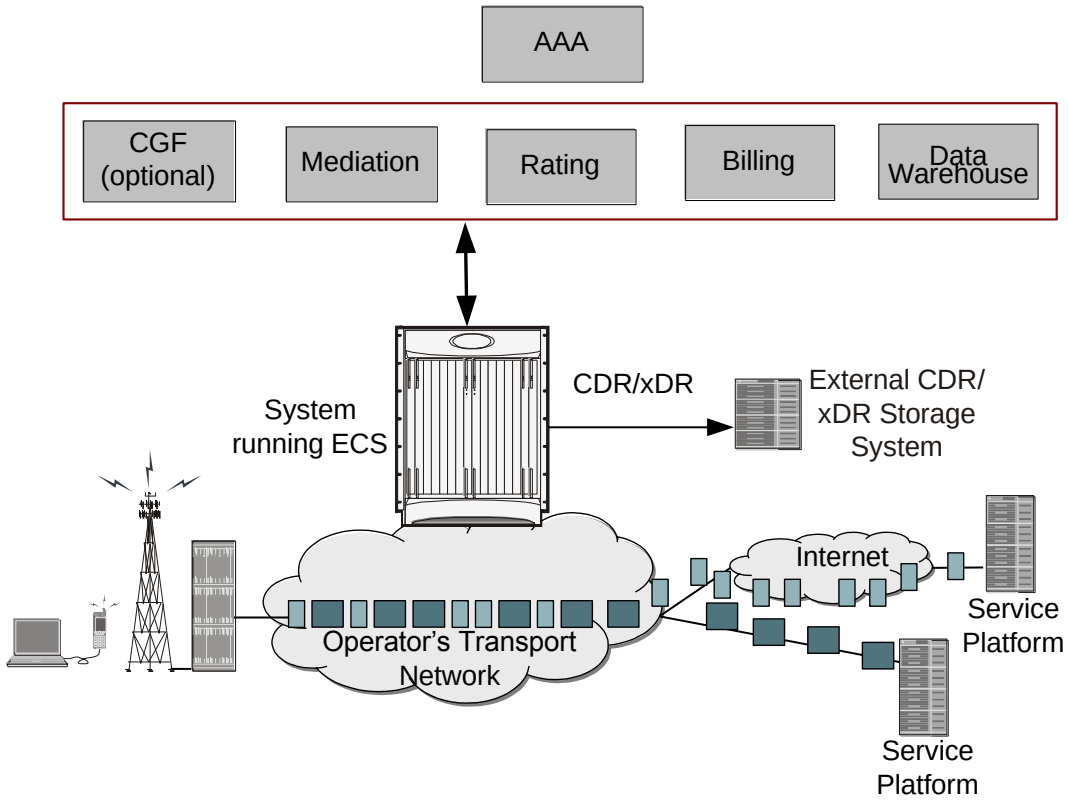
Important: For information on ESS contact your Cisco account representative.

Postpaid Billing in ECS

This section describes the postpaid billing that is used to implement off-line billing processing for a variety of end user services.

The following figure shows a typical deployment of ECS for postpaid billing system.

Figure 86. Postpaid Billing System Scenario with ECS



How ECS Postpaid Billing Works

This section describes how the ECS postpaid billing works in the GPRS/UMTS and CDMA-2000 Networks.

ECS Postpaid Billing in GPRS/UMTS Networks

The following figure and steps describe how ECS works in a GPRS/UMTS network for postpaid billing.

Figure 87. Postpaid Billing with ECS in GPRS/UMTS Network

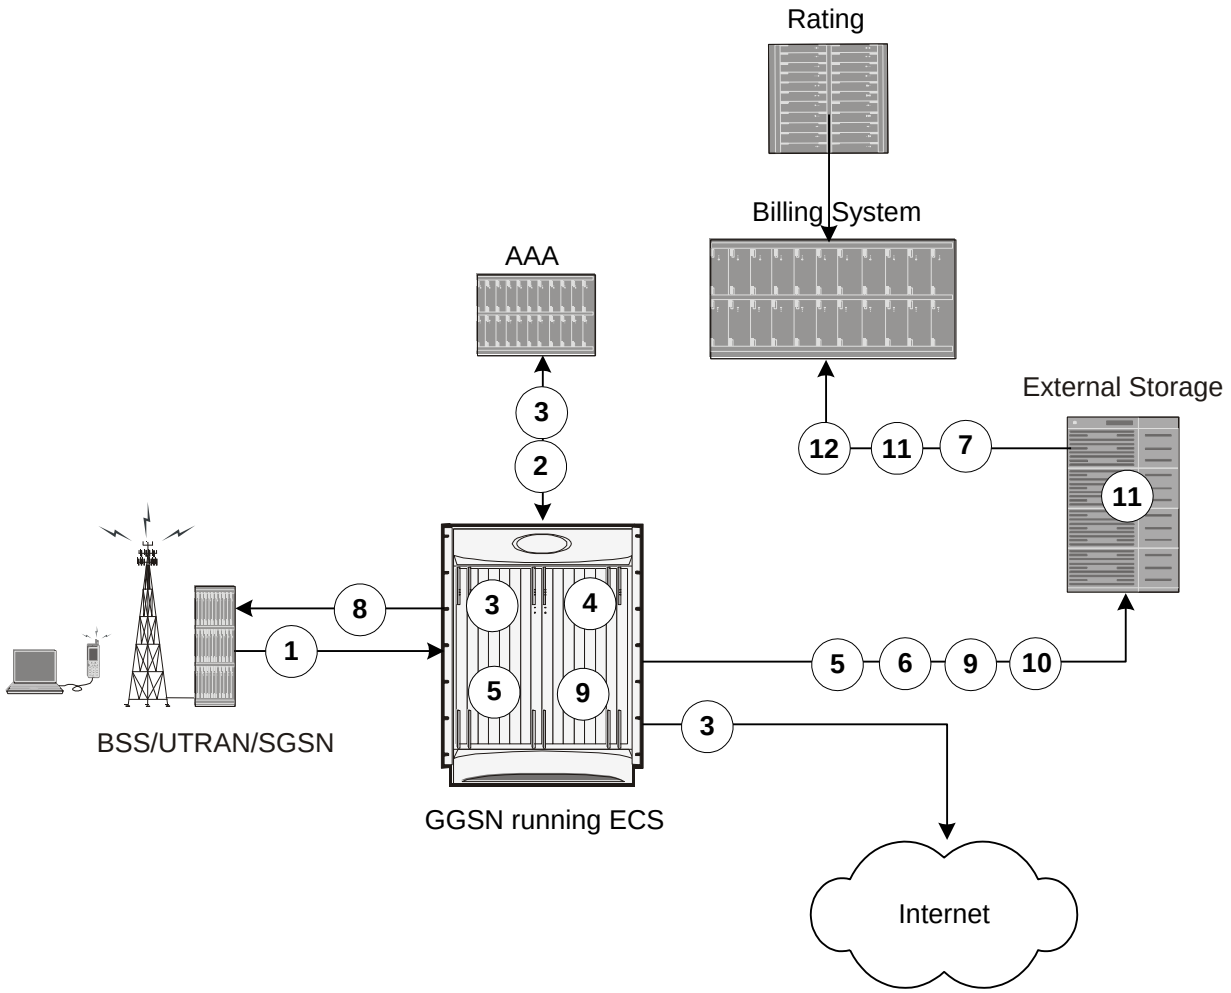


Table 56. Postpaid Billing with ECS in GPRS/UMTS Network

Step No.	Description
1	The subscriber initiates the session.
2	After subscriber authentication and authorization, the system starts the session.
3	Data packet flow and accounting starts.
4	System periodically generates xDRs and stores them to the system memory.
5	System generates G-CDRs/eG-CDRs and sends them to billing system as they are generated.
6	The billing system picks up the CDR files periodically.
7	Subscriber session ends after subscriber terminates the session.

Step No.	Description
8	The system stores the last of the xDRs to the system memory and final xDRs are SFTPd from system memory to ESS/external storage, if deployed or to billing system directly.
9	System sends the last of the G-CDRs/eG-CDRs to the billing system.
10	File Generation Utility, FileGen in external storage periodically runs to generate G-CDRs/eG-CDRs files for billing system and send them to the billing system.
11	The billing system picks up the xDR files from the ESS/external storage periodically.

Postpaid Billing in CDMA-2000 Networks

The following figure and steps describe how ECS works within a CDMA-2000 network for postpaid billing.

Figure 88. Postpaid Billing with ECS in CDMA-2000 Network

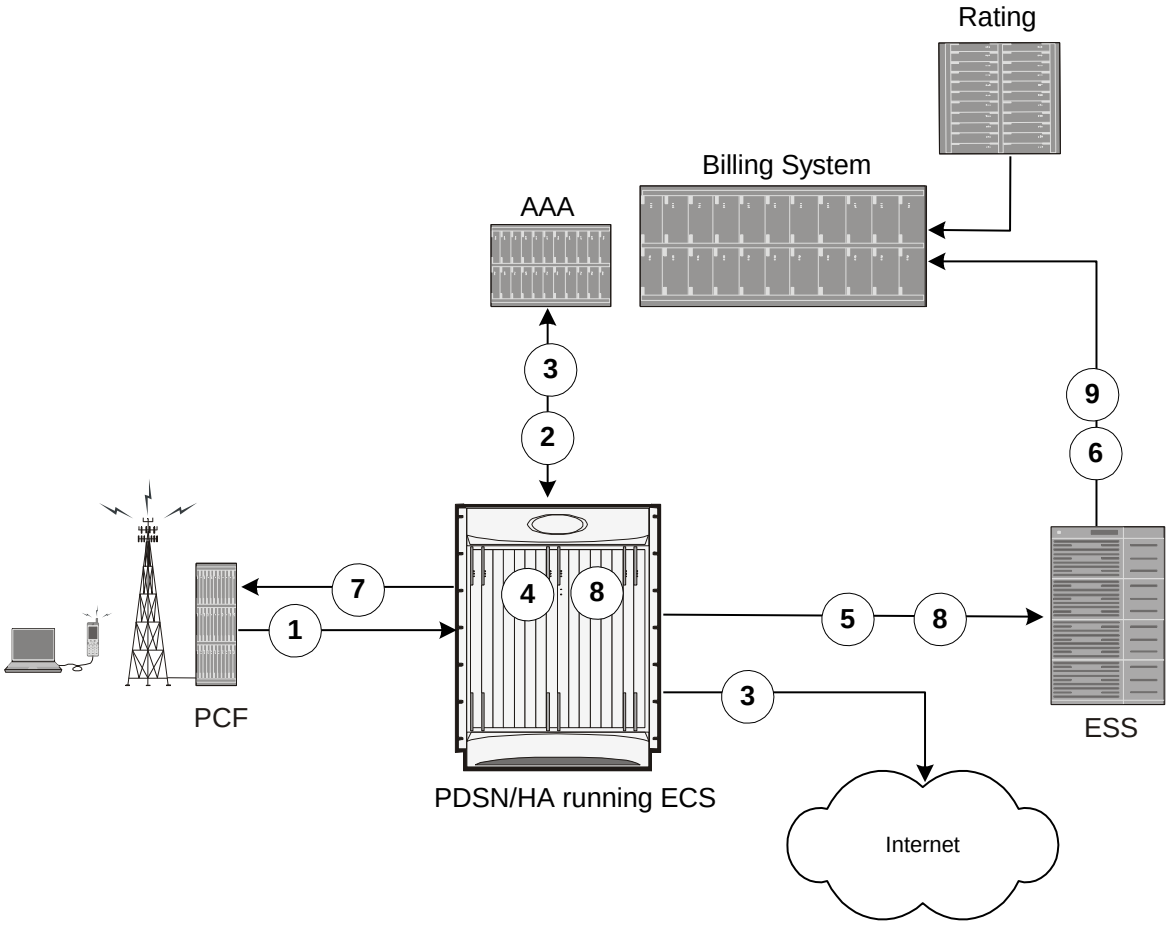


Table 57. Postpaid Billing with ECS in GPRS/UMTS Network

Step No.	Description
1	The subscriber initiates the session.
2	After subscriber authentication and authorization, the system starts the session.
3	Data packet flow and accounting starts.
4	System periodically generates xDRs and stores them to the system memory.
5	EDRs/UDRs are periodically SFTPd from system memory to ESS/external storage, if deployed or to billing system directly as they are generated.
6	The billing system picks up the xDR files from the ESS/external storage periodically.
7	Subscriber session ends after subscriber terminates the session.
8	The system stores the last of the xDRs to the system memory and final xDRs are SFTPd from system memory to the ESS/external storage, if deployed or to billing system directly.
9	The ESS/external storage finally sends xDRs to the billing system.

External Storage System

Important: For information on availability/support for ESS, contact your Cisco account representative.

The External Storage System (ESS) is a high availability, fault tolerant, redundant solution for short-term storage of files containing detail records (UDRs/EDRs/FDRs (xDRs)). To avoid loss of xDRs on the chassis due to overwriting, deletion, or unforeseen events such as power or network failure or unplanned chassis switchover, xDRs are off-loaded to ESS for storage and analysis to avoid loss of charging and network analysis information contained in the xDRs.

The xDR files can be pulled by the ESS from the chassis, or the chassis can push the xDR files to the ESS using SFTP protocol. In the Push mode, the ESS URL to which the CDR files need to be transferred to is specified. The configuration allows a primary and a secondary server to be configured. Configuring the secondary server is optional. Whenever a file transfer to the primary server fails for four consecutive times, the files will be transferred to the secondary server. The transfer will switch back to the original primary server when:

- Four consecutive transfer failures to the secondary server occur
- After switching from the primary server, 30 minutes elapses

In the push transfer mode, the following can be configured:

- Transfer interval—A time interval, in seconds, after which the CDRs are pushed to the configured IP periodically. All the files that are completed before the PUSH timer expires are pushed.
- Remove file after transfer—An option to keep or remove the CDR files on the hard disk after they are transferred to the ESS successfully.

The system running with ECS stores xDRs on an ESS, and the billing system collects the xDRs form the ESS and correlates them with the AAA accounting messages using 3GPP2 -Correlation-IDs (for PDSN) or Charging IDs (for GGSN).

Important: For more information on the ESS, refer to the *ESS Installation and Administration Guide*.

System Resource Allocation

ECS does not require manual resource allocation. The ECS subsystem automatically allocates the resources when ECS is enabled on the chassis. ECS must be enabled on the chassis before configuring services.

Redundancy Support in ECS

This section describes the redundancy support available in ECS to recover user sessions and charging records in the event of software/hardware failure.

Caution: Persistent data flows are NOT recoverable during session recovery.

Important: Redundancy is not available in the current version of the Cisco XT2 platform.

Intra-chassis Session Recovery Interoperability

Intra-chassis session recovery is coupled with SessMgr recovery procedures.

Intra-chassis session recovery support is achieved by mirroring the SessMgr and AAAMgr processes. The SessMgrs are paired one-to-one with the AAAMgrs. The SessMgr sends checkpointed session information to the AAAMgr. ECS recovery is accomplished using this checkpointed information.

Important: In order for session recovery to work there should be at least four packet processing cards, one standby and three active. Per active CPU with active SessMgrs, there is one standby SessMgr, and on the standby CPU, the same number of standby SessMgrs as the active SessMgrs in the active CPU.

There are two modes of session recovery, one from task failure and another on failure of CPU or packet processing card.

Recovery from Task Failure

When a SessMgr failure occurs, recovery is performed using the mirrored "standby-mode" SessMgr task running on the active packet processing card. The "standby-mode" task is renamed, made active, and is then populated using checkpointed session information from the AAAMgr task. A new "standby-mode" SessMgr is created.

Recovery from CPU or Packet Processing Card Failure

When a PSC, PSC2, or PPC hardware failure occurs, or when a planned packet processing card migration fails, the standby packet processing card is made active and the "standby-mode" SessMgr and AAAMgr tasks on the newly activated packet processing card perform session recovery.

Inter-chassis Session Recovery Interoperability

The system supports the simultaneous use of ECS and the Inter-chassis Session Recovery feature. (For more information on the Inter-chassis Session Recovery feature, refer to the *System Administration Guide*.) When both features are enabled, ECS session information is regularly checkpointed from the active chassis to the standby as part of normal Service Redundancy Protocol processes.

In the event of a manual switchover, there is no loss of accounting information. All xDR data from the active chassis is moved to a customer-configured ESS before switching over to the standby. This data can be retrieved at a later time.

Upon completion of the switchover, the ECS sessions are maintained and the "now-active" chassis recreates all of the session state information including the generation of new xDRs.

In the event of an unplanned switchover, all accounting data that has not been written to the external storage is lost. (Note that either the ESS can pull the xDR data from the chassis, or the chassis can push the xDR files to a configured ESS at user-configured intervals. For more information, see External Storage System section.) Upon completion of switchover, the ECS sessions are maintained and the "now-active" chassis recreates all of the session state information including the generation of new xDRs.

Regardless of the type of switchover that occurred, the names of the new xDR files will be different from those stored in the /records directory of packet processing card RAM on the "now-standby" chassis. Also, in addition to the file name, the content of many of the fields within the xDR files created by the "now-active" chassis will be different. ECS manages this impact with recovery mechanism. For more information on the differences and how to correlate the two files and other recovery information, see the Impact on xDR File Naming section.

Inter-chassis Session Recovery Architecture

Inter-chassis redundancy in ECS uses Flow Detail Records (FDRs) and UDRs to manage the switchover between Active-Standby system. xDRs are moved between redundant external storage server and Active-Standby systems.

Impact on xDR File Naming

The xDR file name is limited to 256 characters with the following syntax:

basename_ChargSvcName_timestamp_SeqNumResetIndicator_FileSeqNumber

where:

- *basename*—A global configurable text string that is unique per system that uniquely identifies the global location of the system running ECS.
- *ChargSvcName*—A system context-based configurable text string that uniquely identifies a specific context-based charging service.
- *timestamp*—Date and time at the instance of file creation. Date and time in the form of "MMDDYYYYHHmmSS" where HH is a 24-hour value from 00-23.
- *SeqNumResetIndicator*—A one-byte counter used to discern the potential for duplicated FileSeqNumber with a range of 0 through 255, which is incremented by a value of 1 for the following conditions:
  - Failure of an ECS software process on an individual packet processing card
  - Failure of a system such that a second system takes over according to the Inter-chassis Session Recovery feature
  - File Sequence Number (FileSeqNumber) rollover from 999999999 to 0
- *FileSeqNumber*—Unique file sequence number for the file with nine-digit integer having range from 000000000 to 999999999. It is unique on each system.

With inter-chassis session recovery, only the first two fields in the xDR file names remain consistent between the active and standby chassis as these are parameters that are configured locally on the chassis. Per inter-chassis session recovery implementation requirements, the two chassis systems must be configured identically for all parameters not associated with physical connectivity to the distribution node.

The fields "timestamp", "SeqNumResetIndicator", and "FileSeqNumber" are all locally generated by the specific system through CDR subsystem, regardless of whether they are in an Inter-chassis Session Recovery arrangement or not.

- The "timestamp" value is unique to the system generating the actual xDRs and generated at the time the file is opened on the system.
- The SeqNumResetIndicator is a unique counter to determine the number of resets applied to FileSeqNumber. This counter is generated by CDR subsystem and increment the counter in event of resets in FileSeqNumber. This is required as "timestamp" field is not sufficient to distinguish between a unique and a duplicate xDR.

As such, the "SeqNumResetIndicator" field is used to distinguish between xDR files which have the same "FileSeqNumber" as a previously generated xDR as a result of:

- Normal operation, for example a rollover of the "FileSeqNumber" from maximum limit to 0.
- Due to a failure of one of the ECS processes running on a packet processing card card.
- Failure of the system (that is, Inter-chassis Session Recovery switchover).

In any scenario where the "FileSeqNumber" is reset to 0, the value of the "SeqNumResetIndicator" field is incremented by 1.

- The value of the "FileSeqNumber" is directly linked to the ECS process that is generating the specific xDRs. Any failure of this specific ECS process results in resetting of this field to 0.

Impact on xDR File Content

The following scenarios impact the xDR file content:

- On failure of an active chassis:

  On system startup, xDR files are generated in accordance with the standard processes and formats. If the system fails at any time it results in an inter-chassis session recovery switchover from active to standby and the following occurs depending on the state of the call/flow records and xDR file at the time of failure:

  - Call/flow records that were being generated and collected in system memory prior to being written out to /records directory on packet processing card RAM are not recoverable and therefore are lost.
  - Closed xDRs that have been written out to records directory on packet processing card RAM but that have yet to be retrieved by the ESS are recoverable.
  - Closed xDRs that have been retrieved and processed by the ESS have no impact.

- On the activation of a Standby chassis:

  Upon detection of a failure of the original active chassis, the standby chassis transits to the active state and begins serving the subscriber sessions that were being served by the now failed chassis. Any subsequent new subscriber session will be processed by this active chassis and will generate xDRs per the standard processes and procedures.

  However, this transition impacts the xDRs for those subscribers that are in-progress at the time of the transition. For in progress subscribers, a subset of the xDR fields and their contents are carried over to the newly active chassis via the SRP link. These fields and their contents, which are carried over after an Inter-chassis Session Recovery switchover, are as follows:

  - HA-CORRELATION-ID
  - PDSN-CORRELATION-ID (PDSN only)
  - PDSN-NAS-IP-ADDRESS (PDSN only)
  - PDSN-NAS-ID (PDSN only)
  - USERNAME
  - MSID

- RADIUS-NAS-IP-ADDRESS

All remaining fields are populated in accordance with the procedures associated with any new flow with the exceptions that, the field "First Packet Direction" is set to "Unknown" for all in-progress flows that were interrupted by the switchover and the field "FDR Reason" is marked as a PDSN Handoff and therefore is set to a value of "1" and corresponding actions are taken by the billing system to assure a proper and correct accounting of subscriber activities.

Chapter 13
External Storage System Overview

The External Storage System (ESS) is used to collect, store, and report billing information from the Enhanced Charging Service running on the ASR 5x00 chassis. This guide contains information on installing, configuring, and maintaining the ESS.

This chapter consists of the following topics:

- ESS Overview
- System Requirements

ESS Overview

Important: The ESS is not a part of the ASR5x00 platform or the Enhanced Charging Service (ECS) in-line service. It is an external server.

Important: For information on compatibility between ESS and StarOS releases, contact your Cisco account representative.

On the ASR 5x00 chassis, the CDR subsystem provides 512 MB of volatile memory on the packet processing card RAM on the ASR 5000 and the data processing card RAM on the ASR 5500 to store accounting information. This on-board memory is intended as a short-term buffer for accounting information so that billing systems can periodically retrieve the buffered information for bill generation purposes. However if network outages or other failures cause billing systems to lose contact with the system, it is possible that the CDR subsystem storage area can be filled with non-retrieved accounting information. When the storage is filled the CDR subsystem starts deleting the oldest files to make sure that there is room for new billing files and non-retrieved accounting information can be lost. Using an external storage server with a large storage volume in close proximity to the chassis ensures room for storing a large amount of billing data that is not lost by any failure.

The ESS has the capability of simultaneously fetching any types of files from one or more chassis. That is, it can fetch xDRs like CDR, EDR, NBR, UDR file, etc.

In case of Hard Disk Drive (HDD) support on the chassis, the platform has the capability to push the xDR files to ESS, and ESS forwards these files to the required destinations. If HDD is not configured on the platform, ESS pulls the files from the system and forwards them to the destinations.

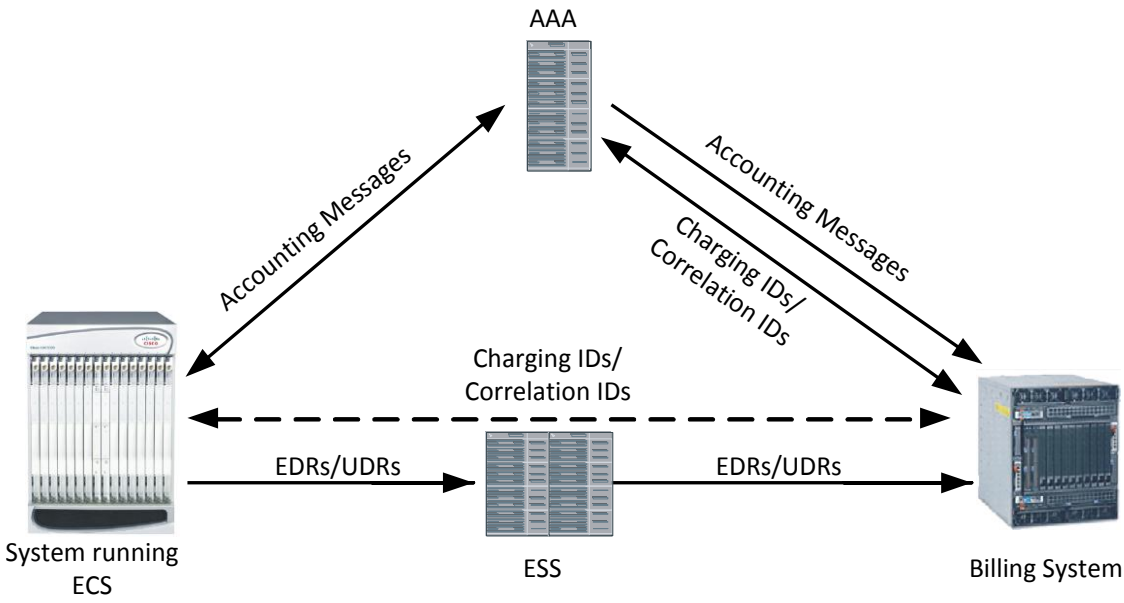
The ESS is designed to be used as a safe storage area. A mediation or billing server within your network must be configured to collect accounting records from the ESS once it retrieves them.

The ESS supports a high level of redundancy for secure charging and billing information for post-processing of xDRs. This system can store charging data of up to 30 days.

Important: The procedures in this guide assume that you have already configured your chassis with ECS as described in the *Enhanced Charging Services Administration Guide*.

The following figure shows a typical organization of ESS and billing system with chassis having a AAA server.

Figure 89. ESS Architecture with ECS



The system running with ECS stores xDR files on an ESS and billing system collects the files from the ESS, and correlates them with the AAA accounting messages using either 3GPP2-Correlation-IDs on a PDSN system or Charging IDs on a GGSN system.

ESS also pushes xDR files to external applications for post-processing, reporting, subscriber profiling, and trend analysis.

ESS Features and Functions

The ESS is a storage server logically connected with the ASR 5x00 and acts as an integrated network system.

The following are some of the important features of an ESS:

- High speed dedicated redundant connections to chassis to pull xDR files.
- High-speed dedicated and redundant connection with billing system to transfer xDR files.
- Different management addresses than the management addresses of the chassis and billing system.
- Management interface with support of multiple VLANs.
- Redundancy support with two or more geographically co-located or isolated chassis to pull xDRs.

In general ESS provides the following functions:

- Stores copy of records pulled from chassis.
- Supports storage of up to 7 days worth of records.
- Supports storage capacity of carrier-class redundant.
- Provides a means of limiting the amount of bandwidth, in term of kbps, used for the file transfer between chassis and ESS.
- Provides a means of archiving/compression of the pulled xDR files for the purpose of extending the storage capacity.

- Provides xDR files to the billing system.

System Requirements

The requirements described in this section must be met in order to ensure proper operation of the ESS system.

ASR 5x00 System Requirements

The following configurations must be implemented, as described in *Configuring Enhanced Charging Services* chapter of the *Enhanced Charging Services Administration Guide*:

- ECS must be configured for generating billing records.
- An administrator or config-administrator account that is enabled for FTP must be configured.
- SSH keys must be generated.
- The SFTP subsystem must be enabled.

ESS System Requirements

Important: System requirement recommendation is dependent of different parameters including xDR generation, compression, deployment scenario, etc. Contact your sales representative for system requirements specific to your ESS deployment.

ESS System Recommendations for Cluster Deployment

This section identifies the minimum system requirements recommended for the the cluster deployment of the ESS application in 9.0 and earlier releases:

- Sun Microsystems Netra™ T5220 server
  - 1 x 1.2GHz 4 core UltraSPARC T2 processor with 8GB RAM
  - 2 x 146GB SAS hard drives
  - Quad Gigabit Ethernet interfaces

    Important: It is recommended that you have separate interfaces (in IPMP) for mediation device and chassis. Also, for given IPMP, the two interfaces should be on different cards.

  - Internal CDROM drive
  - AC or DC power supplies depending on your application
- Fiber channel (FC) based Common Storage System for Servers (Sun Storage Tek 2540)
- PCI Dual FC 4GB HBA
- Dual RAID Controllers
- 5 x 300GB 15K drives
- AC or DC power supplies depending upon your application

Chapter 14
Femto Network Gateway Overview

This chapter contains general overview information about the Femto Network Gateway (FNG), including:

- Product Description
- Summary of FNG Features and Functions
- Network Deployment(s) and Interfaces
- Features and Functionality
- How the FNG Works
- Supported Standards

Product Description

The Cisco® Femto Network Gateway (FNG) enables mobile operators with CDMA2000 networks to provide 3G services to subscribers with wireless handsets via Femtocell Access Points (FAPs). The FNG makes it possible for operators to provide secure access to the operator's 3G network from a non-secure network, extend wireless service coverage indoors, especially where access would otherwise be limited or unavailable, reduce the load on the macro wireless network, and make use of existing backhaul infrastructure to reduce the cost of carrying wireless calls.

The FNG functions as a security gateway that allows the FAPs in the access network to connect to circuit, packet, and IMS core networks. The FNG implements an IPSec interface to provide a secure, encrypted IPSec tunnel to each FAP in the access network as it connects to the operator's core network. In addition, the FNG provides a highly scalable femtocell solution by allowing a large number of FAPs to interoperate with legacy core network elements that are typically not designed to interface with such a large number of elements.

The FNG splits voice and data traffic flows into and out of the core network. It forwards all voice traffic to the operator's IMS core network and all data traffic to the PDSN/HA toward the packet data network. This network configuration fully isolates the traditional MSC from IP attacks, because all backhauled traffic is secure and offloaded to a convergence server in the IMS core network.

Platform Requirements

The FNG service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the installation guide for the chassis and/or contact your Cisco account representative.

Licenses

The FNG is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the "Managing License Keys" section of the "Software Management Operations" chapter in the *System Administration Guide*.

Summary of FNG Features and Functions

The FNG features and functions include:

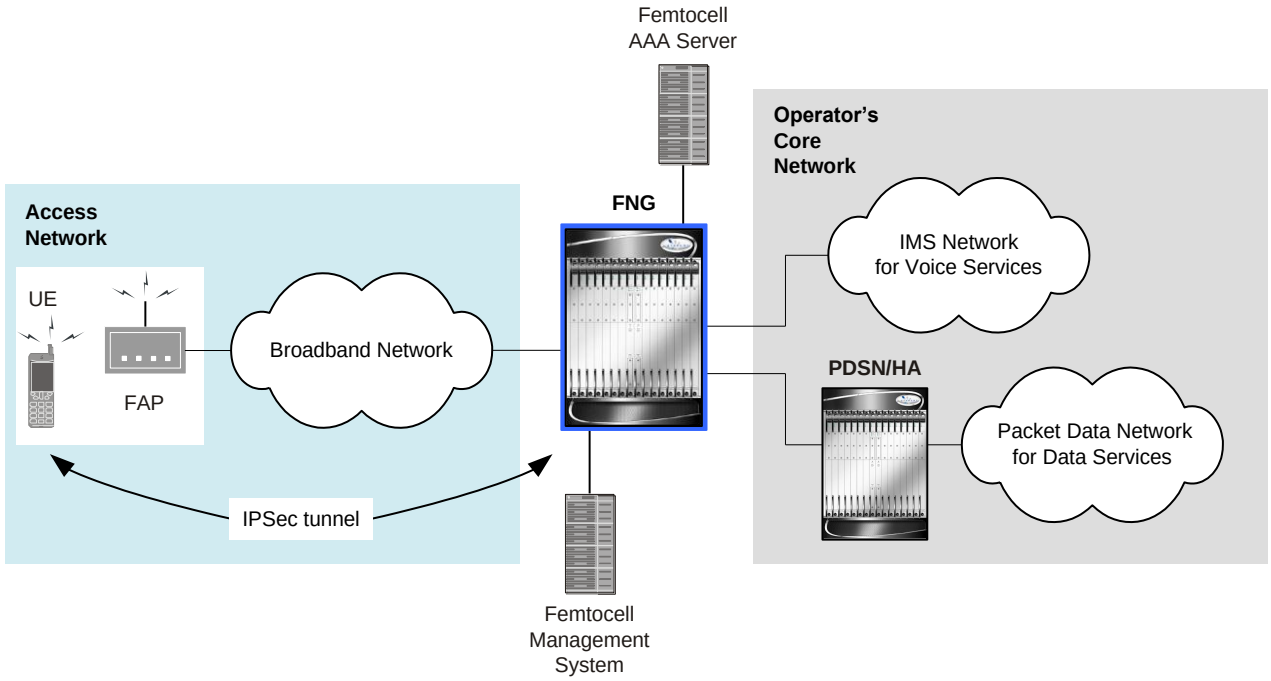
- FNG service
- IKEv2 and IP Security (IPSec) encryption
- A12 aggregation
- X.509 certificate-based peer authentication
- RADIUS Support
- AAA server group selection
- FAP ID-based duplicate session detection
- Child SA rekey support
- Multiple Child SAs
- DoS protection cookie challenge
- IKEv2 keep-alive messages (dead peer detection)
- DSCP marking
- Custom DNS handling
- Session recovery support
- Congestion control
- Bulk statistics
- Threshold crossing alerts (TCAs)

Network Deployment(s) and Interfaces

This section describes the FNG as it functions in a CDMA2000 network.

The figure below shows how the FNG functions both as a security gateway and a femtocell gateway between the FAPs in the access network and the operator's IMS core network for voice services and the PDSN/HA and the packet data network for data services.

Figure 90. FNG Network Architecture



Network Elements

This section provides a description of the network elements in an FNG network.

Femtocell Access Point

The Femtocell Access Point (FAP) is a SIP-based CDMA2000 wireless access point that provides coverage in a small area, usually a private residence or small office, and connects the subscriber UEs to an operator's core network via a broadband connection (e.g., DSL or cable). A FAP allows operators to extend wireless service coverage indoors, especially where access would otherwise be limited or unavailable.

Femtocell Management System

The Femtocell Management System (FMS) is a network element that resides in the operator's network and facilitates the provisioning, activation, and operational management of the FAPs in the network based on industry standards such as TR-069. The FMS helps to ensure the scalability of the FAP network to potentially millions of devices.

Femto Network Gateway

The Femto Network Gateway (FNG) is a network element that resides in the operator's network and functions as both a security gateway and a femtocell gateway. The security gateway functions provide secure access for the FAPs to access services within the operator's core network. The femtocell gateway functions provide aggregation and proxy capabilities for the FAPs. The FNG forwards all voice traffic to the operator's IMS core network and all data traffic to the PDSN/HA.

Femtocell AAA Server

The Femtocell AAA Server provides a FAP authorization function. It sends authorization policy information to the FNG.

IMS Core Network Elements

An operator's IMS core network may include the following elements to enable voice services:

- **P-CSCF:** The P-CSCF (Proxy Call/Session Control Function) is the entry point into the IMS domain and serves as the outbound proxy server for SIP messaging for the subscriber UEs. The UEs attach to the P-CSCF prior to performing IMS registrations and initiating SIP sessions. All SIP signaling traffic to and from the FAPs and the IMS core is handled by the P-CSCF. The P-CSCF provides message manipulation, breakout of emergency call services, QoS (Quality of Service) authorization, and signaling compression. Once the P-CSCF completes all of the functions for which it is responsible, it forwards the call to the I-CSCF.
- **I-CSCF:** The I-CSCF (Interrogating Call/Session Control Function) functions as a location server in the IMS core network. Its major functions are to select the appropriate registrar server for the subscriber UEs by consulting the HSS (Home Subscriber Server) and forwarding the request to the IMS registrar (the S-CSCF). The HSS returns a set of required S-CSCF capabilities for initial registration requests by the UE. Based on these capabilities, the I-CSCF selects the appropriate S-CSCF.
- **S-CSCF:** The S-CSCF (Serving Call/Session Control Function) provides session control and registration services for the subscriber UEs and FAPs in the network. It is responsible for all aspects of session control, handling all subscriber requests, which it relays to the appropriate application server. The S-CSCF routes mobile-terminating traffic to the P-CSCF and routes mobile-originating traffic to the convergence server based on iFC (initial Filter Criteria) downloaded from the HSS.
- **HSS:** The HSS (Home Subscriber Server), is the master user database that supports the IMS network entities that handle calls. It contains subscription-related information (subscriber profiles), performs authentication and authorization of the user, and provides information about the subscriber's location and IP information.
- **Femtocell Convergence Server:** The Femtocell Convergence Server (FCS) is an IMS application server that provides legacy Telephony Application Services (TAS) to 1x femtocell subscribers via SIP, including voice services and voice feature delivery. The femtocell convergence server also manages idle and active mode mobility for 1x subscribers as they move into and out of range of FAP coverage. It functions as an MFIF (MAP-Femtocell Interworking Function) and interfaces with the HLR for 1x subscriber authentication. It appears as an IMS application server to the S-CSCF and as a serving MSC to the HLR.

- **Media Gateway:** The Media Gateway terminates bearer channels from the circuit-switched network and media streams from the packet-switched network. It can support media conversion, bearer control, and payload processing (e.g., using codecs, echo cancellers, and conference bridges).

PDSN/HA

The PDSN/HA enables femtocell subscribers to receive packet data services in the mobile operator's core network. In most cases, these services are the same as those available via the mobile operator's macro network.

Basic Operation

When a FAP powers up, it uses DNS resolution to resolve its pre-configured FQDN of the FNG and obtain the FNG's IP address. It then initiates IPSec tunnel establishment over the broadband access network. The IPSec tunnel terminates at the FNG.

The FAP receives an IPv4 address known as the Tunnel Inner Address (TIA) from the FNG during the first IPSec tunnel establishment. The FNG assigns the TIA from its own IPv4 address pool. Once an IPSec tunnel is established, the FAP uses the TIA in all its SIP messages and to obtain configuration data from its FMS. The FNG is agnostic in regard to the protocol used between the FAP and its FMS, and simply forwards packets between the FAP and the FMS over the secure connection.

The 1x FAP performs P-CSCF discovery via a DHCP server per RFC 3319, or it may receive the IP address of the P-CSCF from the FMS. Once the FAP gets the P-CSCF address, it initiates SIP registration with the IMS core network. When a UE attaches to the FAP, it performs 1x registration with the IMS core network.

Network Interfaces

The following table provides descriptions of the network interfaces supported by the FNG in a CDMA2000 network.

Table 58. Network Interfaces in a CDMA2000 Network

Interface	Description
FAP Interface	The secure interface to the FAPs in the network is an IPSec tunnel. The FNG uses IKEv2 for establishing the IPSec tunnel. Note that the FNG does not have a direct interface to the UEs in the network. The FNG receives all voice and data traffic from the UEs via secure IPSec tunnels between the FNG and the FAPs and sends the traffic to the operator's IMS core or PDSN/HA.
RADIUS Interface	The interface to the RADIUS server is used for FAP device authentication. The FAP can use one of the following authentication methods: • EAP-AKA (Extensible Authentication Protocol - Authentication and Key Agreement) authentication • PSK (Pre-Shared Key) authentication • X.509 certificate-based peer (client) authentication
Interface with the IMS Core	The FNG sends all SIP signaling and bearer traffic from the FAPs to the IMS core to access voice services.

Interface	Description
Interface with the PDSN/HA	The FNG sends all signaling and bearer traffic from the FAPs to the PDSN/HA to access packet data services.

Features and Functionality

This section describes the features and functions supported by the FNG.

The following features are supported and described in this section:

- FNG Service
- IKEv2 and IP Security (IPSec) Encryption
- X.509 Certificate-based Peer Authentication
- A12 Aggregation
- RADIUS Support
- AAA Server Group Selection
- FAP ID-based Duplicate Session Detection
- Child SA Rekey Support
- Multiple Child SAs
- DoS Protection Cookie Challenge
- IKEv2 Keep-Alive Messages (Dead Peer Detection)
- DSCP Marking
- Custom DNS Handling
- Session Recovery Support
- Congestion Control
- Bulk Statistics
- Threshold Crossing Alerts

FNG Service

The FNG service and its associated processes enable the system to function as a femtocell gateway. The FNG service enables the FAPs in the network to connect to the core network elements via a secure IPSec interface. During configuration, you create the FNG service in an FNG context, which is a routing domain on the ASR 5x00. FNG context and service configuration includes the following main steps:

- **Configure the IPv4 address for the service:** This is the IP address of the FNG to which the FAPs in the network attempt to connect, sending IKEv2 messages to this IP address to establish IPSec tunnels.
- **Configure the name of the crypto template for IKEv2/IPSec:** A crypto template is used to configure an IKEv2/IPSec policy. It includes most of the IKEv2 and IPSec parameters for keep-alive, lifetime, NAT-T, and cryptographic and authentication algorithms. There must be one crypto template per FNG service.
- **The name of the EAP profile:** This profile defines the EAP authentication method and associated parameters. If the PSK (Pre-Shared Key) authentication method is used, this configuration is not needed.
- **IKEv2 and IPSec transform sets:** Transform sets define the negotiable algorithms for IKE SAs and Child SAs to enable calls to connect to the FNG.

- **The setup timeout value:** This parameter specifies the session setup timeout timer value. The FNG terminates a connection attempt if the FAP does not establish a successful connection within the specified timeout period.
- **Max-sessions:** This parameter sets the maximum number of subscriber sessions allowed by this FNG service.
- **FNG supports a domain template for storing domain-related configuration:** The domain name is taken from the received Network Address Identifier (NAI) and searched in the domain template database.
- **Duplicate session detection parameters:** The FNG supports the FAP ID in the form of an NAI for duplicate session detection. This setting enables duplicate session detection for the FNG service.

When the FNG service is configured in the system with the IP address, crypto template, and so on, the FNG is ready to accept IKEv2 control packets for establishing IKEv2 sessions.

IKEv2 and IP Security (IPSec) Encryption

The FNG supports IKEv2 and IPSec encryption using IPv4 addressing. IKEv2 and IPSec encryption enables network domain security for all IP packet-switched networks in order to provide confidentiality, integrity, authentication, and anti-replay protection.

At the beginning of IKEv2 session setup, the FNG and the FAP exchange capabilities for authentication. IKEv2 and IPSec transform sets configured in the crypto template define the negotiable algorithms for IKE SA and Child SA setup to connect calls to the FNG by creating a single IPSec tunnel, called the Tunnel Inner Address (TIA), which is intended for user traffic coming from the FAP. There can be multiple UEs connecting to a single FAP at the same time, and the traffic from all of the connected UEs passes through the same IPSec tunnel. The FAP to which a UE is connected can request one of the following authentication methods:

- EAP-AKA (Extensible Authentication Protocol - Authentication and Key Agreement) authentication
- PSK (Pre-Shared Key) authentication
- X.509 certificate-based peer (client) authentication

The FNG partially supports the EAP MD5 (Extensible Authentication Protocol Message-Digest 5) authentication method.

X.509 Certificate-based Peer Authentication

In addition to the EAP-AKA (Extensible Authentication Protocol - Authentication and Key Agreement) and PSK (Pre-Shared Key) peer authentication methods, the FNG supports X.509 certificate-based peer authentication.

The FNG checks the network policy on whether a FAP is authorized to provide service. If the network policy states that all FAPs that pass device authentication are authorized to provide service, no further authorization check may be required. If the network policy requires that each FAP be individually authorized for service (in the case where the FEID is associated with a valid subscription), the FNG sends a RADIUS Access-Request message to the AAA server. If the AAA server sends a RADIUS Access-Accept message, the FNG proceeds with device authentication. Otherwise, the FNG terminates the IPSec tunnel setup by sending an IKEv2 Notification message indicating authentication failure.

For a detailed presentation of X.509 certificate-based peer authentication, see the section *How the FNG Works* later in this chapter.

A12 Aggregation

The Access Network AAA (AN-AAA) servers in 1x networks are not designed to handle a large numbers of FAPs attempting A12 authentication to access the network. The A12 aggregation feature reduces the number of source addresses in the A12 Access-Request messages sent to the AN-AAA servers by the FNG, which simplifies the configuration of the AN-AAA server's database.

A12 authentication is a CHAP-based authentication method used by CDMA2000 AN-AAA servers to provide High Rate Packet Data (HRPD) access authentication between the AN function in the FAPs and the AN-AAA servers in the network.

When the FNG receives an A12 Access-Request message from a FAP, it validates the source address of the FAP, then substitutes the source address (and, optionally, the NAS IP address/port number) in the Access-Request message with its own source address before sending the message to the AN-AAA server. When the FNG receives the Access-Accept message from the AN-AAA server, the FNG sends it back to the FAP. In this way, the number of AAA sessions required by the AN-AAA server is reduced.

RADIUS Support

RADIUS support on the FNG provides a mechanism for performing authentication, authorization, and accounting (AAA) for subscribers. The benefits of using AAA are:

- Higher flexibility for subscriber access control
- Better accounting, charging, and reporting options
- Industry standard RADIUS authentication

The Remote Authentication Dial-In User Service (RADIUS) protocol can be used to provide AAA functionality for subscribers. The AAA functionality on the FNG provides a wide range of configuration options via AAA server groups, which allow a number of RADIUS parameters to be configured in support of the FNG service.

Currently, two types of authentication load-balancing methods are supported: first-server and round-robin. The first-server method sends requests to the highest priority active server. A request will be sent to a different server only if the highest priority server is not reachable. With the round-robin method, requests are sent to all active servers in a round-robin fashion.

The FNG can detect the status of the AAA servers. Status checking is enabled by configuration in the AAA Server Group Configuration Mode of the system's CLI. Once an AAA server is detected to be down, it is kept in the down state up to a configurable duration of time called the dead-time period. After the dead-time period expires, the AAA server is eligible to be retried. If a subsequent request is directed to that server and the server properly responds to the request, the system makes the server active again.

Important: In 12.3 and earlier releases, refer to the *AAA and GTPP Interface Administration and Reference* for more information on RADIUS AAA configuration.

AAA Server Group Selection

This feature provides a maximum of 64 AAA groups on the ASR 5x00. This could be spread across multiple contexts or all groups can be configured within a single context. A maximum of 320 RADIUS servers is allowed on the chassis, unless the **`aaa-large-configuration`** command is issued, and this number becomes a maximum of 800 AAA groups and 1600 RADIUS servers allowed to be configured per chassis.

Important: In 12.3 and earlier releases, refer to the *AAA and GTPP Interface Administration and Reference* for more information on AAA server groups.

FAP ID-based Duplicate Session Detection

When this feature is enabled and a FAP sets up a new session, the FNG automatically checks for any remnants of abandoned calls, and if found, clears them. Clearing the old session and establishing the new session in parallel optimizes FNG processing functions.

With every new session setup, the FNG verifies whether there are any old sessions that are bound to the Femtocell Access Point Identifiers (FAP IDs). For example, when a FAP reboots, it may initiate a new session with the FNG. After authentication, if the FNG detects an old session with the same FAP ID, the FNG clears the old IPSec tunnel and establishes a new IPSec tunnel with the FAP. This feature is designed with the assumption that not more than one call with duplicate FAP IDs is in the setup stage at any one time.

You enable FAP ID-based duplicate session detection in the FNG Service Configuration Mode of the system's CLI. This feature should be enabled in the boot-time configuration before any calls are established.

Tunnel Cleanup on FAP Reboot

The FNG supports initial contact handling in IKE_AUTH messages as per RFC 4306 and cleans up the original tunnel if a FAP initiates a new tunnel after a reboot. The CLI command for duplicate session detection is not needed to enable this detection. Initial contact notification asserts that this IKE_SA is the only IKE_SA currently active between the authenticated identities. It may be sent when an IKE_SA is established after a crash, and the recipient may use this information to delete any other IKE_SAs it has for the same authenticated identity without waiting for a timeout.

Child SA Rekey Support

Rekeying of an IKEv2 Child Security Association (SA) occurs for an already established Child SA whose lifetime (either time-based or data-based) is about to exceed a maximum limit. The FNG initiates rekeying to replace the existing Child SA. During rekeying, two Child SAs exist momentarily (500ms or less) to ensure that transient packets from the original Child SA are processed by the FNG and not dropped.

FNG-initiated Child SA rekeying is disabled by default, and rekey requests are ignored. You can enable this feature in the Crypto Configuration Payload Mode of the system's CLI.

Multiple Child SAs

The FNG supports the instantiation, termination, and rekeying of multiple simultaneous Child SAs derived from an IKE SA, as defined in RFC 4306.

As specified in the IKEv2 policy, which controls the behavior of encrypted tunnels, the first Child SA is instantiated during the IKE_AUTH exchange between the FAP and the FNG, and any additional Child SAs are instantiated during subsequent CREATE_CHILD_SA exchanges that may occur between the FAP and the FNG.

An IKEv2 policy may be terminated via operator intervention or be terminated when a service is terminated. In these scenarios, all objects derived from the IKEv2 policy, including the IKE SA and all Child SAs, are terminated.

The FNG maintains two maximum Child SA values per IKEv2 policy. The first is a system-enforced maximum value, which is four Child SAs per IKEv2 policy. The second is a configurable maximum value, which can be a value between one and four, and which is specified via the system's CLI in the Crypto Template Configuration Mode.

If the system maximum value or the configured maximum value is reached and the FNG receives a CREATE_CHILD_SA Request for an additional Child SA, the FNG returns a CREATE_CHILD_SA Response that contains a Notify payload of the type NO_ADDITIONAL_SAS. Note that the maximum value does not apply to interim Child SAs that may exist during transitional phases such as during Child SA rekeying. For example, if a maximum of two simultaneous Child SAs are specified, the FNG allows a burst of four during Child SA rekeying.

DoS Protection Cookie Challenge

There are several known types of Denial of Service (DoS) attacks associated with IKEv2. Through a configurable option in the Crypto Template Configuration Mode in the system's CLI, the FNG can implement the IKEv2 cookie challenge payload method per RFC 4306. This method is intended to protect against the FNG creating too many half-opened sessions or other similar mechanisms.

This feature is disabled by default. When enabled, and when the number of half-opened IPSec sessions exceeds the configured limit of any integer between 0 and 100,000 (or the trigger point with other detection mechanisms), the FNG invokes the cookie challenge payload mechanism to insure that only legitimate subscribers are initiating IKEv2 tunnel requests, as follows:

1. The FAP connects to the FNG and sends an IKE_SA_INIT Request message.
2. The FNG sends a Notify (cookie) payload to the FAP to request retransmission of the IKE_SA_INIT Request message with the received Notify (cookie) payload in the message.
3. Upon receipt of the retransmitted message, the FNG verifies the cookie payload and ensures that it is the same cookie payload as the one it had sent.
4. If the cookie challenge is met, setup continues as normal with the FNG sending an IKE_SA_INIT Response message.

IKEv2 Keep-Alive Messages (Dead Peer Detection)

The FNG supports IKEv2 keep-alive messages, also known as Dead Peer Detection (DPD), originating from both the FAPs and the FNG. You configure DPD per FNG service. You can also disable DPD, and the FNG will not initiate DPD exchanges with the FAPs. However, the FNG always responds to DPD availability checks initiated by a FAP regardless of the FNG configuration.

DSCP Marking

If different classes of traffic are sent on the same SA and if the FAPs in the network and the FNG are employing the optional anti-replay feature in the Encapsulating Security Payload (ESP), this could result in inappropriate discarding of lower priority packets due to the windowing mechanism used by this feature. Therefore, it is recommended that multiple Child SAs are used to provide the appropriate QoS services. This handling can be applied to different types of traffic (voice and data) coming from the same UE behind a FAP, or from multiple UEs belonging to the same QoS class. The FNG will determine the traffic type and provide a QoS treatment based on configured rules.

Custom DNS Handling

The custom DNS feature provides a mechanism whereby the FNG sends the DNS address specified in the FNG configuration file to the FAP only if the FAP requests it. The FNG considers an address of 0.0.0.0 invalid and does not include it.

Session Recovery Support

The session recovery feature provides seamless failover and nearly instantaneous reconstruction of subscriber session information in the event of a hardware or software fault within the same chassis, preventing a fully-connected user session from being dropped.

Important: Use of the session recovery feature requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

Session recovery is performed by mirroring key software processes (the IPSec manager, session manager, and AAA manager, for example) on the FNG. These mirrored processes remain in an idle state (in standby mode), where they perform no processing until they may be needed in the case of a software failure (a session manager task aborts, for example). The system spawns new instances of standby mode sessions and AAA managers for each active control processor being used.

Additionally, other key system-level software tasks such as VPN manager are performed on a physically separate PSC/PSC2 to ensure that a double software fault (the session manager and the VPN manager fail at same time on same card, for example) cannot occur. The PSC/PSC2 used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

Important: For more information about session recovery support, refer to the *System Administration Guide*.

Congestion Control

The congestion control feature allows you to set policies and thresholds and specify how the system reacts when faced with a heavy load condition.

The congestion control feature monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are resolved quickly. However, continuous or large numbers of these conditions within a specific time interval may have an impact on the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establishes limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operation thresholds that are configured for the system as described in the *Thresholding Configuration Guide*. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated. A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.

- **Port Utilization Thresholds:** If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
- **Port-specific Thresholds:** If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
- **Service Congestion Policies:** Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

Important: For more information on congestion control, refer to the *System Administration Guide*.

Bulk Statistics

Bulk statistics allow operators to choose to view not only statistics that are of importance to them, but to also configure the format in which they are presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics and send them to a collection server called a receiver. Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. The following is a partial list of supported schemas:

- **System:** Provides system-level statistics.
- **Card:** Provides card-level statistics.
- **Port:** Provides port-level statistics.
- **FNG:** Provides FNG service statistics.

The system supports the configuration of up to four sets (primary/secondary) of receivers. Each set can be configured to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, system host name, system uptime, the IP address of the system generating the statistics (available for headers and footers only), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics Server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Important: For more information on bulk statistic configuration, refer to the "Configuring and Maintaining Bulk Statistics" chapter of the *System Administration Guide*.

Threshold Crossing Alerts

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outages. Typically, these conditions are temporary (i.e., high CPU utilization or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to avoid and/or minimize system downtime.

The system supports threshold crossing alerts for certain key resources such as CPU, memory, IP pool addresses, and so on. With this capability, the operator can configure a threshold on these resources whereby, should the resource depletion cross the configured threshold, an SNMP trap would be sent.

The following thresholding models are supported by the system:

- **Alert:** A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated, then generated and/or sent again at the end of the polling interval.
- **Alarm:** Both high and low thresholds are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated, then generated and/or sent again at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP Traps:** SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values. Generation of specific traps can be enabled or disabled on the chassis, ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.
- **Logs:** The system provides a thresholding facility for which active and event logs can be generated. As with other system facilities, logs are generated messages pertaining to the condition of a monitored value are generated with a severity level of WARNING. Logs are supported in both Alert and Alarm modes.
- **Alarm System:** High threshold alarms generated within the specified polling interval are considered outstanding until a condition no longer exists or a condition clear alarm is generated. Outstanding alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

How the FNG Works

This section describes the FNG functioning as a security gateway during IPSec tunnel establishment.

IPSec Tunnel Establishment

The figure below shows the message flow during IPSec tunnel establishment. The table that follows the figure describes each step in the message flow.

Figure 91. IPSec Tunnel Establishment

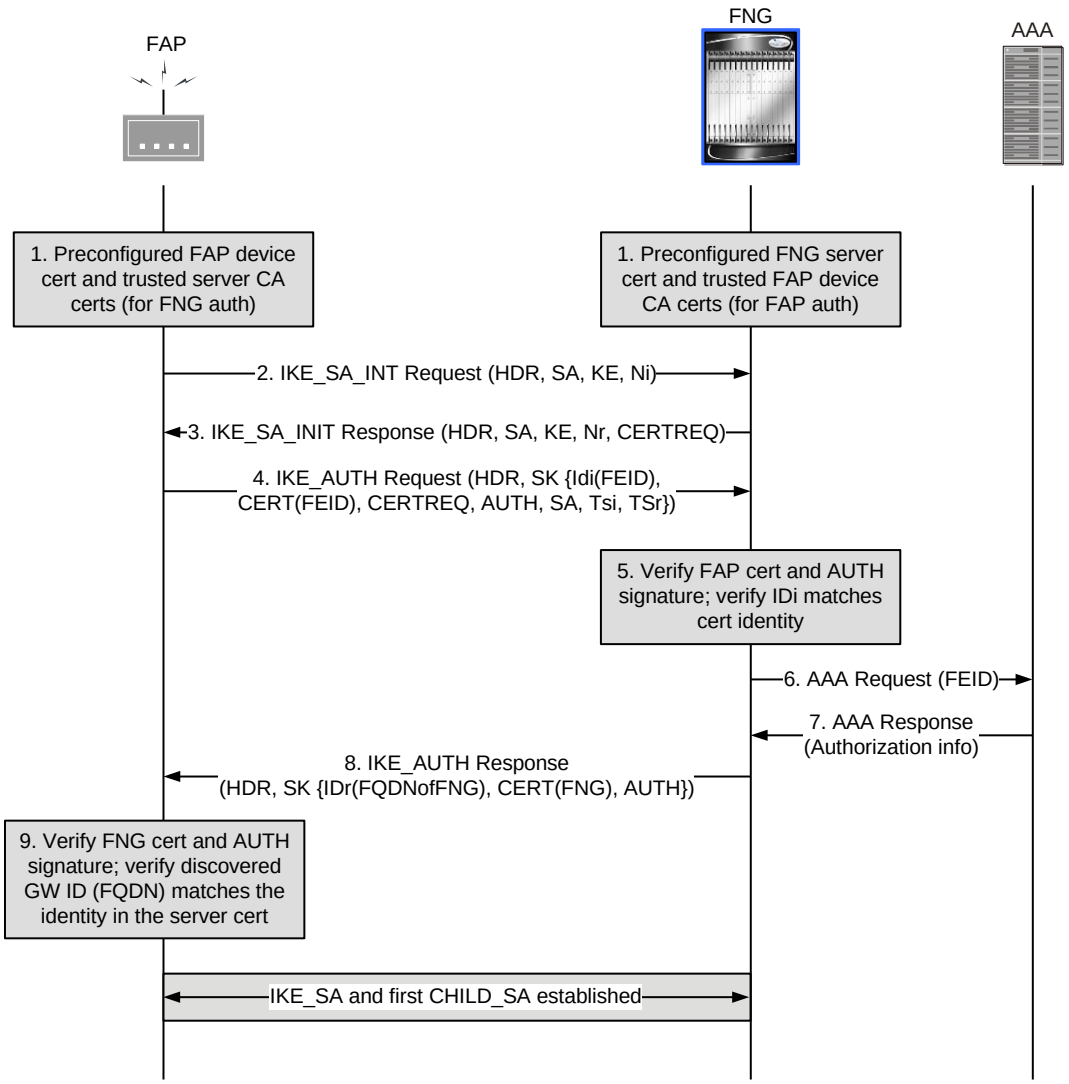


Table 59. IPSec Tunnel Establishment

Step	Description
1.	The FAP is assigned a device certificate during it's manufacturing. The private key for the certificate is stored securely at the FAP. Similarly, the FNG is assigned a server certificate. The FNG is also configured with a list of root CA certificates corresponding to the trusted device CA certificates.
2.	The FAP initiates an IKEv2 exchange with the FNG, known as the IKE_SA_INIT exchange, by issuing an IKE_SA_INIT Request to negotiate cryptographic algorithms, exchange nonces, and perform a Diffie-Hellman exchange with the FNG.
3.	The FNG responds with an IKE_SA_INIT Response by choosing a cryptographic suite from the initiator's offered choices, completing the Diffie-Hellman and nonce exchanges with the FAP. In addition, the FNG includes the list of FAP CA certificates that it will accept in its CERTREQ payload. For successful FAP authentication, the CERTREQ payload has to contain at least one CA certificate that is in the trust chain of the FAP device certificate. At this point in the negotiation, the IKE_SA_INIT exchange is complete and all but the headers of all the messages that follow are encrypted and integrity-protected.
4.	The FAP initiates an IKE_AUTH exchange with the FNG by setting the IDi payload to the FEID, the CERT payload set to the FAP device certificate corresponding to the FEID, and the AUTH payload containing the signature of the previous IKE_SA_INIT Request message (in step 2) generated using the private key of the FAP device certificate. The authentication algorithm used to generate the AUTH payload is also included in the AUTH payload.
5.	Using the CA certificate corresponding to the FAP device certificate, the FNG first verifies that the FAP device certificate in the CERT payload has not been modified and the identity included in the IDi corresponds to the identity in the FAP device certificate. If the verification is successful, using the public key of the FAP device certificate, the FNG generates the expected AUTH payload and compares it with the received AUTH payload. If they match, the authentication of the FAP is successful. Otherwise, the FNG sends an IKEv2 Notification message indicating authentication failure.
6.	If the network policy requires femtocell subscription authorization, the FNG contacts the AAA server to verify that the FAP identified by the FEID is authorized to provide service.
7.	The AAA server responds with the authorization result. If the authorization is not successful, the FNG sends an IKEv2 Notification message indicating authorization failure. Otherwise, the FNG proceeds with server authentication.
8.	The FNG responds with the IKE_AUTH Response by setting the IDr payload to the FQDN of the FNG, setting the CERT payload to the FNG server certificate corresponding to the FQDN, and including the AUTH payload containing the signature of the IKE_SA_INIT Response message (in step 3) generated using the private key of the FNG server certificate. The authentication algorithm used to generate the AUTH payload is also included in the AUTH payload.
9.	Using the CA certificate corresponding to the FNG server certificate, the FAP first verifies that the FNG server certificate in the CERT payload has not been modified and the identity included in the IDi corresponds to the identity in the server certificate and contains the expected FNG value as discovered during the FNG discovery procedures. If the verification is successful, using the public key of the FNG server certificate, the FAP generates the expected AUTH payload and compares it with the received AUTH payload. If they match, the FNG server authentication is successful. This completes the IKE_AUTH exchange. An IPSec SA with the first CHILD_SA pair is established between the FAP and the FNG.

IPSec Tunnel Establishment with EAP-AKA Authentication

The figure below shows the message flow during IPSec tunnel establishment with EAP-AKA authentication. The table that follows the figure describes each step in the message flow.

Figure 92.IPSec Tunnel Establishment with EAP-AKA Authentication

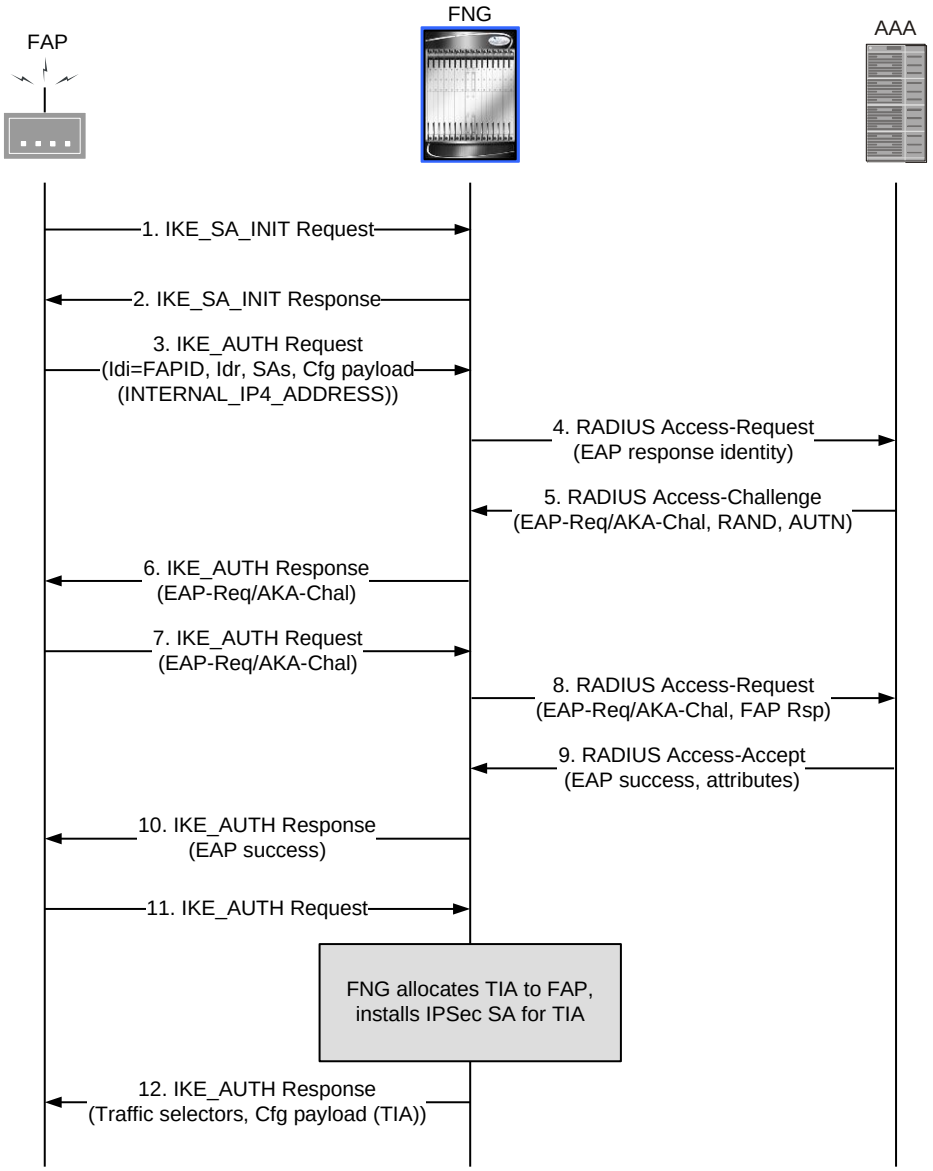


Table 60. IPSec Tunnel Establishment with EAP-AKA Authentication

Step	Description
1.	The FAP initiates an IKEv2 exchange with the FNG, known as the IKE_SA_INIT exchange, by issuing an IKE_SA_INIT Request to negotiate cryptographic algorithms, exchange nonces, establish NAT traversal, and perform a Diffie-Hellman exchange with the FNG.
2.	The FNG responds with an IKE_SA_INIT Response by choosing a cryptographic suite from the initiator's offered choices, completing the Diffie-Hellman and nonce exchanges with the FAP.

Step	Description
3.	The FAP initiates an IKE_AUTH exchange with the FNG. The FAP omits the AUTH payload, indicating that it wants to use an EAP exchange over IKEv2. The FAP includes its identity in the IDi payload of the IKE_AUTH Request. The IDi is set to the FAP ID. The FAP ID is a string in the format id@domain. The FAP also includes the IKEv2 CFG_REQUEST payload in the IKE_AUTH Request. The INTERNAL_IP4_ADDRESS attribute is included in the CFG_REQUEST payload with the length set to 0.
4.	The FNG receives the IKE_AUTH Request and sends the FAPID as the EAP Response identity to the AAA server using a RADIUS Access-Request message with an EAP-Message attribute.
5.	The AAA server verifies the FAP's identity and generates a random value RAND and AUTN based on the shared CHAP-key and a sequence number. The AAA server sends the EAP-Request/AKA-Challenge to the FNG via a RADIUS Access-Challenge message. The EAP-Request/AKA-Challenge contains the RAND and AUTN to protect the integrity of the EAP message.
6.	The FNG sends an IKE_AUTH Response to the FAP that contains the EAP-Request/AKA-Challenge message received from the AAA server.
7.	The FAP verifies the authentication parameters in the EAP-Request/AKA-Challenge message and if the verification is successful, it responds to the challenge with an IKE_AUTH Request message to the FNG.
8.	The FNG forwards the EAP-Response/AKA-Challenge message to the AAA server via a RADIUS Access-Request message.
9.	If the authentication is successful, the AAA server sends a RADIUS Access-Accept message with an EAP-Message attribute containing EAP Success. The AAA server sends the EAP Success and the MSK generated during the EAP-AKA authentication process to the FNG. In addition, the AAA server also sends other attributes that it normally sends to the PDSN for a simple IP session. These attributes include at a minimum the Framed-Pool (if required), so that the FNG can assign a TIA from the correct IP address pool, the Session-Timeout, and the Idle-Timeout.
10.	The FNG forwards the EAP Success message to the FAP in an IKE_AUTH Response message.
11.	The FAP calculates the MSK according to RFC 4187 and uses it as an input to generate the AUTH payload to authenticate the first IKE_SA_INIT message. The FAP sends the AUTH payload to the FNG in an IKE_AUTH Request message.
12.	The FNG uses the MSK to check the validity of the AUTH payload received from the FAP and calculates its own AUTH payload for the FAP to verify per RFC 4306. The FNG sends the AUTH payload to the FAP together with the configuration payload containing SAs and the rest of the IKEv2 parameters in an IKE_AUTH Response message. This completes the IKEv2 negotiation. The configuration payload contains the TIA. It is up to the FAP implementation to establish separate Child SAs for configuration management and VoIP traffic, or to use the same Child SA for all traffic types. The FNG supports both options. Once the IPSec tunnel is established, the FAP uses the TIA assigned by the FNG for each 1x UE (in the SIP headers or as the RTP IP address).

X.509 Certificate-based Peer Authentication

The figure below shows the message flow during X.509 certificate-based peer authentication. The table that follows the figure describes each step in the message flow.

Figure 93. X.509 Certificate-based Peer Authentication

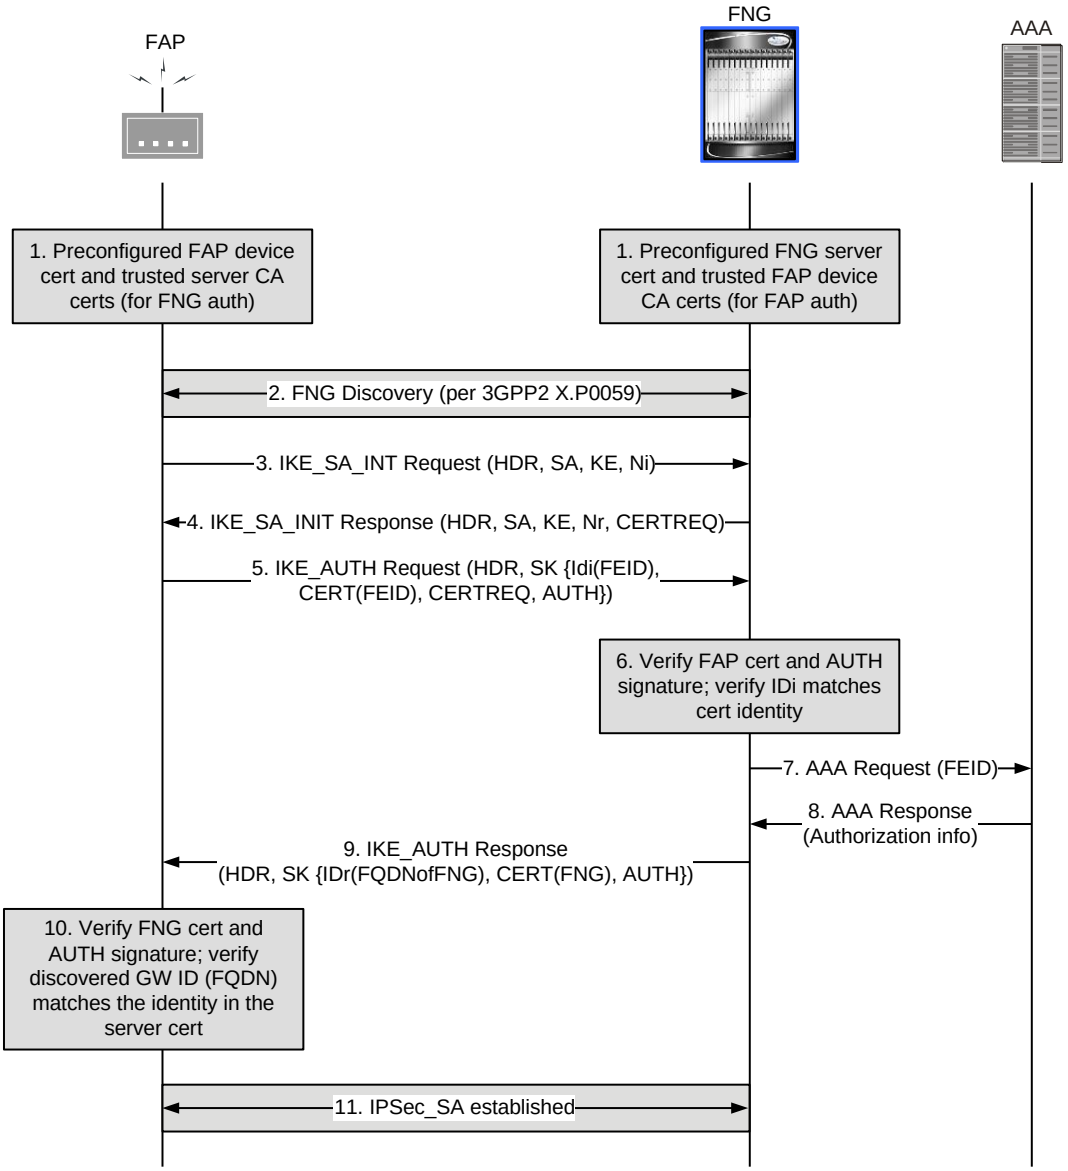


Table 61. X.509 Certificate-based Peer Authentication

Step	Description
1.	The FAP is assigned a device certificate during it's manufacturing. The FAP device certificate is signed by a Certificate Authority (device certificate CA) trusted by the operator. The private key for the certificate is stored securely at the FAP. Similarly, the FNG is assigned a server certificate. The private key of the FNG is stored securely at the FNG. In addition, the FNG is configured with a list of root CA certificates corresponding to the trusted device certificate CAs. The FAP is also configured with a list of root CA certificates corresponding to the server certificates that the FAP will accept from the FNG.

Step	Description
2.	Upon FAP power-up, using the FNG discovery procedures such as DNS discovery, the FAP determines the FQDN/IP address of the appropriate FNG.
3.	The FAP initiates an IKEv2 exchange with the FNG, known as the IKE_SA_INIT exchange, by issuing an IKE_SA_INIT Request to negotiate cryptographic algorithms, exchange nonces, and perform a Diffie-Hellman exchange with the FNG. In addition, using the NAT Traversal procedures, the FAP includes NAT_DETECTION_SOURCE_IP and NAT_DETECTION_DESTINATION_IP payloads to negotiate support for UDP encapsulation.
4.	The FNG responds with an IKE_SA_INIT Response by choosing a cryptographic suite from the initiator's offered choices, completing the Diffie-Hellman and nonce exchanges with the FAP. In addition, the FNG includes the list of FAP CA certificates that it will accept in its CERTREQ payload. For successful FAP authentication, the CERTREQ payload must contain at least one CA certificate that is in the trust chain of the FAP device certificate. At this point in the negotiation, the IKE_SA_INIT exchange is complete and all but the headers of all the messages that follow are encrypted and integrity-protected.
5.	The FAP initiates an IKE_AUTH exchange with the FNG by setting the IDi payload to the FEID in FQDN format (from the subjectAltName extension of the FAP certificate), setting the CERT payload to the FAP device certificate corresponding to the FEID, and including the AUTH payload containing the signature of the previous IKE_SA_INIT Request message (in step 3) generated using the private key of the FAP device certificate. The authentication algorithm used to generate the AUTH payload is also included in the AUTH payload. The FAP also includes the CERTREQ payload containing the list of SHA-1 hash algorithms for server authentication. For successful server authentication, the CERTREQ payload must contain at least one CA certificate that is in the trust chain of the FNG server certificate.
6.	Using the CA certificate corresponding to the FAP device certificate, the FNG first verifies that the FAP device certificate in the CERT payload has not been modified and the identity included in the IDi corresponds to the identity in the FAP device certificate. If the verification is successful, using the public key of the FAP device certificate, the FNG generates the expected AUTH payload and compares it with the received AUTH payload. If they match, the authentication of the FAP is successful. Otherwise, the FNG sends an IKEv2 Notification message indicating authentication failure.
7.	If the network policy requires femtocell subscription authorization, the FNG contacts the AAA server to verify that the FAP identified by the FEID is authorized to provide service.
8.	The AAA server responds with the authorization result. If the authorization is not successful, the FNG sends an IKEv2 Notification message indicating authorization failure. Otherwise, the FNG proceeds with server authentication.
9.	The FNG responds with the IKE_AUTH Response by setting the IDr payload to the FQDN (or IP address) of the FNG, setting the CERT payload to the FNG server certificate corresponding to the FQDN (or IP address), and including the AUTH payload containing the signature of the IKE_SA_INIT Response message (in step 4) generated using the private key of the FNG server certificate. The authentication algorithm used to generate the AUTH payload is also included in the AUTH payload.
10.	Using the CA certificate corresponding to the FNG server certificate, the FAP first verifies that the FNG server certificate in the CERT payload has not been modified and the identity included in the IDi corresponds to the identity in the server certificate and contains the expected FNG value as discovered during the FNG discovery procedures. If the verification is successful, using the public key of the FNG server certificate, the FAP generates the expected AUTH payload and compares it with the received AUTH payload. If they match, FNG server authentication is successful. This completes the IKE_AUTH exchange.
11.	An IPSec SA is established between the FAP and the FNG. If more IPSec SAs are needed, either the FAP or the FNG can initiate the creation of additional Child SAs using a CREATE_CHILD_SA exchange.

Supported Standards

The FNG service complies with the following standards:

- 3GPP2 References
- IETF References

3GPP2 References

- 3GPP2 X.S0059-000-0 (V1.0): “cdma2000 Femtocell Network: Overview”.
- 3GPP2 X.S0059-100-0 (V1.0): “cdma2000 Femtocell Network: Packet Data Network Aspects”.
- 3GPP2 X.P0059-200-0_v0.C_1x_Femto_R&F.pdf

IETF References

- RFC 2401 (November 1998): “Security Architecture for the Internet Protocol”.
- RFC 2402 (November 1998): “IP Authentication Header”.
- RFC 2403 (November 1998): “The Use of HMAC-MD5-96 within ESP and AH”.
- RFC 2404 (November 1998): “The Use of HMAC-SHA1-96 within ESP and AH”.
- RFC 2405 (November 1998): “The ESP DES-CBC Cipher Algorithm With Explicit IV”.
- RFC 2406 (November 1998): “IP Encapsulating Security Payload (ESP)”.
- RFC 2410 (November 1998): “The NULL Encryption Algorithm and Its Use With IPsec”.
- RFC 3168 (September 2001): “The Addition of Explicit Congestion Notification (ECN) to IP”.
- RFC 3526 (May 2003): “More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)”.
- RFC 3602 (May 2003): “The AES-CBC Cipher Algorithm and Its Use with IPsec”.
- RFC 3706 (February 2004): “A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers”.
- RFC 3715 (March 2004): “IPsec-Network Address Translation (NAT) Compatibility Requirements”.
- RFC 3748 (June 2004): “Extensible Authentication Protocol (EAP)”.
- RFC 3947 (January 2005): “Negotiation of NAT-Traversal in the IKE”.
- RFC 3948 (January 2005): “UDP Encapsulation of IPsec ESP Packets”.
- RFC 4187 (January 2006): “Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA)”.
- RFC 4306 (December 2005): “Internet Key Exchange (IKEv2) protocol”.
- RFC 4308 (December 2005): “Cryptographic Suites for IPsec”.
- RFC 4764 (January 2007): “The EAP-PSK Protocol: A Pre-Shared Key Extensible Authentication Protocol (EAP) Method”.
- RFC 4894 (May 2007): “Use of Hash Algorithms in Internet Key Exchange (IKE)”.

Chapter 15
GGSN Support in GPRS/UMTS Wireless Data Services

The Cisco systems provides wireless carriers with a flexible solution that functions as a Gateway GPRS Support Node (GGSN) in General Packet Radio Service (GPRS) or Universal Mobile Telecommunications System (UMTS) wireless data networks.

This overview provides general information about the GGSN including:

- Product Description
- Product Specification
- Network Deployment and Interfaces
- Features and Functionality - Base Software
- Features and Functionality - Optional Enhanced Feature Software
- How GGSN Works
- Supported Standards

Product Description

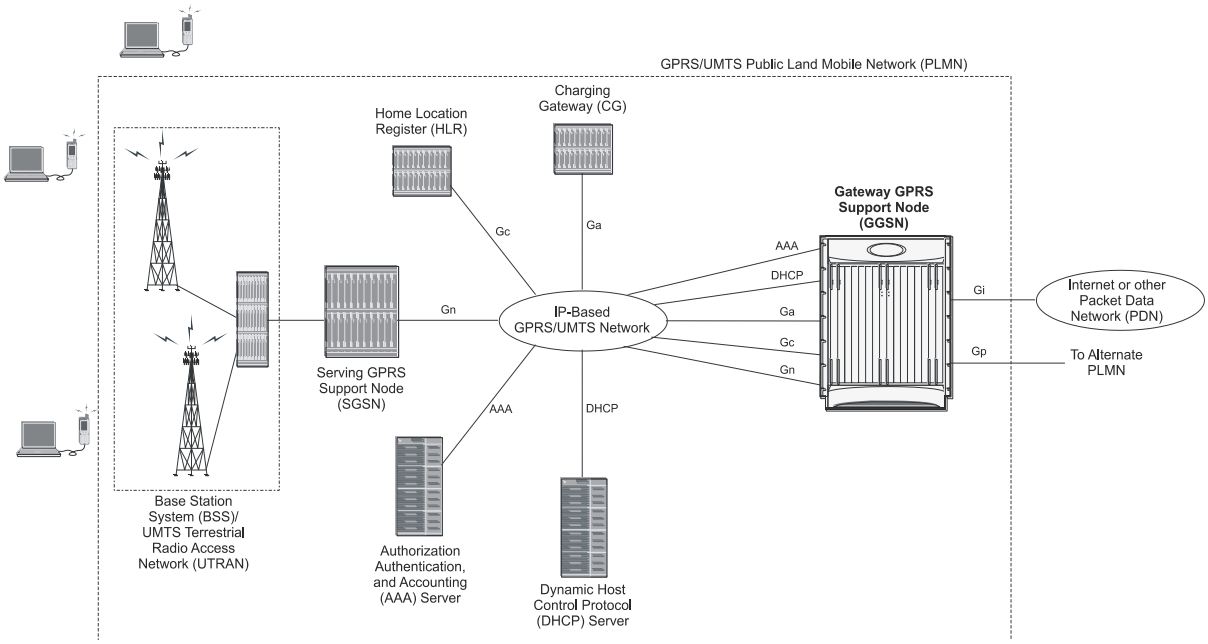
The GGSN works in conjunction with Serving GPRS Support Nodes (SGSNs) within the network to perform the following functions:

- Establish and maintain subscriber Internet Protocol (IP) or Point-to-Point Protocol (PPP) type Packet Data Protocol (PDP) contexts originated by either the mobile or the network
- Provide charging detail records (CDRs) to the charging gateway (CG, also known as the Charging Gateway Function (CGF))
- Route data traffic between the subscriber's Mobile Station (MS) and a Packet Data Networks (PDNs) such as the Internet or an intranet

PDNs are associated with Access Point Names (APNs) configured on the system. Each APN consists of a set of parameters that dictate how subscriber authentication and IP address assignment is to be handled for that APN.

In addition, to providing basic GGSN functionality as described above, the system can be configured to support Mobile IP and/or Proxy Mobile IP data applications in order to provide mobility for subscriber IP PDP contexts. When supporting these services, the system can be configured to either function as a GGSN and Foreign Agent (FA), a stand-alone Home Agent (HA), or a GGSN, FA, and HA simultaneously within the carrier's network.

Figure 94. Basic GPRS/UMTS Network Topology



In accordance with RFC 2002, the FA is responsible for mobile node registration with, and the tunneling of data traffic to/from the subscriber's home network. The HA is also responsible for tunneling traffic, but also maintains subscriber location information in Mobility Binding Records (MBRs).

Product Specification

This section describes the hardware and software requirement for GGSN service.

The following information is located in this section:

- Licenses
- Platform Requirements
- Operating System Requirements

Licenses

The GGSN is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Platform Requirements

The GGSN service runs on a Cisco® ASR 5x00 Series chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the Installation Guide for the chassis and/or contact your Cisco account representative.

Operating System Requirements

The GGSN is available for chassis running StarOS™ Release 7.1 or later.

Network Deployment and Interfaces

This section describes the supported interfaces and deployment scenario of GGSN in GPRS/UMTS network.

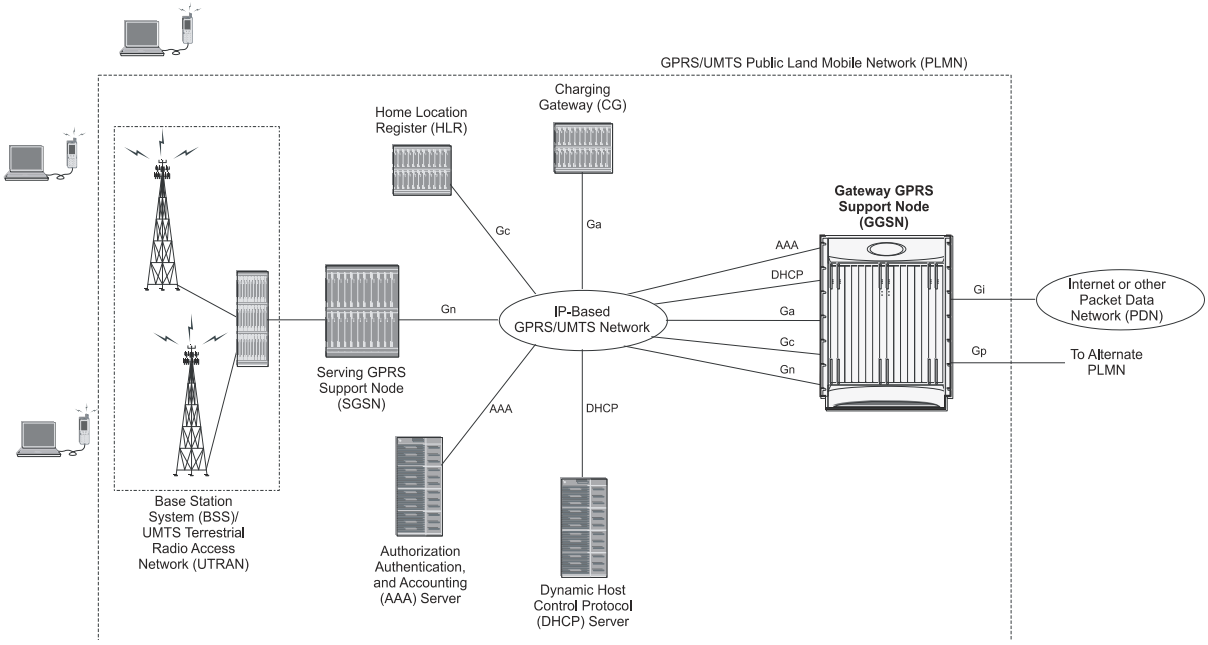
The following information is provided in this section:

- GGSN in the GPRSUMTS Data Network
- Supported Interfaces

GGSN in the GPRS/UMTS Data Network

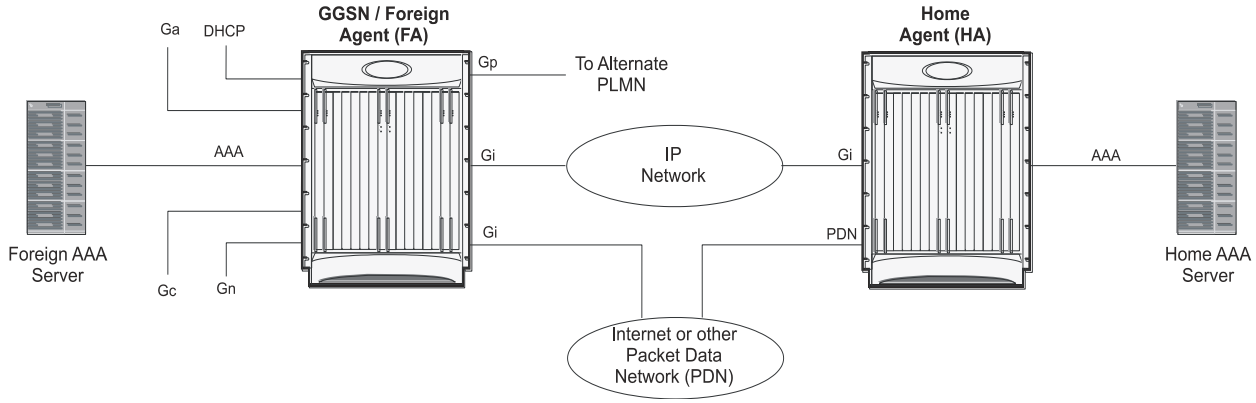
The figures that follow display simplified network views of the GGSN in a GPRS/UMTS network and the system supporting Mobile IP and Proxy Mobile IP function both the GGSN/Foreign Agent (FA) and GGSN/FA/Home Agent (HA) combinations respectively.

Figure 95. Basic GPRS/UMTS Network Topology



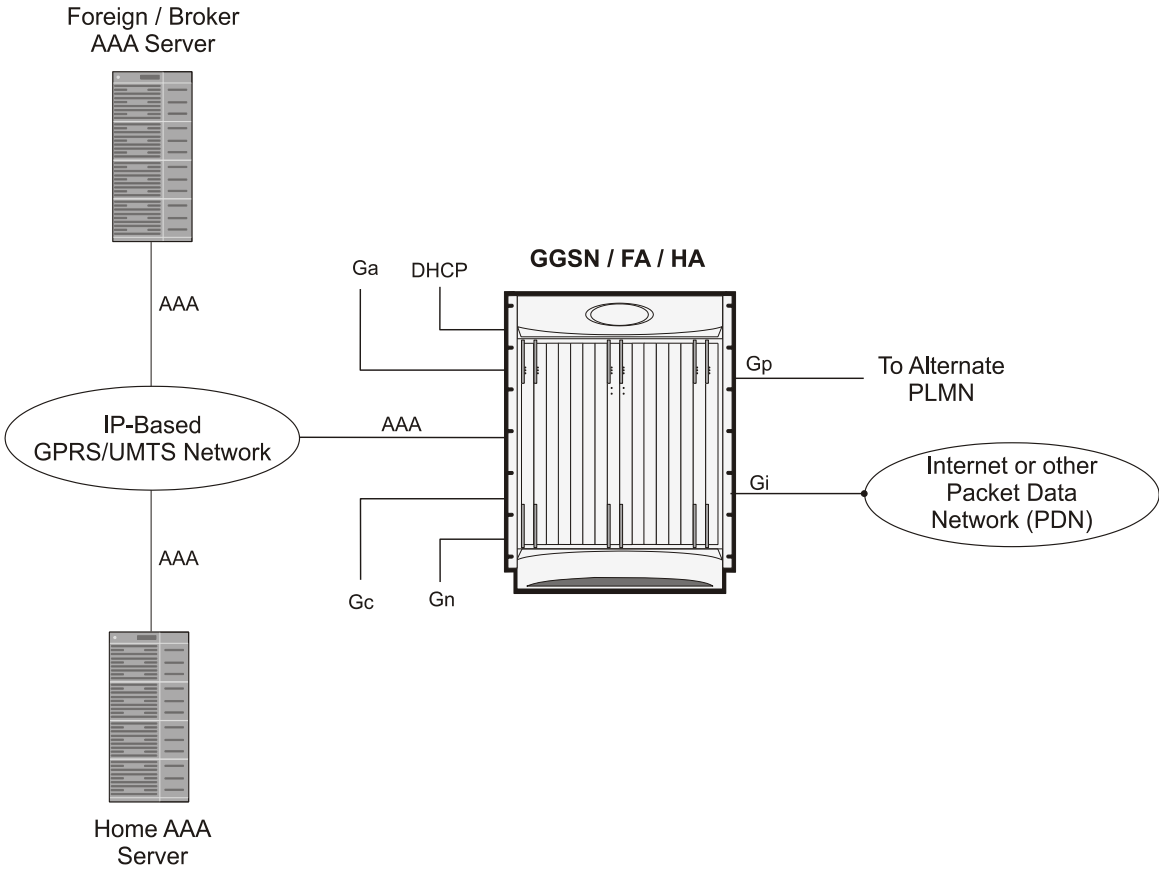
The figures that follow display simplified network views of the GGSN in a GPRS/UMTS network and the system supporting Mobile IP and Proxy Mobile IP function both the GGSN/Foreign Agent (FA) and GGSN/FA/Home Agent (HA) combinations respectively.

Figure 96. Combined GGSN/FA Deployment for Mobile IP and/or Proxy Mobile IP Support



The figures that follow display simplified network views of the GGSN in a GPRS/UMTS network and the system supporting Mobile IP and Proxy Mobile IP function both the GGSN/Foreign Agent (FA) and GGSN/FA/Home Agent (HA) combinations respectively.

Figure 97. Combined GGSN/FA/HA Deployment for Mobile IP and/or Proxy Mobile IP Support



Supported Interfaces

In support of both mobile and network originated subscriber PDP contexts, the system GGSN provides the following network interfaces:

- **Gn**: This is the interface used by the GGSN to communicate with SGSNs on the same GPRS/UMTS Public Land Mobile Network (PLMN). This interface serves as both the signaling and data path for establishing and maintaining subscriber PDP contexts.

  The GGSN communicates with SGSNs on the PLMN using the GPRS Tunnelling Protocol (GTP). The signaling or control aspect of this protocol is referred to as the GTP Control Plane (GTPC) while the encapsulated user data traffic is referred to as the GTP User Plane (GTPU).

  One or more Gn interfaces can be configured per system context.

- **Ga**: This is the interface used by the GGSN to communicate with the Charging Gateway (CG). The charging gateway is responsible for sending GGSN Charging Data Records (G-CDRs) received from the GGSN for each PDP context to the billing system. System supports TCP and UDP as transport layer for this interface.

  The GGSN communicates with the CGs on the PLMN using GTP Prime (GTPP).

  One or more Ga interfaces can be configured per system context.

- **Gc**: This is the interface used by the GGSN to communicate with the Home Location Register (HLR) via a GTP-to-MAP (Mobile Application Part) protocol convertor. This interface is used for network initiated PDP contexts.

  For network initiated PDP contexts, the GGSN will communicate with the protocol convertor using GTP. The convertor, in turn, will communicate with the HLR using MAP over Signaling System 7 (SS7).

  One Gc interface can be configured per system context.

- **Gi**: This is the interface used by the GGSN to communicate with Packet Data Networks (PDNs) external to the PLMN. Examples of PDNs are the Internet or corporate intranets.

  Inbound packets received on this interface could initiate a network requested PDP context if the intended MS is not currently connected.

  For systems configured as a GGSN/FA, this interface is used to communicate with HAs for Mobile IP and Proxy Mobile IP support.

  One or more Gi interfaces can be configured per system context. For Mobile IP and Proxy Mobile IP, at least one Gi interface must be configured for each configured FA service. Note that when the system is simultaneously supporting GGSN, FA, and HA services, traffic that would otherwise be routed over the Gi interface is routed inside the chassis.

- **Gp**: This is the interface used by the GGSN to communicate with GPRS Support Nodes (GSNs, e.g. GGSNs and/or SGSNs) on different PLMNs. Within the system, a single interface can serve as both a Gn and a Gp interface.

  One or more Gn/Gp interfaces can be configured per system context.

- **AAA**: This is the interface used by the GGSN to communicate with an authorization, authentication, and accounting (AAA) server on the network. The system GGSN communicates with the AAA server using the Remote Authentication Dial In User Service (RADIUS) protocol.

  This is an optional interface that can be used by the GGSN for subscriber PDP context authentication and accounting.

- **DHCP**: This is the interface used by the GGSN to communicate with a Dynamic Host Control Protocol (DHCP) Server. The system can be configured as DHCP-Proxy or DHCP Client to provide IP addresses to MS on PDP contexts activation the DHCP server dynamically.

- **Gx**: This is an optional Diameter protocol-based interface over which the GGSN communicates with a Charging Rule Function (CRF) for the provisioning of charging rules that are based on the dynamic analysis of flows used for an IP Multimedia Subsystem (IMS) session. The system provides enhanced support for use of Service Based Local Policy (SBLP) to provision and control the resources used by the IMS subscriber. It also provides Flow based Charging (FBC) mechanism to charge the subscriber dynamically based on content usage.

  Important: The Gx interface is a license-enabled support. For more information on this support, refer *Gx Interface Support* in this guide.

- **Gy**: This is an optional Diameter protocol-based interface over which the GGSN communicates with a Charging Trigger Function (CTF) server that provides online charging data. Gy interface support provides an online charging interface that works with the ECS deep packet inspection feature. With Gy, customer traffic can be gated and billed in an "online" or "prepaid" style. Both time- and volume-based charging models are supported. In all of these models, differentiated rates can be applied to different services based on shallow or deep packet inspection.

  Important: This interface is supported through Enhanced Charging Service. For more information on this support, refer *Enhanced Charging Service Administration Guide*.

- **GRE**: This new protocol interface in GGSN platform adds one additional protocol to support mobile users to connect to their enterprise networks: Generic Routing Encapsulation (GRE). GRE Tunneling is a common technique to enable multi-protocol local networks over a single-protocol backbone, to connect non-contiguous networks and allow virtual private networks across WANs. This mechanism encapsulates data packets from one protocol inside a different protocol and transports the data packets unchanged across a foreign network. It is important to note that GRE tunneling does not provide security to the encapsulated protocol, as there is no encryption involved (like IPSEC offers, for example).

  Important: The GRE protocol interface is a license-enabled support. For more information on this support, refer *GRE Protocol Interface Support* in this guide.

- **S6b**: This is an optional Diameter protocol-based interface over which the GGSN communicates with 3G AAA/HSS in LTE/SAE network for subscriber authorization.

  Important: The S6b interface can still be disabled via the CLI per the existing MOPs in the event of a long-term AAA outage

  Important: This interface is supported through license-enabled feature. For more information on this support, refer *Common Gateway Access Support* section of this guide.

Important: GGSN Software also supports additional interfaces. For more information on additional interfaces, refer *Features and Functionality - Optional Enhanced Feature Software* section.

Features and Functionality - Base Software

This section describes the features and functions supported by default in base software on GGSN service and do not require any additional licenses.

Important: To configure the basic service and functionality on the system for GGSN service, refer configuration examples provide in *GGSN Administration Guide*.

This section describes following features:

- 16,000 SGSN Support
- AAA Server Groups
- Access Control List Support
- ANSI T1.276 Compliance
- APN Support
- Bulk Statistics Support
- Direct Tunnel Support
- DHCP Support
- DSCP Marking
- Framed-Route Attribute Support
- Generic Corporate APN
- GnGp Handoff Support
- GTPP Support
- Host Route Advertisement
- IP Policy Forwarding
- IP Header Compression - Van Jacobson
- IPv6 Support
- Management System Overview
- MPLS Forwarding with LDP
- Overlapping IP Address Pool Support
- Per APN Configuration to Swap out Gn to Gi APN in CDRs
- Port Insensitive Rule for Enhanced Charging Service
- Quality of Service Support
- RADIUS Support
- PDP Context Support
- RADIUS VLAN Support

- Routing Protocol Support
- Subscriber Session Trace Support
- Support of Charging Characteristics Provided by AAA Server
- Support of all GGSN generated causes for partial G-CDR closure
- Threshold Crossing Alerts (TCA) Support

16,000 SGSN Support

With growing roaming agreements, many more GPRS/UMTS networks support certain APNs and therefore the number of SGSNs that could connect to the GGSN increases. This feature increases the number of connected SGSNs thereby allowing a single GGSN service to support a much larger roaming network.

The GGSN service supports a maximum of 16,000 SGSN IP addresses. The chassis limit for bulk statistics collection is also limit to 16,000. No change in configuration is needed to support this feature.

AAA Server Groups

Value-added feature to enable VPN service provisioning for enterprise or MVNO customers. Enables each corporate customer to maintain its own AAA servers with its own unique configurable parameters and custom dictionaries.

This feature provides support for up to 800 AAA (RADIUS and Diameter) server groups and 800 NAS IP addresses that can be provisioned within a single context or across the entire chassis. A total of 128 servers can be assigned to an individual server group. Up to 1,600 accounting, authentication and/or mediation servers are supported per chassis and may be distributed across a maximum of 1,000 APNs. This feature also enables the AAA servers to be distributed across multiple APN within the same context.

Important: In 12.3 and earlier releases, refer to the *AAA and GTPP Interface Administration and Reference* for more information on AAA Server Group configuration.

Access Control List Support

Access Control Lists provide a mechanism for controlling (i.e permitting, denying, redirecting, etc.) packets in and out of the system.

IP access lists, or Access Control Lists (ACLs) as they are commonly referred to, are used to control the flow of packets into and out of the system. They are configured on a per-context basis and consist of "rules" (ACL rules) or filters that control the action taken on packets that match the filter criteria

Once configured, an ACL can be applied to any of the following:

- An individual interface
- All traffic facilitated by a context (known as a policy ACL)
- An individual subscriber
- All subscriber sessions facilitated by a specific context

There are two primary components of an ACL:

- Rule: A single ACL consists of one or more ACL rules. As discussed earlier, the rule is a filter configured to take a specific action on packets matching specific criteria. Up to 128 rules can be configured per ACL.

  Each rule specifies the action to take when a packet matches the specifies criteria. This section discusses the rule actions and criteria supported by the system.
- Rule Order: A single ACL can consist of multiple rules. Each packet is compared against each of the ACL rules, in the order in which they were entered, until a match is found. Once a match is identified, all subsequent rules are ignored.

Important: For more information on Access Control List configuration, refer *IP Access Control List* in *System Administration Guide*.

ANSI T1.276 Compliance

ANSI T1.276 specifies security measures for Network Elements (NE). In particular it specifies guidelines for password strength, storage, and maintenance security measures.

ANSI T1.276 specifies several measures for password security.

These measures include:

- Password strength guidelines
- Password storage guidelines for network elements
- Password maintenance, e.g. periodic forced password changes

These measures are applicable to the chassis and the Web Element Manager since both require password authentication. A subset of these guidelines where applicable to each platform will be implemented. A known subset of guidelines, such as certificate authentication, are not applicable to either product. Furthermore, the platforms support a variety of authentication methods such as RADIUS and SSH which are dependent on external elements. ANSI T1.276 compliance in such cases will be the domain of the external element. ANSI T1.276 guidelines will only be implemented for locally configured operators.

APN Support

The GGSN's Access Point Name (APN) support offers several benefits:

- Extensive parameter configuration flexibility for the APN.
- Creation of subscriber tiers for individual subscribers or sets of subscribers within the APN.
- Virtual APNs to allow differentiated services within a single APN.

Up to 1024 APNs can be configured in the GGSN. An APN may be configured for any type of PDP context, i.e., PPP, IPv4, IPv6 or both IPv4 and IPv6. Many dozens of parameters may be configured independently for each APN.

Here are a few highlights of what may be configured:

- **Accounting**: RADIUS, GTPP or none. Server group to use. Charging characteristics. Interface with mediation servers.
- **Authentication**: Protocol, such as, CHAP or PAP or none. Default username/password. Server group to use. Limit for number of PDP contexts.
- **Enhanced Charging**: Name of rulebase to use, which holds the enhanced charging configuration (e.g., eG-CDR variations, charging rules, prepaid/postpaid options, etc.).

- **IP**: Method for IP address allocation (e.g., local allocation by GGSN, Mobile IP, DHCP, DHCP relay, etc.). IP address ranges, with or without overlapping ranges across APNs.
- **Tunneling**: PPP may be tunneled with L2TP. IPv4 may be tunneled with GRE, IP-in-IP or L2TP. Load-balancing across multiple tunnels. IPv6 is tunneled in IPv4. Additional tunneling techniques, such as, IPsec and VLAN tagging may be selected by the APN, but are configured in the GGSN independently from the APN.
- **QoS**: IPv4 header ToS handling. Traffic rate limits for different 3GPP traffic classes. Mapping of R98 QoS attributes to work around particular handset defections. Dynamic QoS renegotiation (described elsewhere).

After an APN is determined by the GGSN, the subscriber may be authenticated/authorized with an AAA server. The GGSN allows the AAA server to return VSAs (Vendor Specific Attributes) that override any/all of the APN configuration. This allows different subscriber tier profiles to be configured in the AAA server, and passed to the GGSN during subscriber authentication/authorization.

The GGSN's Virtual APN feature allows the carrier to use a single APN to configure differentiated services. The APN that is supplied by the SGSN is evaluated by the GGSN in conjunction with multiple configurable parameters. Then the GGSN selects an APN configuration based on the supplied APN and those configurable parameters. The configurable parameters are: the access gateway IP address, bearer access service name, charging characteristics (CC)-profile index, subscribers within an MSISN range, subscriber's mcc/mnc, whether the subscriber is home/visiting/roaming, subscriber's domain name and the radio access (RAT) type including gen, geran, hspa, eutran, utran, and wlan.

Important: For more information on APN configuration, refer *APN Configuration* in *GGSN Service Configuration*.

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema.

The following schemas are supported for GGSN service:

- **System**: Provides system-level statistics
- **Card**: Provides card-level statistics
- **Port**: Provides port-level statistics
- **FA**: Provides FA service statistics
- **HA**: Provides HA service statistics
- **IP Pool**: Provides IP pool statistics
- **PPP**: Provides Point-to-Point Protocol statistics
- **GTPC**: Provides GPRS Tunneling Protocol - Control message statistics
- **GTPP**: Provides GPRS Tunneling Protocol - Prime message statistics
- **APN**: Provides Access Point Name statistics
- **RADIUS**: Provides per-RADIUS server statistics
- **ECS**: Provides Enhanced Charging Service Statistics

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, system host name, system uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

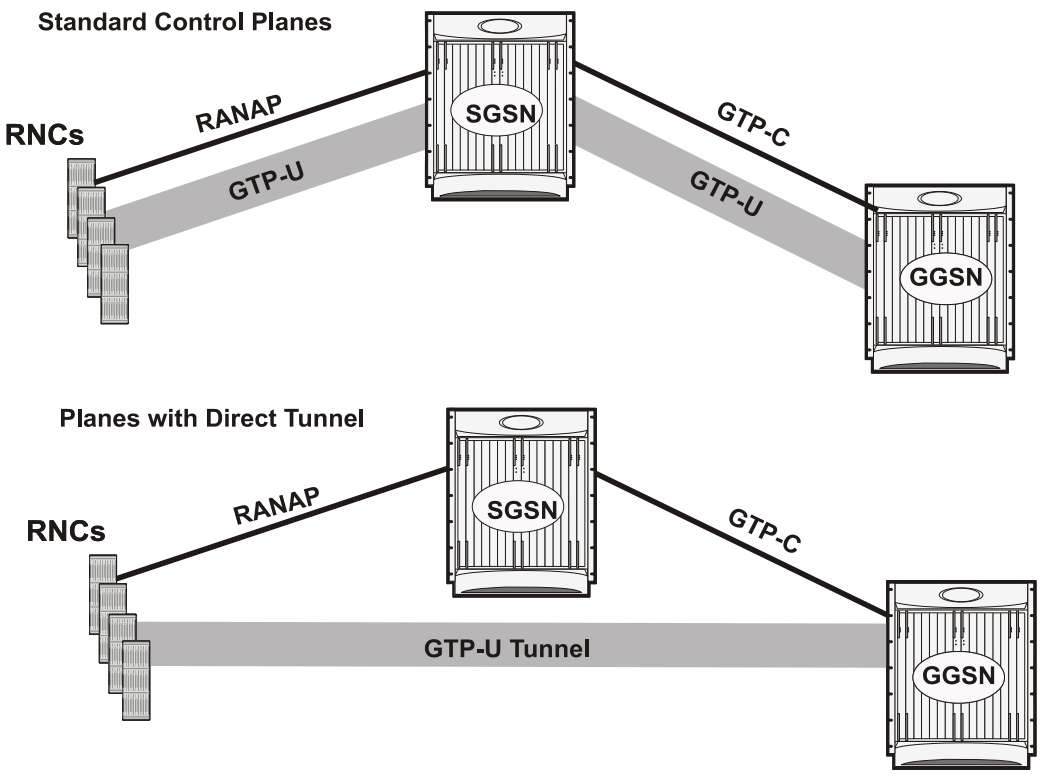
Direct Tunnel Support

Direct tunnel improves the user experience (e.g. expedited web page delivery, reduced round trip delay for conversational services, etc.) by eliminating SGSN tunnel 'switching' latency from the user plane. An additional advantage of Direct Tunnel from an operational and capital expenditure perspective is that direct tunnel optimizes the usage of user plane resources by removing the requirement for user plane processing on the SGSN.

The Direct Tunnel architecture allows the establishment of a direct user plane tunnel between the RAN and the GGSN, bypassing the SGSN. The SGSN continues to handle the control plane signalling and typical makes the decision to establish Direct Tunnel at PDP Context Activation. A Direct Tunnel is achieved at PDP context activation by the SGSN establishing a user plane (GTP-U) tunnel directly between RNC and GGSN (using an Update PDP Context Request towards the GGSN).

The following figure illustrates the working of Direct Tunnel between RNC and GGSN.

Figure 98. Direct Tunnel Support in GGSN



A major consequence of deploying Direct Tunnel is that it produces a significant increase in control plane load on both the SGSN and GGSN components of the packet core. It is therefore of paramount importance to a wireless operator to ensure that the deployed GGSNs are capable of handling the additional control plane loads introduced of part of Direct Tunnel deployment. The Cisco GGSN and SGSN offers massive control plane transaction capabilities, ensuring system control plane capacity will not be a capacity limiting factor once Direct Tunnel is deployed.

DHCP Support

Dynamic IP address assignment to subscriber IP PDP contexts using the Dynamic Host Control Protocol as defined by the following standards:

- RFC 2131, Dynamic Host Configuration Protocol
- RFC 2132, DHCP Options and BOOTP Vendor Extensions

As described in the PDP Context Support section of this document, the method by which IP addresses are assigned to a PDP context is configured on an APN-by-APN basis. Each APN template dictates whether it will support static or dynamic addresses.

Dynamically assigned IP addresses for subscriber PDP contexts can be assigned through the use of DHCP.

The system can be configured to support DHCP using either of the following mechanisms:

- **DHCP-proxy**: The system acts as a proxy for client (MS) and initiates the DHCP Discovery Request on behalf of client (MS). Once it receives an allocated IP address from DHCP server in response to DHCP Discovery Request, it assigns the received IP address to the MS. This allocated address must be matched with the an address configured in an IP address pool on the system. This complete procedure is not visible to MS.

- **DHCP-relay**: The system acts as a relay for client (MS) and forwards the DHCP Discovery Request received from client (MS). Once it receives an allocated IP address from DHCP server in response to DHCP Discovery Request, it assigns the received IP address to the MS.

Important: For more information on DHCP service configuration, refer *DHCP Configuration* section in *GGSN Service Configuration* chapter.

DSCP Marking

Provides support for more granular configuration of DSCP marking.

For different Traffic class, the GGSN supports per-GGSN service and per-APN configurable DSCP marking for Uplink and Downlink direction based on Allocation/Retention Priority in addition to the current priorities.

Framed-Route Attribute Support

The Framed-Route attribute provides routing information to be configured for the user on the network access server (NAS). The Framed-Route information is returned to the RADIUS server in the Accounting Access-Accept message.

Mobile Router enables a router to create a PDP context which the GGSN authorizes using RADIUS server. The RADIUS server authenticates this router and includes a Framed-Route attribute in the access-accept response packet. Framed-Route attribute also specifies the subnet routing information to be installed in the GGSN for the "mobile router." If the GGSN receives a packet with a destination address matching the Framed-Route, the packet is forwarded to the mobile router through the associated PDP context.

Generic Corporate APN

Any operator may not be aware of the IP address that a corporation may assign to subscribers through AAA or DHCP and the traffic is sent from the GGSN to the corporation over a tunnel, this feature allows the operator to terminate such users.

Normally the GGSN validates the IP address assigned by RADIUS, however this feature removes the need for this, but does assume that the subscriber traffic is forwarded out of the GGSN through a tunnel.

When the IP address is statically assigned, i.e., either MS provided, RADIUS provided or DHCP provided, the IP address validation is not performed if the address policy is set to disable address validation.

ACL and Policy Group Info processing would still be performed.

Additionally, there is support for Virtual APN selection based on RADIUS VSA returned during Authentication.

The existing Virtual APN selection mechanism is being enhanced to select the Virtual APN based on RADIUS VSA returned during authentication.

The selected V-APN may further require AAA authentication (and accounting) with its own servers.

GnGp Handoff Support

In LTE deployments, the smooth handover support is required between 3G/2G and LTE networks, and Evolved Packet Core (EPC) is designed to be a common packet core for different access technologies. Since support for seamless

handover across different access technologies is basic requirement for EPC, PGW needs to support handovers as user equipment (UE) moves across different access technologies.

Cisco's PGW supports inter-technology mobility handover between 4G and 3G/2G access. Interworking is supported between the 4G and 2G/3G SGSNs which provide only Gn and Gp interfaces but no S3, S4 or S5/S8 interfaces. Therefore these Gn/Gp SGSNs provide no functionality introduced specifically for the evolved packet system (EPS) or for interoperation with the E-UTRAN. These handovers are supported only with a GTP-based S5/S8 and PGW supports handovers between GTPv2 based S5/S8 and GTPv1 based Gn/Gp tunneled connections. In this scenario, the PGW works as an IP anchor for the EPC.

GnGp Handoff in Non-Roaming Scenario

Depending on the existing deployments, PLMN may operate Gn/Gp 2G and/or 3G SGSNs as well as MME and SGW for E-UTRAN access. In such cases, the PGW works as an anchor point for both GERAN/UTRAN and E-UTRAN access. Depending on APN, MME/SGSN select a PGW for each call.

In the home network (non-roaming) when UE firstly attaches to the E-UTRAN, it sets up a PDN connection with some EPS bearers and when the UE moves to Gn/Gp SGSN served GERAN/UTRAN access, handover is initiated from MME to the Gn/Gp SGSN. Gn/Gp SGSN then notifies PGW (with GGSN functionality) about the handoff of EPS bearers. During this handover, each EPS bearer in the PDN connection is converted into a PDP context.

The other way, when the UE first attaches on to Gn/Gp SGSN served GERAN/UTRAN, it sets up PDP contexts, and when the UE moves to E-UTRAN access, handover is initiated from Gn/Gp SGSN to the MME. MME then notifies the PGW (through SGW) about the handoff of PDP contexts to the E-UTRAN access. During this handover, all PDP contexts sharing the same APN and IP address are converted to EPS bearers of same PDN connection. Here one of the PDP context is selected as a Default bearer and rest of the PDP contexts are designated as Dedicated bearers.

GnGp Handoff in Roaming Scenario

In the roaming scenario, the vPLMN (Virtual PLMN) operates Gn/Gp 2G and/or 3G SGSNs as well as MME and SGW for E-UTRAN access and hPLMN (Home PLMN) operates a PGW. Other remaining things work as in non-roaming scenario.

Important: For more information on configuration of Gn-Gp Handoff, refer the *Gn-Gp Support Configuration* section of *GGSN Service Configuration Procedures* chapter.

GTPP Support

Support for the GPRS Tunnelling Protocol Prime (GTPP) in accordance with the following standards:

- **3GPP TS 32.015 v3.12.0 (2003-12)**: 3rd Generation Partnership project; Technical Specification Group Services and System Aspects; Telecommunication Management; Charging and billing; GSM call and event data for the Packet Switched (PS) domain (Release 1999) for support of Charging on GGSN
- **3GPP TS 32.215 v5.9.0 (2005-06)**: 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Charging management; Charging data description for the Packet Switched (PS) domain (Release 4)
- **3GPP TS 29.060 v7.9.0 (2008-09)**: Technical Specification; 3rd Generation Partnership Project; Technical Specification Group Core Network; General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface (Release 6)

The system supports the use of GTPP for PDP context accounting. When the GTPP protocol is used, accounting messages are sent to the Charging Gateways (CGs) over the Ga interface. The Ga interface and GTPP functionality are typically configured within the system's source context. As specified by the standards, a CDR is not generated when a session starts. CDRs are generated according to the interim triggers configured using the charging characteristics

configured for the GGSN, and a CDR is generated when the session ends. For interim accounting, STOP/START pairs are sent based on configured triggers.

GTPP version 2 is always used. However, if version 2 is not supported by the CGF, the system reverts to using GTPP version 1. All subsequent CDRs are always fully-qualified partial CDRs. All CDR fields are R4.

Whether or not the GGSN accepts charging characteristics from the SGSN can be configured on a per-APN basis based on whether the subscriber is visiting, roaming or, home.

By default, the GGSN always accepts the charging characteristics from the SGSN. They must always be provided by the SGSN for GTPv1 requests for primary PDP contexts. If they are not provided for secondary PDP contexts, the GGSN re-uses those from the primary.

If the system is configured to reject the charging characteristics from the SGSN, the GGSN can be configured with its own that can be applied based on the subscriber type (visiting, roaming, or home) at the APN level. GGSN charging characteristics consist of a profile index and behavior settings. The profile indexes specify the criteria for closing accounting records based specific criteria.

Important: For more information on GTPP group configuration, refer *GTPP Accounting Configuration* in *GGSN Service Configuration* chapter.

Host Route Advertisement

When subscribers are assigned IP addresses from RADIUS or HLR, yet are allowed to connect to multiple GGSNs through the use of DNS round robin or failover, the IP addresses of the subscribers can be advertised on a per user (host) basis to the Gi network using dynamic routing, thereby providing IP reachability to these users.

IP address pools are configured on the GGSN for many reasons, although one of them is so that the pool subnets can be automatically advertised to the network. These are connected routes and are advertised for all non-tunneling pools.

A configuration **`explicit-route-advertise`** is provided to the IP pool configuration and when this option is enabled, the subnet(s) of the pool are not added to routing table and routing protocols like OSPF and BGP do not know of these addresses and hence do not advertise the subnet(s).

As calls come up, and addresses from this pool (with the "explicit-route-advertise" flag) are used, the assigned addresses are added to the routing table and these addresses can be advertised by OSPF or BGP through the network or the "redistribute connected" command.

Example

A subscriber connecting to GGSN A with an IP address from a pool P1 will be assigned the IP address and the routing domain will be updated with the host route. When a subscriber connects to GGSN B with an IP address from the same pool, the subscriber will be assigned the requested IP address and the routing domain will then learn its host route. When the subscriber disconnects, the route is removed from the routing table and the routing domain is updated.

The explicit-route-advertise option can be applied and removed from the pool at any time and the routing tables are updated automatically.

The overlap and resource pool behavior does not change therefore it does not make sense to configure an overlap/resource pool with the "explicit-route-advertise" option.

IP Policy Forwarding

IP Policy Forwarding enables the routing of subscriber data traffic to specific destinations based on configuration. This functionality can be implemented in support of enterprise-specific applications (i.e. routing traffic to specific enterprise domains) or for routing traffic to back-end servers for additional processing.

The system can be configured to automatically forward data packets to a predetermined network destination. This can be done in one of three ways:

- **IP Pool-based Next Hop Forwarding** - Forwards data packets based on the IP pool from which a subscriber obtains an IP address.
- **ACL-based Policy Forwarding** - Forwards data packets based on policies defined in Access Control Lists (ACLs) and applied to contexts or interfaces.
- **Subscriber specific Next Hop Forwarding** - Forwards all packets for a specific subscriber.

The simplest way to forward subscriber data is to use IP Pool-based Next Hop Forwarding. An IP pool is configured with the address of a next hop gateway and data packets from all subscribers using the IP pool are forward to that gateway.

Subscriber Next Hop forwarding is also very simple. In the subscriber configuration a nexthop forwarding address is specified and all data packets for that subscriber are forwarded to the specified nexthop destination.

ACL-based Policy Forwarding gives you more control on redirecting data packets. By configuring an Access Control List (ACL) you can forward data packets from a context or an interface by different criteria, such as; source or destination IP address, ICMP type, or TCP/UDP port numbers.

ACLs are applied first. If ACL-based Policy Forwarding and Pool-based Next Hop Forwarding or Subscriber are configured, data packets are first redirected as defined in the ACL, then all remaining data packets are redirected to the next hop gateway defined by the IP pool or subscriber profile.

Important: For more information on IP Policy Forwarding configuration, refer *Policy Forwarding* in this guide.

IP Header Compression - Van Jacobson

Implementing IP header compression provides the following benefits:

- Improves interactive response time
- Allows the use of small packets for bulk data with good line efficiency
- Allows the use of small packets for delay sensitive low data-rate traffic
- Decreases header overhead
- Reduces packet loss rate over lossy links

The system supports the Van Jacobson (VJ) IP header compression algorithms by default for subscriber traffic.

The VJ header compression is supported as per RFC 1144 (CTCP) header compression standard developed by V. Jacobson in 1990. It is commonly known as VJ compression. It describes a basic method for compressing the headers of IPv4/TCP packets to improve performance over low speed serial links.

By default IP header compression using the VJ algorithm is enabled for subscribers. You can also turn off IP header compression for a subscriber.

Important: For more information on IP header compression support, refer *IP Header Compression* in this guide.

IPv6 Support

Native IPv6 support allows for the configuration of interfaces/routes with IPv6 (128 bit) addressing. The increased address space allows for future subscriber growth beyond what is currently possible in IPv4. Native IPv6 support on the Gi interface allows support for packets coming from or destined to a mobile over the Gi interface. IPv6 address assignment is supported from a dynamic or static pool via standard 3GPP attributes. The GGSN can communicate using DIAMETER as the transport protocol for Gx to the AAA. Overlapping address space or resource pools are supported if they are in different VPNs. The VPN subsystem is responsible for the configuration and recovery of IP interfaces and routes. IP resources are grouped into separate routing domains know as contexts. The VPN subsystem creates and maintains each context and the resources associated with them. The existing IPv4 model of interface and route notification will be extended to support IPv6.

This feature allows IPv6 subscribers to connect via the GPRS/UMTS infrastructure in accordance with the following standards:

- RFC 2460: Internet Protocol, Version 6 (IPv6) Specification
- RFC 2461: Neighbor Discovery for IPv6
- RFC 2462: IPv6 Stateless Address Autoconfiguration
- RFC 3314: Recommendations for IPv6 in 3GPP Standards
- RFC 3316: Internet Protocol Version 6 (IPv6) for Some Second and Third Generation Cellular Hosts
- RFC 3056: Connection of IPv6 domains via IPv4 clouds
- 3GPP TS 23.060: General Packet Radio Service (GPRS) Service description
- 3GPP TS 27.060: Mobile Station Supporting Packet Switched Services
- 3GPP TS 29.061: Interworking between the Public Land Mobile Network (PLMN) supporting Packet Based Services and Packet Data Networks (PDN)

IP version 6 is enhanced version of IP version 4 with following modifications:

- Expanded addressing capabilities with 128 bit for address as compared to 32 bits in IPv4.
- Header format simplification
- Improved support of extensions and options
- Flow labeling capability
- Authentication and Privacy capabilities

IPv6 Neighbor Discovery protocol is used to dynamically discover the directly attached devices on IPv6 Interfaces. It facilitates the mapping of MAC addresses to IPv6 Addresses. The GGSN supports a subset of IPv6 Neighbor Discovery as defined by RFC 2461, including the following:

- The GGSN uses IPv6 Neighbor Discovery to learn the Ethernet link-layer addresses of the directly connected next-hop gateway.
- The GGSN supports configuration of the static IPv6 neighbor (next-hop gateway).
- Link-local addresses will be automatically added to Ethernet type interfaces.
- The GGSN performs Unsolicited Neighbor Advertisement on line card switchover.
- The GGSN will reply to neighbor discovery requests for the node's IPv6 addresses.

ICMPv6 is a protocol for IPv6 networks to allow error reporting and check connectivity via echo messages. The GGSN supports a subset of ICMPv6 as defined by [RFC-4443]. The GGSN replies to the link-local, configured IP address, and the all-hosts IP address.

Native IPv6 Routing allows the forwarding of IPv6 packets between IPv6 Networks. The forwarding lookup is based on a longest prefix match of the destination IPv6 address. The GGSN supports configuration of IPv6 routes to directly attached next hops via an IPv6 Interface.

Management System Overview

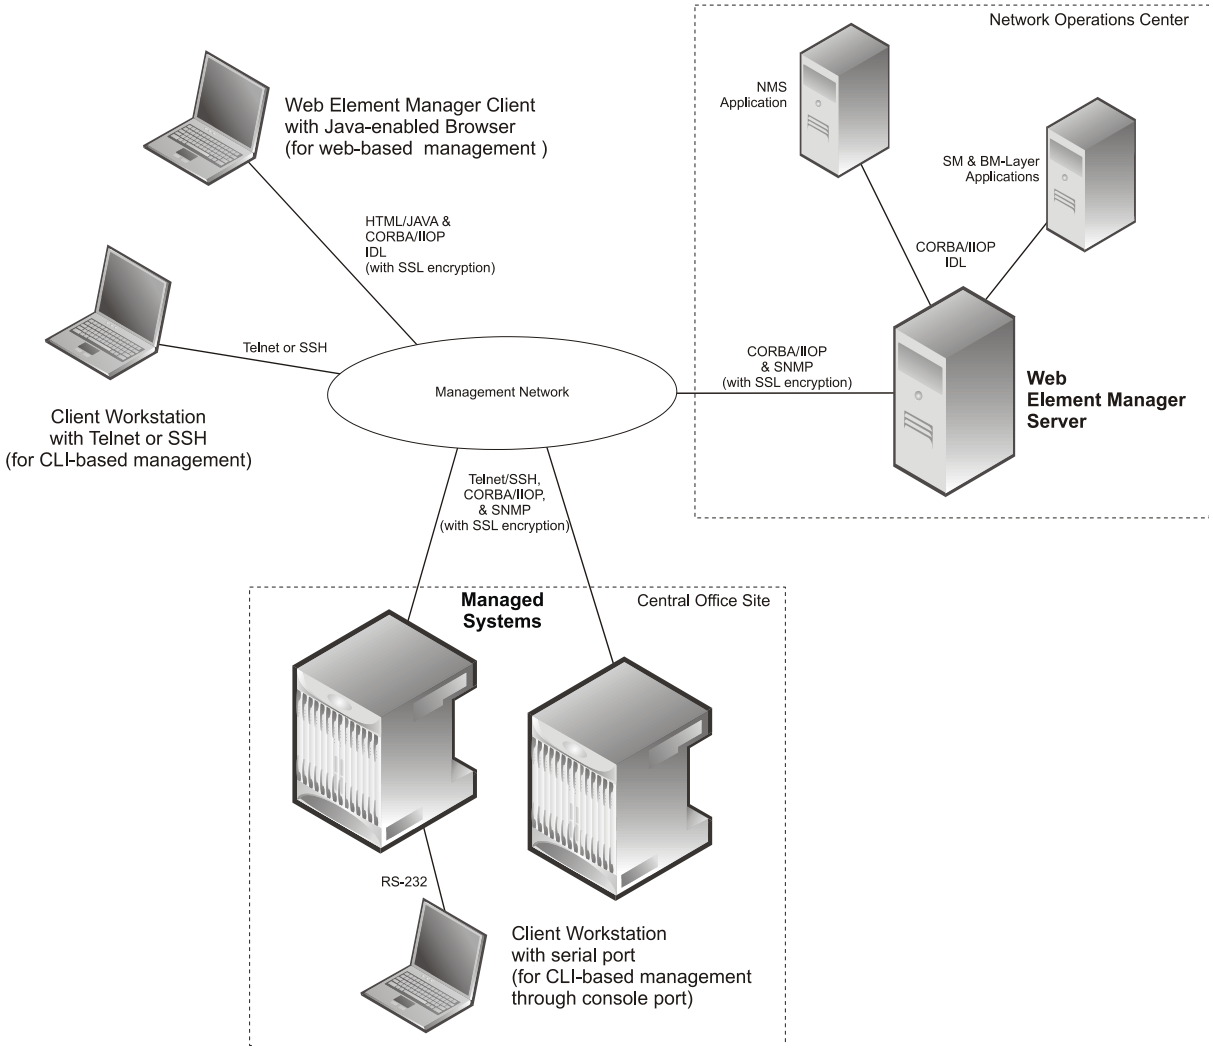
The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management -- focusing on providing superior quality Network Element (NE) and element management system (Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems -- giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

The Operation and Maintenance module of system offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces. These include:

- Using the Command Line Interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager application
- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000
- Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (i.e. Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network.

Figure 99. Element Management Methods



Important: GGSN management functionality is enabled by default for console-based access. For GUI-based management support, refer *Web Element Management System* section.

Important: For more information on command line interface based management, refer *GPRS/UMTS Command Line Interface Reference* and *GGSN Administration Guide*.

MPLS Forwarding with LDP

Multi Protocol Label Switching (MPLS) is an operating scheme or a mechanism that is used to speed up the flow of traffic on a network by making better use of available network paths. It works with the routing protocols like BGP and OSPF and therefore it is not a routing protocol.

It generates a fixed-length label to attach or bind with the IP packet's header to control the flow and destination of data. The binding of the labels to the IP packets is done by the label distribution protocol (LDP). All the packets in a

forwarding equivalence class (FEC) are forwarded by a label-switching router (LSR) which is also called an MPLS node. The LSR uses the LDP in order to signal its forwarding neighbors and distribute its labels for establishing a label switching path (LSP).

In order to support the increasing number of corporate APNs which have a number of different addressing models and requirements, MPLS is deployed to fulfill at least following two requirements:

- The corporate APN traffic must remain segregated from other APNs for security reasons.
- Overlapping of IP addresses in different APNs.

When deployed, MPLS backbone automatically negotiates the routes using the labels binded with the IP packets. Cisco GGSN as an LSR learns the default route from the connected provider edge (PE) while the PE populates its routing table with the routes provided by the GGSN.

Overlapping IP Address Pool Support

Overlapping IP Address Pools provides a mechanism for allowing operators to more flexibly support multiple corporate VPN customers with the same private IP address space without the expensive investments in physically separate routers, or expensive configurations using virtual routers.

The system supports two type of overlapping pools: resource and overlap. Resource pools are designed for dynamic assignment only, and use a VPN tunnel, such as a GRE tunnel, to forward and receive the private IP addresses to and from the VPN. Overlapping type pools can be used for both dynamic and static, and use VLANs and a next hop forwarding address to connect to the VPN customer.

To forward downstream traffic to the correct PDP context, the GGSN uses either the GRE tunnel ID, or the VLAN ID to match the packet. When forwarding traffic upstream, the GGSN uses the tunnel and forwarding information in the IP pool configuration, so overlapping pools must be configured in the APN for this feature to be used.

When a PDP context is created, the IP addresses is either assigned from the IP pool, in this case the forwarding rules are also configured into the GGSN at this point. If the address is assigned statically, when the GGSN confirms the IP address from the pool configured in the APN, the forwarding rules are also applied.

The GGSN can scale to as many actual overlapping pools as there are VLAN interfaces per context, and there can be multiple contexts per GGSN, or when using resource then the limit is the number of IP pools. This scalability allows operators, who wish to provide VPN services to customers using the customer's private IP address space, need not be concerned about escalating hardware costs, or complex configurations.

Important: For more information on IP pool overlapping configuration, refer *VLANs* in *System Administration Guide*.

PDP Context Support

Support for subscriber primary and secondary Packet Data Protocol (PDP) contexts in accordance with the following standards:

- **3GPP TS 23.060 v7.4.0 (2007-9)**: 3rd Generation Partnership project; Technical Specification Group Services and System Aspects; General Packet Radio Service (GPRS); Service description (Release 1999) as an additional reference for GPRS/UMTS procedures
- **3GPP TS 29.061 v7.6.0 (2008-09)**: 3rd Generation Partnership Project; Technical Specification Group Core Network; Packet Domain; Interworking between the Public Land Mobile Network (PLMN) supporting Packet Based Services and Packet Data Networks (PDN) (Release 4)

PDP context processing is based on the APN that the subscriber is attempting to access. Templates for all of the possible APNs that subscribers will be accessing must be configured within the system. Up to 1024 APNs can be configured on the system.

Each APN template consists of parameters pertaining to how PDP contexts are processed such as the following:

- Type (IPv4, IPv6, and/or PPP)
- Accounting protocol (GTPP or RADIUS)
- Authentication protocol (CHAP, MSCHAP, PAP, Allow-NOAUTH, IMSI-based, MSISDN-based)
- Charging characteristics (use SGSN-supplied or use configured)
- IP address allocation method (static or dynamic)
- PDP Context timers
- Quality of Service

A total of 11 PDP contexts are supported per subscriber. These could be all primaries, or 1 Primary and 10 secondaries or any combination of primary and secondary. Note that there must be at least one primary PDP context in order for secondaries to come up.

Per APN Configuration to Swap out Gn to Gi APN in CDRs

In order to allow for better correlation of CDRs with the network or application used by the subscriber, a configuration option has been added to the GGSN replace the Gn APN with the Gi (virtual) APN in emitted G-CDRs.

When virtual APNs are used, the operator can specify via EMS or a configuration command that the Gi APN should be used in the "Access Point Name Network Identifier" field of emitted G-CDRs, instead of the Gn APN.

Port Insensitive Rule for Enhanced Charging Service

This feature allows a single host or url rule to be applied to two different addresses, one with and one without the port number appended. As adding the port to the address is optional, this means that the number of rules could be halved.

Browser applications can sometimes appended the port number to the host or url when sending the host or URL fields. RFC 2616 for example states that port should be appended but if it is omitted then 80 should be assumed.

When configuring rules to define the content, as the web browser may provide the port number, even if it is the default one of 80 for HTTP, then two of each URL are needed.

Example

```
host = www.w3.org host = www.w3.org:80orhttp url =
http://213.229.187.118:80/chat/c/wel.w.wml http url =
http://213.229.187.118/chat/c/wel.w.wml
```

This feature provides a means to configure the rule such that the traffic is matched irrespective of the presence of a port number.

A new configurable has been added to the rulebase configuration that will ignore the port numbers embedded in the application headers of HTTP, RTSP, SIP, and WSP protocols.

When this feature is enabled, a single rule, such as "host = www.w3.org" would be matched even if the port number is appended and in this case the host field has the value www.w3.org:80, thereby cutting the number of rules needed by up to a half.

Important: For more information on enhanced charging service, refer *Enhanced Charging Service Administration Guide*.

Quality of Service Support

Provides operator control over the prioritization of different types of traffic.

Quality of Service (QoS) support provides internal processing prioritization based on needs, and DiffServ remarking to allow external devices to perform prioritization.

Important: The feature described here is internal prioritization and DiffServ remarking for external prioritization. For additional QoS capabilities of the GGSN, refer Features and Functionality - Optional Enhanced Feature Software section.

External prioritization (i.e., the value to use for the DiffServ marking) is configured for the uplink and downlink directions. In the uplink direction, each APN is configurable for the DiffServ ToS value to use for each of the 3GPP traffic classes. Alternatively, you can configure "pass-through", whereby the ToS value will pass through unchanged.

In the downlink direction, the ToS value of the subscriber packet is not changed, but you can configure what to use for the ToS value of the outer GTP tunnel. The value for ToS is configurable for each of the 3GPP traffic classes. In addition, the connections between the GGSN and one or more SGSNs can be configured as a "GGSN Service", and different values for ToS for the same 3GPP traffic class may be configured for different GGSN Services.

RADIUS Support

Provides a mechanism for performing authorization, authentication, and accounting (AAA) for subscriber PDP contexts based on the following standards:

- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000
- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000

The Remote Authentication Dial-In User Service (RADIUS) protocol is used to provide AAA functionality for subscriber PDP contexts. (RADIUS accounting is optional since GTPP can also be used.)

Within context contexts configured on the system, there are AAA and RADIUS protocol-specific parameters that can be configured. The RADIUS protocol-specific parameters are further differentiated between RADIUS Authentication server RADIUS Accounting server interaction.

Among the RADIUS parameters that can be configured are:

- **Priority**: Dictates the order in which the servers are used allowing for multiple servers to be configured in a single context.
- **Routing Algorithm**: Dictate the method for selecting among configured servers. The specified algorithm dictates how the system distributes AAA messages across the configured AAA servers for new sessions. Once

a session is established and an AAA server has been selected, all subsequent AAA messages for the session will be delivered to the same server.

In the event that a single server becomes unreachable, the system attempts to communicate with the other servers that are configured. The system also provides configurable parameters that specify how it should behave should all of the RADIUS AAA servers become unreachable.

The system provides an additional level of flexibility by supporting the configuration RADIUS server groups. This functionality allows operators to differentiate AAA services for subscribers based on the APN used to facilitate their PDP context.

In general, 128 AAA Server IP address/port per context can be configured on the system and it selects servers from this list depending on the server selection algorithm (round robin, first server). Instead of having a single list of servers per context, this feature provides the ability to configure multiple server groups. Each server group, in turn, consists of a list of servers.

This feature works in following way:

- All RADIUS authentication/accounting servers configured at the context-level are treated as part of a server group named "default". This default server group is available to all subscribers in that context through the realm (domain) without any configuration.
- It provides a facility to create "user defined" RADIUS server groups, as many as 399 (excluding "default" server group), within a context. Any of the user defined RADIUS server groups are available for assignment to a subscriber through the APN configuration within that context.

Since the configuration of the APN can specify the RADIUS server group to use as well as IP address pools from which to assign addresses, the system implements a mechanism to support some in-band RADIUS server implementations (i.e. RADIUS servers which are located in the corporate network, and not in the operator's network) where the NAS-IP address is part of the subscriber pool. In these scenarios, the GGSN supports the configuration of the first IP address of the subscriber pool for use as the RADIUS NAS-IP address.

Important: In 12.3 and earlier releases, refer to the *AAA and GTPP Interface Administration and Reference* for more information on RADIUS AAA configuration.

RADIUS VLAN Support

VPN customers often use private address space which can easily overlap with other customers. The subscriber addresses are supported with overlapping pools which can be configured in the same virtual routing context.

This feature now allows Radius Server and NAS IP addresses to also overlapping without the need to configure separate contexts, thereby simplifying APN and RADIUS configuration and network design.

This feature now allows Radius Server and NAS IP addresses to also overlapping without the need to configure separate contexts, thereby simplifying APN and RADIUS configuration and network design.

This feature supports following scenarios to be defined in the same context:

- Overlapping RADIUS NAS-IP address for various RADIUS server groups representing different APNs.
- Overlapping RADIUS server IP address for various RADIUS servers groups.

Previously, the above scenarios were supported, albeit only when the overlapping addresses were configured in different contexts. Moreover a static route was required in each context for IP connectivity to the RADIUS server.

The new feature utilizes the same concept as overlapping IP pools such that every overlapping NAS-IP address is giving a unique next-hop address which is then bound to an interface that is bound to a unique VLAN, thereby allowing the configuration to exist within the same context.

RADIUS access requests and accounting messages are forwarded to the next hop defined for that NAS-IP and it is then up to the connected router's forward the messages to the RADIUS server. The next hop address determines the interface and VLAN to use. Traffic from the server is identified as belonging to a certain NAS-IP by the port/VLAN combination.

The number of Radius NAS-IP addresses that can be configured is limited by the number of loopback addresses that can be configured.

Important: For more information on VLAN support, refer *VLANs* in *System Administration Guide*.

Routing Protocol Support

The system's support for various routing protocols and routing mechanism provides an efficient mechanism for ensuring the delivery of subscriber data packets.

GGSN node supports Routing Protocol in different way to provide an efficient mechanism for delivery of subscriber data.

The following routing mechanisms and protocols are supported by the system:

- **Static Routes**: The system supports the configuration of static network routes on a per context basis. Network routes are defined by specifying an IP address and mask for the route, the name of the interface in the currant context that the route must use, and a next hop IP address.
- **Open Shortest Path First (OSPF) Protocol**: A link-state routing protocol, OSPF is an Interior Gateway Protocol (IGP) that routes IP packets based solely on the destination IP address found in the IP packet header using the shortest path first. IP packets are routed "as is", meaning they are not encapsulated in any further protocol headers as they transit the network.

  Variable length subnetting, areas, and redistribution into and out of OSPF are supported.

  OSPF routing is supported in accordance with the following standards:

  - RFC-1850, OSPF Version 2 Management Information Base, November 1995
  - RFC-2328, OSPF Version 2, April 1998
  - RFC-3101 OSPF-NSSA Option, January 2003
- **Border Gateway Protocol version 4 (BGP-4)**: The system supports a subset of BGP (RFC-1771, A Border Gateway Protocol 4 (BGP-4)), suitable for eBGP support of multi-homing typically used to support geographically redundant mobile gateways, is supported.

  EBGP is supported with multi-hop, route filtering, redistribution, and route maps. The network command is support for manual route advertisement or redistribution.

  BGP route policy and path selection is supported by the following means:

  - Prefix match based on route access list
  - AS path access-list
  - Modification of AS path through path prepend
  - Origin type
  - MED
  - Weight

- **Route Policy**: Routing policies modify and redirect routes to and from the system to satisfy specific routing needs. The following methods are used with or without active routing protocols (i.e. static or dynamic routing) to prescribe routing policy:
  - **Route Access Lists**: The basic building block of a routing policy, route access lists filter routes based upon a specified range of IP addresses.
  - **IP Prefix Lists**: A more advanced element of a routing policy. An IP Prefix list filters routes based upon IP prefixes
  - **AS Path Access Lists**: A basic building block used for Border Gateway Protocol (BGP) routing, these lists filter Autonomous System (AS) paths.
- **Route Maps**: Route-maps are used for detailed control over the manipulation of routes during route selection or route advertisement by a routing protocol and in route redistribution between routing protocols. This detailed control is achieved using IP Prefix Lists, Route Access Lists and AS Path Access Lists to specify IP addresses, address ranges, and Autonomous System Paths.
- **Equal Cost Multiple Path (ECMP)**: ECMP allows distribution of traffic across multiple routes that have the same cost to the destination. In this manner, throughput load is distributed across multiple path, typically to lessen the burden on any one route and provide redundancy. The mobile gateway supports from four to ten equal-cost paths.

Important: For more information on IP Routing configuration, refer *Routing* in *System Administration Guide*.

Subscriber Session Trace Support

The Subscriber Level Trace provides a 3GPP standards-based session-level trace function for call debugging and testing new functions and access terminals in an UMTS environment.

In general, the Session Trace capability records and forwards all control activity for the monitored subscriber on the monitored interfaces. This is typically all the signaling and authentication/subscriber services messages that flow when a UE connects to the access network.

The UMTS network entities like SGSN and GGSN support 3GPP standards based session level trace capabilities to monitor all call control events on the respective monitored interfaces including **Gn**, **Gi**, **Gx**, and **Gmb** interface on GGSN. The trace can be initiated using multiple methods:

- Management initiation via direct CLI configuration
- Management initiation at AAA with trace activation via authentication response messages over **Gx** reference interface
- Signaling based activation through signaling from subscriber access terminal

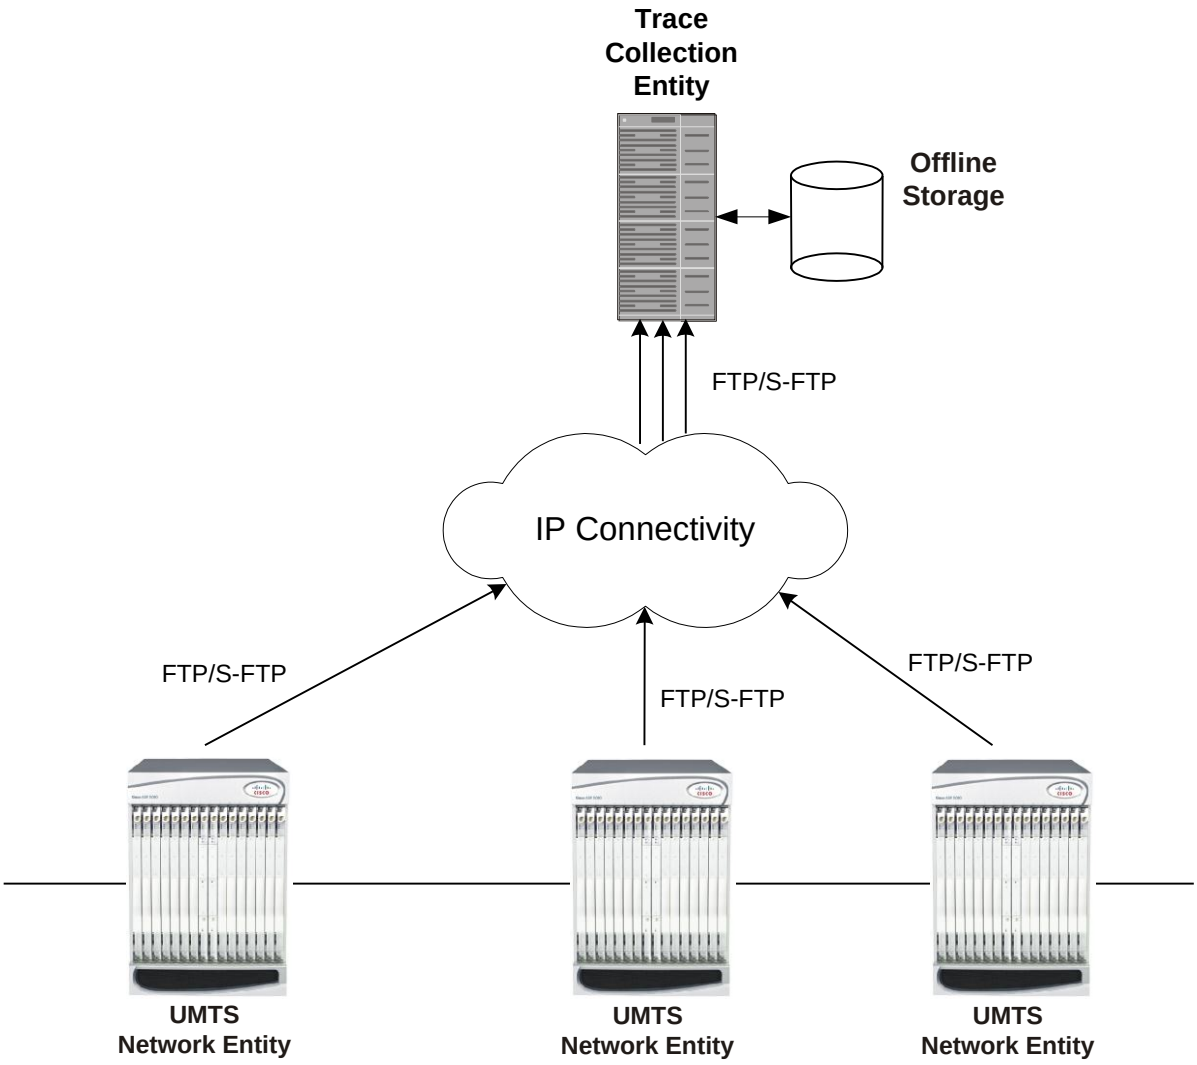
Important: Once the trace is provisioned it can be provisioned through the access cloud via various signaling interfaces.

The session level trace function consists of trace activation followed by triggers. The time between the two events is treated much like Lawful Intercept where the UMTS network element buffers the trace activation instructions for the provisioned subscriber in memory using camp-on monitoring. Trace files for active calls are buffered as XML files using non-volatile memory on the local dual redundant hard drives on the system. The Trace Depth defines the granularity of data to be traced. Six levels are defined including Maximum, Minimum and Medium with ability to configure additional levels based on vendor extensions.

Important: Only Maximum Trace Depth is supported in the current release.

The following figure shows a high-level overview of the session-trace functionality and deployment scenario:

Figure 100. Session Trace Function and Interfaces



All call control activity for active and recorded sessions is sent to an off-line Trace Collection Entity (TCE) using a standards-based XML format over a FTP or secure FTP (SFTP) connection.

Note: In the current release the IPv4 interfaces are used to provide connectivity to the TCE. Trace activation is based on IMSI or IMEI.

Support of Charging Characteristics Provided by AAA Server

This feature provides the ability for operators to apply Charging Characteristics (CC) from the AAA server instead of a hard coded local profile during access authentication.

The RADIUS attribute **3GPP-Chrg-Char** can be used to get the charging characteristics from RADIUS in Access-Accept message. Accepting the RADIUS returned charging characteristic profile must be enabled per APN. The CC profile returned by AAA will override any CC provided by the SGSN, the GGSN or per APN configuration. All 16 profile behaviors can be defined explicitly or the default configuration for that profile is used.

Support of all GGSN generated causes for partial G-CDR closure

Provides more detailed eG-CDR and/or G-CDR closure causes as per 3GPP TS 32.298.

System handles the GGSN generated causes for partial closure of CDRs. It supports various type of causes including Radio Access Technology Change, MS Time Zone Change, Cell update, inter-PLMN SGSN change, PLMN id change, QoS, Routing-Area update etc.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert**: A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm**: Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps**: SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored value.

  Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.
- **Logs**: The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING

  Logs are supported in both the Alert and the Alarm models.
- **Alarm System**: High threshold alarms generated within the specified polling interval are considered "outstanding" until a the condition no longer exists or a condition clear alarm is generated. "Outstanding"

alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer *Thresholding Configuration Guide*.

Features and Functionality - Optional Enhanced Feature Software

This section describes the optional enhanced features and functions for GGSN service.

Each of the following features require the purchase of an additional license to implement the functionality with the GGSN service.

This section describes following features:

- Common Gateway Access Support
- Dynamic RADIUS Extensions (Change of Authorization)
- GRE Protocol Interface Support
- Gx Interface Support
- Inter-Chassis Session Recovery
- IP Security (IPSec)
- L2TP LAC Support
- L2TP LNS Support
- Lawful Intercept
- Mobile IP Home and Foreign Agents
- Mobile IP NAT Traversal
- Multimedia Broadcast Multicast Services Support
- Overcharging Protection on Loss of Coverage
- Proxy Mobile IP
- Session Persistence
- Session Recovery Support
- Traffic Policing and Rate Limiting
- Web Element Management System

Common Gateway Access Support

Common Gateway Access support is a consolidated solution that combines 3G and 4G access technologies in a common gateway supporting logical services of HA, PGW, and GGSN to allow users to have the same user experience, independent of the access technology available.

In todays scenario an operator must have multiple access networks (CDMA, eHRPD and LTE) plus a GSM/UMTS solution for international roaming. Therefore, operator requires a solution to allow customers to access services with the same IP addressing behavior and to use a common set of egress interfaces, regardless of the access technology (3G or 4G).

This solution allows static customers to access their network services with the same IP addressing space assigned for wireless data, regardless of the type of connection (CDMA, eHRPD/LTE or GSM/UMTS). Subscribers using static IP addressing will be able to get the same IP address regardless of the access technology.

For more information on this product, refer *Common Gateway Access Support* section in GGSN Service Administration Guide.

Dynamic RADIUS Extensions (Change of Authorization)

Dynamic RADIUS extension support provide operators with greater control over subscriber PDP contexts by providing the ability to dynamically redirect data traffic, and or disconnect the PDP context.

This functionality is based on the RFC 3576, Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS), July 2003 standard.

The system supports the configuration and use of the following dynamic RADIUS extensions:

- **Change of Authorization**: The system supports CoA messages from the AAA server to change data filters associated with a subscriber session. The CoA request message from the AAA server must contain attributes to identify NAS and the subscriber session and a data filter ID for the data filter to apply to the subscriber session.
- **Disconnect Message**: The DM message is used to disconnect subscriber sessions in the system from a RADIUS server. The DM request message should contain necessary attributes to identify the subscriber session.

The above extensions can be used to dynamically re-direct subscriber PDP contexts to an alternate address for performing functions such as provisioning and/or account set up. This functionality is referred to as Session Redirection, or Hotlining.

Session redirection provides a means to redirect subscriber traffic to an external server by applying ACL rules to the traffic of an existing or a new subscriber session. The destination address and optionally the destination port of TCP/IP or UDP/IP packets from the subscriber are rewritten so the packet is forwarded to the designated redirected address.

Return traffic to the subscriber has the source address and port rewritten to the original values. The redirect ACL may be applied dynamically by means of the Radius Change of Authorization (CoA) extension.

Important: For more information on dynamic RADIUS extensions support, refer *CoA, RADIUS, And Session Redirection (Hotlining)* in this guide.

GRE Protocol Interface Support

GGSN supports GRE generic tunnel interface support in accordance with RFC-2784, Generic Routing Encapsulation (GRE).

GRE protocol functionality adds one additional protocol on the system to support mobile users to connect to their enterprise networks through Generic Routing Encapsulation (GRE).

GRE tunnels can be used by the enterprise customers of a carrier 1) To transport AAA packets corresponding to an APN over a GRE tunnel to the corporate AAA servers and, 2) To transport the enterprise subscriber packets over the GRE tunnel to the corporation gateway.

The corporate servers may have private IP addresses and hence the addresses belonging to different enterprises may be overlapping. Each enterprise needs to be in a unique virtual routing domain, known as VRF. To differentiate the tunnels between same set of local and remote ends, GRE Key will be used as a differentiation.

GRE Tunneling is a common technique to enable multi-protocol local networks over a single-protocol backbone, to connect non-contiguous networks and allow virtual private networks across WANs. This mechanism encapsulates data packets from one protocol inside a different protocol and transports the data packets unchanged across a foreign network. It is important to note that GRE tunneling does not provide security to the encapsulated protocol, as there is no encryption involved (like IPSEC offers, for example).

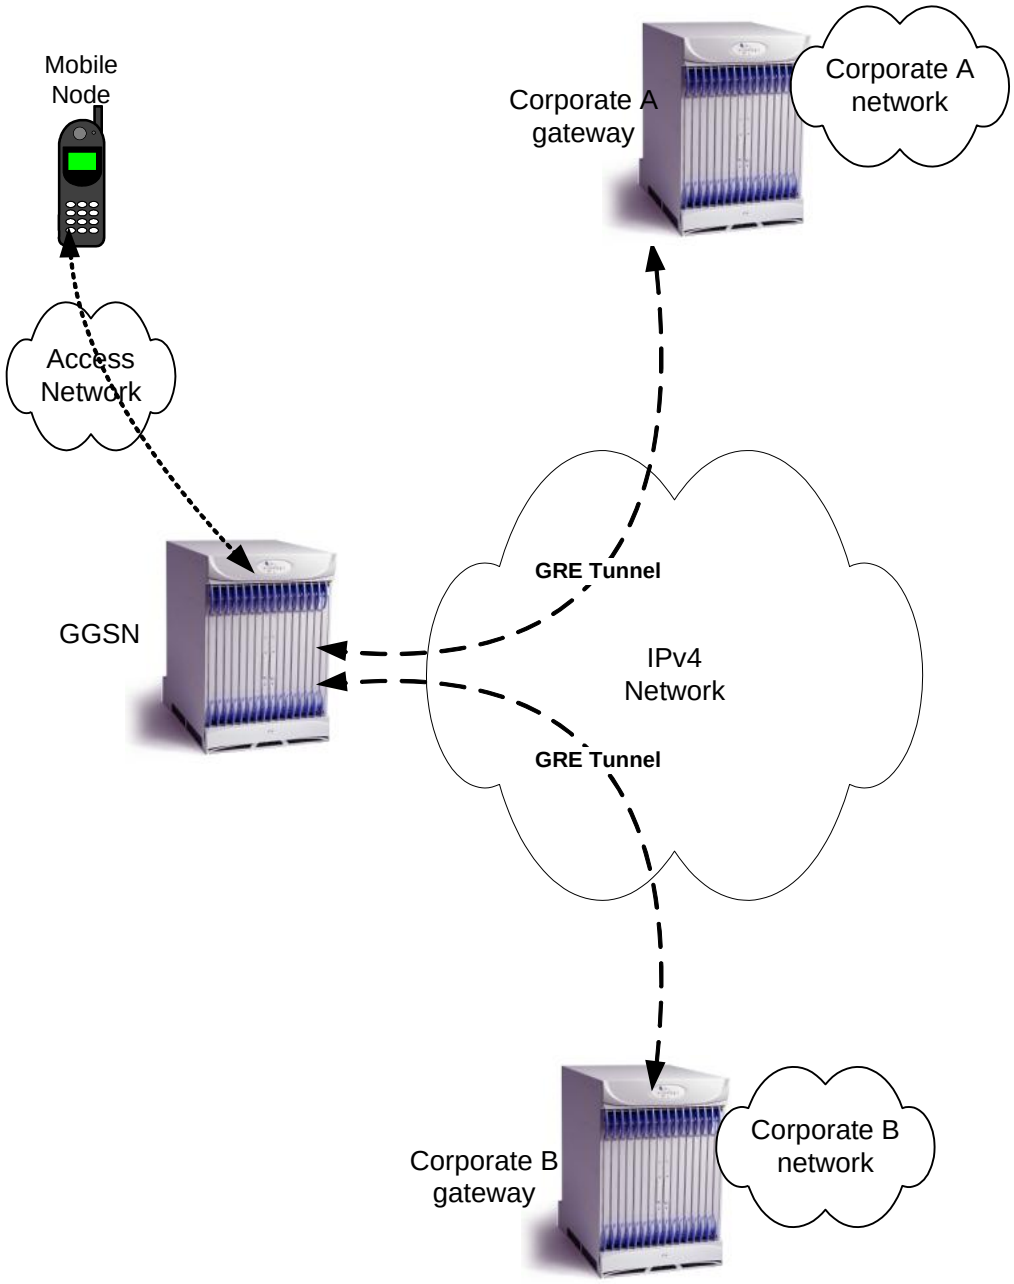
GRE Tunneling consists of three main components:

- Passenger protocol-protocol being encapsulated. For example: CLNS, IPv4 and IPv6.
- Carrier protocol-protocol that does the encapsulating. For example: GRE, IP-in-IP, L2TP, MPLS and IPSEC.
- Transport protocol-protocol used to carry the encapsulated protocol. The main transport protocol is IP.

The most simplified form of the deployment scenario is shown in the following figure, in which GGSN has two APNs talking to two corporate networks over GRE tunnels.

The following figure shows a high-level overview of the GRE deployment scenario:

Figure 101. GRE Deployment Scenario



Gx Interface Support

Gx interface support on the system enables the wireless operator to:

- Implement differentiated service profiles for different subscribers
- Intelligently charge the services accessed depending on the service type and parameters

This interface is particularly suited to control and charge multimedia applications and IMS services. This interface support is compliant to following standards:

- 3GPP TS 23.203 V7.6.0 (2008-03): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Policy and charging control architecture (Release 7)
- 3GPP TS 29.210 V6.2.0 (2005-06): 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Charging rule provisioning over Gx interface; (Release 6)
- 3GPP TS 29.212 V7.4.0 (2008-03): 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Policy and Charging Control over Gx reference point (Release 7)
- 3GPP TS 29.213 V7.4.0 (2008-03): 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Policy and Charging Control signalling flows and QoS parameter mapping; (Release 7)
- RFC 3588, Diameter Base Protocol
- RFC 4006, Diameter Credit-Control Application

In addition to the above RFCs and standards IMS authorization partially supports 3GPP TS 29.212 for Policy and Charging Control over Gx reference point functionality.

The goal of the Gx interface is to provide network based QoS control as well as dynamic charging rules on a per bearer basis. The Gx interface is in particular needed to control and charge multimedia applications.

QoS Parameter ARP Setting via Gx Interface: GGSN controls the assignment of different radio interface QoS priorities (gold/silver/bronze) via the PCRF Gx interface during PDP context setup (CCR/CCA-I). This is performed using the Allocation Retention Priority (ARP) parameter (AVP code 1034) as specified in 3GPP TS 29.212, with values = 0-3; ARP values from the PCRF other than 0-3 are ignored. During PDP context setup the PCRF returns the ARP value in CCA-I and this ARP is then assigned/negotiated with the SGSN and RNC.

The Gx interface is located between the GGSN and the E-PDF / PCRF. It is a Diameter- based interface and provides the functions provided earlier by the Gx and Go interfaces:

- QoS control based on either a token-based or token-less mechanism. In the token-based mechanism, the E-PDF or PCRF dynamically assign network resources to the different bearers used by the subscriber. These resource assignments are transmitted in Tokens carried over the Gx interface. The authorization tokens are allocated by the network (E-PDF/PCRF), hence the network is in full control of the mechanism since it only authorizes resources. The token-less mechanism is for further study.
- Dynamic rules for Flexible Bearer Charging. These dynamic charging rules are carried in the resource assignment tokens and provide 5-tuple type charging rules that enables to implement a specific charging policy for each subscriber bearer. These charging rules will be applied by the FBC function of the GGSN, and produce the appropriate eG-CDRs or the appropriate messages on the Gy interface to the OCS.

Important: For more information on Gx interface support, refer *Gx Interface Support* in this guide.

Inter-Chassis Session Recovery

The chassis provides industry leading carrier class redundancy. The systems protects against all single points of failure (hardware and software) and attempts to recover to an operational state when multiple simultaneous failures occur.

The system provides several levels of system redundancy:

- Under normal N+1 packet processing card hardware redundancy, if a catastrophic packet processing card failure occurs all affected calls are migrated to the standby packet processing card if possible. Calls which cannot be

migrated are gracefully terminated with proper call-termination signaling and accounting records are generated with statistics accurate to the last internal checkpoint

- If the Session Recovery feature is enabled, any total packet processing card failure will cause a packet processing card switchover and all established sessions for supported call-types are recovered without any loss of session.

Even though chassis provides excellent intra-chassis redundancy with these two schemes, certain catastrophic failures which can cause total chassis outages, such as IP routing failures, line-cuts, loss of power, or physical destruction of the chassis, cannot be protected by this scheme. In such cases, the GGSN Inter-Chassis Session Recovery feature provides geographic redundancy between sites. This has the benefit of not only providing enhanced subscriber experience even during catastrophic outages, but can also protect other systems such as the RAN from subscriber re-activation storms.

The Inter-chassis Session Recovery feature allows for continuous call processing without interrupting subscriber services. This is accomplished through the use of redundant chassis. The chassis are configured as primary and backup with one being active and one in recovery mode. A checkpoint duration timer is used to control when subscriber data is sent from the active chassis to the inactive chassis. If the active chassis handling the call traffic goes out of service, the inactive chassis transitions to the active state and continues processing the call traffic without interrupting the subscriber session. The chassis determines which is active through a propriety TCP-based connection called a redundancy link. This link is used to exchange **Hello** messages between the primary and backup chassis and must be maintained for proper system operation.

Interchassis Session Recovery uses following for failur handling and communication:

- **Interchassis Communication**:

  Chassis configured to support Interchassis Session Recovery communicate using periodic **Hello** messages. These messages are sent by each chassis to notify the peer of its current state. The **Hello** message contains information about the chassis such as its configuration and priority. A dead interval is used to set a time limit for a **Hello** message to be received from the chassis' peer. If the standby chassis does not receive a **Hello** message from the active chassis within the dead interval, the standby chassis transitions to the active state. In situations where the redundancy link goes out of service, a priority scheme is used to determine which chassis processes the session. The following priority scheme is used:

  - router identifier
  - chassis priority
  - SPIO MAC address

- **Checkpoint Message**:

  Checkpoint messages are sent from the active chassis to the inactive chassis. Checkpoint messages are sent at specific intervals and contain all the information needed to recreate the sessions on the standby chassis, if that chassis were to become active. Once a session exceeds the checkpoint duration, checkpoint data is collected on the session. The checkpoint parameter determines the amount of time a session must be active before it is included in the checkpoint message.

Important: For more information on inter-chassis session recovery support, refer *Interchassis Session Recovery* in *System Administration Guide*.

IP Security (IPSec)

IP Security provides a mechanism for establishing secure tunnels from mobile subscribers to pre-defined endpoints (i.e. enterprise or home networks) in accordance with the following standards:

- RFC 2401, Security Architecture for the Internet Protocol

- RFC 2402, IP Authentication Header (AH)
- RFC 2406, IP Encapsulating Security Payload (ESP)
- RFC 2409, The Internet Key Exchange (IKE)
- RFC-3193, Securing L2TP using IPSEC, November 2001

IP Security (IPSec) is a suite of protocols that interact with one another to provide secure private communications across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways.

IPSec tunnel supports AAA and DHCP address overlapping. Address overlapping is meant for multiple customers using the same IP address for AAA/DHCP servers. The AAA and DHCP control messages are sent over IPSec tunnels and AAA/DHCP packets required to be encrypted are decided as per the ACL configuration done for specific session.

IPSec can be implemented on the system for the following applications:

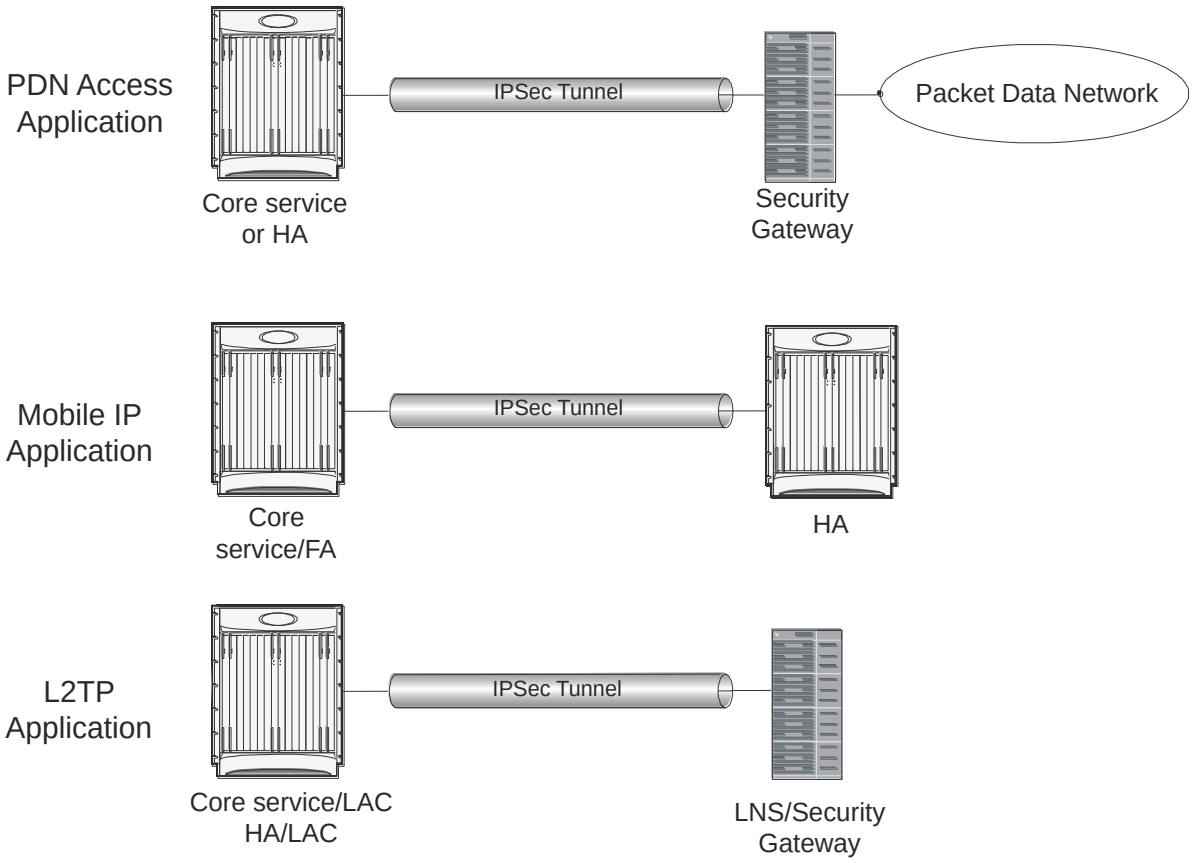
- **PDN Access**: Subscriber IP traffic is routed over an IPSec tunnel from the system to a secure gateway on the Packet Data Network (PDN) as determined by Access Control List (ACL) criteria.
- **Mobile IP**: Mobile IP control signals and subscriber data is encapsulated in IPSec tunnels that are established between Foreign Agents (FAs) and Home Agents (HAs) over the Pi interfaces.

  Important: Once an IPSec tunnel is established between an FA and HA for a particular subscriber, all new Mobile IP sessions using the same FA and HA are passed over the tunnel regardless of whether or not IPSec is supported for the new subscriber sessions. Data for existing Mobile IP sessions will be unaffected.

- **L2TP**: L2TP-encapsulated packets are routed from the system to an LNS/secure gateway over an IPSec tunnel.

The following figure shows a high-level overview of the IPSec application deployment scenario:

Figure 102. IPSec Application Deployment



Important: For more information on IPSec support, refer *IP Security* in this guide.

L2TP LAC Support

The system configured as a Layer 2 Tunneling Protocol Access Concentrator (LAC) enables communication with L2TP Network Servers (LNSs) for the establishment of secure Virtual Private Network (VPN) tunnels between the operator and a subscriber's corporate or home network.

The use of L2TP in VPN networks is often used as it allows the corporation to have more control over authentication and IP address assignment. An operator may do a first level of authentication, however use PPP to exchange user name and password, and use IPCP to request an address. To support PPP negotiation between the GGSN and the corporation, an L2TP tunnel must be setup in the GGSN running a LAC service.

L2TP establishes L2TP control tunnels between LAC and LNS before tunneling the subscriber PPP connections as L2TP sessions. The LAC service is based on the same architecture as the GGSN and benefits from dynamic resource allocation and distributed message and data processing. This design allows the LAC service to support over 4000 setups per second or a maximum of over 3G of throughput. There can be a maximum up to 65535 sessions in a single tunnel and as many as 500,000 L2TP sessions using 32,000 tunnels per system.

The LAC sessions can also be configured to be redundant, thereby mitigating any impact of hardware of software issues. Tunnel state is preserved by copying the information across processor cards.

Important: For more information on this feature support, refer *L2TP Access Concentrator* in this guide.

L2TP LNS Support

The system configured as a Layer 2 Tunneling Protocol Network Server (LNS) supports the termination secure Virtual Private Network (VPN) tunnels between from L2TP Access Concentrators (LACs).

The LNS service takes advantage of the high performance PPP processing already supported in the system design and is a natural evolution from the LAC. The LNS can be used as a standalone, or running alongside a GGSN service in the same platform, terminating L2TP services in a cost effective and seamless manner.

L2TP establishes L2TP control tunnels between LAC and LNS before tunneling the subscriber PPP connections as L2TP sessions. There can be a maximum of up to 65535 sessions in a single tunnel and up to 500,000 sessions per LNS.

The LNS architecture is similar to the GGSN and utilizes the concept of a de-multiplexer to intelligently assign new L2TP sessions across the available software and hardware resources on the platform without operator intervention.

Important: For more information on this feature support, refer *L2TP Network Server* in this guide.

Lawful Intercept

The system supports the Lawful Interception (LI) of subscriber session information. This functionality provides Telecommunication Service Providers (TSPs) with a mechanism to assist Law Enforcement Agencies (LEAs) in the monitoring of suspicious individuals (referred to as targets) for potential criminal activity.

The following standards were referenced for the system's LI implementation:

- TR-45 Lawfully Authorized Electronic Surveillance TIA/EIA J-STD-025 PN4465 RV 1.7
- 3GPP TS 33.106 V6.1.0 (2004-06): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G security; Lawful Interception requirements (Release 6)
- 3GPP TS 33.107 V6.2.0 (2004-06): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G security; Lawful interception architecture and functions (Release 6)
- 3GPP TS 33.108 V9.0.0 (2009-09): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G security; Handover interface for Lawful Interception (LI) (Release 9)
- Technical Directive: Requirements for implementing statutory telecommunications interception measures (TR TKÜ), Version 4.0

LEAs provide one or more TSPs with court orders or warrants requesting the monitoring of a particular target. The target is identified by information such as their mobile station Integrated Services Digital Network (MSISDN) number, or their International Mobile Subscriber Identification (IMSI) number.

Once the target has been identified, the system, functioning as either a GGSN or HA, serves as an Access Function (AF) and performs monitoring for both new PDP contexts or PDP contexts that are already in progress. While monitoring, the system intercepts and duplicates Content of Communication (CC) and/or Intercept Related Information (IRI) and forwards it to a Delivery Function (DF) over an extensible, proprietary interface. Note that when a target establishes multiple, simultaneous PDP contexts, the system intercepts CC and IRI for each of them. The DF, in turn, delivers the intercepted content to one or more Collection Functions (CFs).

Lawful intercept supports TCP transport on node interfaces along with support for IPv6 address link between chassis and LI server.

On the system with StarOS version 9.0 or later, this feature enhanced to allow 20,000 LI targets to be provisioned as well as monitored.

Caution: This capacity improvement impacts performance over various network scenario and in order to reach the full target of 20000 LI targets, it is required that the used platform have at least 12 active packet processing cards installed.

Important: For more information on this feature support, refer *Lawful Intercept Configuration Guide*.

Mobile IP Home and Foreign Agents

Consolidation of GGSN, HA and/or FA services on the same platform eliminates CapEx and OpEx requirements for separate network elements and devices under management. Service integration also enables seamless mobility and inter-technology roaming between 1xEV-DO and UMTS/W-CDMA/GPRS/EDGE radio access networks. This shared configuration also enables common address pools to be applied across all service types. In addition, this combination of collapsed services does not create dependencies for Mobile IP client software on the user access device and consequently does not introduce additional requirements for Mobile IP signaling in the 3GPP radio access network.

This functionality provides the following benefits:

- Timely release of Mobile IP resources at the FA and/or HA
- Accurate accounting
- Timely notification to mobile node of change in service

The system is capable of supporting both GGSN and Mobile IP functions on a single chassis. For Mobile IP applications, the system can be configured to provide the function of a Gateway GPRS Support Node/Foreign Agent (GGSN/FA) and/or a Home Agent (HA).

HA and FA components are defined by RFC 2002 in support of Mobile IP. Mobile IP provides a network-layer solution that allows Mobile Nodes (MNs, i.e. mobile phones, wireless PDAs, and other mobile devices) to receive routed IP packets from their home network while they are connected to any visitor network using their permanent or home IP address. Mobile IP allows mobility in a dynamic method that allows nodes to maintain ongoing communications while changing links as the user traverses the global Internet from various locations outside their home network.

When configured to support HA functionality, the system is capable of supporting following enhanced features:

- **Mobile IP HA Session Rejection/Redirection**: Enables the HA service to either reject new calls or redirect them to another HA when a destination network connection failure is detected. When network connectivity is re-established, the HA service begins to accept calls again in the normal manner. This feature provides the benefit of reducing OpEx through increased operational efficiency and limiting of system downtime.
- **Mobile IP Registration Revocation**: Registration Revocation is a general mechanism whereby the HA providing Mobile IP or Proxy Mobile IP functionality to a mobile node can notify the GGSN/FA of the termination of a binding. Mobile IP Registration Revocation can be triggered at the HA by any of the following:
  - Administrative clearing of calls
  - Session Manager software task outage resulting in the loss of FA sessions (sessions that could not be recovered)
  - Session Idle timer expiry (when configured to send Revocation)
  - Any other condition under which a binding is terminated due to local policy (duplicate IMSI detected, duplicate home address requested)

Important: For more information on Mobile IP HA service and FA service configuration, refer *HA Administration Guide* and *GGSN Administration Guide* respectively

Mobile IP NAT Traversal

This functionality enables converged WiFi-cellular data deployments in which the system is used to concentrate and switch traffic between WiFi hotspots. UDP/IP tunneling enables NAT firewalls in WLAN hotspots to maintain state information for address translation between NATed public address/UDP ports and addresses that are privately assigned for the mobile access device by a local DHCP server.

The Mobile IP protocol does not easily accommodate subscriber mobile nodes that are located behind WLAN or WAN-based NAT devices because it assumes that the addresses of mobile nodes or FA's are globally routable prefixes. However, the mobile node's co-located care of address (CCoA/CoA) is a private address. This presents a problem when remote hosts try to reach the mobile node via the public advertised addresses. The system provides a solution that utilizes UDP tunneling subject to subscriber reservation requests. In this application, the HA uses IP UDP tunneling to reach the mobile subscriber and includes the same private address that was provided in original reservation request in the encapsulated IP payload packet header.

Important: For more information on this feature, refer *MIP NAT Traversal* in *System Administration Guide*.

Multimedia Broadcast Multicast Services Support

Multimedia services are taking on an ever-increasing role in the wireless carriers' plans for an application centric service model. As such, any next generation GGSN platform must be capable of supporting the requirements of multimedia service delivery, including:

- Higher bandwidth requirements of streaming audio and video delivery
- Efficient broadcast and multicast mechanisms, to conserve resources in the RAN

MBMS represents the evolutionary approach to multicast and broadcast service delivery. MBMS uses spectrum resources much more efficiently than Multicast-over-Unicast by optimizing packet replication across all critical components in the bearer path. Thus, services requiring largely uni-directional multicast flows towards the UE are particularly well suited to the MBMS approach. These would include news, event streaming, suitably encoded/compressed cable/radio programs, video-on-demand, multi-chat / group-push-to-talk/video-conferencing sessions with unicast uplink and multicast downlink connections, and other applications.

For MBMS functionality, the system supports the Gmb interface, which is used signal to the BM-SC

Important: For more information on this feature, refer *Multicast Broadcast Service* in this guide.

Overcharging Protection on Loss of Coverage

This solution provides the ability to configure mobile carriers to maximize their network solutions and balancing the requirements to accurately bill their customer.

Considering a scenario where a mobile is streaming or downloading very large files from external sources and the mobile goes out of radio coverage. If this download is happening on Background/Interactive traffic class then the GGSN

is unaware of such loss of connectivity as SGSN does not perform the Update PDP Context procedure to set QoS to 0kbps (this is done when traffic class is either Streaming or Conversational only). The GGSN continues to forward the downlink packets to SGSN. In the loss of radio coverage, the SGSN will do paging request and find out that the mobile is not responding; SGSN will then drops the packets. In such cases, the G-CDR will have increased counts but S-CDR will not. This means that when operators charge the subscribers based on G-CDR the subscribers may be overcharged. This feature is implemented to avoid the overcharging in such cases.

This implementation is based on Cisco-specific private extension to GTP messages and/or any co-relation of G-CDRs and S-CDRs. It also does not modify any RANAP messages.

Important: For more information on this feature, refer *Subscriber Overcharging Protection* in this guide.

Proxy Mobile IP

Mobility for subscriber sessions is provided through the Mobile IP protocol as defined in RFCs 2002-2005. However, some older Mobile Nodes (MNs) do not support the Mobile IP protocol. The Proxy Mobile IP feature provides a mobility solution for these MNs.

For IP PDP contexts using Proxy Mobile IP, the MN establishes a session with the GGSN as it normally would. However, the GGSN/FA performs Mobile IP operations with an HA (identified by information stored in the subscriber's profile) on behalf of the MN (i.e. the MN is only responsible for maintaining the IP PDP context with the GGSN, no Agent Advertisement messages are communicated with the MN).

The MN is assigned an IP address by either the HA, an AAA server, or on a static-basis. The address is stored in a Mobile Binding Record (MBR) stored on the HA. Therefore, as the MN roams through the service provider's network, each time a hand-off occurs, the MN will continue to use the same IP address stored in the MBR on the HA.

Proxy Mobile IP can be performed on a per-subscriber basis based on information contained in their user profile, or for all subscribers facilitated by a specific APN. In the case of non-transparent IP PDP contexts, attributes returned from the subscriber's profile take precedence over the configuration of the APN.

Important: For more information on this feature, refer *Proxy Mobile IP* in this guide.

Session Persistence

Important: Other licenses (i.e. IP Security and L2TP) may be additionally required depending on your network deployment and implementation.

Provides seamless mobility to mobile subscribers as they roam between WLAN and 3G cellular access networks. This type of inter-technology roaming is ordinarily not possible as wireline access networks do not include SGSNs to permit inter-SGSN call hand-offs with cellular access networks.

The Cisco Session Persistence Solution maintains consistent user identities and application transparency for your mobile subscribers as they roam across bearer access networks. This is accomplished through the integration of Home Agent (HA) and GGSN functionality on the wireless access gateway in the packet network and the use of standards-based protocols such as Mobile IP and Mobile IP NAT Traversal. The solution also includes Session Persistence client software that runs on dual-mode WiFi/GPRS/EDGE and/or UMTS/W-CDMA access devices including cellular phones and laptop computers with wireless data cards.

The Session Persistence client is designed to permit Mobile IP tunneling over the applicable underlying network including cellular access connections and cable or XDSL broadband access networks. When the user is attached to a WiFi access network, the Session Persistence client utilizes a Mobile IP Co-located Care of Address Foreign Agent Service (CCoA FA) and establishes a MIP tunnel to the HA service in the platform. This scenario is completely transparent to the GGSN service that operates in the same system. The Mobile IP protocol requires a publicly addressable FA service; however, this is a problem when the mobile subscriber is located behind a NAT firewall. In this case, the NAT firewall has no way of maintaining state to associate the public NAT ed address with the private address assigned to the user by local DHCP server. Mobile IP NAT Traversal solves this problem by establishing a UDP/IP tunnel between the subscriber access device and Home Agent. The NAT firewall uses the UDP port address to build state for the subscriber session. During this Mobile IP transaction, the HA establishes a mobility binding record for the subscriber session.

When the subscriber roams to a 3GPP cellular access network, it uses the IP address from normal PDP IP context establishment as its new Mobile IP Care of Address to refresh the mobility binding record at the Home Agent. For reduced latency between access hand-offs, it is also possible to utilize a permanent 'always-on' PDP IP context with the IP address maintained in the MIP session persistence client. In this scenario, the mobile access device only needs to re-establish the dormant RAB wireless connection with the 3GPP access network prior to transmitting a new Mobile IP registration.

The system also enables network-provisioned VPNs for Session Persistence applications by permitting use of overlapping address pools on the HA and using various tunneling protocols including IPSEC, Layer 2 Tunneling Protocol (L2TP) and Ethernet IEEE 802.1Q VLANs for separation of subscriber traffic. This application may be further augmented by additional features such as 800 RADIUS Server Groups to permit use of enterprise controlled AAA servers and custom dictionaries.

Session Recovery Support

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

Session recovery is performed by mirroring key software processes (e.g. session manager and AAA manager) within the system. These mirrored processes remain in an idle state (in standby-mode), wherein they perform no processing, until they may be needed in the case of a software failure (e.g. a session manager task aborts). The system spawns new instances of "standby mode" session and AAA managers for each active Control Processor (CP) being used.

Additionally, other key system-level software tasks, such as VPN manager, are performed on a physically separate packet processing card to ensure that a double software fault (e.g. session manager and VPN manager fails at same time on same card) cannot occur. The packet processing card used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

The additional hardware resources required for session recovery include a standby System Processor Card (SPC) and a standby packet processing card.

There are following modes of Session Recovery:

- **Task recovery mode**: Wherein one or more session manager failures occur and are recovered without the need to use resources on a standby packet processing card. In this mode, recovery is performed by using the mirrored "standby-mode" session manager task(s) running on active packet processing cards. The "standby-mode" task is renamed, made active, and is then populated using information from other tasks such as AAA manager.
- **Full packet processing card recovery mode**: Used when a packet processing card hardware failure occurs, or when a packet processing card migration failure happens. In this mode, the standby packet processing card is

made active and the "standby-mode" session manager and AAA manager tasks on the newly activated packet processing card perform session recovery.

Session/Call state information is saved in the peer AAA manager task because each AAA manager and session manager task is paired together. These pairs are started on physically different packet processing cards to ensure task recovery.

Important: For more information on this feature, refer *Session Recovery* in *System Administration Guide*.

Traffic Policing and Rate Limiting

Allows the operator to proportion the network and support Service-level Agreements (SLAs) for customers.

The Traffic-Policing/Shaping feature enables configuring and enforcing bandwidth limitations on individual PDP contexts of a particular 3GPP traffic class. Values for traffic classes are defined in 3GPP TS 23.107 and are negotiated with the SGSN during PDP context activation using the values configured for the APN on the GGSN. Configuration and enforcement is done independently on the downlink and the uplink directions for each of the 3GPP traffic classes. Configuration is on a per-APN basis, but may be overridden for individual subscribers or subscriber tiers during RADIUS authentication/authorization.

A Token Bucket Algorithm (a modified trTCM, as specified in RFC2698) is used to implement the Traffic-Policing feature. The algorithm measures the following criteria when determining how to mark a packet:

- **Committed Data Rate (CDR)**: The guaranteed rate (in bits per second) at which packets may be transmitted/received for the subscriber during the sampling interval.
- **Peak Data Rate (PDR)**: The maximum rate (in bits per second) that packets may be transmitted/received for the subscriber during the sampling interval.
- **Burst-size**: The maximum number of bytes that may be transmitted/received for the subscriber during the sampling interval for both committed (CBS) and peak (PBS) rate conditions. This represents the maximum number of tokens that can be placed in the subscriber's "bucket". Note that the committed burst size (CBS) equals the peak burst size (PBS) for each subscriber.

Tokens are removed from the subscriber's bucket based on the size of the packets being transmitted/received. Every time a packet arrives, the system determines how many tokens need to be added (returned) to a subscriber's CBS (and PBS) bucket. This value is derived by computing the product of the time difference between incoming packets and the CDR (or PDR). The computed value is then added to the tokens remaining in the subscriber's CBS (or PBS) bucket. The total number of tokens can not be greater than the configured burst-size. If the total number of tokens is greater than the burst-size, the number is set to equal the burst-size. After passing through the Token Bucket Algorithm, the packet is internally classified with a color, as follows:

- There are not enough tokens in the PBS bucket to allow a packet to pass, then the packet is considered to be in violation and is marked "red" and the violation counter is incremented by one.
- There are enough tokens in the PBS bucket to allow a packet to pass, but not in the CBS "bucket", then the packet is considered to be in excess and is marked "yellow", the PBS bucket is decremented by the packet size, and the exceed counter is incremented by one.
- There are more tokens present in the CBS bucket than the size of the packet, then the packet is considered as conforming and is marked "green" and the CBS and PBS buckets are decremented by the packet size.

The APN on the GGSN can be configured with actions to take for red and yellow packets. Any of the following actions may be specified:

- **Drop**: The offending packet is discarded.
- **Transmit**: The offending packet is passed.

- **Lower the IP Precedence**: The packet's ToS octet is set to “0”, thus downgrading it to Best Effort, prior to passing the packet.
- **Buffer the Packet**: The packet stored in buffer memory and transmitted to subscriber once traffic flow comes in allowed bandwidth.

Different actions may be specified for red and yellow, as well as for uplink and downlink directions and different 3GPP traffic classes.

Important: For more information on this feature, refer *Traffic Policing and Shaping* in this guide.

Web Element Management System

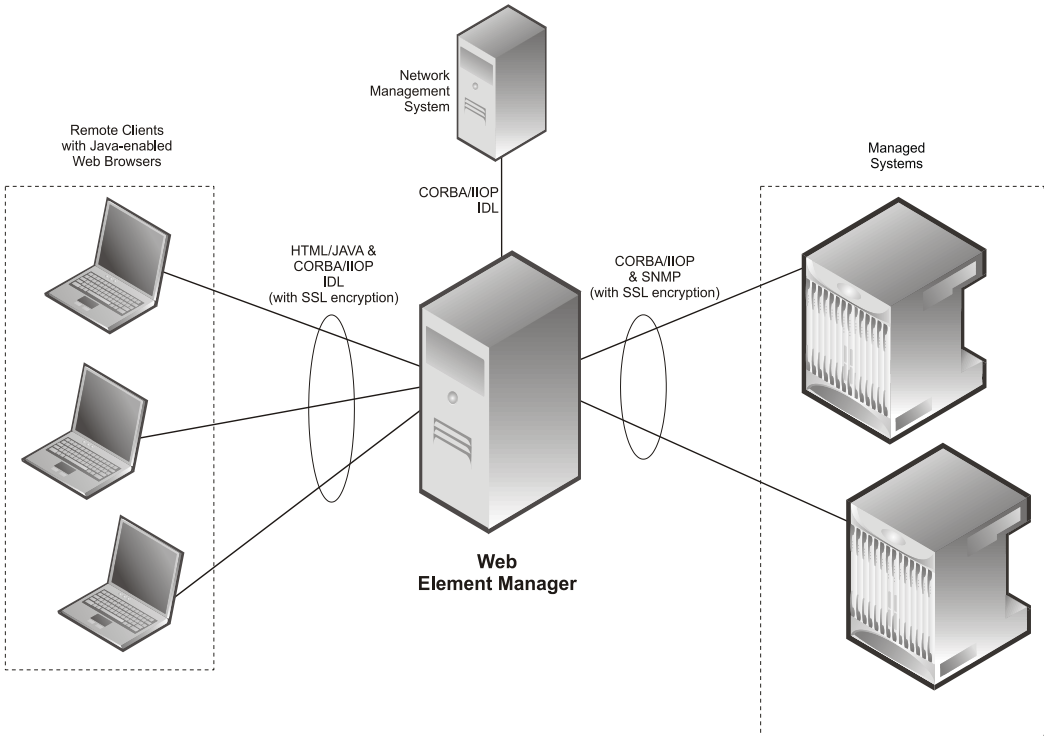
Provides a Graphical User Interface (GUI) for performing Fault, Configuration, Accounting, Performance, and Security (FCAPS) management of the system.

The Web Element Manager is a Common Object Request Broker Architecture (CORBA)-based application that provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) management capability for the system.

For maximum flexibility and scalability, the Web Element Manager application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the fault-tolerant Sun® Solaris® operating system.

The following figure demonstrates various interfaces between the Cisco Web Element Manager and other network components.

Figure 103. Web Element Manager Network Interfaces



Important: For more information on WEM support, refer *WEM Installation and Administration Guide*.

How GGSN Works

This section provides information on the function of the GGSN in a GPRS/UMTS network and presents call procedure flows for different stages of session setup.

The following topics and procedure flows are included:

- PDP Context Processing
- Dynamic IP Address Assignment
- Subscriber Session Call Flows

PDP Context Processing

PDP context processing is based on the APN that the subscriber is attempting to access. Templates for all of the possible APNs that subscribers will be accessing must be configured within the system. Up to 1024 APNs can be configured on the system.

Each APN template consists of parameters pertaining to how PDP contexts are processed such as the following:

- **Type**: The system supports IPv4, IPv6, and PPP PDP contexts. For IPv6 PDP configuration to work, at least one IPv6 interface needs to be configured in the destination context.
- **Accounting protocol**: Support is provided for using either the GTPP or Remote Authentication Dial-In User Service (RADIUS) protocols. In addition, an option is provided to disable accounting if desired.
- **Authentication protocol**: Support is provided for using any of the following:
  - Challenge Handshake Authentication Protocol (CHAP)
  - Microsoft CHAP (MSCHAP)
  - Password Authentication Protocol (PAP)
  - IMSI-based authentication
  - MSISDN-based authentication

    In addition, an option is provided to disable authentication if desired.
- **Charging characteristics**: Each APN template can be configured to either accept the charging characteristics it receives from the SGSN for a PDP context or use it's own characteristics.
- **IP address allocation method**: IP addresses for PDP contexts can be assigned using one of the following methods:
  - **Statically**: The APN template can be configured to provide support for MS-requested static IP addresses. Additionally, a static address can be configured in a subscriber's profile on an authentication server and allocated upon successful authentication.
  - **Dynamically**: The APN template can be configured to dynamically assign an IP address from locally configured address pools or via a Dynamic Host Control Protocol (DHCP) server. Additional information on dynamic address assignment can be found in the *Dynamic IP Address Assignment* section that follows.

    Important: Static IP addresses configured in subscriber profiles must also be part of a static IP address pool configured locally on the system.

- **Selection mode**: The MS's right to access the APN can be either verified or unverified. For verified access, the SGSN specifies the APN that should be used. For unverified access, the APN can be specified by either the SGSN or the MS.
- **Timeout**: Absolute and idle session timeout values specify the amount of time that an MS can remain connected.
- **Mobile IP configuration**: Mobile IP requirements, HA address, and other related parameters are configured in the APN template.
- **Proxy Mobile IP support**: Mobile IP support can be enabled for all subscribers facilitated by the APN. Alternatively, it can be enabled for individual subscribers via parameters in their RADIUS or local-user profiles.
- **Quality of Service**: Parameters pertaining to QoS feature support such as for Dynamic Renegotiation, Traffic Policing, and DSCP traffic class.

A total of 11 PDP contexts are supported per subscriber. These could be all primaries, or 1 Primary and 10 secondaries or any combination of primary and secondary. Note that there must be at least one primary PDP context in order for secondaries to come up.

Dynamic IP Address Assignment

IP addresses for PDP contexts can either be static—an IP address is permanently assigned to the MS—or dynamic—an IP address is temporarily assigned to the MS for the duration of the PDP context.

As previously described in the *PDP Context Processing* section of this chapter, the method by which IP addresses are assigned to a PDP context is configured on an APN-by-APN basis. Each APN template dictates whether it will support static or dynamic addresses. If dynamic addressing is supported, the following methods can be implemented:

- **Local pools:** The system supports the configuration of public or private IP address pools. Addresses can be allocated from these pools as follows:
  - **Public pools:** Provided that dynamic assignment is supported, a parameter in the APN configuration mode specifies the name of the local public address pool to use for PDP contexts facilitated by the APN.
  - **Private pools:** Provided that dynamic assignment is supported, the name of the local private pool can be specified in the subscriber's profile. The receipt of a valid private pool name will override the APN's use of addresses from public pools.
- **Dynamic Host Control Protocol (DHCP):** The system can be configured to use DHCP PDP context address assignment using either of the following mechanisms:
  - **DHCP-proxy**: The system acts as a proxy for client (MS) and initiates the DHCP Discovery Request on behalf of client (MS). Once it receives an allocated IP address from DHCP server in response to DHCP Discovery Request, it assigns the received IP address to the MS. This allocated address must be matched with the an address configured in an IP address pool on the system. This complete procedure is not visible to MS.
  - **DHCP-relay**: The system acts as a relay for client (MS) and forwards the DHCP Discovery Request received from client (MS). Once it receives an allocated IP address from DHCP server in response to DHCP Discovery Request, it assigns the received IP address to the MS.

In addition to the above methods, IP addresses for subscriber Mobile IP sessions are also dynamically assigned by the subscriber's home network upon registration. The GGSN/FA, in turn, provide the assigned address to the mobile station.

Subscriber Session Call Flows

This section provides information on how GPRS/UMTS subscriber data sessions are processed by the system GGSN. The following data session scenarios are provided:

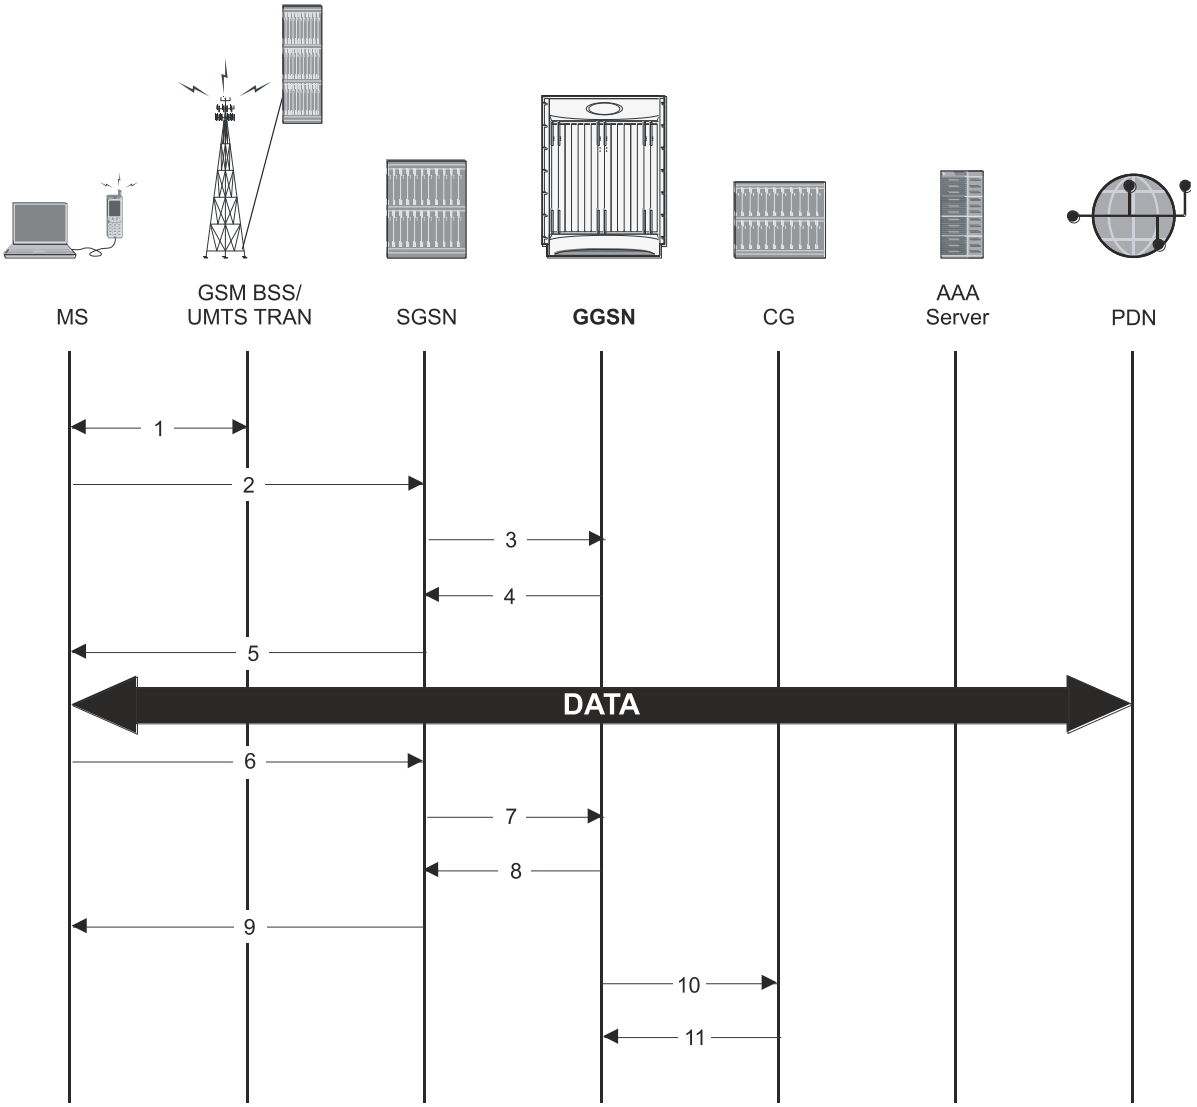
- **Transparent IP:** The subscriber is provided basic access to a PDN without the GGSN authenticating the subscriber. Either a static or dynamic IP address can be assigned to the MS in this scenario.
- **Non-transparent IP**: The GGSN provides subscriber authentication services for the data session. Either a static or dynamic IP address can be assigned to the MS in this scenario.
- **Network-initiated**: An IP Packet Data Unit (PDU) is received by the GGSN from the PDN for a specific subscriber. If configured to support network-initiated sessions, the GGSN, will initiate the process of paging the MS and establishing a PDP context.
- **PPP Direct Access**: The GGSN terminates the subscriber's PPP session and provides subscriber authentication services for the data session. Either a static or dynamic IP address can be assigned to the MS in this scenario.
- **Virtual Dialup Access**: The GGSN functions as an LAC, encapsulates subscriber packets using L2TP, and tunnels them directly to an LNS for processing.
- **Corporate IP VPN Connectivity**: Similar to the Virtual Dialup Access model, however, the GGSN is configured to tunnel subscriber packets to a corporate server using IP-in-IP.
- **Mobile IP**: Subscriber traffic is routed to their home network via a tunnel between the GGSN/FA and an HA. The subscriber's IP PDP context is assigned an IP address from the HA.
- **Proxy Mobile IP**: Provides a mobility solution for subscribers whose Mobile Nodes (MNs) do not support the Mobile IP protocol. The GGSN/FA proxy the Mobile IP tunnel with the HA on behalf of the MS. The subscriber receives an IP address from their home network. As the subscriber roams through the network, the IP address is maintained providing the subscriber with the opportunity to use IP applications that require seamless mobility such as transferring files.
- **IPv6 Stateless Address Auto Configuration**: The mobile station may select any value for the interface identifier portion of the address. The only exception is the interface identifier for the link-local address used by the mobile station. This interface identifier is assigned by the GGSN to avoid any conflict between the mobile station link-local address and the GGSN address. The mobile station uses the interface ID assigned by the GGSN during stateless address auto-configuration procedure (e.g., during the initial router advertisement messages). Once this is over, the mobile can select any interface ID for further communication as long as it does not conflict with the GGSN's interface ID (that the mobile would learn through router advertisement messages from the GGSN).

Additionally, information about the process used by the system to dynamically assign IP addresses to the MS is provided in following sections.

Transparent Session IP Call Flow

The following figure and the text that follows describe the call flow for a successful transparent data session.

Figure 104. Transparent IP Session Call Flow



1. The Mobile Station (MS) goes through the process of attaching itself to the GPRS/UMTS network.
2. The MS sends an Activate PDP Context Request message that is received by an SGSN. The message contains information about the subscriber such as the Network layer Service Access Point Identifier (NSAPI), PDP Type, PDP Address, Access Point Name (APN), Quality of Service (QoS) requested, and PDP configuration options.
3. The SGSN authenticates the request message and sends a Create PDP Context Request message to a GGSN using the GPRS Tunneling Protocol (GTPC, "C" indicates the control signaling aspect of the protocol). The recipient GGSN is selected based on either the request of the MS or is automatically selected by the SGSN. The message consists of various information elements including: PDP Type, PDP Address, APN, charging characteristics, and Tunnel Endpoint Identifier (TEID, if the PDP Address was static).
4. The GGSN determines if it can facilitate the session (in terms of memory or CPU resources, configuration, etc.) and creates a new entry in its PDP context list and provides a Charging ID for the session.

If the MS required the dynamic assignment of an IP address (i.e., the PDP Address received from the mobile was null), the GGSN will assign one. The IP address assignment methods supported by the system GGSN are described in the *Dynamic IP Address Assignment* section of this guide.

The GGSN replies with an affirmative Create PDP Context Response using GTPC. The response will contain information elements such as the PDP Address representing either the static address requested by the MS or the address assigned by the GGSN, the TEID used to reference PDP Address, and PDP configuration options specified by the GGSN.

5. The SGSN returns an Activate PDP Context Accept response to the MS.

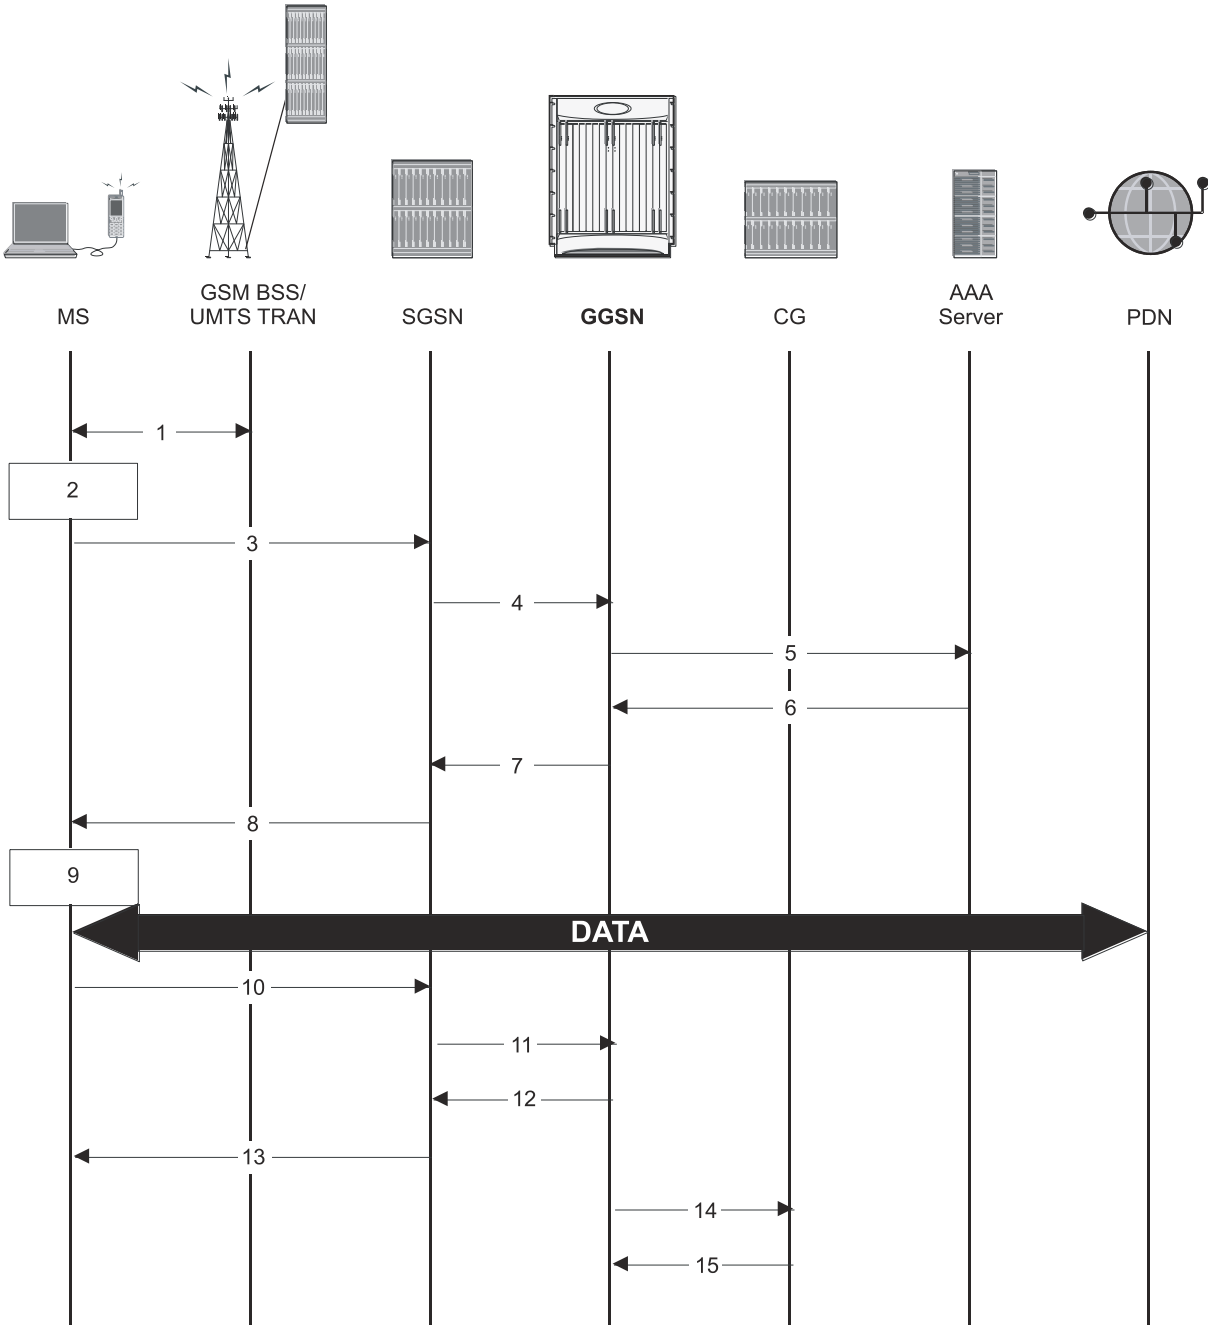
   The MS can now send and receive data to or from the PDN until the session is closed or times out. The MS can initiate multiple PDP contexts if desired and supported by the system. Each additional PDP context can share the same IP address or use alternatives.

6. The MS can terminate the data session at any time. To terminate the session, the MS sends a Deactivate PDP Context Request message that is received by the SGSN.
7. The SGSN sends a Delete PDP Context Request message to the GGSN facilitating the data session. The message includes the information elements necessary to identify the PDP context (i.e., TEID, and NSAPI).
8. The GGSN removes the PDP context from memory and returns a Delete PDP Context Response message to the SGSN. If the PDP context was the last associated with a particular dynamically assigned PDP Address, the GGSN will re-claim the IP address for use by subsequent PDP contexts.
9. The SGSN returns a Deactivate PDP Context Accept message to the MS.
10. The GGSN delivers the GGSN Charging Detail Records (G-CDRs) to a Charging Gateway (CG) using GTP Prime (GTPP). Note that, though not shown in this example, the GGSN could optionally be configured to send partial CDRs while the PDP context is active.
11. For each accounting message received from the GGSN, the CG responds with an acknowledgement.

Non-Transparent IP Session Call Flow

The following figure and the text that follows describe the call flow for a successful non-transparent data session.

Figure 105. Non-Transparent IP Session Call Flow



1. The Mobile Station (MS) goes through the process of attaching itself to the GPRS/UMTS network.
2. The Terminal Equipment (TE) aspect of the MS sends AT commands to the Mobile Terminal (MT) aspect of the MS to place it into PPP mode.

   The Link Control Protocol (LCP is then used to configure the Maximum-Receive Unit size and the authentication protocol (Challenge-Handshake Authentication Protocol (CHAP), Password Authentication Protocol (PAP), or none). If CHAP or PAP is used, the TE will authenticate itself to the MT, which, in turn, stores the authentication information.

Upon successful authentication, the TE sends an Internet Protocol Control Protocol (IPCP) Configure-Request message to the MT. The message will either contain a static IP address to use or request that one be dynamically assigned.

3. The MS sends an Activate PDP Context Request message that is received by an SGSN. The message contains information about the subscriber such as the Network layer Service Access Point Identifier (NSAPI), PDP Type, PDP Address, Access Point Name (APN), Quality of Service (QoS) requested, and PDP configuration options.
4. The SGSN authenticates the request message and sends a Create PDP Context Request message to a GGSN using the GPRS Tunneling Protocol (GTPC, "C" indicates the control signaling aspect of the protocol). The recipient GGSN is selected based on either the request of the MS or is automatically selected by the SGSN. The message consists of various information elements including: PDP Type, PDP Address, APN, charging characteristics, and tunnel endpoint identifier (TEID, if the PDP Address was static).
5. The GGSN determines if it can facilitate the session (in terms of memory or CPU resources, configuration, etc.) and creates a new entry in its PDP context list and provides a Charging ID for the session.

   From the APN specified in the message, the GGSN determines whether or not the subscriber is to be authenticated, how an IP address should be assigned if using dynamic allocation, and how to route the session.

   If authentication is required, the GGSN attempts to authenticate the subscriber locally against profiles stored in memory or send a RADIUS Access-Request message to an AAA server.

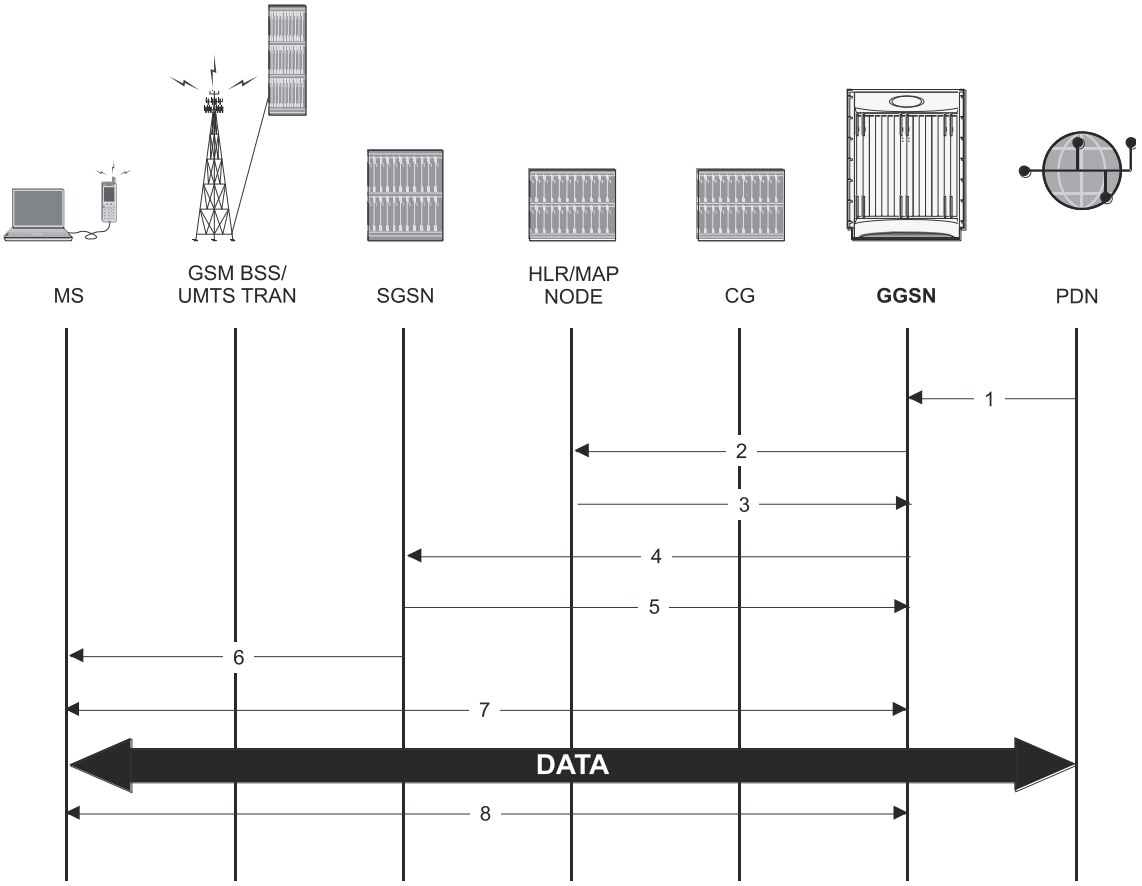
   If the MS required the dynamic assignment of an IP address (i.e., the PDP Address received from the mobile was null), the GGSN will assign one. The IP address assignment methods supported by the system GGSN are described in the *Dynamic IP Address Assignment* section of this chapter.
6. If the GGSN authenticated the subscriber to an AAA server, the AAA server responds with a RADIUS Access-Accept message indicating successful authentication.
7. The GGSN replies with an affirmative Create PDP Context Response using GTPC. The response will contain information elements such as the PDP Address representing either the static address requested by the MS or the address assigned by the GGSN, the TEID used to reference PDP Address, and PDP configuration options specified by the GGSN.
8. The SGSN returns an Activate PDP Context Accept message to the MS. The message includes response to the configuration parameters sent in the initial request.
9. The MT, will respond to the TE's IPCP Config-request with an IPCP Config-Ack message.

   The MS can now send and receive data to or from the PDN until the session is closed or times out. The MS can initiate multiple PDP contexts if desired and supported by the system. Each additional PDP context can share the same IP address or use alternatives.
10. The MS can terminate the data session at any time. To terminate the session, the MS sends a Deactivate PDP Context Request message that is received by the SGSN.
11. The SGSN sends a Delete PDP Context Request message to the GGSN facilitating the data session. The message includes the information elements necessary to identify the PDP context (i.e., TEID, and NSAPI).
12. The GGSN removes the PDP context from memory and returns a Delete PDP Context Response message to the SGSN. If the PDP context was the last associated with a particular dynamically assigned PDP Address, the GGSN will re-claim the IP address for use by subsequent PDP contexts.
13. The SGSN returns a Deactivate PDP Context Accept message to the MS.
14. The GGSN delivers the GGSN Charging Detail Records (G-CDRs) to a Charging Gateway (CG) using GTP Prime (GTPP). Note that, though not shown in this example, the GGSN could optionally be configured to send partial CDRs while the PDP context is active.
15. For each accounting message received from the GGSN, the CG responds with an acknowledgement.

Network-Initiated Session Call Flow

The following figure and the text that follows describe the call flow for a successful network-initiated data session.

Figure 106. Network-initiated Session Call Flow



1. An IP Packet Data Unit (PDU) is received by the GGSN from the PDN. The GGSN determines if it is configured to support network-initiated sessions. If not, it will discard the packet. If so, it will begin the Network-Requested PDP Context Activation procedure.
2. The GGSN may issue a Send Routing Information for GPRS request to the HLR to determine if the MS is reachable. The message includes the MS's International Mobile Subscriber Identity (IMSI).
3. If the MS is reachable, the HLR returns a Send Routing Information for GPRS Ack containing the address of the SGSN currently associated with the MS's IMSI.
4. The GGSN sends a PDU Notification Request message to the SGSN address supplied by the HLR. This message contains the IMSI, PDP Type, PDP Address, and APN associated with the session.
5. The SGSN sends a PDU Notification Response to the GGSN indicating that it will attempt to page the MS requesting that it activate the PDP address indicated in the GGSN's request.
6. The SGSN sends a Request PDP Context Activation message to the MS containing the information supplied by the GGSN.
7. The MS begins the PDP Context Activation procedure as described in *step 2* through *step 5* of the *Transparent Session IP Call Flow* section of this chapter.

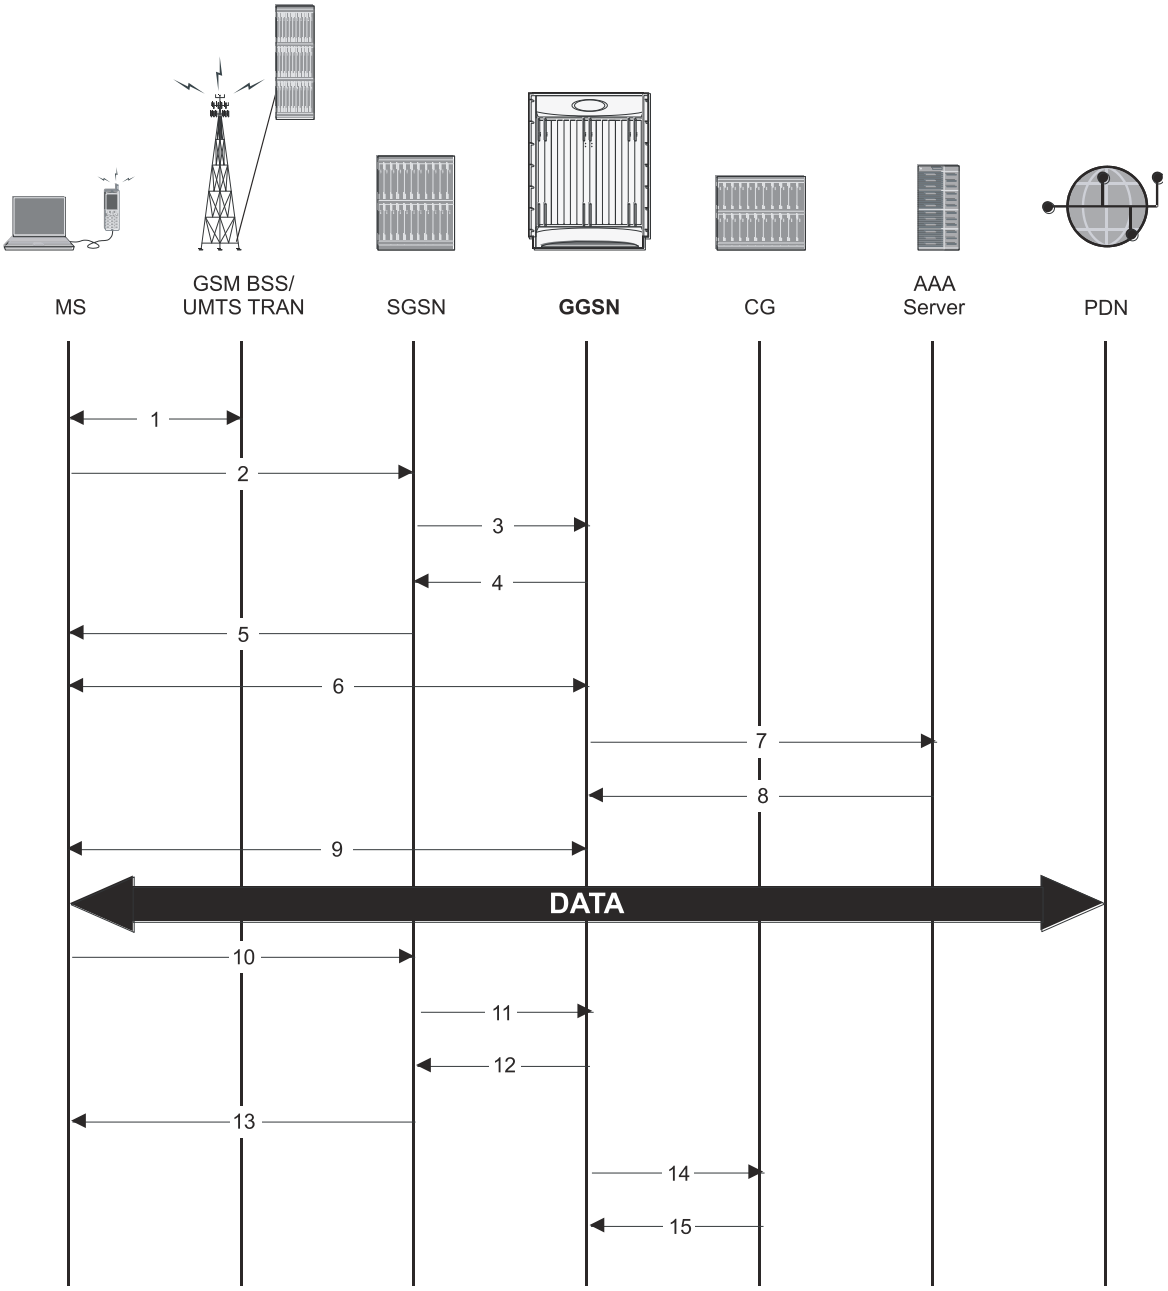
Upon PDP context establishment, the MS can send and receive data to or from the PDN until the session is closed or times out.

8. The MS can terminate the data session at any time. To terminate the session, the MS begins the PDP Context De-Activation procedure as described in *step 6* through *step 11* of the *Transparent Session IP Call Flow* section of this chapter.

PPP Direct Access Call Flow

The following figure and the text that follows describe the call flow for a successful PPP Direct Access data session.

Figure 107. PPP Direct Access Call Flow



1. The Mobile Station (MS) goes through the process of attaching itself to the GPRS/UMTS network.
2. The MS sends an Activate PDP Context Request message that is received by an SGSN. The message contains information about the subscriber such as the Network layer Service Access Point Identifier (NSAPI), PDP Type, PDP Address, Access Point Name (APN), Quality of Service (QoS) requested, and PDP configuration options.
3. The SGSN authenticates the request message and sends a Create PDP Context Request message to a GGSN using the GPRS Tunneling Protocol (GTPC, "C" indicates the control signaling aspect of the protocol). The

recipient GGSN is automatically selected by the SGSN. The message consists of various information elements including: PDP Type, PDP Address, APN, and charging characteristics.

4. The GGSN determines if it can facilitate the session (in terms of memory or CPU resources, configuration, etc.) and creates a new entry in its PDP context list and provides a Charging ID for the session. It determines that the PDP context type is PPP and based on the APN, what authentication protocol to use and how to perform IP address assignment.

   The GGSN replies with an affirmative Create PDP Context Response using GTPC.

5. The SGSN returns an Activate PDP Context Accept response to the MS.
6. The MS and the GGSN negotiate PPP.
7. The GGSN forwards authentication information received from the MS as part of PPP negotiation to the AAA server in the form of an Access-Request.
8. The AAA server authenticates the MS and sends an Access-Accept message to the GGSN.
9. The GGSN assigns an IP address to the MS and completes the PPP negotiation process. More information about IP addressing for PDP contexts is located in the *PDP Context Processing* and *Dynamic IP Address Assignment* sections of this chapter.

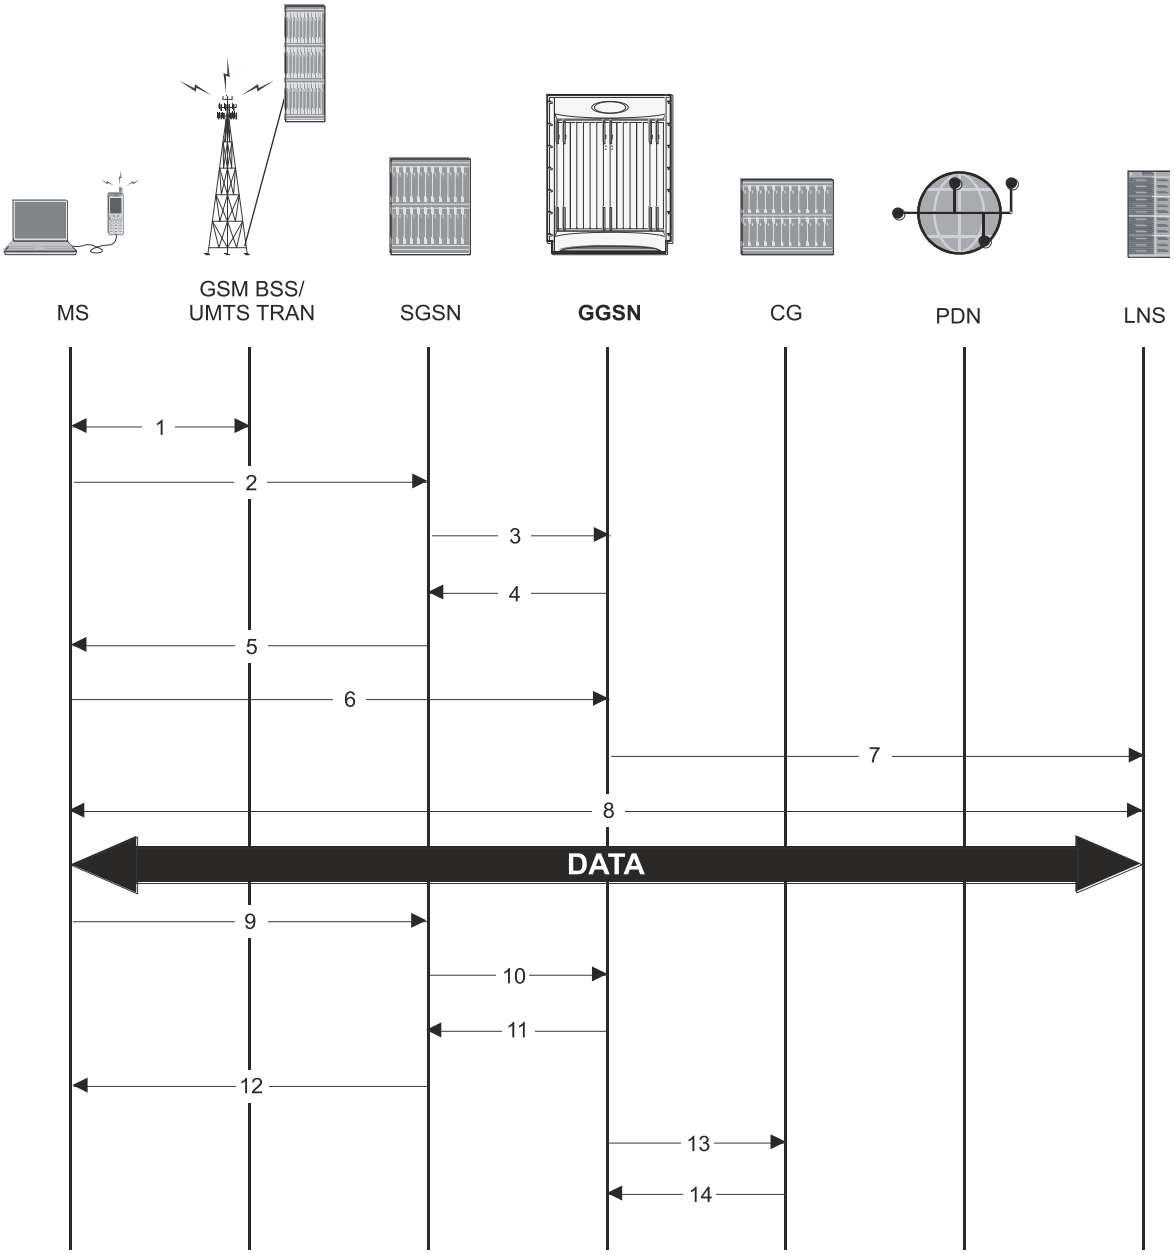
   Once the PPP negotiation process is complete, the MS can send and receive data.

10. The MS can terminate the data session at any time. To terminate the session, the MS sends a Deactivate PDP Context Request message that is received by the SGSN.
11. The SGSN sends a Delete PDP Context Request message to the GGSN facilitating the data session. The message includes the information elements necessary to identify the PDP context.
12. The GGSN removes the PDP context from memory and returns a Delete PDP Context Response message to the SGSN. If the PDP context was the last associated with a particular dynamically assigned PDP Address, the GGSN will re-claim the IP address for use by subsequent PDP contexts.
13. The SGSN returns a Deactivate PDP Context Accept message to the MS.
14. The GGSN delivers the GGSN Charging Detail Records (G-CDRs) to a Charging Gateway (CG) using GTP Prime (GTPP). Note that, though not shown in this example, the GGSN could optionally be configured to send partial CDRs while the PDP context is active.
15. For each accounting message received from the GGSN, the CG responds with an acknowledgement.

Virtual Dialup Access Call Flow

The following figure and the text that follows describe the call flow for a successful VPN Dialup Access data session.

Figure 108. Virtual Dialup Access Call Flow



1. The Mobile Station (MS) goes through the process of attaching itself to the GPRS/UMTS network.
2. The MS sends an Activate PDP Context Request message that is received by an SGSN. The message contains information about the subscriber such as the Network layer Service Access Point Identifier (NSAPI), PDP Type, PDP Address, Access Point Name (APN), Quality of Service (QoS) requested, and PDP configuration options.
3. The SGSN authenticates the request message and sends a Create PDP Context Request message to a GGSN using the GPRS Tunneling Protocol (GTPC, "C" indicates the control signaling aspect of the protocol). The recipient GGSN is automatically selected by the SGSN. The message consists of various information elements including: PDP Type, PDP Address, APN, and charging characteristics.

4. The GGSN determines if it can facilitate the session (in terms of memory or CPU resources, configuration, etc.) and creates a new entry in its PDP context list and provides a Charging ID for the session. It determines the PDP context type and based on the APN, what authentication protocol to use and how to perform IP address assignment.

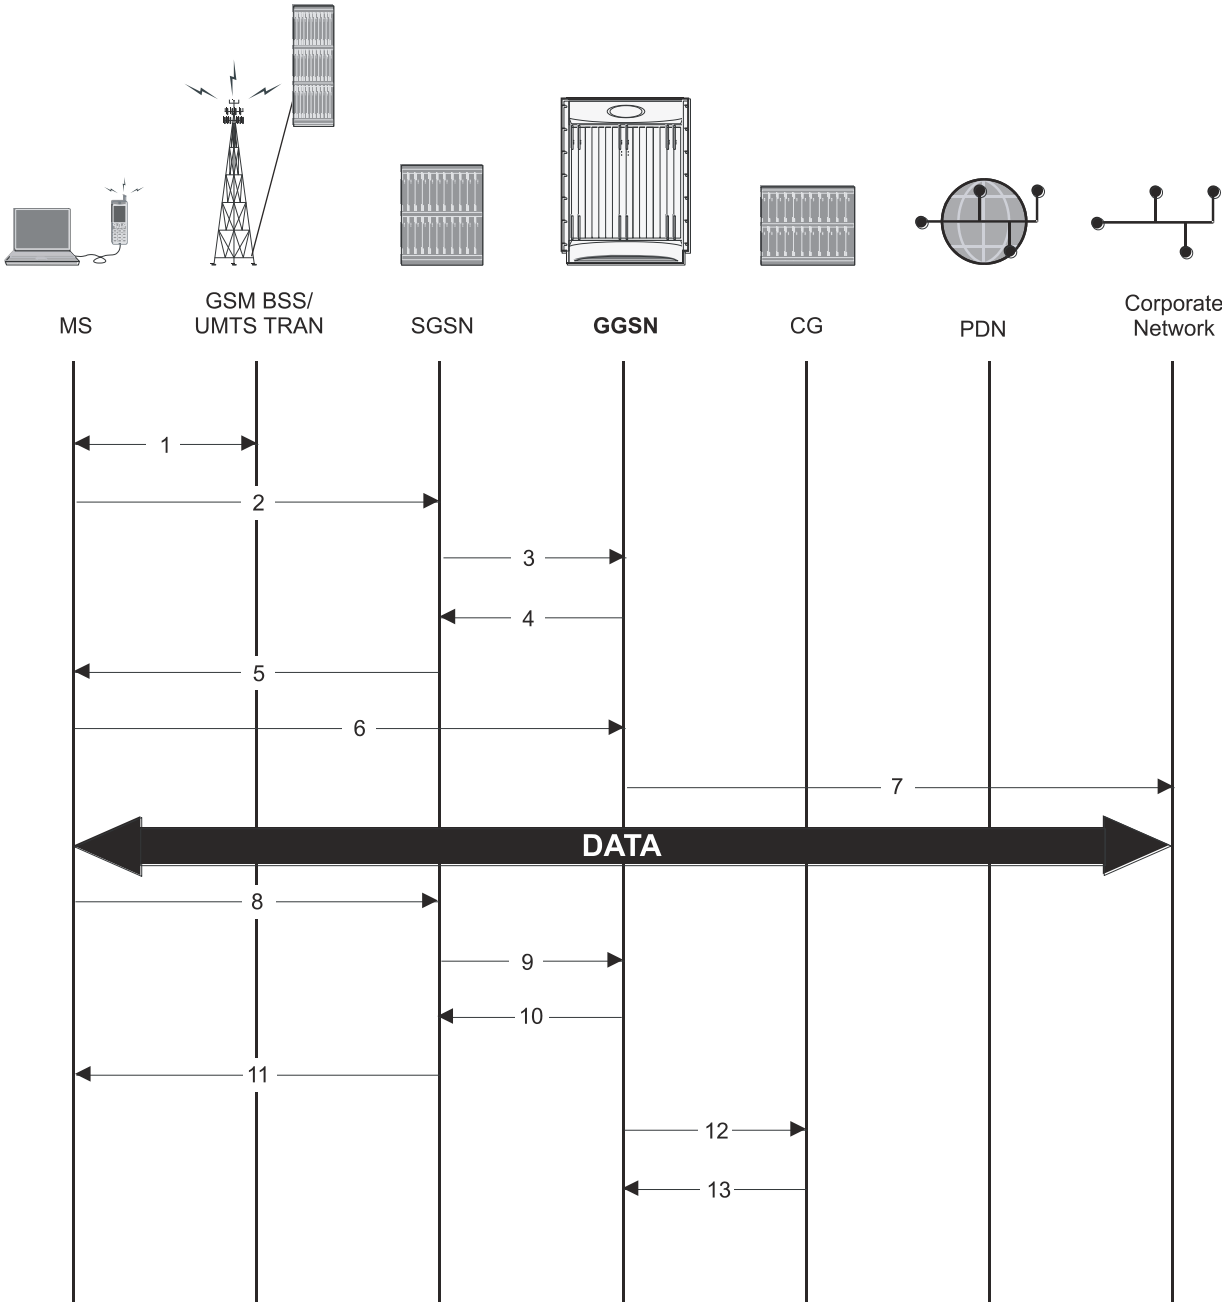
   The GGSN replies with an affirmative Create PDP Context Response using GTPC.
5. The SGSN returns an Activate PDP Context Accept response to the MS.
6. The MS sends packets which are received by the GGSN.
7. The GGSN encapsulates the packets from the MS using L2TP and tunnels them to the LNS.
8. The LNS terminates the tunnel and un-encapsulates the packets.

   The MS can send and receive data over the L2TP tunnel facilitated by the GGSN.
9. The MS can terminate the data session at any time. To terminate the session, the MS sends a Deactivate PDP Context Request message that is received by the SGSN.
10. The SGSN sends a Delete PDP Context Request message to the GGSN facilitating the data session. The message includes the information elements necessary to identify the PDP context.
11. The GGSN removes the PDP context from memory and returns a Delete PDP Context Response message to the SGSN. If the PDP context was the last associated with a particular dynamically assigned PDP Address, the GGSN will re-claim the IP address for use by subsequent PDP contexts.
12. The SGSN returns a Deactivate PDP Context Accept message to the MS.
13. The GGSN delivers the GGSN Charging Detail Records (G-CDRs) to a Charging Gateway (CG) using GTP Prime (GTPP). Note that, though not shown in this example, the GGSN could optionally be configured to send partial CDRs while the PDP context is active.
14. For each accounting message received from the GGSN, the CG responds with an acknowledgement.

Corporate IP VPN Connectivity Call Flow

The following figure and the text that follows describe the call flow for a successful Corporate IP Connectivity data session.

Figure 109. Corporate IP VPN Connectivity Call Flow



1. The Mobile Station (MS) goes through the process of attaching itself to the GPRS/UMTS network.
2. The MS sends an Activate PDP Context Request message that is received by an SGSN. The message contains information about the subscriber such as the Network layer Service Access Point Identifier (NSAPI), PDP Type, PDP Address, Access Point Name (APN), Quality of Service (QoS) requested, and PDP configuration options.
3. The SGSN authenticates the request message and sends a Create PDP Context Request message to a GGSN using the GPRS Tunneling Protocol (GTPC, "C" indicates the control signaling aspect of the protocol). The

recipient GGSN is automatically selected by the SGSN. The message consists of various information elements including: PDP Type, PDP Address, APN, and charging characteristics.

4. The GGSN determines if it can facilitate the session (in terms of memory or CPU resources, configuration, etc.) and creates a new entry in its PDP context list and provides a Charging ID for the session. It determines the PDP context type and based on the APN, what authentication protocol to use and how to perform IP address assignment.

   If the MS required the dynamic assignment of an IP address (i.e., the PDP Address received from the mobile was null), the GGSN will assign one. The IP address assignment methods supported by the system GGSN are described in the *Dynamic IP Address Assignment* section of this chapter.

   The GGSN replies with an affirmative Create PDP Context Response using GTPC.

5. The SGSN returns an Activate PDP Context Accept response to the MS.
6. The MS sends IP packets which are received by the GGSN.
7. The GGSN encapsulates the IP packets from the MS using IP-in-IP and tunnels them to the subscriber's corporate network.

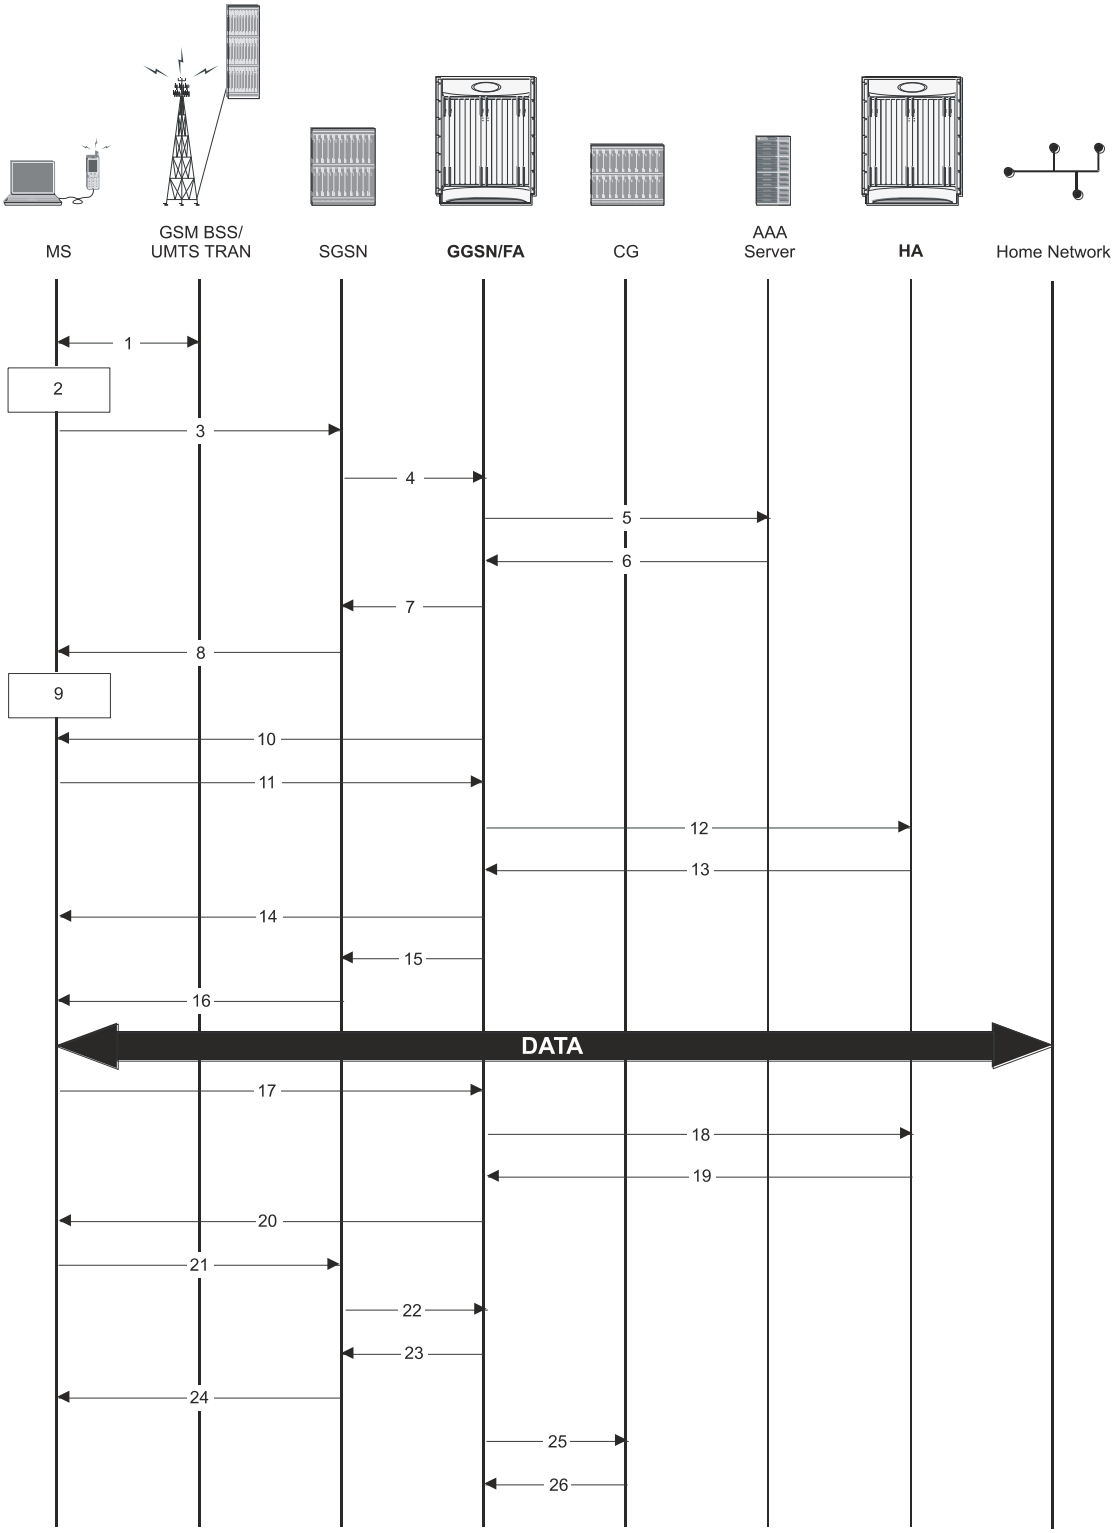
   All data sent and received by the MS over the IP-in-IP tunnel facilitated by the GGSN.

8. The MS can terminate the data session at any time. To terminate the session, the MS sends a Deactivate PDP Context Request message that is received by the SGSN.
9. The SGSN sends a Delete PDP Context Request message to the GGSN facilitating the data session. The message includes the information elements necessary to identify the PDP context.
10. The GGSN removes the PDP context from memory and returns a Delete PDP Context Response message to the SGSN. If the PDP context was the last associated with a particular dynamically assigned PDP Address, the GGSN will re-claim the IP address for use by subsequent PDP contexts.
11. The SGSN returns a Deactivate PDP Context Accept message to the MS.
12. The GGSN delivers the GGSN Charging Detail Records (G-CDRs) to a Charging Gateway (CG) using GTP Prime (GTPP). Note that, though not shown in this example, the GGSN could optionally be configured to send partial CDRs while the PDP context is active.
13. For each accounting message received from the GGSN, the CG responds with an acknowledgement.

Mobile IP Call Flow

The following figure and the text that follows describe the call flow for a successful Corporate IP Connectivity data session.

Figure 110. Mobile IP Call Flow



1. The Mobile Station (MS) goes through the process of attaching itself to the GPRS/UMTS network.

2. The Terminal Equipment (TE) aspect of the MS sends AT commands to the Mobile Terminal (MT) aspect of the MS to place it into PPP mode.

   The Link Control Protocol (LCP is then used to configure the Maximum-Receive Unit size and the authentication protocol (Challenge-Handshake Authentication Protocol (CHAP), Password Authentication Protocol (PAP), or none). If CHAP or PAP is used, the TE will authenticate itself to the MT, which, in turn, stores the authentication information.

   Upon successful authentication, the TE sends an Internet Protocol Control Protocol (IPCP) Configure-Request message to the MT. The message will either contain a static IP home address to use or request that one be dynamically assigned.

3. The MS sends an Activate PDP Context Request message that is received by an SGSN. The message contains information about the subscriber such as the Network layer Service Access Point Identifier (NSAPI), PDP Type, Access Point Name (APN), Quality of Service (QoS) requested, and PDP configuration options.

   Note that regardless of whether or not the MS has a static address or is requesting a dynamic address, the "Requested PDP Address" field is omitted from the request when using Mobile IP.

4. The SGSN authenticates the request message and sends a Create PDP Context Request message to a GGSN using the GPRS Tunneling Protocol (GTPC, "C" indicates the control signaling aspect of the protocol). The recipient GGSN is selected based on either the request of the MS or is automatically selected by the SGSN. The message consists of various information elements including: PDP Type, Requested PDP con, APN, charging characteristics, and Tunnel Endpoint Identifier (TEID).

5. The GGSN determines if it can facilitate the session (in terms of memory or CPU resources, configuration, etc.) and creates a new entry in its PDP context list and provides a Charging ID for the session.

   From the APN specified in the message, the GGSN determines how to handle the PDP context including whether or not Mobile IP should be used.

   If authentication is required, the GGSN attempts to authenticate the subscriber locally against profiles stored in memory or send a RADIUS Access-Request message to an AAA server.

6. If the GGSN authenticated the subscriber to an AAA server, the AAA server responds with a RADIUS Access-Accept message indicating successful authentication.

7. The GGSN replies to the SGSN with a PDP Context Response using GTPC. The response will contain information elements such as the PDP Address, and PDP configuration options specified by the GGSN. Note that for Mobile IP, the GGSN returns a PDP Address of 0.0.0.0 indicating that it will be reset with a Home address after the PDP context activation procedure.

8. The SGSN returns an Activate PDP Context Accept message to the MS. The message includes response to the configuration parameters sent in the initial request.

9. The MT, will respond to the TE's IPCP Config-request with an IPCP Config-Ack message. This ends the PPP mode between the MT and TE components of the MS.

   Data can now be transmitted between the MS and the GGSN.

10. The FA component of the GGSN sends an Agent Advertisement message to the MS. The message contains the FA parameters needed by the mobile such as one or more card-of addresses. The message is sent as an IP limited broadcast message (i.e. destination address 255.255.255.255), however only on the requesting MS's TEID to avoid broadcast over the radio interface.

11. The MS sends a Mobile IP Registration request to the GGSN/FA. This message includes either the MS's static home address or it can request a temporary address by sending 0.0.0.0 as its home address. Additionally, the request must always include the Network Access Identifier (NAI) in a Mobile-Node-NAI Extension.

12. The FA forwards the registration request from the MS to the HA while the MS's home address or NAI and TEID are stored by the GGSN.

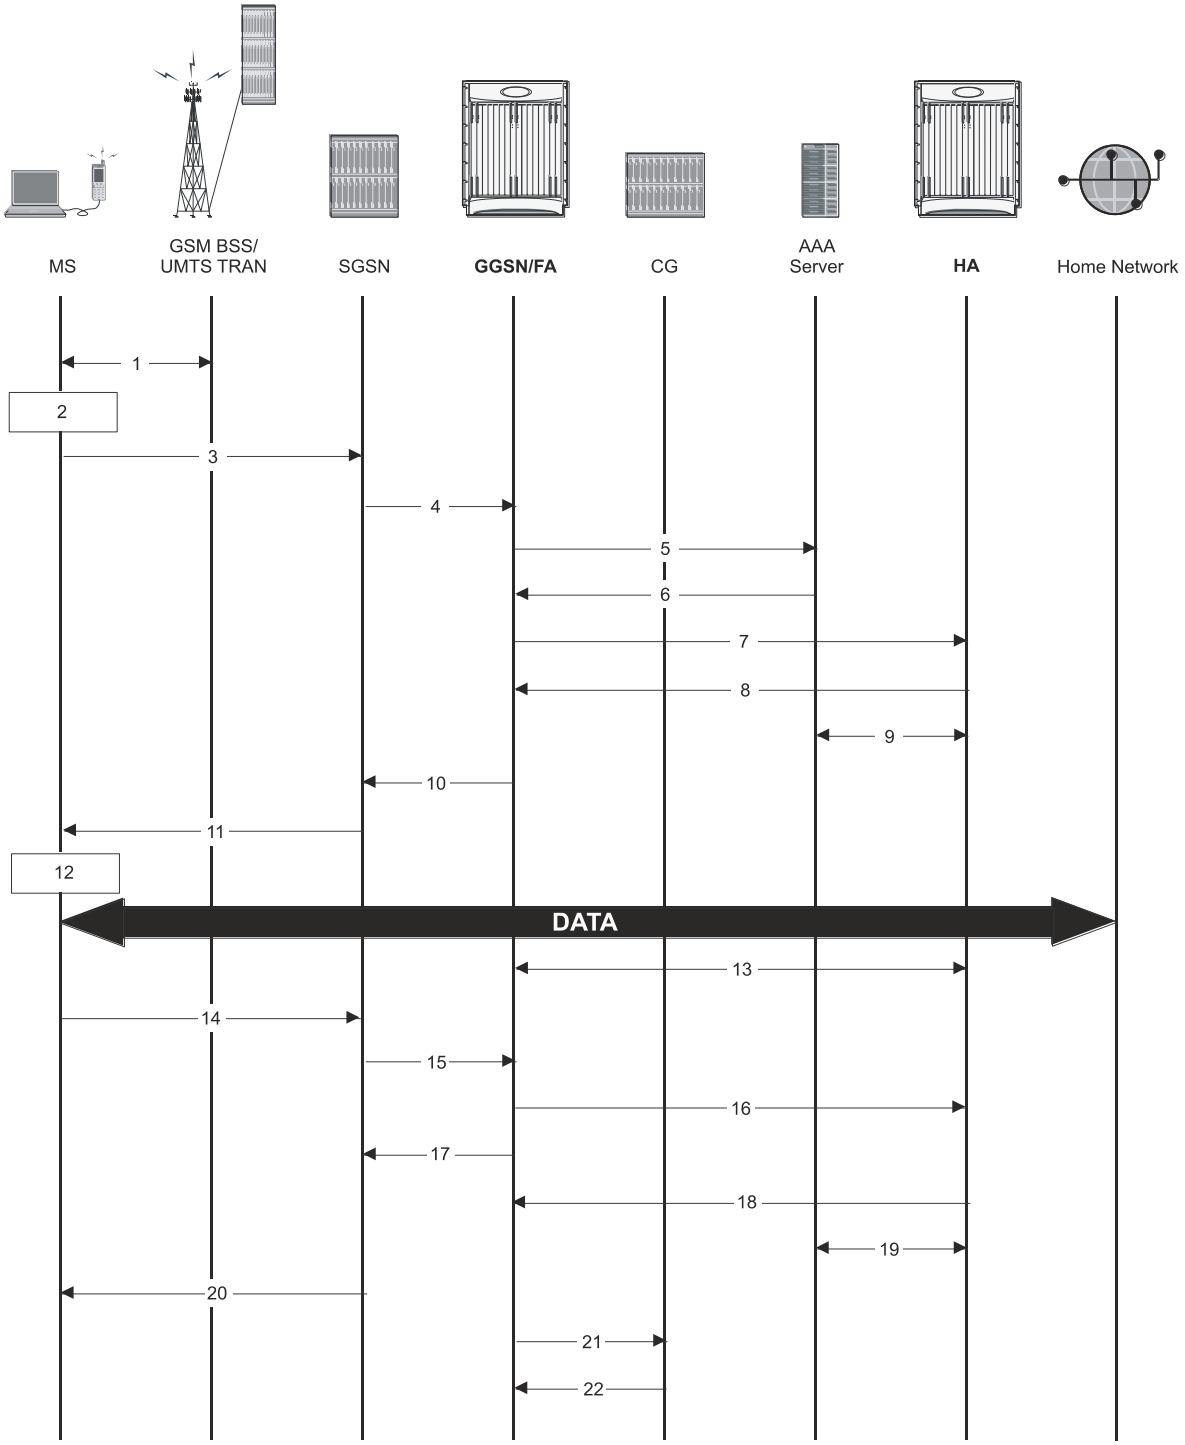
13. The HA sends a registration response to the FA containing the address assigned to the MS.
14. The FA extracts the home address assigned to the MS by the HA from the response and the GGSN updates the associated PDP context. The FA then forwards it to the MS (identified by either the home address or the NAI and TEID).
15. The GGSN issues a PDP context modification procedure to the SGSN in order to update the PDP address for the MS.
16. The SGSN forwards the PDP context modification message to the MS.

    The MS can now send and receive data to or from their home network until the session is closed or times out. Note that for Mobile IP, only one PDP context is supported for the MS.
17. The MS can terminate the Mobile IP data session at any time. To terminate the Mobile IP session, the MS sends a Registration Request message to the GGSN/FA with a requested lifetime of 0.
18. The FA component forwards the request to the HA.
19. The HA sends a Registration Reply to the FA accepting the request.
20. The GGSN/FA forwards the response to the MN.
21. The MS sends a Deactivate PDP Context Request message that is received by the SGSN.
22. The SGSN sends a Delete PDP Context Request message to the GGSN facilitating the data session. The message includes the information elements necessary to identify the PDP context.
23. The GGSN removes the PDP context from memory and returns a Delete PDP Context Response message to the SGSN.
24. The SGSN returns a Deactivate PDP Context Accept message to the MS.
25. The GGSN delivers the GGSN Charging Detail Records (G-CDRs) to a Charging Gateway (CG) using GTP Prime (GTPP). Note that, though not shown in this example, the GGSN could optionally be configured to send partial CDRs while the PDP context is active.
26. For each accounting message received from the GGSN, the CG responds with an acknowledgement.

Proxy Mobile IP Call Flows

The following figure and the text that follows describe a sample successful Proxy Mobile IP session setup call flow in which the MS receives its IP address from the HA.

Figure 111. HA Assigned IP Address Proxy Mobile IP Call Flow



1. The Mobile Station (MS) goes through the process of attaching itself to the GPRS/UMTS network.
2. The Terminal Equipment (TE) aspect of the MS sends AT commands to the Mobile Terminal (MT) aspect of the MS to place it into PPP mode.

The Link Control Protocol (LCP is then used to configure the Maximum-Receive Unit size and the authentication protocol (Challenge-Handshake Authentication Protocol (CHAP), Password Authentication Protocol (PAP), or none). If CHAP or PAP is used, the TE will authenticate itself to the MT, which, in turn, stores the authentication information.

Upon successful authentication, the TE sends an Internet Protocol Control Protocol (IPCP) Configure-Request message to the MT. The message will either contain a static IP address to use or request that one be dynamically assigned.

3. The MS sends an Activate PDP Context Request message that is received by an SGSN. The message contains information about the subscriber such as the Network layer Service Access Point Identifier (NSAPI), PDP Type, PDP Address, Access Point Name (APN), Quality of Service (QoS) requested, and PDP configuration options.
4. The SGSN authenticates the request message and sends a Create PDP Context Request message to a GGSN using the GPRS Tunneling Protocol (GTPC, "C" indicates the control signaling aspect of the protocol). The recipient GGSN is selected based on either the request of the MS or is automatically selected by the SGSN. The message consists of various information elements including: PDP Type, PDP Address, APN, charging characteristics, and Tunnel Endpoint Identifier (TEID, if the PDP Address was static).
5. The GGSN determines if it can facilitate the session (in terms of memory or CPU resources, configuration, etc.) and creates a new entry in its PDP context list and provides a Charging ID for the session.

   From the APN specified in the message, the GGSN determines whether or not the subscriber is to be authenticated, if Proxy Mobile IP is to be supported for the subscriber, and if so, the IP address of the HA to contact.

   Note that Proxy Mobile IP support can also be determined by attributes in the user's profile. Attributes in the user's profile supersede APN settings.

   If authentication is required, the GGSN attempts to authenticate the subscriber locally against profiles stored in memory or send a RADIUS Access-Request message to an AAA server.
6. If the GGSN authenticated the subscriber to an AAA server, the AAA server responds with a RADIUS Access-Accept message indicating successful authentication and any attributes for handling the subscriber PDP context.
7. If Proxy Mobile IP support was either enabled in the APN or in the subscriber's profile, the GGSN/FA forwards a Proxy Mobile IP Registration Request message to the specified HA. The message includes such things as the MS's home address, the IP address of the FA (the care-of-address), and the FA-HA extension (Security Parameter Index (SPI)).
8. The HA responds with a Proxy Mobile IP Registration Response. The response includes an IP address from one of its locally configured pools to assign to the MS (its Home Address). The HA also creates a Mobile Binding Record (MBR) for the subscriber session.
9. The HA sends a RADIUS Accounting Start request to the AAA server which the AAA server responds to.
10. The GGSN replies with an affirmative Create PDP Context Response using GTPC. The response will contain information elements such as the PDP Address representing either the static address requested by the MS or the address assigned by the GGSN, the TEID used to reference PDP Address, and PDP configuration options specified by the GGSN.
11. The SGSN returns an Activate PDP Context Accept message to the MS. The message includes response to the configuration parameters sent in the initial request.
12. The MT, will respond to the TE's IPCP Config-request with an IPCP Config-Ack message.

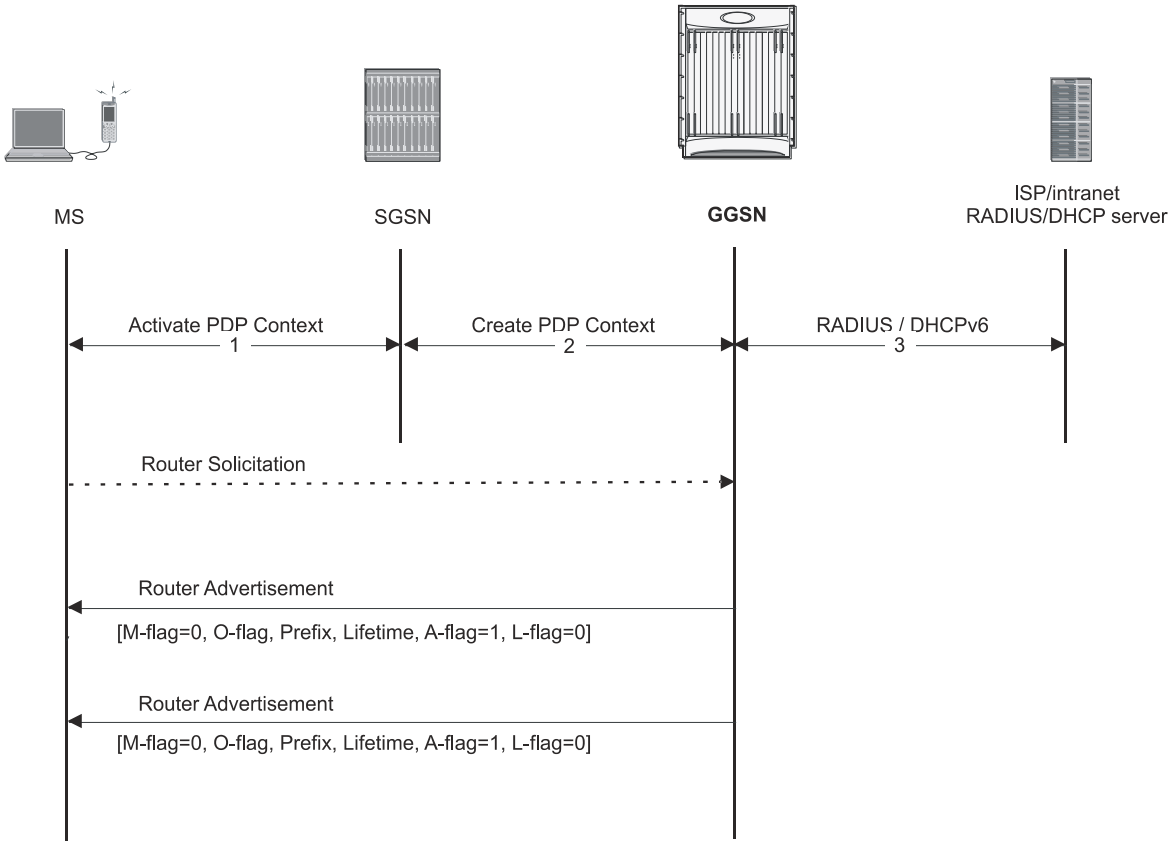
    The MS can now send and receive data to or from the PDN until the session is closed or times out. Note that for Mobile IP, only one PDP context is supported for the MS.

13. The FA periodically sends Proxy Mobile IP Registration Request Renewal messages to the HA. The HA sends responses for each request.
14. The MS can terminate the data session at any time. To terminate the session, the MS sends a Deactivate PDP Context Request message that is received by the SGSN.
15. The SGSN sends a Delete PDP Context Request message to the GGSN facilitating the data session. The message includes the information elements necessary to identify the PDP context (i.e., TEID, and NSAPI).
16. The GGSN removes the PDP context from memory and the FA sends a Proxy Mobile IP Deregistration Request message to the HA.
17. The GGSN returns a Delete PDP Context Response message to the SGSN.
18. The HA replies to the FA with a Proxy Mobile IP Deregistration Request Response.
19. The HA sends a RADIUS Accounting Stop request to the AAA server which the AAA server responds to.
20. The SGSN returns a Deactivate PDP Context Accept message to the MS.
21. The GGSN delivers the GGSN Charging Detail Records (G-CDRs) to a Charging Gateway (CG) using GTP Prime (GTPP). Note that, though not shown in this example, the GGSN could optionally be configured to send partial CDRs while the PDP context is active.
22. For each accounting message received from the GGSN, the CG responds with an acknowledgement.

IPv6 Stateless Address Auto Configuration Flows

The following figure and the text that follows describe a sample IPv6 stateless address auto configuration session setup call flow in which the MS receives its IP address from the RADIUS DHCP server.

Figure 112. IPv6 Stateless Address Auto Configuration Flow



1. The MS uses the IPv6 interface identifier provided by the GGSN to create its IPv6 link-local unicast address. Before the MS communicates with other hosts or mobile stations on the intranet/ISP, the MS must obtain an IPv6 global or site-local unicast address.
2. After the GGSN sends a create PDP context response message to the SGSN, it starts sending router advertisements periodically on the new MS-GGSN link established by the PDP context.
3. When creating a global or site-local unicast address, the MS may use the interface identifier received during the PDP context activation or it generates a new interface identifier. There is no restriction on the value of the interface identifier of the global or site-local unicast address, since the prefix is unique.

Supported Standards

The GGSN complies with the following standards for 3GPP wireless data services.

- 3GPP References
- IETF References
- Object Management Group (OMG) Standards

3GPP References

- 3GPP TS 09.60 v7.10.0 (2001-09): 3rd Generation Partnership project; Technical Specification Group Core Network; General Packet Radio Services (GPRS); GPRS Tunneling Protocol (GTP) across the Gn and Gp Interface (Release 1998) for backward compatibility with GTPv0
- 3GPP TS 23.060 v7.6.0 (2007-9): 3rd Generation Partnership project; Technical Specification Group Services and System Aspects; General Packet Radio Service (GPRS); Service description (Release 1999) as an additional reference for GPRS/UMTS procedures
- 3GPP TS 23.107 v7.1.0 (2007-09): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; QoS Concept and Architecture
- 3GPP TS 23.203 V7.7.0 (2006-08): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Policy and charging control architecture (Release 7)
- 3GPP TS 23.246 v7.4.0 (2007-09): 3rd Generation Partnership Project, Technical Specification Group Services and System Aspects; Multimedia Broadcast/Multicast Service (MBMS); Architecture and functional description (Release 7)
- 3GPP TS 24.008 v7.11.0 (2001-06): Mobile radio interface layer 3 specification; Core Network Protocols - Stage 3 (Release 1999) as an additional reference for GPRS/UMTS procedures
- 3GPP TS 29.060 v7.9.0 (2008-09): 3rd Generation Partnership Project; Technical Specification Group Core Network; General Packet Radio Services (GPRS); GRPS Tunneling Protocol (GTP) across the Gn and Gp Interface (Release 4) for the Core GTP Functionality
- 3GPP TS 29.061 v7.7.0 (2008-09): 3rd Generation Partnership Project; Technical Specification Group Core Network; Packet Domain; Interworking between the Public Land Mobile Network (PLMN) supporting Packet Based Services and Packet Data Networks (PDN)
- 3GPP 29.212 v7.6.0 (2008-09) 3rd Generation Partnership Project, Technical Specification Group Core Network and Terminals; Policy and Charging Control over Gx reference point (Release 7)
- 3GPP TS 29.213 V7.5.0 (2005-08): 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Policy and Charging Control signalling flows and QoS parameter mapping; (Release 7)
- 3GPP TR 29.846 6.0.0 (2004-09) 3rd Generation Partnership Project, Technical Specification Group Core Networks; Multimedia Broadcast/Multicast Service (MBMS); CN1 procedure description (Release 6)
- 3GPP TS 32.015 v3.12.0 (2003-12): 3rd Generation Partnership Project, Technical Specification Group Services and System Aspects; Telecommunication Management; Charging management; Call and event data for the Packet Switched (PS) domain (Release 1999) for support of Charging on GGSN

- 3GPP TS 32.215 v5.9.0 (2005-06): 3rd Generation Partnership Project, Technical Specification Group Services and System Aspects; Telecommunication Management; Charging Management; Charging data description for the Packet Switched (PS) domain (Release 5)
- 3GPP 32.251 v7.5.1 (2007-10) 3rd Generation Partnership Project, Technical Specification Group Services and System Aspects; Telecommunication management; Charging management; Packet Switched (PS) domain charging (Release 7)
- 3GPP TS 32.298 v7.4.0 (2007-09): 3rd Generation Partnership Project; Technical Specification Group Service and System Aspects; Telecommunication management; Charging management; Charging Data Record (CDR) parameter description
- 3GPP TS 32.299 v7.7.0 (2007-10): 3rd Generation Partnership Project; Technical Specification Group Service and System Aspects; Telecommunication management; Charging management; Diameter charging applications (Release 7)
- 3GPP TS 32.403 V7.1.0: Technical Specification Performance measurements - UMTS and combined UMTS/GSM
- 3GPP TS 33.106 V7.0.1 (2001-06): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G security; Lawful Interception requirements (Release 7)
- 3GPP TS 33.107 V7.7.0 (2007-09): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G security; Lawful interception architecture and functions (Release 7)

IETF References

- RFC-768, User Datagram Protocol (UPD), August 1980
- RFC-791, Internet Protocol (IP), September 1982
- RFC-793, Transmission Control Protocol (TCP), September 1981
- RFC-894, A Standard for the Transmission of IP Datagrams over Ethernet Networks, April 1984
- RFC-1089, SNMP over Ethernet, February 1989
- RFC-1144, Compressing TCP/IP headers for low-speed serial links, February 1990
- RFC-1155, Structure & identification of management information for TCP/IP-based Internets, May 1990
- RFC-1157, Simple Network Management Protocol (SNMP) Version 1, May 1990
- RFC-1212, Concise MIB Definitions, March 1991
- RFC-1213, Management Information Base for Network Management of TCP/IP-based Internets: MIB-II, March 1991
- RFC-1215, A Convention for Defining Traps for use with the SNMP, March 1991
- RFC-1224, Techniques for managing asynchronously generated alerts, May 1991
- RFC-1256, ICMP Router Discovery Messages, September 1991
- RFC-1305, Network Time Protocol (Version 3) Specification, Implementation and Analysis, March 1992
- RFC-1332, The PPP Internet Protocol Control Protocol (IPCP), May 1992
- RFC-1398, Definitions of Managed Objects for the Ethernet-Like Interface Types, January 1993
- RFC-1418, SNMP over OSI, March 1993
- RFC-1570, PPP LCP Extensions, January 1994
- RFC-1643, Definitions of Managed Objects for the Ethernet-like Interface Types, July 1994

- RFC-1661, The Point to Point Protocol (PPP), July 1994
- RFC-1662, PPP in HDLC-like Framing, July 1994
- RFC-1701, Generic Routing Encapsulation (GRE), October 1994
- RFC-1850, OSPF Version 2 Management Information Base, November 1995
- RFC-1901, Introduction to Community-based SNMPv2, January 1996
- RFC-1902, Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1903, Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1904, Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1905, Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1906, Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1907, Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1908, Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework, January 1996
- RFC-1918, Address Allocation for Private Internets, February 1996
- RFC-1919, Classical versus Transparent IP Proxies, March 1996
- RFC-1962, The PPP Compression Control Protocol (CCP), June 1996
- RFC-1974, PPP STAC LZS Compression Protocol, August 1996
- RFC-2002, IP Mobility Support, May 1995
- RFC-2003, IP Encapsulation within IP, October 1996
- RFC-2004, Minimal Encapsulation within IP, October 1996
- RFC-2005, Applicability Statement for IP Mobility Support, October 1996
- RFC-2118, Microsoft Point-to-Point Compression (MPPC) Protocol, March 1997
- RFC 2131, Dynamic Host Configuration Protocol
- RFC 2132, DHCP Options and BOOTP Vendor Extensions
- RFC-2136, Dynamic Updates in the Domain Name System (DNS UPDATE)
- RFC-2211, Specification of the Controlled-Load Network Element Service
- RFC-2246, The Transport Layer Security (TLS) Protocol Version 1.0, January 1999
- RFC-2290, Mobile-IPv4 Configuration Option for PPP IPCP, February 1998
- RFC-2328, OSPF Version 2, April 1998
- RFC-2344, Reverse Tunneling for Mobile IP, May 1998
- RFC-2394, IP Payload Compression Using DEFLATE, December 1998
- RFC 2401, Security Architecture for the Internet Protocol
- RFC 2402, IP Authentication Header (AH)

- RFC 2406, IP Encapsulating Security Payload (ESP)
- RFC 2409, The Internet Key Exchange (IKE)
- RFC-2460, Internet Protocol Version 6 (IPv6)
- RFC-2461, Neighbor Discovery for IPv6
- RFC-2462, IPv6 Stateless Address Autoconfiguration
- RFC-2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, December 1998
- RFC-2475, An Architecture for Differentiated Services, December 1998
- RFC-2484, PPP LCP Internationalization Configuration Option, January 1999
- RFC-2486, The Network Access Identifier (NAI), January 1999
- RFC-2571, An Architecture for Describing SNMP Management Frameworks, April 1999
- RFC-2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP), April 1999
- RFC-2573, SNMP Applications, April 1999
- RFC-2574, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3), April 1999
- RFC-2597, Assured Forwarding PHB Group, June 1999
- RFC-2598, Expedited Forwarding PHB, June 1999
- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2661, Layer Two Tunneling Protocol "L2TP", August 1999
- RFC-2697, A Single Rate Three Color Marker, September 1999
- RFC-2698, A Two Rate Three Color Marker, September 1999
- RFC-2784, Generic Routing Encapsulation (GRE) - March 2000, IETF
- RFC-2794, Mobile IP Network Access Identifier Extension for IPv4, March 2000
- RFC-2809, Implementation of L2TP Compulsory Tunneling via RADIUS, April 2000
- RFC-2845, Secret Key Transaction Authentication for DNS (TSIG), May 2000
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000
- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000
- RFC-3007, Secure Domain Name System (DNS) Dynamic Update, November 2000
- RFC-3012, Mobile IPv4 Challenge/Response Extensions, November 2000
- RFC-3056, Connection of IPv6 Domains via IPv4 Clouds, February 2001
- RFC-3101 OSPF-NSSA Option, January 2003
- RFC-3143, Known HTTP Proxy/Caching Problems, June 2001
- RFC-3193, Securing L2TP using IPSEC, November 2001

- RFC-3314, Recommendations for IPv6 in Third Generation Partnership Project (3GPP) Standards, September 2002
- RFC-3316, Internet Protocol Version 6 (IPv6) for Some Second and Third Generation Cellular Hosts, April 2003
- RFC-3706, A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers, February 2004
- RFC-3543, Registration Revocation in Mobile IPv4, August 2003
- RFC 3588, Diameter Base Protocol, September 2003
- RFC 4006, Diameter Credit-Control Application, August 2005
- Draft, Route Optimization in Mobile IP
- Draft, Generalized Key Distribution Extensions for Mobile IP
- Draft, AAA Keys for Mobile IP

Object Management Group (OMG) Standards

CORBA 2.6 Specification 01-09-35, Object Management Group

Chapter 16
HA Overview

The Home Agent (HA) allows mobile nodes to be reached, or served, by their home network through its home address even when the mobile node is not attached to its home network. The HA performs this function through interaction with a Foreign Agent (FA) that the mobile node is communicating with using the Mobile IP (MIP) standard. Such transactions are performed through the use of virtual private networks that create MIP tunnels between the HA and FA.

When functioning as an HA, the system can either be located within the carrier's 3G network or in an external enterprise or ISP network. Regardless, the FA terminates the mobile subscriber's PPP session, and then routes data to and from the appropriate HA on behalf of the subscriber.

This chapter includes the following sections:

- Product Specifications
- Network Deployment Configurations
- Understanding Mobile IP

Product Specifications

The following application and line cards are required to support CDMA2000 wireless data services on the system:

Hardware Requirements

Platforms

The Home Agent service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the Installation Guide for the chassis and/or contact your Cisco account representative.

Components

The following application and line cards are required to support HA functionality on an ASR 5x00 platform:

- **System Management Cards (SMCs):** Provides full system control and management of all cards within the ASR 5x00 platform. Up to two SMC can be installed; one active, one redundant.
- **Packet Processing Cards (PSC, PSC2, PPC):** Within the ASR 5x00 platform, packet processing cards provide high-speed, multi-threaded PPP processing capabilities to support HA services. Up to 14 packet processing cards can be installed, allowing for multiple active and/or redundant cards.
- **Switch Processor Input/Outputs (SPIO):** Installed in the upper-rear chassis slots directly behind the SMCs, SPIOs provide connectivity for local and remote management, Central Office (CO) alarms. Up to two SPIOs can be installed; one active, one redundant.
- **Ethernet 10/100 and/or Ethernet 1000/Quad Ethernet 1000 Line Cards:** Installed directly behind processing cards, these cards provide the RP, AAA, PDN, and Pi interfaces to elements in the data network. Up to 26 line cards should be installed for a fully loaded system with 13 active processing cards, 13 in the upper-rear slots and 13 in the lower-rear slots for redundancy. Redundant processing cards do no not require line cards.
- **Redundancy Crossbar Cards (RCCs):** Installed in the lower-rear chassis slots directly behind the SMCs, RCCs utilize 5 Gbps serial links to ensure connectivity between Ethernet 10/100 or Ethernet 1000/Quad Ethernet 1000 line cards and every processing card in the system for redundancy. Two RCCs can be installed to provide redundancy for all line cards and processing cards.

Important: Additional information pertaining to each of the application and line cards required to support CDMA2000 wireless data services is located in the Product Overview Guide.

Operating System Requirements

The HA is available for all Cisco ASR 5x00 platforms running StarOS Release 10.0 or later.

MPLS Forwarding with LDP

Multi Protocol Label Switching (MPLS) is an operating scheme or a mechanism that is used to speed up the flow of traffic on a network by making better use of available network paths. It works with the routing protocols like BGP and OSPF and therefore it is not a routing protocol.

It generates a fixed-length label to attach or bind with the IP packet's header to control the flow and destination of data. The binding of the labels to the IP packets is done by the label distribution protocol (LDP). All the packets in a forwarding equivalence class (FEC) are forwarded by a label-switching router (LSR) which is also called an MPLS node. The LSR uses the LDP in order to signal its forwarding neighbors and distribute its labels for establishing a labelswitching path (LSP).

In order to support the increasing number of corporate APNs which have a number of different addressing models and requirements, MPLS is deployed to fulfill at least following two requirements:

- The corporate APN traffic must remain segregated from other APNs for security reasons.
- Overlapping of IP addresses in different APNs.

When deployed, MPLS backbone automatically negotiates the routes using the labels binded with the IP packets. Cisco GGSN as an LSR learns the default route from the connected provider edge (PE) while the PE populates its routing table with the routes provided by the GGSN.

Supported Standards

The system supports the following industry standards for 1x/CDMA2000/EV-DO devices.

Requests for Comments (RFCs)

- RFC-768, User Datagram Protocol (UPD), August 1980
- RFC-791, Internet Protocol (IP), September 1982
- RFC-793, Transmission Control Protocol (TCP), September 1981
- RFC-894, A Standard for the Transmission of IP Datagrams over Ethernet Networks, April 1984
- RFC-1089, SNMP over Ethernet, February 1989
- RFC-1144, Compressing TCP/IP headers for low-speed serial links, February 1990
- RFC-1155, Structure and Identification of Management Information for TCP/IP-based Internets, May 1990
- RFC-1157, Simple Network Management Protocol (SNMP) Version 1, May 1990
- RFC-1212, Concise MIB Definitions, March 1991
- RFC-1213, Management Information Base for Network Management of TCP/IP-based Internets: MIB-II, March 1991
- RFC-1215, A Convention for Defining Traps for use with the SNMP, March 1991
- RFC-1224, Techniques for Managing Asynchronously Generated Alerts, May 1991
- RFC-1256, ICMP Router Discovery Messages, September 1991
- RFC-1305, Network Time Protocol (Version 3) Specification, Implementation and Analysis, March 1992
- RFC-1332, The PPP Internet Protocol Control Protocol (IPCP), May 1992
- RFC-1398, Definitions of Managed Objects for the Ethernet-Like Interface Types, January 1993
- RFC-1418, SNMP over OSI, March 1993
- RFC-1570, PPP LCP Extensions, January 1994
- RFC-1643, Definitions of Managed Objects for the Ethernet-like Interface Types, July 1994
- RFC-1661, The Point to Point Protocol (PPP), July 1994
- RFC-1662, PPP in HDLC-like Framing, July 1994
- RFC-1701, Generic Routing Encapsulation (GRE), October 1994
- RFC-1771, A Border Gateway Protocol 4 (BGP-4)
- RFC-1850, OSPF Version 2 Management Information Base, November 1995
- RFC-1901, Introduction to Community-based SNMPv2, January 1996
- RFC-1902, Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1903, Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996

- RFC-1904, Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1905, Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1906, Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1907, Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1908, Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework, January 1996
- RFC-1918, Address Allocation for Private Internets, February 1996
- RFC-1919, Classical versus Transparent IP Proxies, March 1996
- RFC-1962, The PPP Compression Control Protocol (CCP), June 1996
- RFC-1974, PPP STAC LZS Compression Protocol, August 1996
- RFC-2002, IP Mobility Support, May 1995
- RFC-2003, IP Encapsulation within IP, October 1996
- RFC-2004, Minimal Encapsulation within IP, October 1996
- RFC-2005, Applicability Statement for IP Mobility Support, October 1996
- RFC-2118, Microsoft Point-to-Point Compression (MPPC) Protocol, March 1997
- RFC-2136, Dynamic Updates in the Domain Name System (DNS UPDATE)
- RFC-2211, Specification of the Controlled-Load Network Element Service
- RFC-2246, The Transport Layer Security (TLS) Protocol Version 1.0, January 1999
- RFC-2290, Mobile IPv4 Configuration Option for PPP IPCP, February 1998
- RFC-2328, OSPF Version 2, April 1998
- RFC-2344, Reverse Tunneling for Mobile IP, May 1998
- RFC-2394, IP Payload Compression Using DEFLATE, December 1998
- RFC-2401, Security Architecture for the Internet Protocol, November 1998
- RFC-2402, IP Authentication Header (AH), November 1998
- RFC-2406, IP Encapsulating Security Payload (ESP), November 1998
- RFC-2408, Internet Security Association and Key Management Protocol (ISAKMP), November 1998
- RFC-2409, The Internet Key Exchange (IKE), November 1998
- RFC-2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, December 1998
- RFC-2475, An Architecture for Differentiated Services, December 1998
- RFC-2484, PPP LCP Internationalization Configuration Option, January 1999
- RFC-2486, The Network Access Identifier (NAI), January 1999
- RFC-2571, An Architecture for Describing SNMP Management Frameworks, April 1999
- RFC-2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP), April 1999

- RFC-2573, SNMP Applications, April 1999
- RFC-2574, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3), April 1999
- RFC-2597, Assured Forwarding PHB Group, June 1999
- RFC2598 - Expedited Forwarding PHB, June 1999
- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2661, Layer Two Tunneling Protocol "L2TP", August 1999
- RFC-2697, A Single Rate Three Color Marker, September 1999
- RFC-2698, A Two Rate Three Color Marker, September 1999
- RFC-2784, Generic Routing Encapsulation (GRE) - March 2000, IETF
- RFC-2794, Mobile IP Network Access Identifier Extension for IPv4, March 2000
- RFC-2809, Implementation of L2TP Compulsory Tunneling via RADIUS, April 2000
- RFC-2845, Secret Key Transaction Authentication for DNS (TSIG), May 2000
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000
- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000
- RFC-3007, Secure Domain Name System (DNS) Dynamic Update, November 2000
- RFC-3012, Mobile IPv4 Challenge/Response Extensions, November 2000
- RFC-3095, Robust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP and uncompressed, July 2001
- RFC-3101, OSPF NSSA Option, January 2003.
- RFC-3141, CDMA2000 Wireless Data Requirements for AAA, June 2001
- RFC-3143, Known HTTP Proxy/Caching Problems, June 2001
- RFC-3193, Securing L2TP using IPSEC, November 2001
- RFC-3241 Robust Header Compression (ROHC) over PPP, April 2002
- RFC-3409, Lower Layer Guidelines for Robust (RTP/UDP/IP) Header Compression, December 2002
- RFC-3519, NAT Traversal for Mobile IP, April 2003
- RFC-3543, Registration Revocation in Mobile IPv4, August 2003
- RFC 3576 - Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS), July 2003
- RFC-3706, A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers, February 2004
- RFC-3759, Robust Header Compression (ROHC): Terminology and Channel Mapping Examples, April 2004
- RFC-3588, Diameter Based Protocol, September 2003
- RFC-4005, Diameter Network Access Server Application, August 2005
- RFC-4006, Diameter Credit-Control Application, August 2005

- Draft, Generalized Key Distribution Extensions for Mobile IP
- Draft, AAA Keys for Mobile IP

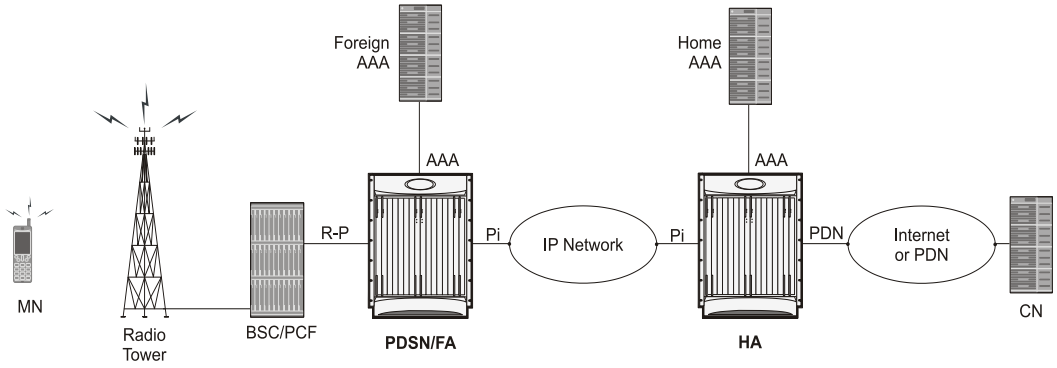
Network Deployment Configurations

This section provides examples of how the system can be deployed within a wireless carrier's network. As noted previously in this chapter, the system can be deployed in standalone configurations, serving as a Home Agent (HA) and a Packet Data Serving Node/Foreign Agent (PDSN/FA), or in a combined PDSN/FA/HA configuration providing all services from a single chassis.

Standalone PDSN/FA and HA Deployments

The following figure depicts a sample network configuration wherein the HA and the PDSN/FA are separate systems.

Figure 113. PDSN/FA and HA Network Deployment Configuration Example



The HA allows mobile nodes to be reached, or served, by their home network through its home address even when the mobile node is not attached to its home network. The HA performs this function through interaction with an FA that the mobile node is communicating with using the Mobile IP protocol. Such transactions are performed through the use of virtual private networks that create Mobile IP tunnels between the HA and FA.

Interface Descriptions

This section describes the primary interfaces used in a CDMA2000 wireless data network deployment.

Pi Interfaces

The Pi interface provides connectivity between the HA and its corresponding FA. The Pi interface is used to establish a Mobile IP tunnels between the PDSN/FA and HA.

PDN Interfaces

PDN interface provide connectivity between the PDSN and/or HA to packet data networks such as the Internet or a corporate intranet.

AAA Interfaces

Using the LAN ports located on the Switch Processor I/O (SPIO) and Ethernet line cards, these interfaces carry AAA messages to and from RADIUS accounting and authentication servers. The SPIO supports RADIUS-capable management interfaces using either copper or fiber Ethernet connectivity through two auto-sensing 10/100/1000 Mbps Ethernet interfaces or two SFP optical gigabit Ethernet interfaces. User-based RADIUS messaging is transported using the Ethernet line cards.

While most carriers will configure separate AAA interfaces to allow for out-of-band RADIUS messaging for system administrative users and other operations personnel, it is possible to use a single AAA interface hosted on the Ethernet line cards to support a single RADIUS server that supports both management users and network users.

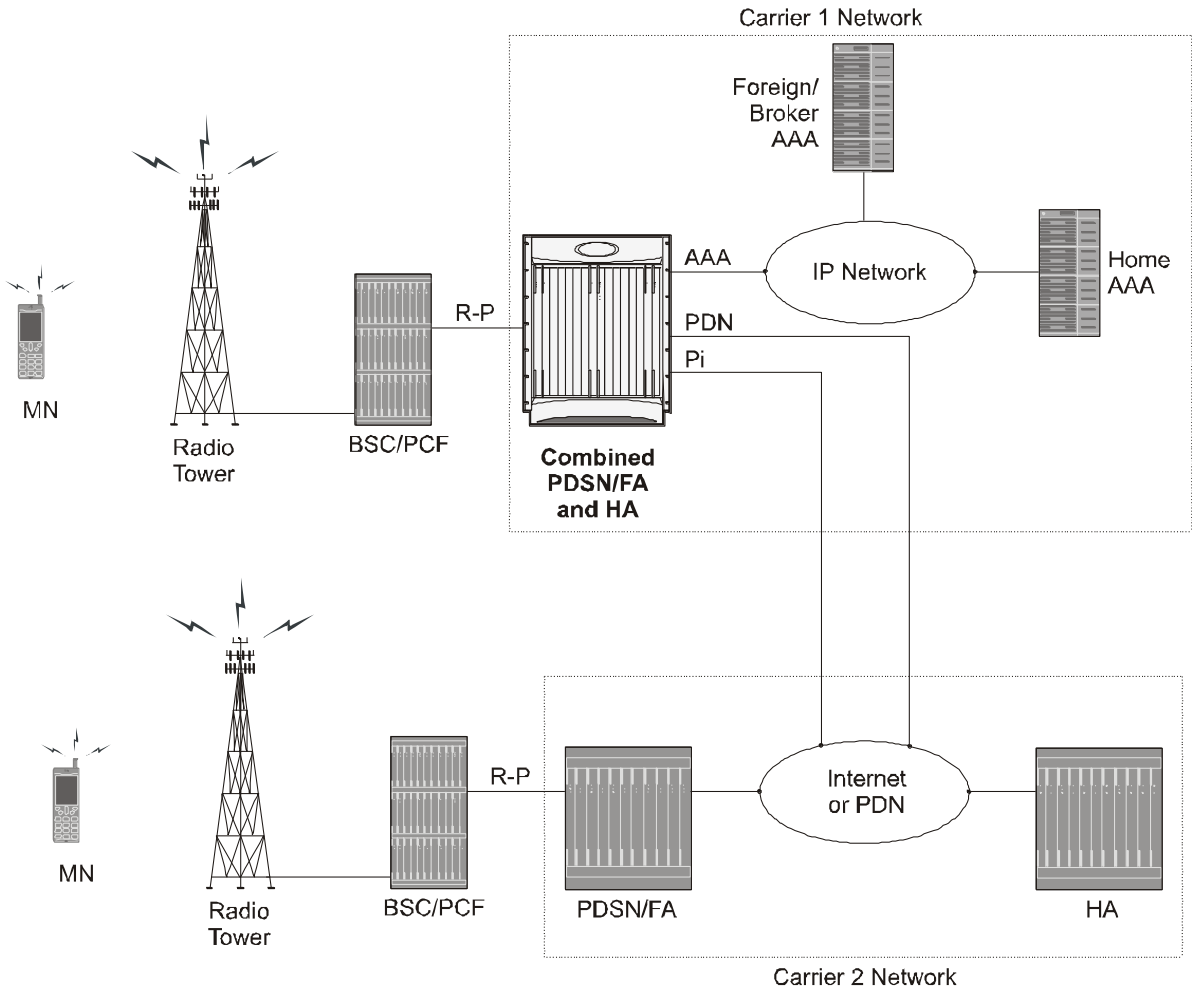
Important: Subscriber AAA interfaces should always be configured using Ethernet line card interfaces for the highest performance. The local context should not be used for service subscriber AAA functions.

Co-Located Deployments

An advantage of the system is its ability to support both high-density HA and PDSN/FA configurations within the same chassis. The economies of scale presented in this configuration example provide for both improved session handling and reduced cost in deploying a CDMA2000 data network.

The following figure depicts a sample co-located deployment.

Figure 114. Co-located PDSN/FA and HA Configuration Example



It should be noted that all interfaces defined within the 3GPP2 standards for 1x deployments exist in this configuration as they are described in the two previous sections. This configuration can support communications to external, or standalone, HAs and/or PDSNs/FAs using all prescribed standards.

Mobile IP Tunneling Methods

Tunneling by itself is a technology that enables one network to send its data via another network's connections. Tunneling works by encapsulating a network protocol within a packet, carried by the second network. Tunneling is also called encapsulation. Service providers typically use tunneling for two purposes; first, to transport otherwise un-routable packets across the IP network and second, to provide data separation for Virtual Private Networking (VPN) services. In Mobile IP, tunnels are used to transport data packets between the FA and HA.

The system supports the following tunneling protocols, as defined in the IS-835-A specification and the relevant Request For Comments (RFCs) for Mobile IP:

IP in IP tunnels

IP in IP tunnels basically encapsulate one IP packet within another using a simple encapsulation technique. To encapsulate an IP datagram using IP in IP encapsulation, an outer IP header is inserted before the datagram's existing IP header. Between them are other headers for the path, such as security headers specific to the tunnel configuration. Each header chains to the next using IP Protocol values. The outer IP header Source and Destination identify the "endpoints" of the tunnel. The inner IP header Source and Destination identify the original sender and recipient of the datagram, while the inner IP header is not changed by the encapsulator, except to decrement the TTL, and remains unchanged during its delivery to the tunnel exit point. No change to IP options in the inner header occurs during delivery of the encapsulated datagram through the tunnel. If needed, other protocol headers such as the IP Authentication header may be inserted between the outer IP header and the inner IP header.

The Mobile IP working group has specified the use of encapsulation as a way to deliver datagrams from an MN's HA to an FA, and conversely from an FA to an HA, that can deliver the data locally to the MN at its current location.

GRE tunnels

The Generic Routing Encapsulation (GRE) protocol performs encapsulation of IP packets for transport across disparate networks. One advantage of GRE over earlier tunneling protocols is that any transport protocol can be encapsulated in GRE. GRE is a simple, low overhead approach—the GRE protocol itself can be expressed in as few as eight octets as there is no authentication or tunnel configuration parameter negotiation. GRE is also known as IP Protocol 47.

Important: The chassis simultaneously supports GRE protocols with key in accordance with RFC-1701/RFC-2784 and "Legacy" GRE protocols without key in accordance to RFC-2002.

Another advantage of GRE tunneling over IP-in-IP tunneling is that GRE tunneling can be used even when conflicting addresses are in use across multiple contexts (for the tunneled data).

Communications between the FA and HA can be done in either the forward or reverse direction using the above protocols. Additionally, another method of routing information between the FA and various content servers used by the HA exists. This method is called Triangular Routing. Each of these methods is explained below.

Forward Tunneling

In the wireless IP world, forward tunneling is a tunnel that transports packets from the packet data network towards the MN. It starts at the HA and ends at the MN's care-of address. Tunnels can be as simple as IP-in-IP tunnels, GRE tunnels, or even IP Security (IPSec) tunnels with encryption. These tunnels can be started automatically, and are selected based on the subscriber's user profile.

Reverse Tunneling

A reverse tunnel starts at the MN's care-of address, which is the FA, and terminates at the HA.

When an MN arrives at a foreign network, it listens for agent advertisements and selects an FA that supports reverse tunnels. The MN requests this service when it registers through the selected FA. At this time, the MN may also specify a delivery technique such as Direct or the Encapsulating Delivery Style.

Using the Direct Delivery Style, which is the default mode for the system, the MN designates the FA as its default router and sends packets directly to the FA without encapsulation. The FA intercepts them, and tunnels them to the HA.

Using the Encapsulating Delivery Style, the MN encapsulates all its outgoing packets to the FA. The FA then de-encapsulates and re-tunnels them to the HA, using the FA's care-of address as the entry-point for this new tunnel.

Following are some of the advantages of reverse tunneling:

- All datagrams from the mobile node seem to originate from its home network

- The FA can keep track of the HA that the mobile node is registered to and tunnel all datagrams from the mobile node to its HA

Triangular Routing

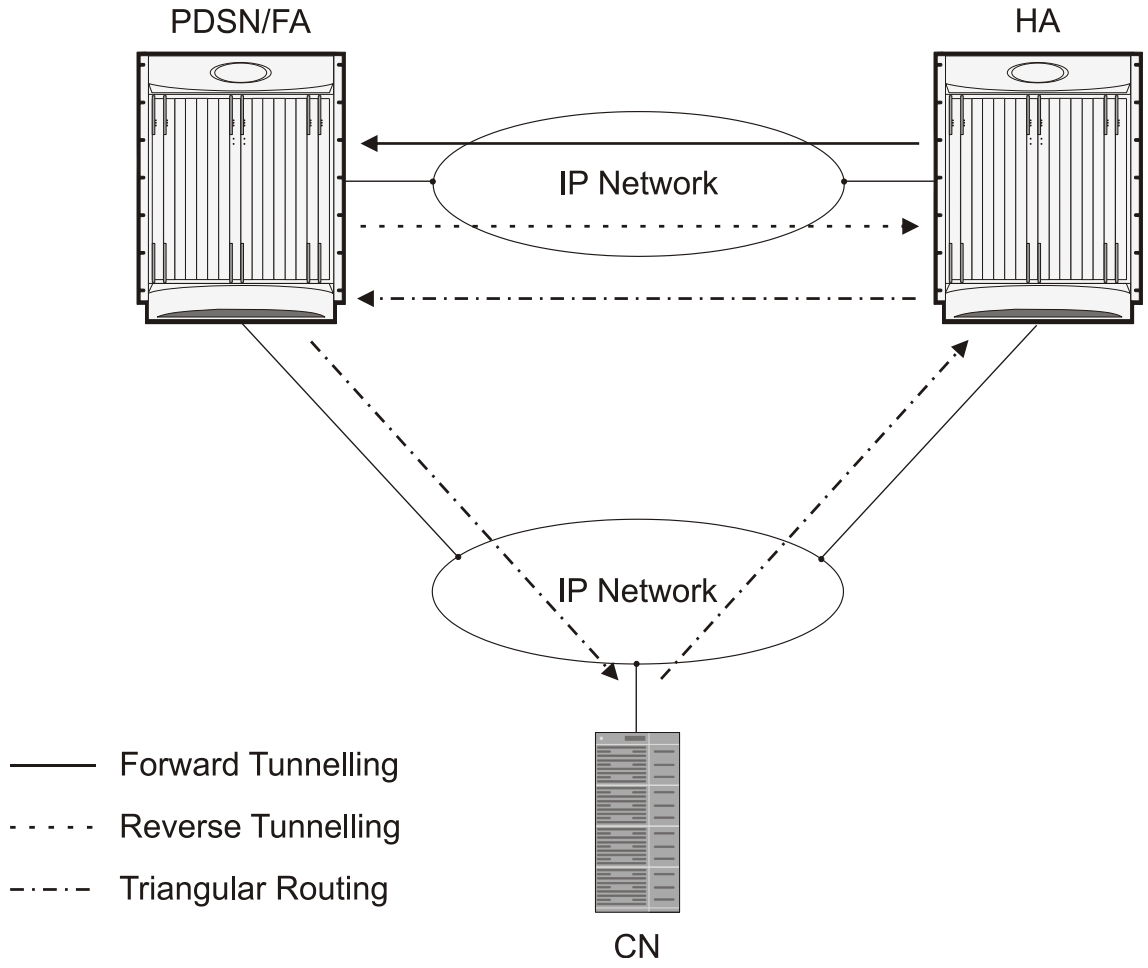
Triangular routing is the path followed by a packet from the MN to the Correspondent Node (CN) via the FA. In this routing scenario, the HA receives all the packets destined to the MN from the CN and redirects them to the MN's care-of-address by forward tunneling. In this case, the MN sends packets to the FA, which are transported using conventional IP routing methods.

A key advantage of triangular routing is that reverse tunneling is not required, eliminating the need to encapsulate and de-capsulate packets a second time during a Mobile IP session since only a forward tunnel exists between the HA and PDSN/FA.

A disadvantage of using triangular routing is that the HA is unaware of all user traffic for billing purposes. Also, both the HA and FA are required to be connected to a private network. This can be especially troublesome in large networks, serving numerous enterprise customers, as each FA would have to be connected to each private network.

The following figure shows an example of how triangular routing is performed.

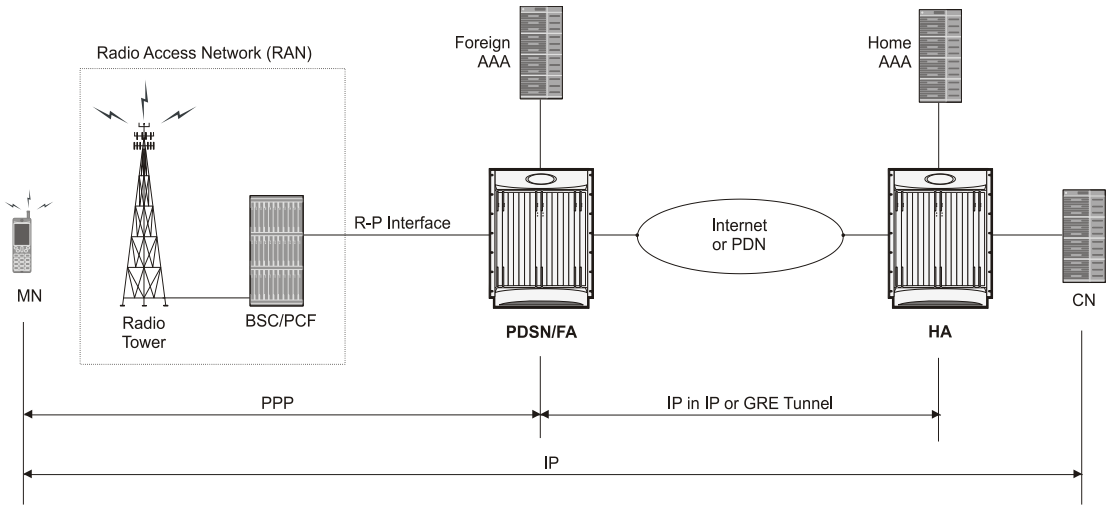
Figure 115. Mobile IP, FA and HA Tunneling/Transport Methods.



How Mobile IP Works

As described earlier, Mobile IP uses three basic communications protocols; PPP, IP, and Tunneled IP in the form of IP-in-IP or GRE tunnels. The following figure depicts where each of these protocols are used in a basic Mobile IP call.

Figure 116. Mobile IP Protocol Usage.



As depicted above, PPP is used to establish a communications session between the MN and the FA. Once a PPP session is established, the MN can communicate with the HA, using the FA as a mediator or broker. Data transport between the FA and HA use tunneled IP, either IP-in-IP or GRE tunneling. Communication between the HA and End Host can be achieved using the Internet or a private IP network and can use any IP protocol.

The following figure provides a high-level view of the steps required to make a Mobile IP call that is initiated by the MN to a HA. The following table explains each step in detail. Users should keep in mind that steps in the call flow related to the Radio Access Node (RAN) functions are intended to show a high-level overview of radio communications iterations, and as such are outside the scope of packet-based communications presented here.

Figure 117. Mobile IP Call Flow

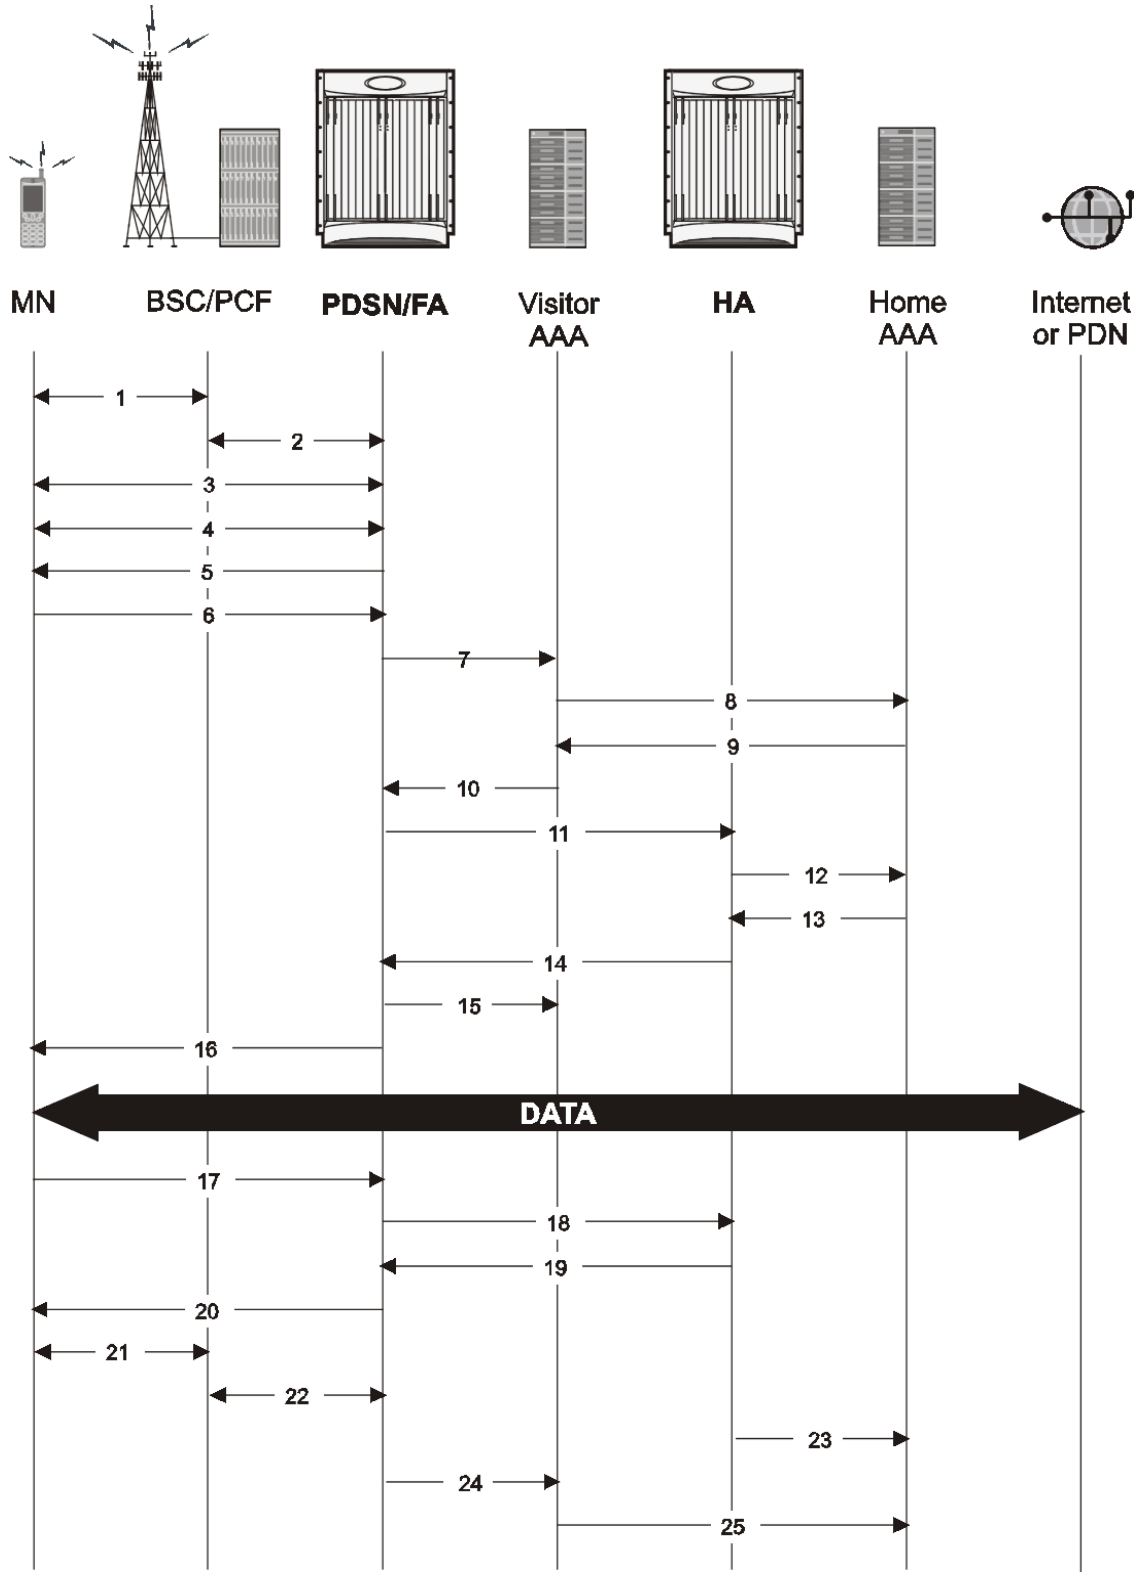


Table 62. Mobile IP Call Flow Description

Step	Description
1	Mobile Node (MN) secures a traffic channel over the airlink with the RAN through the BSC/PCF.
2	The PCF and PDSN establish the R-P interface for the session.
3	The PDSN and MN negotiate Link Control Protocol (LCP).
4	The PDSN and MN negotiate the Internet Protocol Control Protocol (IPCP).
5	The PDSN/FA sends an Agent Advertisement to the MN.
6	The MN sends a Mobile IP Registration Request to the PDSN/FA.
7	The PDSN/FA sends an Access Request message to the visitor AAA server.
8	The visitor AAA server proxies the request to the appropriate home AAA server.
9	The home AAA server sends an Access Accept message to the visitor AAA server.
10	The visitor AAA server forwards the response to the PDSN/FA.
11	Upon receipt of the response, the PDSN/FA forwards a Mobile IP Registration Request to the appropriate HA.
12	The HA sends an Access Request message to the home AAA server to authenticate the MN/subscriber.
13	The home AAA server returns an Access Accept message to the HA.
14	Upon receiving response from home AAA, the HA sends a reply to the PDSN/FA establishing a forward tunnel. Note that the reply includes a Home Address (an IP address) for the MN.
15	The PDSN/FA sends an Accounting Start message to the visitor AAA server. The visitor AAA server proxies messages to the home AAA server as needed.
16	The PDSN return a Mobile IP Registration Reply to the MN establishing the session allowing the MN to send/receive data to/from the PDN.
17	Upon session completion, the MN sends a Registration Request message to the PDSN/FA with a requested lifetime of 0.
18	The PDSN/FA forwards the request to the HA.
19	The HA sends a Registration Reply to the PDSN/FA accepting the request.
20	The PDSN/FA forwards the response to the MN.
21	The MN and PDSN/FA negotiate the termination of LCP effectively ending the PPP session.
22	The PCF and PDSN/FA close terminate the R-P session.
23	The HA sends an Accounting Stop message to the home AAA server.
24	The PDSN/FA sends an Accounting Stop message to the visitor AAA server.
25	The visitor AAA server proxies the accounting data to the home AAA server.

Understanding Mobile IP

Mobile IP provides a network-layer solution that allows Mobile Nodes (MNs, i.e. mobile phones, wireless PDAs, and other mobile devices) to receive routed IP packets from their home network while they are connected to any visitor network using their permanent or home IP address. Mobile IP allows mobility in a dynamic method that allows nodes to maintain ongoing communications while changing links as the user traverses the global Internet from various locations outside their home network.

In Mobile IP, the Mobile Node (MN) receives an IP address, either static or dynamic, called the "home address" assigned by its Home Agent (HA). A distinct advantage with Mobile IP is that MNs can hand off between different radio networks that are served by different PDSNs.

In this scenario, the Network Access Function (such as a PDSN) in the visitor network performs as a Foreign Agent (FA), establishing a virtual session with the MN's HA. Each time the MN registers with a different PDSN/FA, the FA assigns the MN a care-of-address. Packets are then encapsulated into IP tunnels and transported between FA, HA, and the MN.

Session Continuity Support for 3GPP2 and WiMAX Handoffs

HA provides this feature for seamless session mobility for WiMAX subscriber and other access technology subscribers as well. By implementation of this feature HA can be configured for:

- 3GPP2 HA Service
- 3GPP HA Service
- WiMAX HA Service
- Combination of 3GPP2 and WiMAX HA Services for Dual mode device

The above configurations provide the session continuity capability that enables a dual mode device (a multi radio device) to continue its active data session as it changes its active network attachment from 3GPP2 to Wimax and vice versa with no perceived user impacts from a user experience perspective. This capability brings the following benefits:

- common billing and customer care
- accessing home 3GPP2 service through Wimax network and vice versa
- better user experience with seamless session continuity

Chapter 17
HNB Gateway in Wireless Network

The Cisco® provides 3GPP wireless carriers with a flexible solution that functions as a Home NodeB Gateway (HNB-GW) in HNB Access Network to connect UEs with existing UMTS networks.

The Home NodeB Gateway works as a gateway for Home NodeBs (HNBs) to access the core networks. The HNB-GW concentrates connections from a large amount of HNBs through IuH interface and terminates the connection to existing Core Networks (CS or PS) using standard Iu (IuCS or IuPS) interface.

This overview provides general information about the HNB Gateway including:

- Product Description
- Network Deployment and Interfaces
- Features and Functionality - Base Software
- Features and Functionality - Optional Enhanced Feature Software
- How HNB-GW Works
- Supported Standards

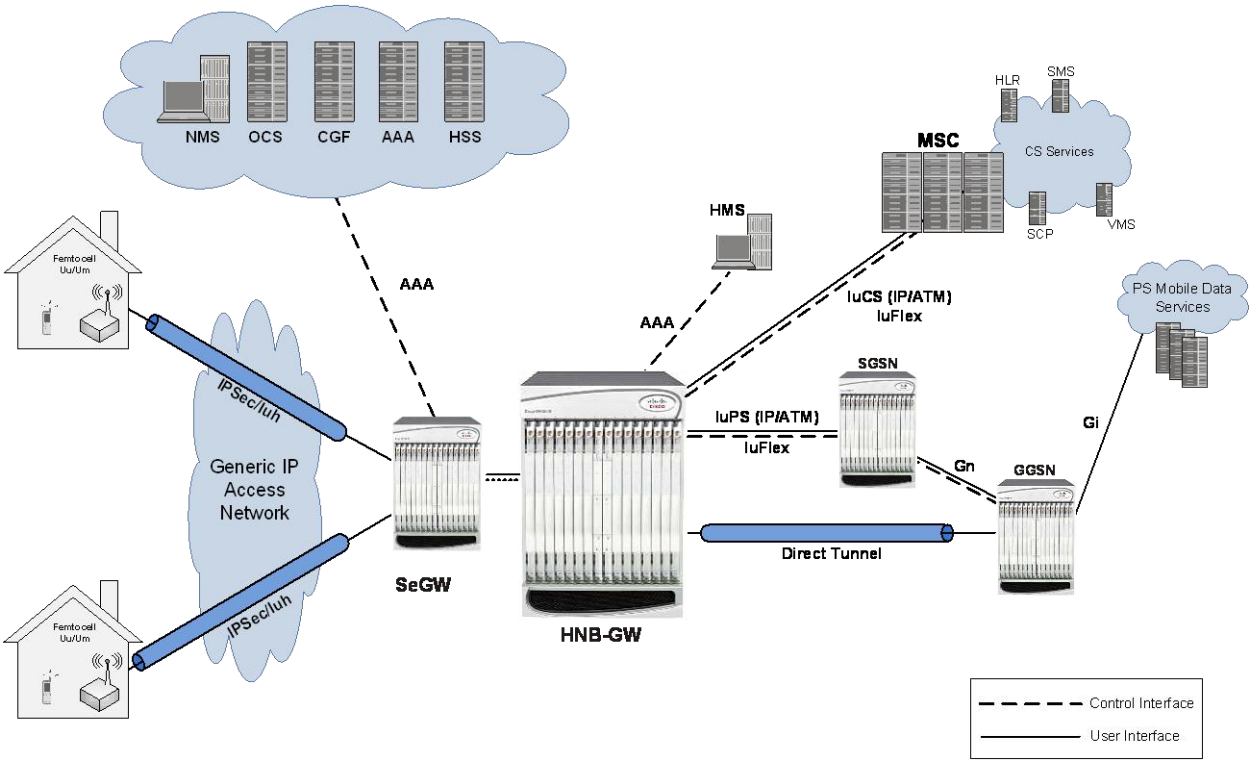
Product Description

The Home NodeB Gateway is the HNB network access concentrator used to connect the Home NodeBs (HNBs)/Femto Access Point (FAP) to access the UMTS network through HNB Access Network. It aggregates Home Node-B or Femto Access Points to a single network element and then integrates them into the Mobile Operators Voice, Data and Multimedia networks.

Femtocell is an important technology and service offering that enables new Home and Enterprise service capabilities for Mobile Operators and Converged Mobile Operators (xDSL/Cable/FFTH plus Wireless). The Femtocell network consists of a plug-n-play customer premise device generically called a Home NodeB (HNB) with limited range radio access in home or Enterprise. The HNB will auto-configure itself with the Operators network and the user can start making voice, data and multimedia calls.

The figure given describes a high level view of UMTS network with Femtocell and HNB-GW.

Figure 118. HNB-GW Deployment in 3G UMTS Network



Once a secure tunnel has been established between the HNB and the SeGW and the HNB has been configured by the HMS, the Operator has to connect the Femtocell network to their Core Network and services. There are several interworking approaches to Circuit Switch (CS) and Packet Switch (PS) domains. One approach is to make the Femtocell network appear as a standard Radio Access Network (RAN) to the Core Network. In addition to the HNB, SeGW and HMS the RAN approach requires a network element generically called a Femto Gateway (FGW/HNB-GW). The HNB-GW provides interworking and aggregation of large amount of Femtocell sessions toward standard CN interfaces (IuPS/IuCS). In this approach services and mobility are completely transparent to CN elements (e.g. MSC, xGSN).

The other approach is to connect the Femtocell to an IMS Network to provide CS services to subscribers when on the Femtocell and deploy a new network element generically called a Convergence Server to provide service continuity and mobility over standard interfaces at the MSC layer (e.g GSM-MAP, IS-41). These two approaches are clearly different in how CS based services and mobility are achieved.

In accordance with 3GPP standard, the HNB-GW provides following functions and procedures in UMTS core network:

- HNB Registration/De-registration Function
- UE Registration/De-registration Function for HNB
- IuH User-plane Management Functions
- IuH User-plan Transport Bearer Handling
- Iu Link Management Functions

Important: Some of the features may not be available in this release. Kindly contact your local Cisco representative for more information on supported features.

HNB Access Network Elements

This section provides the brief description and functionality of various network elements involved in the UMTS Femto access network. The HNB access network includes the following functional entities:

- Home NodeB
- Security Gateway (SeGW)
- HNB Gateway (HNB-GW)
- HNB Management System (HMS)

Home NodeB

A Home NodeB (HNB) is the a customer premise equipment that offers Uu interface to UE and IuH over IPSec tunnel to HNB-GW for accessing UMTS Core Network (PS or CS) in Femtocell access network.

It also provides the support to HNB registration and UE registration over IuH with HNB-GW. Apart from these functions HNB also supports some RNC like functions as given below:

- RAB management functions
- Radio Resource Management functions
- Iu Signalling Link management
- GTP-U Tunnels management
- Buffer Management
- Iu U-plane frame protocol initialization
- Mobility management functions
- Security Functions
- Service and Network Access functions
- Paging co-ordination functions
- UE Registration for HNB

- IuH user-plane Management functions

Security Gateway (SeGW)

Security Gateway is a logical entity in Cisco HNB-GW.

Basic function of this entity are:

- Authentication of HNB
- Providing access to HMS and HNB-GW

This entity terminates the secure tunnelling for IuH and TR-069 between HNB and HNB-GW and HMS respectively. In this implementation it is an optional element which is situated on HNB-GW.

HNB Gateway (HNB-GW)

The HNB-GW provides the access to Femto user to UMTS core network. It acts as an access gateway to HNB and concentrates connections from a large amount of HNBs. The IuH interface is used between HNB and HNB-GW and HNB-GW connects with the Core Networks (CS or PS) using the generic Iu (IuCS or IuPS) or Gn interface.

It also terminates Gn and other interfaces from UMTS core networks to provide mobile data services to HNB and to interact with HMS to perform HNB authentication and authorization.

HNB Management System (HMS)

It is a network element management system for HNB access. Management interface between HNB and HMS is based on TR-069 family of standards.

It performs following functions while managing HNB access network:

- Facilitates HNB-GW discovery for HNB
- Provision of configuration data to the HNB
- Performs location verification of HNB and assigns appropriate serving elements (HMS, Security Gateway and HNB-GW)

The HNB Management System (HMS) comprises of the following functional entities:

- File Server: used for file upload or download, as instructed by TR-069 manager
- TR-069 Manager: Performs CM, FM and PM functionality to the HNB through Auto-configuration server (HMS)

Licenses

The HNB-GW is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Platform Requirements

The HNB-GW service runs on a Cisco® ASR 5x00 chassis running StarOS Rel. 10 or later. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

Network Deployment and Interfaces

This section describes the supported interfaces and deployment scenario of HNB-GW in 3G Femto access network.

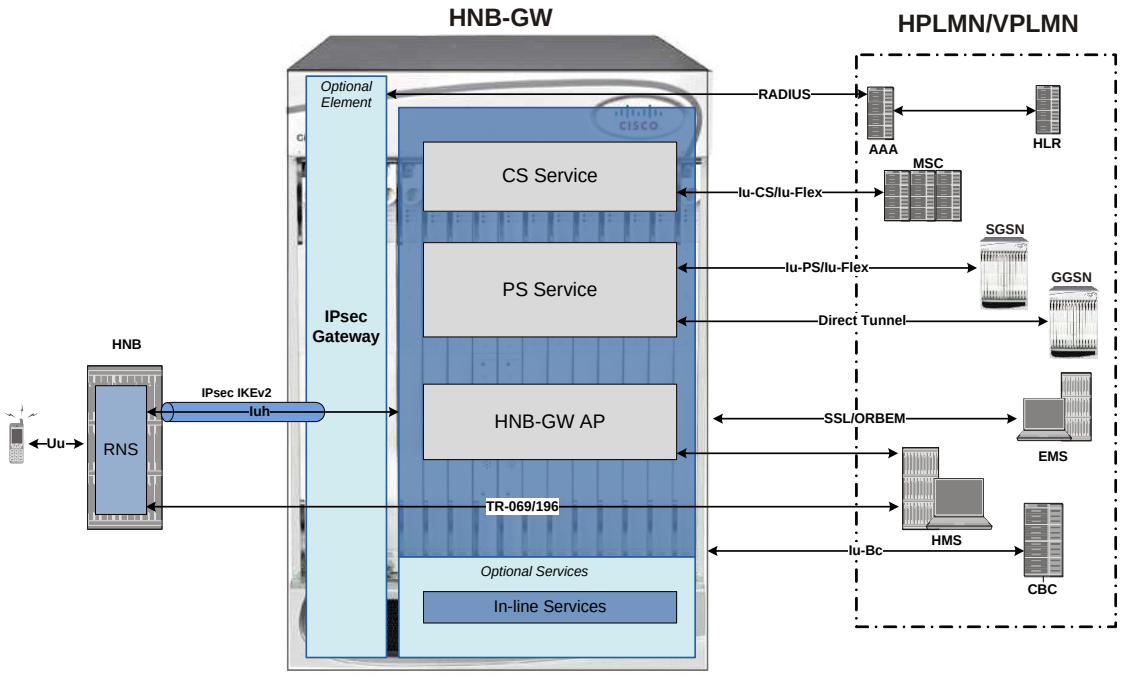
The following information is provided in this section:

- HNB Gateway in 3G UMTS Network
- Supported Logical Interfaces

HNB Gateway in 3G UMTS Network

The following figure displays simplified network views of the HNB-GW in an Femto access network accessing UMTS PS or CS Core Network.

Figure 119. HNB-GW in UMTS Network and Interfaces



Supported Logical Interfaces

This section provides the brief information on supported interfaces on HNB-GW node.

In support of both mobile and network originated subscriber UE contexts, the HNB-GW provides the following network interface support:

- **IuH Interface**: This interface is the reference point for the control plane protocol between Home NodeB and HNB-GW. IuH uses SCTP over IPSec IKEv2 tunnel as the transport layer protocol for guaranteed delivery of signaling messages between HNB-GW and Home NodeB.

This is the interface used by the HNB-GW to communicate with HNB on the same Femtocell Access Network. This interface serves as path for establishing and maintaining subscriber UE contexts.

One or more IuH interfaces can be configured per system context.

- **IuCS**: This interface is the reference point in UMTS which links the HNB-GW, which acts as an RNC (Radio Network Controller), with a Mobile Switching Centre (3G MSC) in the 3G UMTS Femtocell Access Network. This interface provides an IuCS over IP or IuCS over ATM (IP over AAL5 over ATM) interface between the MSC and the RNC (HNB-GW) in the 3G UMTS Femtocell Access Network. RAN Application Part (RANAP) is the control protocol that sets up the data plane (GTP-U) between these nodes. SIGTRAN (M3UA/SCTP) or QSAAL (MTP3B/QSAAL) handle IuCS (control) for the HNB-GW.

  This is the interface used by the HNB-GW to communicate with 3G MSC on the same Public Land Mobile Network (PLMN). This interface serves as path for establishing and maintaining the CS access for Femtocell UE to circuit switched UMTS core networks

  One or more IuCS interfaces can be configured per system context.

- **IuPS**: This interface is the reference point between HNB-GW and SGSN. This interface provides an IuPS over IP or IuPS over ATM (IP over AAL5 over ATM) interface between the SGSN and the RNC (HNB-GW) in the 3G UMTS Femtocell Access Network. RAN Application Part (RANAP) is the control protocol that sets up the data plane (GTP-U) between these nodes. SIGTRAN (M3UA/SCTP) or QSAAL (MTP3B/QSAAL) handle IuPS-C (control) for the HNB-GW.

  This is the interface used by the HNB-GW to communicate with SGSN on the same Public Land Mobile Network (PLMN). This interface serves as path for establishing and maintaining the PS access for Femtocell UE to packet switched UMTS core networks.

  One or more IuPS interfaces can be configured per system context.

- **Gi**: This interface is the reference point between HNB-GW and IP Offload Gateway. It is used by the HNB-GW to communicate with Packet Data Networks (PDNs) through IP Offload Gateway in the H-PLMN/V-PLMN. Examples of PDNs are the Internet or corporate intranets.

  One or more Gi interfaces can be configured per system context.

- **Gn**: This interface is the reference point between HNB-GW and GGSN. It is used by the HNB-GW to communicate with GGSNs on the same GPRS/UMTS Public Land Mobile Network (PLMN).

  One or more Gn interfaces can be configured per system context.

- **RADIUS**: This interface is the reference point between a Security Gateway (SeGW) and a 3GPP AAA Server or 3GPP AAA proxy (OCS/CGF/AAA/HSS) over RADIUS protocol for AAA procedures for Femto user.

  In the roaming case, the 3GPP AAA Proxy can act as a stateful proxy between the SeGW and 3GPP AAA Server.

  The AAA server is responsible for transfer of subscription and authentication data for authenticating/authorizing user access and UE authentication. The SeGW communicates with the AAA on the PLMN using RADIUS protocol.

  One or more RADIUS interfaces can be configured per system context.

- **TR-069**: This interface is an application layer protocol which is used for remote configuration of terminal devices, such as DSL modems, HNBs and STBs. TR-069 provides an auto configuration mechanism between the HNB and a remote node in the service provider network termed the Auto Configuration Server. The standard also uses a combination of security measures including IKEv2 (Internet Key Exchange v2) and IPsec (IP Security) protocols to authenticate the operator and subscriber and then guarantee the privacy of the data exchanged.

  One TR-069 interface can be configured per HNB node.



Important: DHCP interface support is available in StarOS Release 14.0 and later only.

Features and Functionality - Base Software

This section describes the features and functions supported by default in base software on HNB-GW service and do not require any additional license to implement the functionality with the HNB-GW service.

Following features and supports are discussed in this section:

- AAA Server Group Support
- AAL2 Establish and Release Support
- Access Control List Support
- ANSI T1.276 Compliance
- ATM VC Management Support
- Congestion Control and Management Support
- Emergency Call Handling
- GTP-U Tunnels Management Support
- HNB-UE Access Control
- HNB Management Function
- Intra-Domain Multiple CN Support Through Iu-Flex
- Iu Signalling Link Management Support
- IuH User-Plane Transport Bearer Handling Support
- Multiple MSC Selection without Iu-Flex
- Network Access Control Functions through SeGW
- Open Access Mode Support
- QoS Management with DSCP Marking
- RADIUS Support
- System Management Features
- UE Management Function for Pre-Rel-8 UEs

AAA Server Group Support

Value-added feature to enable VPN service provisioning for enterprise or MVNO customers. Enables each corporate customer to maintain its own AAA servers with its own unique configurable parameters and custom dictionaries.

This feature provides support for up to 800 AAA (RADIUS and Diameter) server groups and 800 NAS IP addresses that can be provisioned within a single context or across the entire chassis. A total of 128 servers can be assigned to an individual server group. Up to 1,600 accounting, authentication and/or mediation servers are supported per chassis and may be distributed across a maximum of 1,000 nodes. This feature also enables the AAA servers to be distributed across multiple nodes within the same context.

Important: For more information on AAA Server Group configuration, refer to the *AAA and GTPP Interface Administration and Reference*.

AAL2 Establish and Release Support

Support to establish and release of ATM adaptation layer 2 (AAL2) channel within an ATM virtual connection by the HNB-GW in complete or partial compliance with the following standards:

- **3GPP TS 25.414 V9.0.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface data transport and transport signalling (Release 9)
- **3GPP TS 25.415 V8.0.0 (2008-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface user plane protocols (Release 8)
- **3GPP TS 25.467 V8.0.0. (2008-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home NodeB; Stage 2 (Release 8)
- **3GPP TS 25.467 V9.1.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home Node B (HNB); Stage 2 (Release 9)
- **ITU-T Recommendation Q.2630.1**: AAL type2 signalling protocol (Capability Set 1)
- **ITU-T Recommendation Q.2630.2**: AAL type2 signalling protocol (Capability Set 2)
- **ITU-T Recommendation I.363.2 B**: ISDN ATM Adaptation Layer (AAL) Specification: Type 2 AAL
- **ITU-T Recommendation I.366.1**: Segmentation and Reassembly Service Specific Convergence Sublayer for the AAL type 2

The HNB-GW connects to core network elements like MSC and SGSN over IuCS and IuPS interfaces respectively. The Iu interface towards core network elements could either by IP based or ATM based. To provide ATM based interface support, Cisco HNB-GW provides AAL2 support on system in order to establish a voice bearer with MSC.

Access Control List Support

Access Control Lists provide a mechanism for controlling (i.e permitting, denying, redirecting, etc.) packets in and out of the system.

IP access lists, or Access Control Lists (ACLs) as they are commonly referred to, are used to control the flow of packets into and out of the system. They are configured on a per-context basis and consist of "rules" (ACL rules) or filters that control the action taken on packets that match the filter criteria

Once configured, an ACL can be applied to any of the following:

- An individual interface
- All traffic facilitated by a context (known as a policy ACL)
- An individual subscriber
- All subscriber sessions facilitated by a specific context

There are two primary components of an ACL:

- Rule: A single ACL consists of one or more ACL rules. As discussed earlier, the rule is a filter configured to take a specific action on packets matching specific criteria. Up to 128 rules can be configured per ACL.

Each rule specifies the action to take when a packet matches the specifies criteria. This section discusses the rule actions and criteria supported by the system.

- Rule Order: A single ACL can consist of multiple rules. Each packet is compared against each of the ACL rules, in the order in which they were entered, until a match is found. Once a match is identified, all subsequent rules are ignored.

Important: For more information on Access Control List configuration, refer *IP Access Control List* chapter in *System Administration Guide*.

ANSI T1.276 Compliance

ANSI T1.276 specifies security measures for Network Elements (NE). In particular it specifies guidelines for password strength, storage, and maintenance security measures.

ANSI T1.276 specifies several measures for password security.

These measures include:

- Password strength guidelines
- Password storage guidelines for network elements
- Password maintenance, e.g. periodic forced password changes

These measures are applicable to the systems and the Web Element Manager since both require password authentication. A subset of these guidelines where applicable to each platform will be implemented. A known subset of guidelines, such as certificate authentication, are not applicable to either product. Furthermore, the platforms support a variety of authentication methods such as RADIUS and SSH which are dependent on external elements. ANSI T1.276 compliance in such cases will be the domain of the external element. ANSI T1.276 guidelines will only be implemented for locally configured operators.

ATM VC Management Support

Support for Asynchronous Transfer Mode (ATM) virtual circuits (VC) management function of AAL2 and AAL5 protocol by the HNB-GW in accordance with the following standards:

- **3GPP TR 29.814 V7.1.0 (2007-06)**: 3rd Generation Partnership Project; Technical Specification Group Core Networks and Terminals Feasibility Study on Bandwidth Savings at Nb Interface with IP transport (Release 7)

HNBGW supports PVC (permanent virtual circuits) connections with CN nodes for AAL2 and AAL5 type of traffic. The Common Part Sublayer (CPS) payload which is carried out by the AAL2 protocol over ATM is also configurable with this feature. It provides the dynamic Common Part Sublayer (CPS) payload configuration for AAL2 protocol traffic over ATM for negotiation between HNB-GW and MSC during call. Default size for payload is 45 but values may range from 1 to 64 Bytes. This feature makes the operator to choose the CPS payload size dynamically.

Congestion Control and Management Support

Congestion Control monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are quickly resolved. However, continuous or large numbers of these conditions within a specific time interval may have an

impact the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establishes limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operation thresholds that are configured for the system as described in the *Thresholding Configuration Guide*. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated.

  A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.

  - **Port Utilization Thresholds**: If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
  - **Port-specific Thresholds**: If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
- **Service Congestion Policies:** Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

Important: For more information on Congestion Control support, refer *Congestion Control* chapter in *System Administration Guide*.

Emergency Call Handling

The HNB-GW supports the handling of Emergency call in accordance with the following standards:

- **3GPP TS 25.467 V9.3.0 (2010-06)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home Node B (HNB); Stage 2 (Release 9)
- **3GPP TS 33.102 V9.1.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G Security; Security architecture Release 9)

The HNB-GW provides access for all UE/HNB when emergency call initiated. In case of non CSG UEs or non CSG HNBs, after Emergency call is finished, the context established between the HNB and operator's core network entities for UEs, who can not get access over the HNB, will be de-registered to prevent the UE from accessing non-emergency services. However if the UE remains idle for value equal to ue-idle-time, then it will be de-registered.

HNB-GW handles the emergency call in following way:

- **Authentication**: In case of emergency call, HNB sends a UE REGISTRATION REQUEST message with "Registration cause" as emergency call and excludes the "UE Permanent identity" (i.e IMSI) and HNB-GW does not perform access control for emergency call case.
- **Single Iu and Single RAB**: In case of emergency call, HNB-GW does not allow multiple RABs for UE. This means that UE must have only one Iu connection, either CS or PS, and have only one RAB on that Iu connection. HNB-GW implements "Single IU, Single RAB policy" when UE registration comes with Emergency.

  The RUA-CONNECT has an IE called "establishment cause" which can take values as "Normal" or "Emergency". If UE-registration was due to emergency then RUA-CONNECT must contain "Emergency". If RUA-CONNECT contains "normal" then HNB-GW rejects it.

  While rejecting RUA connection or RAB connection the HNB-GW uses following reject cause:

  - RUA - Misc: unspecified
  - RAB - Misc: unspecified
- If UE-registration is normal then both (normal and emergency) RUA-CONNECT is allowed.

GTP-U Tunnels Management Support

Support to manage the GTP-U tunnels between HNB-GW and GSNs by in accordance with the following standards:

- **3GPP TS 25.467 V9.1.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home Node B (HNB); Stage 2 (Release 9)
- **3GPP TS 25.468 V9.0.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh Interface RANAP User Adaptation (RUA) signalling (Release 9)
- **3GPP TS 25.469 V9.0.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh interface Home Node B Application Part (HNBAP) signalling (Release 9)
- **3GPP TS 29.060 V9.0.0 (2009-09)**: 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface (Release 9)

HNB-GW supports establishment of GTPU tunnels for each RAB over the IuPS interface. HNB-GW terminates the GTP-U tunnels coming from CN (SGSN) and initiates separate GTP-U tunnel towards HNB.

HNB-UE Access Control

UE/HNB access control support in 3G UMTS HNB Access Network is provided on HNB-GW through IMSI White list database and AAA attribute processing. This feature is in accordance with following standards:

- **3GPP TS 23.003 V8.9.0 (2010-06)**: 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Numbering, addressing and identification (Release 8)
- **3GPP TS 25.467 V9.3.0 (2010-06)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home Node B (HNB); Stage 2 (Release 9)
- **3GPP TS 25.469 V9.2.0 (2010-06)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh interface Home Node B (HNB) Application Part (HNBAP) signalling (Release 9)
- IETF RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000

The HNB-GW provides UE registration and de-registration procedure for the HNB to convey Rel-8 UE identification data to the HNB-GW in order to perform access control for the UE in the HNB-GW. The UE Registration also establishes a UE specific context identifier to be used between HNB and HNB-GW. The procedure triggered when the UE attempts to access the HNB via an initial NAS message and there is no context in the HNB allocated for that UE.

For pre-Release 8 UEs, which do not support CSG and does not listen for CSG-ID, the HNB-GW ensures that a UE is authorized to access a particular Femtocell. To perform access control check for pre-Release 8 UE, HNB-GW maintains a per-HNB Whitelist. This whitelist consists of IMSIs which are allowed to access that particular HNB. The whitelist is stored in the HMS and is downloaded to HNB-GW when HNB-REGISTRATION procedure happens.

HNB Management Function

Support for HNB registration and de-registration in 3G UMTS HNB Access Network accordance with the following standards:

- **3GPP TS 25.469 V8.1.0 (2009-03)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh interface Home Node B Application Part (HNBAP) signalling (Release 8)
- IETF RFC 4960, Stream Control Transmission Protocol, December 2007

The HNB-GW provides HNB registration and de-registration procedure to register the HNB with the HNB-GW. This procedure enables the HNB-GW to provide service and core network connectivity for the HNB. On HNB-GW node this procedure is the first HNBAP procedure triggered after the SCTP association has become operational between HNB and HNB-GW.

HNB management function processes the HNB/UE access control procedure through White-List processing on HNB-GW node. Dynamic update of White-List gives the dynamic HNB management ability to HNB-GW.

Intra-Domain Multiple CN Support Through Iu-Flex

Iu-Flex is the routing functionality for intra domain connection of HNB-GW nodes to multiple CN nodes (MSC/SGSN). It provides a routing mechanism and related functionality on HNB-GW to enable it to route information of different Core Network (CN) nodes with in the CS or PS domain. It is implemented in accordance with the following standards:

- **3GPP TS 23.236 V9.0.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Intra-domain connection of Radio Access Network (RAN) nodes to multiple Core Network (CN) nodes (Release 9)
- **3GPP TS 25.468 V9.2.0 (2010-06)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh Interface RANAP User Adaptation (RUA) signalling (Release 9)

HNB-GW supports Iu-Flex routing mechanism and other applications like many-to-many relation and load-sharing between CN nodes with HNB-GW and CN node pooling. This mechanism provides following benefits to network operator:

- Eliminates the single point of failure between an RNC/HNB-GW and a CN Node.
- Ensures geographical redundancy, as a pool can be distributed across sites.
- Minimizes subscriber impact during service, maintenance, or node additions or replacements.
- Increases overall capacity via load sharing across the MSCs/SGSNs in a pool.
- Reduces the need/frequency for inter-CN node RAUs. This substantially reduces signaling load and data transfer delays.
- Supports load redistribution with the MSC/SGSN offloading procedure.

To incorporate the concept of multiple CN nodes, Iu-Flex introduces the concept of "pool-areas" which is enabled by the routing mechanism in HNB GW. A pool-area is served by multiple CN nodes (MSCs or SGSNs) in parallel which share the traffic of this area between each other. Furthermore, pool-areas may overlap. From a RAN perspective a pool-area comprises all LA(s)/RA(s) of one or more RNC/BSC or HNBGW that are served by a certain group of CN nodes in parallel. One or more of the CN nodes in this group may in addition serve LAs/RAs outside this pool-area or may also serve other pool-areas. This group of CN nodes is also referred to as MSC pool or SGSN pool respectively.

The Iu-Flex enables a few different application scenarios with certain characteristics. The service provision by multiple CN nodes within a pool-area enlarges the served area compared to the service area of one CN node. This results in reduced inter CN node updates, handovers and relocations and it reduces the HLR/HSS update traffic. The configuration

of overlapping pool-areas allows to separate the overall traffic into different UE moving pattern, e.g. pool-areas where each covers a separate residential area and all the same city centre. Other advantages of multiple CN nodes in a pool-area are the possibility of capacity upgrades by additional CN nodes in the pool-area or the increased service availability as other CN nodes may provide services in case one CN node in the pool-area fails.

Iu Signalling Link Management Support

Support for Iu signal link management function for HNB-GW in accordance with the following standards:

- **3GPP TS 25.412 V8.0.0 (2008-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface signalling transport (Release 8)
- **3GPP TS 25.413 V7.9.0 (2008-06)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface RANAP signalling (Release 7)
- **3GPP TS 25.414 V9.0.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface data transport and transport signalling (Release 9)

HNBGW supports RANAP protocol for management of IuPS/IuCS connections. The IU connection on the IuPS/IuCS interface is realized using an SCCP connection towards SGSN/MSC. The SCCP could be over SIGTRAN or ATM.

IuH User-Plane Transport Bearer Handling Support

Support for transfer of CS as well as PS data over IP on the IuH interface:

- **3GPP TS 25.467 V8.0.0. (2008-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home NodeB; Stage 2 (Release 8)

HNB-GW supports GTP-U v1 for PS traffic transport and RTP/RTCP for CS traffic transport on IuH interface. HNB-GW terminates the GTPU tunnels and RTP sessions at itself for each tunnel/session between CN and HNB.

Multiple MSC Selection without Iu-Flex

Support for multiple MSC selection in a CS core network is provided with this feature support.

HNBGW can connect to multiple MSC and SGSN through Iu-Flex or LAC mapping. This feature implements the multiple MSC selection using LAC.

For this support the HNB-GW uses HNB's LAC, received during registration procedure in HNB_REGISTER_REQUEST message, to distribute RANAP-Initial UE message to an MSC. It maps the LAC with MSC point code and a set of LACs configured for each MSC, connected to the HNB-GW.

In the HNBGW, to select an MSC based on the LAC the following algorithm is used:

- If both Iu-Flex and LACs are configured for a MSC, then Iu-Flex is used to select a MSC.
- If only Iu-Flex is configured then Iu-Flex is used for selecting MSC.
- If only LACs are configured then MSC is selected using LAC from HNB.
- If both Iu-Flex and LACs are not configured in the HNBGW, it selects default MSC.

Network Access Control Functions through SeGW

These functions enable secure user and device level authentication between the authenticator component of the HNB-GW and a 3GPP HSS/AuC and RADIUS-based AAA interface support.

This section describes following features:

- Authentication and Key Agreement (AKA)
- 3GPP AAA Server Support
- X.509 Certificate-based Authentication Support

Authentication and Key Agreement (AKA)

HNB-GW provides Authentication and Key Agreement mechanism for user authentication procedure over the HNB Access Network. The Authentication and Key Agreement (AKA) mechanism performs authentication and session key distribution in networks. AKA is a challenge- response based mechanism that uses symmetric cryptography.

The AKA is the procedure that take between the user and network to authenticate themselves towards each other and to provide other security features such as integrity and confidentiality protection.

In a logical order this follows the following procedure:

1. **Authentication**: Performs authentication by, identifying the user to the network; and identifying the network to the user.
2. **Key agreement**: Performs key agreement by, generating the cipher key; and generating the integrity key.
3. **Protection**: When the AKA procedure is performed it protects, the integrity of messages; confidentiality of signalling data; and confidentiality of user data

3GPP AAA Server Support

This interface between the SeGW and AAA Server provides a secure connection carrying authentication, authorization, and related information. in accordance with the following standards:

- 3GPP TS 33.320 V9.1.0 (2010-03): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Security of Home Node B (HNB) / Home evolved Node B (HeNB) (Release 9)

This reference point is located between 3GPP AAA Server/Proxy and HNB-GW. The functionality of this reference point is to enable following requirements on SeGW:

- The SeGW shall be authenticated by the HNB using a SeGW certificate.
- The SeGW shall authenticate the HNB based on HNB certificate.
- The SeGW authenticates the hosting party of the HNB in cooperation with the AAA server using EAP-AKA.
- The SeGW shall allow the HNB access to the core network only after successful completion of all required authentications.
- Any unauthenticated traffic from the HNB shall be filtered out at the SeGW

X.509 Certificate-based Authentication Support

HNB-GW supports X.509 Certificate-based authentication to HNB/UE for a public key infrastructure (PKI) for single sign-on (SSO) and Privilege Management Infrastructure (PMI). X.509 specifies the standard formats for public key certificates, certificate revocation lists, attribute certificates, and a certification path validation algorithm.

Open Access Mode Support

An Open HNB provides its services to any mobile subscriber in Femto network. This feature is intended to provide Open access mode support on an UMTS HNB-GW. Open access HNBs can be deployed in public places like airports to increase the indoor coverage or to offload the traffic from the macro cell.

This feature provides following procedure for Open Access HNB registration on HNB-GW:

- **OPEN HNB registration:**
  1. On receiving HNB-REGISTER-REQ from HNB, HNB-GW sends RADIUS-Access-Request to AAA-server. HNB-GW will not care whether HNB has sent cell-access-mode in the register request or not.
  2. AAA-server performs authentication and authorization. If this is successful AAA-server sends RADIUS-Access-Accept to HNB-GW. Then AAA-server includes the Whitelist attribute in the response. AAA-server prepare Whitelist attribute in the following manner:
     - Cell-access-mode field in the Whitelist attribute will be set to "Open".
     - Number-of-IMSIs field in the Whitelist attribute will be set to 0 (zero).
     - IMSIs will not be included in the Whitelist attribute.
  3. HNB-GW overrides the cell-access-mode value received from HNB by the one received from AAA-server.
  4. HNB-GW discards "IMSI List" received for an Open Access-mode HNB from AAA-server in Access-Accept or COA message.
  5. AAA server sends access-mode as 0 (Closed mode) or 1 (Hybrid mode) or 2 (Open mode). If it sends any other value in Access-Accept, then HNB-GW shall send HNB-REGISTER-REJECT with Unauthorised-HNB cause.

  Important: HNB-GW support both "Open" mode and "Closed" mode HNBs simultaneously. In case HNB registration is disabled the HNB-GW sends HNB-REGISTER-REJECT with O&M Intervention cause.

- **UE registration**:
  1. HNB-GW does not perform access control check of any UE registration request received from an Open Access-mode HNB.
  2. HNB-GW also does not perform any check on the UE-identity received in the registration request from an Open Access mode HNB.
- **RUA Connect**: HNB-GW does not check IMSI received in RANAP-CommonId if the same was not received from HNB in UE-REGISTER-REQ.
- **Paging**:
  1. Paging received without paging-area is sent to all Open HNBs. CS Paging received with paging-area is sent to those Open HNBs whose LAC (received in HNB-REG-REQ) matches with the one received in paging message.
  2. PS Paging received with paging-area is sent to those Open HNBs whose LAC and RAC (received in HNB-REG-REQ) matches with those received in paging message.
  3. If sending paging towards open HNBs is disabled the paging with or without paging-area will not be sent to any Open HNB.

On receiving a relocation request, relocation proceeds only if target Closed HNB is found. If UE is currently registered via Open HNB, then that registration will be cleaned up.

Important: For more information on Open Access Mode support configuration, refer *Open Access Mode Configuration* section of *HNB-GW Administration Guide*.

QoS Management with DSCP Marking

Differentiated Services Code Point (DSCP) marking over IuH interface support in 3G UMTS HNB Access Network is provided on HNB-GW for traffic quality management in accordance with following standards:

- **3GPP TS 25.414 V9.0.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface data transport and transport signalling (Release 9)
- **3GPP TS 25.468 V9.2.0 (2010-06)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh Interface RANAP User Adaptation (RUA) signalling (Release 9)
- **3GPP TS 25.469 V9.2.0 (2010-06)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh interface Home Node B (HNB) Application Part (HNBAP) signalling (Release 9)
- IETF RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
- IETF RFC 4594, Configuration Guidelines for DiffServ Service Classes
- IETF RFC 4960, Stream Control Transmission Protocol

In a fixed line-mobile convergence scenario, the user data and signaling traffic from a UE is forwarded by an HNB to HNB-GW over IuH interface. IP is used as network layer for IuH. RTP/ RTCP or GTP over UDP/IP form transport for user data. SCTP/IP is used for control signaling over IuH.

These data and control packets traverse public Internet before reaching HNB-GW and vice-a-versa for the downlink traffic. RTP typically carries jitter-sensitive real-time media data such as voice and video. RTCP carries media reception/transmit feedback that is not delay sensitive. GTP carries generic, non-media data. These various traffic types, each, deserve different QoS handling by the IP nodes they traverse between HNB and HNB-GW. Thus DSCP codes are assigned in the IP headers of the traffic such that intermediate IP nodes can provide differentiated QoS treatment to the traffic for an acceptable end-user experience.

HNB-GW supports DSCP marking of the traffic on IuH and Iu interface for downlink traffic towards HNB and for uplink traffic towards CN when IP transport is used for IuCS or IuPS.

Important: For more information on DSCP marking configuration, refer *DSCP Marking Configuration* section of *HNB-GW Administration Guide*.

RADIUS Support

In HNB-GW the RADIUS support provides a mechanism for performing authorization and authentication for subscriber sessions based on the following standards:

- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000

- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000

Within context configured on the system, there are AAA and RADIUS protocol-specific parameters that can be configured. The RADIUS protocol-specific parameters are further differentiated between RADIUS Authentication server RADIUS Accounting server interaction.

Among the RADIUS parameters that can be configured are:

- **Priority**: Dictates the order in which the servers are used allowing for multiple servers to be configured in a single context.
- **Routing Algorithm**: Dictate the method for selecting among configured servers. The specified algorithm dictates how the system distributes AAA messages across the configured AAA servers for new sessions. Once a session is established and an AAA server has been selected, all subsequent AAA messages for the session will be delivered to the same server.

In the event that a single server becomes unreachable, the system attempts to communicate with the other servers that are configured. The system also provides configurable parameters that specify how it should behave should all of the RADIUS AAA servers become unreachable.

Important: For more information on RADIUS AAA configuration, refer to the *AAA and GTPP Interface Administration and Reference*.

UE Management Function for Pre-Rel-8 UEs

Support for Pre-Rel-8 UE registration and de-registration in 3G UMTS HNB Access Network in accordance with the following standards:

- **3GPP TS 25.467 V8.0.0. (2008-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home NodeB; Stage 2 (Release 8)
- **3GPP TS 25.469 V8.1.0 (2009-03)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh interface Home Node B Application Part (HNBAP) signalling (Release 8)
- IETF RFC 4960, Stream Control Transmission Protocol, December 2007

The HNB-GW provides UE registration and de-registration procedure for the HNB to convey pre-Rel-8 UE identification data to the HNB-GW in order to perform access control for the UE in the HNB-GW. The UE Registration also establishes a UE specific context identifier to be used between HNB and HNB-GW. The procedure triggered when the UE attempts to access the HNB via an initial NAS message and there is no context in the HNB allocated for that UE.

System Management Features

This section describes following features:

- Management System Overview
- Bulk Statistics Support
- Threshold Crossing Alerts (TCA) Support
- ANSI T1.276 Compliance

Management System Overview

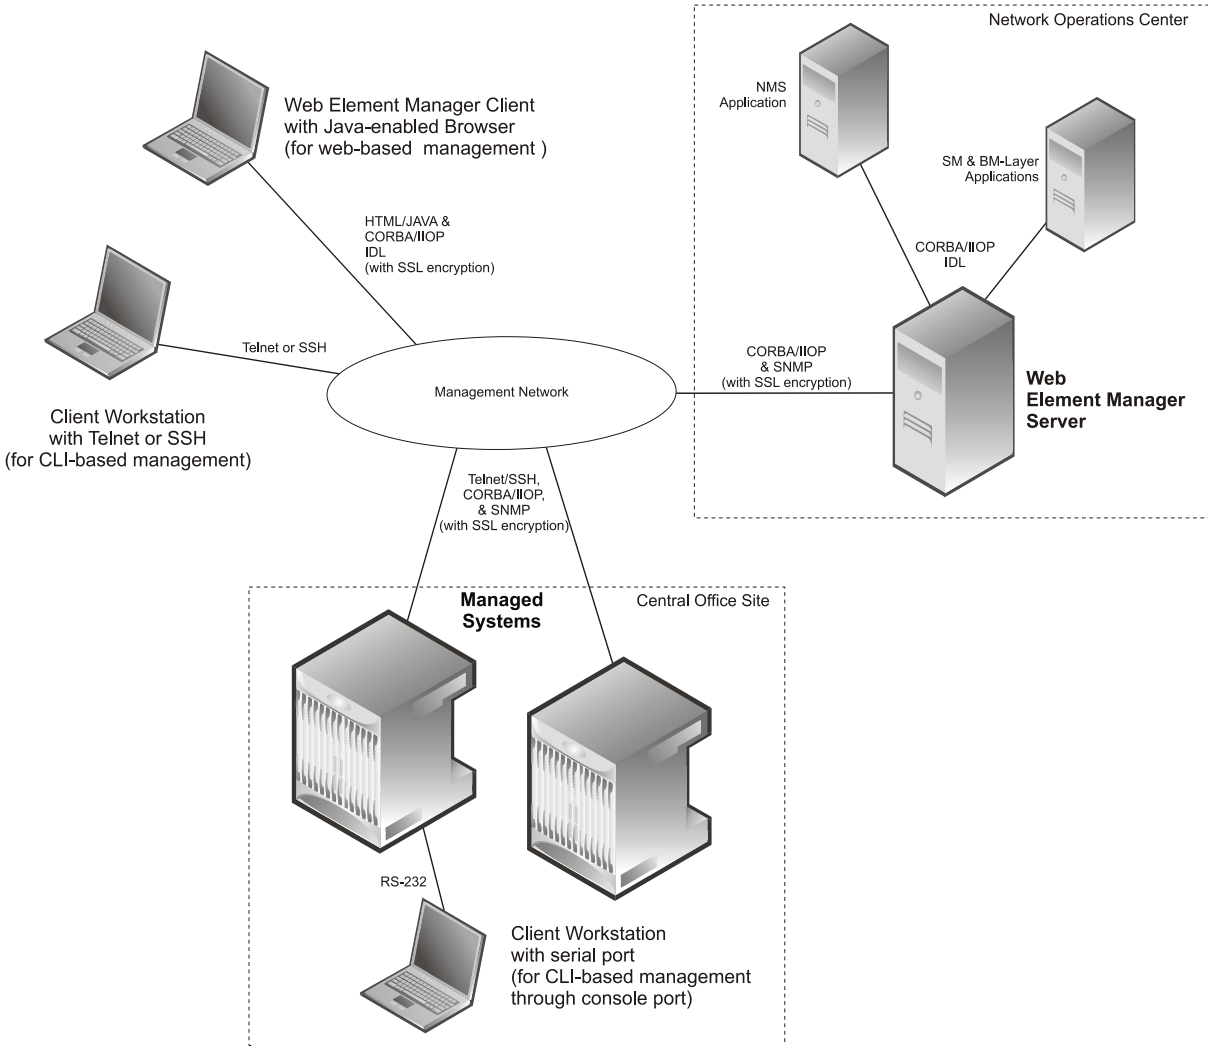
The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management - focusing on providing superior quality network element (NE) and element management system (Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems - giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

Operation and Maintenance module of chassis offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces. These include:

- Using the command line interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager application
- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000
- Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (i.e. Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network.

Figure 120. Element Management System



Important: HNB-GW management functionality is enabled for console-based access by default. For GUI-based management support, refer *WEM Installation and Administration Guide*.

Important: For more information on command line interface based management, refer *Command Line Interface Reference*.

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Following is a partial list of supported schemas:

- **System**: Provides system-level statistics
- **Card**: Provides card-level statistics
- **Port**: Provides port-level statistics
- **GTP-U**: Provides GPRS Tunneling Protocol - User message statistics
- **HNB-AAL2**: Provides ATM adaptation layer 2 (AAL2) protocol level-statistics
- **HNB-ALCAP**: Provides Access Link Control Application Part (ALCAP) service-level statistics
- **CS-Network-RANAP**: Provides RANAP-level statistics for HNB-CS network
- **CS-Network-RTP**: Provides RTP protocol-level statistics for HNB-CS network
- **HNB-GW-HNBAP**: Provides HNBAP-level statistics for HNB-GW service
- **HNB-GW-HNBAP-ACCESS-CLOSED**: Provides HNBAP-level statistics filtered for HNBs registered for Closed access mode with HNB-GW service
- **HNB-GW-HNBAP-ACCESS-HYBRID**: Provides HNBAP-level statistics filtered for HNBs registered for Hybrid access mode with HNB-GW service
- **HNB-GW-HNBAP-ACCESS-OPEN**: Provides HNBAP-level statistics filtered for HNBs registered for Open access mode with HNB-GW service
- **HNB-GW-RANAP**: Provides RANAP-level statistics for HNB-GW service
- **HNB-GW-RANAP-ACCESS-CLOSED**: Provides RANAP-level statistics filtered for HNBs registered for Closed access mode with HNB-GW service
- **HNB-GW-RANAP-ACCESS-HYBRID**: Provides RANAP-level statistics filtered for HNBs registered for Hybrid access mode with HNB-GW service
- **HNB-GW-RANAP-ACCESS-OPEN**: Provides RANAP-level statistics filtered for HNBs registered for Open access mode with HNB-GW service
- **HNB-GW-RTP**: Provides RTP protocol-level statistics for HNB-GW service
- **HNB-GW-RUA**: Provides RUA protocol-level statistics for HNB-GW service
- **HNB-GW-RUA-ACCESS-CLOSED**: Provides RUA protocol-level statistics filtered for HNBs registered for Closed access mode with HNB-GW service
- **HNB-GW-RUA-ACCESS-HYBRID**: Provides RUA protocol-level statistics filtered for HNBs registered for Hybrid access mode with HNB-GW service
- **HNB-GW-RUA-ACCESS-OPEN**: Provides RUA protocol-level statistics filtered for HNBs registered for Open access mode with HNB-GW service
- **HNB-GW-SCTP**: Provides HNB -SCTP protocol-level statistics
- **PS-Network--RANAP**: Provides RANAP-level statistics for HNB-PS network
- **SCCP**: Provides SCCP service-level statistics at system-level
- **SS7Link**: Provides SS7 link configuration related statistics at system-level
- **SS7 Routing Domain**: Provides SS7 Routing domain configuration related statistics at system level

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the IMG or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, IMG host name, IMG uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, number of sessions etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert**: A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm**: Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps**: SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values.

  Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.

- **Logs**: The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

  Logs are supported in both the Alert and the Alarm models.

- **Alarm System**: High threshold alarms generated within the specified polling interval are considered "outstanding" until a the condition no longer exists or a condition clear alarm is generated. "Outstanding" alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer *Thresholding Configuration Guide*.

ANSI T1.276 Compliance

ANSI T1.276 specifies security measures for Network Elements (NE). In particular it specifies guidelines for password strength, storage, and maintenance security measures.

ANSI T1.276 specifies several measures for password security. These measures include:

- Password strength guidelines
- Password storage guidelines for network elements
- Password maintenance, e.g. periodic forced password changes

These measures are applicable to the systems and the Web Element Manager since both require password authentication. A subset of these guidelines where applicable to each platform will be implemented. A known subset of guidelines, such as certificate authentication, are not applicable to either product. Furthermore, the platforms support a variety of authentication methods such as RADIUS and SSH which are dependent on external elements. ANSI T1.276 compliance in such cases will be the domain of the external element. ANSI T1.276 guidelines will only be implemented for locally configured operators.

Features and Functionality - Optional Enhanced Feature Software

This section describes the optional enhanced features and functions support with HNB-GW service.

Important: Some of the following features may require the purchase of an additional license to implement the functionality with the HNB-GW service.

This section describes following features:

- Dynamic RADIUS Extensions (Change of Authorization)
- IP Security (IPSec)
- Session Recovery
- Web Element Management System

Dynamic RADIUS Extensions (Change of Authorization)

Dynamic RADIUS extension support provide operators with greater control over subscriber PDP contexts by providing the ability to dynamically redirect data traffic, and or disconnect the PDP context.

This functionality is based on the RFC 3576, Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS), July 2003 standard.

The system supports the configuration and use of the following dynamic RADIUS extensions:

- **Change of Authorization**: The system supports CoA messages from the AAA server to change data filters associated with a subscriber session. The CoA request message from the AAA server must contain attributes to identify NAS and the subscriber session and a data filter ID for the data filter to apply to the subscriber session.
- **Disconnect Message**: The DM message is used to disconnect subscriber sessions in the system from a RADIUS server. The DM request message should contain necessary attributes to identify the subscriber session.

The above extensions can be used to dynamically re-direct subscriber PDP contexts to an alternate address for performing functions such as provisioning and/or account set up. This functionality is referred to as Session Redirection, or Hotlining.

Session redirection provides a means to redirect subscriber traffic to an external server by applying ACL rules to the traffic of an existing or a new subscriber session. The destination address and optionally the destination port of TCP/IP or UDP/IP packets from the subscriber are rewritten so the packet is forwarded to the designated redirected address.

Return traffic to the subscriber has the source address and port rewritten to the original values. The redirect ACL may be applied dynamically by means of the Radius Change of Authorization (CoA) extension.

Important: For more information on dynamic RADIUS extensions support, refer *CoA, RADIUS, And Session Redirection (Hotlining)* in this guide.

IP Security (IPSec)

IP Security provides a mechanism for establishing secure tunnels from mobile subscribers to pre-defined endpoints (i.e. enterprise or home networks) in accordance with the following standards:

- RFC 2401, Security Architecture for the Internet Protocol
- RFC 2402, IP Authentication Header (AH)
- RFC 2406, IP Encapsulating Security Payload (ESP)
- RFC 2409, The Internet Key Exchange (IKE)
- RFC-3193, Securing L2TP using IPSEC, November 2001

IP Security (IPSec) is a suite of protocols that interact with one another to provide secure private communications across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways.

IPSec tunnel supports AAA and DHCP address overlapping. Address overlapping is meant for multiple customers using the same IP address for AAA/DHCP servers. The AAA and DHCP control messages are sent over IPSec tunnels and AAA/DHCP packets required to be encrypted are decided as per the ACL configuration done for specific session.

Important: For more information on IPSec configuration, refer *HNB-GW Service Configuration* section.

Session Recovery

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

Session recovery is performed by mirroring key software processes (e.g. session manager and AAA manager) within the system. These mirrored processes remain in an idle state (in standby-mode), wherein they perform no processing, until they may be needed in the case of a software failure (e.g. a session manager task aborts). The system spawns new instances of "standby mode" session and AAA managers for each active Control Processor (CP) being used.

Additionally, other key system-level software tasks, such as VPN manager, are performed on a physically separate packet processing card to ensure that a double software fault (e.g. session manager and VPN manager fails at same time on same card) cannot occur. The packet processing card used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

The additional hardware resources required for session recovery include a standby System Processor Card (SPC) and a standby packet processing card.

There are two modes for Session Recovery.

- **Task recovery mode**: Wherein one or more session manager failures occur and are recovered without the need to use resources on a standby packet processing card. In this mode, recovery is performed by using the mirrored "standby-mode" session manager task(s) running on active packet processing cards. The "standby-mode" task is renamed, made active, and is then populated using information from other tasks such as AAA manager.
- **Full packet processing card recovery mode**: Used when a packet processing card hardware failure occurs, or when a packet processing card migration failure happens. In this mode, the standby packet processing card is made active and the "standby-mode" session manager and AAA manager tasks on the newly activated packet processing card perform session recovery.

Session/Call state information is saved in the peer AAA manager task because each AAA manager and session manager task is paired together. These pairs are started on physically different packet processing cards to ensure task recovery.

Important: For more information on this feature, refer *Session Recovery* chapter in *System Administration Guide*.

Web Element Management System

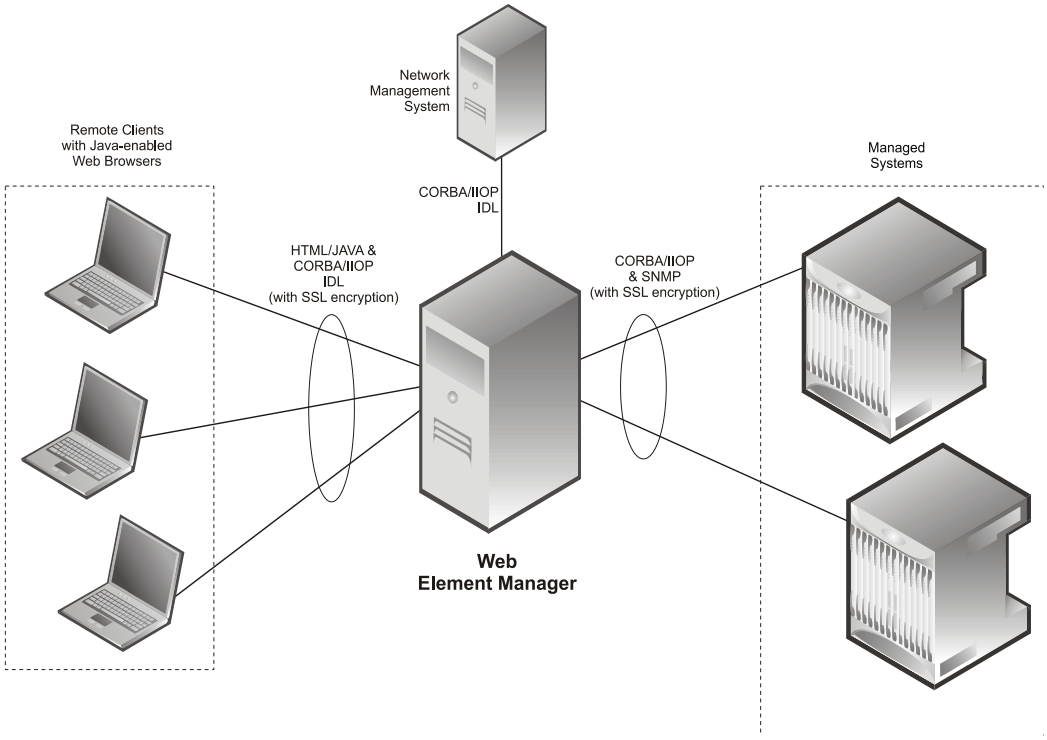
Provides a Graphical User Interface (GUI) for performing Fault, Configuration, Accounting, Performance, and Security (FCAPS) management of the system.

The Web Element Manager is a Common Object Request Broker Architecture (CORBA)-based application that provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) management capability for the system.

For maximum flexibility and scalability, the Web Element Manager application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the fault-tolerant Sun® Solaris® operating system.

The following figure demonstrates various interfaces between the Cisco Web Element Manager and other network components.

Figure 121. Web Element Manager Network Interfaces



Important: For more information on WEM support, refer *WEM Installation and Administration Guide*.

How HNB-GW Works

This section provides information on the function and procedures of the HNB-GW in a wireless network and presents message flows for different stages of session setup.

The following procedures are supported in this release:

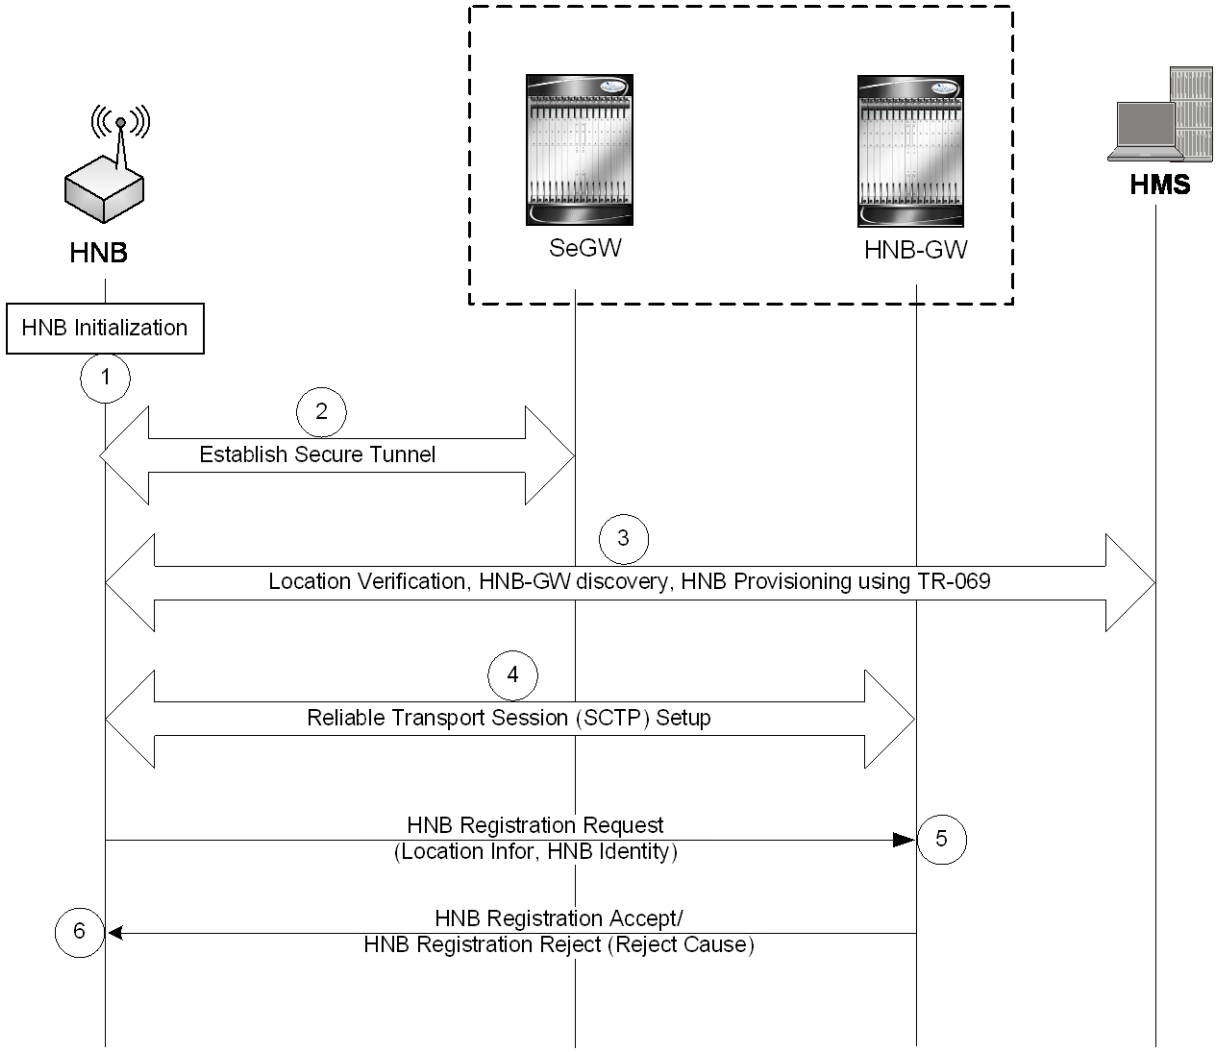
- HNB Provisioning and Registration Procedure
- UE Registration Procedure
- Iu Connection Procedures
- Paging and Serving RNS Relocation Procedures
- RANAP Reset Procedures

HNB Provisioning and Registration Procedure

This section describes the call flow for HNB provisioning and registration procedure.

The following figure and the text that follows describe the message flow for HNB provisioning and registration with HNB-GW procedure.

Figure 122. HNB Provisioning and Registration Setup Call Flow



1. HNB initialization is performed to obtain HNB configuration from the HNB Management System (HMS). Similarly, HNB-GW discovery is performed to obtain the initial serving HNB-GW information.
2. A secure tunnel is established from the HNB to the Security Gateway.
3. Location verification shall be performed by the HMS based on information sent by the HNB (e.g. macro neighbor cell scans, global navigational satellite system type of information etc.). HMS determines the serving elements and provides the HNB-GW, HMS and Security Gateway to the HNB. The HMS also provisions configuration parameters to the HNB only after successful location verification in the HMS.
4. Reliable transport setup (SCTP) completed and the HNB sets up a SCTP transport session to a well-defined port on the serving HNB-GW. HNB Registration procedure started.
5. The HNB attempts to register with the serving HNB-GW using a HNB-REGISTER-REQUEST message. This message may contains:
    - **HNB Location Information**: The HNB provides location information via use of one or more of the following mechanisms:
        - detected macro coverage information (e.g. GERAN and/or UMTS cell information)
        - geographical co-ordinates (e.g. via use of GPS, etc)

- Internet connectivity information (e.g. IP address).

- **HNB Identity**: the HNB has a globally unique and permanent identity.
- **HNB Operating Parameters**: Such as the selected LAC, RAC, SAC, etc.

6. The HNB-GW uses the information from the HNB-REGISTER-REQUEST message to perform access control of the HNB (e.g. whether a particular HNB is allowed to operate in a given location, etc). If the HNB-GW accepts the registration attempt the PLMN-ID received in the request shall be used to lookup the PLMN to RNC id mapping table and corresponding RNC-ID shall be returned in the HNB-REGISTER-ACCEPT message else the HNB-GW may reject the registration request (e.g. due to network congestion, blacklisted HNB, unauthorized HNB location, etc). In reject case, the HNB-GW shall respond with a HNB-REGISTER-REJECT indicating the reject cause.

Important: The HNB shall start broadcasting only after successful registration with the HNB-GW.

UE Registration Procedure

This section describes the UE registration procedures for HNB provides means for the HNB to convey UE identification data to the HNB-GW in order to perform access control for the UE in the HNB GW. The UE Registration also informs the HNB-GW of the specific HNB where the UE is located.

The UE registration procedure generally triggers when the UE attempts to access the HNB through an initial NAS message and there is no context id in the HNB for specific UE.

UE Registration procedure is described for following scenarios:

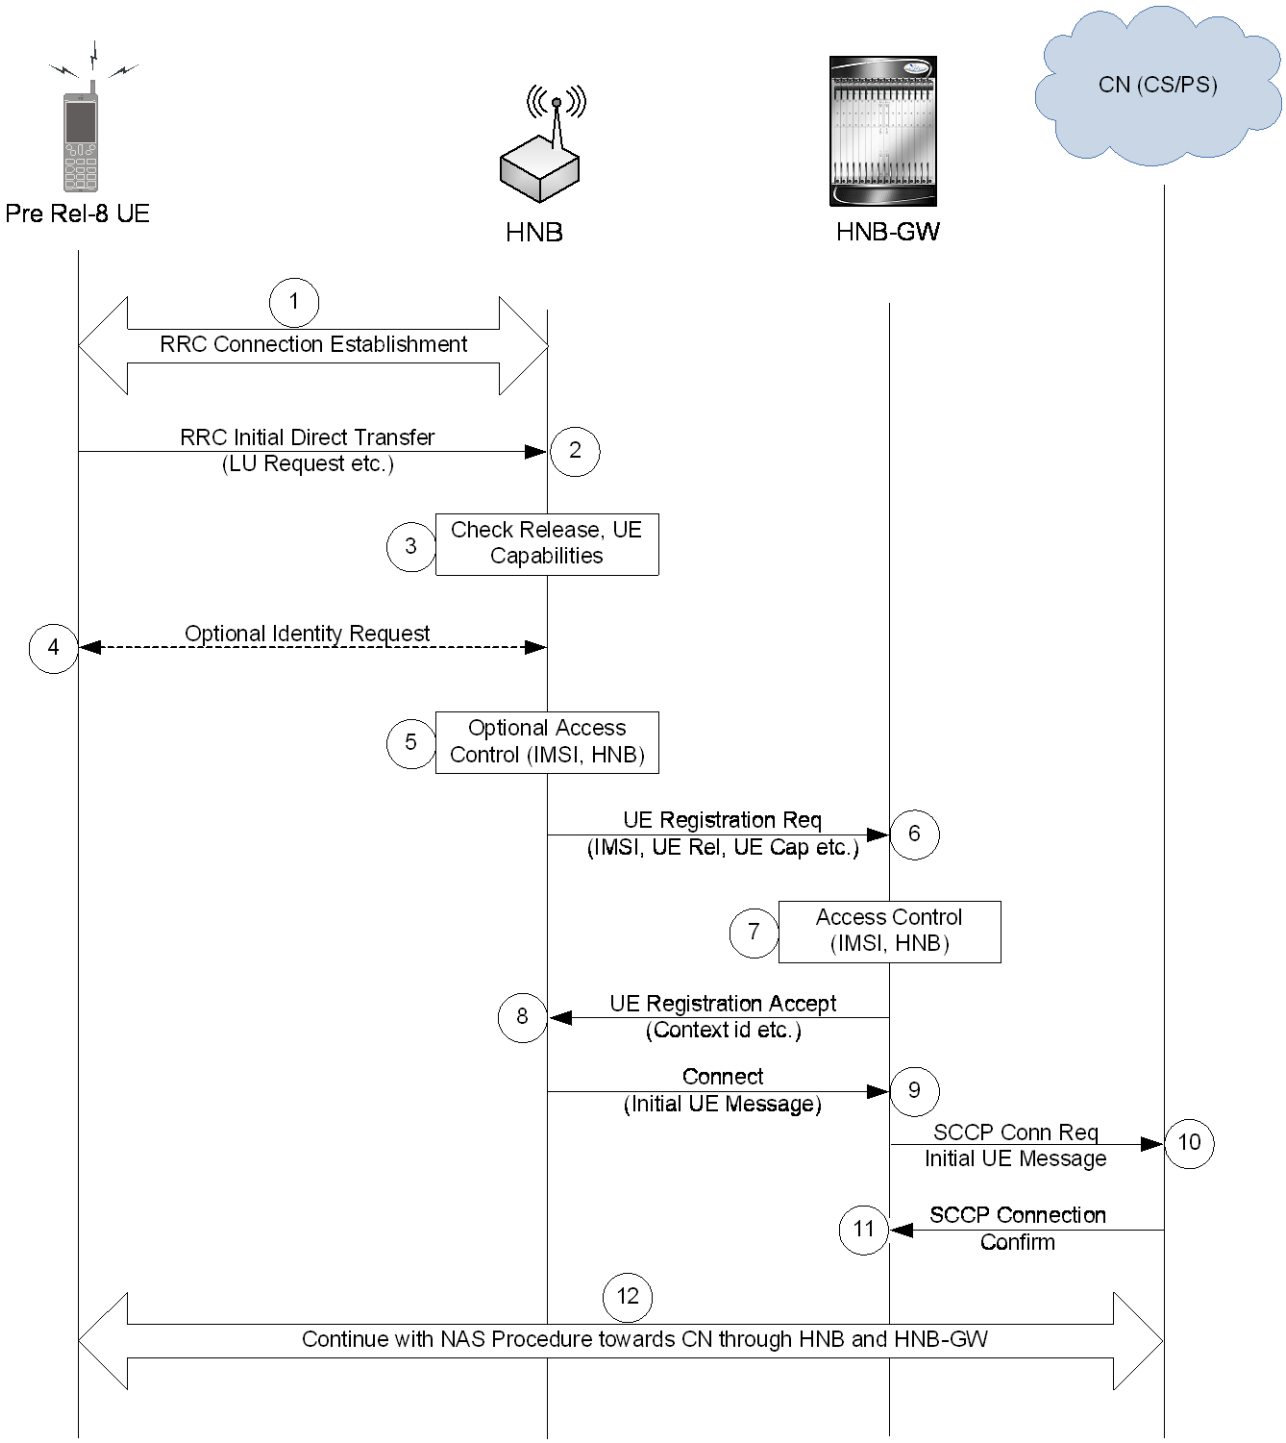
- UE Registration Procedure of Non-CSG UEs or Non-CSG HNBs

UE Registration Procedure of Non-CSG UEs or Non-CSG HNBs

This procedure is applicable for non-CSG UEs or HNBs.

The following figure and the text that follows describe the message flow for UE registration procedure of Non-CSG UEs or Non-CSG HNBs:

Figure 123. UE Registration Call Flow for Non-CSG UEs or Non-CSG HNBs



1. Upon camping on the HNB, the UE initiates an initial NAS procedure (e.g. LU Procedure) by establishing an RRC connection with the HNB. UE capabilities are reported to the HNB as part of the RRC Connection establishment procedure.

2. The UE then transmits a RRC Initial Direct Transfer message carrying the initial NAS message (e.g. Location Updating Request message) with identity (IMSI or TMSI).
3. The HNB checks UE capabilities provided in step 1, if these indicate that CSG is not supported and if the identity of the UE (provided during RRC Connection Establishment) is unknown at the HNB being accessed, i.e. no Context id exists for the UE, the HNB initiates UE registration towards HNB-GW (step 6-8).
4. Before starting the UE Registration procedure, HNB optionally triggers the Identification procedure asking for the UE IMSI, if such identity is not provided during the RRC Connection Establishment. If the HNB has a context id for the UE, the UE registration procedure is not performed nor the Identification procedure.
5. The HNB may optionally perform access control based on IMSI and provided access control list.
6. The HNB attempts to register the UE on the HNB-GW by transmitting the UE-REGISTER-REQUEST. The message contains at a minimum:
    - **UE Identity**: IMSI of the (U)SIM associated with the UE and the indication about UE capabilities provided in step 1.

    Important: The UE IMSI provided in the UE-REGISTER message is unauthenticated.
7. The HNB-GW checks UE capabilities and if these indicate that CSG is not supported the HNB-GW shall perform access control for the particular UE attempting to utilize the specific HNB.
8. If the HNB-GW accepts the UE registration attempt it shall allocate a context-id for the UE and respond with a UE-REGISTER-ACCEPT message, including the context-id, to the HNB. If the HNB-GW chooses to not accept the incoming UE registration request then the HNB-GW shall respond with a UE-REGISTRATION-REJECT message.
9. The HNB then sends a RUA (RANAP User Adaptation) CONNECT message containing the RANAP Initial UE message to HNB-GW.
10. The reception of the RUA CONNECT message at the HNB-GW triggers the setup of SCCP connection by the HNB-GW towards the CN. HNB-GW forwards the Initial UE Message to CN.
11. The CN response with a SCCP Connection Confirm message to HNB-GW.
12. The UE then continue with the NAS procedure (e.g. Location Updating procedure) towards the CN, via HNB and the HNB-GW.

Iu Connection Procedures

This section describes call flow for Iu connection procedures on HNB-GW.

Following procedure call flows are described for Iu connection procedures between HNB, HNB-GW, and SGSN/MSC in core network:

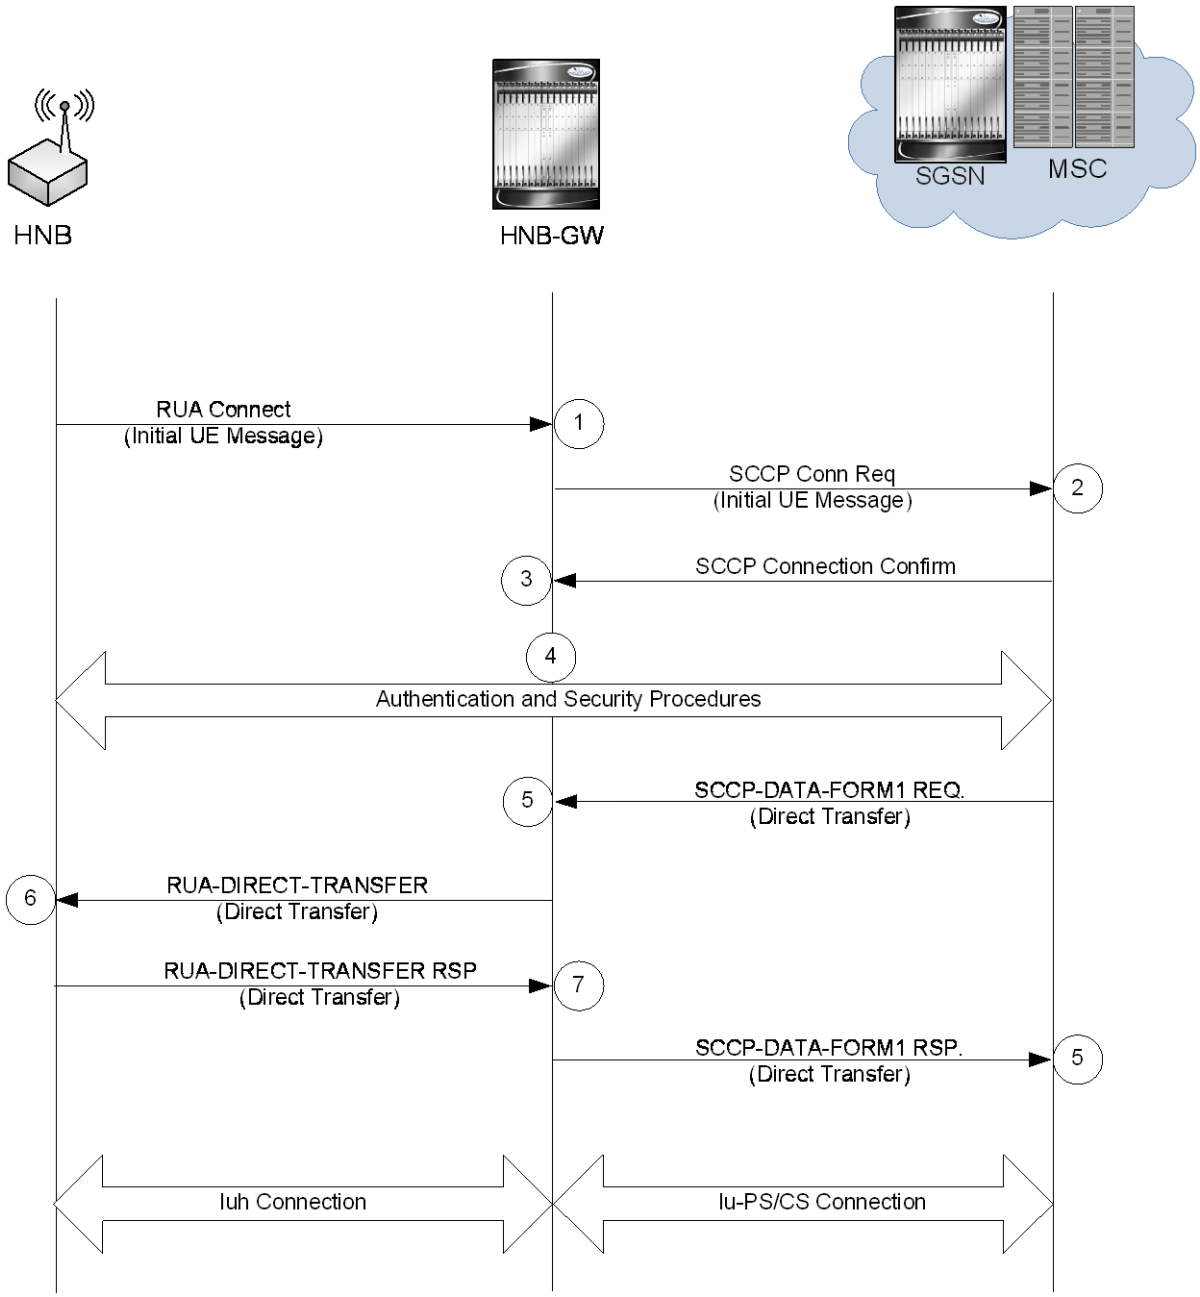
- Iu Connection Establishment Procedure
- Network Initiated Iu Connection Release Procedure

Iu Connection Establishment Procedure

This procedure is applicable for establishment of IuH and IuPS/IuCS connection between HNB to HNB-GW and HNB-GW to SGSN/MSC in core network.

The following figure and the text that follows describe the message flow for an Iu connection establishment procedure.

Figure 124. Iu Connection Establishment Call Flow



1. Upon receiving of UE-REGISTER-ACCEPT message from HNB-GW, the HNB then sends a RUA CONNECT message to HNB-GW containing the RANAP Initial UE message.
2. The reception of the RUA CONNECT message at the HNB-GW triggers the setup of SCCP connection by the HNB-GW towards the CN (SGSN/MSC). HNB-GW forwards the Initial UE Message.
3. The CN responses with a SCCP Connection Confirm message.
4. The UE then continue with the authentication and security procedures towards the CN, via HNB and the HNB-GW.
5. The SGSN/MSC performs Direct Transfer procedure with HNB-GW and sends SCCP-DATA-FORM1 REQ to HNB-GW.

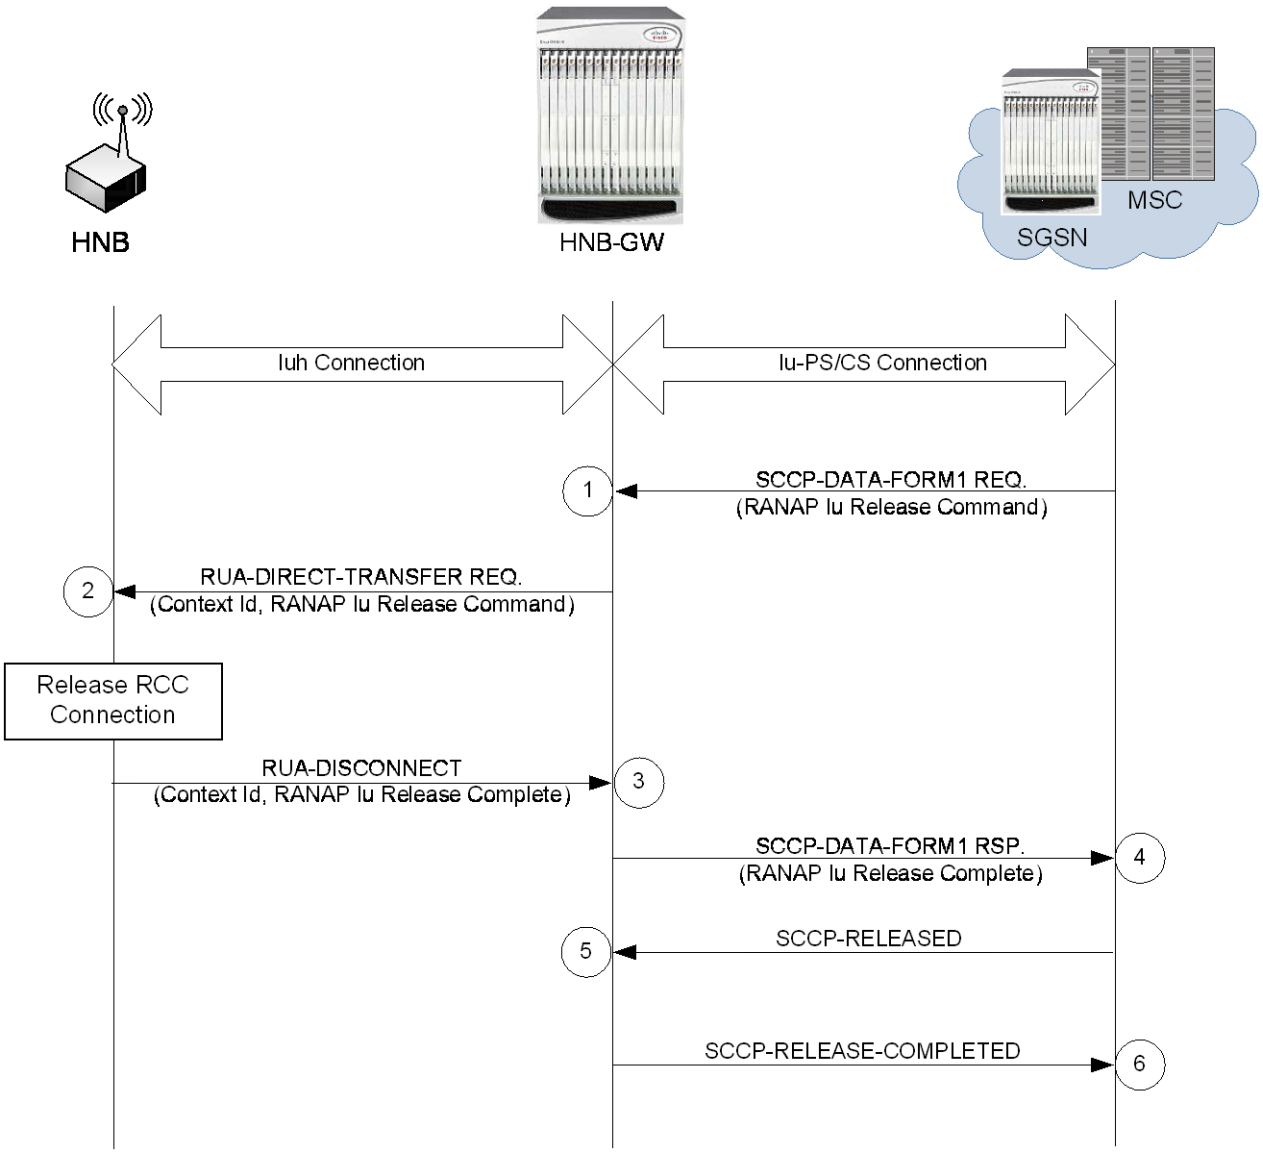
6. The HNB-GW uses the information received in Direct Transfer procedure from CN and forwards the same to HNB through RUA-DIRECT-TRANSFER message.
7. On successful acceptance of RUA-DIRECT-TRANSFER message the HNB responds to HNB-GW and sends RUA-DIRECT-TRANSFER Response message to HNB-GW.
8. On reception of successful acceptance of RUA-DIRECT-TRANSFER message from HNB, the HNB-GW sends SCCP-DATA-FORM1 (Direct Transfer) Response message to CN (SGSN/MSC). This completes the establishment of IuH and IuPS/IuCS connection through HNB, HNB-GW, and SGSN/MSC in core network.

Network Initiated Iu Connection Release Procedure

This procedure is applicable for release of IuH and IuPS/IuCS connection between HNB to HNB-GW and HNB-GW to SGSN/MSC in core network.

The following figure and the text that follows describe the message flow for an Iu connection release procedure initiated by CN (SGSN/MSC).

Figure 125. Network Initiated Iu Connection Release Call Flow



1. User session is established between UE and CN via HNB and HNB-GW over Iu interface and CN (SGSN/MSC) starts RANAP Iu Release procedure with HNB-GW and sends SCCP-DATA-FORM1 REQ with RANAP Iu Release command to HNB-GW.
2. The HNB-GW uses the information received in SCCP-DATA-FORM1 REQ with RANAP Iu Release procedure from CN and forwards the same to HNB through RUA-DIRECT-TRANSFER message with RANAP Iu Release command.
3. On reception of RANAP Iu Release command in RUA-DIRECT-TRANSFER message the HNB triggers the RCC Connection Release procedure and responds to HNB-GW with RANAP Iu Release Complete command in RUA-DISCONNECT Response message.
4. On reception of successful RANAP Iu Release Complete command in RUA-DISCONNECT Response message from HNB, the HNB-GW sends RANAP Iu Release Complete command in SCCP-DATA-FORM1 Response message to CN (SGSN/MSC).
5. On reception of RANAP Iu Release Complete command in SCCP-DATA-FORM1 Response message from HNB-GW, CN sends SCCP-RELEASED message to HNB-GW and triggers the associated SCCP connection.

On reception of SCCP-RELEASED message from CN, the HNB-GW sends RUA-DISCONNECT message to HNB and disconnect the IuH connection with HNB.

6. After successful completion of RUA-DISCONNECT procedure and IuH connection release, HNB-GW sends SCCP-RELEASE-COMPLETE message to CN and HNB-GW confirms the IuPS/IuCS connection released between HNB-GW and CN.

Paging and Serving RNS Relocation Procedures

This section describes the call flow for network-initiated paging and SRNS relocation procedures on HNB-GW.

Following procedure call flows are described for Paging and SRNS relocation procedures between HNB, HNB-GW, and SGSN/MSC in core network:

- Paging Procedure
- SRNS Relocation Procedure

Paging Procedure

This procedure is applicable for establishment of IuH and IuPS/IuCS connection between HNB to HNB-GW and HNB-GW to SGSN/MSC in core network.

The following text describes the call flow for Paging procedure on HNB-GW:

1. HNB-GW receives Paging from SGSN/MSC. HNB-GW finds out if any UE is registered with that IMSI.
2. If a UE is registered then HNB-GW sends the Paging message to the HNB through which the UE is registered.
3. If no registered UE is found then HNB-GW finds out the list of HNBs which have IMSI received in the message in their respective Whitelist.
4. If one or more HNBs were found, and Paging message contained LAI, then HNB-GW compares the HNB's PLMN-ID and LAC values against LAI received in the Paging. The HNB which do not have matching values is dropped from this list.
5. If one or more HNBs were found, and Paging message contained RAI, then HNB-GW compares the HNB's PLMN-ID, LAC and RAC values against RAI received in the Paging. The HNB which do not have matching values is dropped from this list.
6. If Paging message did not have Paging-area then list of HNBs is same as what was found in step 1 otherwise list of HNBs is as found in step 2 or step 3.

   If this list is empty then Paging message is dropped. Otherwise HNB-GW sends Paging message to these HNBs.

SRNS Relocation Procedure

This procedure is applicable for intra-CN or inter-CN handover procedure between HNB to HNB-GW and HNB-GW to SGSN/MSC in core network.

The following text describes the call flow for SRNS relocation procedure on HNB-GW:

1. HNB-GW receives Relocation-Request from SGSN/MSC in case subscriber moves from Macrocell to Femtocell in a connected mode.
2. If the request does not contain IMSI (i.e. for an emergency call), HNB-GW sends Relocation-Request-Reject with an appropriate cause.
3. . If the request contains IMSI, HNB-GW finds the list of registered HNBs which have this IMSI in their white-list. If there is no such HNB found, HNB-GW sends Relocation-Request-Reject with appropriate cause.
4. If there is only one such HNB found which has this IMSI in its white-list, HNB-GW sends Relocation-Request to this HNB.

5. If there are more than one such HNBs found which have this IMSI in their whitelist, then HNBGW looks for Home-HNB for this IMSI. If there are more than one Home-HNB found then HNB-GW sends Relocation-Request-Reject with appropriate cause.
6. If there are multiple HNBs registered which have this IMSI in their whitelist but only one Home-HNB found, HNBGW sends Relocation-Request to this HNB.

RANAP Reset Procedures

This section describes the call flow for various RANAP Reset procedures supported in HNB-GW.

Following procedure call flows are described for RANAP Reset procedures between HNB, HNB-GW, and SGSN/MSC in core network:

- HNB Initiated RANAP Reset Procedure
- CN Initiated RANAP Reset Procedure
- HNB-GW Initiated RANAP Reset Procedure

HNB Initiated RANAP Reset Procedure

This procedure is applicable for HNB-initiated RANAP Reset procedure between HNB, HNB-GW, and SGSN/MSC in core network.

The following text describes the call flow for HNB-initiated RANAP Reset procedure:

1. HNB sends RANAP-RESET command message to HNB-GW for a session.
2. HNB-GW identifies the all affected Iu connection for particular HNB and sends RESET-ACK message to HNB.
3. HNB-GW sends SCCP_Released (SCCP-RLSD) message to CN to release the SCCP connection for each affected Iu connection for particular HNB.
4. CN (SGSN/MSC) sends the SCCP_Release_Complete (SCCP-RLC) message to HNB-GW and release the SCCP connection for requested HNB.

CN Initiated RANAP Reset Procedure

This procedure is applicable for HNB-initiated RANAP Reset procedure between HNB, HNB-GW, and SGSN/MSC in core network.

The following text describes the call flow for HNB-initiated RANAP Reset procedure:

1. CN (SGSN/MSC) sends RANAP-RESET command message to HNB-GW for a session.
2. On receiving RANAP-RESET from CN, the HNB-GW starts Guard timer for configured timeout duration.
3. HNB-GW identifies the all affected Iu connections and sends RUA-DISCONNECT message to HNB.
4. On expiry of Guard timer the HNB-GW sends the RESET-ACK message to CN.

HNB-GW Initiated RANAP Reset Procedure

This procedure is applicable for HNB-GW-initiated RANAP Reset procedure between HNB, HNB-GW, and SGSN/MSC in core network.

The HNB-GW initiates RESET towards CN node in following scenarios:

- The HNB-GW is reloaded or service restarted and SCCP Subsystem Number (SSN) allowed from CN (SGSN/MSC) node is received.

- The received SSN Prohibited or Point-code Address Inaccessible indication comes for a CN node, HNB-GW start a configurable timer.
    - If SSN allowed indication comes before timer expires, the timer is stopped.
    - On timer expiry HNB-GW deletes all SCCP connections towards the CN node.
    - If SSN Allowed indication comes after timer expiry, HNB-GW sends RANAP-RESET command message to the CN node.

The RANAP-RESET from HNB-GW is sent only if HNB-GW-initiated RANAP-RESET is configured in HNB-GW service.

Supported Standards

The HNB-GW complies with the following standards for 3G UMTS Femto wireless data services.

- 3GPP References
- IETF References
- ITU-T Recommendations
- Object Management Group (OMG) Standards

3GPP References

- 3GPP TS 23.003 V8.9.0 (2010-06): 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Numbering, addressing and identification (Release 8)
- 3GPP TS 23.041 V10.3.0 (2012-03): 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Technical realization of Cell Broadcast Service (CBS) (Release 10)
- 3GPP TS 25.412 V8.0.0 (2008-12): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface signalling transport (Release 8)
- 3GPP TS 25.413 V7.9.0 (2008-06): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface RANAP signalling (Release 7)
- 3GPP TS 25.414 V9.0.0 (2009-12): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface data transport and transport signalling (Release 9)
- 3GPP TS 25.415 V8.0.0 (2008-12): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface user plane protocols (Release 8)
- 3GPP TS 25.467 V8.0.0. (2008-12): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home NodeB; Stage 2 (Release 8)
- 3GPP TS 25.467 V9.1.0 (2009-12): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home Node B (HNB); Stage 2 (Release 9)
- 3GPP TS 25.467 V9.3.0 (2010-06): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN architecture for 3G Home Node B (HNB); Stage 2 (Release 9)
- 3GPP TS 25.468 V8.0.0 (2008-12): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN IuH Interface RANAP User Adaptation (RUA) signalling (Release 8)
- 3GPP TS 25.468 V9.0.0 (2009-12): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh Interface RANAP User Adaptation (RUA) signalling (Release 9)
- 3GPP TS 25.468 V9.2.0 (2010-06): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh Interface RANAP User Adaptation (RUA) signalling (Release 9)
- 3GPP TS 25.469 V8.1.0 (2009-03): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh interface Home Node B Application Part (HNBAP) signalling (Release 8)
- 3GPP TS 25.469 V9.0.0 (2009-12): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh interface Home Node B Application Part (HNBAP) signalling (Release 9)

- 3GPP TS 25.469 V9.2.0 (2010-06): 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iuh interface Home Node B (HNB) Application Part (HNBAP) signalling (Release 9)
- 3GPP TS 29.060 V9.0.0 (2009-09): 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface (Release 9)
- 3GPP TR 29.814 V7.1.0 (2007-06): 3rd Generation Partnership Project; Technical Specification Group Core Networks and Terminals Feasibility Study on Bandwidth Savings at Nb Interface with IP transport (Release 7)
- 3GPP TS 33.320 V9.1.0 (2010-13): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Security of Home Node B (HNB) / Home evolved Node B (HeNB) (Release 9)
- 3GPP TS 23.236 V8.0.0 (2008-12): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Intra-domain connection of Radio Access Network(RAN) nodes to multiple Core Network(CN) nodes (Release 8)

IETF References

- RFC-768, User Datagram Protocol (UPD), August 1980
- RFC-791, Internet Protocol (IP), September 1982
- RFC-793, Transmission Control Protocol (TCP), September 1981
- RFC-894, A Standard for the Transmission of IP Datagrams over Ethernet Networks, April 1984
- RFC-1089, SNMP over Ethernet, February 1989
- RFC-1144, Compressing TCP/IP headers for low-speed serial links, February 1990
- RFC-1155, Structure & identification of management information for TCP/IP-based internets, May 1990
- RFC-1157, Simple Network Management Protocol (SNMP) Version 1, May 1990
- RFC-1212, Concise MIB Definitions, March 1991
- RFC-1213, Management Information Base for Network Management of TCP/IP-based Internets: MIB-II, March 1991
- RFC-1215, A Convention for Defining Traps for use with the SNMP, March 1991
- RFC-1224, Techniques for managing asynchronously generated alerts, May 1991
- RFC-1256, ICMP Router Discovery Messages, September 1991
- RFC-1305, Network Time Protocol (Version 3) Specification, Implementation and Analysis, March 1992
- RFC-1398, Definitions of Managed Objects for the Ethernet-Like Interface Types, January 1993
- RFC-1418, SNMP over OSI, March 1993
- RFC-1570, PPP LCP Extensions, January 1994
- RFC-1643, Definitions of Managed Objects for the Ethernet-like Interface Types, July 1994
- RFC-1701, Generic Routing Encapsulation (GRE), October 1994
- RFC-1850, OSPF Version 2 Management Information Base, November 1995
- RFC-1901, Introduction to Community-based SNMPv2, January 1996
- RFC-1902, Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996

- RFC-1903, Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1904, Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1905, Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1906, Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1907, Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1908, Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework, January 1996
- RFC-1918, Address Allocation for Private Internets, February 1996
- RFC-1919, Classical versus Transparent IP Proxies, March 1996
- RFC-2002, IP Mobility Support, May 1995
- RFC-2003, IP Encapsulation within IP, October 1996
- RFC-2004, Minimal Encapsulation within IP, October 1996
- RFC-2005, Applicability Statement for IP Mobility Support, October 1996
- RFC-2118, Microsoft Point-to-Point Compression (MPPC) Protocol, March 1997
- RFC 2131, Dynamic Host Configuration Protocol
- RFC-2136, Dynamic Updates in the Domain Name System (DNS UPDATE)
- RFC-2211, Specification of the Controlled-Load Network Element Service
- RFC-2246, The Transport Layer Security (TLS) Protocol Version 1.0, January 1999
- RFC-2328, OSPF Version 2, April 1998
- RFC-2344, Reverse Tunneling for Mobile IP, May 1998
- RFC-2394, IP Payload Compression Using DEFLATE, December 1998
- RFC 2401, Security Architecture for the Internet Protocol
- RFC 2402, IP Authentication Header (AH)
- RFC 2406, IP Encapsulating Security Payload (ESP)
- RFC 2409, The Internet Key Exchange (IKE)
- RFC-2460, Internet Protocol Version 6 (IPv6)
- RFC-2461, Neighbor Discovery for IPv6
- RFC-2462, IPv6 Stateless Address Autoconfiguration
- RFC-2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
- RFC-2486, The Network Access Identifier (NAI), January 1999
- RFC-2571, An Architecture for Describing SNMP Management Frameworks, April 1999
- RFC-2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP), April 1999
- RFC-2573, SNMP Applications, April 1999

- RFC-2574, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3), April 1999
- RFC-4594, Configuration Guidelines for DiffServ Service Classes
- RFC-2597, Assured Forwarding PHB Group, June 1999
- RFC-2598, Expedited Forwarding PHB, June 1999
- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2661, Layer Two Tunneling Protocol "L2TP", August 1999
- RFC-2697, A Single Rate Three Color Marker, September 1999
- RFC-2698, A Two Rate Three Color Marker, September 1999
- RFC-2784, Generic Routing Encapsulation (GRE) - March 2000, IETF
- RFC-2794, Mobile IP Network Access Identifier Extension for IPv4, March 2000
- RFC-2809, Implementation of L2TP Compulsory Tunneling via RADIUS, April 2000
- RFC-2845, Secret Key Transaction Authentication for DNS (TSIG), May 2000
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000
- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000
- RFC-4960, Stream Control Transmission Protocol
- RFC-3007, Secure Domain Name System (DNS) Dynamic Update, November 2000
- RFC-3012, Mobile IPv4 Challenge/Response Extensions, November 2000
- RFC-3056, Connection of IPv6 Domains via IPv4 Clouds, February 2001
- RFC-3101 OSPF-NSSA Option, January 2003
- RFC-3143, Known HTTP Proxy/Caching Problems, June 2001
- RFC-3193, Securing L2TP using IPSEC, November 2001
- RFC-3314, Recommendations for IPv6 in Third Generation Partnership Project (3GPP) Standards, September 2002
- RFC-3316, Internet Protocol Version 6 (IPv6) for Some Second and Third Generation Cellular Hosts, April 2003
- RFC-3706, A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers, February 2004
- RFC-3543, Registration Revocation in Mobile IPv4, August 2003
- RFC 3588, Diameter Base Protocol, September 2003
- RFC 4006, Diameter Credit-Control Application, August 2005
- RFC-4306, Internet Key Exchange (IKEv2) Protocol, December 2005

ITU-T Recommendations

- ITU-T Recommendation Q.2630.1 - AAL type2 signalling protocol (Capability Set 1)

- ITU-T Recommendation Q.2630.2 - AAL type2 signalling protocol (Capability Set 2)
- ITU-T Recommendation I.361 B-ISDN ATM layer specification
- ITU-T Recommendation I.363.2 B-ISDN ATM Adaptation Layer (AAL) Specification: Type 2 AAL
- ITU-T Recommendation I.366.1 Segmentation and Reassembly Service Specific Convergence Sublayer for the AAL type 2
- ITU-T Recommendation Q.2150.1 AAL type 2 signaling transport converter on broadband MTP
- ITU-T Recommendation E.164 - The international public telecommunication numbering plan
- ITU-T Recommendation E.191 - B-ISDN addressing

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35,Object Management Group

Chapter 18
HRPD Serving Gateway Overview

The ASR 5x00 provides wireless carriers with a flexible solution that functions as an HRPD Serving Gateway (HSGW) in 3GPP2 evolved High Rate Packet Data (eHRPD) wireless data networks.

This overview provides general information about the HSGW including:

- Product Description
- Network Deployment(s)
- Features and Functionality - Base Software
- Features and Functionality - External Application Support
- Features and Functionality - Optional Enhanced Feature Software
- Call Session Procedure Flows
- Supported Standards

Product Description

The HSGW terminates the HRPD access network interface from the Evolved Access Network/Evolved Packet Core Function (eAN/ePCF) and routes UE-originated or terminated packet data traffic.

The HSGW functionality provides interworking of the AT with the 3GPP Evolved Packet System (EPS) architecture and protocols specified in 3GPP 23.402 (mobility, policy control (PCC), and roaming). It supports efficient (seamless) inter-technology mobility between Long Term Evolution (LTE) and HRPD with the following requirements:

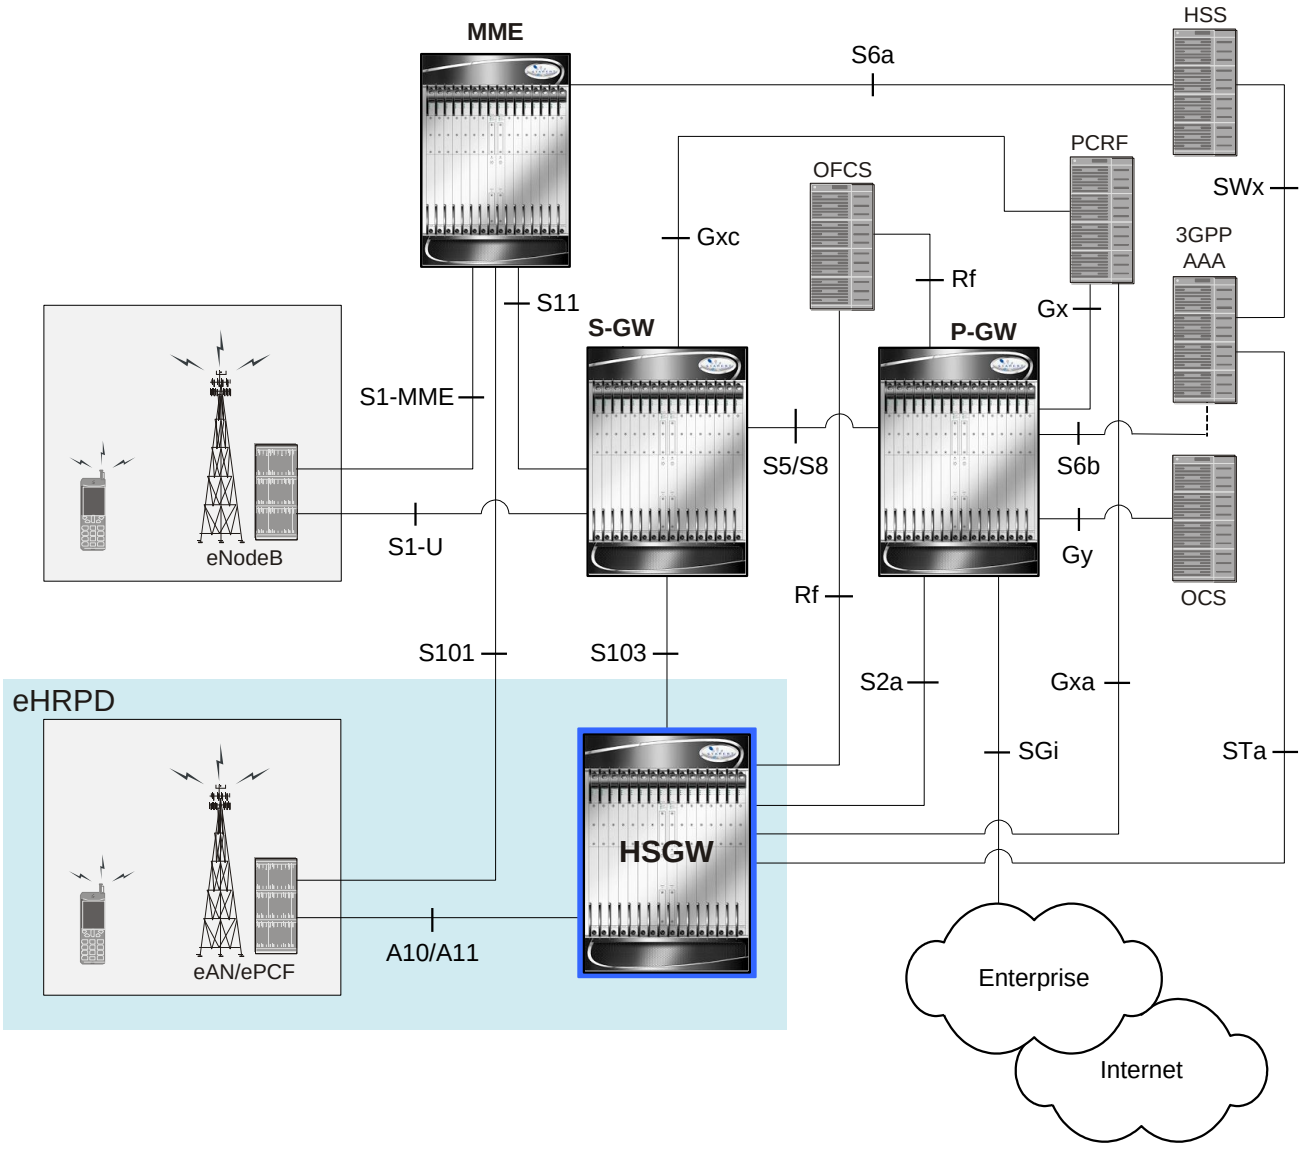
- Sub 300ms bearer interruption
- Inter-technology handoff between 3GPP Enhanced UMTS Terrestrial Radio Access Network (E-UTRAN) and HRPD
- Intra-technology handoff between an HSGW and an existing PDSN
- Support for inter-HSGW fast handoff via Proxy Mobile IPv6 (PMIPv6) Binding Update

The HSGW provides interworking with the eAN/ePCF and the PDN Gateway (P-GW) within the Evolved Packet Core (EPC) or LTE/SAE (4G System Architecture Evolution) core network and performs the following functions:

- Mobility anchoring for inter-eAN handoffs
- Transport level packet marking in the uplink and the downlink, e.g., setting the DiffServ Code Point, based on the QCI of the associated EPS bearer
- Uplink and downlink charging per UE, PDN, and QCI
- Downlink bearer binding based on policy information
- Uplink bearer binding verification with packet dropping of UL traffic that does not comply with established uplink policy
- MAG functions for S2a mobility (i.e., Network-based mobility based on PMIPv6)
- Support for IPv4 and IPv6 address assignment
- EAP Authenticator function
- Policy enforcement functions defined for the Gxa interface
- Robust Header Compression (RoHC)
- Support for VSNCP and VSNP with UE
- Support for packet-based or HDLC-like framing on auxiliary connections
- IPv6 SLACC support, generating RAs responding to RSs

An HSGW also establishes, maintains and terminates link layer sessions to UEs. The HSGW functionality provides interworking of the UE with the 3GPP EPS architecture and protocols. This includes support for mobility, policy control and charging (PCC), access authentication, and roaming. The HSGW also manages inter-HSGW handoffs.

Figure 126. eHRPD Basic Network Topology



Basic Features

Authentication

The HSGW supports the following authentication features:

- EAP over PPP
- UE and HSGW negotiates EAP as the authentication protocol during LCP
- HSGW is the EAP authenticator

- EAP-AKA' (trusted non-3GPP access procedure) as specified in TS 33.402
- EAP is performed between UE and 3GPP AAA over PPP/STa

For more information on authentication features, refer to the Features and Functionality - Base Software section in this overview.

IP Address Allocation

The HSGW supports the following IP address allocation features:

- Support for IPv4 and IPv6 addressing
- Types of PDNs - IPv4, IPv6 or IPv4v6
- IPv6 addressing
  - Interface Identifier assigned during initial attach and used by UE to generate it's link local address
  - HSGW sends the assigned /64 bit prefix in RA to the UE
  - Configure the 128-bits IPv6 address using IPv6 SLAAC (RFC 4862)
  - Optional IPv6 parameter configuration via stateless DHCPv6(Not supported)
- IPv4 address
  - IPv4 address allocation during attach
  - Deferred address allocation using DHCPv4 (Not supported)
  - Option IPv4 parameter configuration via stateless DHCPv4 (Not supported)

Quality of Service

The HSGW supports the following QoS features:

- DSCP Marking
- HRPD Profile ID to QCI Mapping
- QCI to DSCP Mapping
- UE Initiated Dedicated Bearer Resource Establishment

For more information on QoS features, refer to the Features and Functionality - Base Software section in this overview.

AAA, Policy and Charging

The HSGW supports the following AAA, policy and charging features:

- AAA Server Groups
- Dynamic Policy and Charging: Gxa Reference Interface
- EAP Authentication (STa)
- Intelligent Traffic Control
- Rf Diameter Accounting

For more information on policy and charging features, refer to the Features and Functionality - Base Software section in this overview.

Platform Requirements

The HSGW service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the Installation Guide for the chassis and/or contact your Cisco account representative.

Licenses

The HSGW is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

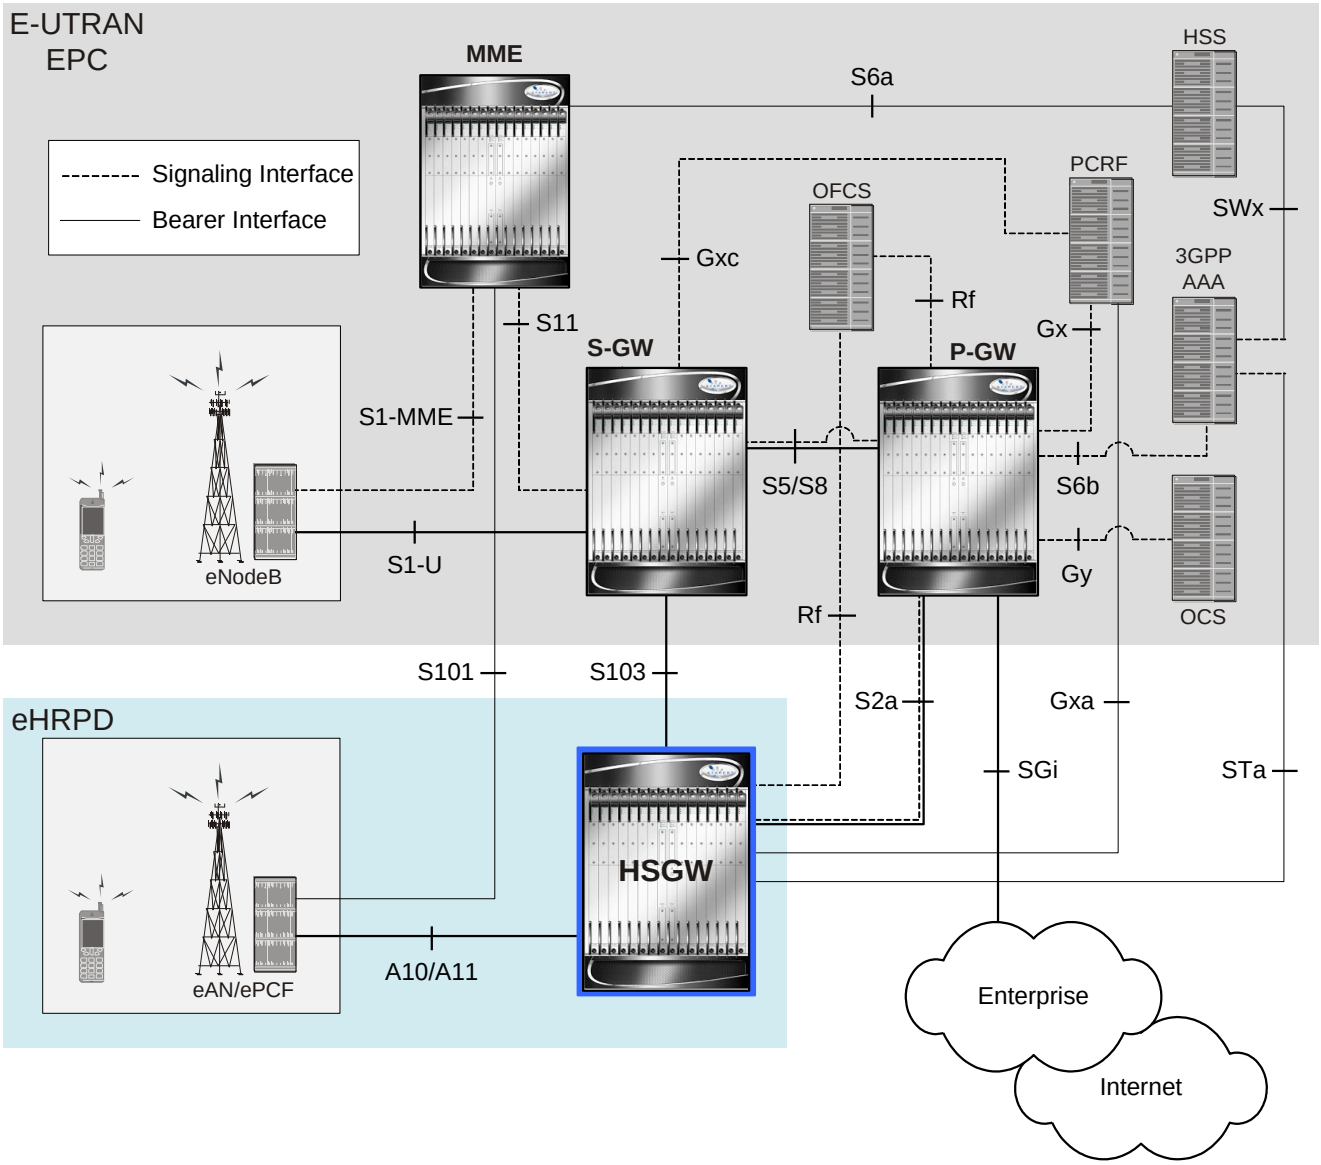
Network Deployment(s)

This section describes the supported interfaces and the deployment scenario of an HSGW in an eHRPD network.

HRPD Serving Gateway in an eHRPD Network

The following figure displays a simplified network view of the HSGW in an eHRPD network and how it interconnects with a 3GPP Evolved-UTRAN/Evolved Packet Core network. The interfaces shown in the following graphic are standards-based and are presented for informational purposes only. For information on interfaces supported by Cisco Systems' HSGW, refer to the next section,.

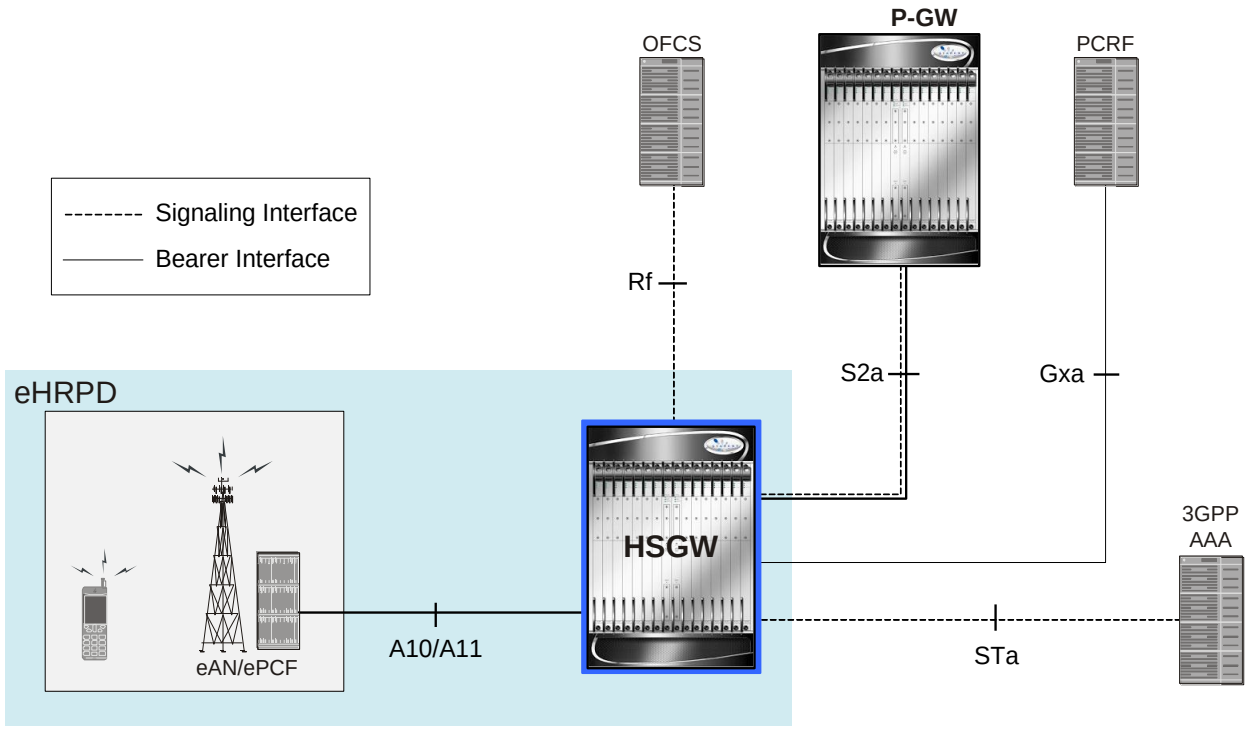
Figure 127. HSGW in an eHRPD Network Architecture



Supported Logical Network Interfaces (Reference Points)

The HSGW supports many of the standards-based logical network interfaces or reference points. The graphic below and following text define the supported interfaces. Basic protocol stacks are also included.

Figure 128. HSGW Supported Network Interfaces

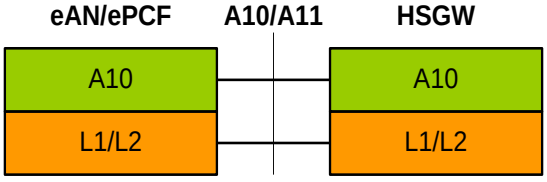


In support of both mobile and network originated subscriber PDP contexts, the HSGW provides the following n etwork interfaces:

A10/A11

This interface exists between the Evolved Access Network/Evolved Packet Control

Function (eAN/ePCF) and the HSGW and implements the A10 (bearer) and A11 (signaling) protocols defined in 3GPP2 specifications.



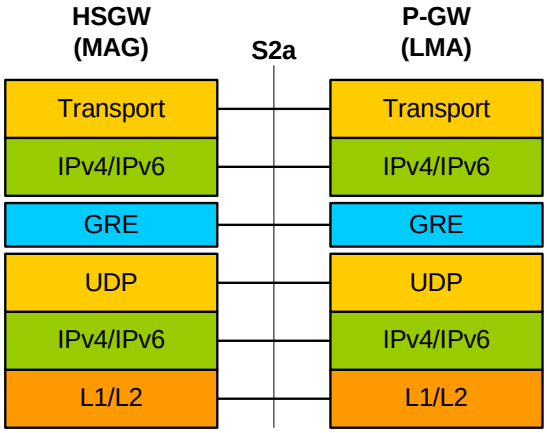
S2a Interface

This reference point supports the bearer interface by providing signaling and mobility support between a trusted non-3GPP access point (HSGW) and the PDN Gateway. It is based on Proxy Mobile IP but also supports Client Mobile IPv4 FA mode which allows connectivity to trusted non-3GPP IP access points that do not support PMIP.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: GRE

- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

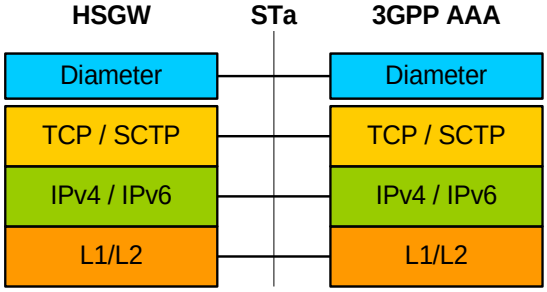


STa Interface

This signaling interface supports Diameter transactions between a 3GPP2 AAA proxy and a 3GPP AAA server. This interface is used for UE authentication and authorization.

Supported protocols:

- Transport Layer: TCP, SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

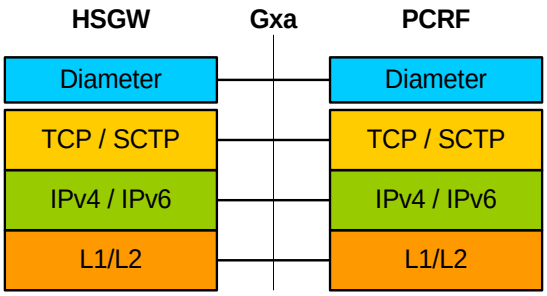


Gxa Interface

This signalling interface supports the transfer of policy control information (QoS) between the HSGW (BBERF) and a PCRF.

Supported protocols:

- Transport Layer: TCP, SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet


HSGW
Gxa
PCRF
Diameter
TCP / SCTP
IPv4 / IPv6
L1/L2
Diameter
TCP / SCTP
IPv4 / IPv6
L1/L2

Features and Functionality - Base Software

This section describes the features and functions supported by default in the base software for the HSGW service and do not require any additional licenses to implement the functionality.

Important: To configure the basic service and functionality on the system for the HSGW service, refer to the configuration examples provided in the *Cisco ASR 5x00 HRPD Serving Gateway Administration Guide*.

The following features are supported and described in this section:

- A10A11
- AAA Server Groups
- ANSI T1.276 Compliance
- Bulk Statistics Support
- Congestion Control
- DSCP Marking
- Dynamic Policy and Charging: Gxa Reference Interface
- EAP Authentication (STa)
- IP Access Control Lists
- Management System
- Mobile IP Registration Revocation
- Multiple PDN Support
- Non-Optimized Inter-HSGW Session Handover
- P-GW Selection (Discovery)
- PPP VSNCP
- Proxy Mobile IPv6 (S2a)
- Rf Diameter Accounting
- Threshold Crossing Alerts (TCA) Support
- UE Initiated Dedicated Bearer Resource Establishment

A10/A11

Provides a lighter weight PPP network control protocol designed to reduce connection set-up latency for delay sensitive multimedia services. Also provides a mechanism to allow user devices in an evolved HRPD network to request one or more PDN connections to an external network.

The HRPD Serving Gateway connects the evolved HRPD access network with the Evolved Packet Core (EPC) as a trusted non-3GPP access network. In an e-HRPD network the A10'/A11' reference interfaces are functionally equivalent to the comparable HRPD interfaces. They are used for connection and bearer establishment procedures. In contrast to the conventional client-based mobility in an HRPD network, mobility management in the e-HRPD application is

network based using Proxy Mobile IPv6 call anchoring between the MAG function on HSGW and LMA on PDN GW. Connections between the UE and HSGW are based on Simple IPv6. A11' signaling carries the IMSI based user identity.

The main A10' connection (SO59) carries PPP traffic including EAP-over-PPP for network authentication. The UE performs LCP negotiation with the HSGW over the main A10' connection. The interface between the e-PCF and HSGW uses GRE encapsulation for A10's. HDLC framing is used on the Main A10 and SO64 auxiliary A10's while SO67 A10 connections use packet based framing. After successful authentication, the HSGW retrieves the QoS profile from the 3GPP HSS and transfers this information via A11' signaling to the e-PCF.

AAA Server Groups

Value-added feature to enable VPN service provisioning for enterprise or MVNO customers. Enables each corporate customer to maintain its own AAA servers with its own unique configurable parameters and custom dictionaries.

This feature provides support for up to 800 AAA server groups and 800 NAS IP addresses that can be provisioned within a single context or across the entire chassis. A total of 128 servers can be assigned to an individual server group. Up to 1,600 accounting, authentication and/or mediation servers are supported per chassis.

ANSI T1.276 Compliance

ANSI T1.276 specifies security measures for Network Elements (NE). In particular it specifies guidelines for password strength, storage, and maintenance security measures.

ANSI T1.276 specifies several measures for password security. These measures include:

- Password strength guidelines
- Password storage guidelines for network elements
- Password maintenance, e.g. periodic forced password changes

These measures are applicable to the ASR 5x00 and the Web Element Manager since both require password authentication. A subset of these guidelines where applicable to each platform will be implemented. A known subset of guidelines, such as certificate authentication, are not applicable to either product. Furthermore, the platforms support a variety of authentication methods such as RADIUS and SSH which are dependent on external elements. ANSI T1.276 compliance in such cases will be the domain of the external element. ANSI T1.276 guidelines will only be implemented for locally configured operators.

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Following is a list of supported schemas for HSGW:

- **Card**: Provides card-level statistics
- **Context**: Provides context-level statistics
- **Diameter-acct**: Provides Diameter Accounting statistics

- **Diameter-auth**: Provides Diameter Authentication statistics
- **ECS**: Provides Enhanced Charging Service statistics
- **HSGW**: Provides HSGW statistics
- **IMSA**: Provides IMS Authorization statistics
- **IP Pool**: Provides IP pool statistics
- **MAG**: Provides Mobile Access Gateway statistics
- **Port**: Provides port-level statistics
- **PPP**: Provides Point-to-Point Protocol statistics
- **RADIUS**: Provides per-RADIUS server statistics
- **RP**: Provides RP statistics
- **System**: Provides system-level statistics

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the chassis or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, chassis host name, chassis uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Important: For more information on bulk statistic configuration, refer to the *Configuring and Maintaining Bulk Statistics* chapter in the *System Administration Guide*.

Congestion Control

The congestion control feature allows you to set policies and thresholds and specify how the system reacts when faced with a heavy load condition.

Congestion control monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are quickly resolved. However, continuous or large numbers of these conditions within a specific time interval may have an impact the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establishes limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operation thresholds that are configured for the system as described in the Thresholding Configuration Guide. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated.

  A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.

  - **Port Utilization Thresholds**: If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
  - **Port-specific Thresholds**: If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
- **Service Congestion Policies:** Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

Important: For more information on congestion control, refer to the *Congestion Control* chapter in the *System Administration Guide*.

DSCP Marking

Provides support for more granular configuration of DSCP marking.

For Interactive Traffic class, the HSGW supports per-HSGW service and per-APN configurable DSCP marking for Uplink and Downlink direction based on Allocation/Retention Priority in addition to the current priorities.

The following matrix may be used to determine the Diffserv markings used based on the configured traffic class and Allocation/Retention Priority:

Table 63. Default DSCP Value Matrix

Allocation Priority	1	2	3
Traffic Handling Priority			
1	ef	ef	ef
2	af21	af21	af21
3	af21	af21	af21

Dynamic Policy and Charging: Gxa Reference Interface

Enables network initiated policy based usage controls for such functions as service data flow authorization for EPS bearers, QCI mapping, modified QoS treatments and per-APN AMBR bandwidth rate enforcement.

In an e-HRPD application, the Gxa reference point is defined to transfer QoS policy information between the PCRF and Bearer Binding Event Reporting Function (BBERF) on the HSGW. In contrast with an S5/S8 GTP network model where the sole policy enforcement point resides on the PGW, the S2a model introduces the additional BBERF function to map EPS bearers to the main and auxiliary A10 connections. Gxa is sometimes referred to as an off-path signaling

interface because no in-band procedure is defined to convey PCC rules via the PMIPv6 S2a reference interface. Gxa is a Diameter based policy signaling interface.

Gxa signaling is used for bearer binding and reporting of events. It provides control over the user plane traffic handling and encompasses the following functionality:

- Provisioning, update and removal of QoS rules from PCRF to BBERF.
- Bearer binding: Associates Policy Charging and Control (PCC) rules with default or dedicated EPS bearers. For a service data flow that is under QoS control, the Bearer Binding Function (BBF) within the HSGW ensures that the service data flow is carried over the bearer with the appropriate QoS service class.
- Bearer retention and teardown procedures
- Event reporting: Transmission of traffic plane events from BBERF to PCRF.
- Service data flow detection for tunneled and un-tunneled service data flows: The HSGW uses service data flow filters received from the PCRF for service data flow detection.
- QoS interworking/mapping between 3GPP QoS (QCI, GBR, MBR) and 3GPP2 ProfileID's

EAP Authentication (STa)

Enables secure user and device level authentication with a 3GPP AAA server or via 3GPP2 AAA proxy and the authenticator in the HSGW.

In an evolved HRPD access network, the HSGW uses the Diameter based STa interface to authenticate subscriber traffic with the 3GPP AAA server. Following completion of the PPP LCP procedures between the UE and HSGW, the HSGW selects EAP-AKA as the method for authenticating the subscriber session. EAP-AKA uses symmetric cryptography and pre-shared keys to derive the security keys between the UE and EAP server. EAP-AKA user identity information (NAI=IMSI) is conveyed over EAP-PPP between the UE and HSGW.

The HSGW represents the EAP authenticator and triggers the identity challenge-response signaling between the UE and back-end 3GPP AAA server. On successful verification of user credentials the 3GPP AAA server obtains the Cipher Key and Integrity Key from the HSS. It uses these keys to derive the Master Session Keys (MSK) that are returned on EAP-Success to the HSGW. The HSGW uses the MSK to derive the Pair-wise Mobility Keys (PMK) that are returned in the Main A10' connection to the e-PCF. The RAN uses these keys to secure traffic transmitted over the wireless access network to the UE.

After the user credentials are verified by the 3GPP AAA and HSS the HSGW returns the PDN address in the VSNCP signaling to the UE. In the e-HRPD connection establishment procedures the PDN address is triggered based on subscription information conveyed over the STa reference interface. Based on the subscription information and requested PDN-Type signaled by the UE, the HSGW informs the PDN GW of the type of required address (v6 HNP and/or IPv4 Home Address Option for dual IPv4/v6 PDNs).

IP Access Control Lists

IP access control lists allow you to set up rules that control the flow of packets into and out of the system based on a variety of IP packet parameters.

IP access lists, or access control lists (ACLs) as they are commonly referred to, are used to control the flow of packets into and out of the system. They are configured on a per-context basis and consist of "rules" (ACL rules) or filters that control the action taken on packets that match the filter criteria. Once configured, an ACL can be applied to any of the following:

- An individual interface

- All traffic facilitated by a context (known as a policy ACL)
- An individual subscriber
- All subscriber sessions facilitated by a specific context

Important: For more information on IP access control lists, refer to the *IP Access Control Lists* chapter in the *System Administration Guide*.

Management System

The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management - focusing on providing superior quality network element (NE) and element management system (Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems - giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

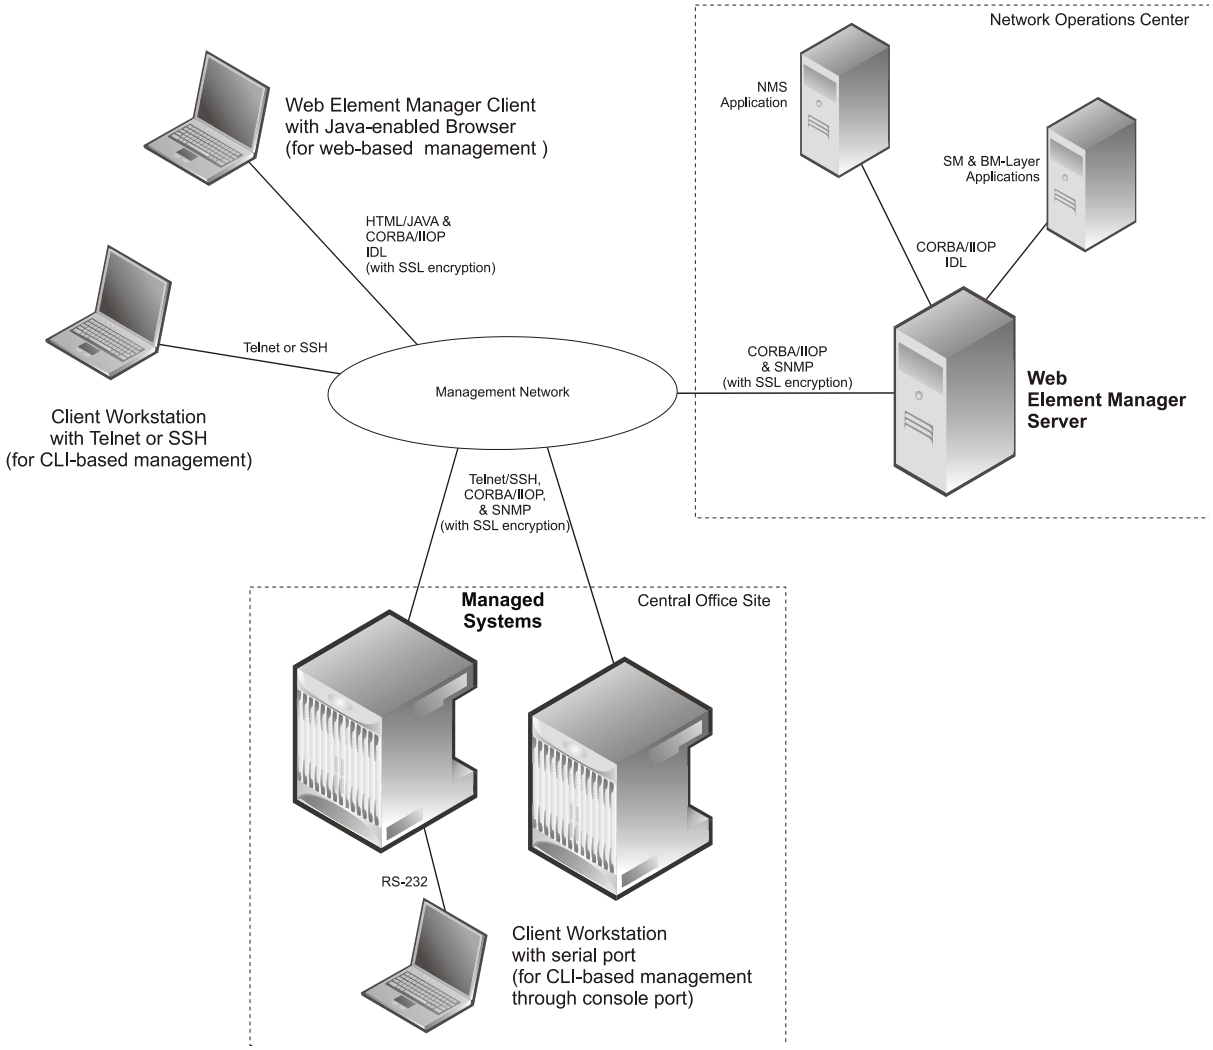
Cisco Systems' O&M module offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces.

These include:

- Using the command line interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager application
- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000
- Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (i.e., Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network.

Figure 129. Element Management Methods



Important: HSGW management functionality is enabled by default for console-based access. For GUI-based management support, refer to the *Web Element Management System* section in this chapter. For more information on command line interface based management, refer to the *Command Line Interface Reference*.

Mobile IP Registration Revocation

Mobile IP registration revocation functionality provides the following benefits:

- Timely release of Mobile IP resources at the HSGW and/or P-GW
- Accurate accounting
- Timely notification to mobile node of change in service

Registration Revocation is a general mechanism whereby either the P-GW or the HSGW providing Mobile IP functionality to the same mobile node can notify the other mobility agent of the termination of a binding. Mobile IP Registration Revocation can be triggered at the HSGW by any of the following:

- Session terminated with mobile node for whatever reason
- Session renegotiation
- Administrative clearing of calls
- Session Manager software task outage resulting in the loss of HSGW sessions (sessions that could not be recovered)

Multiple PDN Support

Enables an APN-based user experience that enables separate connections to be allocated for different services including IMS, Internet, walled garden services, or offdeck content services.

The MAG function on the HSGW can maintain multiple PDN or APN connections for the same user session. The MAG runs a single node level Proxy Mobile IPv6 tunnel for all user sessions toward the LMA function of the PDN GW. When a user wants to establish multiple PDN connections, the MAG brings up the multiple PDN connections over the same PMIPv6 session to one or more PDN GW LMA's. The PDN GW in turn allocates separate IP addresses (Home Network Prefixes) for each PDN connection and each one can run one or multiple EPC default & dedicated bearers. To request the various PDN connections, the MAG includes a common MN-ID and separate Home Network Prefixes, APN's and a Handover Indication Value equal to one in the PMIPv6 Binding Updates.

Non-Optimized Inter-HSGW Session Handover

Enables non-optimized roaming between two eHRPD access networks that lack a relationship of trust and when there are no SLAs in place for low latency hand-offs.

Inter-HSGW hand-overs without context transfers are designed for cases in which the user roams between two eHRPD networks where no established trust relationship exists between the serving and target operator networks. Additionally no H1/H2 optimized hand-over interface exists between the two networks and the Target HSGW requires the UE to perform new PPP LCP and attach procedures. Prior to the hand-off the UE has a complete data path with the remote host and can send and receive packets via the eHRPD access network and HSGW and PGW in the EPC core.

The UE eventually transitions between the Serving and Target access networks in active or dormant mode as identified via A16 or A13 signaling. The Target HSGW receives an A11 Registration Request with VSNCP set to “Hand-Off”. The request includes the IP address of the Serving HSGW, the MSID of the UE and information concerning existing A10 connections. Since the Target HSGW lacks an authentication context for the UE, it sends the LCP config-request to trigger LCP negotiation and new EAP-AKA procedures via the STa reference interface. After EAP success, the UE sends its VSNCP Configure Request with Attach Type equal to “Hand-off”. It also sets the IP address to the previously assigned address in the PDN Address Option. The HSGW initiates PMIPv6 binding update signaling via the S2a interface to the PGW and the PGW responds by sending a PMIPv6 Binding Revocation Indication to the Serving HSGW.

P-GW Selection (Discovery)

Supports the allocation of a P-GW used to provide PDN access to the subscriber. Subscriber information is used via the STa interface from the 3GPP AAA server, which receives subscriber information from the HSS.

The HSGW uses subscriber information provided by the 3GPP AAA server for P-GW selection. PDN subscription contexts provided by the 3GPP AAA server may contain:

1. the IP address of a P-GW

   If the 3GPP AAA server provides the IP address of a P-GW, no further P-GW selection functionality is performed.

2. the identity of a P-GW

   If the P-GW identity is a fully qualified domain name (FQDN) instead of an IP address, the P-GW address is derived by using the Domain Name Service (DNS) function.

3. the identity of an APN

   If only an APN is provided, an APN FQDN constructed for the APN is used to derive the P-GW address through the DNS function. If the DNS function provides a list of P-GW addresses, one P-GW address is selected from this list using the following criteria:

   - topology matching (if enabled)
   - P-GW priority (as configured in DNS records)

PPP VSNCP

VSNCP offers streamlined PPP signaling with fewer messages to reduce connection set-up latency for VoIP services (VORA). VSNCP also includes PDN connection request messages for signaling EPC attachments to external networks.

Vendor Specific Network Control Protocol (VSNCP) provides a PPP vendor protocol in accordance with IETF RFC 3772 that is designed for PDN establishment and is used to encapsulate user datagrams sent over the main A10' connection between the UE and HSGW. The UE uses the VSNCP signaling to request access to a PDN from the HSGW. It encodes one or more PDN-ID's to create multiple VSNCP instances within a PPP connection. Additionally, all PDN connection requests include the requested Access Point Name (APN), PDN Type (IPv4, IPv6 or IPv4/v6) and the PDN address. The UE can also include the Protocol Configuration Options (PCO) in the VSNCP signaling and the HSGW can encode this attribute with information such as primary/secondary DNS server or P-CSCF addresses in the Configuration Acknowledgement response message.

Proxy Mobile IPv6 (S2a)

Provides a mobility management protocol to enable a single LTE-EPC core network to provide the call anchor point for user sessions as the subscriber roams between native EUTRAN and non-native e-HRPD access networks

S2a represents the trusted non-3GPP interface between the LTE-EPC core network and the evolved HRPD network anchored on the HSGW. In the e-HRPD network, network-based mobility provides mobility for IPv6 nodes without host involvement. Proxy Mobile IPv6 extends Mobile IPv6 signaling messages and reuses the HA function (now known as LMA) on PDN Gateway. This approach does not require the mobile node to be involved in the exchange of signaling messages between itself and the Home Agent. A proxy mobility agent (MAG function on HSGW) in the network performs the signaling with the home agent and does the mobility management on behalf of the mobile node attached to the network

The S2a interface uses IPv6 for both control and data. During the PDN connection establishment procedures the PDN Gateway allocates the IPv6 Home Network Prefix (HNP) via Proxy Mobile IPv6 signaling to the HSGW. The HSGW returns the HNP in router advertisement or based on a router solicitation request from the UE. PDN connection release events can be triggered by either the UE, the HSGW or the PGW.

In Proxy Mobile IPv6 applications the HSGW (MAG function) and PDN GW (LMA function) maintain a single shared tunnel and separate GRE keys are allocated in the PMIP Binding Update and Acknowledgement messages to distinguish

between individual subscriber sessions. If the Proxy Mobile IP signaling contains Protocol Configuration Options (PCOs) it can also be used to transfer P-CSCF or DNS server addresses

Rf Diameter Accounting

Provides the framework for offline charging in a packet switched domain. The gateway support nodes use the Rf interface to convey session related, bearer related or service specific charging records to the CGF and billing domain for enabling charging plans.

The Rf reference interface enables offline accounting functions on the HSGW in accordance with 3GPP Release 8 specifications. In an LTE application the same reference interface is also supported on the S-GW and PDN Gateway platforms. The systems use the Charging Trigger Function (CTF) to transfer offline accounting records via a Diameter interface to an adjunct Charging Data Function (CDF) / Charging Gateway Function (CGF). The HSGW and Serving Gateway collect charging information for each mobile subscriber UE pertaining to the radio network usage while the P-GW collects charging information for each mobile subscriber related to the external data network usage.

The ASR 5x00 Charging Trigger Function features dual redundant 140GB RAID hard drives and up to 100GB of capacity on each drive is reserved for writing charging records (CDRs, UDRs, and FDRs) to local file directories with non-volatile persistent memory. The CTF periodically uses the sFTP protocol to push charging files to the CDF/CGF. It is also possible for the CDF/CGF to pull offline accounting records at various intervals or times of the day.

The HSGW, S-GW and P-GW collect information per-user, per IP CAN bearer or per service. Bearer charging is used to collect charging information related to data volumes sent to and received from the UE and categorized by QoS traffic class. Users can be identified by MSISDN or IMSI. Flow Data Records (FDRs) are used to correlate application charging data with EPC bearer usage information. The FDRs contain application level charging information like service identifiers, rating groups, IMS charging identifiers that can be used to identify the application. The FDRs also contain the authorized QoS information (QCI) that was assigned to a given flow. This information is used correlate charging records with EPC bearers.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert**: A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm**: Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps**: SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values.

Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.

- **Logs**: The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

  Logs are supported in both the Alert and the Alarm models.

- **Alarm System**: High threshold alarms generated within the specified polling interval are considered "outstanding" until a the condition no longer exists or a condition clear alarm is generated. "Outstanding" alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

  The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

UE Initiated Dedicated Bearer Resource Establishment

Enables a real-time procedure as applications are started, for the Access Terminal to request the appropriate end-to-end QoS and service treatment to satisfy the expected quality of user experience.

Existing HRPD applications use UE/AT initiated bearer setup procedures. As a migration step toward the EUTRAN-based LTE-SAE network model, the e-HRPD architecture has been designed to support two approaches to resource allocation that include network initiated and UE initiated dedicated bearer establishment. In the StarOS 9.0 release, the HSGW will support only UE initiated bearer creation with negotiated QoS and flow mapping procedures.

After the initial establishment of the e-HRPD radio connection, the UE/AT uses the A11' signaling to establish the default PDN connection with the HSGW. As in the existing EV-DO Rev A network, the UE uses RSVP setup procedures to trigger bearer resource allocation for each additional dedicated EPC bearer. The UE includes the PDN-ID, ProfileID, UL/DL TFT, and ReqID in the reservation.

Each Traffic Flow Template (referred to as Service Data Flow Template in the LTE terminology) consists of an aggregate of one or more packet filters. Each dedicated bearer can contain multiple IP data flows that utilize a common QoS scheduling treatment and reservation priority. If different scheduling classes are needed to optimize the quality of user experience for any service data flows, it is best to provision additional dedicated bearers. The UE maps each TFT packet filter to a Reservation Label/FlowID. The UE sends the TFT to the HSGW to bind the DL SDF IP flows to a FlowID that is in turn mapped to an A10 tunnel toward the RAN. The HSGW uses the RSVP signaling as an event trigger to request Policy Charging and Control (PCC) rules from the PCRF. The HSGW maps the provisioned QoS PCC rules and authorized QCI service class to ProfileID's in the RSVP response to the UE. At the final stage the UE establishes the auxiliary RLP and A10' connection to the HSGW. Once that is accomplished traffic can begin flowing across the dedicated bearer.

Features and Functionality - External Application Support

This section describes the features and functions of external applications supported on the HSGW. These services require additional licenses to implement the functionality.

- Web Element Management System

Web Element Management System

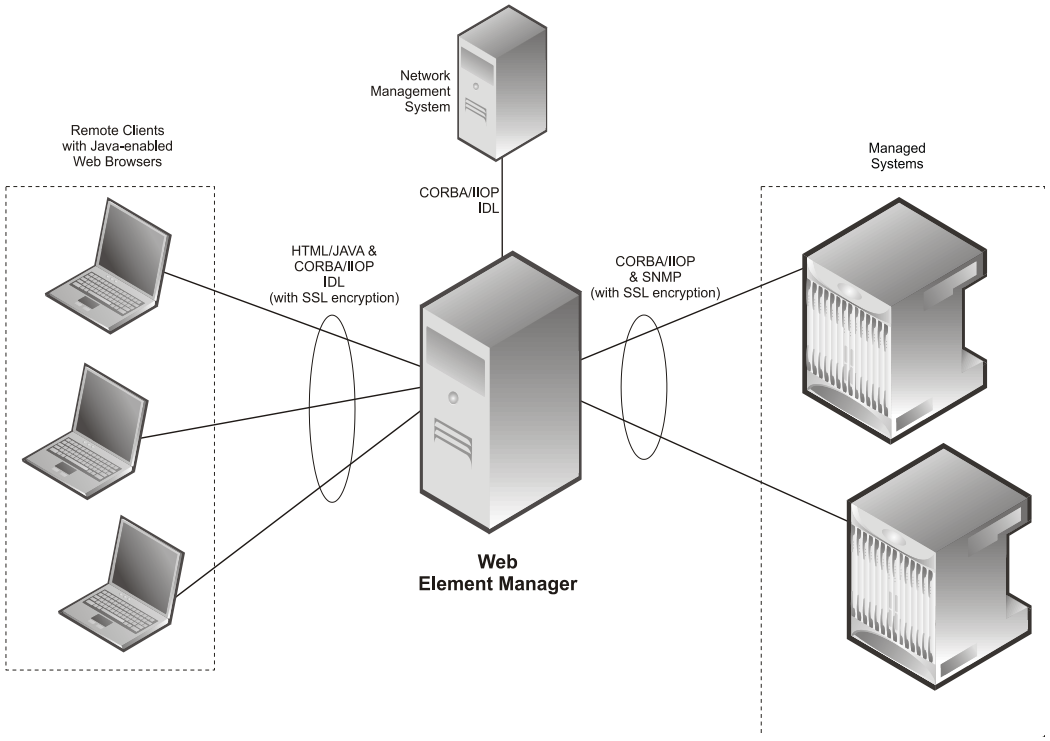
Provides a graphical user interface (GUI) for performing fault, configuration, accounting, performance, and security (FCAPS) management for the ASR 5x00.

The Web Element Manager is a Common Object Request Broker Architecture (CORBA)-based application that provides complete fault, configuration, accounting, performance, and security (FCAPS) management capability for the system.

For maximum flexibility and scalability, the Web Element Manager application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the fault-tolerant Sun® Solaris® operating system.

The following figure demonstrates various interfaces between the Web Element Manager and other network components.

Figure 130. Web Element Manager Network Interfaces



License Keys: A license key is required in order to use the Web Element Manager application. Please contact your local Sales or Support representative for more information.



Important: For more information on WEM support, refer to the *WEM Installation and Administration Guide*.

Features and Functionality - Optional Enhanced Feature Software

This section describes the optional enhanced features and functions for the HSGW service.

Each of the following features require the purchase of an additional license to implement the functionality with the HSGW service.

This section describes following features:

- Intelligent Traffic Control
- IP Header Compression (RoHCv1 for IPv4/IPv6)
- IP Security (IPSec)
- Layer 2 Traffic Management (VLANs)
- Session Recovery Support
- Traffic Policing and Shaping

Intelligent Traffic Control

The feature use license for Intelligent Traffic Control on the HSGW is included in the HSGW session use license.

Intelligent Traffic Control (ITC) supports customizable policy definitions that enforce and manage service level agreements for a subscriber profile, thus enabling differentiated levels of services for native and roaming subscribers.

In 3GPP2, service ITC uses a local policy look-up table and permits either static EV-DO Rev 0 or dynamic EV-DO Rev A policy configuration.

Important: ITC includes the class-map, policy-map and policy-group commands. Currently ITC does not include an external policy server interface.

ITC provides per-subscriber/per-flow traffic policing to control bandwidth and session quotas. Flow-based traffic policing enables the configuring and enforcing bandwidth limitations on individual subscribers, which can be enforced on a per-flow basis on the downlink and the uplink directions.

Flow-based traffic policies are used to support various policy functions like Quality of Service (QoS), and bandwidth, and admission control. It provides the management facility to allocate network resources based on defined traffic-flow, QoS, and security policies.

IP Header Compression (RoHCv1 for IPv4/IPv6)

Use of Robust Header Compression requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

Dynamic header compression contexts enable more efficient memory utilization by allocating and deleting header compression contexts based on the presence/absence of traffic flowing over an S067 A10 bearer connection.

In order to provision VoIP services over an e-HRPD network, the StarOS supports ROHC compression contexts over IPv4 or IPv6 datagrams using the RTP profile over S067 auxiliary A10' connections. The e-HRPD application uses pre-

established SO67 A10' connections for VoIP bearers. A header compression context is allocated for the first time when a new SO67 A10' connection request comes with negotiated ROHC parameters.

In order to optimize memory allocation and system performance, the HSGW uses configured inactivity time of traffic over the bearer to dynamically determine when the ROHC compression context should be removed. This feature is also useful for preserving compression contexts on intra-HSGW call hand-offs. The dynamic header compression context parameters are configured in the ROHC profile that is associated with the subscriber session.

IP Security (IPSec)

Use of Network Domain Security requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

IP Security provides a mechanism for establishing secure tunnels from mobile subscribers to pre-defined endpoints (i.e. enterprise or home networks) in accordance with the following standards:

- RFC 2401, Security Architecture for the Internet Protocol
- RFC 2402, IP Authentication Header (AH)
- RFC 2406, IP Encapsulating Security Payload (ESP)
- RFC 2409, The Internet Key Exchange (IKE)

IP Security (IPSec) is a suite of protocols that interact with one another to provide secure private communications across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways. For IPv4, IKEv1 is used and for IPv6, IKEv2 is supported. IPSec can be implemented on the system for the following applications:

- **PDN Access**: Subscriber IP traffic is routed over an IPSec tunnel from the system to a secure gateway on the packet data network (PDN) as determined by access control list (ACL) criteria.
- **Mobile IP**: Mobile IP control signals and subscriber data is encapsulated in IPSec tunnels that are established between foreign agents (FAs) and home agents (HAs) over the Pi interfaces.

Important: Once an IPSec tunnel is established between an FA and HA for a particular subscriber, all new Mobile IP sessions using the same FA and HA are passed over the tunnel regardless of whether or not IPSec is supported for the new subscriber sessions. Data for existing Mobile IP sessions is unaffected.

Layer 2 Traffic Management (VLANs)

Use of Layer 2 Traffic Management requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

Virtual LANs (VLANs) provide greater flexibility in the configuration and use of contexts and services.

VLANs are configured as "tags" on a per-port basis and allow more complex configurations to be implemented. The VLAN tag allows a single physical port to be bound to multiple logical interfaces that can be configured in different contexts. Therefore, each Ethernet port can be viewed as containing many logical ports when VLAN tags are employed.

Important: For more information on VLAN support, refer to the *VLANs* chapter in the *System Administration Guide*.

Session Recovery Support

The feature use license for Session Recovery on the HSGW is included in the HSGW session use license.

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

This feature is also useful for Software Patch Upgrade activities. If session recovery feature is enabled during the software patch upgrading, it helps to permit preservation of existing sessions on the active PSC during the upgrade process.

Session recovery is performed by mirroring key software processes (e.g. session manager and AAA manager) within the system. These mirrored processes remain in an idle state (in standby-mode), wherein they perform no processing, until they may be needed in the case of a software failure (e.g. a session manager task aborts). The system spawns new instances of "standby mode" session and AAA managers for each active control processor (CP) being used.

Additionally, other key system-level software tasks, such as VPN manager, are performed on a physically separate Packet Service Card (PSC) to ensure that a double software fault (e.g. session manager and VPN manager fails at same time on same card) cannot occur. The PSC used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

The additional hardware resources required for session recovery include a standby system processor card (SPC) and a standby PSC.

There are two modes for Session Recovery.

- **Task recovery mode**: Wherein one or more session manager failures occur and are recovered without the need to use resources on a standby PSC. In this mode, recovery is performed by using the mirrored "standby-mode" session manager task(s) running on active PSCs. The "standby-mode" task is renamed, made active, and is then populated using information from other tasks such as AAA manager.
- **Full PSC recovery mode**: Used when a PSC hardware failure occurs, or when a PSC migration failure happens. In this mode, the standby PSC is made active and the "standby-mode" session manager and AAA manager tasks on the newly activated PSC perform session recovery.

Session/Call state information is saved in the peer AAA manager task because each AAA manager and session manager task is paired together. These pairs are started on physically different PSCs to ensure task recovery.

Important: For more information on session recovery support, refer to the *Session Recovery* chapter in the *System Administration Guide*.

Traffic Policing and Shaping

Use of Per-Subscriber Traffic Policing/Shaping requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

Traffic policing and shaping allows you to manage bandwidth usage on the network and limit bandwidth allowances to subscribers. Shaping allows you to buffer excesses to be delivered at a later time.

Traffic Policing

Traffic policing enables the configuring and enforcing of bandwidth limitations on individual subscribers and/or APNs of a particular traffic class in 3GPP/3GPP2 service.

Bandwidth enforcement is configured and enforced independently on the downlink and the uplink directions.

A Token Bucket Algorithm (a modified trTCM) [RFC2698] is used to implement the Traffic-Policing feature. The algorithm used measures the following criteria when determining how to mark a packet:

- Committed Data Rate (CDR): The guaranteed rate (in bits per second) at which packets can be transmitted/received for the subscriber during the sampling interval.
- Peak Data Rate (PDR): The maximum rate (in bits per second) that subscriber packets can be transmitted/received for the subscriber during the sampling interval.
- Burst-size: The maximum number of bytes that can be transmitted/received for the subscriber during the sampling interval for both committed (CBS) and peak (PBS) rate conditions. This represents the maximum number of tokens that can be placed in the subscriber's "bucket". Note that the committed burst size (CBS) equals the peak burst size (PBS) for each subscriber.

The system can be configured to take any of the following actions on packets that are determined to be in excess or in violation:

- Drop: The offending packet is discarded.
- Transmit: The offending packet is passed.
- Lower the IP Precedence: The packet's ToS bit is set to "0", thus downgrading it to Best Effort, prior to passing the packet. Note that if the packet's ToS bit was already set to "0", this action is equivalent to "Transmit".

Traffic Shaping

Traffic Shaping is a rate limiting method similar to the Traffic Policing, but provides a buffer facility for packets exceeded the configured limit. Once the packet exceeds the data-rate, the packet queued inside the buffer to be delivered at a later time.

The bandwidth enforcement can be done in the downlink and the uplink direction independently. If there is no more buffer space available for subscriber data system can be configured to either drop the packets or kept for the next scheduled traffic session.

Call/Session Procedure Flows

This section provides information on the function of the HSGW in an eHRPD network and presents call procedure flows for different stages of session setup.

The following topics and procedure flows are included:

- Initial Attach with IPv6IPv4 Access
- PMIPv6 Lifetime Extension without Handover
- PDN Connection Release Initiated by UE
- PDN Connection Release Initiated by HSGW
- PDN Connection Release Initiated by P-GW

Initial Attach with IPv6/IPv4 Access

This section describes the procedure of initial attach and session establishment for a subscriber (UE).

Figure 131. Initial Attach with IPv6/IPv4 Access Call Flow

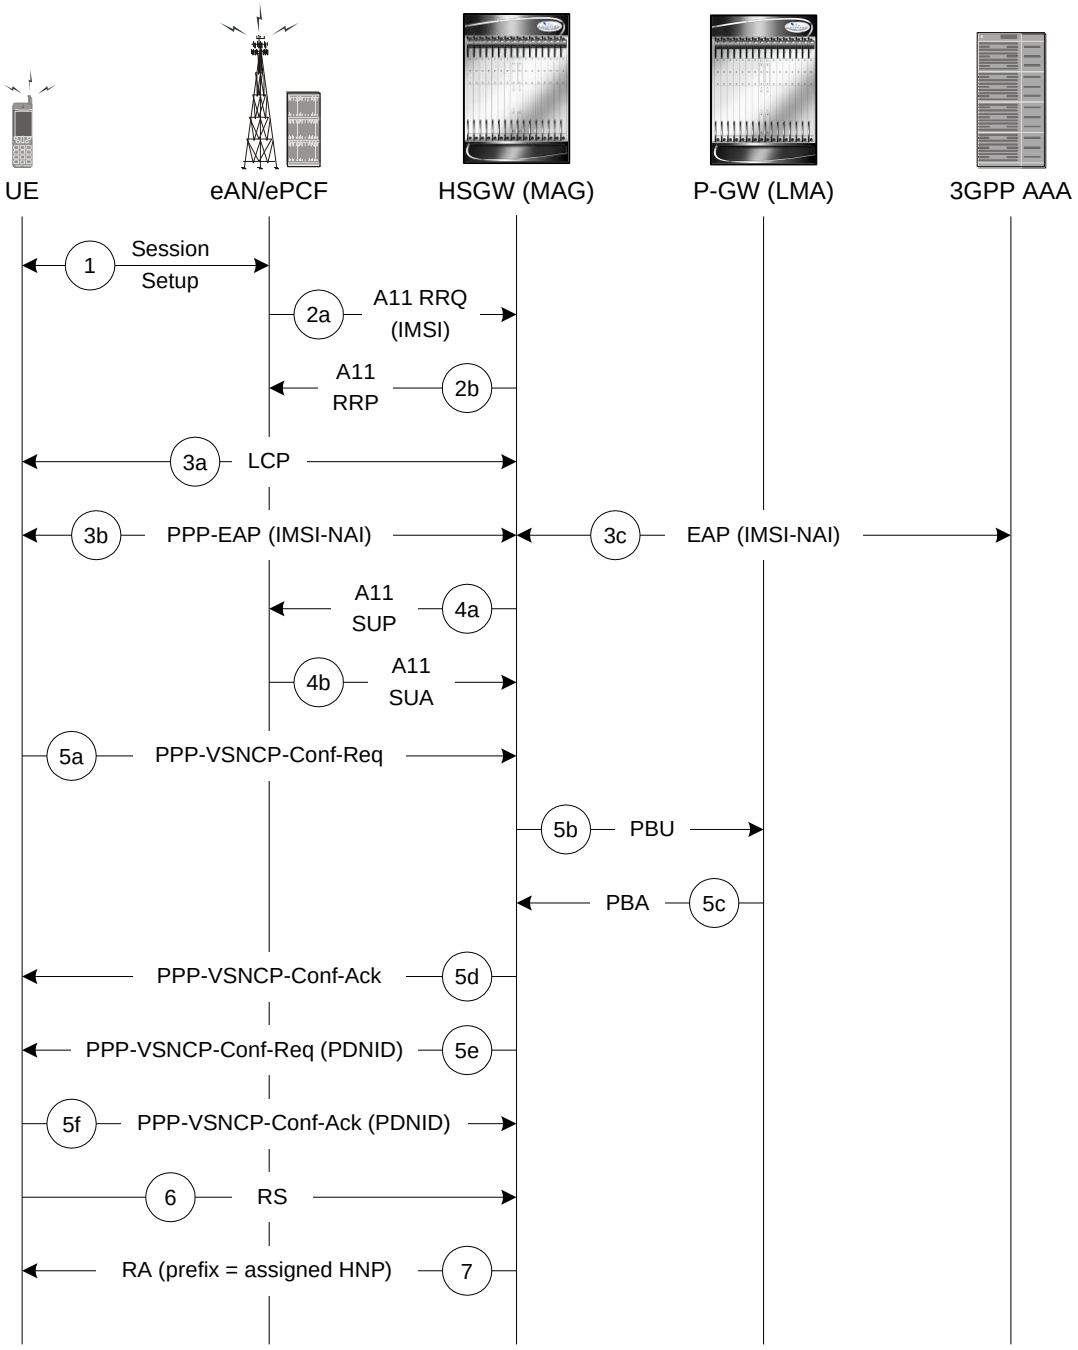


Table 64. Initial Attach with IPv6/IPv4 Access Call Flow Description

Step	Description
1	The subscriber (UE) attaches to the eHRPD network.

Step	Description
2a	The eAN/PCF sends an A11 RRQ to the HSGW. The eAN/PCF includes the true IMSI of the UE in the A11 RRQ.
2b	The HSGW establishes A10s and respond back to the eAN/PCF with an A11 RRP.
3a	The UE performs LCP negotiation with the HSGW over the established main A10.
3b	The UE performs EAP over PPP.
3c	EAP authentication is completed between the UE and the 3GPP AAA. During this transaction, the HSGW receives the subscriber profile from the AAA server.
4a	After receiving the subscriber profile, the HSGW sends the QoS profile in A11 Session Update Message to the eAN/PCF.
4b	The eAN/PCF responds with an A11 Session Update Acknowledgement (SUA).
5a	The UE initiates a PDN connection by sending a PPP-VSNCP-Conf-Req message to the HSGW. The message includes the PDNID of the PDN, APN, PDN-Type=IPv6/[IPv4], PDSN-Address and, optionally, PCO options the UE is expecting from the network.
5b	The HSGW sends a PBU to the P-GW.
5c	The P-GW processes the PBU from the HSGW, assigns an HNP for the connection and responds back to the HSGW with PBA.
5d	The HSGW responds to the VSNCP Conf Req with a VSNCP Conf Ack.
5e	The HSGW sends a PPP-VSNCP-Conf-Req to the UE to complete PPP VSNCP negotiation.
5f	The UE completes VSNCP negotiation by returning a PPP-VSNCP-Conf-Ack.
6	The UE optionally sends a Router Solicitation (RS) message.
7	The HSGW sends a Router Advertisement (RA) message with the assigned Prefix.

PMIPv6 Lifetime Extension without Handover

This section describes the procedure of a session registration lifetime extension by the P-GW without the occurrence of a handover.

Figure 132. PMIPv6 Lifetime Extension (without handover) Call Flow

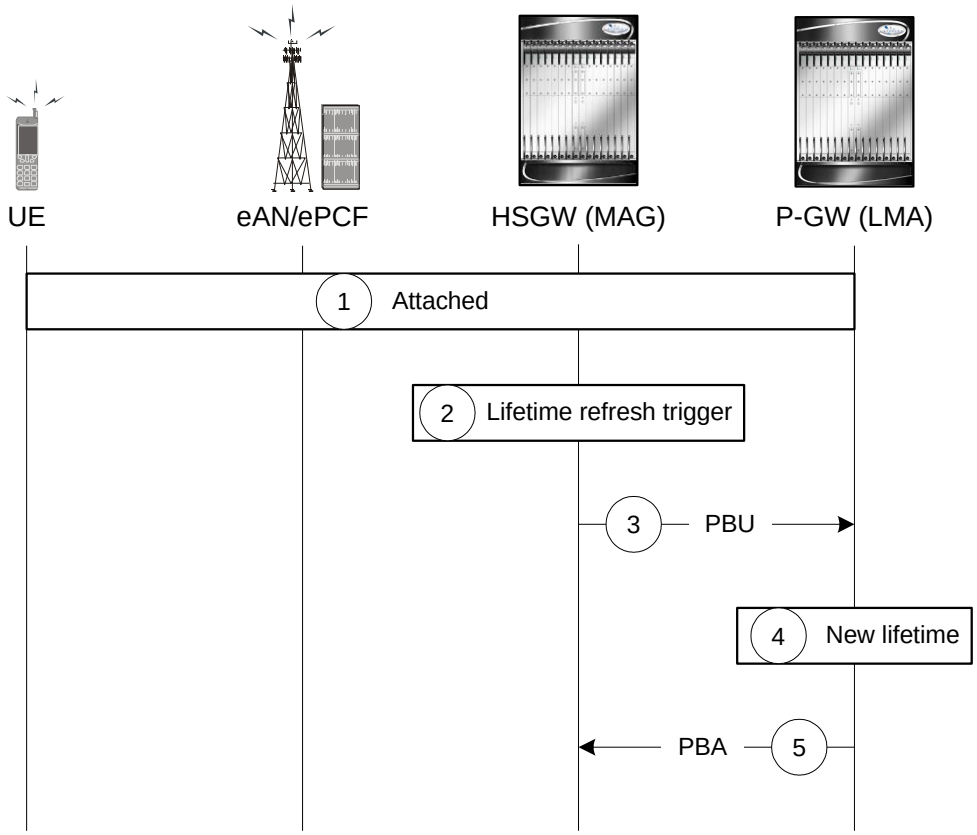


Table 65. PMIPv6 Lifetime Extension (without handover) Call Flow Description

Step	Description
1	The UE is attached to the EPC and has a PDN connection with the P-GW where PDNID=x and an APN with assigned HNP.
2	The HSGW MAG service registration lifetime nears expiration and triggers a renewal request for the LMA.
3	The MAG service sends a Proxy Binding Update (PBU) to the P-GW LMA service with the following attributes: Lifetime, MNID, APN, ATT=HRPD, HNP.
4	The P-GW LMA service updates the Binding Cache Entry (BCE) with the new granted lifetime.
5	The P-GW responds with a Proxy Binding Acknowledgement (PBA) with the following attributes: Lifetime, MNID, APN.

PDN Connection Release Initiated by UE

This section describes the procedure of a session release by the UE.

Figure 133. PDN Connection Release by the UE Call Flow

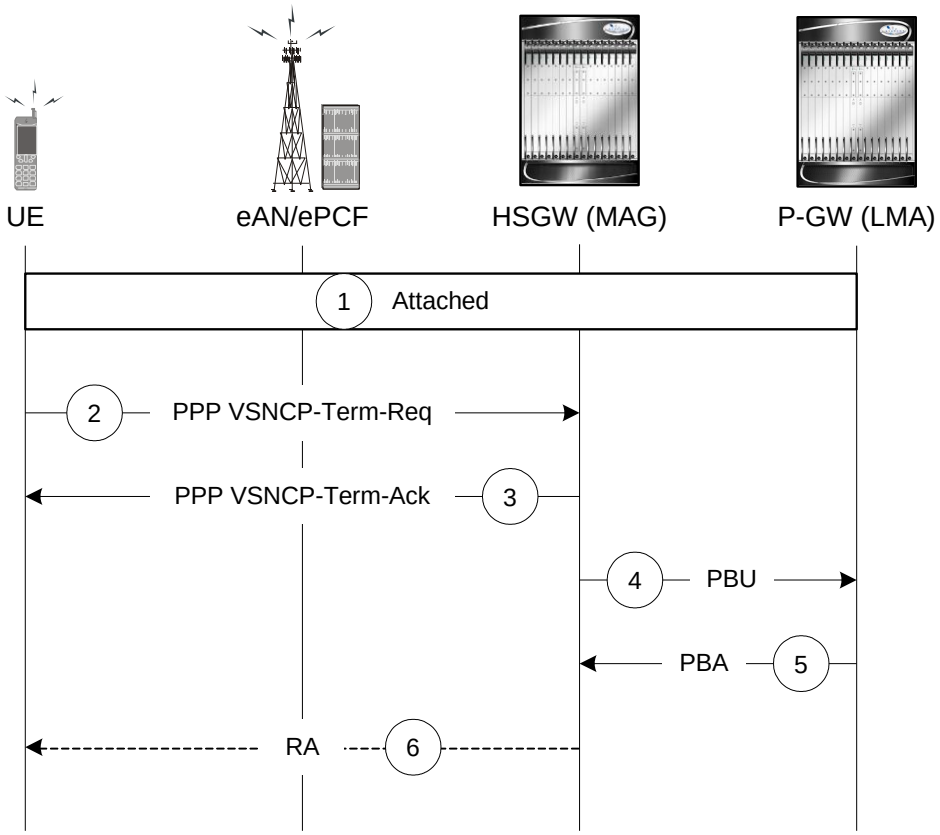


Table 66. PDN Connection Release by the UE Call Flow Description

Step	Description
1	The UE is attached to the EPC and has a PDN connection with the P-GW for PDN-ID=x and APN with assigned HNP.
2	The UE decides to disconnect from the PDN and sends a PPP VSNCP-Term-Req with PDNID=x.
3	The HSGW starts disconnecting the PDN connection and sends a PPP-VSNCP-Term-Ack to the UE (also with PDNID=x).
4	The HSGW begins the tear down of the PMIP session by sending a PBU Deregistration to the P-GW with the following attributes: Lifetime=0, MNID, APN, ATT=HRPD, HNP. The PBU Deregistration message should contain all the mobility options that were present in the initial PBU that created the binding.
5	The P-GW looks up the Binding Cache Entry (BCE) based on the HNP, deletes the binding, and responds to the HSGW with a Deregistration PBA with the same attributes (Lifetime=0, MNID, APN, ATT=HRPD, HNP).
6	The HSGW optionally sends a Router Advertisement (RA) with assigned HNP and prefix lifetime=0.

PDN Connection Release Initiated by HSGW

This section describes the procedure of a session release by the HSGW.

Figure 134. PDN Connection Release by the HSGW Call Flow

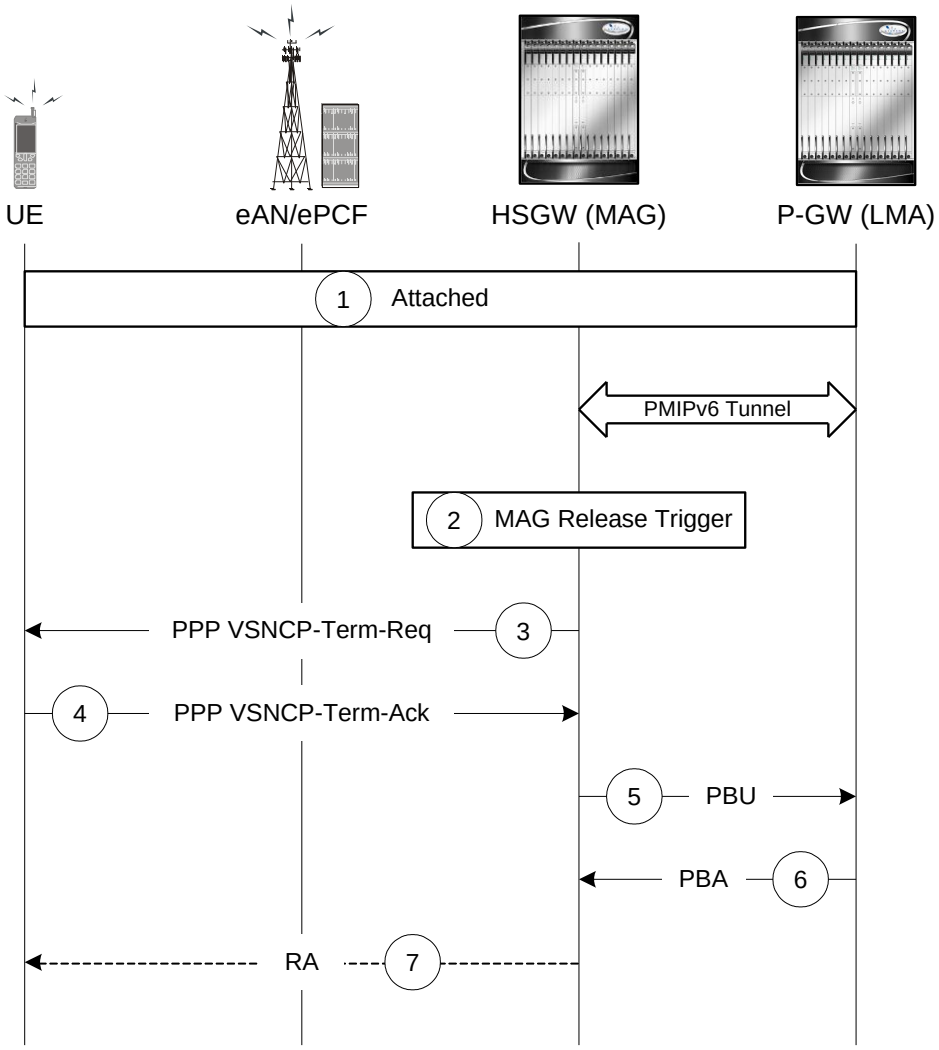


Table 67. PDN Connection Release by the HSGW Call Flow Description

Step	Description
1	The UE is attached to the EPC and has a PDN connection with the P-GW for PDN-ID=x and APN with assigned HNP.
2	The HSGW MAG service triggers a disconnect of the PDN connection for PDNID=x.
3	The HSGW sends a PPP VSNCP-Term-Req with PDNID=x to the UE.
4	The UE acknowledges the receipt of the request with a VSNCP-Term-Ack (PDNID=x).
5	The HSGW begins the tear down of the PMIP session by sending a PBU Deregistration to the P-GW with the following attributes: Lifetime=0, MNID, APN, HNP. The PBU Deregistration message should contain all the mobility options that were present in the initial PBU that created the binding.
6	The P-GW looks up the BCE based on the HNP, deletes the binding, and responds to the HSGW with a Deregistration PBA with the same attributes (Lifetime=0, MNID, APN, ATT=HRPD, HNP).

Step	Description
7	The HSGW optionally sends a Router Advertisement (RA) with assigned HNP and prefix lifetime=0.

PDN Connection Release Initiated by P-GW

This section describes the procedure of a session release by the P-GW.

Figure 135. PDN Connection Release by the HSGW Call Flow

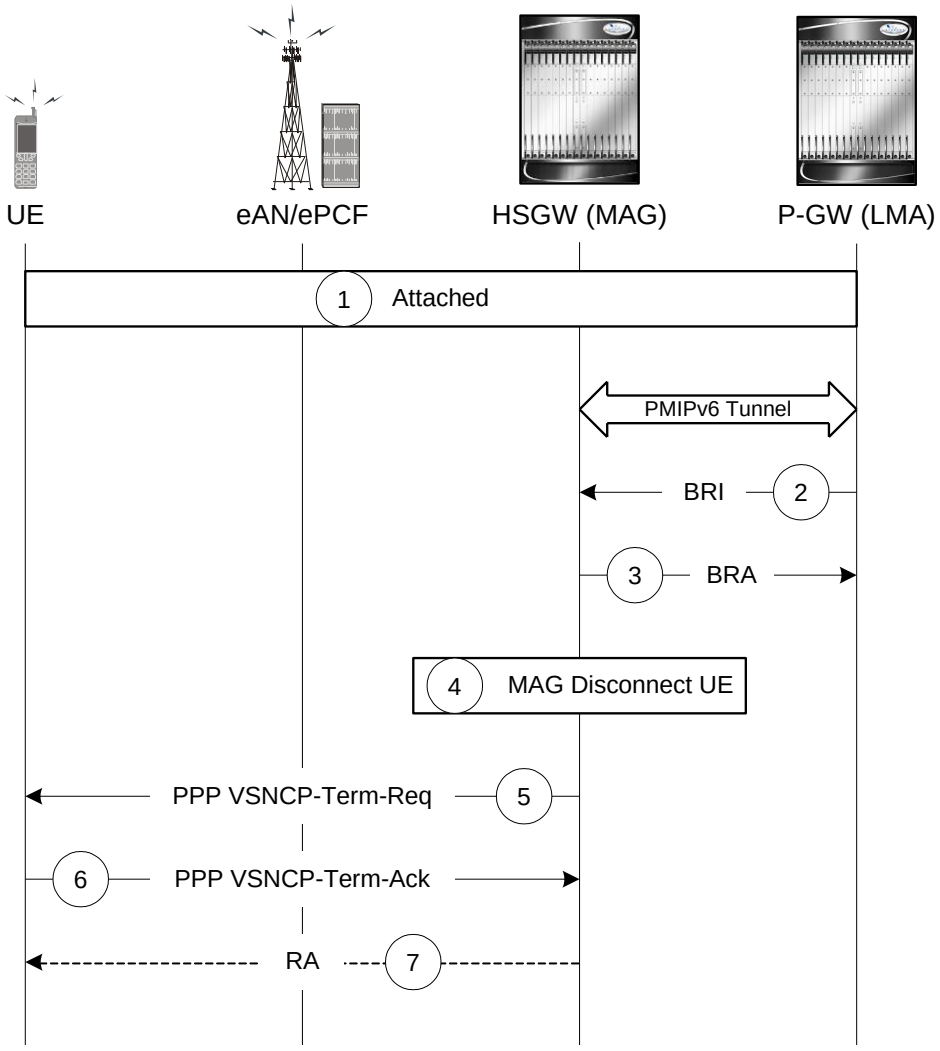


Table 68. PDN Connection Release by the HSGW Call Flow Description

Step	Description

Step	Description
1	The UE is attached to the EPC and has a PDN connection with the P-GW for PDN-ID=x and APN with assigned HNP.
2	A PGW trigger causes a disconnect of the PDN connection for PDNID=x and the PGW sends a Binding Revocation Indication (BRI) message to the HSGW with the following attributes: MNID, APN, HNP.
3	The HSGW responds to the BRI message with a Binding Revocation Acknowledgement (BRA) message with the sane attributes (MNID, APN, HNP).
4	The HSGW MAG service triggers a disconnect of the UE PDN connection for PDNID=x.
5	The HSGW sends a PPP VSNCP-Term-Req with PDNID=x to the UE.
6	The UE acknowledges the receipt of the request with a VSNCP-Term-Ack (PDNID=x).
7	The HSGW optionally sends a Router Advertisement (RA) with assigned HNP and prefix lifetime=0.

Supported Standards

The HSGW complies with the following standards.

- 3GPP References
- Release 8 3GPP References
- IETF References
- Object Management Group (OMG) Standards

Release 8 3GPP References

Important: The HSGW currently supports the following Release 8 3GPP specifications. Most 3GPP specifications are also used for 3GPP2 support; any specifications that are unique to 3GPP2 are listed under *3GPP2 References*.

- 3GPP TR 23.401 General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
- 3GPP TS 23.402 Architecture enhancements for non-3GPP accesses
- 3GPP TS 29.273 Evolved Packet System (EPS);3GPP EPS AAA interfaces
- 3GPP TS 29.275 Proxy Mobile IPv6 (PMIPv6) based Mobility and Tunnelling protocols; Stage 3
- 3GPP TS 32.299 Rf Offline Accounting Interface

3GPP2 References

- A.S0008-C v1.0: Interoperability Specification (IOS) for High Rate Packet Data (HRPD) Radio Access Network Interfaces with Session Control in the Access Network, August 2007. (HRPD IOS)
- A.S0009-C v1.0: Interoperability Specification (IOS) for High Rate Packet Data (HRPD) Radio Access Network Interfaces with Session Control in the Packet Control Function, August 2007. (HRPD IOS)
- A.S0017-D v1.0: Interoperability Specification (IOS) for cdma2000 Access Network Interfaces - Part 7 (A10 and A11 Interfaces), June, 2007.
- A.S0022-0 v1.0: E-UTRAN - HRPD Connectivity and Interworking: Access Network Aspects (E-UTRAN – HRPD IOS), March 2009.
- X.P0057-0 v0.11.0 E-UTRAN - eHRPD Connectivity and Interworking: Core Network Aspects
- X.S0011-001-D v1.0: cdma2000 Wireless IP Network Standard: Introduction, February, 2006.
- X.S0057-0 v3.0: E-UTRAN - eHRPD Connectivity and Interworking: Core Network Aspects, September 17, 2010

IETF References

- RFC 1661 (July 1994): The Point-to-Point Protocol (PPP)
- RFC 2205 (September 1997): Resource Reservation Protocol (RSVP)
- RFC 2473 (December 1998): Generic Packet Tunneling in IPv6 Specification
- RFC 3095 (July 2001): RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed
- RFC 3748 (June 2004): Extensible Authentication Protocol (EAP)
- RFC 3772 (May 2004): PPP Vendor Protocol
- RFC 3775 (June 2004): Mobility Support in IPv6
- RFC 4283 (November 2005): Mobile Node Identifier Option for Mobile IPv6 (MIPv6)
- RFC 5094 (February 2008): Service Selection for Mobile IPv6
- RFC 5149 (December 2007): Mobile IPv6 Vendor Specific Option
- RFC 5213 (August 2008): Proxy Mobile IPv6
- Internet-Draft (draft-ietf-netlmm-pmip6-ipv4-support-09.txt): IPv4 Support for Proxy Mobile IPv6
- Internet-Draft (draft-ietf-netlmm-grekey-option-06.txt): GRE Key Option for Proxy Mobile IPv6
- Internet-Draft (draft-meghana-netlmm-pmipv6-mipv4-00): Proxy Mobile IPv6 and Mobile IPv4 interworking
- Internet-Draft (draft-ietf-mip6-nemo-v4traversal-06.txt): Mobile IPv6 support for dual stack Hosts and Routers (DSMIPv6)
- Internet-Draft (draft-ietf-netlmm-proxymip6-07.txt): Proxy Mobile IPv6
- Internet-Draft (draft arkko-eap-aka-kdf): Improved Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA)
- Internet-Draft (draft-muhanna-mext-binding-revocation-01): Binding Revocation for IPv6 Mobility

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

Chapter 19
InTracer Overview

This chapter provides an overview of the Cisco InTracer application and describes its architecture.

It includes the following sections:

- Introduction
- InTracer Architecture
- InTracer System Requirements
- Supported Features

Introduction

Cisco InTracer is a high-performance subscriber troubleshooting and monitoring solution. It performs call tracing, control data acquisition, processing and analysis of both active and historical subscriber sessions. This provides a framework for operators to analyze and investigate call flows and call events for subscriber sessions in **near real time**.

The InTracer solution consists of 2 basic parts:

- The InTracer Client part runs on the Cisco gateway(s) that needs to be enabled and configured to start sending subscriber traces (control plane information for sessions) to InTracer Server.
- The InTracer Server part runs on an external box sitting close to the gateway to process and store these subscriber traces.

Important: External application is now supported by the Cisco MITG RHEL v5.5 OS on selected Cisco UCS servers. The Cisco MITG RHEL v5.5 OS is a custom image that contains only those software packages required to support compatible Cisco MITG external software applications. Users must not install any other applications on servers running the Cisco MITG RHEL v5.5 OS. For detailed software compatibility information, refer to the "Cisco MITG RHEL v5.5 OS Application Note."

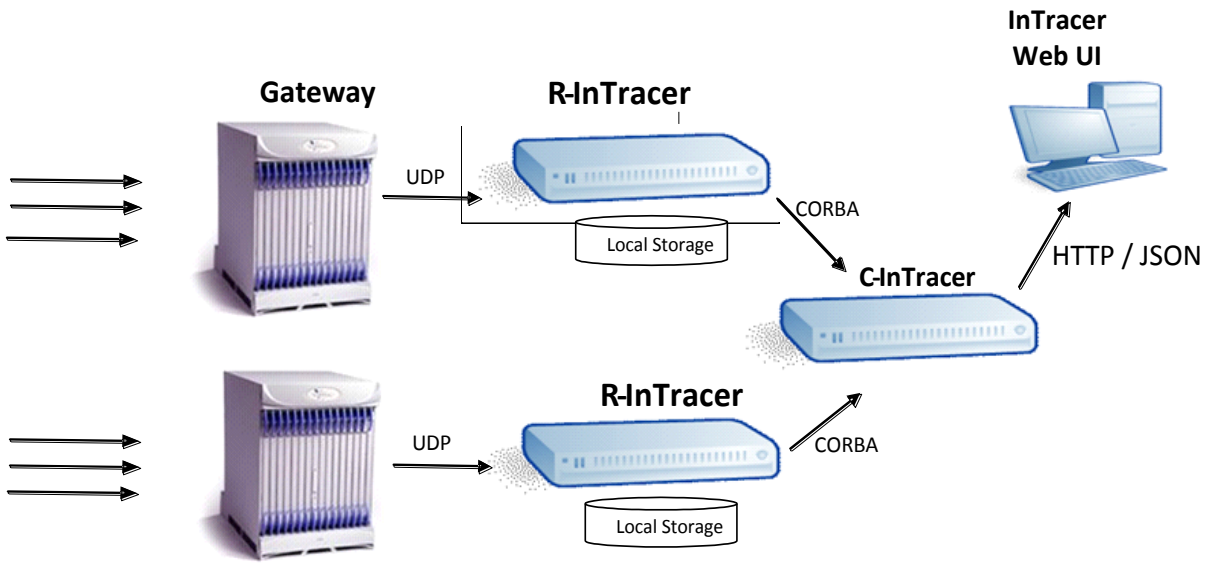
InTracer Architecture

This section describes the InTracer Architecture.

Nodal Trace

The following figure shows a high level view of architecture of InTracer with other components in a deployment scenario.

Figure 136. Nodal Trace

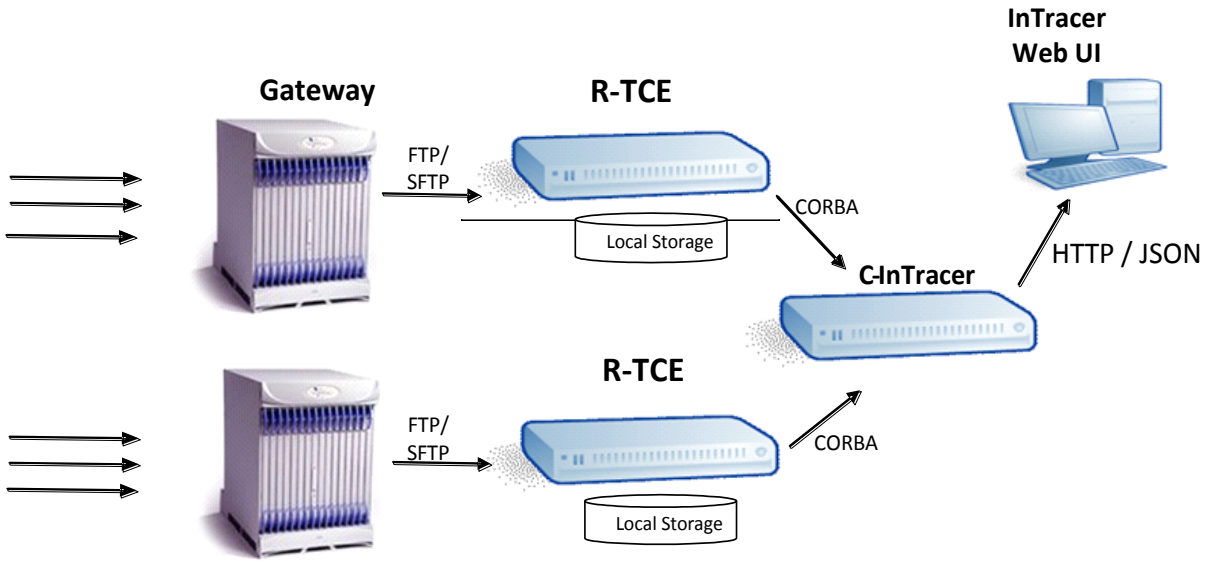


For Nodal trace, the communication between the Gateway and the R-InTracer / R-TCE happens over a proprietary protocol on top of UDP.

3GPP Trace

In 3GPP Trace communication happens through XML files being FTP pushed from the gateway to the R-TCE. R-TCE receives the FTP pushed files from the gateway and parses them to store their data into the database.

Figure 137. 3GPP Trace



InTracer Components

This section provides the physical / logical components of InTracer solution.

Component	Nodal Trace	3GPP Trace
InTracer client	When configured sends a copy of all signaling packets for all sessions to the R-InTracer over UDP.	When configured, sends a copy of all signaling packets for activated trace sessions to the R-TCE by pushing the generated XML trace files via FTP/SFTP.
R-InTracer (Regional InTracer)	Receives packets over UDP from configured InTracer client, processes and stores them for the configured storage period.	N/A
R-TCE(Regional Trace Collection Entity)	N/A	Receives XML trace files via FTP/SFTP from configured InTracer client, parses the files and stores the information for the configured storage period.
C-InTracer (Central InTracer)	This component provides configuration and management in the InTracer system. It is capable of fetching data from all configured R-InTracers and displaying coalesced information to the end-user.	This component provides configuration and management in the InTracer system. It is capable of fetching data from all configured R-TCEs and displaying coalesced information to the end-user. In addition to this, it also provides an interface for activating/de-activating subscriber traces.
InTracer Web User Interface	A unified web-interface of the product providing configuration, management and data-querying options.	A unified web-interface of the product providing configuration, management and data-querying options, along with Activation / De-activation of subscriber traces.

InTracer Client Nodal Trace

Generic Features

The following generic features are supported by the InTracer Client for Nodal Trace:

- To send a copy of all signaling packets for each subscriber session to configured R-InTracer
- To check the health of configured R-InTracer (up and running)

Following are the configuration items for the InTracer Client for Nodal Trace:

- Enabling InTracer
- R-InTracer configuration
- UDP heartbeat configuration parameters

Important: Refer to the appropriate System and Product Administration Guides for detailed information.

Specific Features

Following are the additional configuration items to support the above features:

- License to export SA-keys (on the gateway)
- Enabling SA-keys in InTracer configuration

Important: Refer to the appropriate System and Product Administration Guides for detailed information.

InTracer Client 3GPP Trace

The following generic features are supported by the InTracer Client for 3GPP Trace :

- To send a copy of all transmitted signaling packets for activated trace sessions to the configured R-TCE (3GPP TCE - Trace Collection Entity)
- To send trace activation/deactivation failure notifications to the configured R-TCE

  Following are the configuration items for the InTracer Client for 3GPP Trace:

  - Configuration of of global trace parameters
  - Configuration of TCE profiles (R-TCE configuration)
  - Management-based Activation / Deactivation of Trace sessions

Important: Refer to the appropriate System and Product Administration Guides for detailed information

R-InTracer / R-TCE

The basic functionality of the R-InTracer / R-TCE is to receive signaling packets from the gateways and process them. It stores all the packets along with additional correlation information either received along with the packet or derived after processing the packet.

Following are the features and functions supported by the R-InTracer / R-TCE:

Feature	Configurable	Default	Nodal Trace	3GPP Trace
Event Processing	Yes		Yes	Yes
Event Storage	Yes	30 Days	Yes	Yes
Packet-loss detection	Yes	Enabled	Yes	Yes
Health Monitoring	Yes	Enabled	Yes	Yes

Feature	Configurable	Default	Nodal Trace	3GPP Trace
Resource Monitoring	Yes	Enabled (Default threshold values)	Yes	Yes
Maintenance Activities	Yes	Enabled (Daily midnight)	Yes	Yes
Statistical Counters	No	Always	Yes	Yes

Event Processing

The basic task of the R-InTracer / R-TCE is to receive signaling packets from the gateways and process them. It then persistently stores all these packets along with additional correlation information either received along with the packet or derived after processing the packet.

Event Storage

The R-InTracer / R-TCE stores processed packets along with other correlation information to persistent storage. The storage period is a configurable value between 2 to 236 days, default value is 30 days.

Important: If event information exceeds the storage period, it will be automatically deleted by the system even if the system has sufficient additional storage available.

When the R-InTracer / R-TCE receives packets from the gateway (either through UDP sockets in case of Nodal trace or through XML files in case of 3GPP trace), it requires some time before it persistently stores these packets and makes them available for query. The possible reasons for this delay are:

- It needs to take care of out-of-order packet delivery over UDP in case of Nodal trace
- It needs to detect packet-losses and take appropriate actions in case of Nodal trace
- It needs to perform event correlation

For better search, each query specified by the operator should be bounded by time intervals. The shorter the interval specified, the lesser the amount of data to be searched which results in faster query response.

Important: InTracer is a **best-effort** solution and there can be possible losses of packet data even after successfully processed by the R-InTracer / R-TCE. Such scenarios occur since some packet information is cached and not persistently stored immediately. All efforts are made internally to minimize such scenarios.

Event Filtering

Event filtering allows operators to configure which packets to be processed/ignored by the R-InTracer / R-TCE. Filtering can be done based on the event type and/or protocol type. The R-InTracer / R-TCE maintains statistical counters about the filtered packets, which can later be used to verify this behaviour.

Important: Events that are filtered out are not stored on persistent storage, hence cannot be queried later.

Packet-loss detection

Signaling packets may be received by the R-InTracer / R-TCE over UDP. Since UDP does not guarantee reliable packet delivery, such network losses are detected by the R-InTracer / R-TCE at application level and reported by generating alarms based on configuration.

Health Monitoring

Server heartbeat process is part of Health Monitoring of InTracer components. This is not only responsible for starting/stopping of all the processes, but also has a watchdog function that ensures that:

- All the processes are up
- Every process is functioning normally

In an event where any of the process is not up or not functioning normally as expected, it can stop and restart the process to restore the system to processing state.

It also has a function to guard the system and notify the operator using logs and alarm against persistent errors which need operator intervention.

Resource Monitoring

Resource Monitoring has below two parts to monitor:

- CPU and Memory utilization of the system and other individual processes
- Disk utilization for Export, Data, Log, Crash Directories and for Disk itself

CPU and Memory utilization of the system and other individual processes: Helps monitor complete resource utilization of InTracer.

Below table describes system resources of InTracer which can be monitored.

Process	Description
CPU Usage (%)	CPU utilization in percentage for specific InTracer process.
Memory Usage (%)	System memory utilization in percentage for specific InTracer process.

Disk Utilization Monitoring: Disk utilization monitoring helps monitor Disk Utilization of various components of InTracer.

Disk Utilization monitoring provides for selected device name can be performed on the below components of InTracer:

- **Mount Point**: The name of the mount point of storage device.
- Filesystem ID: The identifier for Filesystem on storage device.

Maintenance Activities

The R-InTracer / R-TCE performs various routine activities for the system maintenance and for persistent storage. These activities include the following:

- Call event information and SA Key information is stored for a configurable period configured at install time
- Running DB verification to ensure data is in proper and valid form
- Backup of InTracer system state information, if enabled
- Compression of processing log data, removal of older logs

These activities are performed in background by UNIX cron jobs that are invoked periodically by the system. Not all these activities are done in a single run. These are shared across various invocations of the jobs at different times of the day when the processing load on the R-InTracer / R-TCE is assumed to be less. The R-InTracer / R-TCE facilitates operators to specify configurations for all these activities. These configurations can be performed through the InTracer clients.

Statistical Counters

The R-InTracer / R-TCE process maintains current statistical information for each of its functions. The InTracer clients provide an interface to view these statistics. These statistics can be used to:

- Perform analysis of call events
- Measure performance of the AGW and the R-InTracer / R-TCE and health of the link between them

C-InTracer Features and Functions

The following are the features and functions supported by the C-InTracer

Feature	Configurable	Default	Nodal Trace	3GPP Trace
Query Processing	Yes	-	Yes	Yes
User Management	Yes	Enabled	Yes	Yes
System Configuration	Yes	Admin only	Yes	Yes
Gateway Configuration	Yes	Admin only	No	Yes
Query Management	Yes	All users	Yes	Yes
Audit Trail	Yes	30 days	Yes	Yes
Health Monitoring	Yes	Enabled	Yes	Yes
Resource Monitoring	Yes	Enabled (default threshold values)	Yes	Yes
Maintenance Activities	Yes	Enabled (daily midnight)	Yes	Yes
Statistical Counters	No	Always	Yes	Yes

Query Processing

For C-InTracer(Nodal installation), Session related queries can be based on IMSI, NAI, Call-Id, Disconnect reasons and so on, limited by a time window.

For R-TCE (3gpp installation), Session related queries can be based on IMSI, IMEI/IMEISV, Trace Session reference, and so on, limited by a time window.

System Configuration

Use the system setup to manage and configure In Tracer components such as the R-InTracer / R-TCE Server, the C-InTracer and the Event Source.

Audit Trail

This keeps track of various system activities such as User session login/logout and configuration change activities. This audit trail is helpful for troubleshooting and monitoring the system activities.

Health Monitoring

Server Hearth beat process is part of Health Monitoring of InTracer components. This is not only responsible for starting/stopping of all the processes, but also has a watchdog function that ensures that:

- All the processes are up
- Every process is functioning normally

In an event where any of the process is not up or not functioning normally as expected, it can stop and restart the process to restore the system to processing state.

It also has a function to guard the system and notify the operator using logs and alarm against persistent errors which need operator intervention.

Resource Monitoring

Resource Monitoring has below two parts to monitor:

- CPU and Memory utilization of the system and other individual processes
- Disk utilization for Export, Data, Log, Crash Directories and for Disk itself

CPU and Memory utilization of the system and other individual processes: Helps monitor complete resource utilization of InTracer. Below table describes system resources of InTracer which can be monitored.

Process	Description
CPU Usage (%)	CPU utilization in percentage for specific InTracer process.
Memory Usage (%)	System memory utilization in percentage for specific InTracer process.

Disk Utilization Monitoring Helps monitor Disk Utilization of various components of InTracer. Disk Utilization monitoring Provides for selected device name can be performed on the below components of InTracer:

- **Mount Point**: The name of the mount point of storage device.
- Filesystem ID: The identifier for Filesystem on storage device.

Maintenance Activities

The C-InTracer performs various routine activities for the system maintenance and for persistent storage. These activities include the following:

- Call event information and SA Key information is stored for a max period of 30 days and information exceeding this storage period is automatically deleted by the server
- Running DB verification to ensure data is in proper and valid form
- Backup of InTracer system state information, if enabled
- Compression of processing log data, removal of older logs

These activities are performed in the background by UNIX cron jobs that are invoked periodically by the system. Not all of these activities are done in a single run. These are shared across various invocations of the jobs at different times of the day when the processing load on the R-InTracer / R-TCE is assumed to be less. The R-InTracer / R-TCE allows operators to specify configurations for all of these activities. These configurations can be performed through the InTracer clients.

Statistical Counters

The C-InTracer process maintains current statistical information for each of its functions. Cumulative statistics are collected on a daily basis and stored in the server statistical database. The InTracer clients provide an interface to view these statistics. These statistics can be used to:

- Perform analysis of call events
- Measure performance of the AGW and the R-InTracer / R-TCE and health of the link between them

InTracer Web User Interface

A unified Web-based user interface for performing configuration and querying in InTracer.

InTracer System Requirements

This section provides InTracer System Requirements.

InTracer Client Pre-requisites

The InTracer Nodal Trace solution is for Cisco ASR5xxx based gateway services. The InTracer Client in this case requires an ASR5xxx based gateway.

The InTracer 3GPP Trace solution is for Cisco 7600 based gateway services. The InTracer Client in this case requires a 7600 SAMI based gateway.

R-InTracer / R-TCE Pre-requisites

Following are the hardware pre-requisites for the R-InTracer / R-TCE:

Hardware Parameters	Server option 1	Server option 2
Server	UCS C460	UCS C460
CPU		
Memory	32GB	32GB
Internal Disk	2 * 300 GB SCSI	2 * 300 GB SCSI
HBA Card	2 * 4 Gbps FC ports	2 * 4 Gbps FC ports
Brocade Switch		
External Disk	Sun StorageTek 2540	Sun StorageTek 2540
Network Interfaces	2 physical interfaces >= 100 Mbps	2 physical interfaces >= 100 Mbps

Following are the software pre-requisites for R-InTracer / R-TCE:

Software Parameters	Description
Operating System	Custom RHEL 64-bit distribution (with all required packages and patches)
File System	XFS

C-InTracer Pre-requisites

Following are the hardware pre-requisites for the C-InTracer:

Hardware Parameters	Server option 1	Server option 2
Server	UCS C460	UCS C460
CPU		
Memory	32 GB	32 GB
Internal Disk	2 * 300 GB SCSI	2 * 300 GB SCSI
HBA Card	2 * 4 Gbps FC ports	2 * 4 Gbps FC ports

Hardware Parameters	Server option 1	Server option 2
Brocade switch		
External Disk	Sun StorageTek 2540	Sun StorageTek 2540
Network Interfaces	1 physical interfaces >= 100 Mbps	1 physical interfaces >= 100 Mbps

Important: The HBA card, Brocade switch and External disk are required for 3GPP Trace only.

Following are the software pre-requisites for C-InTracer:

Software Parameters	Description
Operating System	Custom RHEL 64-bit distribution (with all required packages and patches)
File System	Ext3,Ext4,XFS

InTracer User Work Station Pre-requisites

Important: There are no specific hardware pre-requisites for the InTracer User Work-station

Following are the software pre-requisites for User work station:

Software Parameters	OS Option1	OS Option2
OS	Windows 7, Windows XP	Mac OS
Browsers	Firefox 3.6Internet Explorer	Firefox 3.6

Supported Features

The InTracer solution currently supports two kinds of tracing functionalities.

Nodal Trace

A tracing solution for Cisco ASR5xxx based gateway services. In this tracing solution, once tracing is activated on the gateway, it sends a copy of all signaling packets (transmitted and received) for each subscriber session as part of normal processing to an external server.

- Protocol packets are sent from the lowest layer where subscriber or session information is available.
- Additional call correlation information is also sent along with the protocol packets.
- Correlation of packets is limited to a single gateway.
- Additional Cisco ASR5xxx specific system events like Card / CPU / Task crash etc are sent for hardware / software failures to associate call failures (if any) to these events.

3GPP Trace

A 3GPP standard based tracing solution. In this tracing solution, the gateway sends a copy of all signaling packets (transmitted and received) for only the trace sessions that are explicitly activated (either through management-based activation or signaling-based activation).

- All packets are GTPv2 packets.
- Additional correlation information like the Trace Session Reference is sent along with the protocol packets.
- Correlation of packets using Trace Session Reference is across all gateways on which the same trace is activated.
- Support available for PGW and SGW.

Chapter 20
IP Services Gateway Overview

This chapter provides an overview of the IP Services Gateway (IPSG) product.

This chapter covers the following topics:

- Introduction
- How it Works
- In-line Services
- Enhanced Feature Support

Introduction

The IP Services Gateway (IPSG) is a stand-alone device capable of providing managed services to IP flows. The IPSG is situated on the network side of legacy, non-service capable GGSNs, PDSNs, HAs, and other subscriber management devices. The IPSG can provide per-subscriber services such as Enhanced Charging Service, Application Detection and Control, and others.

The IPSG allows the carrier to roll out advanced services without requiring a replacement of the HA, PDSN, GGSN, or other access gateways and eliminates the need to add multiple servers to support additional services.

Platform Requirements

The IPSG runs on Cisco ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

License Requirements

The IPSG is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on licensing requirements.

For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

How it Works

The IPSG supports the following service modes:

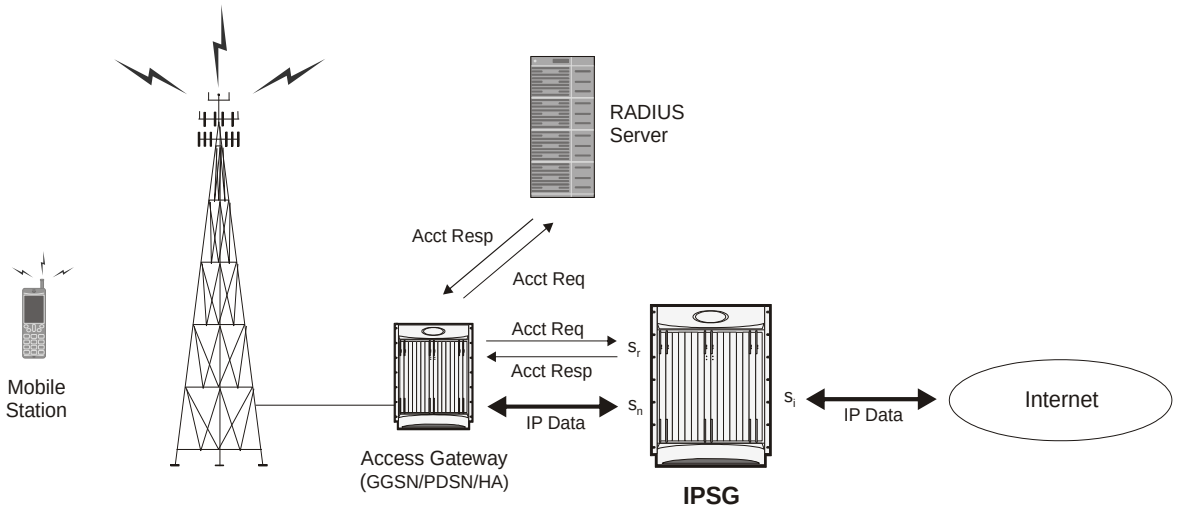
- RADIUS Server Mode
- RADIUS Snoop Mode

RADIUS Server Mode

When configured in RADIUS server mode, the IPSG inspects identical RADIUS accounting request packets sent to the RADIUS accounting server and the IPSG simultaneously.

As shown in the following figure, the IPSG inspects the RADIUS accounting request, extracts the required user information, then sends a RADIUS accounting response message back to the access gateway. The IPSG has three reference points: sn, si, and sr. The sn interface transmits/receives data packets to/from the access gateway (GGSN, HA, PDSN, etc.). The si interface transmits/receives data packets to/from the Internet or a packet data network. The sr interface receives RADIUS accounting requests from the access gateway. The system inspects the accounting request packets and extracts information to be used to determine the appropriate service(s) to apply to the flow.

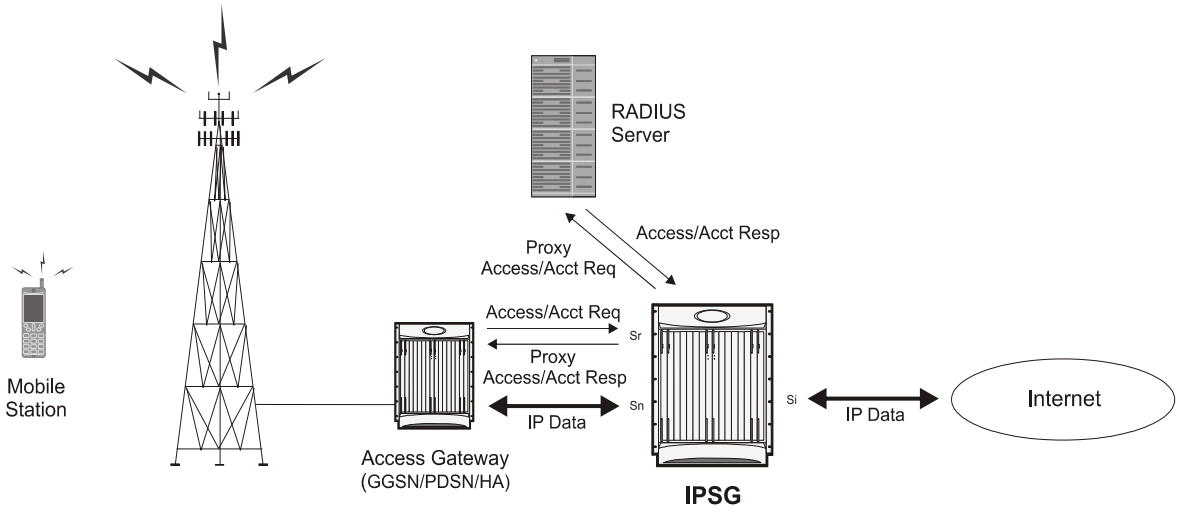
Figure 138. IPSG Message/Data Flow (RADIUS Server Mode)



RADIUS Proxy

In the event that the Access Gateway is incapable of sending two separate RADIUS Start message, the IPSG can be configured as a RADIUS Proxy. As shown in the following figure, the IPSG receives an IPSG RADIUS proxy Access request, then generates the Authentication and Accounting requests to the AAA Server.

Figure 139. IPSG Message/Data Flow (RADIUS Server Mode - RADIUS Proxy)

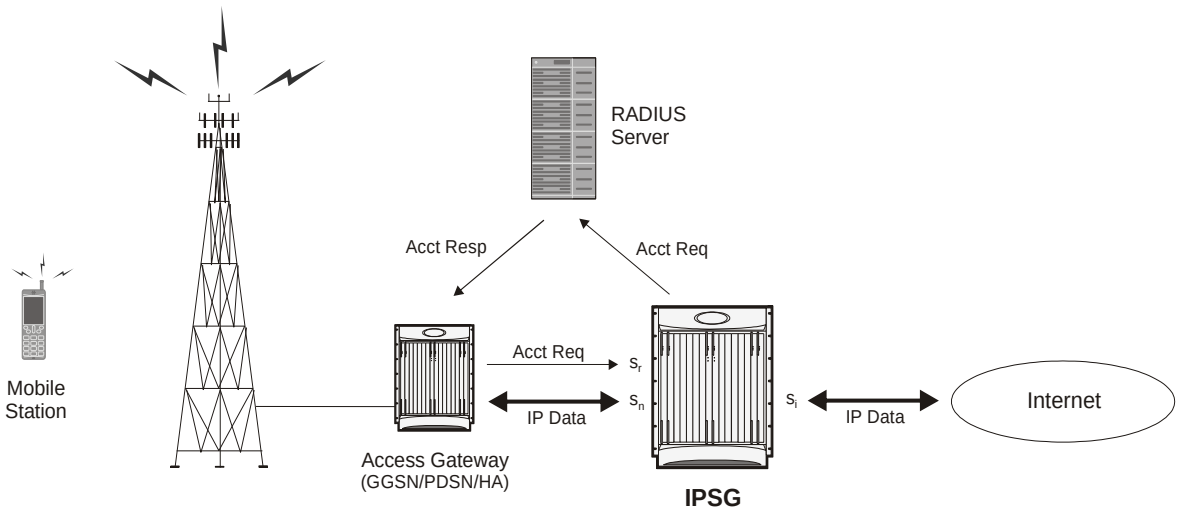


RADIUS Snoop Mode

When configured in RADIUS snoop mode, the IPSG simply inspects RADIUS accounting request packets sent to a RADIUS server through the IPSG.

As shown in the following figure, the IPSG has three reference points: sn, si, and sr. The sn interface transmits/receives data packets to/from the access gateway (GGSN, HA, PDSN, etc.). The si interface transmits/receives data packets to/from the Internet or a packet data network. The sr interface receives RADIUS accounting requests from the access gateway. The system inspects the accounting request packets and extracts information to be used to determine the appropriate service(s) to apply to the flow. Information is not extracted from the RADIUS accounting responses so they are sent directly to the access gateway by the RADIUS Server, but can also be sent back through the IPSG.

Figure 140. IPSG Message/Data Flow (RADIUS Snoop Mode)



In-line Services

As described previously, the IPSG provides a method of inspecting RADIUS packets to discover user identity for the purpose of applying enhanced services to the subsequent data flow. Internal applications such as the Enhanced Charging Service, Content Filtering, and Application Detection and Control are primary features that take advantage of the IPSG service.

Application Detection and Control

Application Detection and Control (ADC) is an in-line service feature that detects peer-to-peer protocols in real time and applies actions such as permitting, blocking, charging, bandwidth control, and TOS marking.

For more information, refer to the *Application Detection and Control Administration Guide*.

Content Filtering

Content Filtering is an in-line service feature that filters HTTP and WAP requests from mobile subscribers based on the URLs in the requests. This enables operators to filter and control the content that an individual subscriber can access, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences.

For more information, refer to the *Content Filtering Services Administration Guide*.

Enhanced Charging Service

Enhanced Charging Service (ECS)/Active Charging Service (ACS) is the primary vehicle performing packet inspection and applying rules to the session which includes the delivery of enhanced services.

For more information, refer to the *Enhanced Charging Service Administration Guide*.

Enhanced Feature Support

This section describes the enhanced features supported by IPSG.

Dynamic RADIUS Extensions (Change of Authorization)

Dynamic RADIUS extension support provide operators with greater control over subscriber PDP contexts by providing the ability to dynamically redirect data traffic, and or disconnect the PDP context.

This functionality is based on the RFC 3576, Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS), July 2003 standard.

The system supports the configuration and use of the following dynamic RADIUS extensions:

- Change of Authorization: The system supports CoA messages from the AAA server to change data filters associated with a subscriber session. The CoA request message from the AAA server must contain attributes to identify NAS and the subscriber session and a data filter ID for the data filter to apply to the subscriber session.
- Disconnect Message: The DM message is used to disconnect subscriber sessions in the system from a RADIUS server. The DM request message should contain necessary attributes to identify the subscriber session.

The above extensions can be used to dynamically re-direct subscriber PDP contexts to an alternate address for performing functions such as provisioning and/or account set up. This functionality is referred to as Session Redirection, or Hotlining.

Session redirection provides a means to redirect subscriber traffic to an external server by applying ACL rules to the traffic of an existing or a new subscriber session. The destination address and optionally the destination port of TCP/IP or UDP/IP packets from the subscriber are rewritten so the packet is forwarded to the designated redirected address.

Return traffic to the subscriber has the source address and port rewritten to the original values. The redirect ACL may be applied dynamically by means of the RADIUS Change of Authorization (CoA) extension.

Important: For more information on dynamic RADIUS extensions support, refer the *CoA, RADIUS, and Session Redirection (Hotlining)* appendix in the *IP Services Gateway Administration Guide*.

Gx Interface Support

To support roaming IMS subscribers in a GPRS/UMTS network, the IPSG must be able to charge only for the amount of resources consumed by the particular IMS application and bandwidth used. The IPSG must also allow for the provisioning and control of the resources used by the IMS subscriber. To facilitate this, the IPSG supports the R7 Gx interface to a Policy Control and Charging Rule Function (PCRF).

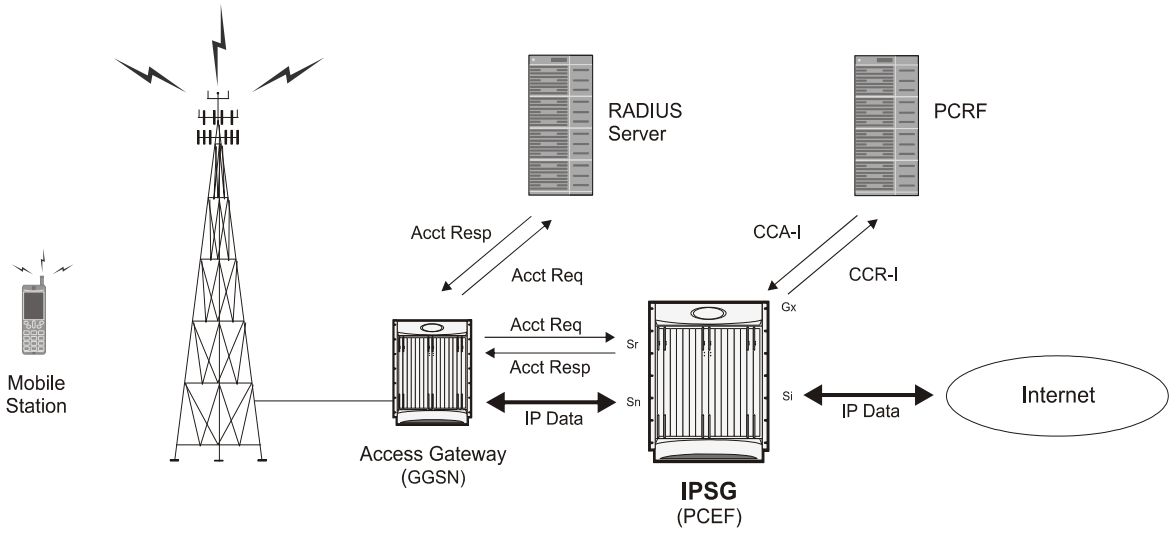
For detailed information on Gx Interface support, refer to the *Gx Interface Support* appendix in the *IP Services Gateway Administration Guide*.

Note the following for IPSG:

- Only single bearer/session concept is supported. Multiple bearer concept is not applicable.
- Only PCRF binding is applicable. PCEF binding is not applicable.

The following figure shows the interface and basic message flow of the Gx interface.

Figure 141. IPSG Message/Data Flow (RADIUS Server Mode - IMS Auth Service)



IPSG also supports IMS Authorization Service Session Recovery with the following limitations:

- Active calls only
- The number of rules recovered is limited to the following:
  - 3 flow-descriptions per charging-rule-definition
  - 3 Charging-rule-definitions per PDP context
- The above are combined limits for opened/closed gates and for uplink and downlink rules. IMSA sessions with rules more than the above are not recoverable.

Gy Interface Support

This is a Diameter protocol-based interface over which the IPSG communicates with a Charging Trigger Function (CTF) server that provides online charging data. Gy interface support provides an online charging interface that works with the ECS deep packet inspection feature. With Gy, customer traffic can be gated and billed in an "online" or "prepaid" style. Both time- and volume-based charging models are supported. In all of these models, differentiated rates can be applied to different services based on shallow or deep packet inspection.

For more information on Gy interface support, refer to the *Gy Interface Support* appendix in the *IP Services Gateway Administration Guide*.

Content Service Steering

Content Service Steering (CSS), defines how traffic is handled by the system based on the content of the data presented by a mobile subscriber. CSS can be used to direct traffic to in-line services that are internal to the system. CSS controls how subscriber data is forwarded to a particular in-line service, but does not control the content.

IPSG supports steering subscriber sessions to Content Filtering Service based on their policy setting. If a subscriber does not have a policy setting (ACL name) requiring Content Filtering, their session will bypass the Content Filtering Service and will be routed on to the destination address.

If subscriber policy entitlements indicate filtering is required for a subscriber, CSS will be used to steer subscriber sessions to the Content Filtering in-line service.

If a subscriber is using a mobile application with protocol type not supported, their session will bypass the Content Filtering Service and will be efficiently routed on to destination address.

For more information regarding CSS, refer to the *Content Service Steering* chapter in the *System Administration Guide*.

Multiple IPSG Services

Multiple IPSG services, can be configured on the system in different contexts. Both source and destination contexts should be different for the different IPSG services. Each such IPSG service functions independently as an IPSG.

Session Recovery

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

Session recovery is performed by mirroring key software processes (for example, Session Manager and AAA Manager) within the system. These mirrored processes remain in an idle state (in standby-mode), wherein they perform no processing, until they may be needed in the case of a software failure (for example, a Session Manager task aborts). The system spawns new instances of "standby mode" session and AAA Managers for each active Control Processor (CP) being used.

Additionally, other key system-level software tasks, such as VPN Manager, are performed on a physically separate packet processing card to ensure that a double software fault (for example, Session Manager and VPN Manager fails at same time on same card) cannot occur. The packet processing card used to host the VPN Manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

For more information on Session Recovery, refer to the *Session Recovery* chapter in the *System Administration Guide*.

Note that the Inter-Chassis Session Recovery feature is not supported in this release.

Chapter 21
Mobility Unified Reporting System Overview

This chapter provides an overview of the Mobility Unified Reporting (MUR) application.

This chapter describes the following topics:

- Introduction
- MUR Architecture
- Distributed Architecture of MUR
- Scalable MUR
- MUR Deployment
- MUR System Requirements
- MUR Ports

Introduction

The Cisco Mobility Unified Reporting (MUR) system is a Web-based application providing a unified reporting interface for diverse data from Cisco Systems In-line service and storage applications.

The MUR application enables:

- Generating customized reports and comparison charts.

  This release of MUR supports generating HTML-based historical canned reports displaying data in graphical—graphs/charts—and tabular formats. Reports for ad-hoc periods are not supported. For information on the various reports supported, see the Report Types section.

- Analyzing the reporting data and enabling the operator to get a full understanding of the performance of the network, enabling operators to optimally configure and plan their network.
- Supporting distributed installation which allows to view reports from multiple sites.
- Rich visualization (Graphs/tabular form).
- Exporting reports in Microsoft Excel, Adobe PDF, and CSV formats.
- Report scheduling, notification, and distribution. The report notification can be in the form of alarms/traps.
- Capacity monitoring and planning of system supporting a suite of products such as PDSN, GGSN, SGSN, and inline service applications like Content Filtering, Stateful Firewall, Application Detection and Control (ADC).

The MUR application is available for report generation only when you install the software application on to your local server. For information on the server recommendations, refer to MUR System Requirements section in this guide. For information on how to install the MUR application, refer to the *Managing Mobility Unified Reporting System Installation* chapter in this guide.

The MUR application provides comprehensive and consistent set of statistics and customized reports, report scheduling and distribution from ASR chassis / in-line service product. For example, a subscriber's Quality of Experience, top 10 sites visited, top 10 users, and so on.

The MUR application provides reporting capability for Content Filtering (CF) data, bulk statistics, Key Performance Indicators (KPIs), EDR data from in-line service and storage applications. The MUR application facilitates and enhances the operators' ability to simply and easily determine the health and usage of the network.

The chassis directly pushes the bulkstat files and EDR data to the reporting server through SFTP. MUR receives the input data from the chassis only when the Enhanced Charging Services (ECS) module is enabled and configured to generate reporting EDRs. To enable this, you must purchase and install ECSv2 license on the chassis.

Important: In RHEL-based deployment of MUR, L-ESS is NOT required as the ASR chassis Enhanced Charging Services (ECS) module can be configured to push the xDRs directly to the MUR reporting server. Push from ASR chassis is the Cisco recommended deployment model. Currently L-ESS is supported only on Solaris platforms. For information on the L-ESS installation instructions, refer to the *ESS Installation and Administration Guide*. Existing deployments where L-ESS is installed, to pull EDRs from the chassis, may continue with their deployment model in the 12.0 version of MUR Software Release and later.

For information on obtaining and installing the license, see *System Administration Guide* and *Enhanced Charging Services Administration Guide*. For information on configuring the ECS module, see *Configuring Chassis for Mobility Unified Reporting System* chapter in this guide.

MUR receives the following types of EDRs for report processing:

- CF-EDRs
- Flow EDRs
- HTTP EDRs

To reduce disk space and improve performance, MUR limits the bucket distribution for EDR data to ONLY last 2 days in case a EDR is spanning across more than 2 days or so.

For example, if the following EDR is received:

```
#sn-start-time,sn-end-time,radius-calling-station-id,ip-subscriber-ip-address,sn-
subscriber-port,ip-server-ip-address,sn-server-port,sn-app-protocol,p2p-protocol,traffic-
type,voip-duration,sn-volume-amt-ip-bytes-uplink,sn-volume-amt-ip-bytes-downlink,sn-
volume-amt-ip-pkts-uplink,sn-volume-amt-ip-pkts-downlink,bearer-3gpp rat-type,radius-
called-station-id,bearer-3gpp imei,ip-protocol,bearer-3gpp sgsn-address,sn-flow-start-
time,sn-flow-end-time
1275330600,1275334200,9689944191,19.19.1.1,35111,1.1.1.1,21,8,,,0,52428800,1048576,100,20
0,1,apn.org1,35302703-090362-52,6,1.1.1.3,1275330600,1275334200
```

MUR determines the difference between the starttime and endtime attributes and limits the bucket distribution as shown here.

```
#starttime,endtime,protocol,rxbytes,txbytes 2011/02/26 10:00:00,2011/02/28
10:00:00,HTTP,100MB,100MB
```

Important: The bucket distribution calculation will remain intact i.e. the volume will be distributed equally among all the half-hour's buckets that fall in the starttime and endtime.

Important: The MUR receives the data in terms of EDRs which are generated based on the flow. As the EDRs are flow-based and the bulkstats is a real-time data, the volumes reported in the EDR are different from the volumes reported by bulkstats.

For more information on using the MUR application to generate reports, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Report Types

The MUR application supports generation of canned statistical reports that can be used to analyze network performance, and decide the policies for users, and identify the customer trends, network usage patterns, network categorization, etc. The reports can be per gateway, or multiple gateways (region), or for the overall network. The reports can be generated for the usage of different entities such as gateway, content type, etc on an hourly, daily, weekly, or monthly basis.

The typical canned reports that are supported for the MUR application include:

- Historical summary reports (Daily/Weekly/Monthly)
    - Half-hourly Reports: Usage reporting for the specified time period
    - Daily Reports: Usage reporting for the past 24-hour period (midnight through midnight)
    - Weekly Reports: Usage reporting for the past seven day period (Monday through Sunday)
    - Monthly Reports: Usage reporting for the past 30-day period (1 day of the month through the last day of the month)

- Top "N" Reports
- Statistical and analytical reports
- Bulkstats and KPI reports

The static report layout comprises the following sections:

- The report name
- The report ownership: the user account that requested the report
- The date and time of generation
- The list of report parameters
- The chart legend (displayed under the chart)

On the interactive layout the user can set a series of preferences in a specific manner. The user has the flexibility to change the type of chart from Bar to Pie (supported output types depend on the selected report). Changing the preferences like the chart type or report parameters will cause the report to refresh in the same window.

The interactive chart layout provides the following list of features:

- Tool tip: When the mouse pointer stops over a chart series, after a short time a tool tip is displayed showing the information of the targeted sample.
- Dynamic legend: The legend is located beneath the chart and is used to recognize the series plotted on the screen. In case of series representing either network services or subscriber packages, the colors are bound to the service/package names. This means that, for example, the HTTP Service will be rendered with a specific color for the reports. The legend is usually displayed with check-boxes associated to each color.

The MUR application provides the following reports:

- Traffic Analysis Report: The Traffic Analysis report provides the total usage traffic (including uplink and downlink traffic) details for the following application categories:
  - Filesharing
  - Web
  - IM
  - VOIP
  - Standard
  - Streaming
  - Tunnel
  - Gaming
  - Unclassified

  The usage traffic is expressed in terms of megabytes (MB) or Megabits per second (Mbps) and percentage (%). The traffic can also be in gigabytes (GB) / kilobytes (KB) / bytes depending on the magnitude.
- Traffic Distribution Report: The Traffic Distribution report provides the summary of total traffic distribution for all the protocols application categories over a specified time period. The usage traffic is represented in GB/MB/KB/Bytes and percentage.
- Active Flow Count Report: The Active Flow Count report provides the details of traffic distribution flow count against the different application categories. This report also provides the summary of maximum number of flows in the EDR records.

> ***Important:*** Active Flow Count report for current date will not be available because daily tables used to fetch this report are generated only at the end of the day. Also when the user selects a date range, for example, 10/1/2011 to 10/5/2011 where 10/5/2011 is the current date, then the report will be shown for the period 10/1/2011 to 10/4/2011 i.e. up till 10/4/2011.

- Unique Subscriber Hits Report: The Unique Subscriber Hits report provides an overview of the usage patterns of the entire subscriber population per protocol, for example, how many people are actually using VoIP.

  > ***Important:*** Unique Subscriber Hits report can be generated ONLY for a single date/week/month and not for any date-range. Also, note that the time selection is also disabled for this report.

  Typically, this report provides the total number of times a subscriber is using a specific protocol. These reports are displayed for all configured gateways.

  > ***Important:*** Unique Subscriber Hits report for current date will always be available on the subsequent date because unique subscribers hits calculation will be performed at the end of the day.

- TopN versus Total Traffic Report: This report provides the summary of total usage traffic and Top N subscriber traffic for all the protocols over a specified time period. The usage traffic is represented in GB/MB/KB/Bytes and packets.
- TopN Subscribers Report: The TopN Subscribers report simply counts the number of bytes per subscriber for different time intervals. It displays the top 10/100/1000 subscribers for each day/week/month. This report is displayed across all configured gateways, per region or per NOC.

  > ***Important:*** This report is not available for a multiple date range selection.

  After identifying the total amount of transferred data per subscriber, and identifying the top users, to understand the protocol and services breakdown for each subscriber, this report allows listing the different applications used by the top 10/100/1000 subscribers based on the selection of top subscriber per day/week/month.
- TopN VCD Subscribers Report: The TopN Voice Call Duration (VCD) Subscribers report displays the top N subscribers based on their voice usage (voice duration) for Yahoo, MSN and Skype voice protocols. The summary report displays the voice summary (voice duration) for VoIP category.

  > ***Important:*** This report is also available per week or month.

- Weekly Report: The weekly report provides details of the following:
  - Total traffic
  - Total traffic by category
  - VOIP Call Duration
  - Total unclassified traffic (TCP and UDP)
  - Top N subscribers

- Monthly Report: The monthly report provides the details of total traffic across the top N protocols / application categories in a month.
- Custom Reporting: MUR supports on-demand offline reporting of subscriber specific information to operators. This ad-hoc request could be a subscriber search request or top N search request.

  Offline Subscriber Report: The MUR aids in searching individual subscribers' data based on certain parameters like IMSI, MSISDN, NAI, IMEI and Public and Private (NAT) subscriber IP address with ports, individually or in combination, and generates a subscriber-specific report showing the list of URLs visited by the subscriber, and other details like QoS, usage traffic, aggregate application/protocol breakdown, etc for the specified time period. MUR mainly supports this search functionality to track a subscriber or a set of subscribers for lawful intercept.

  To use this Offline Reporting feature seamlessly, you must configure the EDR Filename Format appropriately through the Gateway configuration from **ADMIN** tab, and organize the archive directory date-wise. For information on how to manage the archive directory, see the *Managing Archive Directory* section in the *MUR Administration and Management* chapter of this guide.

  Offline Top N Subscribers Report: MUR also facilitates to generate an offline report that covers the % of volume/duration used by top n% subscribers. This report provides information on the absolute number of subscribers and the list of MSISDNs to facilitate correlation with the provisioning data. In this release, this ad-hoc report is available per APN group, Device Group, Location Group, and Service Profile.

  For more information on these features, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

  Reports based on Tethering Configuration: Tethering refers to the use of a mobile smartphone as a USB dongle/modem to provide Internet connectivity to PC devices (laptops, PDAs, tablets, and so on) running on the smartphone's data plan. Typically, for smartphone users, most operators have in place an unlimited data plan, the usage of which is intended to be from the smartphone as a mobile device. However, some subscribers use the low cost / unlimited usage data plan to provide Internet connectivity to their laptops in places where normal Internet connection via broadband/WiFi may be costly, unavailable, or insecure.

  The ASR chassis works in conjunction with the MUR application to facilitate tethering detection on the chassis. The EDRs generated by the chassis will be enhanced to include OS signatures.

  MUR processes flow-EDR files containing OS signature and IMEI field, HTTP files containing OS signature, User Agent, and IMEI field, and populates the tethering data in database files.

  For more information on this feature, see the *Cisco Mobility Unified Reporting System Online Help* documentation and *Enhanced Charging Services Administration Guide*.
- DPI Report: The Deep Packet Inspection (DPI) reports are the canned statistical reports at the gateway level and region level. You can configure the MUR application to generate the reports for any of the available gateways.

  In this release, MUR supports generating daily, weekly and monthly summary details and busy hour traffic usage details for the following report categories:

  - Traffic Analysis Report
  - Traffic Distribution Report
  - Active Flow Count Report
  - Unique Subscribers Hits Report
  - TopN Reports — Report on Top N vs Total Traffic, TopN subscribers, TopN VCD subscribers

  MUR has the capability to report the following details per protocol:

  - Total volume for the day/week/month
  - Volume distribution in the busy hour

- Peak performance for the day/week/month
- Maximum number of unique subscribers

MUR supports additional information breakdown by network characteristics. These include Application Category, Protocol Groups, IP Protocol, Device Group, RAT (Radio Access Type i.e 2G vs 3G), APN (Access Point Name), SGSN group, Service Profile, Roaming Partner, and Location Group. During its development, a device may have several TAC codes and there may be a need to report devices by broader device type such as "Blackberry" or "Smartphone". Device groups allow the operator to combine a range of TACs into a single named group for reporting purposes.

Busy Hour Reporting: Busy Hour (BH) reporting is mainly useful for the users to monitor different traffic flows in their network during the busy hour. BH indicates the sliding 60-minute period during which occurs the maximum total traffic load in a given 24-hour period.

Please note the following key points:

- BH reporting is available ONLY on the GUI and not in xls format.
- BH reporting is available only under the **DPI** tab.
- BH radio button is available on the date panel.
- BH reporting is available for a date, date range, week and month.
- Busy hour reports are currently available ONLY at the NOC level.

DSL Reports: The current release of MUR provides the following details for Digital Subscriber Line (DSL) traffic reports:

- Traffic analysis — uplink DSL, downlink DSL and total DSL traffic including daily weekly, and monthly aggregation/distribution.
- DSL traffic categorization — total P2P traffic over DSL, IP traffic, web traffic, etc.
- Top N% DSL subscribers
- Comparison of total DSL traffic versus total UMTS traffic

For information on additional reports supported through DPI, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

- CF-RE Report: Content Filtering (CF) solution enables operators to filter HTTP and WAP requests from mobile subscribers based on the URLs in the requests, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences.

  The CF-RE report provides the summary of traffic over CF categories, CF actions, and CF ratings. The CF actions that can be taken on the URL are as follows:

  - allow
  - discard
  - redirect-url
  - content-insert
  - terminate-flow
  - reply-code-terminate-flow

  The CF ratings can be one of the following:

  - dynamic
  - static
  - blacklisted

The CF-RE report also provides the list of top N subscribers and URLs based on their unique subscriber's hit count and total usage.

- HTTP Reports: The MUR application parses HTTP EDRs and then provides the following details for any specific day, week, month and date range:
  - Total traffic per HTTP group / host name and HTTP content type
  - URL hits per HTTP group / host name and HTTP content type
  - Unique subscriber count per HTTP group / host name

  Typically, MUR supports the following categories of HTTP reports:
  - Summary reports — Content type/subtype volume report available for daily, weekly, monthly, and date range
  - Top N reports (Daily/Weekly/Monthly)
    - HTTP Group Aggregation — TopN HTTP group by Volume; TopN HTTP group by Hit count; TopN HTTP group by Unique subscriber hits
    - Top N Referrer Group Aggregation by Hit count
    - TopN User Agent (UA) reports available for APN-TAC combination in addition to individual per APN, per TAC reports.

  The top N referrers' report provides details of the total hit count for top N referrers and their sub-domain wise traffic distribution.

  Important: In the distributed model of MUR, the data received from RDP is populated and TopN referrer report is available only at NOC level.

  Important: It is mandatory to configure *http-url* and *http-referer* fields in the EDR records for top N HTTP referrers report generation.

- Bulkstat Report: The Bulkstat report provides details of the processed bulk statistics from any application (PDSN, GGSN, SGSN, and so only) on the managed nodes in a timely manner.

  Important: Make sure that you configure the bulkstats schemas through the GUI to generate bulkstats reports for any of the available gateways. For more information on schema configuration, refer to the *Configuring Bulkstats Schemas Using GUI* section in this guide and also *Cisco Mobility Unified Reporting System Online Help* documentation.

  The bulkstat data is sent from the gateway to the MUR server with GMT (UTC Time stamps). The bulkstat file processing is triggered by the MUR scheduler engine. The scheduler processes the bulkstat files line by line for each gateway, and gets the schema, timestamp, and key index. If the index does not exist, the parser creates index and inserts data into bulkstats data table. Once the processing is complete, this data file is moved to the archive directory. Summarization must happen as the user moves from gateway to higher levels.

  Important: For Bulkstat, there is no support for distributed model and all the bulkstat input files will be parsed by master MUR system only.

  MUR supports generation of busy hour reports, top N Min/Max reports, performance aggregation reports i.e. daily, weekly and monthly summary reports.

Please keep the following key points in mind for bulkstats reporting:

- The gateway(s) and MUR server need to be NTP synced for accurate BS aggregation reports.
- Hourly aggregation reports are triggered at 50th minute of every hour.
- Daily reports are scheduled at 3:45 PM the next day.
- Weekly reports are scheduled at 5:00 PM every Monday.
- Monthly reports are scheduled at 06:15 PM on 1st of every month.

- KPI Report: The KPI report provides details of the KPIs for each selected schema. KPIs are the formula-based calculations of selected bulk statistics counters. You can configure the MUR application to generate the reports for any of the available gateways. For a complete listing of supported KPIs and its associated formulas/descriptions, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Important: Please note that the Bulkstats and KPI reports are displayed based on the gateway's time zone.

Important: Please note that the subscriber's private data like Mobile Station Integrated Services Digital Network (MSISDN) will appear encrypted in all the subscribers reporting. Users with administrative privilege can only decrypt the MSISDNs using a shell script utility. For information on how to use this script, refer to the *MUR Administration and Management* chapter in this guide.

Important: Please note that the availability of any report is typically based on the date/date range configurations and purging interval. If you are trying to view a report beyond the configured purging interval, MUR system will display an error message indicating that the report is unavailable.

For more information on each of these reports, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Exporting Reports to Other File Formats

The MUR application supports exporting reports to the following file formats:

- Microsoft Excel format: To export a report to Microsoft Excel format, use the *get_excel_report* script in the CLI. For more information about this script, refer to the *Generating Reports in Excel Format* section in the *MUR Administration and Management* chapter of this guide.

  Exporting of reports to Excel format is also possible through the GUI by clicking the excel icon present in the tabular view of each of the reports under **HOME** and **DPI** tabs.

- Comma Separated Value (CSV) file format: To view reports in CSV format, in the HOME and DPI tabs, click the csv icon present in the tabular view of each of the reports.

- PDF format: To export a report to PDF format, in the **HOME** and **DPI** tabs of the MUR GUI, click the **PDF** button. The PDF file is displayed in a new window and can be saved for future reference.

  If there is no data available for a report, the **PDF** button is disabled.

- Text File format: This format is applicable only to HTTP User Agent (UA) reports. To export this report in a text file, click **Export to Text** button available in the HTTP UA reporting page.

  For more information, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

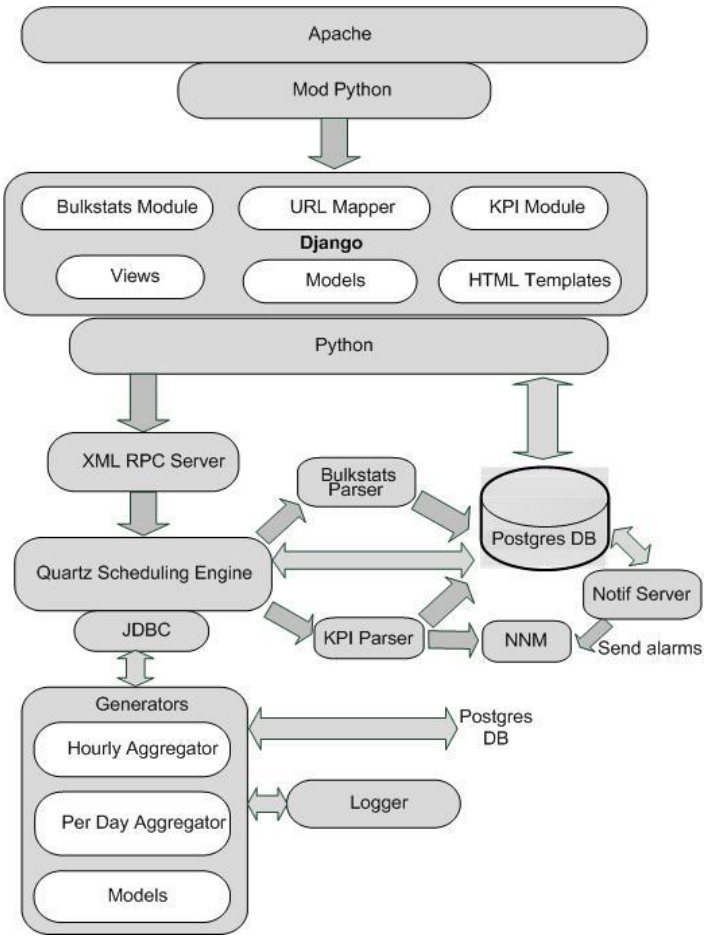
License Requirements

The MUR system is a licensed Cisco product. Contact your Cisco account representative for detailed information on specific licensing requirements.

MUR Architecture

The MUR solution consists of two components — a server and a GUI client. The following figure shows a typical organization of the MUR solution.

Figure 142. Internal Architecture of MUR



The server components include:

- DB Server: This is the standard PostGreSQL 8.3 database server. This is started at the time of application startup.

  MUR uses pgbouncer utility for postgres connection pooling. This utility gets started/stopped with Postgres Server.

- Quartz Scheduling Engine: This is the core of the MUR reporting solution. It is used to schedule different tasks such as parsing of incoming data files (bulkstat, EDR, etc.), trigger various canned reports on a periodic basis, cleaning up of stored outdated data and files, and so on.

- Generators: These are python based scripts that are used for parsing various CSV files. The files are parsed to an extent where generated files (or data in database) themselves represent meaningful data. This is a very powerful concept introduced for faster processing of information.

The generators archive the files once they are parsed. In archival, the files are zipped and placed in the configured location.

- KPI Parser: The KPI Alarm Generator uses the information stored by bulkstat parser in the database for KPI calculations and then, based on the calculations, generates the alarms that are subsequently sent to Network Node Manager (NNM).
- NotifServer: This stands as a separate entity that collects information from the MUR system and generates alarms which are then sent to the NNM for further analysis.
- Loggers: The MUR application uses various loggers so that application logs with various severities are made available for debugging purpose.

Some of the components at the client side include Django and Mod_python.

Distributed Architecture of MUR

MUR supports the distributed model to allow the deployment which enables network wide view or work load balancing. Newly introduced component, Remote Data Processor (RDP), plays the role of pre-processing the input files from gateways. One or more RDPs, installed separately on remote machines can be registered to a master MUR and one RDP can process files from one or more gateways.

RDP periodically sends the intermediate data to registered master MUR. The role of MUR in such deployments is mostly for report generation, report viewing, RDP management and optionally data processing.

Important: RDP installation and registration is required only for network wide deployments. For standalone installation no RDP is required. For information on how to install the RDP, refer to the *Managing MUR Installation* chapter of this guide.

Important: Make sure that you first install the master MUR system and then proceed with the RDP installation. Also, note that the RDP and MUR must be installed, upgraded, and uninstalled separately.

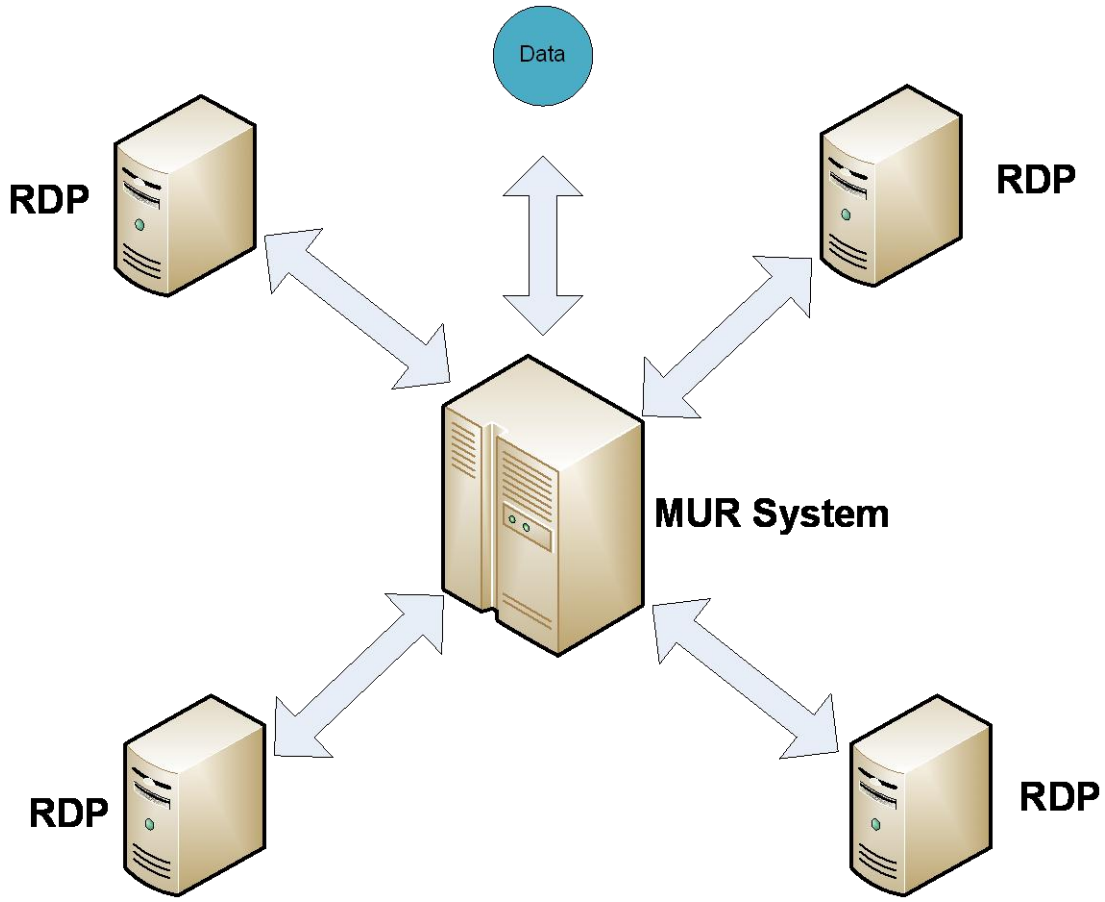
Important: Before registering RDP with the master MUR, ensure that the RDP is installed and running.

Important: The RDP management like configuration and removal is possible from MUR GUI only. For information on managing the RDPs, refer to the *Cisco Mobility Unified Reporting System Online Help* documentation.

Important: For Bulkstat, there is no support for distributed model and all the bulkstat input files will be parsed by master MUR only.

The following figure illustrates the distributed architecture of MUR.

Figure 143. Distributed Architecture of MUR



How RDP works with MUR

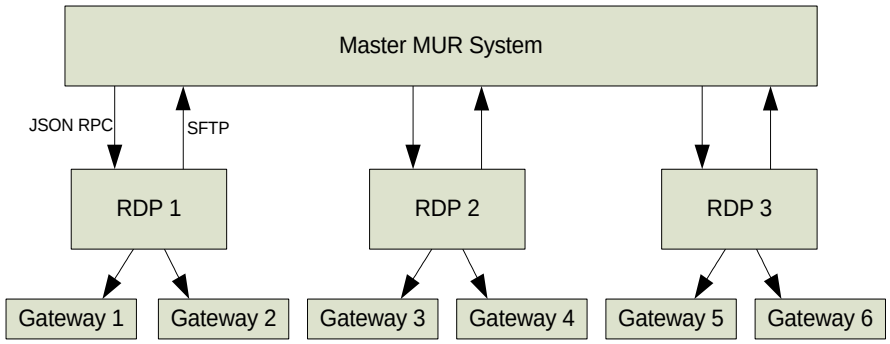
This section describes how the RDP works with the MUR application.

The RDP parses the raw data or EDR files from one or more GGSNs and populates the database for required reports. The RDP pre-processes the data and then periodically forwards them to the master MUR through SFTP for report generation.

Important: If the distributed model of MUR is used, then the SFTP user name and password should be the same as the MUR Administrator user's login name and password provided during installation. For information on configuring SFTP details, see the *Cisco Mobility Unified Reporting System Online Help* documentation.

Each of the RDP and MUR will be assigned a unique ID during installation and will be used for identification of each RDP along with its gateway and data.

Figure 144. MUR with RDPs in Distributed Model



In 12.0 and earlier releases, each of the registered RDPs will form a new region. RDP region can be a child of the root of the MUR (NOC) or can be the child of another region. The gateways associated with a RDP will always be the children of RDP region.

Release 12.2 onwards, users can create individual regions and add RDPs to the regions. All the gateways must be associated with RDP(s) or NOC and not to a region directly.

Important: Only single MUR can communicate with an RDP simultaneously.

Scalable MUR

In case of achieving high throughput (more than 10 GBPS), one RDP is not sufficient to process data. Scalability feature enables load distribution, so that data processing can be performed at multiple RDPs to support high data rate.

Scalability is achieved with the help of clustering of RDP nodes which share a common storage. The gateway pushes the input files to the shared storage through one of the RDPs, and all the RDPs then pick up those files as per certain rules defined for parsing. The parsing of files is based on the filename pattern that is specified while configuring the gateway. The MUR parsers match regular expressions given in the file name pattern and accordingly pick up the files. For example, users can specify *flow_*1 to match all flow EDR files ending with one to be picked up for that gateway. This way, multiple nodes parse the files in parallel. In such a cluster, maximum 16 nodes can be operating in parallel.

RDPs pick up the files based on the gateways that are attached to them. In the scalability model, one main instance of the gateway is configured on one of the RDPs. For the rest of the RDPs, a 'pseudo' instance of the same gateway is configured. This means, while adding another instance of the same gateway on another RDP, this gateway is marked as pseudo (through the MUR GUI) of an existing gateway (gateway configuration on GUI asks for name of a gateway to which gateway being configured is a pseudo). This configuration informs MUR that while processing data for pseudo gateway, this will not be considered as a separate gateway, but as the original gateway only.

In order to enable the RDPs to pick up files distinctly from each other, the configuration 'File Name Pattern' on the gateway configuration GUI is used. User can specify certain patterns for the file names for gateway (e.g. edr_*[1,3,5,7,9] to pick up only odd files and edr_*[0,2,4,6,8] to pick up only even numbered files) during such configuration. So, in the case of two RDPs, the user can specify pattern matching odd numbered files for first gateway that is configured on first RDP, and pattern matching even numbered files for second (pseudo) gateway that is configured on second RDP.

While viewing reports, in order to make sense to users, the pseudo gateway is NOT shown in the gateway hierarchical tree. The gateway tree is typically available under DPI, HTTP, CF, and Bulkstats tabs. However, the pseudo gateway information will be available through the **Admin** tab.

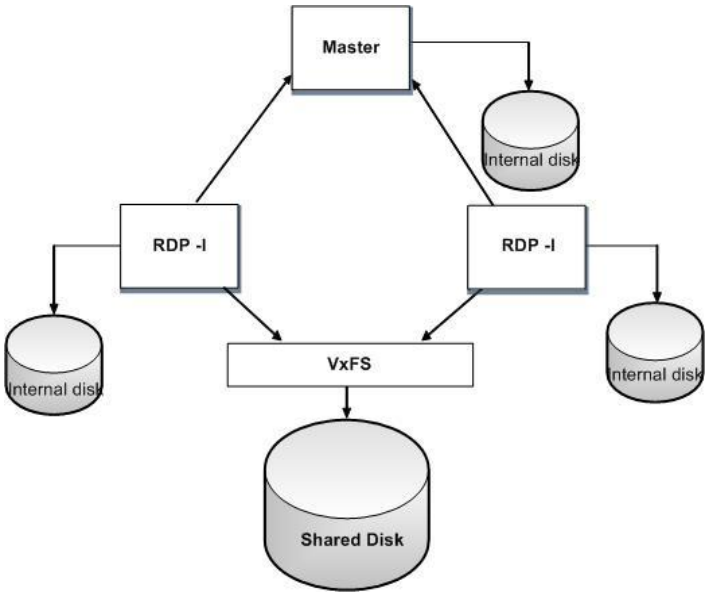
Basic Scalability Model

This section highlights the basic system requirements for the scalability in master MUR and RDPs.

- Master — Internal disk: OS, MUR (RAID-0), Postgres (RAID-5)
- RDP(s) — Internal disk: OS, MUR (RAID-0), Postgres (RAID-5), Archive (archive can be shifted to shared disk based on the requirement and sizing)
- Shared-disk — Only incoming files (RAID-0)

The following figure outlines the basic model of the scalability implementation in MUR.

Figure 145. Scalable MUR



All cluster nodes needs to be interconnected using switch and needs to be in same VLAN/ shared disk. After the cluster installation, the shared disk needs to be configured in all nodes. Check if the nodes are connected using **`multipath -ll`** command.

The following procedure describes the steps to be followed for creating and mounting the shared disk:

1. Check if the disk is available using **`multipath -ll`** command.
2. Use the following VCS command to see the status of the disk:

```
vxdisk -o alldgs list
```

3. To add shared disk group, use the command:

```
vxdg -s init mur_dgA disk=emc_clariion0_22
```

4. Create volume using the following command:

```
vxassist -g mur_dgA make sharedVol1 4000g
```

5. Create file system using the following command:

```
mkfs -t vxfs -o bsize=4096,largefiles /dev/vx/rdsk/mur_dgA/sharedVol1
```

6. Create the directory for shared disk using the following command:

```
mkdir /shared_edr_files
```

7. Create mount point using the following command:

```
mount -o cluster -t vxfs /dev/vx/dsk/mur_dgA/sharedVol1 /shared_edr_files
```

Important: After creating volume and sharing disk on the cluster nodes, directory and mount point should be created.

After completing these configurations, use the output of the command **`df-khT`** to verify if the shared disk directory is created and mounted properly.

Scalability Setup for New Deployments of MUR

The required setup for the scalable MUR will be typically as follows:

1. Setup two or more RDP server nodes in Veritas Cluster with VxFS and Cluster Logical Volume Manager.
2. Connect these RDP nodes to a common shared storage, for example, Sun StorageTek 2540 array.
3. Mount a common partition with VxFS file system from all the nodes.
4. Install MUR in RDP modes on all RDP nodes.
5. While installing MUR, make sure that each instance of the MUR has same user group, same user ID and user name.
6. Specify a path on the shared partition where files will be received from gateway. Create required directory structure. Make sure that when the gateway does FTP, it has appropriate permissions to write the files to specified directory structure.
7. Make sure the final directory where the files will be received, has user ownership (chown), group ownership (chgrp) and all write (777) permissions for RDP's user and group. This would enable the individual MUR processes to create further temporary files under incoming EDR directories.
8. Configure the gateway to push files to one of the RDP's IP address, and the destination path being the above mentioned shared location.
9. Install the MUR as Master on another node.
10. As in the hierarchical model, add RDPs to the master through appropriate regions.

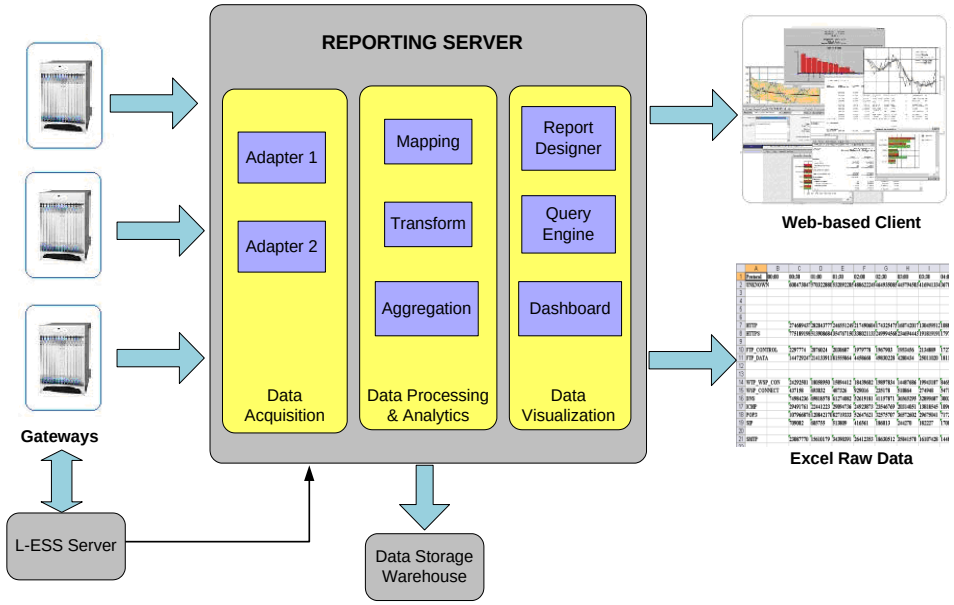
MUR Features

This section provides information on the basic features of MUR application and its implementation.

MUR Deployment

The following figure illustrates how the MUR reporting server interacts with the gateways and generates the reports.

Figure 146. End-to-end Component Mapping



The chassis / gateway supports on board Hard Disk Drive (HDD) for extended storage of the xDR files such as EDR, UDR, CDR, and NBR. If the HDD is configured, then the gateway pushes the files to an external entity like External Storage Server (ESS) for short-term storage. In case of no HDD support on the gateway, the Local, short-term External Storage Server (L-ESS) has the capability of pulling the files from gateways via SFTP, and send it for report processing. For more information on L-ESS, refer to the *ESS Installation and Administration Guide*.

The MUR server collects the EDRs, and bulkstats from gateways or L-ESS server, and processes the incoming data files and presents reports on Web-based GUI. The MUR application can generate reports in Excel, CSV, and PDF formats, and present them to users on a request basis.

Important: L-ESS is NOT required as the ASR5K EDR module can be configured to push the xDRs directly to the MUR reporting server. Push from ASR chassis is the Cisco recommended deployment model. Currently, L-ESS is supported only on Solaris platforms. For information on the L-ESS installation instructions, refer to the *ESS Installation and Administration Guide*. Existing deployments where L-ESS is installed, to pull EDRs from chassis, may continue with their deployment model in the 12.0 version of MUR Software Release and later.

For information on how to configure the chassis to push the xDRs, refer to the *Configuring Chassis for Mobility Unified Reporting System* chapter in this guide.

MUR System Requirements

This section identifies the minimum system requirements that are required for the deployment of MUR at the operator's premises.

Important: The hardware required for MUR may vary depending on incoming EDR generation, subscriber count, and number of gateways.

Server Recommendations for Use in Solaris Environment

This section identifies the minimum system requirements recommended when installing the MUR application in Solaris environment.

NEBS Requirements:

The following are the server specifications for MUR when an additional external storage is required:

- Sun Microsystems Netra™ X4270 server
  - Quad-Core two socket Intel Xeon L5518 processor
  - 32GB RAM
  - 2 * 300GB 10K RPM SAS disks
  - SATA DVD drive
  - 8 internal port SAS HBA
  - Choice of AC or DC power supplies
- Sun StorageTek 2540 SAS Array, Rack-Ready Controller Tray
  - 12 * 300GB 15K RPM SAS drives
  - Two redundant AC power supplies
- Operating system:
  - Sun Solaris 10 with latest patches installed

Non-NEBS Requirements:

The following are the server specifications with only the internal storage used:

- Sun Fire X4270 server
  - Intel Xeon processor 5500 series
  - 32GB RAM
  - 16 * 300GB 10K RPM SAS disks
  - SATA DVD drive
- Operating system:
  - Sun Solaris 10 with latest patches installed

Important: It is strongly recommended to update the Operating System with the latest security patches.

Important: The number of disks recommended is purely based on the throughput of the network and data retention configuration. Please contact Cisco Advanced Service Team for data sizing.

ZFS Pooling Recommendations:

This section provides information on the recommendations for ZFS pooling.

- OS pool: This mirrored ZFS pool shall be created for Solaris OS installation.
- MUR pool: This standard ZFS pool shall be created for MUR i.e. MUR installation, incoming data files.
- Postgres pool: This standard ZFS pool shall be created for MUR postgres database.
- Archive pool: This standard ZFS pool shall be created for retaining archived and data backed up files.

Important: ZFS pool shall NOT be created with RAID-Z since ZFS does not allow attaching an additional disk to an existing RAID-Z pool. Hence, this freezes the chances of data scaling.

Server Recommendations for Use in RHEL Environment

This section identifies the requirements of server recommended when installing the MUR application in RHEL environment.

- UCS C460 M2 server
  - 4 x Intel® Xeon® E7-4860 @ 2.26 GHz, 130W 10 Core CPU / 24 MB Cache
  - 128GB RAM
  - 12 * 600 GB SAS 6G, 10K RPM
  - RAID Controller
  - 4Gb Dual port FC Host Bus Adapter

    Important: The number of disks recommended is purely based on the throughput of the network and data retention configuration. Please contact Cisco Advanced Service Team for data sizing.

- Operating System
  - Cisco UCS running OS version 'Cisco MITG RHEL 5.5'

    For information related to OS installation, refer to the *Cisco MITG RHEL OS v5.5 Application Note*.

    Important: The Cisco MITG RHEL v5.5 OS is a custom image that contains only those software packages required to support compatible Cisco MITG external software applications. Users must not install any other applications on servers running the Cisco MITG v5.5 OS. For detailed software compatibility information, refer to the *Cisco MITG RHEL v5.5 OS Application Note*.

    Important: ZFS Pooling recommendations are applicable ONLY for Solaris hardware.

- XFS/EXT-3 File System Volumes & RAID Recommendations

Hardware Requirements for Scalable Model of MUR

- UCS 460 Server or Oracle x86 4270 Server
- Sun StorageTek 2540 SAS Array
- Other miscellaneous components such as cables, switches, etc same as mentioned in the previous sections of this guide.

Software Requirements for Scalable Model of MUR

- Veritas Cluster 5.1
- Veritas cluster suite
- VxFS as cluster file system
- Veritas Cluster Logical Volume Manager

NOTE: Solaris OS clustering is NOT recommended for the following reasons:

- If there are any deployments that have a Solaris machine running Solaris OS (either as Master or as RDP or both), and if you want to upgrade to scalable model, following cases arise.
  - In a standalone installation —
    - Keep existing machine as master and have Linux machines as RDPs. If the RDPs are Linux, then Veritas cluster would do.
    - Another option is to add all Solaris machines with Solaris OS. This is not recommended as Solaris machine would be much more costlier than UCS.
  - If hierarchical installation exists —
    - Keep master as Solaris OS. For the RDP, there are two possibilities —
    - Add another RDP as Solaris OS and have Solaris cluster. This is not recommended as Solaris machine would be costlier than UCS.
    - If the existing Solaris machine is NOT Linux compatible (e.g. SPARC machine), it would be good to replace current Solaris machine with UCS and purchase another UCS for scalability purpose.
    - If the existing Solaris machine is Linux compatible (e.g. x86 architecture), it would be good to upgrade it to RHEL and have it as cluster.

MUR Ports

This section provides information on various ports and their corresponding port numbers used by the MUR application.

Various ports are used by the MUR for both client-server communication and communication with ASR chassis. If firewalls are used on these interfaces, these ports need to be opened.

The following table lists the ports that are used by MUR.

Table 69. Default Port Utilization

Port Name	Port Number	Usage
TCP Port	22	This port is used by MUR administrator to connect via SSH to UNIX command line on MUR servers for system administration. This port is also used by gateway to upload files via SFTP to MUR servers (stand-alone master and RDPs), and also by RDPs to upload files to the master. In the case of pull model, the L-ESS process on the RDPs or stand-alone master will use SFTP to connect to this port on the gateway. This port is also used between master MUR server and gateway to configure and upload bulkstat files.
TCP Port	25	This port is used to send e-mails to a mail server in case these are configured to deliver reports and alarms.
UDP Port	162	This port is used to send traps to the northbound network management system.
Postgres Port	5432	This port is used by the local processes to access the PostgreSQL server and can be restricted to prevent external access.
Apache Port	8080	For a standalone model: This port is used for communication between client workstation and Apache Webserver on MUR via HTTP. For distributed model: This port is used for both Master to RDP and RDP to Master RPC communication. ***Important:*** When firewall is used, Apache is the only port that should be kept opened.

Typically, MUR starts all its related services with non-root (i.e. muradmin) privileges.

Firewall Settings

When MUR is running on RHEL platform, Firewall is ON by default. In that case, user will NOT be able to get access to MUR GUI. The Firewall MUST be disabled with the following commands:

```
service iptables save

service iptables stop

chkconfig iptables off
```

Using Apache Port

This section provides information on how to configure the Apache port to use in conjunction with the MUR reporting server.

Using Apache in Solaris

In case the user wants to configure Apache port as 80 (i.e. < 1024), it is necessary to run the following command as **root** user so that *muradmin* can start the services on ports < 1024.

```
usermod -K defaultpriv=basic,net_privaddr <mur admin user>
```

Using Apache in RHEL

Important: Make sure that you disable Firewall before using the Apache port in the RHEL environment.

RHEL does not allow port 80 to be used by non-root users. However, Apache Web server requests made on port 80 can be redirected to a port >1024 defined by the operator, with the following two commands:

```
iptables -t nat -A PREROUTING -p tcp --dport 80 -i eth0 -j REDIRECT --to-port <user
defined port> 1024>
```

```
iptables -t nat -AOUTPUT -p tcp -d 127.0.0.1 --dport 80 -j REDIRECT --to-port <user
defined port>  1024>
```

For example, to redirect requests made on port 80 to port 8080:

```
iptables -t nat -A PREROUTING -p tcp  --dport 80 -i eth0 -j REDIRECT --to-port 8080
```

```
iptables -t nat -AOUTPUT -p tcp -d 127.0.0.1 --dport 80 -j REDIRECT --to-port 8080
```

Once this is done, user will be able to access the MUR GUI directly, without specifying the port in the Web browser URL *http://<serveripaddress>*.

Chapter 22
Mobile Video Gateway Overview

This chapter contains general overview information about the Cisco® Mobile Video Gateway, including:

- Product Description
- Network Deployments and Interfaces
- Features and Functionality
- How the Mobile Video Gateway Works

Product Description

The Cisco® Mobile Video Gateway is the central component of the Cisco Mobile Videoscape. It employs a number of video optimization techniques that enable mobile operators with 2.5G, 3G, and 4G wireless data networks to enhance the video experience for their subscribers while optimizing the performance of video content transmission through the mobile network.

Platform Requirements

The Mobile Video Gateway software runs on a Cisco ASR 5000 chassis functioning as a mobile gateway, enabling the ASR 5000 to function as an integrated Mobile Video Gateway. In this software release, the Mobile Video Gateway software can be integrated with the Cisco P-GW (Packet Data Network Gateway), the Cisco GGSN (Gateway GPRS Support Node), and the Cisco HA (Home Agent). The Mobile Video Gateway software runs on the StarOS operating system. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the installation guide for the chassis and/or contact your Cisco account representative.

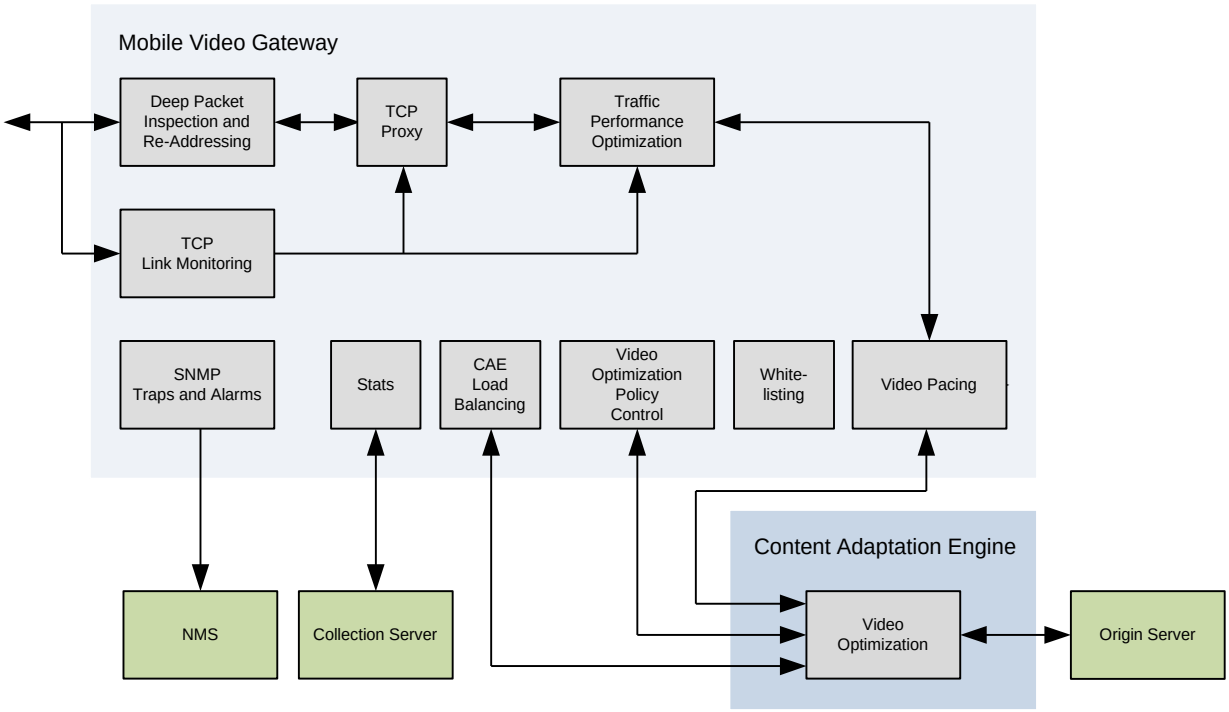
Licenses

The Mobile Video Gateway is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the "Managing License Keys" section of the "Software Management Operations" chapter in the *System Administration Guide*.

Summary of Mobile Video Gateway Features and Functions

The following figure shows the Mobile Video Gateway features and functions.

Figure 147. Mobile Video Gateway Features and Functions



The Mobile Video Gateway features and functions include:

- DPI (Deep Packet Inspection) to identify subscriber requests for video vs. non-video content
- Transparent video re-addressing to the Cisco CAE (Content Adaptation Engine) for retrieval of optimized video content
- CAE load balancing of HTTP video requests among the CAEs in the server cluster
- Video optimization policy control for tiered subscriber services
- Video white-listing, which excludes certain video clips from video optimization
- Video pacing for "just in time" video downloading
- TCP link monitoring
- Dynamically-enabled TCP proxy
- Traffic performance optimization
- N+1 redundancy support
- SNMP traps and alarms (threshold crossing alerts)
- Mobile video statistics
- Bulk statistics for mobile video

The Cisco CAE is an optional component of the Cisco Mobile Videoscape. It runs on the Cisco UCS (Unified Computing System) platform and functions in a UCS server cluster to bring additional video optimization capabilities to

the Mobile Videoscape. For information about the features and functions of the Cisco CAE, see the CAE product documentation.

Network Deployments and Interfaces

This section shows the Mobile Video Gateway as it functions in various wireless networks. The section also includes descriptions of its logical network interfaces.

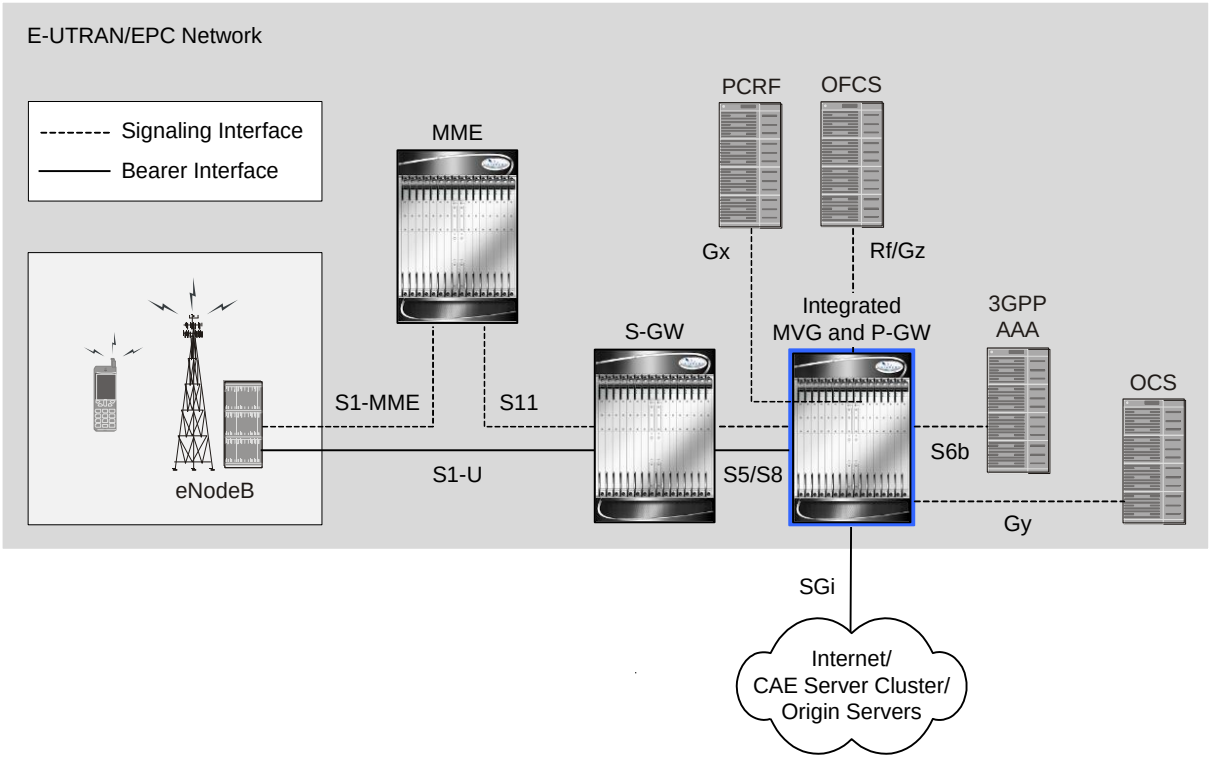
The Mobile Video Gateway in an E-UTRAN/EPC Network

In this software release, the Mobile Video Gateway software can be integrated with the Cisco P-GW in an E-UTRAN/EPC (Evolved UTRAN/Evolved Packet Core) network.

In the EPC (Evolved Packet Core), the Cisco P-GW (Packet Data Network Gateway) is the network node that terminates the SGi interface towards the PDN (Packet Data Network). The P-GW provides connectivity to external PDNs for the subscriber UEs by being the point of exit and entry of traffic for the UEs. A subscriber UE may have simultaneous connectivity with more than one P-GW for accessing multiple PDNs. The P-GW performs policy enforcement, packet filtering for each user, charging support, lawful interception, and packet screening.

The following figure shows the integrated Mobile Video Gateway and P-GW in an E-UTRAN/EPC network.

Figure 148. Mobile Video Gateway in an E-UTRAN/EPC Network



For more information about the Cisco P-GW and its connectivity to related network elements, see the *Packet Data Network Gateway Administration Guide*.

The Mobile Video Gateway in a GPRS/UMTS Network

In this software release, the Mobile Video Gateway software can be integrated with a GGSN (Gateway GPRS Support Node) in a GPRS/UMTS (General Packet Radio Service/Universal Mobile Telecommunications System) network.

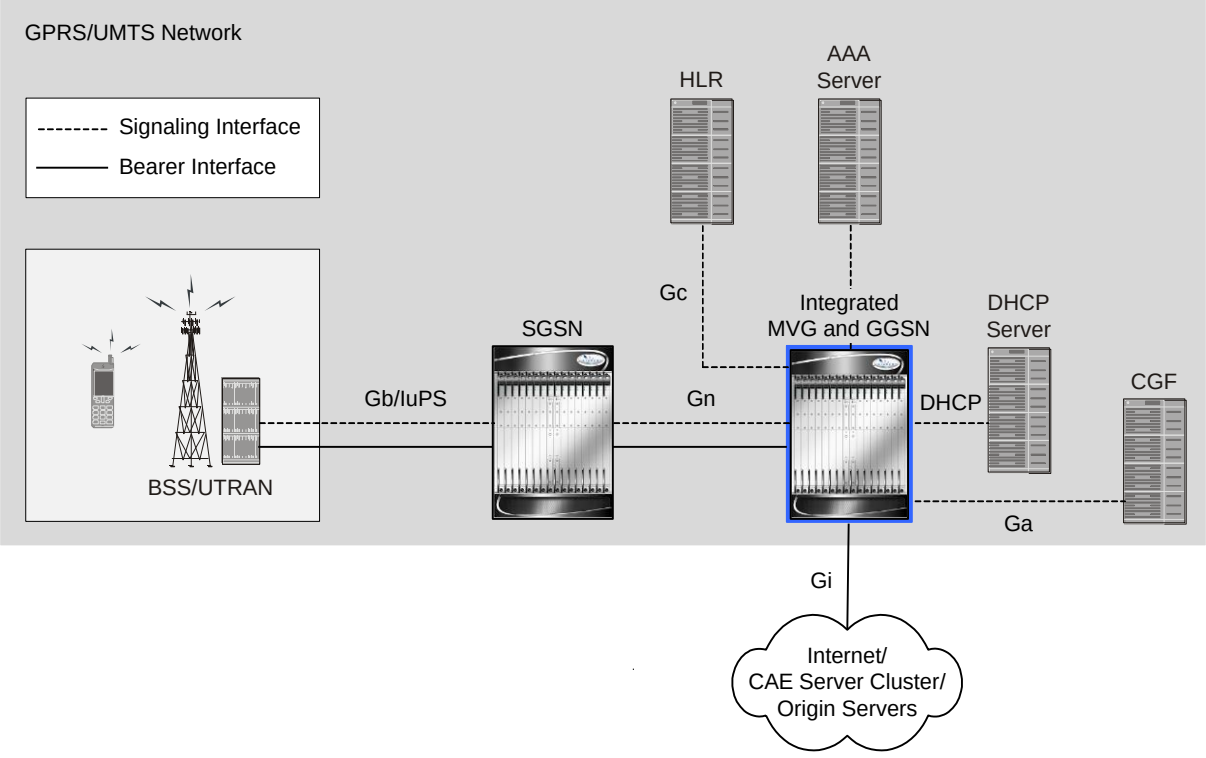
The GGSN works in conjunction with SGSNs (Serving GPRS Support Nodes) in the network to perform the following functions:

- Establish and maintain subscriber IP (Internet Protocol) or PPP (Point-to-Point Protocol) type PDP (Packet Data Protocol) contexts originated by either the MS (Mobile Station) or the network.
- Provide CDRs (Call Detail Records) to the CS (Charging Gateway), also known as the CGF (Charging Gateway Function).
- Route data traffic between the subscriber's MS and a PDN (Packet Data Network) such as the Internet or an intranet.

PDNs are associated with APNs (Access Point Names) configured on the system. Each APN consists of a set of parameters that dictate how subscriber authentication and IP address assignment is to be handled for that APN.

The following figure shows the integrated Mobile Video Gateway and GGSN in a GPRS/UMTS network.

Figure 149. Mobile Video Gateway in a GPRS/UMTS Network

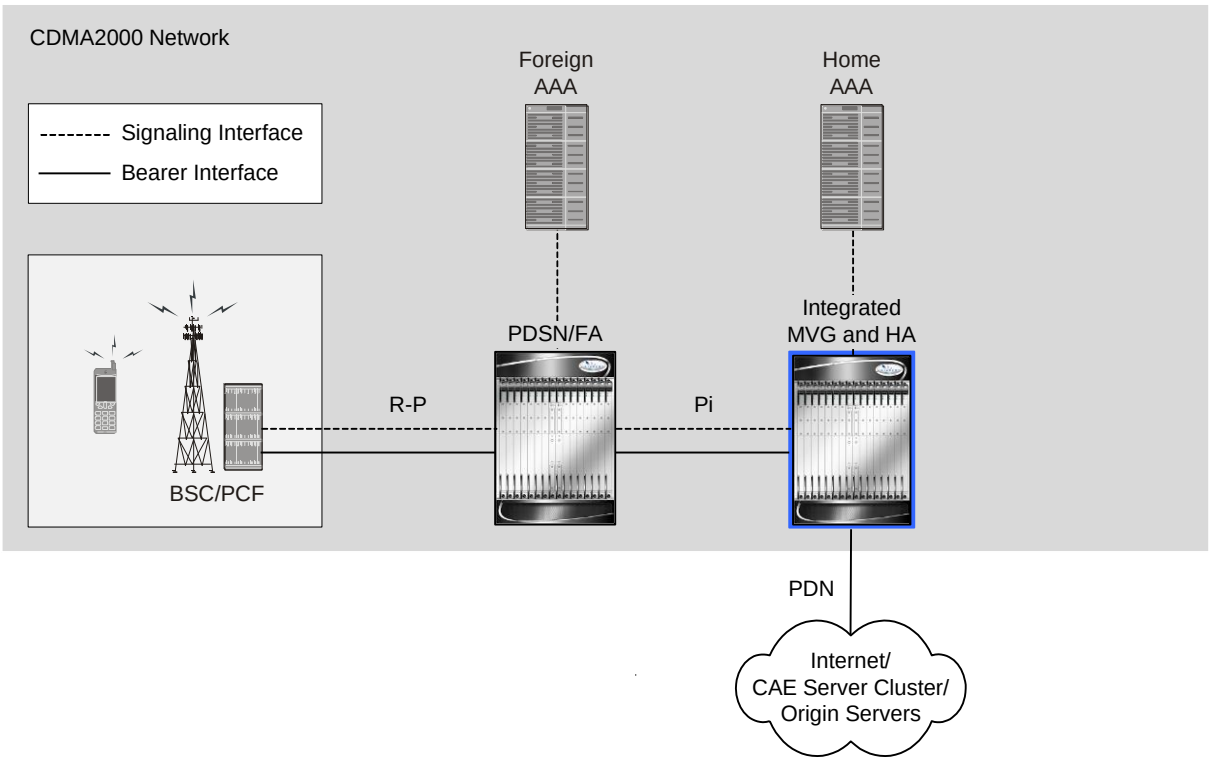


For more information about the Cisco GGSN and its connectivity to related network elements, see the *Gateway GPRS Support Node Administration Guide*.

The Mobile Video Gateway in a CDMA2000 Network

In CDMA2000 networks, the Cisco HA (Home Agent) enables subscribers to be served by their home network even when their mobile devices are not attached to their home network. The Cisco HA performs this function through interaction with the Cisco PDSN/FA (Packet Data Serving Node/Foreign Agent). The PDSN/FA provides the packet processing and redirection to the subscriber's home network via the HA. The following figure shows the integrated Mobile Video Gateway and HA with a PDSN/FA in a CDMA2000 network.

Figure 150. Mobile Video Gateway in a CDMA2000 Network

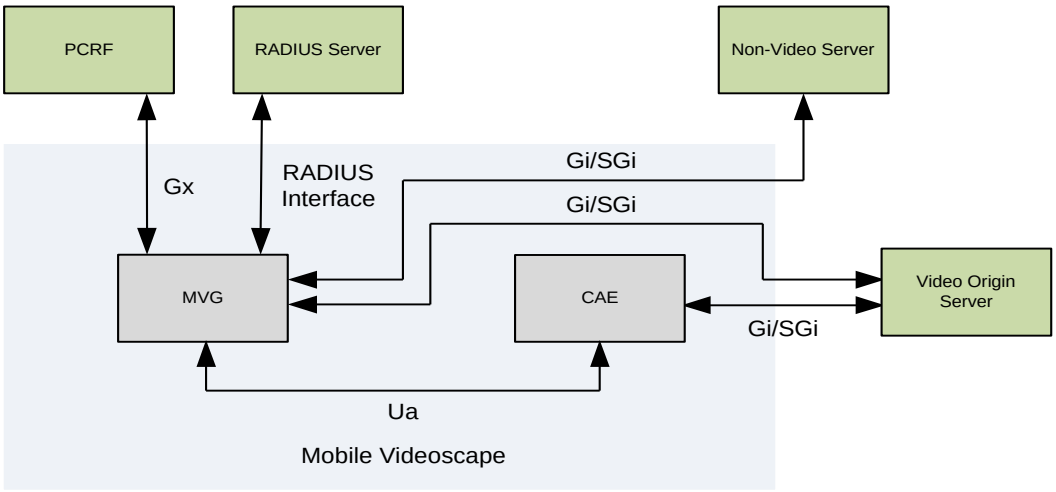


For more information about the Cisco HA and its connectivity to related network elements, see the *Home Agent Administration Guide*.

Mobile Video Gateway Logical Network Interfaces

The following figure shows the logical network interfaces on the Mobile Video Gateway.

Figure 151. Logical Network Interfaces on the Mobile Video Gateway



The following table provides descriptions of the logical network interfaces on the Mobile Video Gateway. The Mobile Video Gateway also supports the logical network interfaces of the Cisco P-GW and Cisco HA when integrated with those products.

Table 70. Logical Network Interfaces on the Mobile Video Gateway

Interface	Description
PCRF Interface	The Mobile Video Gateway can use the Gx interface to connect to a PCRF (Policy and Charging Rules Function) server to receive subscriber policy information and charging rules.
RADIUS Interface	The Mobile Video Gateway uses a RADIUS interface to exchange signaling messages with the external RADIUS server.
Video Origin Server Interface	The Mobile Video Gateway uses the Gi or SGi interface to connect to the video origin servers in the network. The Mobile Video Gateway also uses the Gi or SGi interface to connect to non-video origin servers.
CAE Interface	The Mobile Video Gateway uses a Cisco-enhanced HTTP interface called the Ua interface to connect to the Cisco CAE. The Cisco CAE is an optional component of the Cisco Mobile Videoscape.

Features and Functionality

The following features and functions are supported on the Mobile Video Gateway:

- Deep Packet Inspection
- Transparent Video Re-addressing
- CAE Load Balancing
- Video Optimization Policy Control
- Video White-listing
- Video Pacing
- TCP Link Monitoring
- Dynamically-enabled TCP Proxy
- Traffic Performance Optimization
- N+1 Stateful Redundancy
- Threshold Crossing Alerts
- Mobile Video Statistics
- Bulk Statistics for Mobile Video

Deep Packet Inspection

The Mobile Video Gateway performs DPI (Deep Packet Inspection) of HTTP traffic to identify video vs. non-video traffic based on configured Active Charging Service rule definitions. An Active Charging Service is a component of the Enhanced Charging Services on the Cisco ASR 5000.

While SPI (Shallow Packet Inspection) examines IP headers (Layer 3) and UDP and TCP headers (Layer 4) for an Active Charging Service, DPI on the Mobile Video Gateway examines URI information (Layer 7) for HTTP message information to identify video vs. non-video content based on configured rules. The following information is used for DPI:

- HTTP Request headers for matching hostnames.
- HTTP Request URLs of the destination websites to identify the video content OSs (Origin Servers).
- HTTP Response headers for matching the content type.

For more information about Enhanced Charging Services on the ASR 5000, see the *Enhanced Charging Services Administration Guide*.

Transparent Video Re-addressing

The Mobile Video Gateway can re-address HTTP video requests intended for video content OSs toward the Cisco CAE for retrieval of optimized video content. The Cisco CAE is an optional component of the Cisco Mobile Videoscape. It functions in a video server cluster to bring additional optimization capabilities to the Mobile Videoscape.

Important: The transparent video re-addressing feature is not fully qualified and is not supported for field deployment. It is available for lab demo/lab trial only.

The transparent video re-addressing feature works in conjunction with the dynamic TCP proxy feature to send video requests to the CAE cluster without using HTTP redirection, so that the re-addressing to the CAEs remains transparent to the video clients on the subscriber UEs.

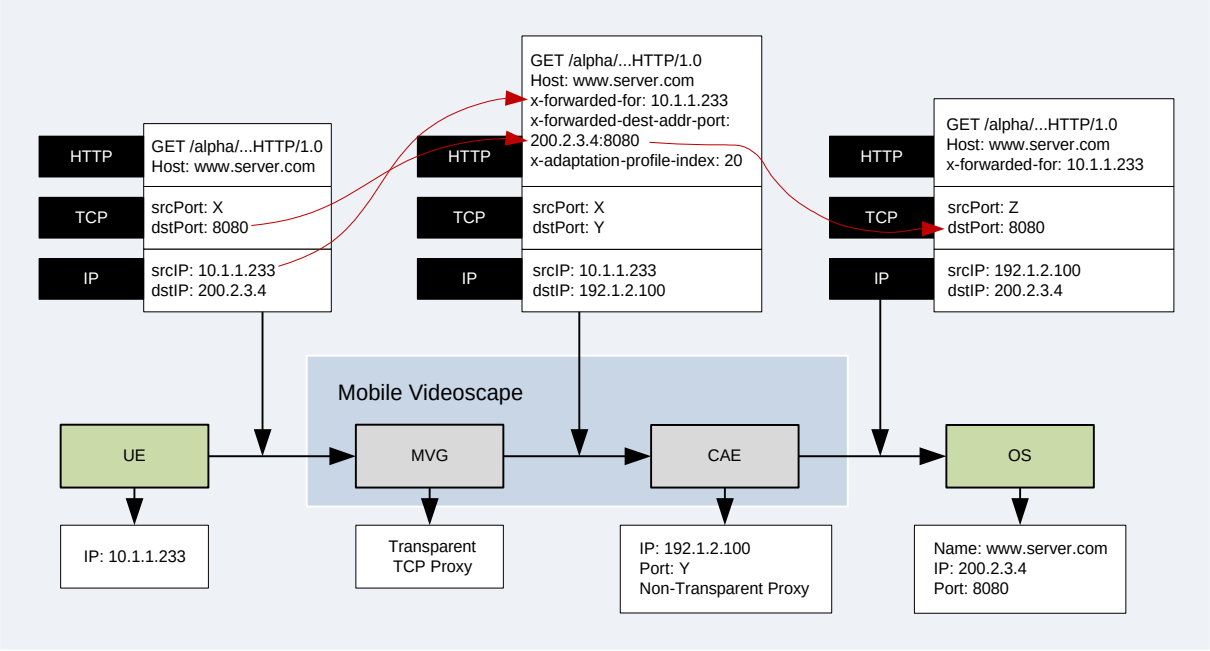
For configuration instructions and a sample configuration, see Chapter 2.

HTTP X-Header Use in Transparent Video Re-addressing

To enable the CAE to reach an OS to retrieve selected video clips for adaptation, the Mobile Video Gateway inserts the Layer 3 destination IP address and Layer 4 destination port number of the OS in a proprietary HTTP x-header in the HTTP video request to the CAE. The CAE uses the information to recreate the Layer 3 and 4 headers to connect to the OS.

The following figure shows how the HTTP x-header is used in transparent video re-addressing to the CAE. In this example, in the original HTTP request from the subscriber UE, the source IP address is 10.1.1.233 and the destination IP address is 200.2.3.4. The destination TCP port is 8080.

Figure 152. HTTP X-Header Use in Transparent Video Re-addressing



Mobile Video Gateway to the CAE

When sending HTTP video requests to the CAE for retrieval of optimized video content, the Mobile Video Gateway inserts the following x-headers:

- **x-forwarded-dest-addr-port:** The IPv4 destination address and TCP port number of the OS.

- **x-adaptation-profile-index:** The index number of the video quality profile for the CAE to use to select the level of video quality for adaptation.

CAE to the OS

When sending HTTP video requests to the OS for video content, the CAE removes the following x-headers:

- **x-forwarded-dest-addr-port:** The IPv4 destination address and TCP port number of the OS.
- **x-adaptation-profile-index:** The index number of the video quality profile for the CAE to use to select the level of video quality for adaptation.

When sending HTTP video requests to the OS for video content, the CAE inserts the following x-header: **x-forwarded-for:** The IPv4 address of the subscriber UE.

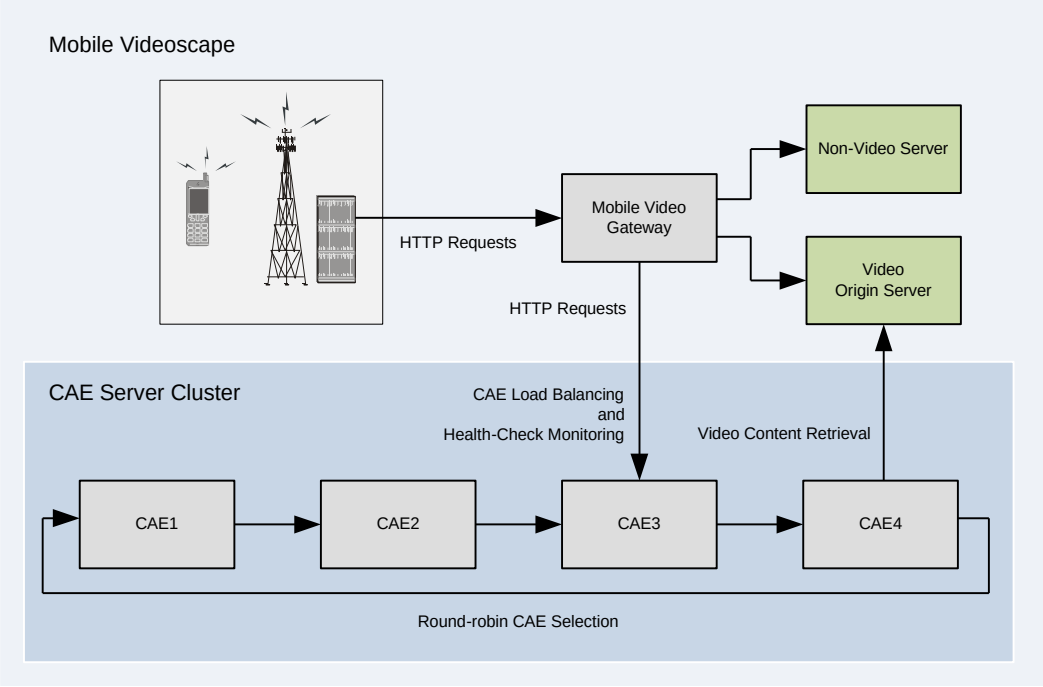
CAE Load Balancing

The optional Cisco CAE runs on the Cisco UCS platform and functions in a UCS server cluster to bring additional optimization capabilities to the Mobile Videoscape. The Mobile Video Gateway interfaces directly with each CAE in the server cluster. The CAE server cluster can serve multiple Mobile Video Gateways simultaneously. In turn, each Mobile Video Gateway is able to support up to 64 CAEs in the server cluster.

Important: The CAE load balancing feature is not fully qualified and is not supported for field deployment. It is available for lab demo/lab trial only.

The following figure shows the CAE in a server cluster.

Figure 153. CAE Server Cluster



The CAE load balancing feature enables the Mobile Video Gateway to distribute HTTP video requests from the subscriber UEs equally among the CAEs in the server cluster.

The CAE load balancing feature is configured and enabled in the context containing the interface to the CAEs, typically the destination context, via system CLI commands. During configuration, each CAE in the server cluster gets defined in a CAE group representing the cluster. Each context on the Mobile Video Gateway can have one and only one CAE group. There can be multiple contexts that contain a CAE group, but there is a system limit of 64 CAEs supported on a Mobile Video Gateway.

In addition to the CAE group configuration above, the CAE load balancing feature gets configured as part of an Active Charging Service, which is a component of the Enhanced Charging Services on the ASR 5000. The feature is configured by creating an Active Charging Service for the Mobile Video Gateway, specifying charging and routing rule definitions, and then creating a charging action for CAE re-addressing, which enables video optimization and CAE load balancing for the CAEs in the CAE group.

For configuration instructions and a sample configuration, see Chapter 2.

CAE Load Balancer Function

When the Mobile Video Gateway identifies a video request during DPI, the CAE load balancer function performs three main operations, as follows:

- It performs CAE load balancing using a round-robin selection of the next available CAE to service the video request.
- It ensures that multiple video flows for a subscriber are serviced by the same CAE once a CAE is selected. This is required for some mobile devices such as the Apple® iPhone®, which can serve video clips using multiple TCP sessions, such as when an iPhone user skips forward in the middle of playback and the iPhone closes the existing TCP session and starts a new one.

- It maintains health-check monitoring for each of the configured CAEs in the server cluster. If a CAE is currently down, the load balancer function prevents video requests from being sent to the down CAE until is up and available again. All of the CAEs in a CAE group optimize the same video content, so the Mobile Video Gateway can direct the video request to any of the other CAEs until the down CAE is up and available again.

CAE Health-Check Monitoring Function

The CAE health-check monitoring function is part of the CAE load balancing feature. It triggers a health-check request sent to the CAEs based on a configurable keep-alive timer. If a CAE does not respond, and after a configurable number of retries and timeouts, it marks the state of the CAE as Down. It also generates an SNMP Server-State-Down trap message, indicating that the CAE is down and unavailable. When a configurable dead-time timer expires, it sends another health-check request to the down CAE, and if the CAE sends a positive response indicating that it is back up, it marks the state of the CAE as Up and generates an SNMP Server-State-Up trap message, indicating that the CAE is back up and available.

Video Optimization Policy Control

The video optimization policy control feature provides the necessary information for the Mobile Video Gateway to select the highest quality video content for a subscriber, based on information received from a PCRF or RADIUS server, or based on the subscriber's policy profile configured on the Mobile Video Gateway. The feature enables mobile operators to offer tiered video services to their subscribers with different levels of service (Gold, Silver, and Bronze levels, for example).

A video policy defines a subscriber's entitlement to the video content provided by the Mobile Video Gateway. A video policy contains various video-specific attributes, including the subscriber's video QoE (Quality of Experience).

In this software release, the video policy includes a CLI **`charging-action`** command option for specifying a suggested maximum bit rate value for video. This value, specified in bits per second (bps), is used by two of the video optimization modules on the Mobile Video Gateway, the video pacing module and the video transrater module.

The following figure shows the flow of information for the video optimization policy control feature on the Mobile Video Gateway.

Figure 154. Video Optimization Policy Control System Flow

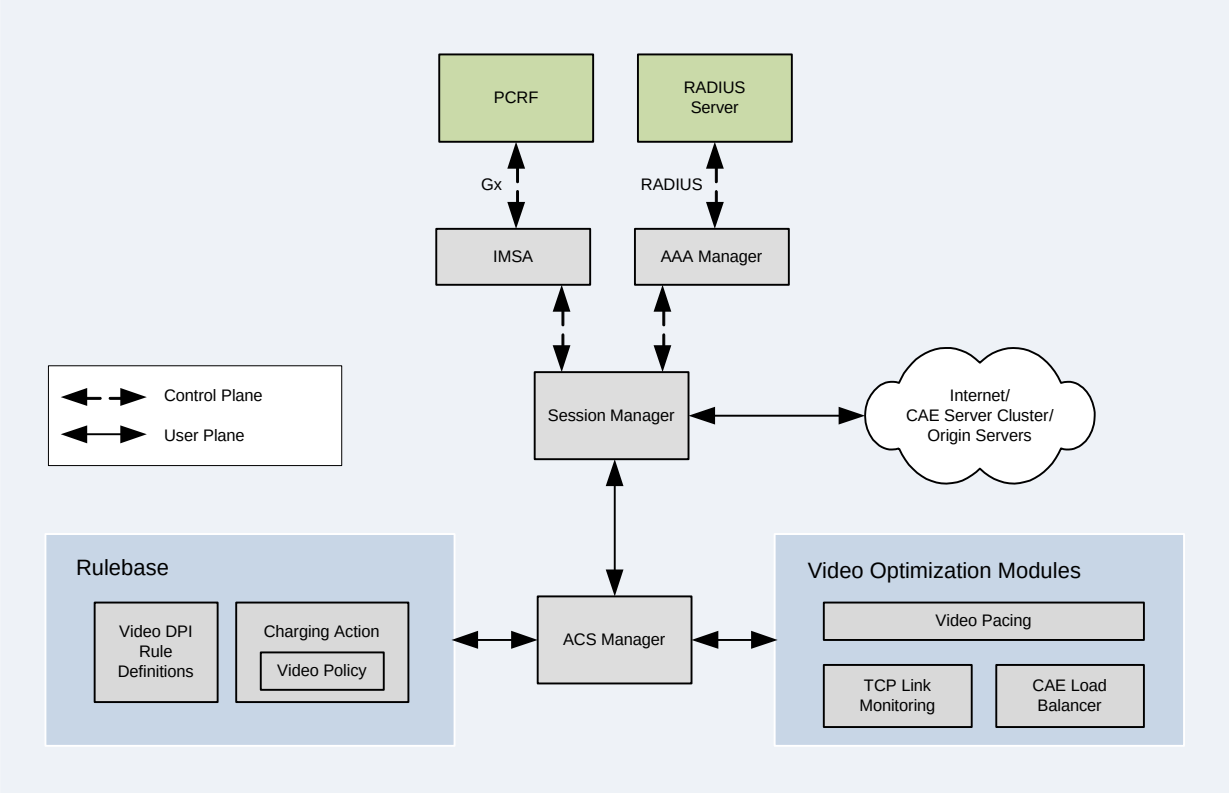


Figure 155. Video Optimization Policy Control System Flow

Functional Overview

The video optimization policy control feature assigns a video policy to a subscriber via one of the following methods:

- **PCRF via the Gx interface:** Acting as a RADIUS endpoint, the Mobile Video Gateway can obtain the video policy for a subscriber using the Gx interface to the PCRF. With this method, the Charging-Rule-Name attribute received in the Charging-Rule-Install AVP in the CCA-I message contains a rule definition name that maps to the video policy. This rule definition is part of the rulebase assigned to the subscriber. The Mobile Video Gateway can assign the rulebase to the subscriber through a static configuration at the subscriber or APN level, or obtained from the RADIUS server in an Access-Accept message.

  Alternately, the Mobile Video Gateway can be configured to obtain the rulebase name itself from the PCRF via the Charging-RuleBase AVP.

- **RADIUS Server via the RADIUS interface:** In the absence of a Gx interface, the Mobile Video Gateway can obtain the video policy from the RADIUS server through the Access-Accept message. With this method, the Mobile Video Gateway applies the RuleBase-Name AVP in the Access-Accept message to the subscriber, and one of the rule definitions in the configured rulebase selected in this manner maps to the video policy. Note that one rulebase gets associated with one level of subscriber entitlement (GOLD_RULEBASE, for example).

- **Static assignment at the subscriber or APN level:** The Mobile Video Gateway can assign a video policy by assigning a rulebase at the subscriber or APN level, so that one of the rule definitions in the configured

rulebase maps to the video policy. As in the RADIUS server method, one rulebase gets associated with one level of subscriber entitlement.

The video optimization policy control feature gets configured as part of an Active Charging Service, which is a component of the Enhanced Charging Services on the ASR 5000. The feature is configured by creating an Active Charging Service for the Mobile Video Gateway, specifying charging and routing rule definitions, and then creating charging actions for the tiered video service levels. Within each service level charging action, the suggested maximum video bit rate is specified.

During configuration, a rulebase is defined for each subscriber or APN and contains multiple rule definitions. When obtaining the video policy from the PCRF via the Gx interface, and when obtaining the video policy via the Charging-Rule-Install AVP, the Mobile Video Gateway enables a particular rule definition when a rule definition name matches the received Charging-Rule-Name attribute. This is achieved by using the **`dynamic-only`** option in the **`action priority`** command when configuring the rulebases. When obtaining the video policy via the RuleBase-Name AVP, note that there can be one and only one rule definition and its corresponding charging action associated with a video policy.

When a rule definition gets matched, the Mobile Video Gateway applies the corresponding charging action. For example, when the VIDEO_GOLD rule definition is matched, the Mobile Video Gateway applies the corresponding GOLD_CHARGING_ACTION. This charging action determines the video policy for the subscriber. If no rule definitions get matched, the Mobile Video Gateway uses the default value for the suggested maximum bit rate.

For configuration instructions and sample configurations, see Chapter 2. For detailed instructions for configuring the Gx interface on the Cisco P-GW, see the *Packet Data Network Gateway Administration Guide*.

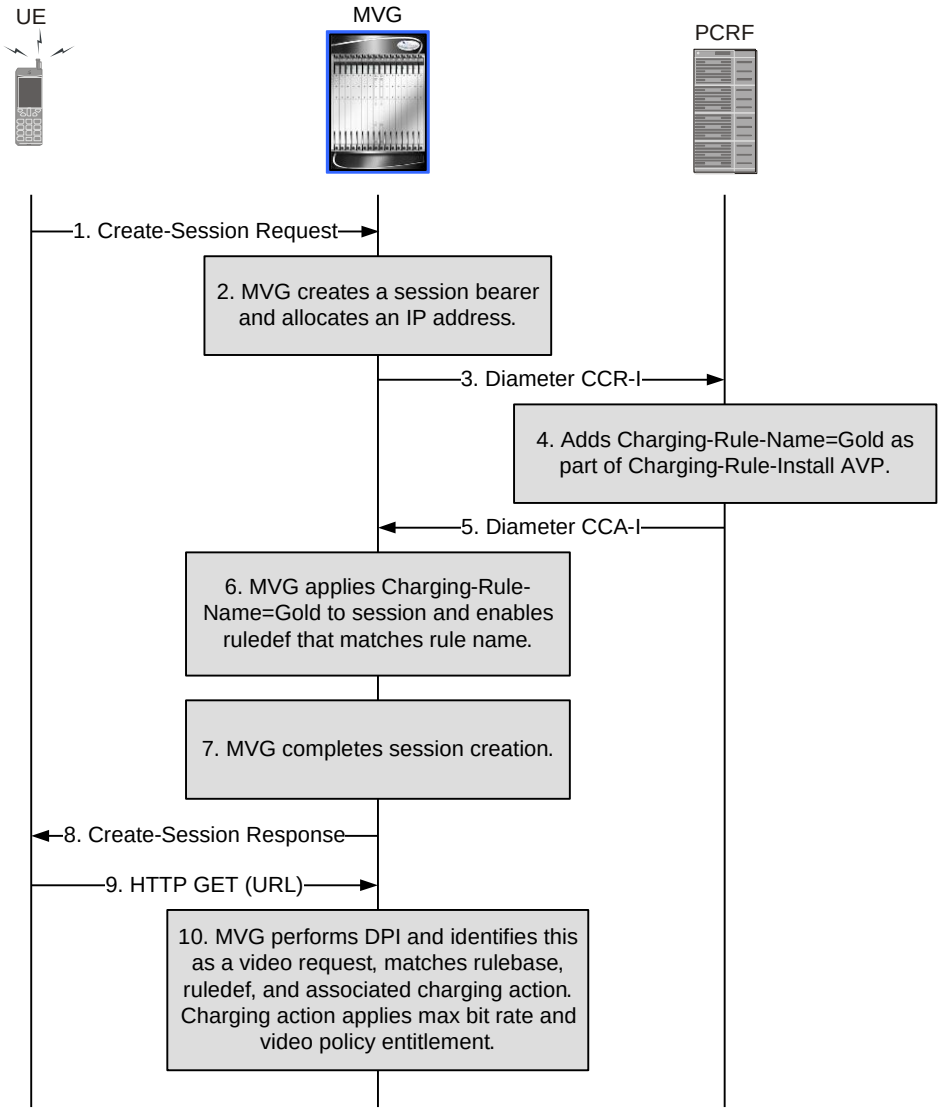
Video Optimization Policy Control Call Flows

This section includes call flows of the Mobile Video Gateway obtaining the video policy for a subscriber in two ways:

- From the PCRF over a Gx interface as it functions as a RADIUS endpoint.
- From the RADIUS server over a RADIUS interface as it functions as a RADIUS proxy.

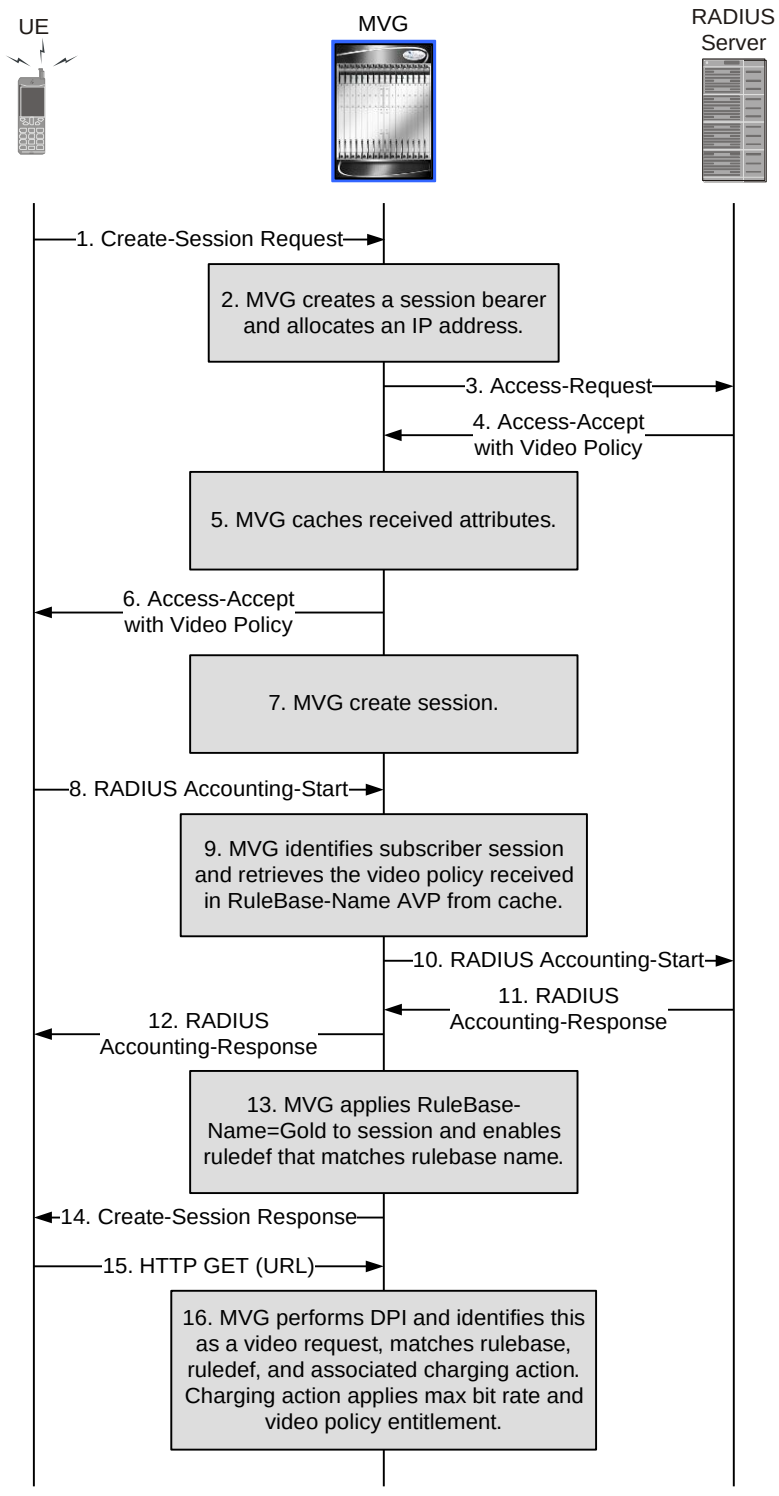
The following figure shows the Mobile Video Gateway functioning as a RADIUS endpoint obtaining the video policy via the PCRF over a Gx interface.

Figure 156. Mobile Video Gateway as a RADIUS Endpoint Obtaining the Video Policy via the PCRF



The following figure shows the Mobile Video Gateway functioning as a RADIUS proxy obtaining the video policy via the RADIUS server over a RADIUS interface.

Figure 157. Mobile Video Gateway as a RADIUS Proxy Obtaining the Video Policy via the RADIUS Server



Video White-listing

Certain video clips can be excluded from video optimization. This is referred to as white-listing. The video white-listing feature can either be configured using empty charging actions that match the white-listed URLs, or using DPI rule definitions that do not match the white-listed URLs.

For configuration instructions, see Chapter 2.

Video Pacing

The video pacing feature enables mobile operators to limit the download speed of over-the-top, progressive download video (video clips provided to subscribers via HTTP downloads over TCP flows) so that their subscribers download just enough video content in time for smooth playback. By limiting the bit rate of progressive downloads to the actual encoded bit rate of each video clip, mobile operators can significantly reduce their air interface bandwidth usage.

Important: The video pacing feature has been qualified to run on the ASR 5500 for the Mobile Video Gateway integrated with the Cisco P-GW (Packet Data Network Gateway) and the Cisco HA (Home Agent).

The video pacing feature determines the optimal download speed for a video by calculating the average bit rate of the video and then, after allowing an initial burst to fill a video buffer on the subscriber UE before playback begins, by enforcing the average bit rate for the duration of the video download.

The video pacing feature is an Active Charging Service, which is a component of the Enhanced Charging Services on the ASR 5000. The video pacing feature is configured using the system CLI commands by creating an Active Charging Service for video pacing, and then specifying charging and routing rule definitions.

For configuration instructions and a sample configuration, see Chapter 2.

Video Pacing Operation

The video pacing feature operates as follows:

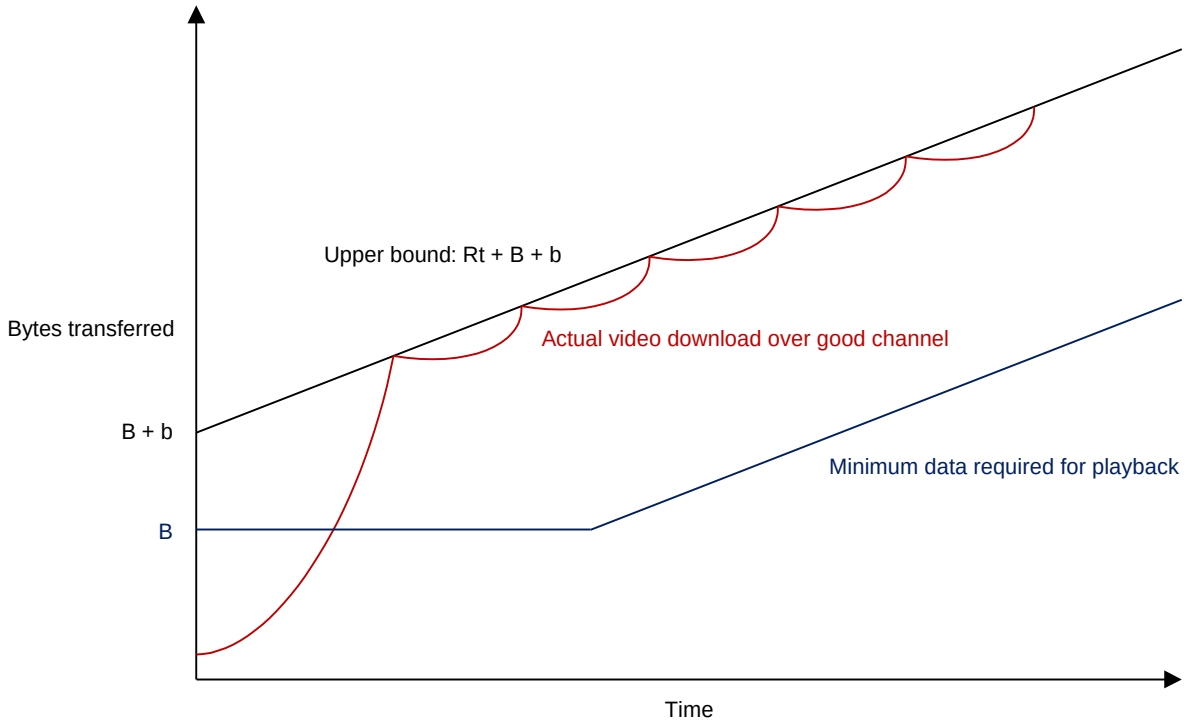
Assume a video-encoding bit rate R and a video playback start time of 0. At time t, the subscriber UE needs to receive Rt bytes of video content just in time for smooth playback. To address fluctuations over the wireless channel, assume that a video buffer is kept on the subscriber UE to accommodate these fluctuations. Assume this buffer size is the standard burst size b.

Because many software media players do not begin playback until a certain amount of video data has been buffered, the video pacing feature allows an initial burst of data, so in addition to the standard burst size b, assume an initial burst size B. This initial burst size is configured based on time duration (as t seconds of video data) and calculated for each video flow based on the determined video bit rate. The video pacing feature allows this initial burst just once, before the video begins playing.

The video pacing feature employs a token bucket algorithm to enforce the permitted video data bytes. When a video download begins, for any given time t, the token bucket algorithm disallows more than (Rt + B + b) data bytes, which is the maximum allowed data bytes. After the initial burst B is completed, the video pacing feature disallows more than (Rt + b) data bytes, and the optimal "just in time" video download rate is achieved.

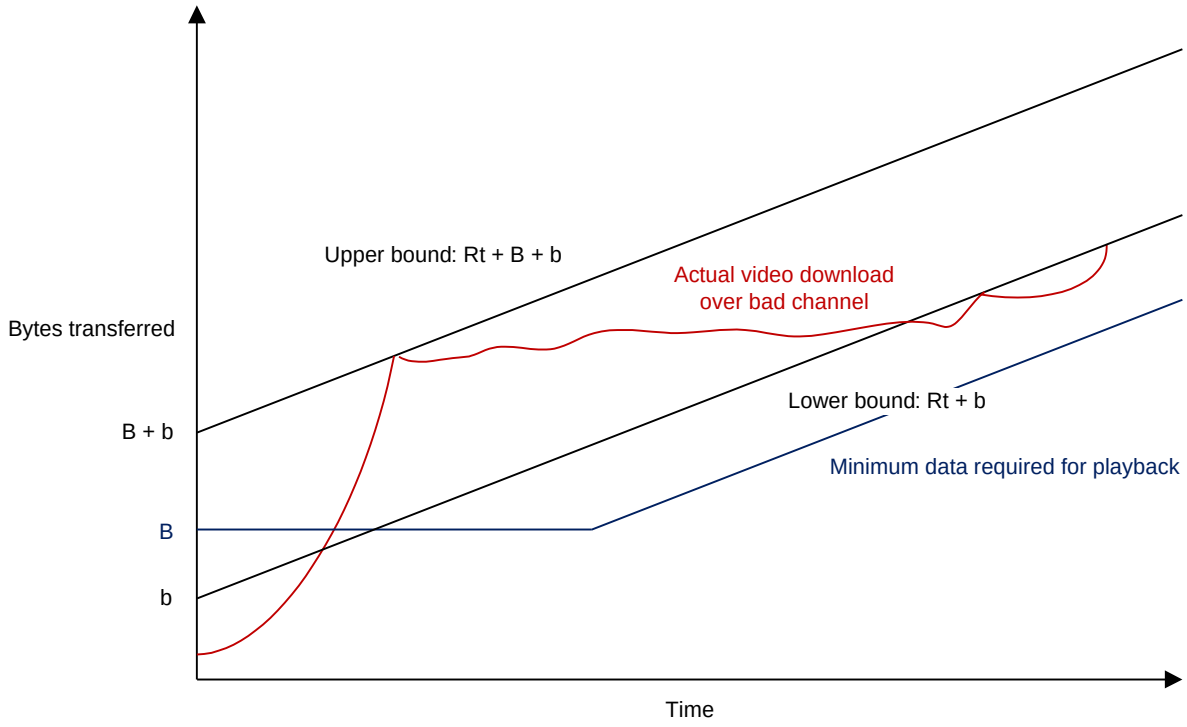
The following figures show video pacing during good and bad channel conditions.

Figure 158. Video Pacing During Good Channel Conditions



In the figure above showing good channel conditions, notice that there is a small difference between the ideal pacing rate (the black line on top) and the actual downloaded video bytes (the red line). This difference is due to network delay, and when the pacing feature begins to take action, the video content OS or Cisco CAE does not respond immediately. Even with this delay, because the video pacing feature allows the standard burst size b, the download rate never falls below the blue line representing the minimum video data required for smooth video playback. Also notice that the media player needs B (not 0) bytes of data for the video to start playing. This is why the video pacing feature allows a bigger initial burst of data (B + b), and then begins enforcing the burst size b until the completion of the download.

Figure 159. Video Pacing During Bad Channel Conditions



In the figure above showing bad channel conditions, when channel conditions worsen, the actual downloaded video bytes cannot keep up with the ideal pacing rate. Nonetheless, if the channel recovers in time, the download rate is still above the blue line representing the minimum data required for smooth playback, and video pacing continues to maintain b bytes of data above this lower limit.

Video Pacing Functions

The video pacing feature includes four main functional components, as follows:

- **Pacing Start Trigger:** The pacing start trigger is part of the Active Charging Service for video pacing. When a rule definition in the Active Charging Service identifies a packet flow as a video flow, and the corresponding charging action for video pacing is enabled, the pacing start trigger invokes video pacing enforcement for the video flow. It sets the video bit rate and initial burst size from the subscriber policy, which is configured for subscribers in the source context as part of the active charging rulebase. It then becomes dormant.

  Some mobile devices such as the Apple iPhone can serve video clips using multiple TCP sessions, such as when an iPhone user skips forward in the middle of playback and the iPhone closes the existing TCP session and starts a new one. When multiple TCP sessions are used to download the same video, the pacing start trigger gets invoked once per video flow, and the video pacing feature correlates these flows to the same video object to continue pacing enforcement from where the last TCP flow left off. When multiple TCP flows are used to download different videos, video pacing is performed independently per flow.

- **Video Pacing Enforcement:** After the initial burst of video content, the video pacing enforcement function sets the optimal video download rate for the incoming downlink packets using a token bucket algorithm. Video pacing occurs based on the settings configured via CLI command options.

- **Video Rate Determination:** The video rate determination function is a software algorithm that examines the initial HTTP RESPONSE packets and video metadata packets to determine the encoded bit rate of the video. It

examines the HTTP RESPONSE headers to determine the content length of the video in total bytes as well as the total video playback duration, and then calculates the average video bit rate as: (Content length/Video playback duration). It then triggers the video pacing enforcement function to enforce the new average bit rate when the next downlink packet is received.

- **CLI Command Options:** The video pacing feature includes a set of CLI command options for the Active Charging Service `charging-action` command.

  For a description of these command options, see the *Command Line Interface Reference*.

Video Pacing Call Flows

When the Mobile Video Gateway receives an HTTP GET request from a subscriber UE, it performs DPI to determine whether it is a request for video content. If the Mobile Video Gateway cannot make this determination by inspecting the HTTP GET request, it performs DPI again when it receives the HTTP RESPONSE from the OS.

The following figures show the message flow during inspection for video content and the subsequent triggering of video pacing functions. The first figure shows the identification of a video request from an HTTP GET request, the second shows the identification of a video request from an HTTP RESPONSE.

Figure 160. DPI of HTTP GET Identifying a Video Request

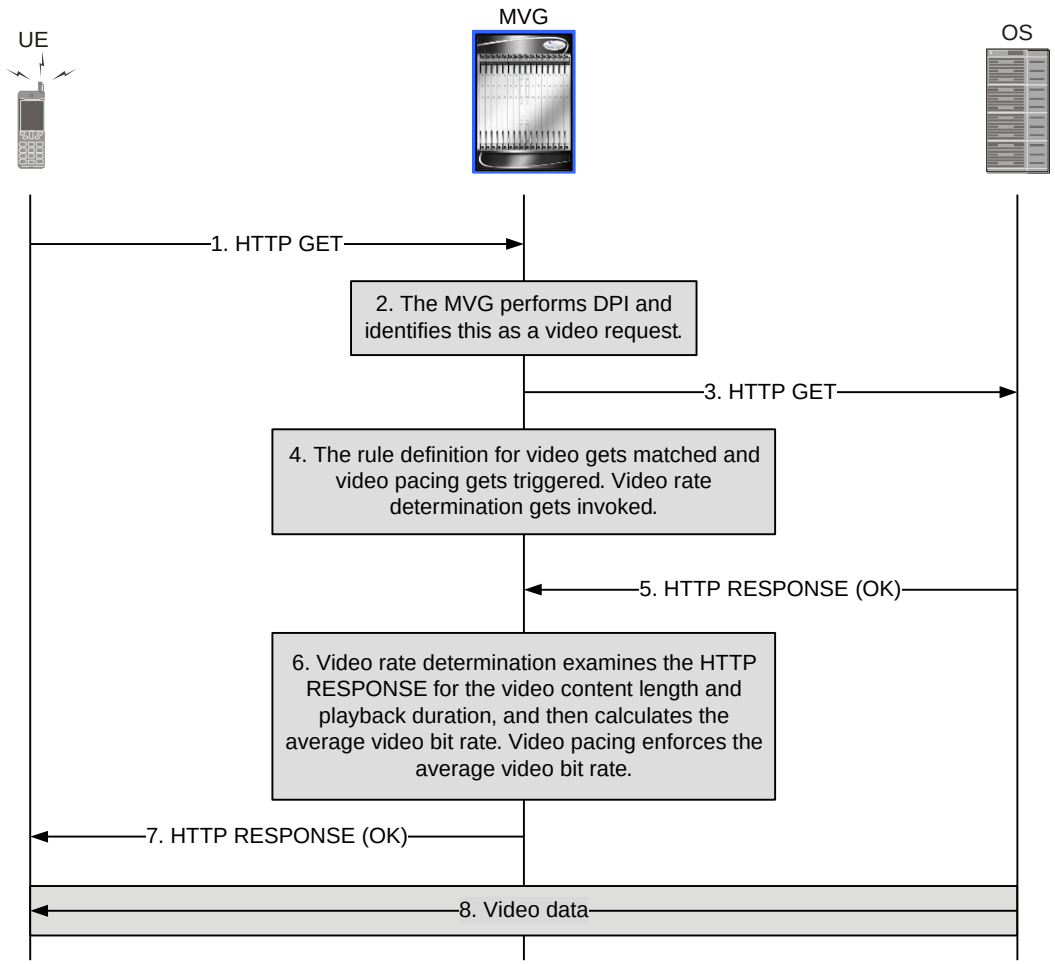
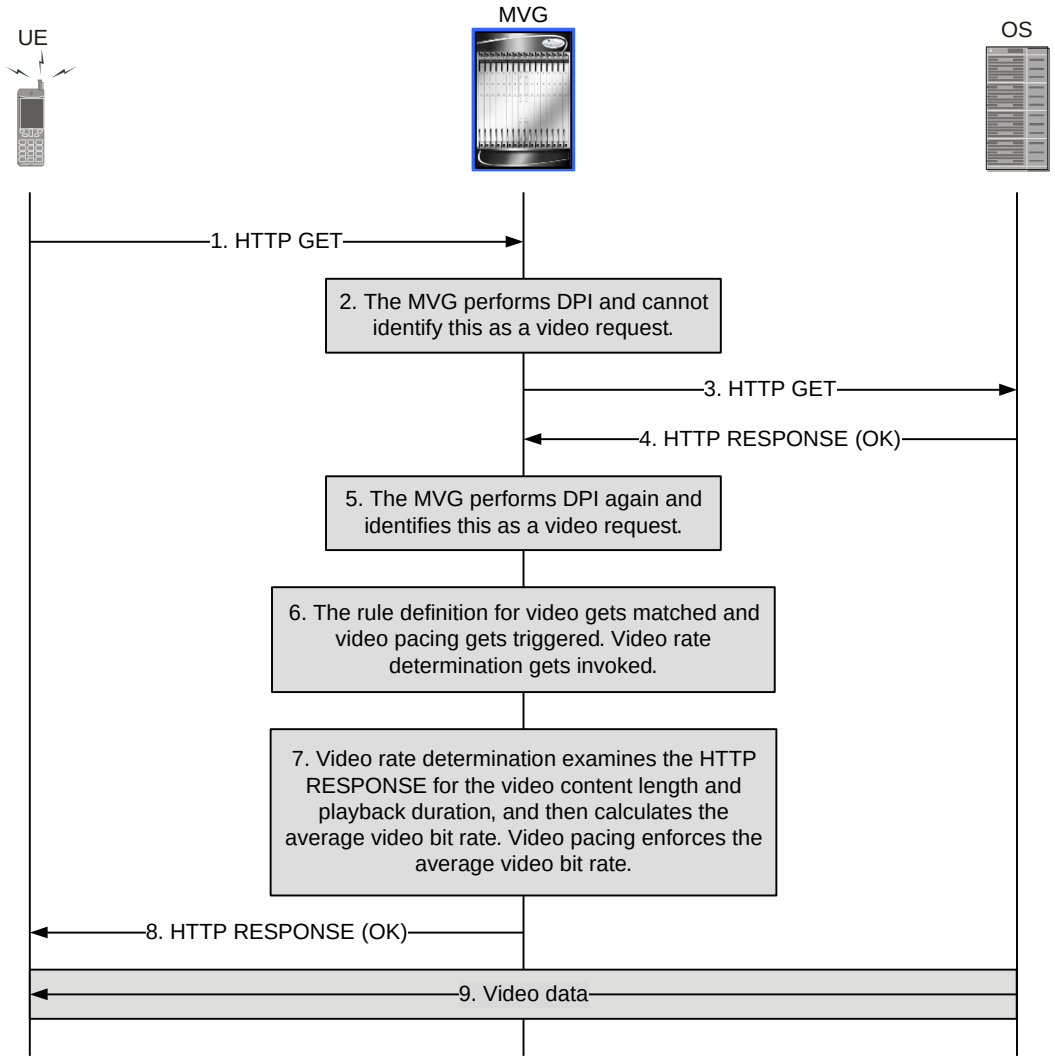


Figure 161. DPI of HTTP RESPONSE Identifying a Video Request



Interactions with Related Functions

The video pacing feature is designed to work with related functional components as follows:

- **Video Pacing and the CAE:** The video pacing feature is an independent software module and has no interface with the Cisco CAE. It performs its function in the same way whether a video is downloaded from the OS or from the CAE. The CAE is an optional component of the Cisco Mobile Videoscape.
- **Video Pacing and the TCP Proxy:** The video pacing feature can be configured to work with or without the TCP proxy feature with no change in its function.
- **Video Pacing and Traffic Performance Optimization:** The traffic performance optimization feature works over the interface on the client side of the TCP proxy. It handles re-transmission, TCP window size adjustment, and so on. Video pacing works over the interface on the video server side of the TCP proxy, and works independent of traffic performance optimization.

Supported Video Container File Formats

In this software release, the video pacing feature supports the following standard video container file formats:

- MP4 File Format
- FLV Files

MP4 follows the ISO Base Media File Format (MPEG-4 Part 12). We provide comprehensive support for progressive download of .FLV files, playable in Adobe® Flash® Player.

TCP Link Monitoring

TCP is the dominant transport protocol for the majority of Internet traffic, including video. For mobile networks, the available transport bandwidth can fluctuate depending on changing conditions over the wireless connections. Knowledge of the available transport bandwidth is especially important for video over mobile networks, since this bandwidth affects video delivery rates, video encoding and compression techniques, and ultimately the video playback experience of the subscribers.

The TCP link monitoring feature adds the capability to enable monitoring and logging of TCP behavior towards the subscriber UEs. Monitoring TCP behavior enables the Mobile Video Gateway to estimate transient bandwidth and identify network congestion for all TCP connections toward the clients on the subscriber UEs.

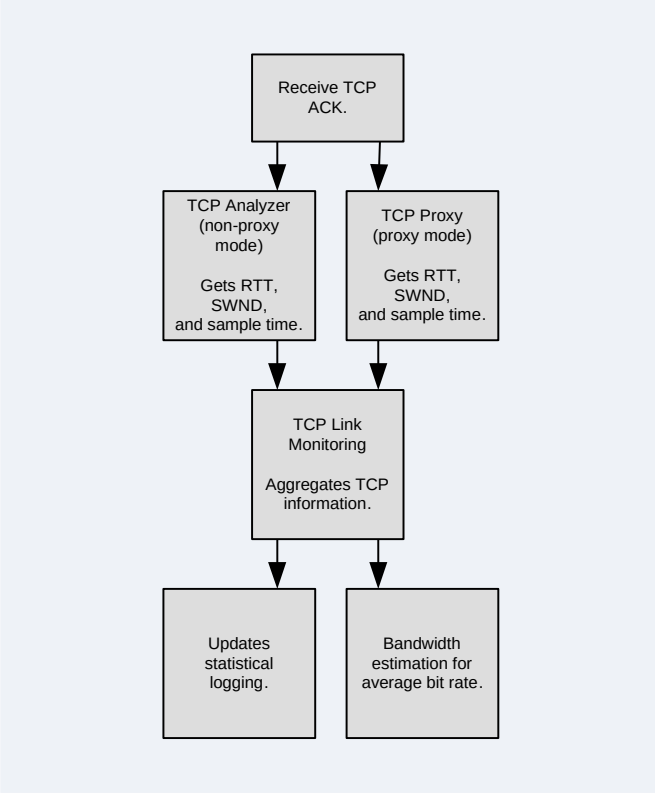
The Mobile Video Gateway services two types of TCP connections. A TCP connection can either pass through the Mobile Video Gateway intact or can be split into two connections by the TCP proxy. For the downlink data towards the subscriber UEs, the TCP link monitoring feature invokes its bandwidth estimation and statistical logging functions, which are enabled for both proxy and non-proxy modes.

TCP link monitoring statistics are gathered on a system-wide basis. This information can be periodically exported to a collection server as bulk statistics, upon which post-processing can be performed.

TCP Link Monitoring System Flow

The following figure shows the flow of information to and from the TCP link monitoring module on the Mobile Video Gateway.

Figure 162. TCP Link Monitoring System Flow



The TCP link monitoring feature calculates the RTT (Round Trip Time) and estimates the link bandwidth based on the downlink data sent towards the UE and the current congestion conditions. It then collects this information at the system level to report to the bulk statistics collection server.

Note that the throughput calculation for the TCP link excludes duplicate, out-of-order, and retransmitted packets.

Functional Overview

The key functions of the TCP link monitoring feature are bandwidth estimation and system-level TCP statistical logging.

Bandwidth Estimation

Because mobile devices are served by a variety of TCP variants, either from the OS or from the Mobile Video Gateway's TCP proxy, the TCP link monitoring feature employs an independent bandwidth estimation technique proposed by TCP Westwood+ (see "Performance Evaluation of Westwood+ TCP Congestion Control" by Mascolo, et al).

Westwood+ estimates bandwidth by calculating the ratio of the number of bytes of acknowledged TCP payload over every RTT. This rate sample is then filtered by a weighted moving average to derive a per-flow average bandwidth estimate for every RTT interval.

Statistical Logging

Statistical logging of TCP traffic supports two types of plots: histogram and time-series.

For histogram logging, the TCP link monitoring feature keeps a counter for every bit rate or RTT range. Whenever a new sample of TCP traffic is generated, a corresponding counter is updated. The collection server retrieves these values based on the configured sampling rate. There are four histogram plots: video bit rate, video RTT, non-video bit rate, and non-video RTT. For each of these plots, a total of 36 counters are used for logging.

For time-series logging, the sampling rate is the same as that of the remote update time for the collection server. Typically, this can be configured in 30-minute intervals. As with histogram logging, there are four time-series counters: video bit rate, video RTT, non-video bit rate, and non-video RTT.

Dynamically-enabled TCP Proxy

The Mobile Video Gateway can act as a dynamically-enabled TCP proxy that provides the following functions:

- Transparent video re-addressing
- Traffic performance optimization

Note that these features require the TCP proxy to function as expected.

The TCP proxy can be dynamically enabled based on Active Charging Service rule definitions. For information about the dynamically-enabled TCP proxy, including configuration instructions, see the *Enhanced Charging Services Administration Guide*.

Traffic Performance Optimization

The Mobile Video Gateway can use traffic performance optimization to improve latency and accelerate the delivery of video content, especially when network congestion and packet drops are present. The feature runs on the dynamically-enabled TCP proxy and can be enabled statically based on the subscriber profile or dynamically based on a DPI match in a charging action.

For information about traffic performance optimization, including configuration instructions, see the *Traffic Performance Optimization Administration Guide*.

N+1 Stateful Redundancy

In the telecommunications industry, over 90 percent of all equipment failures are software-related. With robust hardware failover and redundancy protection, any card-level hardware failures on the system can quickly be corrected. However, software failures can occur for numerous reasons, many times without prior indication.

This software release supports N+1 stateful redundancy for mobile video sessions. N+1 stateful redundancy provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system, preventing fully-connected subscriber sessions from being disconnected. Sessions are maintained over a software failure of a process or hardware failure.

This is an existing feature of the ASR 5000. Note that Layer 4 flows will not be maintained across switch-overs.

Threshold Crossing Alerts

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outages. Typically, these conditions are temporary (high CPU utilization or packet collisions on a network, for example) and are

quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports threshold crossing alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure a threshold on these resources whereby, should the resource depletion cross the configured threshold, an SNMP trap would be sent.

The following thresholding models are supported by the system:

- **Alert:** A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated, then generated and/or sent again at the end of the polling interval.
- **Alarm:** Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated, then generated and/or sent again at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps:** SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values. Generation of specific traps can be enabled or disabled on the chassis, ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.
- **Logs:** The system provides a facility for which active and event logs can be generated. As with other system facilities, logs are generated messages pertaining to the condition of a monitored value and are generated with a severity level of WARNING. Logs are supported in both the Alert and the Alarm models.
- **Alarm System:** High threshold alarms generated within the specified polling interval are considered outstanding until a condition no longer exists or a condition clear alarm is generated. Outstanding alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

For more information about threshold crossing alert configuration, see the *Thresholding Configuration Guide*.

Mobile Video Statistics

The mobile video statistics feature enables mobile operators to collect detailed statistics on mobile video usage to understand how subscribers behave when viewing video content, how much network resources are consumed by video, and what trends develop as video use cases evolve. The mobile video statistics feature collects important statistical data for video and presents this information in three ways: per user device type, per radio access type, and per video container type. With this information, operators can better understand evolving trends in their network and further adapt and fine tune their video optimization solution accordingly.

In this software release, the identification of a video flow is dependent on charging actions defined within the corresponding Active Charging Service. When a flow matches a rule definition for video during DPI, the mobile video statistics feature begins collecting the following statistics for the video flow:

- **Total size of the video file (the HTTP content length):** This is the size given in the HTTP RESPONSE header for the video file, represented in bytes.
- **Total duration of the video clip:** This is the video play duration identified from the video metadata, represented in seconds. If the mobile video statistics feature cannot get this information from the metadata (due to non-standard metadata formatting, etc.), this field shows 0.
- **Total bytes sent to the UE:** This is the payload data bytes (excluding TCP/IP headers) permitted to be sent towards the UE. Note that this counter includes end-to-end (TCP) retransmissions.

- **Total duration that the video object is on:** This is the time it takes for the UE to finish downloading the video, which is from the creation of the first flow to the deletion of the last flow comprising this video.
- **Total number of TCP flows used to download the video:** The total count of TCP sessions used for this video object.

The mobile video statistics feature also derives the following information from the statistics above:

- **Video delivery rate:** Total bytes sent to the UE/Total duration that the video object is on. This is the average bit rate of the video payload bytes being delivered to the UE, represented in bps.
- **Percentage of video download:** Total bytes sent to the UE/Total size of the video file. This is the percentage of the video file that the user actually downloaded. The number reflects whether users tend to watch the entire video or only a small part of it. Note that since "Total bytes sent to the UE" includes retransmissions, this number can be larger than 100%.
- **Video encoding bit rate:** Total size of the video file/Total duration of the video clip. This is the average video encoding bit rate, represented in bps.

The feature collects the information above per video object, in which each video object is defined by a unique URI. When multiple HTTP flows can be used to obtain one video object, as with Apple iOS® devices, the feature combines these flows when collecting statistics and treats them as one video object. The statistics are then aggregated per ACS manager and at the Global system level. This aggregation occurs using the following operations:

- For the first five statistics described above, when each video object terminates, the numbers are added to the aggregator at the ACS manager level. Aggregation among ACS managers happens when triggered by CLI commands or when bulk statistics are generated.
- The three derived statistics are calculated using the first five statistics after aggregation at the ACS manager level and the Global system level.

During aggregation, the mobile video statistics feature categorizes the information above based on UE device type, radio access type, and video container type, as follows:

- **UE device type:** Apple iOS devices (iPhone, iPad®, and iPod®), Android™ devices, laptops, and other devices.
- **Radio access type:** 2G, 3G, 4G-LTE, CDMA, HSPA, WLAN, and other types.
- **Video container type:** flv/f4v, mp4 (includes related types such as m4v, 3gp, 3g2, and mov), and other types.

The statistics include the total video object count for each of these categories, which is the total number of video files downloaded for a particular category.

The feature maintains two statistical arrays. The first array is arranged per UE device type, per radio access type. The second array is arranged per UE device type, per video container type.

For configuration instructions, see Chapter 2. For information about the variables in the MVS schema, see the *Statistics and Counters Reference*.

Bulk Statistics for Mobile Video

Bulk statistics on the ASR 5000 allow operators to choose to view not only statistics that are of importance to them, but to also configure the format in which they are presented. This simplifies the post-processing of statistical data, since it can be formatted to be parsed by external, back-end processors.

The system can be configured to collect bulk statistics and send them to a collection server called a receiver. Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. The following is a partial list of supported schemas:

- **System:** Provides system-level statistics.

- **Card:** Provides card-level statistics.
- **Port:** Provides port-level statistics.
- **MVS:** Provides statistics to support the Mobile Videoscape (MVS).

The system supports the configuration of up to four sets (primary/secondary) of receivers. Each set can be configured to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, system host name, system uptime, the IP address of the system generating the statistics (available for headers and footers only), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics Server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

For configuration instructions, see Chapter 2.

How the Mobile Video Gateway Works

This section shows how the Mobile Video Gateway works during DPI in a number of call scenarios, including scenarios involving the Mobile Video Gateway with the CAE and the Mobile Video Gateway without the CAE.

Mobile Video Gateway with the Content Adaptation Engine

This section shows call scenarios involving the Mobile Video Gateway with the Content Adaptation Engine.

DPI of HTTP GET Request Identifying a Non-Video Request (MVG with the CAE)

When the Mobile Video Gateway receives an HTTP GET request from a subscriber UE, it performs DPI to determine whether it is a request for video content. The figure below shows the Mobile Video Gateway with the CAE performing DPI on an HTTP GET request and identifying it as a non-video request. The table that follows the figure describes each step in the message flow.

Figure 163. DPI of HTTP GET Request Identifying a Non-Video Request

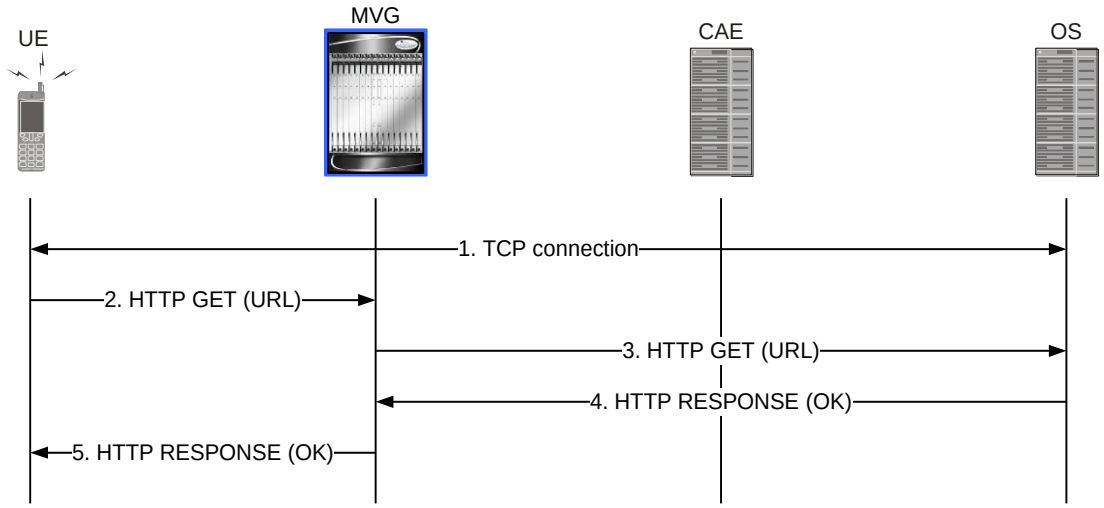


Table 71. DPI of HTTP GET Request Identifying a Non-Video Request

Step	Description
1.	The UE creates a TCP connection with the OS.
2.	The Mobile Video Gateway receives an HTPP GET request from the UE. The Mobile Video Gateway performs DPI and identifies it as a non-video request (the DPI on GET/POST fails).
3.	The Mobile Video Gateway forwards the HTTP request to the OS transparently, using the URL source IP address as the client address and the destination IP address as the OS address.
4.	The OS responds with an HTTP 200 OK, including the content of the page.

Step	Description
5.	The Mobile Video Gateway forwards the HTTP 200 OK to the UE transparently. The connection is not proxied, and the TCP flow continues to the UE.

DPI of HTTP RESPONSE Identifying a Non-Video Request (MVG with the CAE)

When the Mobile Video Gateway cannot determine whether an HTTP GET request is a request for video content during DPI, it performs DPI again when it receives the HTTP RESPONSE from the OS. The figure below shows the Mobile Video Gateway with the CAE performing DPI on an HTTP RESPONSE and identifying it as a response to a non-video request. The table that follows the figure describes each step in the message flow.

Figure 164. DPI of HTTP RESPONSE Identifying a Non-Video Request

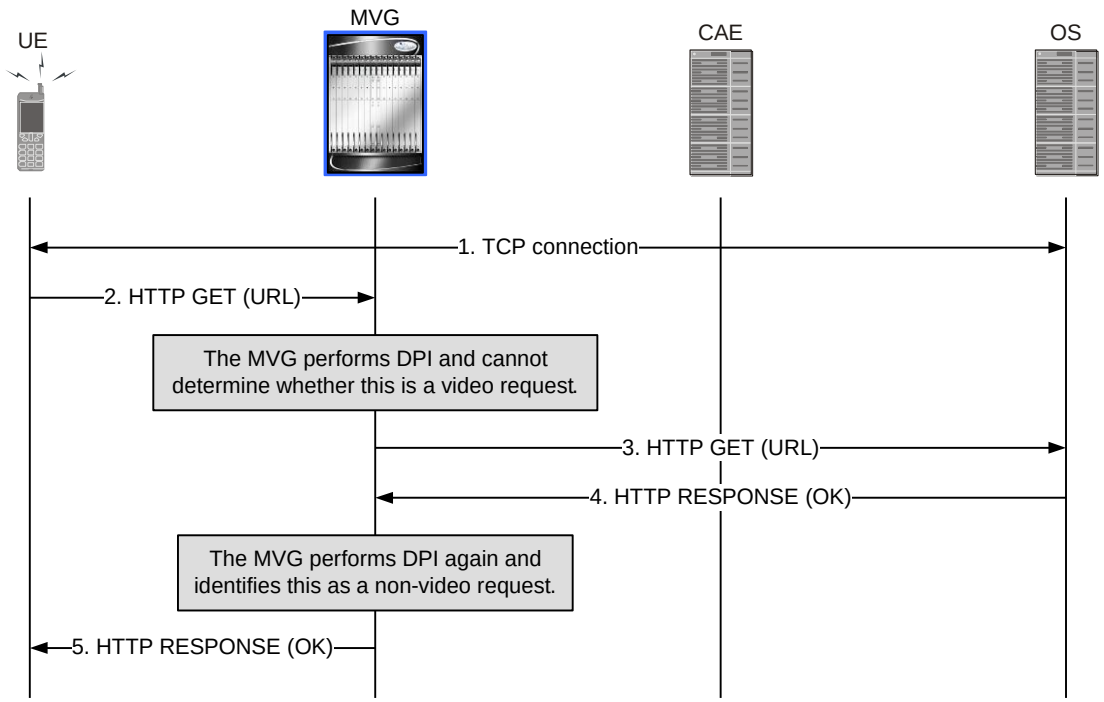


Table 72. DPI of HTTP RESPONSE Identifying a Non-Video Request

Step	Description
1.	The UE creates a TCP connection with the OS.
2.	The Mobile Video Gateway receives an HTPP GET request from the UE. The Mobile Video Gateway performs DPI and cannot determine whether it is a video request.
3.	The Mobile Video Gateway forwards the HTTP request to the OS transparently, using the URL source IP address as the client address and the destination IP address as the OS address.

Step	Description
4.	The OS responds with an HTTP 200 OK, including the content of the page. The Mobile Video Gateway performs DPI again and identifies it as a response to a non-video request (the DPI on the RESPONSE headers fails).
5.	The Mobile Video Gateway forwards the HTTP 200 OK to the UE transparently. The connection is not proxied, and the TCP flow continues to the UE.

DPI of HTTP GET Request Identifying a Video Request (MVG with the CAE)

When the Mobile Video Gateway receives an HTTP GET request from a UE, it performs DPI to determine whether it is a request for video content. The figure below shows the Mobile Video Gateway with the CAE performing DPI on an HTTP GET request and identifying it as a video request. The table that follows the figure describes each step in the message flow.

Figure 165. DPI of HTTP GET Request Identifying a Video Request

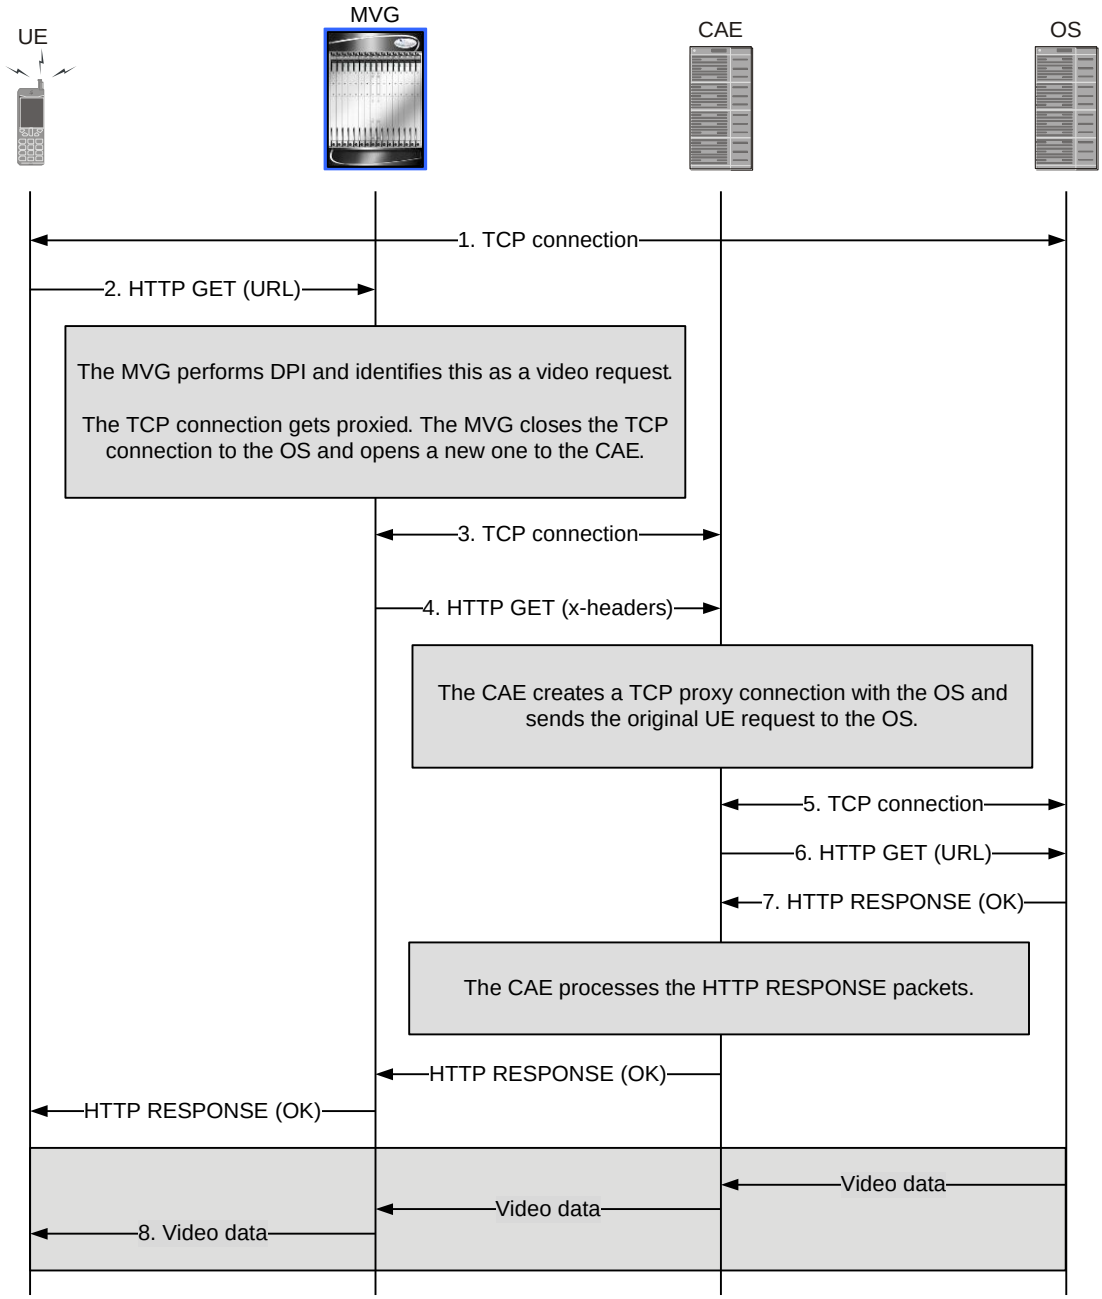


Table 73. DPI of HTTP GET Request Identifying a Video Request

Step	Description
1.	The UE creates a TCP connection with the OS.
2.	The Mobile Video Gateway receives an HTPP GET request from the UE. The Mobile Video Gateway performs DPI and identifies it as a video request (the DPI on GET/POST succeeds).

Step	Description
3.	The TCP connection gets proxied. The Mobile Video Gateway closes the TCP connection with the OS and opens a new one with the CAE.
4.	The Mobile Video Gateway sends the original HTTP GET request to the CAE with x-headers for transport, quality, and UE identity.
5.	The CAE creates a TCP proxy connection with the OS.
6.	The CAE sends the original HTTP GET request from the UE to the OS.
7.	The CAE processes the HTTP RESPONSE packets from the OS and performs video optimization.
8.	The Mobile Video Gateway performs additional video optimization and sends the optimized packets to the UE.

DPI of HTTP RESPONSE Identifying a Video Request (MVG with the CAE)

When the Mobile Video Gateway cannot determine whether an HTTP GET request is a request for video content during DPI, it performs DPI again when it receives the HTTP RESPONSE from the OS. The figure below shows the Mobile Video Gateway with the CAE performing DPI on an HTTP RESPONSE and identifying it as a response to a video request. The table that follows the figure describes each step in the message flow.

Figure 166. DPI of HTTP RESPONSE Identifying a Video Request

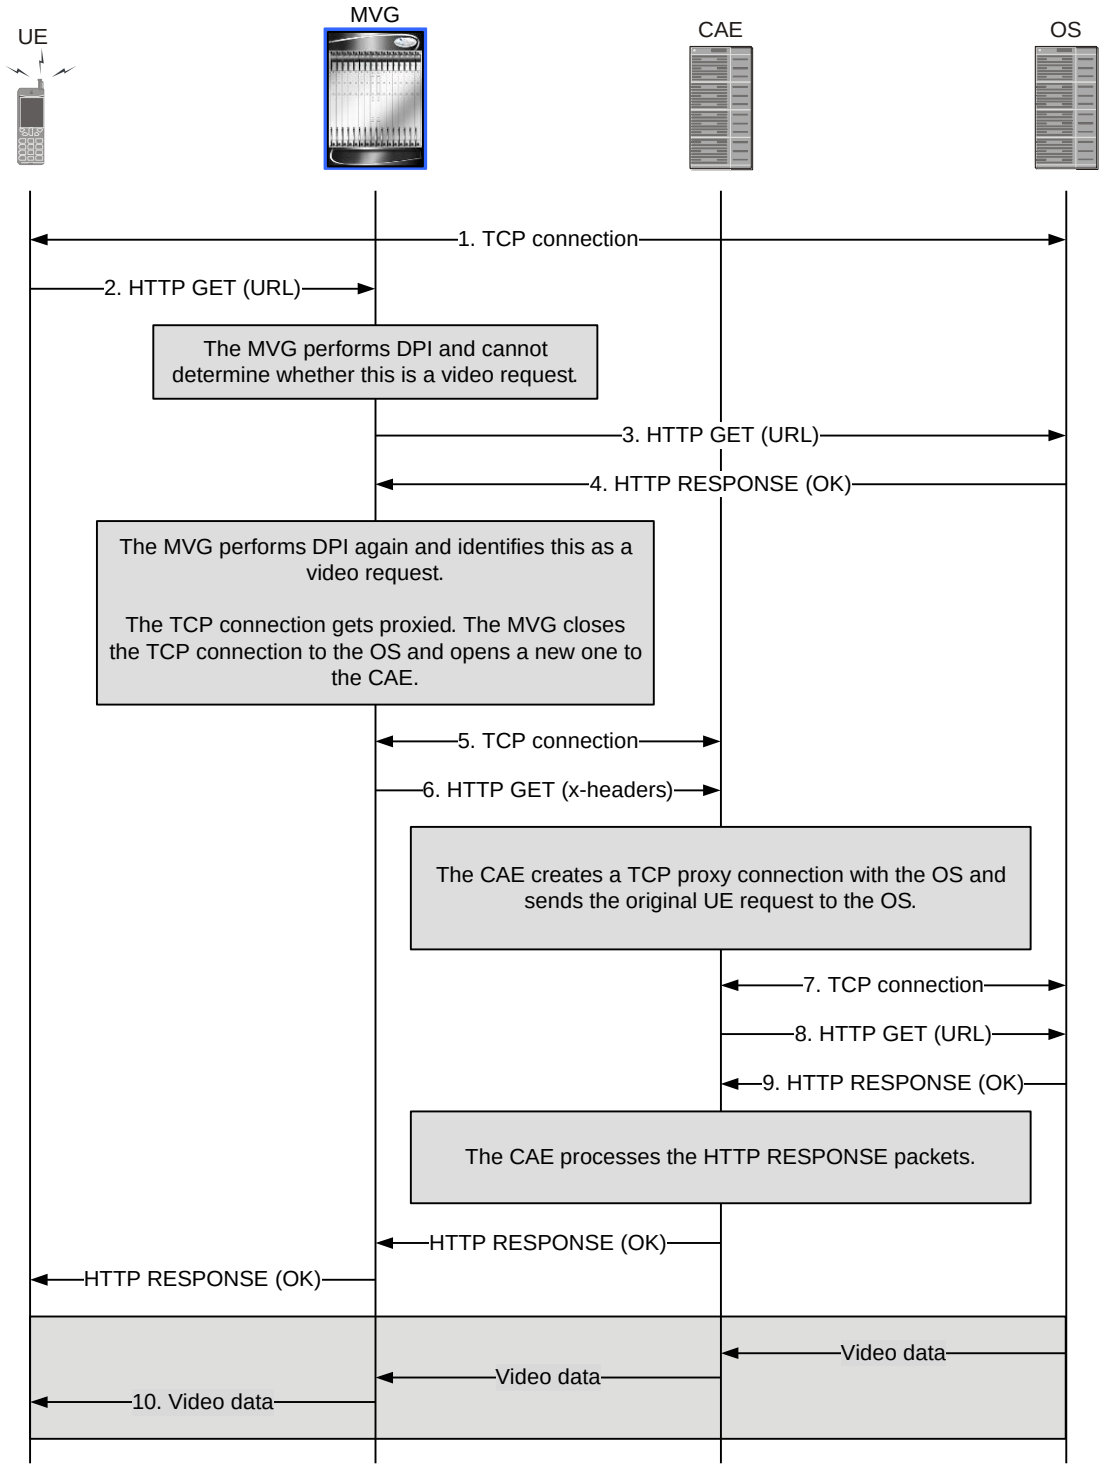


Table 74. DPI of HTTP RESPONSE Identifying a Video Request

Step	Description
1.	The UE creates a TCP connection with the OS.
2.	The Mobile Video Gateway receives an HTPP GET request from the UE. The Mobile Video Gateway performs DPI and cannot determine whether it is a video request.
3.	The Mobile Video Gateway forwards the HTTP request to the OS transparently, using the URL source IP address as the client address and the destination IP address as the OS address.
4.	The OS responds with an HTTP 200 OK. The Mobile Video Gateway performs DPI again and identifies it as a response to a video request (the DPI on the RESPONSE headers succeeds).
5.	The TCP connection gets proxied. The Mobile Video Gateway closes the TCP connection with the OS and opens a new one with the CAE.
6.	The Mobile Video Gateway sends the original HTTP GET request to the CAE with x-headers for transport, quality, and UE identity.
7.	The CAE creates a TCP proxy connection with the OS.
8.	The CAE sends the original HTTP GET request from the UE to the OS.
9.	The CAE processes the HTTP RESPONSE packets from the OS and performs video optimization.
10.	The Mobile Video Gateway performs additional video optimization and sends the optimized packets to the UE.

Mobile Video Gateway without the Content Adaptation Engine

This section shows call scenarios involving a Mobile Video Gateway without the Content Adaptation Engine.

DPI of HTTP GET Request Identifying a Non-Video Request (MVG without the CAE)

When the Mobile Video Gateway receives an HTTP GET request from a UE, it performs DPI to determine whether it is a request for video content. The figure below shows the Mobile Video Gateway performing DPI on an HTTP GET request and identifying it as a non-video request. The table that follows the figure describes each step in the message flow.

Figure 167. DPI of HTTP GET Request Identifying a Non-Video Request

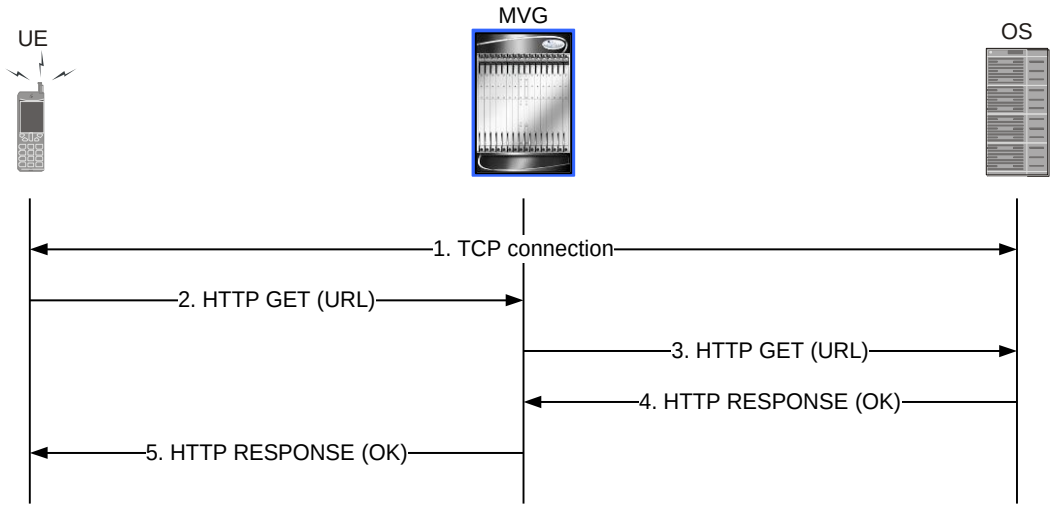


Table 75. DPI of HTTP GET Request Identifying a Non-Video Request

Step	Description
1.	The UE creates a TCP connection with the OS.
2.	The Mobile Video Gateway receives an HTPP GET request from the UE. The Mobile Video Gateway performs DPI and identifies it as a non-video request (the DPI on GET/POST fails).
3.	The Mobile Video Gateway forwards the HTTP request to the OS transparently, using the URL source IP address as the client address and the destination IP address as the OS address.
4.	The OS responds with an HTTP 200 OK, including the content of the page.
5.	The Mobile Video Gateway forwards the HTTP 200 OK to the UE transparently. The connection is not proxied, and the TCP flow continues to the UE.

DPI of HTTP RESPONSE Identifying a Non-Video Request (MVG without the CAE)

When the Mobile Video Gateway cannot determine whether an HTTP GET request is a request for video content during DPI, it performs DPI again when it receives the HTTP RESPONSE from the OS. The figure below shows the Mobile Video Gateway performing DPI on an HTTP RESPONSE and identifying it as a response to a non-video request. The table that follows the figure describes each step in the message flow.

Figure 168. DPI of HTTP RESPONSE Identifying a Non-Video Request

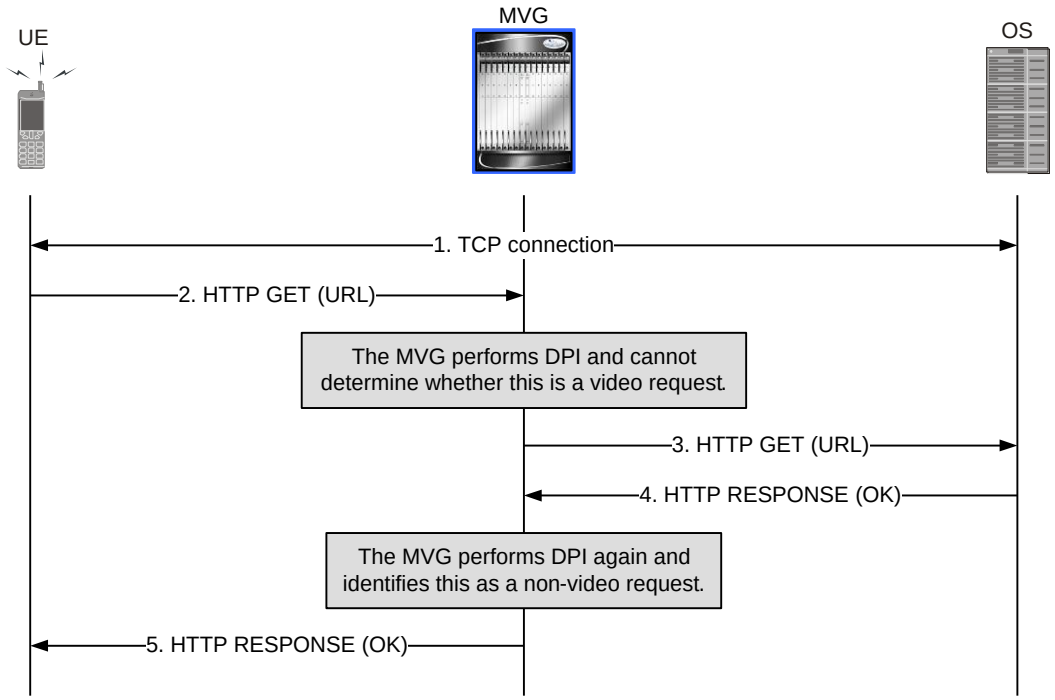


Table 76. DPI of HTTP RESPONSE Identifying a Non-Video Request

Step	Description
1.	The UE creates a TCP connection with the OS.
2.	The Mobile Video Gateway receives an HTPP GET request from the UE. The Mobile Video Gateway performs DPI and cannot determine whether it is a video request.
3.	The Mobile Video Gateway forwards the HTTP request to the OS transparently, using the URL source IP address as the client address and the destination IP address as the OS address.
4.	The OS responds with an HTTP 200 OK, including the content of the page. The Mobile Video Gateway performs DPI again and identifies it as a response to a non-video request (the DPI on the RESPONSE headers fails).
5.	The Mobile Video Gateway forwards the HTTP 200 OK to the UE transparently. The connection is not proxied, and the TCP flow continues to the UE.

DPI of HTTP GET Request Identifying a Video Request (MVG without the CAE)

When the Mobile Video Gateway receives an HTTP GET request from a UE, it performs DPI to determine whether it is a request for video content. The figure below shows the Mobile Video Gateway performing DPI on an HTTP GET request and identifying it as a video request. The table that follows the figure describes each step in the message flow.

Figure 169. DPI of HTTP GET Request Identifying a Video Request

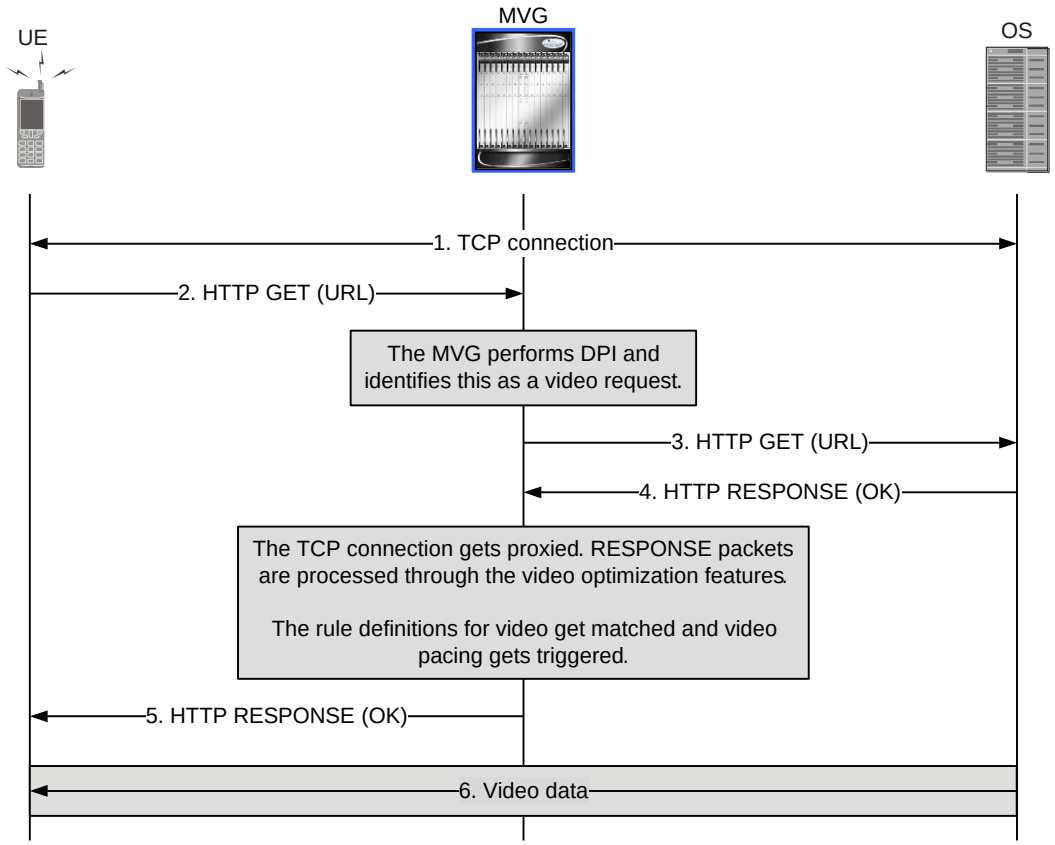


Table 77. DPI of HTTP GET Request Identifying a Video Request

Step	Description
1.	The UE creates a TCP connection with the OS.
2.	The Mobile Video Gateway receives an HTPP GET request from the UE. The Mobile Video Gateway performs DPI and identifies it as a video request (the DPI on GET/POST succeeds).
3.	The Mobile Video Gateway forwards the HTTP request to the OS transparently, using the URL source IP address as the client address and the destination IP address as the OS address.
4.	The OS responds with an HTTP 200 OK. The Mobile Video Gateway proxies the TCP connection and the HTTP RESPONSE packets are processed through the video optimization features.
5.	The Mobile Video Gateway forwards the HTTP 200 OK to the UE.
6.	The optimized TCP video flow continues to the UE.

DPI of HTTP RESPONSE Identifying a Video Request (MVG without the CAE)

When the Mobile Video Gateway cannot determine whether an HTTP GET request is a request for video content during DPI, it performs DPI again when it receives the HTTP RESPONSE from the OS. The figure below shows the Mobile Video Gateway performing DPI on an HTTP RESPONSE and identifying it as a response to a video request. The table that follows the figure describes each step in the message flow.

Figure 170. DPI of HTTP RESPONSE Identifying a Video Request

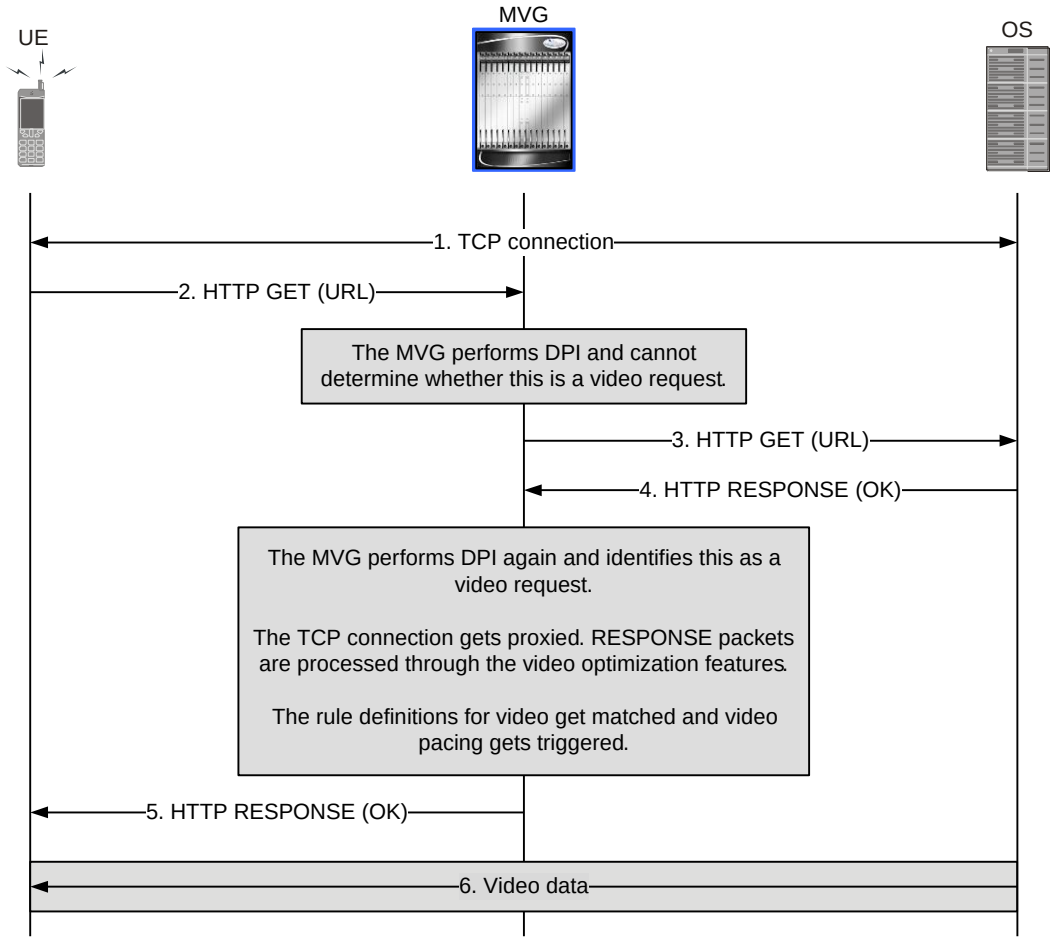


Table 78. DPI of HTTP RESPONSE Identifying a Video Request

Step	Description
1.	The UE creates a TCP connection with the OS.
2.	The Mobile Video Gateway receives an HTPP GET request from the UE. The Mobile Video Gateway performs DPI and cannot determine whether it is a video request.
3.	The Mobile Video Gateway forwards the HTTP request to the OS transparently, using the URL source IP address as the client address and the destination IP address as the OS address.

Step	Description
4.	The OS responds with an HTTP 200 OK. The Mobile Video Gateway performs DPI again and identifies this as a response to a video request (the DPI on the RESPONSE headers succeeds). The Mobile Video Gateway proxies the TCP connection and the HTTP RESPONSE packets are processed through the video optimization features.
5.	The Mobile Video Gateway forwards the HTTP 200 OK to the UE.
6.	The optimized TCP video flow continues to the UE.

Chapter 23
Mobility Management Entity Overview

The Cisco® ASR 5000 chassis provides Long Term Evolution (LTE)/System Architecture Evolution (SAE) wireless carriers with a flexible solution that functions as a Mobility Management Entity (MME) in 3rd Generation Partnership Project (3GPP) LTE/SAE wireless data networks.

This overview provides general information about the MME including:

- Product Description
- Network Deployment and Interfaces
- Features and Functionality - Base Software
- Features and Functionality - External Application Support
- Features and Functionality - Licensed Enhanced Feature Software
- How the MME Works
- Supported Standards

Product Description

This section describes the MME network function and its position in the LTE network.

The MME is the key control-node for the LTE access network. It works in conjunction with the evolved NodeB (eNodeB), Serving Gateway (S-GW) within the Evolved Packet Core (EPC), or LTE/SAE core network to perform the following functions:

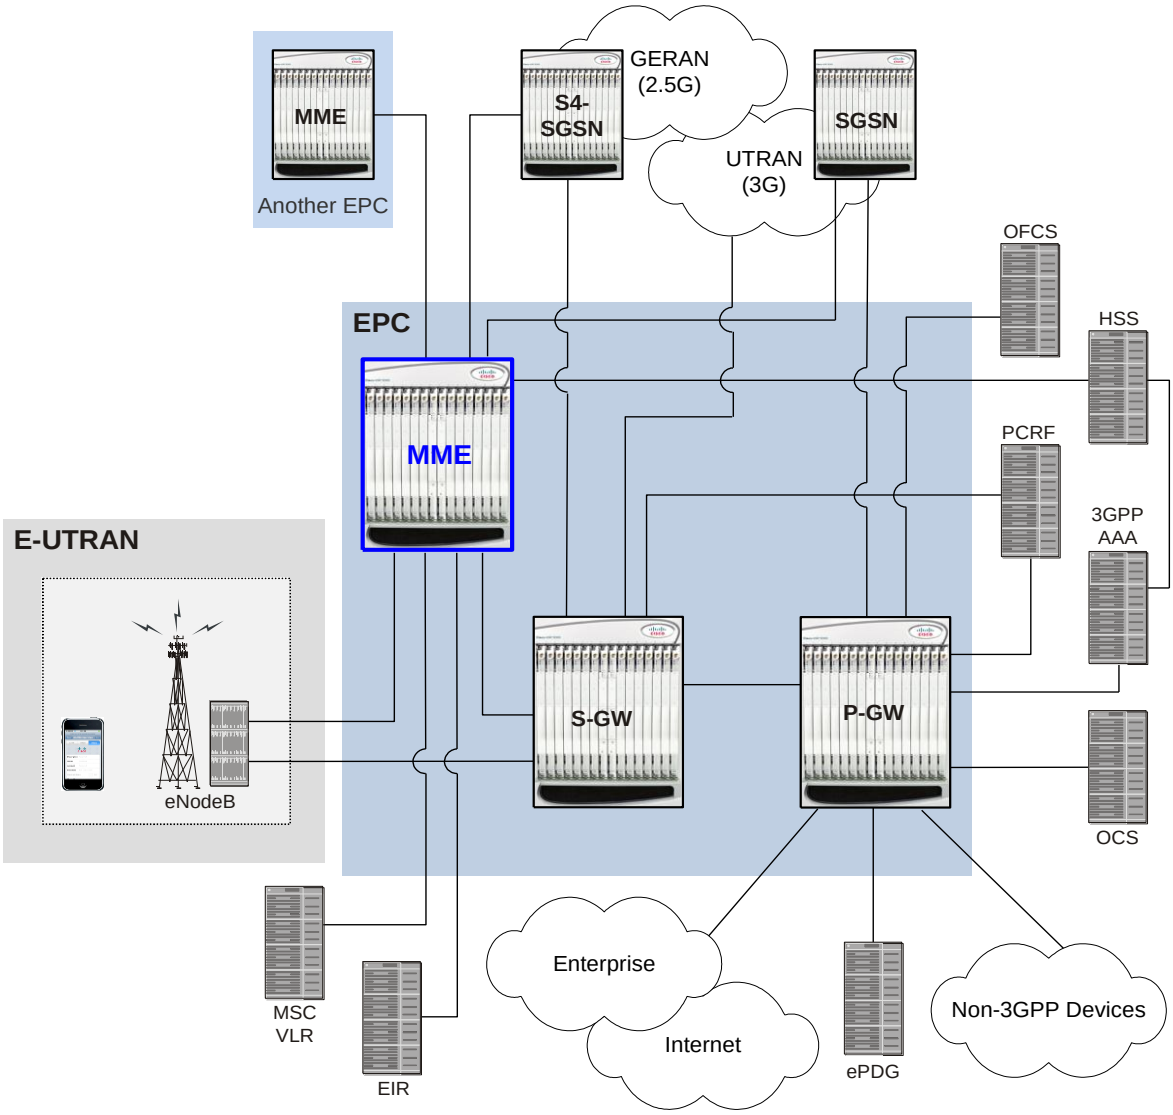
- Involved in the bearer activation/deactivation process and is also responsible for choosing the S-GW and for a UE at the initial attach and at the time of intra-LTE handover involving Core Network (CN) node relocation.
- Provides P-GW selection for subscriber to connect to PDN.
- Provides idle mode UE tracking and paging procedure, including retransmissions.
- Chooses the appropriate S-GW for a UE.
- Responsible for authenticating the user (by interacting with the HSS).
- Works as termination point for Non-Access Stratum (NAS) signaling.
- Responsible for generation and allocation of temporary identities to UEs.
- Checks the authorization of the UE to camp on the service provider's Public Land Mobile Network (PLMN) and enforces UE roaming restrictions.
- The MME is the termination point in the network for ciphering/integrity protection for NAS signaling and handles the security key management.
- Communicates with MMEs in same PLMN or on different PLMNs. The S10 interface is used for MME relocation and MME-to-MME information transfer or handoff.

Besides the above mentioned functions, the lawful interception of signaling is also supported by the MME.

The MME also provides the control plane function for mobility between LTE and 2G/3G access networks with the S3 interface terminating at the MME from the SGSN. In addition, the MME interfaces with SGSN for interconnecting to the legacy network.

The MME also terminates the S6a interface towards the home HSS for roaming UEs.

Figure 171. MME in the E-UTRAN/EPC Network Topology



In accordance with 3GPP standard, the MME provides following functions and procedures in the LTE/SAE network:

- Non Access Stratum (NAS) signalling
- NAS signalling security
- Inter CN node signalling for mobility between 3GPP access networks (terminating S3)
- UE Reachability in ECM-IDLE state (including control and execution of paging retransmission)
- Tracking Area list management
- PDN GW and Serving GW selection
- MME selection for handover with MME change
- SGSN selection for handover to 2G or 3G 3GPP access networks
- Roaming (S6a towards home HSS)

- Authentication
- Bearer management functions including dedicated bearer establishment
- Lawful Interception of signalling traffic
- UE Reachability procedures
- Interfaces with MSC for Voice paging
- Interfaces with SGSN for interconnecting to legacy network

Platform Requirements

The MME service runs on a Cisco® ASR 5000 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

Licenses

The MME is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Network Deployment and Interfaces

This section describes the supported interfaces and deployment scenario of MME in LTE/SAE network.

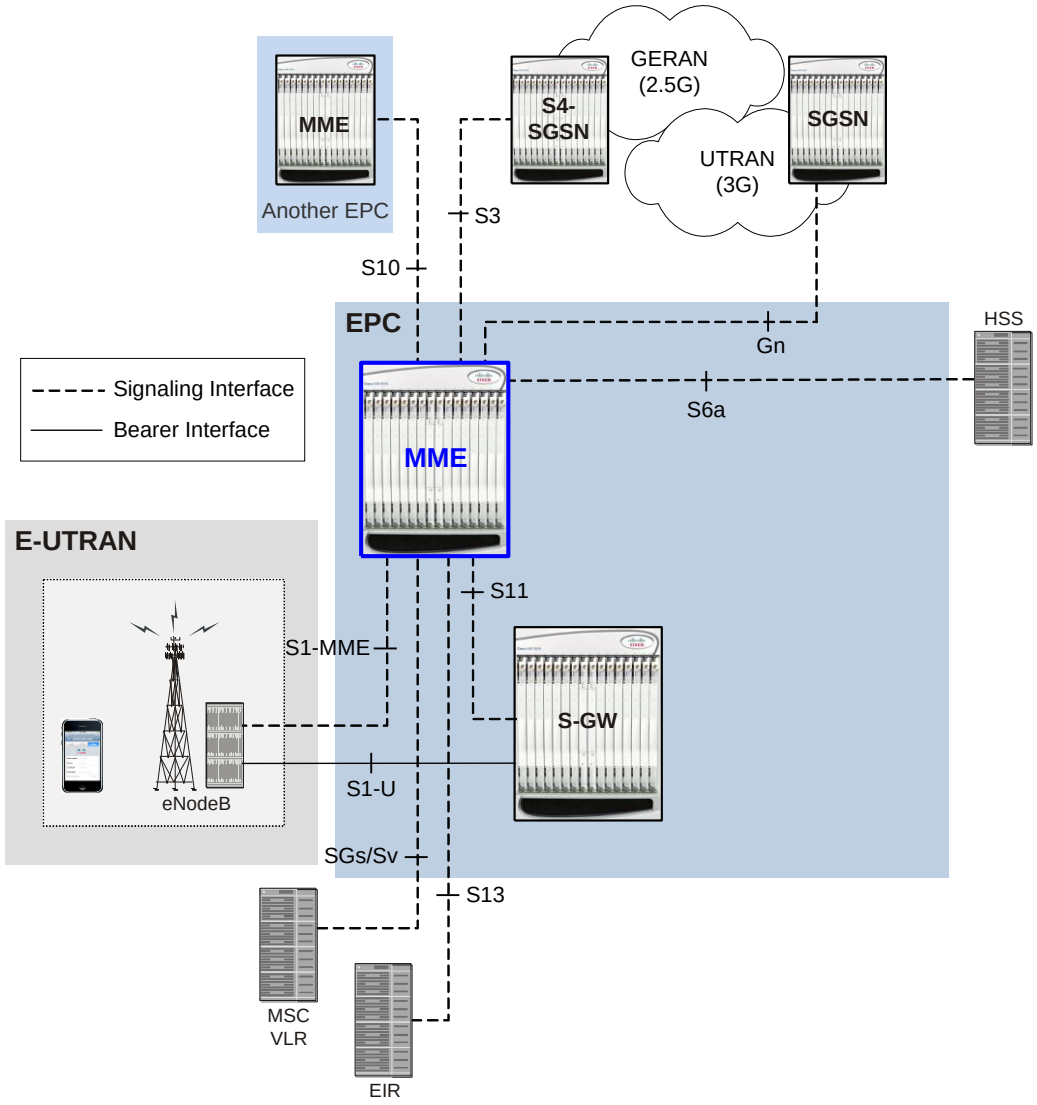
The following information is provided in this section:

- MME in the E-UTRAN/EPC Network
- Supported Logical Network Interfaces (Reference Points)

MME in the E-UTRAN/EPC Network

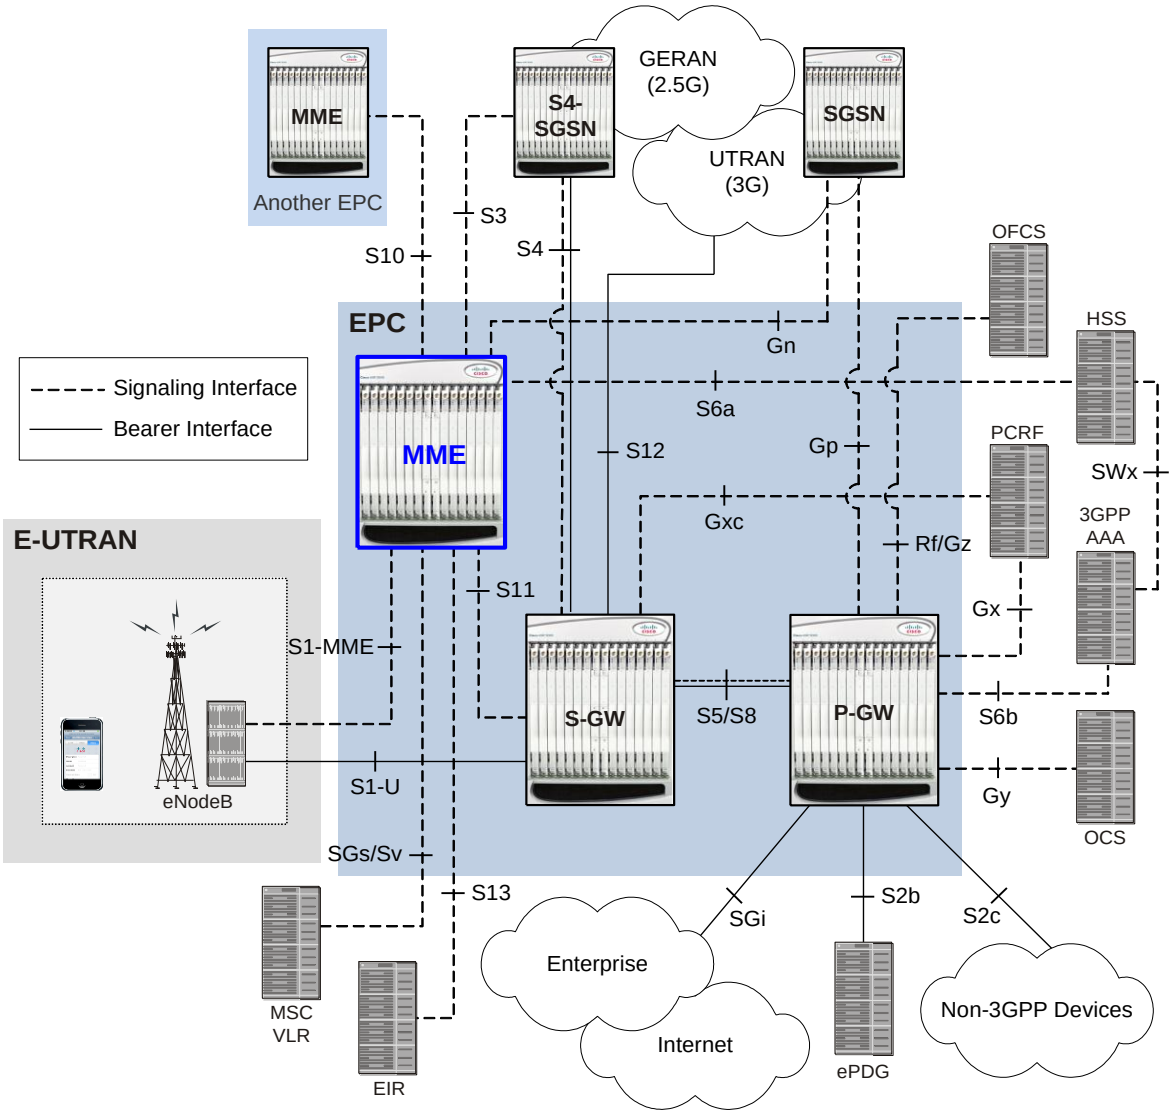
The following figure displays the specific network interfaces supported by the MME. Refer to Supported Logical Network Interfaces (Reference Points) for detailed information about each interface.

Figure 172. Supported MME Interfaces in the E-UTRAN/EPC Network



The following figure displays a sample network deployment of an MME, including all of the interface connections with other 3GPP Evolved-UTRAN/Evolved Packet Core network devices.

Figure 173. E-UTRAN/EPC Network Scenario



Supported Logical Network Interfaces (Reference Points)

The MME supports the following logical network interfaces/reference points:

S1-MME Interface

This interface is the reference point for the control plane protocol between eNodeB and MME. S1-MME uses the S1 Application Protocol (S1-AP) over the Stream Control Transmission Protocol (SCTP) as the transport layer protocol for guaranteed delivery of signaling messages between MME and eNodeB (S1).

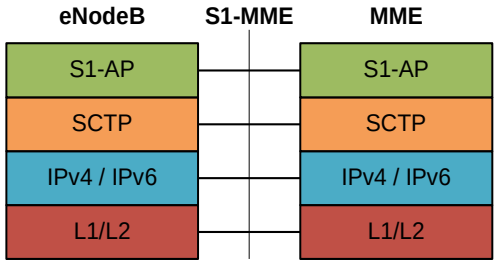
This is the interface used by the MME to communicate with eNodeBs on the same LTE Public Land Mobile Network (PLMN). This interface serves as path for establishing and maintaining subscriber UE contexts.

The S1-MME interface supports IPv4, IPv6, IPSec, and multi-homing.

One or more S1-MME interfaces can be configured per system context.

Supported protocols:

- Application Layer: S1 Application Protocol (S1-AP)
- Transport Layer: SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



S3 Interface

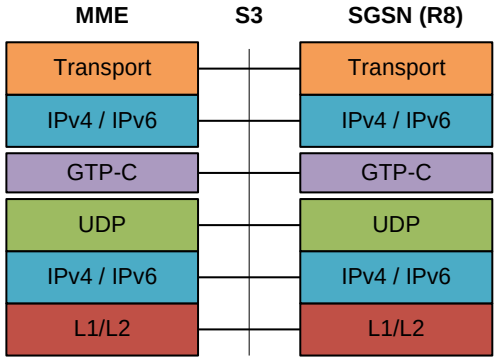
This is the interface used by the MME to communicate with S4-SGSNs on the same Public PLMN for interworking between GPRS/UMTS and LTE network access technologies. This interface serves as the signalling path for establishing and maintaining subscriber UE contexts.

The MME communicates with SGSNs on the PLMN using the GPRS Tunnelling Protocol (GTP). The signalling or control aspect of this protocol is referred to as the GTP Control Plane (GTPC) while the encapsulated user data traffic is referred to as the GTP User Plane (GTPU).

One or more S3 interfaces can be configured per system context.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: IPv4 or IPv6 GTPv2-C (signaling channel)
- Signalling Layer: UDP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



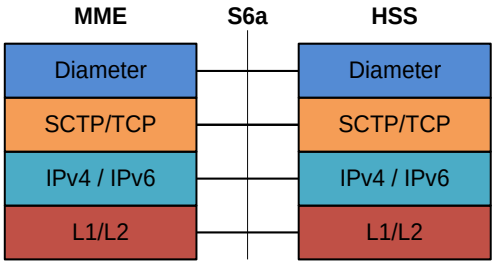
S6a Interface

This is the interface used by the MME to communicate with the Home Subscriber Server (HSS). The HSS is responsible for transfer of subscription and authentication data for authenticating/authorizing user access and UE context authentication. The MME communicates with the HSSs on the PLMN using Diameter protocol.

One or more S6a interfaces can be configured per system context.

Supported protocols:

- Transport Layer: SCTP or TCP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



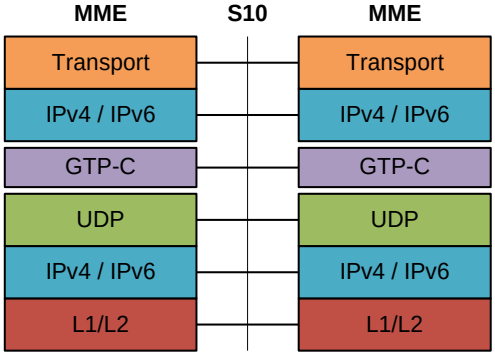
S10 Interface

This is the interface used by the MME to communicate with an MME in the same PLMN or on different PLMNs. This interface is also used for MME relocation and MME-to-MME information transfer or handoff. This interface uses the GTPv2 protocol.

One or more S10 interfaces can be configured per system context.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: IPv4 or IPv6 GTPv2-C (signaling channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



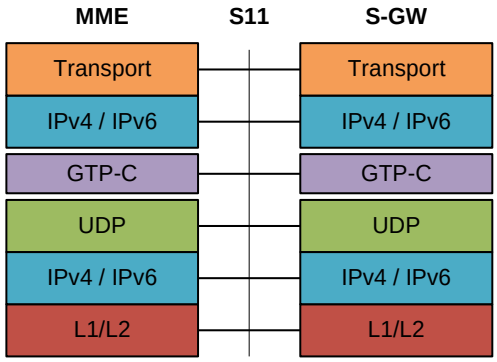
S11 Interface

This interface provides communication between the MME and Serving Gateways (S-GW) for information transfer. This interface uses the GTPv2 protocol.

One or more S11 interfaces can be configured per system context.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: IPv4 or IPv6 GTPv2-C (signaling channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



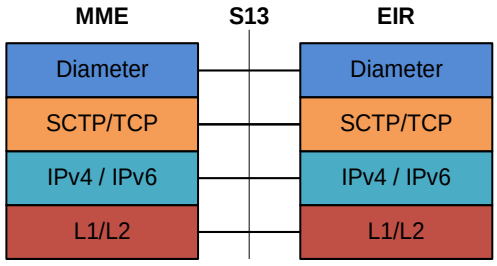
S13 Interface

This interface provides communication between MME and Equipment Identity Register (EIR).

One or more S13 interfaces can be configured per system context.

Supported protocols:

- Transport Layer: SCTP or TCP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

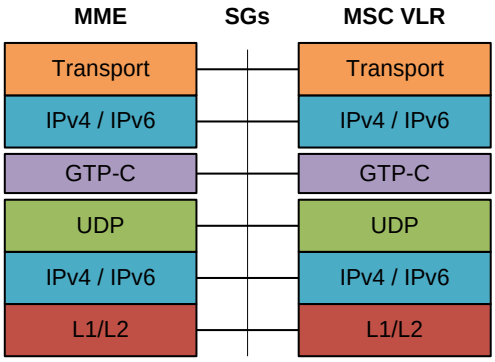


SGs Interface

The SGs interface connects the databases in the VLR and the MME to support circuit switch fallback scenarios.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: IPv4 or IPv6 GTP-C (signaling channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

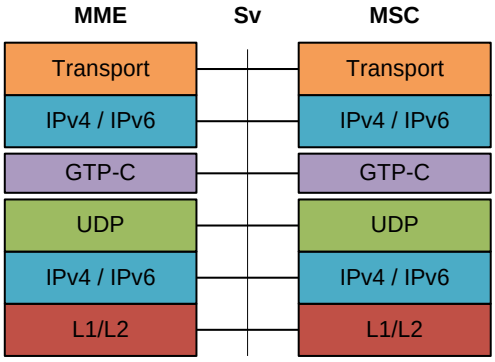


Sv Interface

This interface connects the MME to a Mobile Switching Center to support the exchange of messages during a handover procedure for the Single Radio Voice Call Continuity (SRVCC) feature.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: IPv4 or IPv6 GTP-C (signaling channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



Gn Interface

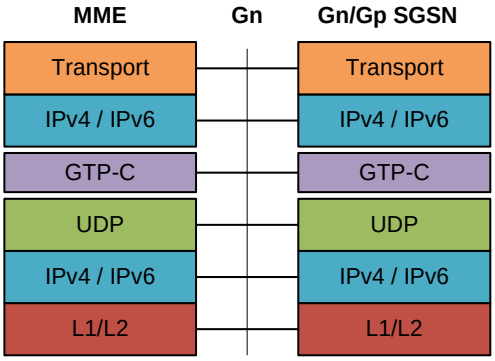
Gn interfaces facilitate user mobility between 2G/3G 3GPP networks. The Gn interface is used for intra-PLMN handovers. The MME supports pre-Release-8 Gn interfaces to allow inter-operation between EPS networks and 2G/3G 3GPP networks.

Roaming and inter access mobility between 2G and/or 3G SGSNs and an MME/S-GW are enabled by:

- Gn functionality, as specified between two SGSNs, which is provided by the MME, and
- Gp functionality, as specified between SGSN and GGSN, that is provided by the P-GW.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: IPv4 or IPv6 GTP-C (signaling channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



Features and Functionality - Base Software

This section describes the features and functions supported by default in the base software on the MME service and do not require any additional licenses.

Important: To configure the basic service and functionality on the system for MME service, refer configuration examples provide in *MME Administration Guide*.

This section describes following features:

- 3GPP R8 Identity Support
- ANSI T1.276 Compliance
- APN Restriction Support
- Authentication and Key Agreement (AKA)
- Bulk Statistics Support
- Congestion Control
- EPS Bearer Context Support
- EPS GTPv2 Support on S11 Interface
- HSS Support Over S6a Interface
- Inter-MME Handover Support
- Interworking Support
- IPv6 Support
- Load Balancing
- Management System Overview
- MME Pooling
- MME Selection
- Mobile Equipment Identity Check
- Mobility Restriction
- Multiple PDN Support
- NAS Protocol Support
- NAS Signalling Security
- Operator Policy Support
- Overload Management in MME
- Packet Data Network Gateway (P-GW) Selection
- Radio Resource Management Functions
- Reachability Management

- SCTP Multi-homing Support
- Serving Gateway Pooling Support
- Serving Gateway Selection
- Subscriber Level Session Trace
- Threshold Crossing Alerts (TCA) Support
- Tracking Area List Management

3GPP R8 Identity Support

Provides the identity allocation of following type:

- EPS Bearer Identity
- Globally Unique Temporary UE Identity (GUTI)
- Tracking Area Identity (TAI)
- MME S1-AP UE Identity (MME S1-AP UE ID)
- **EPS Bearer Identity**: An EPS bearer identity uniquely identifies EPS bearers within a user session for attachment to the E-UTRAN access and EPC core networks. The EPS Bearer Identity is allocated by the MME. There is a one to one mapping between EPS Radio Bearers via the E-UTRAN radio access network and EPS Bearers via the S1-MME interface between the eNodeB and MME. There is also a one-to-one mapping between EPS Radio Bearer Identity via the S1 and X2 interfaces and the EPS Bearer Identity assigned by the MME.
- **Globally Unique Temporary UE Identity (GUTI)**: The MME allocates a Globally Unique Temporary Identity (GUTI) to the UE. A GUTI has; 1) unique identity for MME which allocated the GUTI; and 2) the unique identity of the UE within the MME that allocated the GUTI.

Within the MME, the mobile is identified by the M-TMSI.

The Globally Unique MME Identifier (GUMMEI) is constructed from MCC, MNC and MME Identifier (MMEI). In turn the MMEI is constructed from an MME Group ID (MMEGI) and an MME Code (MMEC).

The GUTI is constructed from the GUMMEI and the M-TMSI.

For paging, the mobile is paged with the S-TMSI. The S-TMSI is constructed from the MMEC and the M-TMSI.

The operator needs to ensure that the MMEC is unique within the MME pool area and, if overlapping pool areas are in use, unique within the area of overlapping MME pools.

The GUTI is used to support subscriber identity confidentiality, and, in the shortened S-TMSI form, to enable more efficient radio signaling procedures (e.g. paging and Service Request).

- **Tracking Area Identity (TAI)**: Provides the function to assign the TAI list to the mobile access device to limit the frequency of Tracking Area Updates in the network. The TAI is the identity used to identify the tracking area or group of cells in which the idle mode access terminal will be paged when a remote host attempts to reach that user. The TAI consists of the Mobile Country Code (MCC), Mobile Network Code (MNC) and Tracking Area Code (TAC).
- **MME S1-AP UE Identity (MME S1-AP UE ID)**: This is the temporary identity used to identify a UE on the S1-MME reference point within the MME. It is unique within the MME per S1-MME reference point instance.

ANSI T1.276 Compliance

ANSI T1.276 specifies security measures for Network Elements (NE). In particular it specifies guidelines for password strength, storage, and maintenance security measures.

ANSI T1.276 specifies several measures for password security. These measures include:

- Password strength guidelines
- Password storage guidelines for network elements
- Password maintenance, e.g. periodic forced password changes

These measures are applicable to the system and the Web Element Manager since both require password authentication. A subset of these guidelines where applicable to each platform will be implemented. A known subset of guidelines, such as certificate authentication, are not applicable to either product. Furthermore, the platforms support a variety of authentication methods such as RADIUS and SSH which are dependent on external elements. ANSI T1.276 compliance in such cases will be the domain of the external element. ANSI T1.276 guidelines will only be implemented for locally configured operators.

APN Restriction Support

The APN-Restriction value may be configured for each APN in the P-GW and transferred to the MME. It is used to determine, on a per-MS basis, whether it is allowed to establish EPS bearers to other APNs.

The APN-Restriction value is defined in clause 15.4 of 3GPP TS 23.060. APN-Restriction affects multiple procedures, such as Initial Attach, TAU, PDN connectivity, and inter-MME handovers. The MME saves the APN-Restriction value received in create session response for an APN and uses the maximum of the values from the currently active PDNs in the next create session request. If a PDN is disconnected, then the maximum APN-Restriction is adjusted accordingly.

Authentication and Key Agreement (AKA)

The MME provides EPS Authentication and Key Agreement mechanism for user authentication procedure over the E-UTRAN. The Authentication and Key Agreement (AKA) mechanism performs authentication and session key distribution in networks. AKA is a challenge-response based mechanism that uses symmetric cryptography. AKA is typically run in a Services Identity Module.

AKA is the procedure that take between the user and network to authenticate themselves towards each other and to provide other security features such as integrity and confidentiality protection.

In a logical order this follows the following procedure:

1. Authentication: Performs authentication by identifying the user to the network and identifying the network to the user.
2. Key agreement: Performs key agreement by generating the cipher key and generating the integrity key.
3. Protection: When the AKA procedure is performed, it protects the integrity of messages, the confidentiality of the signalling data, and the confidentiality of the user data.

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Following is a partial list of supported schemas:

- **System**: Provides system-level statistics
- **Card**: Provides card-level statistics
- **Port**: Provides port-level statistics
- **MME**: Provides MME service statistics
- **GTPC**: Provides GPRS Tunneling Protocol - Control message statistics

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the chassis or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, chassis host name, chassis uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Congestion Control

The congestion control feature allows you to set policies and thresholds and specify how the system reacts when faced with a heavy load condition.

Congestion control monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are quickly resolved. However, continuous or large numbers of these conditions within a specific time interval may have an impact the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establishes limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operation thresholds that are configured for the system as described in the Thresholding

Configuration Guide. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated.

A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.

- **Port Utilization Thresholds**: If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
  - **Port-specific Thresholds**: If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
- **Service Congestion Policies**: Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

Congestion control can be used in conjunction with the load balancing feature provided on the MME. For more information on MME load balancing, refer to the Load Balancing section in this chapter.

Important: For more information on congestion control, refer to the *Congestion Control* chapter in the *Cisco ASR 5000 Series System Administration Guide*.

EPS Bearer Context Support

Provides support for subscriber default and dedicated Evolved Packet System (EPS) bearer contexts in accordance with the following standards:

- **3GPP TS 36.412 V8.6.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Access Network (E-UTRAN); S1 signaling transport (Release 8)
- **3GPP TS 36.413 V8.8.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access Network (E-UTRAN); S1 Application Protocol (S1AP) (Release 8)
- IETF RFC 4960, Stream Control Transmission Protocol, December 2007

EPS bearer context processing is based on the APN that the subscriber is attempting to access. Templates for all of the possible APNs that subscribers will be accessing must be configured within the system. Up to 1024 APNs can be configured on the system.

Each APN template consists of parameters pertaining to how UE contexts are processed such as the following:

- PDN Type: IPv4, IPv6, or IPv4v6
- EPS Bearer Context timers
- Quality of Service

A total of 11 EPS bearer per subscriber are supported. These could be all dedicated, or 1 default and 10 dedicated or any combination of default and dedicated context. Note that there must be at least one default EPS Bearer context in order for dedicated context to come up.

EPS GTPv2 Support on S11 Interface

Support for the EPS GTPv2 on S11 interface in accordance with the following standards:

- **3GPP TS 29.274 V8.4.0 (2009-12)**: 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C); Stage 3 (Release 8)

The system supports the use of GTPv2 for EPS signalling context processing.

When the GTPv2 protocol is used, accounting messages are sent to the charging gateways (CGs) over the Ga interface. The Ga interface and GTPv2 functionality are typically configured within the system's source context. As specified by the standards, a CDR is not generated when a session starts. CDRs are generated according to the interim triggers configured using the charging characteristics configured for the MME, and a CDR is generated when the session ends. For interim accounting, STOP/START pairs are sent based on configured triggers.

GTP version 2 is always used. However, if version 2 is not supported by the CGF, the system reverts to using GTP version 1. All subsequent CDRs are always fully-qualified partial CDRs. All CDR fields are R4.

Whether or not the MME accepts charging characteristics from the SGSN can be configured on a per-APN basis based on whether the subscriber is visiting, roaming or, home.

By default, the MME always accepts the charging characteristics from the SGSN. They must always be provided by the SGSN for GTPv1 requests for primary EPS Bearer contexts. If they are not provided for secondary EPS Bearer contexts, the MME re-uses those from the primary.

If the system is configured to reject the charging characteristics from the SGSN, the MME can be configured with its own that can be applied based on the subscriber type (visiting, roaming, or home) at the APN level. MME charging characteristics consist of a profile index and behavior settings. The profile indexes specify the criteria for closing accounting records based specific criteria.

Important: For more information on GTPv2 configuration, refer to the *Creating and Configuring the eGTP Service and Interface Association* section in the *Mobility Management Entity Configuration* chapter of the *MME Service Administration Guide*.

HSS Support Over S6a Interface

Provides a mechanism for performing Diameter-based authorization, authentication, and accounting (AAA) for subscriber bearer contexts based on the following standards:

- **3GPP TS 23.401 V8.1.0 (2008-03)**: 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access (Release 8)
- **3GPP TS 29.272 V8.1.1 (2009-01)**: 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Evolved Packet System (EPS); Mobility Management Entity (MME) and Serving GPRS Support Node (SGSN) related interfaces based on Diameter protocol (Release 8)
- **3GPP TS 33.401 V8.2.1 (2008-12)**: 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3GPP System Architecture Evolution (SAE): Security Architecture; (Release 8)
- RFC 3588, Diameter Base Protocol, December 2003

The S6a protocol is used to provide AAA functionality for subscriber EPS Bearer contexts through Home Subscriber Server (HSS).

During the initial attachment procedures the MME sends to the USIM on AT via the HSS the random challenge (RAND) and an authentication token AUTN for network authentication from the selected authentication vector. At receipt of this message, the USIM verifies that the authentication token can be accepted and if so, produces a response. The AT and HSS in turn compute the Cipher Key (CK) and Integrity Key (IK) that are bound to Serving Network ID. During the attachment procedure the MME requests a permanent user identity via the S1-MME NAS signaling interface

to eNodeB and inserts the IMSI, Serving Network ID (MCC, MNC) and Serving Network ID it receives in an Authentication Data Request to the HSS. The HSS returns the Authentication Response with authentication vectors to MME. The MME uses the authentication vectors to compute the cipher keys for securing the NAS signaling traffic.

At EAP success, the MME also retrieves the subscription profile from the HSS which includes QoS information and other attributes such as default APN name and S-GW/P-GW fully qualified domain names.

Among the AAA parameters that can be configured are:

- Authentication of the subscriber with HSS
- Subscriber location update/location cancel
- Update subscriber profile from the HSS
- Priority to dictate the order in which the servers are used allowing for multiple servers to be configured in a single context
- Routing Algorithm to dictate the method for selecting among configured servers. The specified algorithm dictates how the system distributes AAA messages across the configured HSS servers for new sessions. Once a session is established and an HSS server has been selected, all subsequent AAA messages for the session will be delivered to the same server.

Inter-MME Handover Support

The S10 interface facilitates user mobility between two MMEs providing for the transfer of the UE context from one to the other. It is a GTPv2 control plane interface that supports the following handover types and features:

- E-UTRAN-to-UTRAN (MME-to-MME) handover through:
  - Tracking Area Update based inter-MME relocation
  - Attach at an eNodeB connected to a different MME
  - S1 handover based inter-MME relocation
- The MME supports handing over multiple bearers and multiple PDNs over to another MME
- Trace functionality, monitor protocol, and monitor subscriber
- DNS client configuration
- IPv4 and IPv6: for peer MME selection, the preference is given to IPv6 addresses. IPv4 addresses are ignored if IPv6 addresses are present.

Interworking Support

This section describes various interworking and handover scenarios supported by the MME. The following interworking types are provided:

- Interworking with SGSNs
- Handover Support for S4 SGSNs

Interworking with SGSNs

This feature enables an integrated EPC core network to anchor calls from multi-mode access terminals and supports seamless mobility on call hand-offs between an LTE or GERAN/UTRAN access network. This provides a valuable

function to enable LTE operators to generate incremental revenue from inbound roaming agreements with 2G/3G roaming partners.

In order to support inter-RAT hand-offs for dual-mode access terminals between LTE and 2G/3G networks with 3GPP Pre-Release 8 SGSN's, the MME will support combined hard handover and SRNS relocation procedures via the GTPv1 Gn/Gp reference interface. In preparation for the handover, the MME sends a Forward Relocation Request to the SGSN and includes subscriber identity and context information including IMSI, Mobility Management context and PDP context. The PDP context includes the GGSN address for the user plane and the uplink Tunnel Endpoint ID. These addresses are equivalent to the PDN GW address. The MME maps the EPS bearer parameters to the PDP contexts.

After sending the forward relocation signaling to the target SGSN, the MME deletes the EPS bearer resources by sending a Delete Bearer Request to the S-GW with a Cause code that instructs the S-GW not to initiate delete procedures toward the P-GW.

When a mobile subscriber roams from an EUTRAN to GERAN/UTRAN access network it must also send a Routing Area Update (RAU) to register its location with the target network. The target SGSN sends a Context Request to the MME with P-TMSI to get the Mobility Management contexts and PDP contexts for the subscriber session. The SGSN uses the Globally Unique Temporary ID (GUTI) from the MME to identify the P-TMSI/RAI.

Handover Support for S4-SGSNs

The S3 interface facilitates user mobility between an MME and an S4-SGSN providing for the transfer of the UE context between the two. It is a GTPv2 control plane interface that supports the following handover types:

- E-UTRAN-to-UTRAN (MME-to-R8 SGSN) handover through:
  - Routing Area Update (RAU) based MME-R8 SGSN relocation where the RAU could be a result of UE movement.
  - Attach at an RNC connected to a R8 SGSN
  - S1 handover/SRNS relocation based MME-R8 SGSN relocation
- UTRAN-to-E-UTRAN (R8 SGSN-to-MME) handover through:
  - Tracking Area Update (TAU) based R8 SGSN-MME relocation where the TAU could be a result of UE movement.
  - Attach at an eNodeB connected to an MME.
  - SRNS relocation/S1 handover based R8 SGSN-MME relocation.

All handover types support handing over multiple bearers and multiple PDNs from the MME to a R8 SGSN and vice versa.

The S3 interface also supports the following features:

- Monitor Protocol and Monitor Subscriber
- Subscriber Session Trace
- IPv4 and IPv6: for peer SGSN selection, the preference is given to IPv6 addresses. IPv4 addresses are ignored if IPv6 addresses are present.
- Operator Policy for SGSN selection
- Session Recovery: all MME sessions established using the S3 interface are capable of being recovered in case of a session manager task failure.

IPv6 Support

This feature allows IPv6 subscribers to connect via the LTE/SAE infrastructure in accordance with the following standards:

- RFC 2460: Internet Protocol, Version 6 (IPv6) Specification
- RFC 2461: Neighbor Discovery for IPv6
- RFC 2462: IPv6 Stateless Address Autoconfiguration
- RFC 3314: Recommendations for IPv6 in 3GPP Standards
- RFC 3316: Internet Protocol Version 6 (IPv6) for Some Second and Third Generation Cellular Hosts
- RFC 3056: Connection of IPv6 domains via IPv4 clouds
- 3GPP TS 27.060: Mobile Station Supporting Packet Switched Services
- 3GPP TS 29.061: Interworking between the Public Land Mobile Network (PLMN) supporting Packet Based Services and Packet Data Networks (PDN)

The MME allows an APN to be configured for IPv6 EPS Bearer contexts. Also, an APN may be configured to simultaneously allow IPv4 EPS Bearer contexts.

The MME supports IPv6 stateless dynamic auto-configuration. The mobile station may select any value for the interface identifier portion of the address. The link-local address is assigned by the MME to avoid any conflict between the mobile station link-local address and the MME address. The mobile station uses the interface identifier assigned by the MME during the stateless address auto-configuration procedure. Once this has completed, the mobile can select any interface identifier for further communication as long as it does not conflict with the MME's interface identifier that the mobile learned through router advertisement messages from the MME.

Control and configuration of the above is specified as part of the APN configuration on the MME, e.g., IPv6 address prefix and parameters for the IPv6 router advertisements. RADIUS VSAs may be used to override the APN configuration.

Following IPv6 EPS Bearer context establishment, the MME can perform either manual or automatic 6to4 tunneling, according to RFC 3056, Connection of IPv6 Domains Via IPv4 Clouds.

MME Interfaces Supporting IPv6 Transport

The following MME interfaces support IPv6 transport:

- S1-MME: runs S1-AP/SCTP over IPv6 and supports IPv6 addresses for S1-U endpoints.
- S3
- S6a
- S10
- S11
- S13
- SGs
- Sv

Load Balancing

Load balancing functionality permits UEs that are entering into an MME pool area to be directed to an appropriate MME in a more efficient manner, spreading the load across a number of MMEs.

Load balancing is achieved by setting a weight factor for each MME so that the probability of the eNodeB selecting an MME is proportional to its weight factor. The weight factor is typically set according to the capacity of an MME node relative to other MME nodes. The weight factor is sent from the MME to the eNodeB via S1-AP messages.

MME load balancing can be used in conjunction with congestion control. For more information on congestion control, refer to the Congestion Control section in this chapter.

Load Re-balancing

The MME load re-balancing functionality permits UEs that are registered on an MME (within an MME pool area) to be moved to another MME.

The rebalancing is triggered using an exec command on the mme-service from which UEs should be offloaded.

When initiated, the MME begins to offload a cross-section of its subscribers with minimal impact on the network and users. The MME avoids offloading only low activity users, and it offloads the UEs gradually (configurable from 1-1000 minutes). The load rebalancing can off-load part of or all the subscribers.

Management System Overview

The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management - focusing on providing superior quality network element (NE) and element management system (Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems - giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

The Operation and Maintenance module of the system offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces.

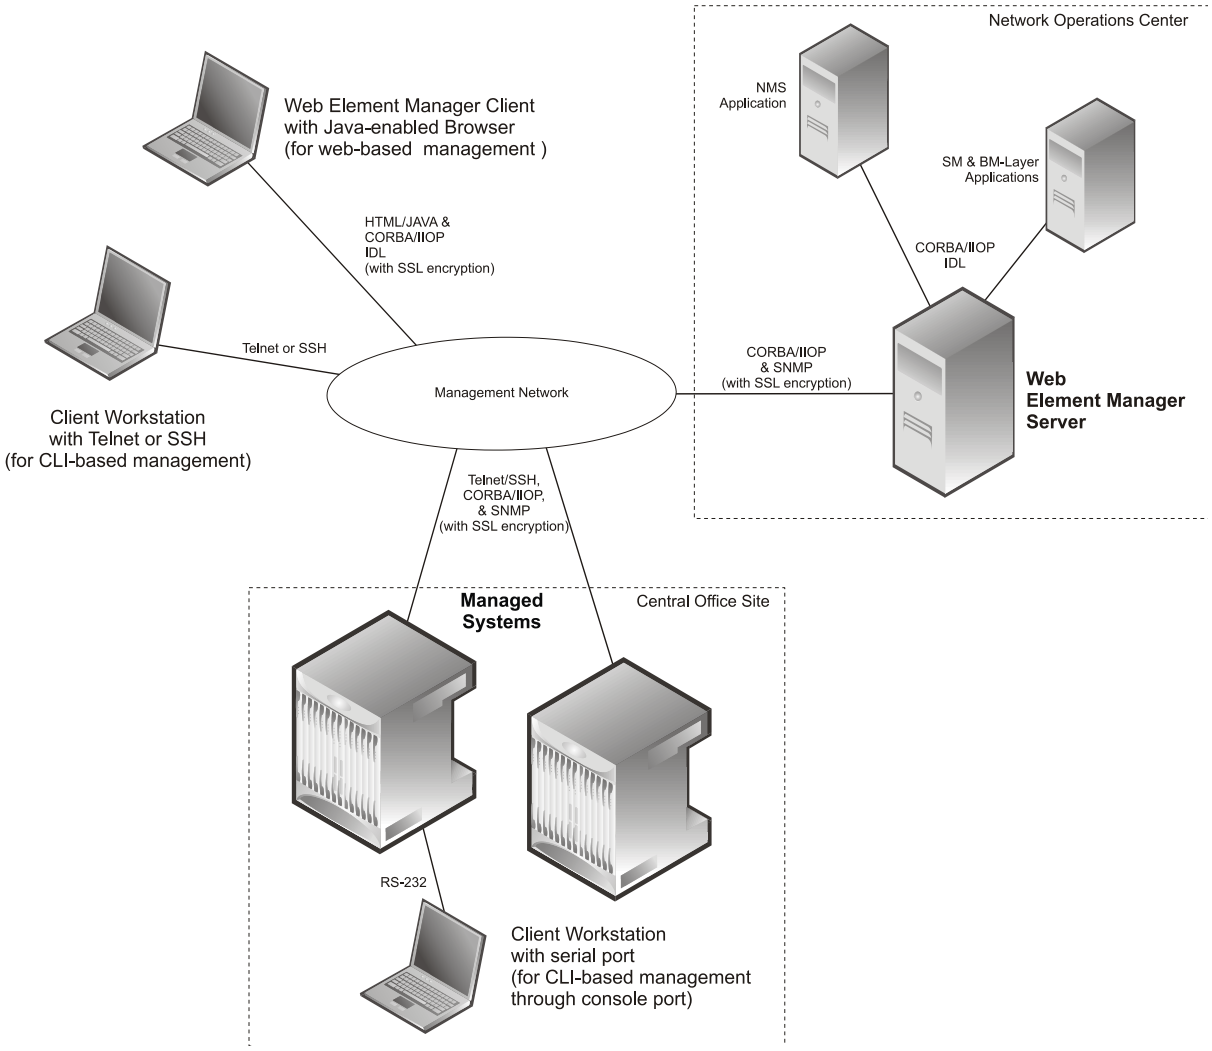
These include:

- Using the command line interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager application
- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000
- Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (i.e. Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities

- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network.

Figure 174. Element Management Methods



Important: MME management functionality is enabled by default for console-based access. For GUI-based management support, refer *Web Element Management System*. For more information on command line interface based management, refer to the *Command Line Interface Reference*.

MME Pooling

Provides support to configure MME pool area consisting multiple MMEs within which a UE may be served without any need to change the serving MME.

The benefits of MME pooling are:

- Enables Geographical Redundancy, as a pool can be distributed across sites.
- Increases overall capacity, as load sharing across the MMEs in a pool is possible (see the Load Balancing feature in this chapter).
- Converts inter-MME Tracking Area Updates (TAUs) to intra-MME TAUs for moves between the MMEs of the same pool. This substantially reduces signaling load as well as data transfer delays.
- Eases introduction of new nodes and replacement of old nodes as subscribers can be moved is a planned manner to the new node.
- Eliminates single point of failure between an eNodeB and MME.
- Enables service downtime free maintenance scheduling.

An MME Pool Area is defined as an area within which a UE may be served without need to change the serving MME. An MME Pool Area is served by one or more MMEs in parallel. MME Pool Areas are a collection of complete Tracking Areas. MME Pool Areas may overlap each other.

The Cisco MME supports MME Pooling functionality as defined in 3GPP TS 23.401. MME pooling allows carriers to load balance sessions among pooled MMEs.

The Cisco MME supports configuration of up to a pool size of 32 nodes.

MME Selection

The MME selection function selects an available MME for serving a UE. This feature is needed for MME selection for handover with minimal MME changes.

MME selection chooses an available MME for serving a UE. Selection is based on network topology, i.e. the selected MME serves the UE's location and in case of overlapping MME service areas, the selection function may prefer MME's with service areas that reduce the probability of changing the MME.

Mobile Equipment Identity Check

The Mobile Equipment Identity Check Procedure permits the operator(s) of the MME and/or the HSS and/or the PDN-GW to check the Mobile Equipment's identity with EIR.

The mobile equipment (ME) identity is checked through the MME by passing it to an Equipment Identity Register (EIR) over the S13 interface and then the MME analyzes the response from the EIR in order to determine its subsequent actions; like rejecting or attaching a UE.

Mobility Restriction

The following types of mobility restriction are supported on the MME:

- Handover Restriction

Handover Restriction

Mobility Restriction comprises the functions for restrictions to mobility handling of a UE in E-UTRAN access. In ECM-CONNECTED state, the core network provides the radio network with a Handover Restriction List.

The MME performs mobility or handover restrictions through the use of handover restriction lists. Handover restriction lists are used by the MME operator policy to specify roaming, service area, and access restrictions. Mobility restrictions at the MME are defined in 3GPP TS 23.401.

Multiple PDN Support

This feature provides multiple PDN connectivity support for UE initiated service requests.

The MME supports an UE-initiated connectivity establishment to separate P-GWs or a single P-GW in order to allow parallel access to multiple PDNs. Up to 11 PDNs are supported per subscriber.

NAS Protocol Support

MME provides this protocol support between the UE and the MME. The NAS protocol includes following elementary procedures for EPS Mobility Management (EMM) and EPS Session Management (ESM):

EPS Mobility Management (EMM)

This feature used to support the mobility of user equipment, such as informing the network of its present location and providing user identity confidentiality. It also provides connection management services to the session management (SM) sublayer.

An EMM context is established in the MME when an attach procedure is successfully completed. The EMM procedures are classified as follows:

- **EMM Common Procedures**: An EMM common procedure can always be initiated when a NAS signalling connection exists.

  Following are the common EMM procedure types:

  - Globally Unique Temporary Identity (GUTI) reallocation
  - Authentication and security mode
  - Identification
  - EMM information
- **EMM Specific Procedures**: This procedure provides Subscriber Detach or de-registration procedure.
- **EMM Connection Management Procedures**: This procedure provides connection management related function like Paging procedure.

EPS Session Management (ESM)

This feature is used to provide the subscriber session management for bearer context activation, deactivation, modification, and update procedures.

NAS Signalling Security

It provides integrity protection and encryption of NAS signalling. The NAS security association is between the UE and the MME.

The MME uses the NAS security mode command procedure to establish a NAS security association between the UE and MME, in order to protect the further NAS signalling messages.

The MME implements AES algorithm (128-EEA1 and 128-EEA2) for NAS signalling ciphering and SNOW 3G algorithm (128-EIA1 and 128-EIA2) for NAS signalling integrity protection.

- 128-EIA1 = SNOW 3G
- 128-EIA2 = AES

Operator Policy Support

The operator policy provides mechanisms to fine tune the behavior of subsets of subscribers above and beyond the behaviors described in the user profile. It also can be used to control the behavior of visiting subscribers in roaming scenarios, enforcing roaming agreements and providing a measure of local protection against foreign subscribers.

An operator policy associates APNs, APN profiles, an APN remap table, and a call-control profile to ranges of IMSIs. These profiles and tables are created and defined within their own configuration modes to generate sets of rules and instructions that can be reused and assigned to multiple policies. In this manner, an operator policy manages the application of rules governing the services, facilities, and privileges available to subscribers. These policies can override standard behaviors and provide mechanisms for an operator to get around the limitations of other infrastructure elements, such as DNS servers and HSSs.

The operator policy configuration to be applied to a subscriber is selected on the basis of the selection criteria in the subscriber mapping at attach time. A maximum of 1,024 operator policies can be configured. If a UE was associated with a specific operator policy and that policy is deleted, the next time the UE attempts to access the policy, it will attempt to find another policy with which to be associated.

A default operator policy can be configured and applied to all subscribers that do not match any of the per-PLMN or IMSI range policies.

Changes to the operator policy take effect when the subscriber re-attaches and subsequent EPS Bearer activations.

Refer to the *Operator Policy* chapter in this guide for more information.

Overload Management in MME

Provides mechanism to handle overload/congestion situation. It can use the NAS signalling to reject NAS requests from UEs on overload or congestion.

MME restricts the load that its eNodeBs are generating on it. This is achieved by the MME invoking the S1 interface overload procedure as per 3GPP TS 36.300 and 3GPP TS 36.413 to a proportion of the eNodeBs with which the MME has S1 interface connections.

Hardware and/or software failures within an MME may reduce the MME's load handling capability. Typically such failures result in alarms which alert the operator or Operation and Maintenance system.

For more information on congestion control management, refer to the *Congestion Control* chapter in the *System Administration Guide*.

Caution: Only if the operator or Operation and Maintenance system is sure that there is spare capacity in the rest of the pool, the operator or Operation and Maintenance system might use the load re-balancing procedure to move some load off an MME. However, extreme care is needed to ensure that this load re-balancing does not overload other MMEs within the pool area (or neighboring SGSNs) as this might lead to a much wider system failure.

Packet Data Network Gateway (P-GW) Selection

Provides a straightforward method based on a default APN provided during user attachment and authentication to assign the P-GW address in the VPLMN or HPLMN. The MME also has the capacity to use a DNS transaction to resolve an APN name provided by a UE to retrieve the PDN GW address.

P-GW selection allocates a P-GW that provides the PDN connectivity for the 3GPP access. The function uses subscriber information provided by the HSS and possibly additional criteria. For each of the subscribed PDNs, the HSS provides:

- an IP address of a P-GW and an APN, or
- an APN and an indication for this APN whether the allocation of a P-GW from the visited PLMN is allowed or whether a P-GW from the home PLMN shall be allocated.

The HSS also indicates the default APN for the UE. To establish connectivity with a PDN when the UE is already connected to one or more PDNs, the UE provides the requested APN for the PDN GW selection function.

If the HSS provides an APN of a PDN and the subscription allows for allocation of a PDN GW from the visited PLMN for this APN, the PDN GW selection function derives a PDN GW address from the visited PLMN. If a visited PDN GW address cannot be derived, or if the subscription does not allow for allocation of a PDN GW from the visited PLMN, then the APN is used to derive a PDN GW address from the HPLMN.

Radio Resource Management Functions

Radio resource management functions are concerned with the allocation and maintenance of radio communication paths, and are performed by the radio access network.

To support radio resource management in E-UTRAN, the MME provides the RAT/Frequency Selection Priority (RFSP) parameter to an eNodeB across S1. The RFSP is a “per UE” parameter that is used by the E-UTRAN to derive UE specific cell reselection priorities to control idle mode camping. The RFSP can also be used by the E-UTRAN to decide on redirecting active mode UEs to different frequency layers or RATs.

The MME receives the RFSP from the HSS during the attach procedure. For non-roaming subscribers, the MME transparently forwards the RFSP to the eNodeB across S1. For roaming subscribers, the MME may alternatively send an RFSP value to the eNodeB across S1 that is based on the visited network policy, such as an RFSP pre-configured per Home-PLMN or a single RFSP’s values to be used for all roamers independent of the Home-PLMN.

Reachability Management

It provides a mechanism to track a UE which is in idle state for EPS connection management.

To reach a UE in idle state the MME initiates paging to all eNodeBs in all tracking areas in the TA list assigned to the UE. The EPS session manager have knowledge about all the eNodeB associations to the MME and generates a list of eNodeBs that needs to be paged to reach a particular UE.

The location of a UE in ECM-IDLE state is known by the network on a Tracking Area List granularity. A UE in ECM-IDLE state is paged in all cells of the Tracking Areas in which it is currently registered. The UE may be registered in multiple Tracking Areas. A UE performs periodic Tracking Area Updates to ensure its reachability from the network.

SCTP Multi-homing Support

This sections describes multi-homing support for specific interfaces on the MME.

SCTP Multi-homing for S6a

The Cisco MME service supports up to four SCTP bind end point IPv4 or IPv6 addresses for the S6a interface.

SCTP Multi-homing for S1-MME

The Cisco MME service supports up to two SCTP bind end point IPv4 or IPv6 addresses for the S1 -MME interface.

Serving Gateway Pooling Support

The S-GW supports independent service areas from MME pooling areas. Each cell is associated to a pool of MMEs and a pool of Serving Gateways. Once a cell selects an MME, that MME is able to select an S-GW which is in an S-GW pool supported by the cell.

Static S-GW pools can be configurable on the MME. Each pool is organized as a set of S-GWs and the Tracking Area Identities (TAIs) supported by them, known as a service area (SA). The incoming TAI is used to select an SA. Then, based on protocol and statistical weight factors, an S-GW is selected from the pool serving that SA. The same list of S-GWs may serve multiple TAIs. Static S-GW pools are used if there is no DNS configured or as a fallback if DNS discovery fails.

For additional Information on TAI lists, refer to the Tracking Area List Management section in this overview.

Serving Gateway Selection

The Serving Gateway (S-GW) selection function selects an available S-GW to serve a UE. This feature reduces the probability of changing the S-GW and a load balancing between S-GWs. The MME uses DNS procedures for S-GW selection.

The selection is based on network topology; the selected S-GW serves the UE's location, and in the case of overlapping S-GW service areas, the selection may prefer S-GWs with service areas that reduce the probability of changing the S-GW. If a subscriber of a GTP-only network roams into a PMIP network, the PDN GWs (P-GWs) selected for local breakout supports the PMIP protocol, while P-GWs for home routed traffic use GTP. This means the S-GW selected for such subscribers may need to support both GTP and PMIP, so that it is possible to set up both local breakout and home routed sessions for these subscribers.

Session and Quality of Service Management

This support provides a foundation for contributing towards improved Quality of User Experience (QoE) by enabling deterministic end-to-end forwarding and scheduling treatments for different services or classes of applications pursuant to their requirements for committed bandwidth resources, jitter and delay. In this way, each application receives the service treatment that users expect.

The MME Operator Policy configuration allows the specification of QoS for each traffic class that can either be used as a default or as an over ride to the HSS settings.

In LTE-EPC 4G architectures, QoS management is network controlled via dynamic policy interactions between the PCRF and PDN GW. EPS bearer management is used to establish, modify or remove dedicated EPC bearers in order to provide service treatments tied to the needs of specific applications/service data flows. The service priority is provisioned based on QoS Class Identifiers (QCI) in the Gx policy signaling. PCRF signaling interaction may also be used to establish or modify the APN-AMBR attribute assigned to the default EPS bearer.

When it is necessary to set-up a dedicated bearer, the PDN GW initiates the Create Dedicated Bearer Request which includes the IMSI (permanent identity of mobile access terminal), Traffic Flow Template (TFT - 5-tuple packet filters) and S5 Tunnel Endpoint ID (TEID) information that is propagated downstream via the S-GW over the S11 interface to the MME. The Dedicated Bearer signaling includes requested QoS information such as QCI, Allocation and Retention Priority (ARP), Guaranteed Bit Rate (GBR - guaranteed minimum sending rate) and Maximum Bit Rate (MBR- maximum burst size).

The MME allocates a unique EPS bearer identity for every dedicated bearer and encodes this information in a Session Management Request that includes Protocol Transaction ID (PTI), TFT's and EPS bearer QoS parameters. The MME signals the Bearer Setup Request in the S1-MME message toward the neighboring eNodeB.

Subscriber Level Session Trace

The Subscriber Level Trace provides a 3GPP standards-based session-level trace function for call debugging and testing new functions and access terminals in an LTE environment.

In general, the Session Trace capability records and forwards all control activity for the monitored subscriber on the monitored interfaces. This is typically all the signaling and authentication/subscriber services messages that flow when a UE connects to the access network.

As a complement to Cisco's protocol monitoring function, the MME supports 3GPP standards based session level trace capabilities to monitor all call control events on the respective monitored interfaces including S6a, S1-MME and S11. The trace can be initiated using multiple methods:

- Management initiation via direct CLI configuration
- Management initiation at HSS with trace activation via authentication response messages over S6a reference interface
- Signaling based activation through signaling from subscriber access terminal

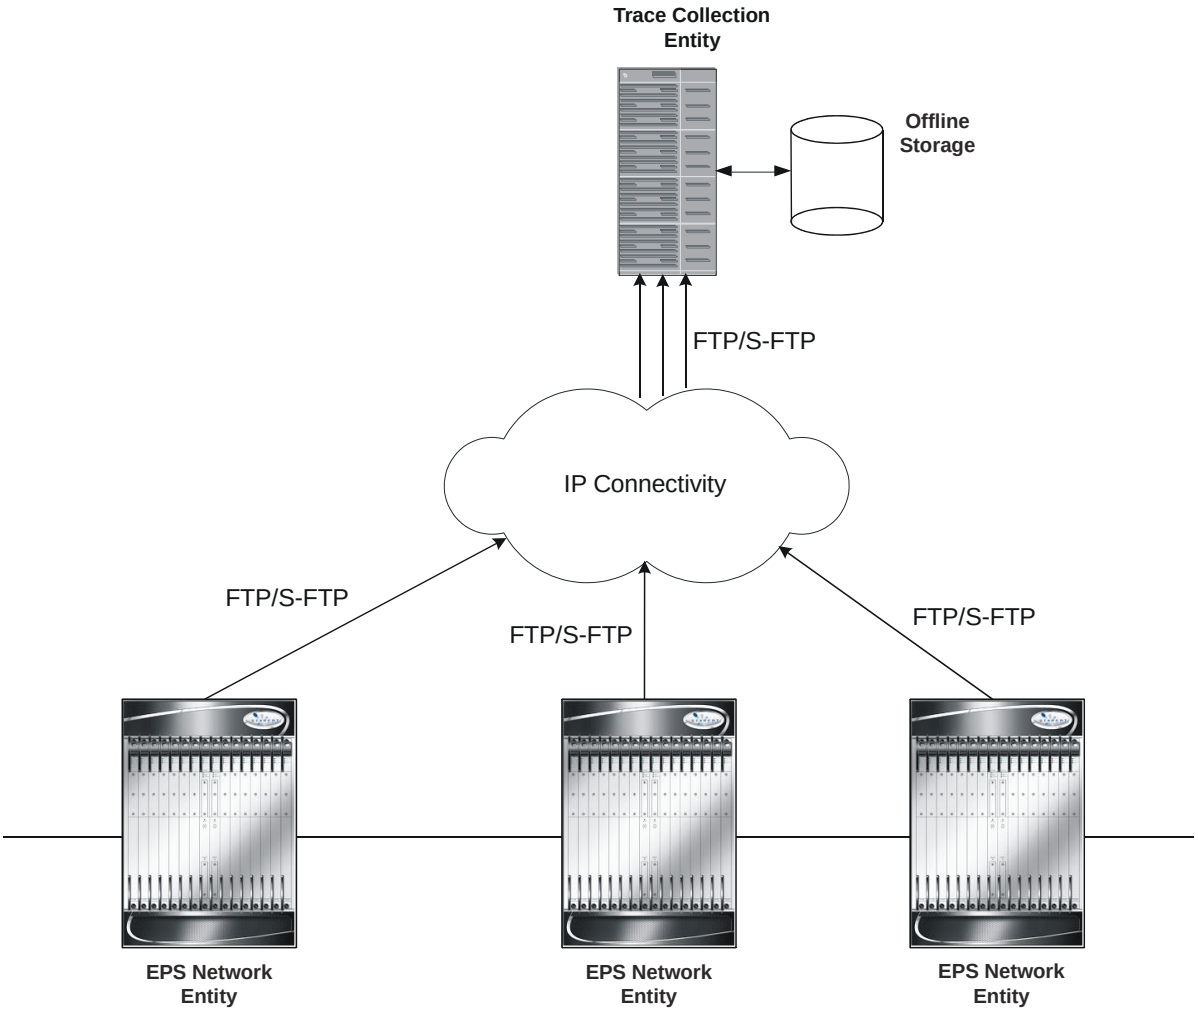
The session level trace function consists of trace activation followed by triggers. The EPC network element buffers the trace activation instructions for the provisioned subscriber in memory using camp-on monitoring. Trace files for active calls are buffered as XML files using non-volatile memory on the local dual redundant hard drives. The Trace Depth defines the granularity of data to be traced. Six levels are defined including Maximum, Minimum and Medium with ability to configure additional levels based on vendor extensions.

All call control activity for active and recorded sessions is sent to an off-line Trace Collection Entity (TCE) using a standards-based XML format over a FTP or secure FTP (SFTP) connection.

Note: In the current release the IPv4 interfaces are used to provide connectivity to the TCE. Trace activation is based on IMSI or IMEI and only *Maximum Trace Depth* is supported in this release.

The following figure shows a high-level overview of the session-trace functionality and deployment scenario:

Figure 175. Session Trace Function and Interfaces



For more information on this feature, refer to the *Configuring Subscriber Session Tracing* chapter in the *MME Service Administration Guide*.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, number of sessions etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert**: A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm**: Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps**: SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values.

Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.

- **Logs**: The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

Logs are supported in both the Alert and the Alarm models.

- **Alarm System**: High threshold alarms generated within the specified polling interval are considered "outstanding" until a the condition no longer exists or a condition clear alarm is generated. "Outstanding" alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

Tracking Area List Management

Provides the functions to allocate and reallocate a Tracking Area Identity (TAI) list to the UE to minimize Tracking Area Updates (TAUs).

The MME assigns the TAI list to a UE so as to minimize the TAUs that are sent by the UE. The TAI list should be kept to a minimum in order to maintain a lower paging load.

To avoid a ping-pong effect, the MME includes the last visited TAI (provided that the tracking area is managed by the MME) in the TAI list assigned to the UE.

Tracking area lists assigned to different UEs moving in from the same tracking area should be different to avoid Tracking Area Update message overflow.

Features and Functionality - External Application Support

This section describes the features and functions of external applications supported on the MME. These services require additional licenses to implement the functionality.

This section describes following external applications:

- Web Element Management System

Web Element Management System

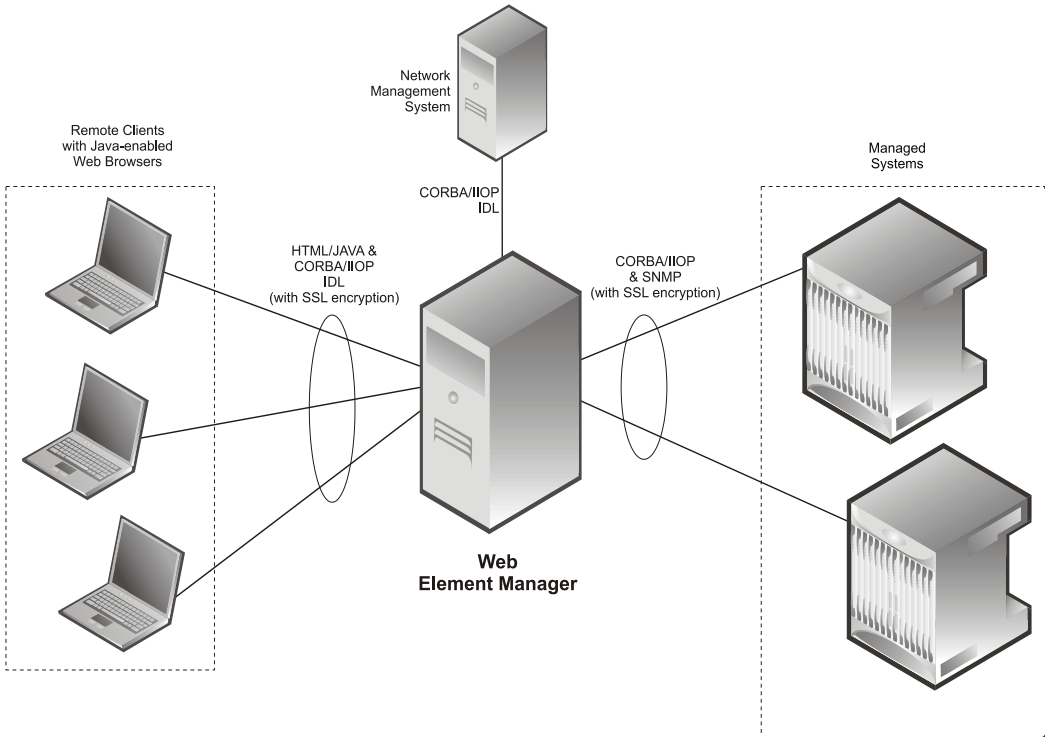
Provides a graphical user interface (GUI) for performing fault, configuration, accounting, performance, and security (FCAPS) management.

The Web Element Manager is a Common Object Request Broker Architecture (CORBA)-based application that provides complete fault, configuration, accounting, performance, and security (FCAPS) management capability for the system.

For maximum flexibility and scalability, the Web Element Manager application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the fault-tolerant Sun® Solaris® operating system.

The following figure demonstrates various interfaces between the Cisco Web Element Manager and other network components.

Figure 176. Web Element Manager Network Interfaces



Important: MME management functionality is enabled by default for console-based access. For GUI-based management support, refer *Web Element Management System*.

Features and Functionality - Licensed Enhanced Feature Software

This section describes the optional enhanced features and functions for MME service.

Important: The following features require the purchase of an additional feature license to implement the functionality with the MME service.

This section describes following enhanced features:

- Circuit Switched Fall Back (CSFB) and SMS over SGs Interface
- IP Security (IPSec)
- Lawful Intercept
- Location Services
- Optimized Paging Support
- Session Recovery Support
- User Location Information Reporting

Circuit Switched Fall Back (CSFB) and SMS over SGs Interface

This feature requires that a valid license key be installed. Contact your Cisco Account or Support representative for information on how to obtain a license.

Circuit Switched Fall Back (CSFB) enables the UE to camp on an EUTRAN cell and originate or terminate voice calls through a forced switchover to the circuit switched (CS) domain or other CS-domain services (e.g., Location Services (LCS) or supplementary services). Additionally, SMS delivery via the CS core network is realized without CSFB. Since LTE EPC networks were not meant to directly anchor CS connections, when any CS voice services are initiated, any PS based data activities on the EUTRAN network will be temporarily suspended (either the data transfer is suspended or the packet switched connection is handed over to the 2G/3G network).

Important: CSFB to CDMA 1x networks is not supported in this release.

CSFB provides an interim solution for enabling telephony and SMS services for LTE operators that do not plan to deploy IMS packet switched services at initial service launch.

CSFB function is realized by reusing Gs interface mechanisms, as defined in 3GPP TS 29.018, on the interface between the MME in the EPS and the VLR. This interface is called the SGs interface. The SGs interface connects the databases in the VLR and the MME.

EPC core networks are designed for all IP services and as such lack intrinsic support for circuit switched voice and telephony applications. This presents challenges for those operators that do not plan to launch packet switched IMS core networks at initial service deployment. CSFB represents an interim solution to address this problem by enabling dual radio mobile devices (LTE/GSM/UMTS or CDMA1xRTT) to fall back to GSM/UMTS or CDMA1x access networks to receive incoming or place outgoing voice calls. Highlights of the CSFB procedure are as follows:

- **Preparation Phase**:

- When the GSM/UMTS/LTE access terminal attaches to the EUTRAN access network, it uses combined attachment procedures to request assistance from the MME to register its presence in the 2G/3G network.
- The MME uses SGs signaling to the MSC/VLR to register on behalf of the AT to the 2G/3G network. The MME represents itself as an SGSN to the MSC and the MSC performs a location update to the SGSN in the target 2G/3G network.
- The MME uses the Tracking Area Identity provided by UE to compute the Location Area Identity it provides to the MSC.

- **Execution Phase: Mobile Terminated Call**:
  - When a call comes in at the MSC for the user, the MSC signals the incoming call via the SGs interface to MME.
  - If the AT is an active state, the MME forwards the request directly to the mobile. If the user wishes to receive the call the UE instructs the MME to hand over the call to the 2G/3G network. The MME then informs the eNodeB to initiate the handoff.
  - If the AT is in dormant state, the MME attempts to page it at every eNodeB within the Tracking Area list to reestablish the radio connection. As no data transfer is in progress, there are no IP data sessions to handover and the mobile switches to its 2G/3G radio to establish the connection with the target access network.
  - If the mobile is active and an IP data transfer is in progress at the time of the handover, the data transfer can either be suspended or the packet switched connection can be handed over if the target network supports Dual Transfer Mode. Note that this is typically only supported on UMTS networks.
  - Once the access terminal attaches to the 2G/3G cell, it answers the initial paging via the target cell.
- **Execution Phase: Mobile Originated Calls**
  - This is very similar to the procedure for Mobile Terminated Calls, except there is no requirement for idle mode paging for incoming calls and the AT has no need to send a paging response to the MSC after it attaches to the target 2G/3G network.

The following CSFB features are supported:

- Release 8 and Release 9 Specification Support
- SGs-AP Encode/Decode of all messages
- SGs-AP Procedure Support
  - Paging
  - Location Update
  - Non-EPS Alert
  - Explicit IMSI Detach
  - Implicit IMSI Detach
  - VLR Failure
  - HSS Failure
  - MM Information
  - NAS Message Tunneling
  - Service Request
  - MME Failure

- SMS
- Mobile Originating Voice Call
- Mobile Terminating Voice Call
- Gn/Gp Handover
- S3 Handover
- Basic and Enhanced TAI to LAI Mapping
- Basic LAI to VLR Mapping
- VLR association distribution among multiple MMEMGRs
- IMSI Paging Procedure
- IPv6 Transport for SGs interface
- SNMP Trap Support (Service/VLR association)
- Operator Policy Support
    - SMS-only
    - Disallow CSFB
- PS Suspend/Resume over S11 (Release 8)
- PS Suspend/Resume over S3/S11 (Release 9)
-
- Support for SGs AP Timers: TS6-1

IP Security (IPSec)

This feature requires that a valid license key be installed. Contact your Cisco Account or Support representative for information on how to obtain a license.

IP Security provides a mechanism for establishing secure tunnels from mobile subscribers to pre-defined endpoints (i.e. enterprise or home networks) in accordance with the following standards:

- RFC 2401, Security Architecture for the Internet Protocol
- RFC 2402, IP Authentication Header (AH)
- RFC 2406, IP Encapsulating Security Payload (ESP)
- RFC 2409, The Internet Key Exchange (IKE)
- RFC-3193, Securing L2TP using IPSEC, November 2001

IP Security (IPSec) is a suite of protocols that interact with one another to provide secure private communications across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways. IPSec can be implemented on the system for the following applications:

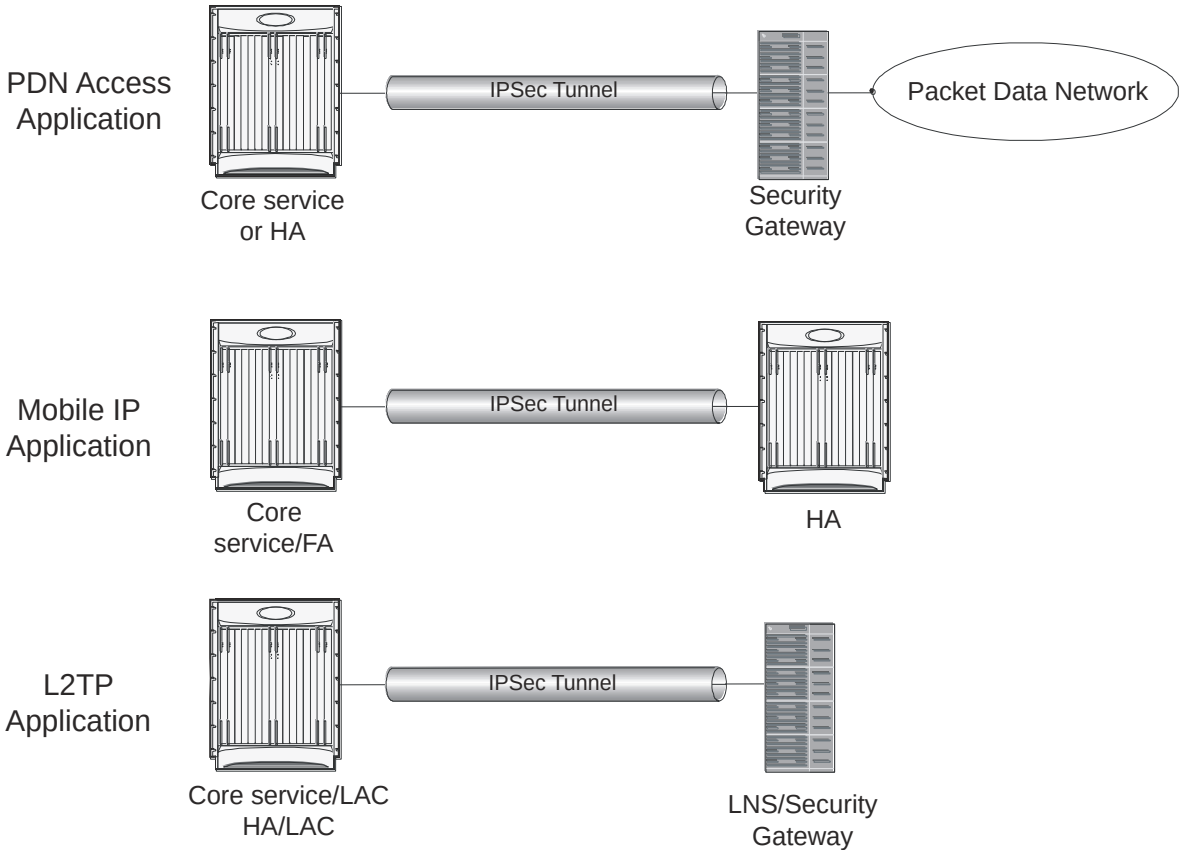
- **PDN Access**: Subscriber IP traffic is routed over an IPSec tunnel from the system to a secure gateway on the packet data network (PDN) as determined by access control list (ACL) criteria.
- **Mobile IP**: Mobile IP control signals and subscriber data is encapsulated in IPSec tunnels that are established between foreign agents (FAs) and home agents (HAs) over the Pi interfaces.

Important: Once an IPSec tunnel is established between an FA and HA for a particular subscriber, all new Mobile IP sessions using the same FA and HA are passed over the tunnel regardless of whether or not IPSec is supported for the new subscriber sessions. Data for existing Mobile IP sessions is unaffected.

- **L2TP**: L2TP-encapsulated packets are routed from the system to an LNS/secure gateway over an IPSec tunnel.

The following figure shows IPSec configurations.

Figure 177. IPSec Applications



Important: For more information on IPSec support, refer to the *IP Security* appendix in the *MME Administration Guide*.

Lawful Intercept

The feature use license for Lawful Intercept on the MME is included in the MME session use license.

The Cisco Lawful Intercept feature is supported on the MME. Lawful Intercept is a license-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your Cisco account representative.

Location Services

This feature requires that a valid license key be installed. Contact your Cisco Account or Support representative for information on how to obtain a license.

LoCation Services (LCS) on the MME and SGSN is a 3GPP standards-compliant feature that enables the system (MME or SGSN) to collect and use or share location (geographical position) information for connected UEs in support of a variety of location services.

Refer to the *Location Services* chapter in the *MME Administration Guide* for more information.

Optimized Paging Support

This feature requires that a valid license key be installed. Contact your Cisco Account or Support representative for information on how to obtain a license.

Also known as heuristic or idle-mode paging, this feature reduces network operations cost through more efficient utilization of paging resources and reduced paging load in the EUTRAN access network.

Idle mode paging over EUTRAN access networks is an expensive operation that causes volumes of signaling traffic between the S-GW and MME/SGSN. This problem is acute in the radio access network, where paging is a shared resource with finite capacity. When a request for an idle mode access terminal is received by the S-GW, the MME floods the paging notification message to all eNodeBs in the Tracking Area List (TAI). To appreciate the magnitude of the problem, consider a network with three million subscribers and a total of 800 eNodeBs in the TAI. If each subscriber was to receive one page during the busy hour, the total number of paging messages would exceed one million messages per second.

To limit the volume of unnecessary paging related signaling, the Cisco MME provides intelligent paging heuristics. Each MME maintains a list of "n" last heard from eNodeBs inside the TAI for the UE. The intent is to keep track of the eNodeBs that the AT commonly attaches to such as the cells located near a person's residence and place of work. During the average day, the typical worker spends the most time attaching to one of these two locations. When an incoming page arrives for the idle mode user, the MME attempts to page the user at the last heard from eNodeB. The MME uses Tracking Area Updates to build this local table. If no response is received within a configurable period, the MME attempts to page the user at the last "n" heard from eNodeBs. If the MME has still not received acknowledgement from the idle mode UE, only then does it flood the paging messages to all eNodeBs in the TAI.

In the majority of instances with this procedure, the UE will be paged in a small set of eNodeBs where it is most likely to be attached.

Session Recovery Support

The feature use license for Session Recovery on the MME is included in the MME session use license.

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

This feature is also useful for Software Patch Upgrade activities. If session recovery feature is enabled during the software patch upgrading, it helps to permit preservation of existing sessions on the active PSC during the upgrade process.

Session recovery is performed by mirroring key software processes (e.g. session manager and AAA manager) within the system. These mirrored processes remain in an idle state (in standby-mode), wherein they perform no processing, until

they may be needed in the case of a software failure (e.g. a session manager task aborts). The system spawns new instances of "standby mode" session and AAA managers for each active control processor (CP) being used.

Additionally, other key system-level software tasks, such as VPN manager, are performed on a physically separate packet processing card to ensure that a double software fault (e.g. session manager and VPN manager fails at same time on same card) cannot occur. The packet processing card used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

The additional hardware resources required for session recovery include a standby system processor card (SPC) and a standby packet processing card.

There are two modes for Session Recovery.

- **Task recovery mode**: Wherein one or more session manager failures occur and are recovered without the need to use resources on a standby packet processing card. In this mode, recovery is performed by using the mirrored "standby-mode" session manager task(s) running on active packet processing cards. The "standby-mode" task is renamed, made active, and is then populated using information from other tasks such as AAA manager.
- **Full packet processing card recovery mode**: Used when a PSC or PSC2 hardware failure occurs, or when a packet processing card migration failure happens. In this mode, the standby packet processing card is made active and the "standby-mode" session manager and AAA manager tasks on the newly activated packet processing card perform session recovery.

Session/Call state information is saved in the peer AAA manager task because each AAA manager and session manager task is paired together. These pairs are started on physically different packet processing cards to ensure task recovery.

Important: For more information on session recovery support, refer to the *Session Recovery* chapter in the *System Administration Guide*.

User Location Information Reporting

This feature requires that a valid license key be installed. Contact your Cisco Account or Support representative for information on how to obtain a license.

User Location Information (ULI) Reporting allows the eNodeB to report the location of a UE to the MME, when requested by a P-GW.

The following procedures are used over the S1-MME interface to initiate and stop location reporting between the MME and eNodeB:

- **Location Reporting Control**: The purpose of Location Reporting Control procedure is to allow the MME to request that the eNodeB report where the UE is currently located. This procedure uses UE-associated signaling.
- **Location Report Failure Indication**: The Location Report Failure Indication procedure is initiated by an eNodeB in order to inform the MME that a Location Reporting Control procedure has failed. This procedure uses UE-associated signalling.
- **Location Report**: The purpose of Location Report procedure is to provide the UE's current location to the MME. This procedure uses UE-associated signalling.

The start/stop trigger for location reporting for a UE is reported to the MME by the S-GW over the S11 interface. The Change Reporting Action (CRA) Information Element (IE) is used for this purpose. The MME updates the location to the S-GW using the User Location Information (ULI) IE.

The following S11 messages are used to transfer CRA and ULI information between the MME and S-GW:

- **Create Session Request**: The ULI IE is included for E-UTRAN Initial Attach and UE-requested PDN Connectivity procedures. It includes ECGI and TAI. The MME includes the ULI IE for TAU/ X2 -Handover procedure if the P-GW has requested location information change reporting and the MME support location information change reporting. The S-GW includes the ULI IE on S5/S8 exchanges if it receives the ULI from the MME. If the MME supports change reporting, it sets the corresponding indication flag in the Create Session Request message.
- **Create Session Response**: The CRA IE in the Create Session Response message can be populated by the S-GW to indicate the type of reporting required.
- **Create Bearer Request**: The CRA IE is included with the appropriate Action field if the Location Change Reporting mechanism is to be started or stopped for the subscriber in the MME.
- **Modify Bearer Request**: The MME includes the ULI IE for TAU/Handover procedures and UE-initiated Service Request procedures if the P-GW has requested location information change reporting and the MME supports location information change reporting. The S-GW includes this IE on S5/S8 exchanges if it receives the ULI from the MME.
- **Modify Bearer Response**: The CRA IE is included with the appropriate Action field if the Location Change Reporting mechanism is to be started or stopped for the subscriber in the MME.
- **Delete Session Request**: The MME includes the ULI IE for the Detach procedure if the P-GW has requested location information change reporting and MME supports location information change reporting. The S-GW includes this IE on S5/S8 exchanges if it receives the ULI from the MME.
- **Update Bearer Request**: The CRA IE is included with the appropriate Action field if the Location Change Reporting mechanism is to be started or stopped for the subscriber in the MME.
- **Change Notification Request**: If no existing procedure is running for a UE, a Change Notification Request is sent upon receipt of an S1-AP location report message. If an existing procedure is running, one of the following messages reports the ULI:
  - Create Session Request
  - Create Bearer Response
  - Modify Bearer Request
  - Update Bearer Response
  - Delete Bearer Response
  - Delete Session Request

  If an existing Change Notification Request is pending, it is aborted and a new one is sent.

Important: Information on configuring User Location Information Reporting support is located in the *Configuring Optional Features on the MME* section of the *Mobility Management Entity Configuration* chapter in this guide.

How the MME Works

This section provides information on the function and procedures of the MME in an EPC network and presents message flows for different stages of session setup.

The following procedures are supported in this release:

- EPS Bearer Context Processing
- Purge Procedure
- Paging Procedure
- Subscriber-initiated Initial Attach Procedure
- Subscriber-initiated Detach Procedure
- Service Request Procedures
    - UE-initiated Service Request Procedure
    - Network-initiated Service Request Procedure

EPS Bearer Context Processing

EPS Bearer context processing is based on the APN that the subscriber is attempting to access. Templates for all of the possible APNs that subscribers will be accessing must be configured within the P-GW system.

Each APN template consists of parameters pertaining to how EPS Bearer contexts are processed such as the following:

- **PDN Type:** The system supports IPv4, IPv6, or IPv4v6.
- **Timeout:** Absolute and idle session timeout values specify the amount of time that an MS can remain connected.
- **Quality of Service:** Parameters pertaining to QoS feature support such as for Traffic Policing and traffic class.

A total of 11 EPS bearer contexts are supported per subscriber. These could be all dedicated, or 1 default and 10 dedicated or any combination of default and dedicated context. Note that there must be at least one default EPS bearer context in order for dedicated context to come up.

Purge Procedure

The purge procedure is employed by the Cisco MME to inform the concerned node that the MME has removed the EPS bearer contexts of a detached UE. This is usually invoked when the number of records exceeds the maximum capacity of the system.

Paging Procedure

Paging is initiated when there is data to be sent to an idle UE to trigger a service request from the UE. Once the UE reaches connected state, the data is forwarded to it.

Paging retransmission can be controlled by configuring a paging-timer and retransmission attempts on system.

Subscriber-initiated Initial Attach Procedure

The following figure and the text that follows describe the message flow for a successful user-initiated subscriber attach procedure.

Figure 178. Subscriber-initiated Attach (initial) Call Flow

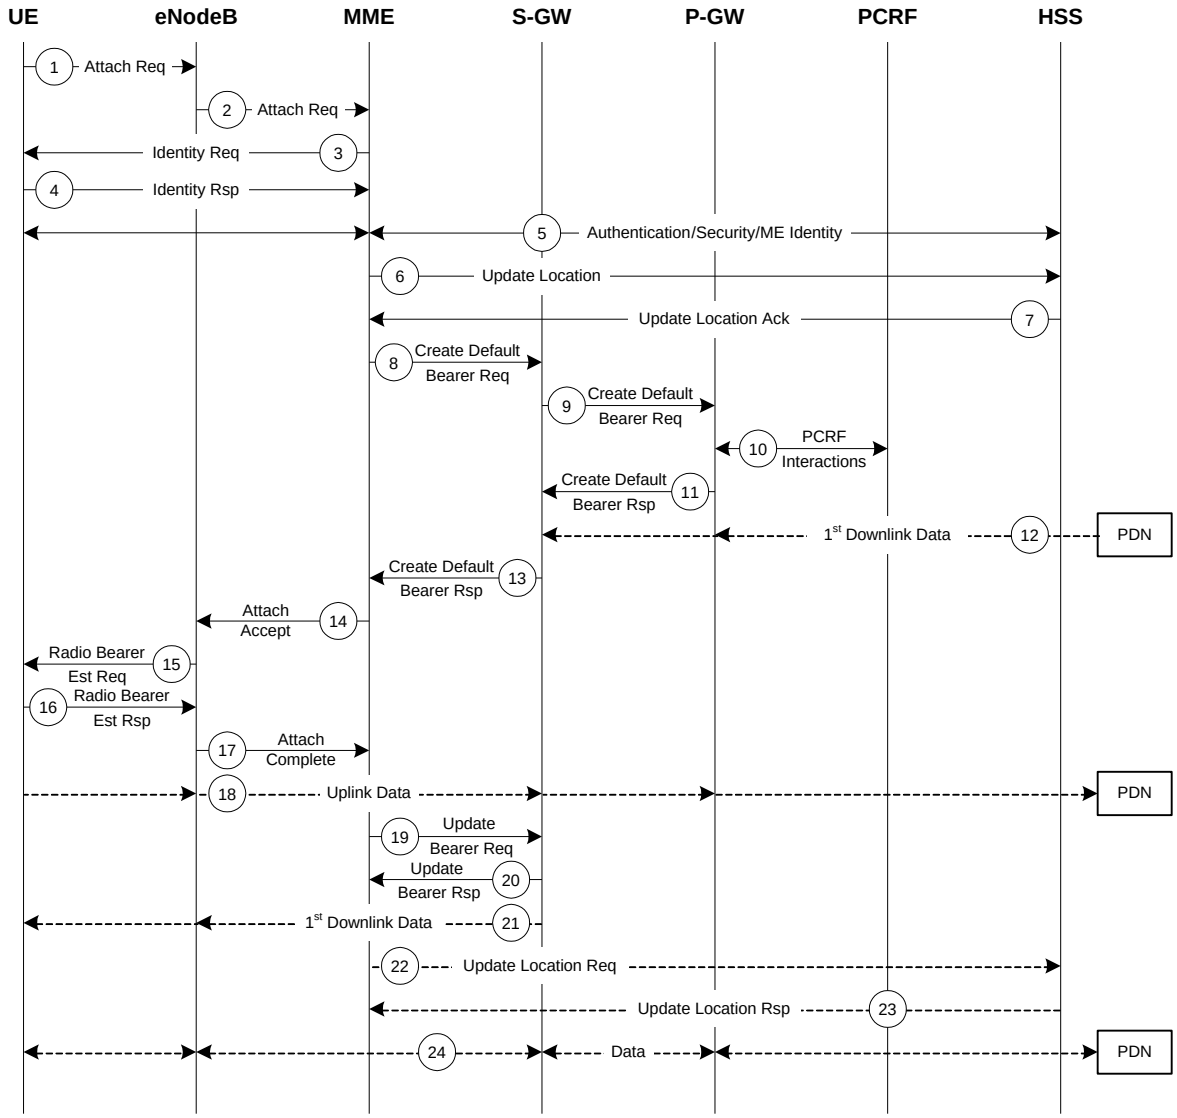


Table 79. Subscriber-initiated Attach (initial) Call Flow Description

Step	Description
1	The UE initiates the Attach procedure by the transmission of an Attach Request (IMSI or old GUTI, last visited TAI (if available), UE Network Capability, PDN Address Allocation, Protocol Configuration Options, Attach Type) message together with an indication of the Selected Network to the eNodeB. IMSI is included if the UE does not have a valid GUTI available. If the UE has a valid GUTI, it is included.

Step	Description
2	The eNodeB derives the MME from the GUTI and from the indicated Selected Network. If that MME is not associated with the eNodeB, the eNodeB selects an MME using an "MME selection function". The eNodeB forwards the Attach Request message to the new MME contained in a S1-MME control message (Initial UE message) together with the Selected Network and an indication of the E-UTRAN Area identity, a globally unique E-UTRAN ID of the cell from where it received the message to the new MME.
3	If the UE is unknown in the MME, the MME sends an Identity Request to the UE to request the IMSI.
4	The UE responds with Identity Response (IMSI).
5	If no UE context for the UE exists anywhere in the network, authentication is mandatory. Otherwise this step is optional. However, at least integrity checking is started and the ME Identity is retrieved from the UE at Initial Attach. The authentication functions, if performed this step, involves AKA authentication and establishment of a NAS level security association with the UE in order to protect further NAS protocol messages.
6	The MME sends an Update Location Request (MME Identity, IMSI, ME Identity) to the HSS.
7	The HSS acknowledges the Update Location message by sending an Update Location Ack to the MME. This message also contains the Insert Subscriber Data (IMSI, Subscription Data) Request. The Subscription Data contains the list of all APNs that the UE is permitted to access, an indication about which of those APNs is the Default APN, and the 'EPS subscribed QoS profile' for each permitted APN. If the Update Location is rejected by the HSS, the MME rejects the Attach Request from the UE with an appropriate cause.
8	The MME selects an S-GW using "Serving GW selection function" and allocates an EPS Bearer Identity for the Default Bearer associated with the UE. If the PDN subscription context contains no P-GW address the MME selects a P-GW as described in clause "PDN GW selection function". Then it sends a Create Default Bearer Request (IMSI, MME Context ID, APN, RAT type, Default Bearer QoS, PDN Address Allocation, AMBR, EPS Bearer Identity, Protocol Configuration Options, ME Identity, User Location Information) message to the selected S-GW.
9	The S-GW creates a new entry in its EPS Bearer table and sends a Create Default Bearer Request (IMSI, APN, S-GW Address for the user plane, S-GW TEID of the user plane, S-GW TEID of the control plane, RAT type, Default Bearer QoS, PDN Address Allocation, AMBR, EPS Bearer Identity, Protocol Configuration Options, ME Identity, User Location Information) message to the P-GW.
10	If dynamic PCC is deployed, the P-GW interacts with the PCRF to get the default PCC rules for the UE. The IMSI, UE IP address, User Location Information, RAT type, AMBR are provided to the PCRF by the P-GW if received by the previous message.
11	The P-GW returns a Create Default Bearer Response (P-GW Address for the user plane, P-GW TEID of the user plane, P-GW TEID of the control plane, PDN Address Information, EPS Bearer Identity, Protocol Configuration Options) message to the S-GW. PDN Address Information is included if the P-GW allocated a PDN address Based on PDN Address Allocation received in the Create Default Bearer Request. PDN Address Information contains an IPv4 address for IPv4 and/or an IPv6 prefix and an Interface Identifier for IPv6. The P-GW takes into account the UE IP version capability indicated in the PDN Address Allocation and the policies of operator when the P-GW allocates the PDN Address Information. Whether the IP address is negotiated by the UE after completion of the Attach procedure, this is indicated in the Create Default Bearer Response.
12	The Downlink (DL) Data can start flowing towards S-GW. The S-GW buffers the data.
13	The S-GW returns a Create Default Bearer Response (PDN Address Information, S-GW address for User Plane, S-GW TEID for User Plane, S-GW Context ID, EPS Bearer Identity, Protocol Configuration Options) message to the new MME. PDN Address Information is included if it was provided by the P-GW.
14	The new MME sends an Attach Accept (APN, GUTI, PDN Address Information, TAI List, EPS Bearer Identity, Session Management Configuration IE, Protocol Configuration Options) message to the eNodeB.
15	The eNodeB sends Radio Bearer Establishment Request including the EPS Radio Bearer Identity to the UE. The Attach Accept message is also sent along to the UE.

Step	Description
16	The UE sends the Radio Bearer Establishment Response to the eNodeB. In this message, the Attach Complete message (EPS Bearer Identity) is included.
17	The eNodeB forwards the Attach Complete (EPS Bearer Identity) message to the MME.
18	The Attach is complete and UE sends data over the default bearer. At this time the UE can send uplink packets towards the eNodeB which are then tunnelled to the S-GW and P-GW.
19	The MME sends an Update Bearer Request (eNodeB address, eNodeB TEID) message to the S-GW.
20	The S-GW acknowledges by sending Update Bearer Response (EPS Bearer Identity) message to the MME.
21	The S-GW sends its buffered downlink packets.
22	After the MME receives Update Bearer Response (EPS Bearer Identity) message, if an EPS bearer was established and the subscription data indicates that the user is allowed to perform handover to non-3GPP accesses, and if the MME selected a P-GW that is different from the P-GW address which was indicated by the HSS in the PDN subscription context, the MME sends an Update Location Request including the APN and P-GW address to the HSS for mobility with non-3GPP accesses.
23	The HSS stores the APN and P-GW address pair and sends an Update Location Response to the MME.
24	Bidirectional data is passed between the UE and PDN.

Subscriber-initiated Detach Procedure

The following figure and the text that follows describe the message flow for a user-initiated subscriber de-registration procedure.

Figure 179. Subscriber-initiated Detach Call Flow

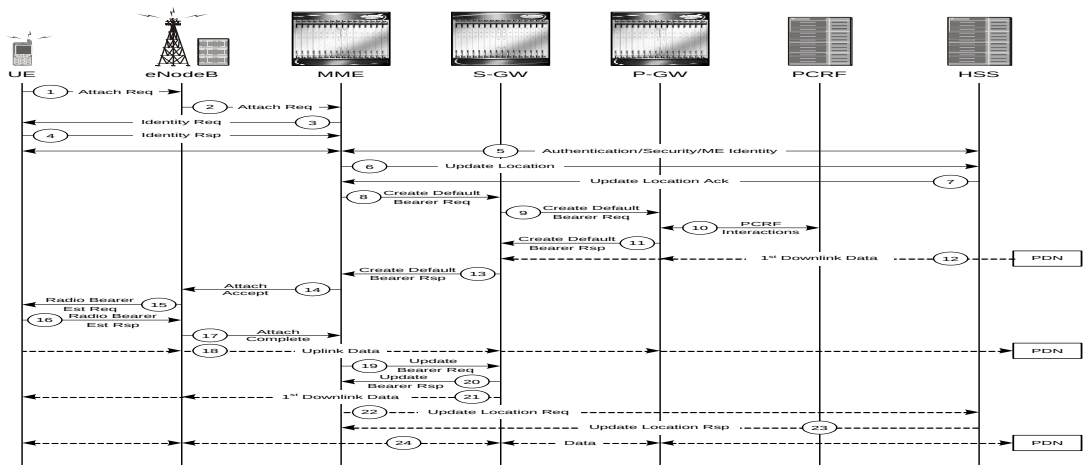


Table 80. Subscriber-initiated Detach Call Flow Description

Step	Description

Step	Description
1	The UE sends NAS message Detach Request (GUTI, Switch Off) to the MME. Switch Off indicates whether detach is due to a switch off situation or not.
2	The active EPS Bearers in the S-GW regarding this particular UE are deactivated by the MME sending a Delete Bearer Request (TEID) message to the S-GW.
3	The S-GW sends a Delete Bearer Request (TEID) message to the P-GW.
4	The P-GW acknowledges with a Delete Bearer Response (TEID) message.
5	The P-GW may interact with the PCRF to indicate to the PCRF that EPS Bearer is released if PCRF is applied in the network.
6	The S-GW acknowledges with a Delete Bearer Response (TEID) message.
7	If Switch Off indicates that the detach is not due to a switch off situation, the MME sends a Detach Accept message to the UE.
8	The MME releases the S1-MME signalling connection for the UE by sending an S1 Release command to the eNodeB with Cause = Detach.

Service Request Procedures

Service Request procedures are used to establish a secure connection to the MME as well as request resource reservation for active contexts. The MME allows configuration of the following service request procedures:

- UE-initiated Service Request Procedure
- Network-initiated Service Request Procedure

UE-initiated Service Request Procedure

The call flow in this section describes the process for re-connecting an idle UE.

The following figure and the text that follows describe the message flow for a successful UE-initiated service request procedure.

Figure 180. UE-initiated Service Request Message Flow

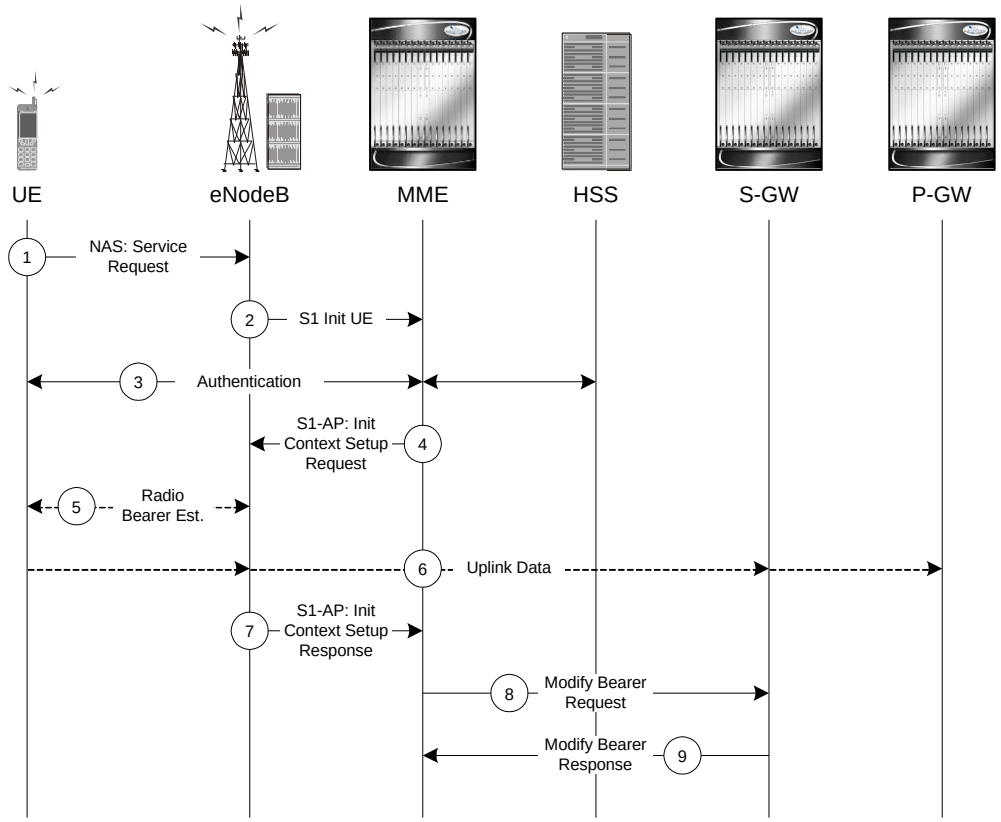


Table 81. UE-initiated Service Request Message Flow Description

Step	Description
1	(NAS) The UE sends a Network Access Signaling (NAS) message Service Request (S-TMSI) towards the MME encapsulated in an RRC message to the eNodeB.
2	The eNodeB forwards NAS message to the MME. The NAS message is encapsulated in an S1-AP: Initial UE message (NAS message, TAI+ECGI of the serving cell).
3	NAS authentication procedures may be performed.
4	The MME sends an S1-AP Initial Context Setup Request (S-GW address, S1-TEID(s) (UL), EPS Bearer QoS(s), Security Context, MME Signalling Connection Id, Handover Restriction List) message to the eNodeB. This step activates the radio and S1 bearers for all the active EPS Bearers. The eNodeB stores the Security Context, MME Signalling Connection Id, EPS Bearer QoS(s) and S1-TEID(s) in the UE RAN context.
5	The eNodeB performs the radio bearer establishment procedure.
6	The uplink data from the UE can now be forwarded by eNodeB to the S-GW. The eNodeB sends the uplink data to the S-GW address and TEID provided in step 4.
7	The eNodeB sends an S1-AP message Initial Context Setup Complete message (eNodeB address, List of accepted EPS bearers, List of rejected EPS bearers, S1 TEID(s) (DL)) to the MME.

Step	Description
8	The MME sends a Modify Bearer Request message (eNodeB address, S1 TEID(s) (DL) for the accepted EPS bearers, RAT Type) to the S-GW. The S-GW is now able to transmit downlink data towards the UE.
9	The S-GW sends a Modify Bearer Response message to the MME.

Network-initiated Service Request Procedure

The call flow in this section describes the process for re-connecting an idle UE when a downlink data packet is received from the PDN.

The following figure and the text that follows describe the message flow for a successful network-initiated service request procedure:

Figure 181. Network-initiated Service Request Message Flow

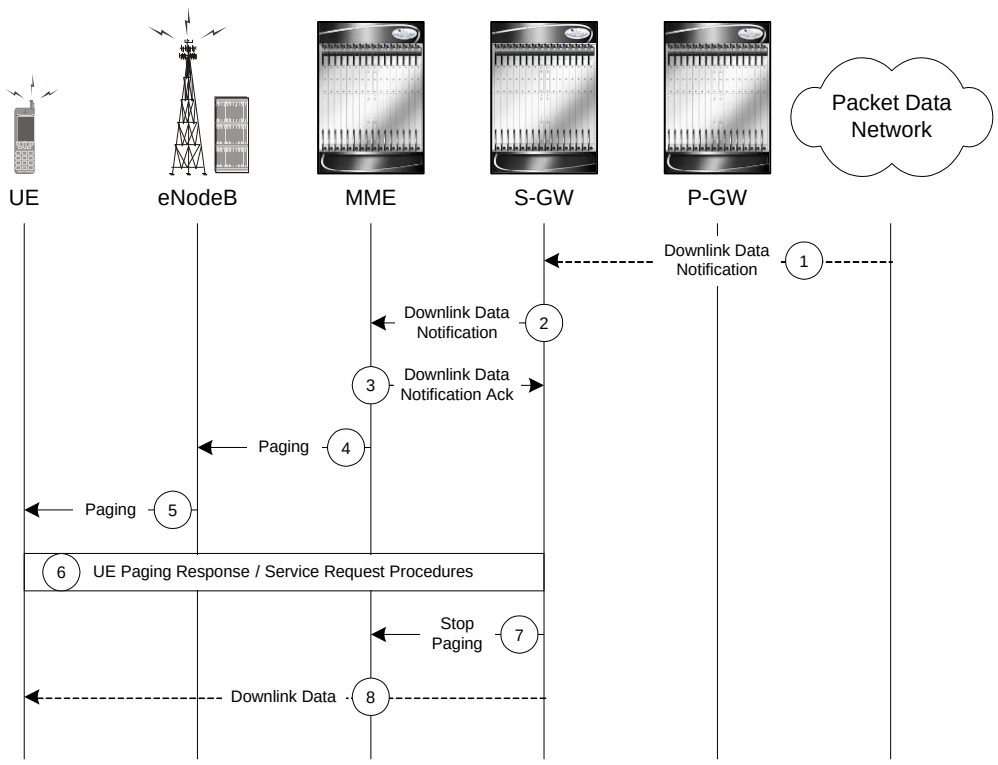


Table 82. Network-initiated Service Request Message Flow Description

Step	Description
1	A downlink data packet is received on the S-GW from PDN for the targeted UE. The S-GW checks to see if the UE is user-plane connected (the S-GW context data indicates that there is no downlink user plane (TEID)). The downlink data is buffered and the S-GW identifies which MME is serving the intended UE.
2	The S-GW sends a Downlink Data Notification message to the MME for the targeted UE.

Step	Description
3	The MME responds with a Downlink Data Notification Acknowledgement message to the S-GW.
4	The MME send a Paging Request to the eNodeB for the targeted UE. The Paging Request contains the NAS ID for paging, TAI(s), the UE identity based DRX index, and the Paging DRX length. The Paging Request is sent to each eNodeB belonging to the tracking area(s) where the UE is registered.
5	The eNodeB broadcasts the Paging Request in its coverage area for the UE. ***Important:*** Steps 4 and 5 are skipped if the MME has a signalling connection over the S1-MME towards the UE.
6	Upon receipt of the Paging indication in the E-UTRAN access network, the UE initiates the UE-triggered Service Request procedure and the eNodeB starts messaging through the UE Paging Response. The MME supervises the paging procedure with a timer. If the MME receives no Paging Response from the UE, it retransmits the Paging Request. If the MME receives no response from the UE after the retransmission, it uses the Downlink Data Notification Reject message to notify the S-GW about the paging failure.
7	The S-GW sends a Stop Paging message to MME.
8	The buffered downlink data is sent to the identified UE.

Supported Standards

The MME complies with the following standards for 3GPP LTE/EPS wireless networks.

- 3GPP References
- IETF References
- Object Management Group (OMG) Standards

3GPP References

IETF References

- RFC-768, User Datagram Protocol (UPD), August 1980
- RFC-791, Internet Protocol (IP), September 1982
- RFC-793, Transmission Control Protocol (TCP), September 1981
- RFC-894, A Standard for the Transmission of IP Datagrams over Ethernet Networks, April 1984
- RFC-1089, SNMP over Ethernet, February 1989
- RFC-1144, Compressing TCP/IP headers for low-speed serial links, February 1990
- RFC-1155, Structure & identification of management information for TCP/IP-based internets, May 1990
- RFC-1157, Simple Network Management Protocol (SNMP) Version 1, May 1990
- RFC-1212, Concise MIB Definitions, March 1991
- RFC-1213, Management Information Base for Network Management of TCP/IP-based Internets: MIB-II, March 1991
- RFC-1215, A Convention for Defining Traps for use with the SNMP, March 1991
- RFC-1224, Techniques for managing asynchronously generated alerts, May 1991
- RFC-1256, ICMP Router Discovery Messages, September 1991
- RFC-1305, Network Time Protocol (Version 3) Specification, Implementation and Analysis, March 1992
- RFC-1332, The PPP Internet Protocol Control Protocol (IPCP), May 1992
- RFC-1398, Definitions of Managed Objects for the Ethernet-Like Interface Types, January 1993
- RFC-1418, SNMP over OSI, March 1993
- RFC-1570, PPP LCP Extensions, January 1994
- RFC-1643, Definitions of Managed Objects for the Ethernet-like Interface Types, July 1994
- RFC-1701, Generic Routing Encapsulation (GRE), October 1994
- RFC-1850, OSPF Version 2 Management Information Base, November 1995
- RFC-1901, Introduction to Community-based SNMPv2, January 1996
- RFC-1902, Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996

- RFC-1903, Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1904, Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1905, Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1906, Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1907, Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2), January 1996
- RFC-1908, Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework, January 1996
- RFC-1918, Address Allocation for Private Internets, February 1996
- RFC-1919, Classical versus Transparent IP Proxies, March 1996
- RFC-2002, IP Mobility Support, May 1995
- RFC-2003, IP Encapsulation within IP, October 1996
- RFC-2004, Minimal Encapsulation within IP, October 1996
- RFC-2005, Applicability Statement for IP Mobility Support, October 1996
- RFC-2118, Microsoft Point-to-Point Compression (MPPC) Protocol, March 1997
- RFC 2131, Dynamic Host Configuration Protocol
- RFC-2136, Dynamic Updates in the Domain Name System (DNS UPDATE)
- RFC-2211, Specification of the Controlled-Load Network Element Service
- RFC-2246, The Transport Layer Security (TLS) Protocol Version 1.0, January 1999
- RFC-2328, OSPF Version 2, April 1998
- RFC-2344, Reverse Tunneling for Mobile IP, May 1998
- RFC-2394, IP Payload Compression Using DEFLATE, December 1998
- RFC 2401, Security Architecture for the Internet Protocol
- RFC 2402, IP Authentication Header (AH)
- RFC 2406, IP Encapsulating Security Payload (ESP)
- RFC 2409, The Internet Key Exchange (IKE)
- RFC-2460, Internet Protocol Version 6 (IPv6)
- RFC-2461, Neighbor Discovery for IPv6
- RFC-2462, IPv6 Stateless Address Autoconfiguration
- RFC-2486, The Network Access Identifier (NAI), January 1999
- RFC-2571, An Architecture for Describing SNMP Management Frameworks, April 1999
- RFC-2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP), April 1999
- RFC-2573, SNMP Applications, April 1999

- RFC-2574, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3), April 1999
- RFC-2597, Assured Forwarding PHB Group, June 1999
- RFC-2598, Expedited Forwarding PHB, June 1999
- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2661, Layer Two Tunneling Protocol "L2TP", August 1999
- RFC-2697, A Single Rate Three Color Marker, September 1999
- RFC-2698, A Two Rate Three Color Marker, September 1999
- RFC-2784, Generic Routing Encapsulation (GRE) - March 2000, IETF
- RFC-2794, Mobile IP Network Access Identifier Extension for IPv4, March 2000
- RFC-2809, Implementation of L2TP Compulsory Tunneling via RADIUS, April 2000
- RFC-2845, Secret Key Transaction Authentication for DNS (TSIG), May 2000
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000
- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000
- RFC-3007, Secure Domain Name System (DNS) Dynamic Update, November 2000
- RFC-3012, Mobile IPv4 Challenge/Response Extensions, November 2000
- RFC-3056, Connection of IPv6 Domains via IPv4 Clouds, February 2001
- RFC-3101 OSPF-NSSA Option, January 2003
- RFC-3143, Known HTTP Proxy/Caching Problems, June 2001
- RFC-3193, Securing L2TP using IPSEC, November 2001
- RFC-3314, Recommendations for IPv6 in Third Generation Partnership Project (3GPP) Standards, September 2002
- RFC-3316, Internet Protocol Version 6 (IPv6) for Some Second and Third Generation Cellular Hosts, April 2003
- RFC-3706, A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers, February 2004
- RFC-3543, Registration Revocation in Mobile IPv4, August 2003
- RFC 3588, Diameter Base Protocol, September 2003
- RFC 4006, Diameter Credit-Control Application, August 2005
- Draft, Route Optimization in Mobile IP
- Draft, Generalized Key Distribution Extensions for Mobile IP
- Draft, AAA Keys for Mobile IP

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

Chapter 24
Network Address Translation Overview

This chapter provides an overview of Network Address Translation (NAT) in-line service feature.

The following topics are covered in this chapter:

- NAT Overview
- How NAT Works

NAT Overview

This section provides an overview of the NAT in-line service feature.

NAT translates non-routable private IP address(es) to routable public IP address(es) from a pool of public IP addresses that have been designated for NAT. This enables to conserve on the number of public IP addresses required to communicate with external networks, and ensures security as the IP address scheme for the internal network is masked from external hosts, and each outgoing and incoming packet goes through the translation process.

The NAT in-line service works in conjunction with the following products:

- GGSN
- HA
- PDSN
- P-GW

NAT works by inspecting both incoming and outgoing IP datagrams and, as needed, modifying the source IP address and port number in the IP header to reflect the configured NAT address mapping for outgoing datagrams. The reverse NAT translation is applied to incoming datagrams.

NAT can be used to perform address translation for simple IP and mobile IP. NAT can be selectively applied/denied to different flows (5-tuple connections) originating from subscribers based on the flows' L3/L4 characteristics—Source-IP, Source-Port, Destination-IP, Destination-Port, and Protocol.

Important: NAT works only on flows originating internally. Bi-directional NAT is not supported.

Important: NAT is supported only for TCP, UDP, and ICMP flows. For other flows NAT is bypassed. For GRE flows, NAT is supported only if the PPTP ALG is configured. For more information on ALGs, please refer to the NAT Application Level Gateway section.

Important: If a subscriber is assigned with a public IP address, NAT is not applied.

Important: To get NATed, the private IP addresses assigned to subscribers must be from the following ranges: Class A 10.0.0.0 – 10.255.255.255, Class B 172.16.0.0 – 172.31.255.255, and Class C 192.168.0.0 – 192.168.255.255

NAT supports the following mappings:

- One-to-One: In one-to-one NAT each private IP address is mapped to a unique public NAT IP address. The private source ports do not change.

  When a private IP address (IP1:port1) is mapped to a public IP address (IP2:port1), any packets from IP1:port1 will be sent as though via IP2:port1. The external host can only send packets to IP2:port1, which are translated to IP1:port1. The NAT port number will be the same as the source private port.
- Many-to-One: In many-to-one NAT, multiple private IP addresses are mapped to a single public NAT IP address. In order to distinguish between different subscribers and different connections originating from same subscriber, internal private L4 source ports are translated to pre-assigned L4 NAT ports. Ports are allocated in chunks such that each private IP address is reserved a set of ports for future use. This is also known as Network Address Port Translation (NAPT).

Once a flow is marked to use a specific NAT IP address the same NAT IP address is used for all packets originating on that flow. The NAT IP address is released only when all flows and subscribers associated with it are released.

When all NAT IP addresses are in use, and a subscriber with a private IP address fails to get a NAT IP address for a specific flow, that specific flow will not be allowed and will fail.

All downlink—inbound from external networks—IP packets that do not match one of the existing NAT bindings are discarded by the system.

Platform Requirements

The NAT in-line service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the Installation Guide for the chassis and/or contact your Cisco account representative.

License Requirements

The NAT is a licensed Cisco feature. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

NAT Realms

A NAT realm is a pool of unique public IP addresses available for translation from private source IP addresses. IP addresses in a NAT IP pool are contiguous, and assignable as a subnet or a range that constitutes less than an entire subnet. IP addresses configured in NAT IP pools within a context must not overlap. At any time, within a context, a NAT IP address must be configured in any one NAT IP pool. IP addresses can be added to a NAT IP pool as a range of IP addresses. Based on the chosen port chunk-size, the number of subscribers that can shared per IP varies.

Important: The minimum number of public IP addresses that must be allocated to each NAT IP pool must be greater than or equal to the number of Session Managers (SessMgrs) available on the system. On the ASR 5x00, it is >= 84 public IP addresses. This can be met by a range of 84 host addresses from a single Class C. The remaining space from the Class C can be used for other allocations. Each address has available its port range ~64K ports.

Up to 2000 unique "IP pools + NAT IP pools" can be configured per context. A maximum of three NAT IP pools/NAT IP pool groups can be configured in a Firewall-and-NAT policy. At any time a subscriber can be associated with a maximum of three different NAT IP pools/NAT IP pool groups and can have NATed flows on three different NAT IP addresses at the same time.

Allocation of NAT IP addresses in NAT IP pools to subscriber traffic is based on the L3/L4 characteristics—IP addresses, ports, and protocol—of the subscriber flows. It is possible to configure the system to perform or not perform NAT based on one or more L3/L4 parameters. This feature is also known as Target-based NAT. For more information, see the Target-based NAT Configuration section.

The following table lists the miminum requirements for configuring NAT pools.

Card Type	No. of SMGRs	No. of NAT Translations per SMGR	No. of Active Cards	Total no. of NAT Translations
PSC	7	317000	12	26628000 (26.628M)

Card Type	No. of SMGRs	No. of NAT Translations per SMGR	No. of Active Cards	Total no. of NAT Translations
PSC2	16	270000	12	51840000 (51.84M)
PSC3	24	400000	12	115200000(115.2M)

NAT IP pools have the following configurable parameters. These parameters are applicable to all IP addresses in a NAT IP pool.

- NAT IP Address Allocation Mode: Specifies when to allocate a NAT IP address to a subscriber; either at call setup or during data flow based on the allocation mode.
  - Not-on-demand Allocation Mode: This is the default mode. In this mode, the NAT IP address is allocated to the subscriber at call setup. If there are three NAT IP pools/NAT IP pool groups (maximum possible) configured in the subscriber's Firewall-and-NAT policy, the subscriber is allocated three NAT IP addresses, one from each NAT IP pool/NAT IP pool group based on rule matching.
  - On-demand Allocation Mode: In this mode NAT resources are assigned and allocated dynamically based on subscriber flows. The NAT IP address is allocated to the subscriber when the data traffic flows in and not at call setup.

    In case of on-demand pools, since the NAT IP address is not allocated to the subscriber at call setup, the subscriber may not have a NAT IP address allocated when the first packet is received. Until the successful allocation of a NAT IP address, based on the configuration, the packets can either be buffered or dropped. Once a free NAT IP address is available, it is allocated to the subscriber to be used for flows matching the pool.
- NAT Binding Timer: Specifies the timeout period, in seconds, to deallocate NAT resources that were allocated to subscriber flows. When a subscriber flow stops the timer starts counting down, and on expiry the NAT resources are deallocated to be made available for other subscriber flows.
  - In one-to-one allocation, for a given NAT IP address, the NAT Binding Timer starts counting down when there are no active flows using that NAT IP address. When the NAT Binding Timer expires, the NAT IP address gets deallocated.
  - In many-to-one allocation, wherein subscribers are allocated port-chunks rather than individual ports, as long as a port-chunk is allocated to a subscriber, all ports from that port-chunk are reserved for that subscriber. When all flows using ports from that port-chunk get timed out/cleared, the NAT Binding Timer starts counting down. If any new flows come up before the NAT Binding Timer expires, ports are once again allocated from that port-chunk, and the NAT Binding Timer gets cancelled. As long as there are active flows using the port-chunk it cannot be deallocated. But, if no new flows come and the NAT Binding Timer expires, the port-chunk gets deallocated. In the case of on-demand NAT, if it is the last port-chunk for the NAT IP address, on NAT Binding Timer expiry, the NAT IP address gets deallocated along with the last port-chunk.
- Maximum Users per NAT IP Address: Applicable only to many-to-one NAT IP pools. Specifies the maximum number of subscribers sharing one NAT IP address. A maximum of 2016 subscribers can be configured per NAT IP address.
- Port Chunk Size: Applicable only to many-to-one NAT IP pools. Specifies the block size of contiguous ports to be assigned to a many-to-one NAT subscriber. This number has to be divisible by 32 up to a maximum of 32,256.
- Maximum Port-chunks per User: Applicable only to many-to-one NAT IP pools. Specifies the maximum number of port-chunks allowed for an individual subscriber from the same NAT IP address. This will limit subscribers from dominating all the available ports in a many-to-one NAT IP. A maximum of 2016 port-chunks can be configured per subscriber.

Consider a case where a single TCP flow is active in a port-chunk. When this connection gets cleared, the TCP NAT port goes to Time Wait state. Since it is the last flow of the port-chunk, the NAT Binding Timer also gets started. Assume NAT Binding Timer >= TCP 2MSL Timer. Once the 2MSL Timer expires, the TCP port would go to Free state. However, the NAT Binding Timer keeps running. On NAT Binding Timer expiry, the port-chunk is deallocated. If this was the last port-chunk for that subscriber, the NAT IP address is also deallocated along with this port-chunk.

In case NAT Binding Timer < TCP 2MSL Timer, at NAT Binding Timer expiry, the TCP port is forcefully moved to Free state from Time Wait state and the port-chunk deallocated.

- Port Chunk Thresholds: Applicable only to many-to-one NAT IP pools. Specifies threshold in terms of percentage of allocated port-chunks against total port-chunks available. Once the threshold is reached, new subscribers will not be allocated the same NAT IP address.
- AAA Binding Update Message Required: Applicable only to one-to-one NAT IP pools. Enables AAA binding messages for one-to-one NAT IP pools. This is not supported for many-to-one NAT IP pools.
- Alert Thresholds: Threshold limits can be specified to trigger alarms for NAT IP pools for pool-used, pool-free, pool-hold, and pool-release cases.
- SRP-Activate: Applicable to both one-to-one and many-to-one NAT IP pools. When configured, the NAT IP pool will become usable only when the SRP state is active.

NAT IP Pool Groups

Similar NAT IP pools can be grouped into NAT IP pool groups. This enables to bind discontinuous IP address blocks in individual NAT IP pools to a single NAT IP pool group.

When configuring a NAT IP pool group, note that only those NAT IP pools that have similar characteristics can be grouped together. The similarity is determined by the NAT IP pool Type (One-to-One / Many-to-One), users configured per NAT IP address (applicable only to many-to-one NAT IP pools), NAT IP Address Allocation Mode (On-demand/Not-on-demand), and Port Chunk Size (applicable only to many-to-one NAT IP pools) parameters. Dissimilar NAT IP pools cannot be grouped together.

It is recommended that all the NAT IP pools in a NAT IP pool group be configured with the same values for the other parameters, so that the NAT behavior is predictable across all NAT IP pools in that NAT IP pool group.

The NAT IP pool from which a NAT IP address is assigned will determine the actual values to use for all parameters.

It is recommended that in a Firewall-and-NAT policy all the realms configured either be NAT IP pools or NAT IP pool groups. If both NAT IP pool(s) and NAT IP pool group(s) are configured, ensure that none of the NAT IP pool(s) are also included in the NAT IP pool group.

NAT IP Address Allocation and Deallocation

Cisco System's implementation of NAPT is Endpoint-independent Mapping, wherein NAT reuses the same NAT source port mapping for subsequent packets sent from the same private IP address and port, and with the same protocol to any public destination host IP address and port.

That is, all flows coming from the subscriber for the current session with the same protocol and same source IP address and source port (X:x) would get the same NAT IP address and NAT port (X:x) irrespective of the destination IP address and port. NAT will not allow any inbound packets to the NAT IP address and NAT port (X:x) from an external host IP address and host port (Y:y), unless the internal host (MS) had previously sent a packet of the same protocol type to that external IP address and Port (Y:y). However, this behavior changes if NAT ALG is enabled. The ALG creates pin holes

/ dynamic routes in the NAT and allows downlink packets that match the pin holes / dynamic routes towards the internal host (MS) given that there was already a parent connection from MS towards the external host.

The advantage of endpoint-independent mapping is that applications are unaffected by NAT translations.

Inbound connection to the NAT IP address can be allowed in one-to-one pools based on configuration.

NAT IP Address Allocation

The NAT IP address is allocated based on the following parameters:

- Maximum Users per NAT IP Address: The maximum number of subscribers sharing a NAT IP address. Once the number of active subscribers using a NAT IP address reaches this limit, that NAT IP address will not be allocated to new subscribers.
- Port-chunk Thresholds: The threshold is configured in percentage of total number of port-chunks. If the number of port-chunks already allocated from a given NAT IP address is less than the configured threshold limit of port-chunks, then the NAT IP address can be chosen for a new subscriber provided the "Maximum Users per NAT IP Address" is not reached. But if the number of chunks allocated is greater than or equal to the threshold limit of port-chunks, then the NAT IP address will not be chosen for a new subscriber. The remaining free port-chunks will be used for existing subscribers using the NAT IP address.

NAT IP Address Deallocation

Whenever a NAT IP address is deallocated, all the port-chunks associated with the subscriber are released back to the pool.

In case there is only one port-chunk associated with the subscriber:

- In case of many-to-one not-on-demand NAT IP pools, the last port-chunk is not released back to the pool even after NAT Binding Timer expires. Only when the call gets disconnected, the port-chunk is released along with the NAT IP address.
- In case of many-to-one on-demand NAT IP pools, when the last flow using the port-chunk gets cleared, the NAT Binding Timer is started. When the NAT Binding Timer expires, the port-chunk along with the NAT IP address is released back to the pool.
- In case of one-to-one on-demand NAT IP pools, when there are no active flows using a NAT IP address, the NAT Binding Timer is started. When the NAT Binding Timer expires, the NAT IP address gets deallocated.

NAT Port-chunk Allocation and Deallocation

This section describes the Port-chunk Allocation and Deallocation feature for many-to-one NAT.

NAT Port-chunk Allocation

Subscribers sharing a NAT IP address are allocated NAT ports in chunks. The ports in a port-chunk are always used for the subscriber to whom that port-chunk is allocated irrespective of the protocol.

Whenever a NAT IP address gets allocated to a subscriber, the first port-chunk gets allocated along with the NAT IP address. Thus, for not-on-demand pools, the first port-chunk gets allocated during call setup, and for on-demand pools during data flow.

A subscriber's TCP and UDP data traffic is NATed with ports chosen in a random fashion from the port-chunk allocated to that subscriber. For other protocol traffic, the first available port is allocated. When all the ports in a port-chunk are in

use, a free port-chunk is requested for. A new port-chunk is only allocated if the "Maximum Port-chunks Per User" limit is not reached.

NAT Port-chunk Deallocation

A port-chunk gets deallocated in the following cases:

- NAT Binding Timer expiry
- Subscriber session disconnect

NAT Binding Timer

When all flows using ports from a particular port-chunk get timed out/cleared, the port-chunk gets freed. When the last port of that port-chunk gets freed, the NAT Binding Timer starts counting. Before the NAT Binding Timer expires, if any new flows come up, ports are reallocated from the port-chunk, and the timer gets cancelled. The port-chunk cannot be deallocated as long as there are active flows using that port-chunk. But, if no new flows come and the NAT Binding Timer expires, the port-chunk gets deallocated.

In case of not-on-demand pools, the additional port-chunks that were allocated on demand will be deallocated based on the NAT binding timeout. However, the last port-chunk will not be deallocated even after the Binding Timer expires. This last port-chunk will only be deallocated when the NAT IP address is deallocated from the subscriber.

In case of on-demand pools, the port-chunks are deallocated based on the NAT binding timeout. When the last port-chunk gets freed, the NAT IP address also gets deallocated from the subscriber.

It is ensured that a port-chunk is associated with the subscriber as long as a valid NAT IP address is allocated to the subscriber.

Subscriber Session Disconnect

When a subscriber disconnects, all port-chunks associated with that subscriber are freed.

If the NAT Binding Timer has not expired, the port-chunks will not be usable immediately, only on NAT Binding Timer expiry will the port-chunks become available for new subscribers.

NAT IP Address/Port Allocation Failure

When a packet cannot be translated, the application can be notified by way of ICMP error messages, if configured. Translation failures may be due to no NAT IP address or port being available for translation.

Important: In the case of P-GW, NAT IP Address/Port Allocation Failure notification is not applicable.

TCP 2MSL Timer

NAT does port management only for many-to-one pools. Hence, The TCP 2MSL timer is only available for many-to-one NAT. It is necessary to ensure that a TCP NAT port in Time Wait state is not reused if there are other free ports available for the subscriber. If such a reuse happens, then there is a possibility that connections might get terminated by the server. To avoid such issues, whenever a many-to-one NAT TCP flow gets cleared, the NAT port goes to Time Wait state (2MSL started for that port). Once 2MSL timer expires, the NAT port becomes usable. The 2MSL timer is started for every TCP NAT port as soon as the TCP connection gets cleared. This ensures that a NAT TCP port gets reused only after expiry of the configured TCP 2MSL timer.

Consider a case where a single TCP flow is active in a port-chunk. When this connection gets cleared, the TCP NAT port goes to Time Wait state. Since this is the last flow of the port-chunk, the NAT Binding Timer also gets started. Assume NAT Binding timer >= TCP 2MSL timer. Once the 2MSL timer expires, the TCP port becomes usable. However, the NAT Binding Timer keeps counting, and on expiry, the port-chunk is released.

In case the NAT Binding Timer < TCP 2MSL Timer, on NAT Binding Timer expiry, the TCP port is forcefully moved to Free State (made usable) from Time Wait state and the port-chunk released.

Flow Mapping Timer

The Flow Mapping timer is a new timer implemented as an extension to the existing idle-timeout in ECS, and is supported only for TCP and UDP flows. This flow mapping applies only for NAT enabled calls.

The purpose of this timer is to hold the resources such as NAT IP, NAT port, and Private IP NPU flow associated with a 5-tuple ECS flow until Mapping timeout expiry. If the feature is disabled, the Flow mapping timeout will not get triggered for TCP/UDP idle timed out flows. The resources such as NAT mapping will be released with the 5 -tuple flow itself.

NAT Binding Records

Whenever a NAT IP address or NAT port-chunk is allocated/deallocated to/from a subscriber, NAT Binding Records (NBR) can be generated. Generation of NBRs is configurable in the Firewall -and-NAT policy configuration.

NBRs are supported for both on-demand and not-on-demand NAT IP pools. For a one-to-one NAT IP pool, an NBR is generated whenever a NAT IP address is allocated/deallocated to/from a subscriber. For a many -to-one NAT IP pool, an NBR is generated when a port-chunk is allocated/deallocated to/from a subscriber for a NAT IP address. It is also possible to configure generation of NBRs only when a port-chunk is allocated, or deallocated, or in both cases.

The following is the list of attributes that can be present in NBRs. You can configure a subset of these attributes or all of them to be logged in NBRs. If an attribute is not available, while logging records that field is populated with NULL.

- ip subscriber-ip-address: The private IP address.
- radius-calling-station-id: The IMSI of the mobile node.
- radius-fa-nas-identifier: A string that identifies PDSN. This field is optional if PDSN-NAS-IP address field is present.
- radius-fa-nas-ip-address:
- radius-user-name: NAI of the mobile node.
- sn-correlation-id: If available. The HA-Correlation-ID identifying the entire MIP session.
- sn-fa-correlation-id: If available. The PDSN-Correlation-ID as sent by the PDSN using the same format and length.
- sn-nat-binding-timer: Optional. The NAT Binding Timer assigned to the Realm.
- sn-nat-gmt-offset: Optional. The offset from GMT to correlate timestamps of records; GMT offset of the node generating this record. For example: -5.00, +5.30
- sn-nat-ip: The NAT IP address of mobile node.
- sn-nat-last-activity-time-gmt: The time the last flow in a specific NAT set of flows was seen in GMT time.
- sn-nat-port-block-end: The NAT Port Block End of the mobile node.
- sn-nat-port-block-start: The NAT Port Block Start of the mobile node.

- sn-nat-port-chunk-alloc-dealloc-flag: 1: allocate; 0: deallocate
- sn-nat-port-chunk-alloc-time-gmt: The NAT Port Chunk Allocation Timestamp (Sample time format: 03/11/2009 10:38:35)
- sn-nat-port-chunk-dealloc-time-gmt: The NAT Port Chunk Deallocation Timestamp (Sample time format: 03/11/2009 10:38:35)
- sn-nat-realm-name: Optional. The name of the locally configured NAT Realm.
- sn-nat-subscribers-per-ip-address: Optional. NAT Multiplier assigned to the Realm.
- bearer 3gpp charging-id: The charging ID for the PDN Session.
- bearer 3gpp sgsn-address: The SGW/SGSN address.
- bearer ggsn-address: The PGW/GGSN address.
- bearer 3gpp imsi: The IMSI value of the subscriber.

Important: The NBR attributes: sn-correlation-id, sn-fa-correlation-id, radius-fa-nas-ip-address, radius-fa-nas-identifier are not applicable for PGW and GGSN.

NAT Binding Updates

Whenever a NAT IP address or NAT port-chunk is allocated/deallocated to/from a subscriber, to update NAT binding information for that subscriber in the AAA, a NAT Binding Update (NBU) can be sent to the AAA server.

Important: In this release, P-GW and GGSN do not support the NBU feature.

Since port-chunk allocation/deallocation happens on a per-call basis, this ensures that AAA messaging is reduced to a great extent. NBUs are sent to the AAA server in accounting-interim messages. To send or not to send NBUs to the AAA server is configurable in the NAT IP pool configuration.

NBUs are supported for both one-to-one and many-to-one NAT IP pools.

An NBU contains the following attributes:

- Alloc-Flag
- Binding-Timer
- Correlation-Id
- Loading-Factor
- NAT-IP-Address
- NAT-Port-Block-End: In the case of one-to-one NAT, the value is 65535
- NAT-Port-Block-Start: In the case of one-to-one NAT, the value is 1

CoA NAT Query

If the NAT binding information is not available at the AAA, the AAA server can query the chassis for the information. This query uses the Change of Authorization (CoA) format, wherein the AAA sends a one-to-one NAT IP address as a query, and in the CoA query response the NBU is obtained if available at the time of query.

Important: In this release, CoA query for NAT binding information is only supported for one-to-one NAT.

The CoA query request must contain the following attributes:

- Event-Timestamp
- NAS-IP-Address
- SN1-NAT-IP-Address

Important: For SN1-NAT-IP-Address, this release supports VSA-Type values 0 and 1.

For a successful query, the CoA ACK response contains the following attributes:

- Acct-Session-Id
- Correlation-Id
- Framed-IP-Address
- NAT-IP-Address
- NAT-Port-Block-End
- NAT-Port-Block-Start
- User-Name

Important: For more information on the AVPs/VSAs, please refer to the *AAA and GTPP Interface Administration and Reference*.

Firewall-and-NAT Policy

Firewall-and-NAT policies are configured in the CLI Firewall-and-NAT Policy Configuration Mode. Each policy contains a set of access ruledefs with priorities and actions, and the NAT configurations. On a system, multiple such policies can be configured, however at any point of time only one policy is associated to a subscriber.

Important: In release 8.x, NAT for CDMA and early UMTS releases used rulebase-based configurations, whereas in later UMTS releases NAT used policy-based configurations. In 9.0 and later releases, NAT for UMTS and CDMA releases both use policy-based configurations. For more information, please contact your local service representative.

Important: In a Firewall-and-NAT policy, a maximum of three NAT IP pools/NAT IP pool groups can be configured. A subscriber can be allocated only one NAT IP address per NAT IP pool/NAT IP pool group, hence at anytime, there can only be a maximum of three NAT IP addresses allocated to a subscriber.

New NAT IP pools/NAT IP pool groups cannot be added to a policy if the maximum allowed is already configured in it. However, a pool/pool group can be removed and then a new one added. When a pool/pool group is removed and a new one added, the pool/pool group that was removed will stay associated with the subscriber as long as the subscriber has active flows using that pool/pool group. If the subscriber is already associated with three NAT IP pools (maximum allowed), any new flows from that subscriber for the newly added pool will be dropped. A deleted pool is disassociated

from the subscriber on termination of all flows from that subscriber using that pool. The new pool/pool group is associated with the subscriber only when the subscriber sends a packet to the newly added pool.

In the Firewall-and-NAT policy configuration, the NAT policy must be enabled. Once NAT is enabled for a subscriber, the NAT IP address to be used is chosen from the NAT IP pools/NAT IP pool groups specified in matching access rules configured in the Firewall-and-NAT policy.

The Firewall-and-NAT policy used for a subscriber can be changed either from the command line interface, or through dynamic update of policy name in Diameter and RADIUS messages. In both the cases, NAT status on the active call remains unchanged.

The Firewall-and-NAT policy to be used for a subscriber can be configured in:

- ECS Rulebase: The default Firewall-and-NAT policy configured in the ECS rulebase has the least priority. If there is no policy configured in the APN/subscriber template, and/or no policy to use is received from the AAA/OCS, only then the default policy configured in the ECS rulebase is used.
- APN/Subscriber Template: The Firewall-and-NAT policy configured in the APN/subscriber template overrides the default policy configured in the ECS rulebase. To use the default policy configured in the ECS rulebase, in the APN/subscriber configuration, the command to use the default rulebase policy must be configured.
- AAA/OCS: The Firewall-and-NAT policy to be used can come from the AAA server or the OCS. If the policy comes from the AAA/OCS, it will override the policy configured in the APN/subscriber template and/or the ECS rulebase.

Important: The Firewall-and-NAT policy received from the AAA and OCS have the same priority. Whichever comes latest, either from AAA/OCS, is applied.

The Firewall-and-NAT policy to use can also be received from RADIUS during authentication.

Disabling NAT Policy

Important: By default, NAT processing for subscribers is disabled.

NAT processing for subscribers is disabled in the following cases:

- If the AAA/OCS sends the SN-Firewall-Policy AVP with the string “disable”, the locally configured Firewall-and-NAT policy does not get applied.
- If the SN-Firewall-Policy AVP is received with the string “NULL”, the existing Firewall-and-NAT policy will continue.
- If the SN-Firewall-Policy AVP is received with a name that is not configured locally, the subscriber session is terminated.

Updating Firewall-and-NAT Policy in Mid-session

The Firewall-and-NAT policy can be updated mid-session provided the policy was enabled during call setup.

Important: When the firewall AVP contains “disable” during mid-session firewall policy change, there will be no action taken as the Firewall-and-NAT policy cannot be disabled dynamically. The policy currently applied will continue.

Important: For all NAT-enabled subscribers, when the Firewall-and-NAT policy is deleted, the call is dropped.

In a Firewall-and-NAT policy, you can change the NAT enabled/disabled status at any time. However, the updated NAT status will only be applied to new calls, active calls using that Firewall-and-NAT policy will remain unaffected.

Target-based NAT Configuration

A NAT IP pool can be selected based on the L3/L4 characteristics of a subscriber's flows. NAT can be configured such that all subscriber traffic coming towards specific public IP address(es) always selects a specific NAT IP pool based on the L3/L4 traffic characteristics.

Important: A subscriber can be allocated only one NAT IP address per NAT IP pool/NAT IP pool group from a maximum of three NAT IP pools/NAT IP pool groups. Hence, at anytime, there can only be a maximum of three NAT IP addresses allocated to a subscriber.

This association is done with the help of access ruledefs configured in the Firewall-and-NAT policy. The NAT IP pool/NAT IP address to be used for a subscriber flow is decided during rule match. When packets match an access ruledef, NAT is applied using the NAT IP address allocated to the subscriber from the NAT IP pool/NAT IP pool group configured in that access ruledef.

If no NAT IP pool/NAT IP pool group name is configured in the access ruledef matching the packet, and if there is a NAT IP pool/NAT IP pool group configured for "no ruledef matches", a NAT IP address from the NAT IP pool/NAT IP pool group configured for "no ruledef matches" is allocated to the flow.

If no NAT IP pool/NAT IP pool group is configured for "no ruledef matches" and if there is a default NAT IP pool/NAT IP pool group configured in the rulebase, a NAT IP address from this default NAT IP pool/NAT IP pool group is allocated to the flow.

If a NAT IP pool/NAT IP pool group is not configured in any of the above cases, no NAT will be performed for the flow. Or, if bypass NAT is configured in a matched access rule or for "no ruledef matches" then NAT will not be applied even if the default NAT IP pool/NAT IP pool group is configured. The order of priority is:

1. Bypass NAT
2. NAT IP pool/NAT IP pool group in ruledef
3. NAT IP pool/NAT IP pool group for "no-ruledef-matches"
4. Default NAT IP pool/NAT IP pool group

When a new NAT IP pool/NAT IP pool group is added to a Firewall-and-NAT policy, it is associated with the active subscriber (call) only if that call is associated with less than three (maximum limit) NAT IP pools/NAT IP pool groups. If the subscriber is already associated with three NAT IP pools/NAT IP pool groups, any new flows referring to the newly added NAT IP pool/NAT IP pool group will get dropped. The newly added NAT IP pool/NAT IP pool group is associated to a call only when one of the previously associated NAT IP pools/NAT IP pool groups is freed from the call.

NAT Application Level Gateway

Some network applications exchange IP/port information of the host endpoints as part of the packet payload. This information is used to create new flows, by server or client.

As part of NAT ALGs, the IP/port information is extracted from the payload, and the flows are allowed dynamically (through pinholes). IP and port translations are done accordingly. However, the sender application may not be aware of these translations since these are transparent, so they insert the private IP or port in the payload as usual.

For example, FTP NAT ALG interprets "PORT" and "PASV reply" messages, and NAT translates the same in the payload so that FTP happens transparently through NAT. This payload-level translation is handled by the NAT ALG module.

The NAT module will have multiple NAT ALGs for each individual application or protocol.

Supported NAT ALGs

This release supports NAT ALGs only for the following protocols:

- H323
- File Transfer Protocol (FTP)
- Point-to-Point Tunneling Protocol (PPTP): If PPTP ALG is enabled, NAT is supported for GRE flows that are generated by PPTP.
- Real Time Streaming Protocol (RTSP)
- Session Initiation Protocol (SIP)
- Trivial File Transfer Protocol (TFTP)

For NAT ALG processing, in the rulebase, routing rules must be configured to route packets to the corresponding analyzers.

This release now supports session recovery for SIP ALG. Only one contact pinhole, and only one connected call and its associated media pinholes will be recovered for a subscriber. Any subscriptions, ongoing transactions, or unconnected calls will not be recovered. SIP ALG recovery data will be check-pointed using the variable length micro checkpointing mechanism.

H323 ALG Support

This release provides support for H323 ALG that is designed to traverse NAT by inspecting and altering information contained in existing H323 messages as they pass through the NAT. It can alter address and port information in registration, call signaling and automatically open pinholes in the NAT to allow media flow.

H323 ALG performs the following functions:

- Communicates with the core for binding management
- Uses H323 stack for parsing and encoding the H323 messages
- Communicates with NAT for signaling messages
- Performs protocol specific processing if required

The following supplementary services are currently supported in H323 ALG:

- Call Transfer: The Call Transfer supplementary service enables the served user (User A) to transform an existing call with a User B (primary call) into a new call between current User B and a new User C (transferred-to user) selected by served user A.
- Call Hold: The Call Hold supplementary service allows the served user, which may be the originally calling or the called user, to interrupt communications on an existing call and then subsequently, if desired, re-establish (i.e. retrieve) communications with the held user.

- Call Diversion: Call Diversion supplementary service permits a served user to have incoming calls addressed to the served user's number redirected to another number; on busy service, it enables a served user to have calls redirected to another endpoint; on No Answer, it enables a served user to have calls addressed to the served endpoint's number and redirected to another endpoint if the connection is not established within a defined period of time.
- Call Waiting: The Call Waiting supplementary service permits a busy user to be informed of an incoming call while being engaged with one or more other calls.
- Call Offering: The Call Offering supplementary service on request from the calling user, enables a call to be offered to a busy user and to wait for that called user to accept the call, after the necessary resources have become available.

NAT Aware H323 Clients

An application layer gateway, at the Firewall/NAT, examines all the H323 packets and modifies the packet such that all the private addresses are replaced by public addresses. It also opens all the pinholes required for successful call establishment. A NAT aware endpoint establishes end-to-end media session through FW/NAT without the need of ALG. Any TCP connection or UDP packet sent from the internal network through the firewall opens a pinhole dynamically in the firewall. This pinhole allows incoming messages to be sent from the destination of the TCP connection or the UDP packet. The pinhole stays open as long as the network sends information through the pinhole to the same destination.

If an end point supports NAT traversal, H323 ALG disables itself so that end point directly opens required pinhole and establishes media path between them. The ALG will not manage any pinhole for media traversal across Firewall/NAT for NAT aware clients. By default, the ALG will bypass all the clients that support H460.18/19 and H460.23/24.

EDRs and UDRs

This section describes the NAT-specific attributes supported in EDRs and UDRs.

EDRs

The following NAT-specific attributes are supported in regular EDRs:

- sn-nat-subscribers-per-ip-address: Subscriber(s) per NAT IP address
- sn-subscriber-nat-flow-ip: NAT IP address of NAT-enabled subscribers
- sn-subscriber-nat-flow-port: NAT port number of NAT-enabled subscribers

UDRs

The following NAT-specific attribute is supported in regular UDRs:

sn-subscriber-nat-flow-ip: NAT IP addresses that are being used by NAT-enabled subscribers. The NAT IP addresses assigned from each of the associated pool for the call are logged. A space is used as a separator between individual IP addresses.

Bulk Statistics

The NAT realms are configured in a context and statistics are stored per context per realm. These statistic variables, both cumulative and snapshot, are available in the nat-realm schema.

Bulkstats are only generated for the first 100 NAT IP pools from an alphabetical list of all NAT IP pools, which is based on the context name and pool name. Therefore, to generate bulkstats for a specific NAT IP pool it must be named such that it gets selected in the first 100 bulkstats.

The following are cumulative statistics that can be part of NAT bulkstats:

- vpnname: Context name
- realmname: Realm name
- nat-bind-updates: Total interim AAA NBU sent.
- nat-rlm-bytes-tx: Total number of bytes transferred by realm (uplink + downlink).
- nat-rlm-flows: Total number of flows used by the realm.
- nat-rlm-ip-denied: Total number of flows denied NAT IP address.
- nat-rlm-port-denied: Total number of flows denied ports.
- nat-rlm-max-port-chunk-subs: Total number of subscribers who used maximum number of port chunks.
- nat-rlm-max-port-chunk-used: Maximum port chunks used.

The following are snapshot statistics that can be part of NAT bulkstats:

- vpnname: Context name
- realmname: Realm name
- nat-rlm-ttl-ips: Total number of NAT public IP addresses, per context per NAT realm. Is a static value.
- nat-rlm-ips-in-use: Total number of NAT IP addresses currently in use, per context per NAT realm.
- nat-rlm-current-users: Total number of subscribers currently using the NAT realm.
- nat-rlm-ttl-port-chunks: Total number port-chunks, per context per NAT realm. Is a static value.
- nat-rlm-chunks-in-use: Total number of port-chunks currently in use, per context per NAT realm.
- nat-rlm-max-cur-port-chunk-subs: Current number of subscribers using maximum number of port chunks.
- nat-rlm-max-cur-port-chunk-used: Maximum port chunks used by active subscribers.
- nat-rlm-port-chunk-size: Size of the port chunk in the NAT realm.
- nat-rlm-port-chunk-average-usage-tcp: Average TCP port usage in the allocated TCP ports, i.e. out of allocated TCP ports how many got used. Not percentage value.
- nat-rlm-port-chunk-average-usage-udp: Average UDP port usage in the allocated UDP ports, i.e. out of allocated UDP ports how many got used. Not percentage value.
- nat-rlm-port-chunk-average-usage-others: Average other (ICMP or GRE) port usage in the allocated other ports, i.e. out of allocated 'other' ports how many got used. Not percentage value.

Alarms

Alert threshold values can be specified to generate alarms for NAT IP pools. To specify realm-specific threshold limits (pool-used, pool-free, pool-release, and pool-hold) "alert-threshold" NAT IP pool parameter can be used, or it can also be specified across context. These thresholds can be specified to any number of NAT IP pools.

In case of many-to-one NAT, it is possible to specify port-chunks usage threshold per NAT IP pool. This threshold value is applicable to all many-to-one NAT IP pools across the system. However, note that alarms are only generated for the first 100 many-to-one NAT IP pools from an alphabetical list of all NAT IP pools.

Session Recovery and ICSR

In session recovery, as part of the Private IP assigned to the subscriber:

- The public IP address used for the subscriber is recovered. The NAT IP address being used by the subscriber can be on-demand or not-on-demand. In case of many-to-one NAT, the port-chunks associated with the NAT IP address for the subscriber needs to check-pointed as well.
- In case Bypass NAT feature is used, then the private IP flow needs to be recovered.

To be recovered the NAT IP addresses need to be checkpointed. The checkpointing can be:

- Full Checkpoint
- Micro Checkpoint

To recover the bypass NAT flow, the bypass flow needs to be checkpointed. The checkpointing of Bypass NAT flow can be:

- Full Checkpoint
- Micro Checkpoint

In case of not-on-demand, the NAT IP address being used by the subscriber is known after call setup. This gets checkpointed as part of the normal full checkpoint. In case of on-demand NAT, the NAT IP address being used by the subscriber is known only in the data-path. This will be checkpointed as part of micro checkpoint.

In case of many-to-one NAT, the port-chunks being used will always be checkpointed as part of micro checkpoint.

In case of bypass NAT flow, in most cases the flow gets checkpointed as part of micro checkpoint.

Any information that is checkpointed as part of full checkpoint is always recovered. Data checkpointed through micro checkpoint cannot be guaranteed to be recovered. The timing of switchover plays a role for recovery of data done through micro checkpoint. If failover happens after micro checkpoint is completed, then the micro checkpointed data will get recovered. If failover happens during micro checkpoint, then the data recovered will be the one obtained from full checkpoint.

Once NAT IP/and Port-Chunks/Bypass NAT flow are recovered, the following holds good:

- One-to-one NAT: Since NAT IP address being used for one-to-one NAT is recovered, on-going flows will be recovered as part of Firewall Flow Recovery algorithm as one-to-one NAT does not change the port.
- Many-to-one NAT: On-going flows will not be recovered as the port numbers being used for flows across chassis peers/SessMgr peers are not preserved.

  It is now possible to enable/disable the checkpointing of NATed flows and control the type of flows to be checkpointed based on criteria. Check pointing is done only for TCP and UDP flows.
- Bypass NAT Flow: On-going flows will be recovered as part of Firewall Flow Recovery algorithm.

All of the above items is applicable for ICSR as well.

Category	Event	Impacted	Details
One-to-One NAT	Session	No	Session recovered.
	New Traffic	No	NAT will be applied.

Category	Event	Impacted	Details
	Ongoing Traffic	Yes	Cannot differentiate between ongoing traffic and unsolicited traffic. A rule-match is done and if allowed, NAT will be applied accordingly on the packet.
	Unsolicited Traffic (downlink packets)	Yes	Cannot differentiate between ongoing traffic and unsolicited traffic. Translation will be done and packet action taken based on the rule-match.
Many-to-One NAT	Session	No	Session recovered.
	New Traffic	No	NAT will be applied.
	Ongoing Traffic — TCP	Yes	Packet will be dropped.
	Ongoing Traffic — UDP	Yes and No	If it is downlink packet, it will be dropped. If it is uplink packet, NAT will be applied with a new port.
	Ongoing Traffic — ICMP	Yes and No	If it is downlink packet, it will be dropped. If it is uplink packet, NAT will be applied with a new port.
	Unsolicited Traffic (downlink packets)	No	Packet will be dropped.
Bypass NAT	Session	No	Session recovered.
	New Traffic	No	Traffic will be NAT bypassed.
	Ongoing Traffic	No	Traffic will be NAT bypassed.
	Unsolicited Traffic (downlink packets)	No	Traffic will be NAT bypassed.

For more information, in the *System Administration Guide*, see the *Session Recovery* and *Interchassis Session Recovery* chapters.

Supported Standards

The NAT feature supports the following RFCs:

- RFC 1631: The IP Network Address Translator (NAT); May 1994
- RFC 1918: Address Allocation for Private Internets; February 1996
- RFC 2663: IP Network Address Translator (NAT) Terminology and Considerations; August 1999
- RFC 2765: Stateless IP/ICMP Translation Algorithm (SIIT); February 2000
- RFC 2766: Network Address Translation - Protocol Translation (NAT-PT); February 2000
- RFC 3022: Traditional IP Network Address Translator (Traditional NAT); January 2001
- RFC 3027: Protocol Complications with the IP Network Address Translator; January 2001
- RFC 4787: Network Address Translation (NAT) Behavioral Requirements for Unicast UDP; January 2007
- RFC 4966: Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status; July 2007
- RFC draft-nishitani-cgn-00.txt: Carrier Grade Network Address Translator (NAT) Behavioral Requirements for Unicast UDP, TCP and ICMP; July 2, 2008

How NAT Works

The following steps describe how NAT works:

Step 1 In the subscriber profile received from the AAA Manager, the SessMgr checks for the following:

- Enhanced Charging Service subsystem must be enabled
- In the Firewall-and-NAT policy, NAT must be enabled
- The Firewall-and-NAT policy must be valid
- For Many-to-One NAT, at least one valid NAT IP pool must be configured in the Firewall-and-NAT policy, and that NAT IP pool must be configured in the context

Step 2 If all of the above is true, once a private IP address is allocated to the subscriber, the NAT resource to be used for the subscriber is determined. This is only applicable for not-on-demand allocation mode.

Important: The private IP addresses assigned to subscribers must be from the following ranges for them to get translated: Class A 10.0.0.0 – 10.255.255.255, Class B 172.16.0.0 – 172.31.255.255, and Class C 192.168.0.0 – 192.168.255.255

Important: A subscriber can be allocated only one NAT IP address per NAT IP pool/NAT IP pool group from a maximum of three pools/pool groups. Hence, at any point, there can be a maximum of three NAT IP addresses allocated to a subscriber.

Step 3 Flow setup is based on the NAT mapping configured for the subscriber:

- In case of one-to-one NAT mapping, the subscriber IP address is mapped to a public IP address. The private source ports do not change. The SessMgr installs a flow using the NAT IP address and a fixed port range (1 – 65535).
- In case of many-to-one NAT mapping, a NAT IP address and a port from a port-chunk, are allocated for each connection originating from the subscriber. In order to identify a particular subscriber call line, the SessMgr installs a flow using NAT (public) IP address + NAT ports allocated for the subscriber.

The following figures illustrate the flow of packets in NAT processing.

Figure 182. NAT Processing Flow

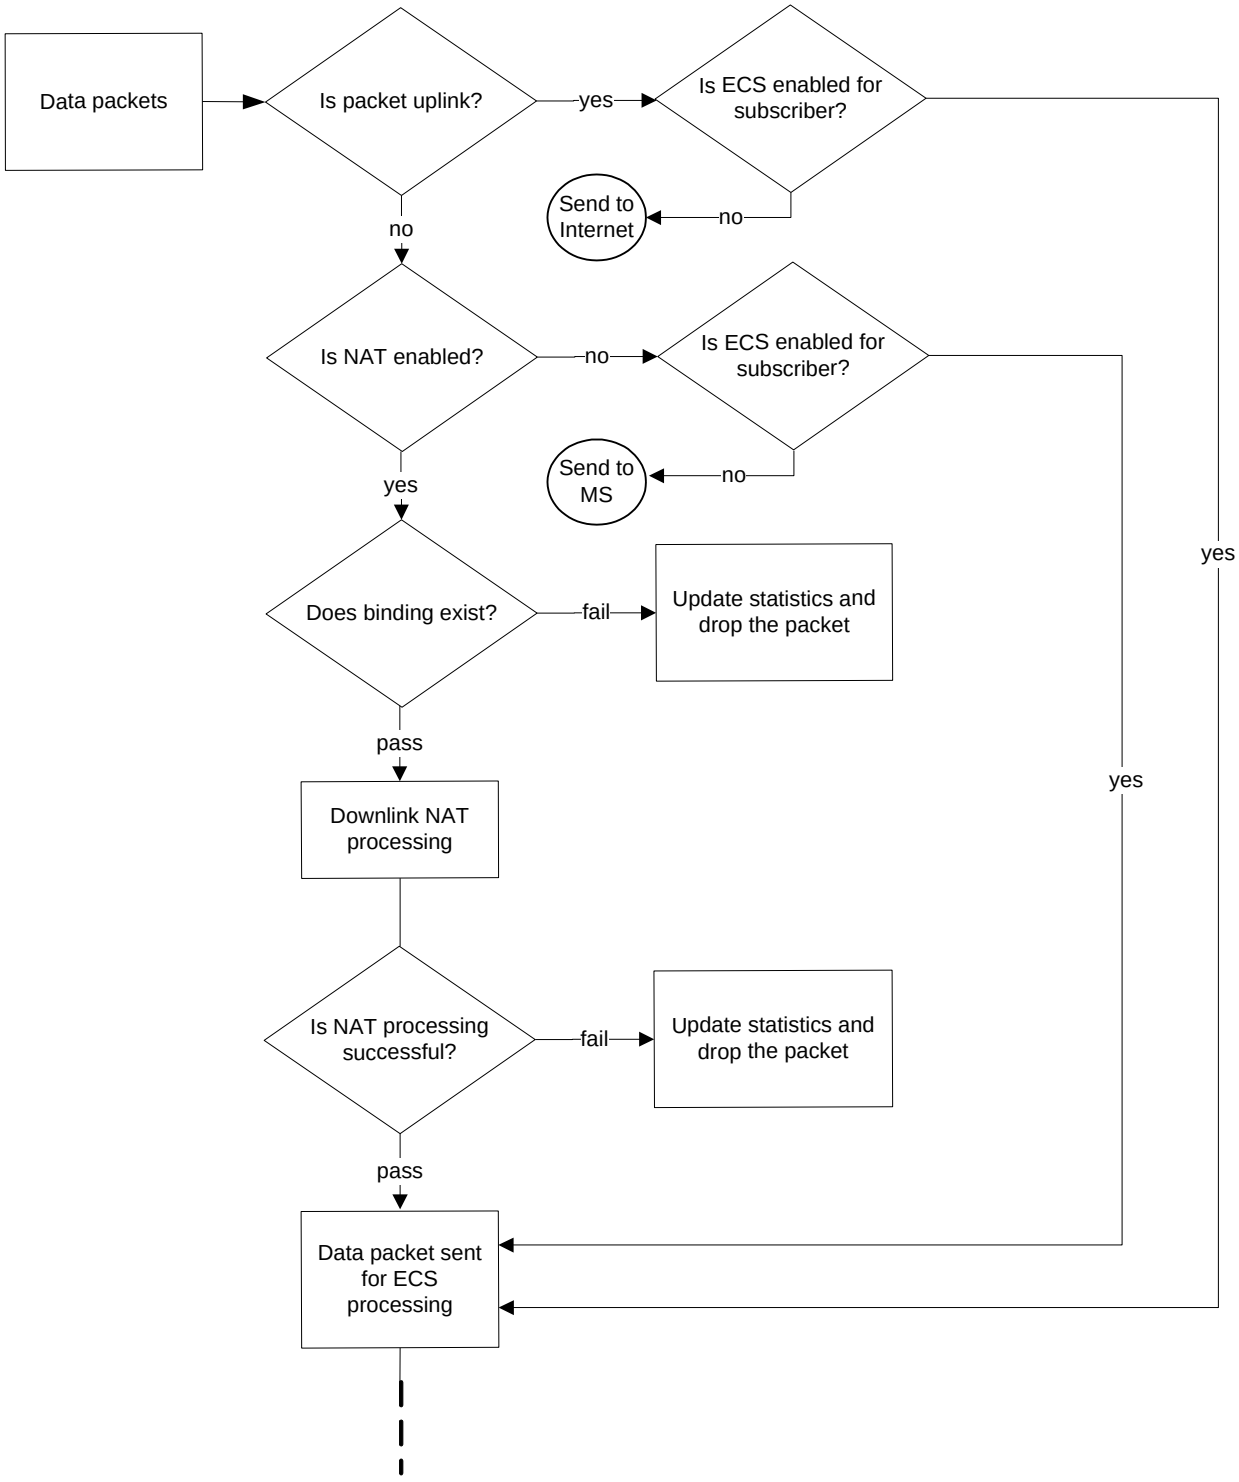


Figure 183. ... NAT Processing Flow

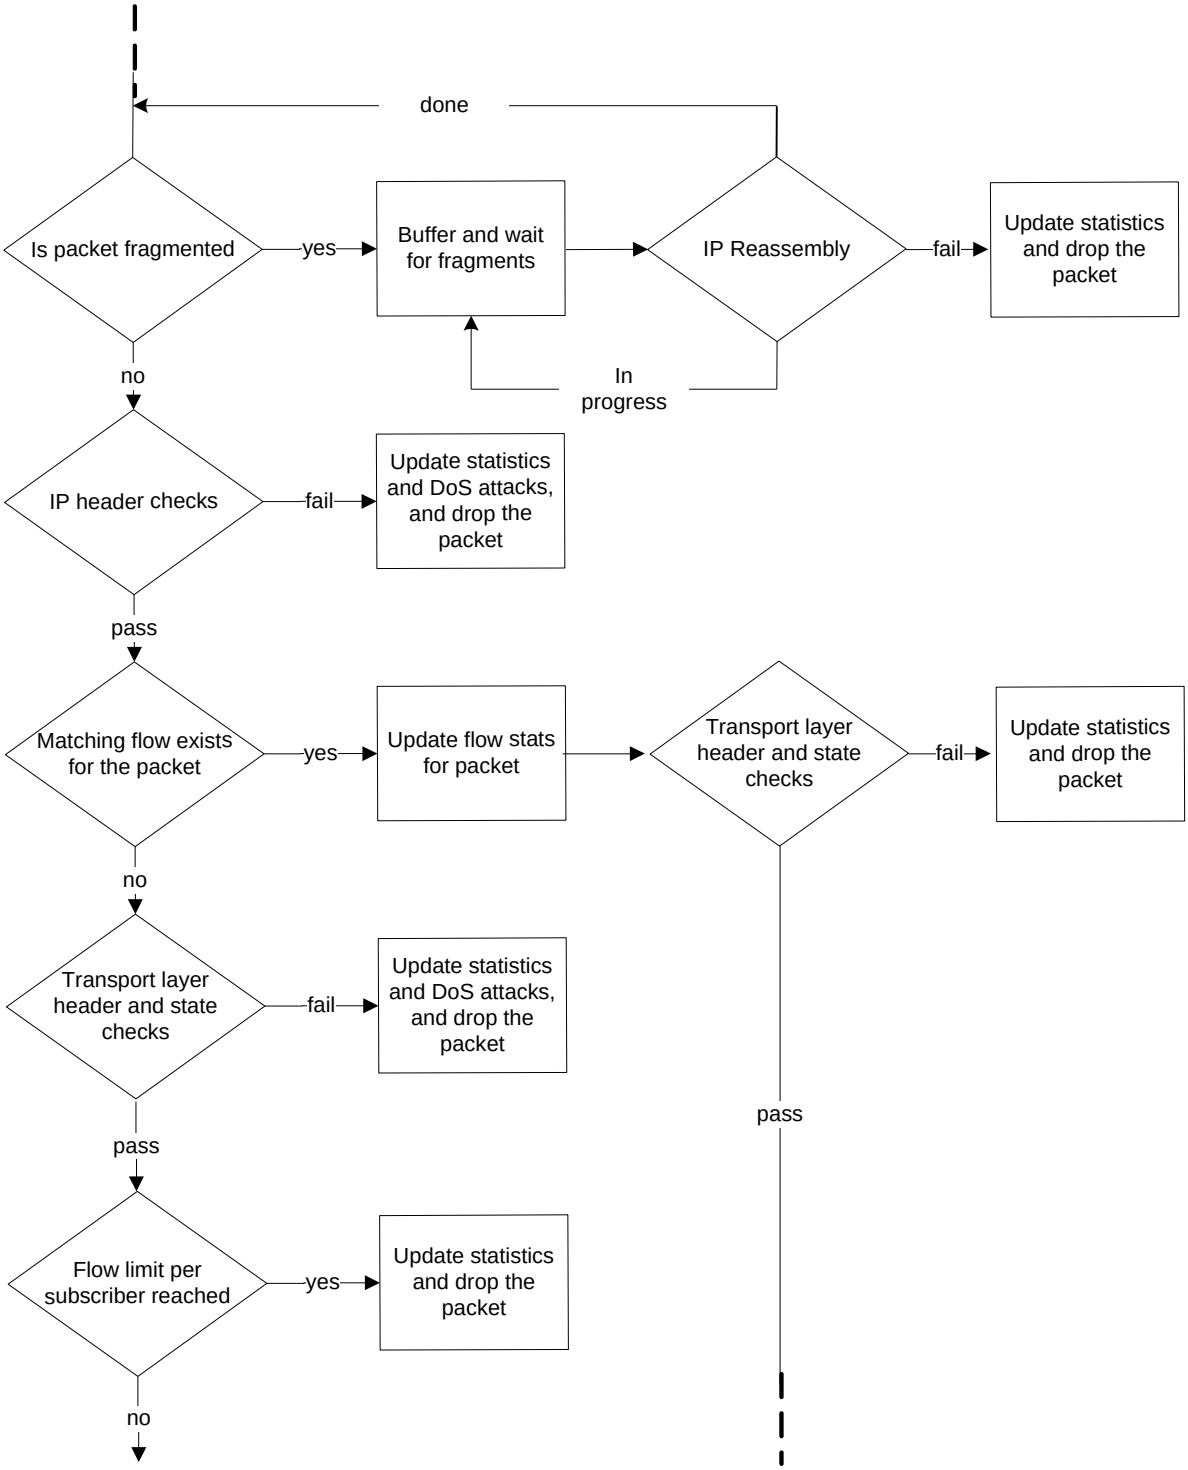


Figure 184. ... NAT Processing Flow

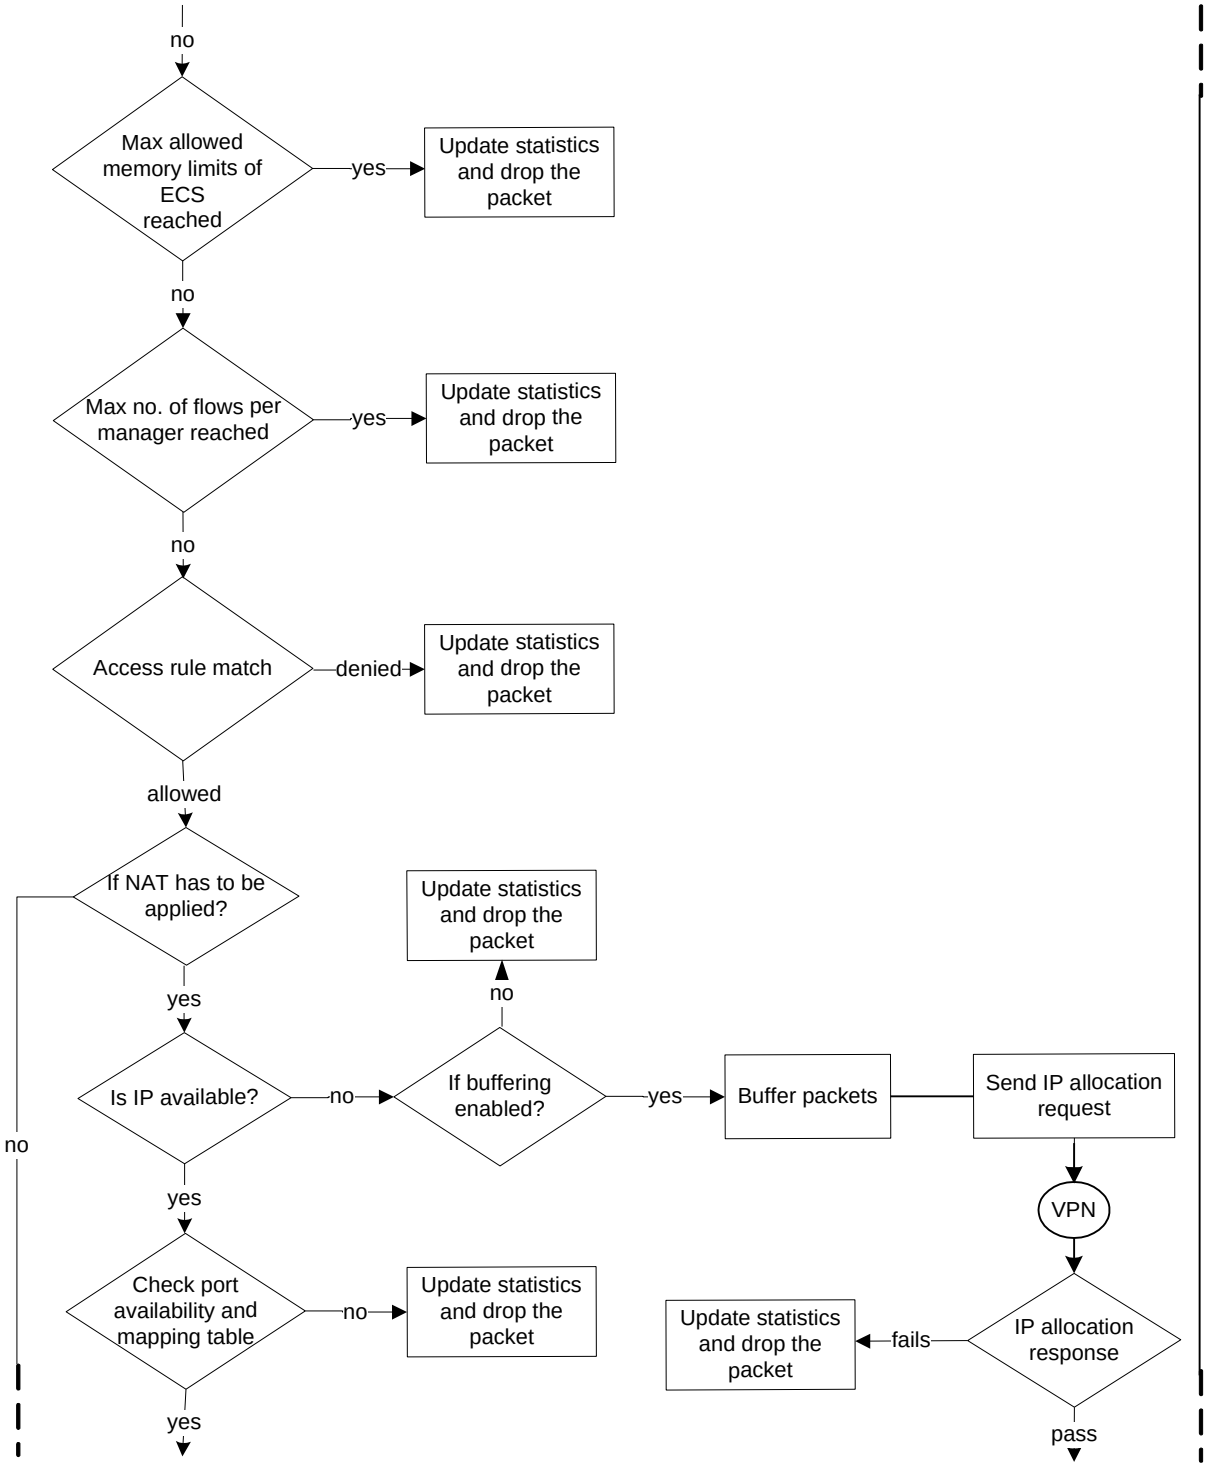
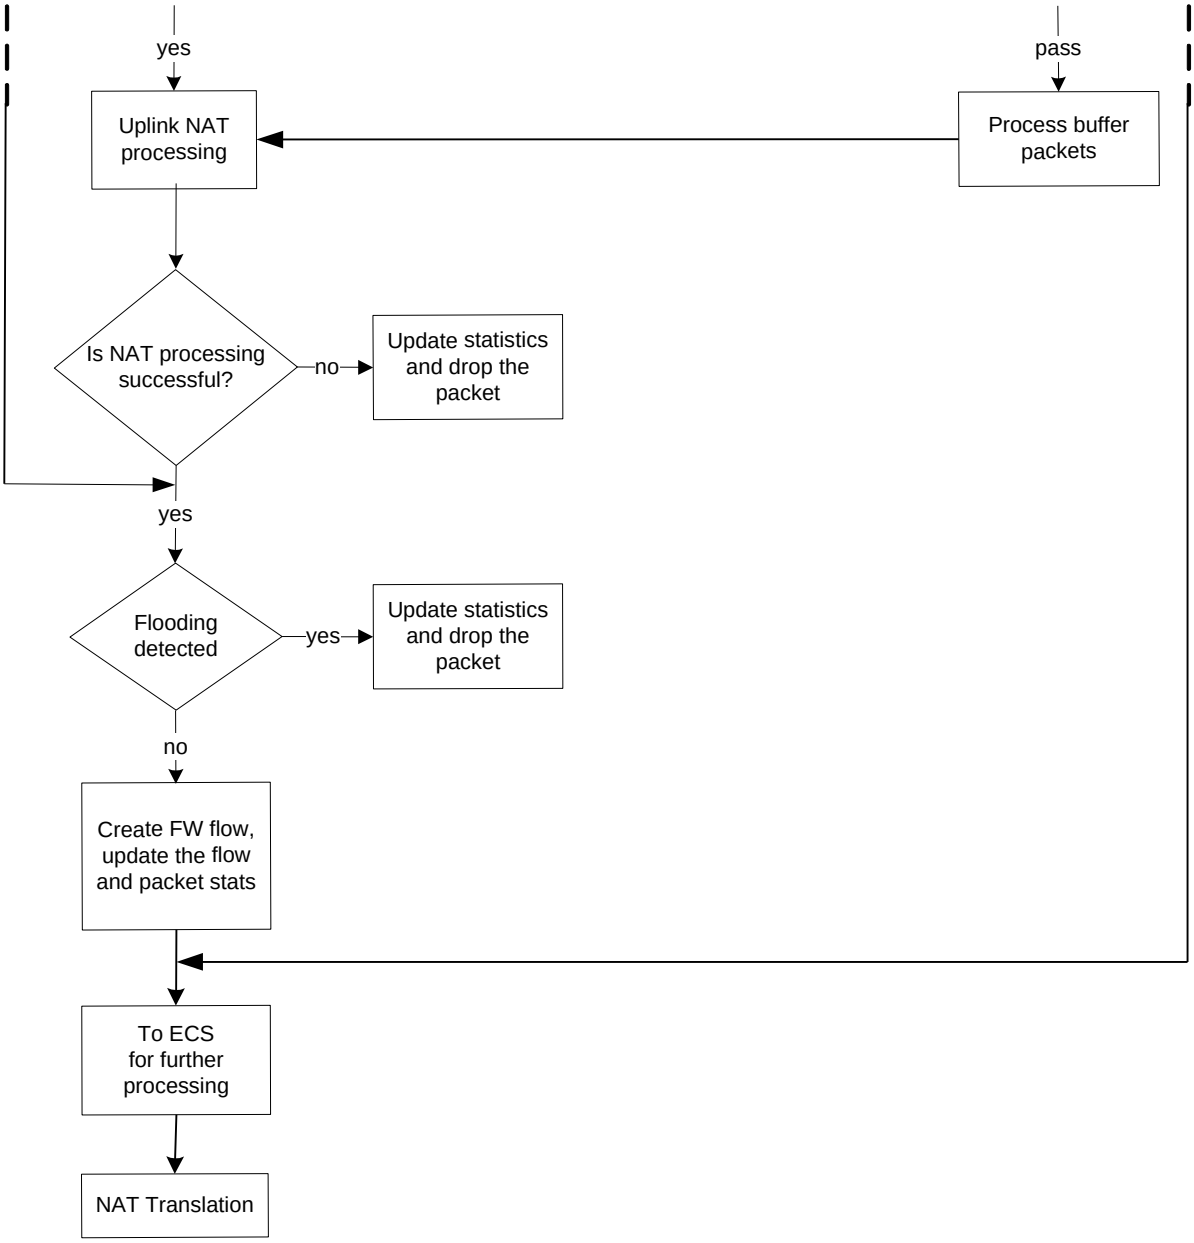


Figure 185. ... NAT Processing Flow



Chapter 25
Packet Data Interworking Function Overview

This chapter discusses the features and functions of Packet Data Interworking Function (PDIF) software. It includes the following topics:

- Product Description
- Interfaces
- Sample Deployments
- Features and Functionality - Base Software
- Features and Functionality - Licensed Enhanced Feature Support
- Supported Standards and RFCs

Product Description

The goal of the Fixed Mobile Convergence (FMC) application is to enhance the in-building cellular coverage for FMC subscribers, to reduce the cost of the infrastructure required to carry these calls, and to provide secure access to the carrier's network from a non-secure network. Designed for use exclusively on the Cisco® ASR 5x00 Chassis, the Packet Data Interworking Function (PDIF) is a network function based on the 3GPP2 X.S0028-200 standard defining CDMA2000 Packet Data Services over an 802.11 WLAN.

A PDIF allows mobile devices to access the Internet over an all-IP WLAN using IKEv2 as the signaling interface. The IKEv2 control path exists between the mobile station (MS) (a dual-mode handset (DMH)) and the PDIF establishing an IPSec tunnel. PDIF also acts as a security gateway protecting CDMA network resources and data. The PDIF is tightly integrated with a collocated Foreign Agent (FA) service, and the PDIF is known throughout this manual as PDIF/FA.

For handsets that do not support mobile IP, PDIF supports proxy mobile IP. If the MS is not suitable for proxy mobile IP registration, it may still be allowed to establish a simple IP session, in which case the traffic is directly routed to the Internet or corporate network from the PDIF. This behavior is controlled through the **`proxy-mip-required`** configuration in the domain, local default subscriber, or the corresponding Diameter AVP or RADIUS Access Accept. If this is not present, establishing a simple IP session is permitted. Although not required for Proxy-MIP, this manual documents Proxy-MIP with a custom-designed feature called multiple authentication (Multi-Auth). Instead of the more usual subscriber authentication, Multi-Auth requires both the device and the subscriber be authenticated using EAP/AKA authentication for the first stage (the device authentication) and GTC/MD5 for the second stage (the subscriber authentication). For this installation, neither GTC nor MD5 is supported, which means authentication is done using PAP/CHAP instead.

When the subscriber is mobile, the MS operates as a normal mobile phone, sending voice and data over the CDMA network. When the FMC subscriber returns home, or encounters a WiFi hotspot, the MS detects the presence of the WiFi network, and automatically establishes an IPSec session with the PDIF/FA. When the secure connection has been established and mobile IP registration procedures successfully finished, the PDIF/FA works with other network elements to provide the MS with access to packet data services.

From here, all voice and data communication is carried over the IPSec tunnel and the PDIF/FA functions as a pass-through for the authentication and accounting information on a RADIUS and/or Diameter server. The MS continues operating over the IPSec tunnel until such time as it can no longer access the WiFi Access Point (AP). At this point, the MS switches back to the CDMA network for normal mobile operation.

Platform Requirements

The PDIF service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the installation guide for the chassis and/or contact your Cisco account representative.

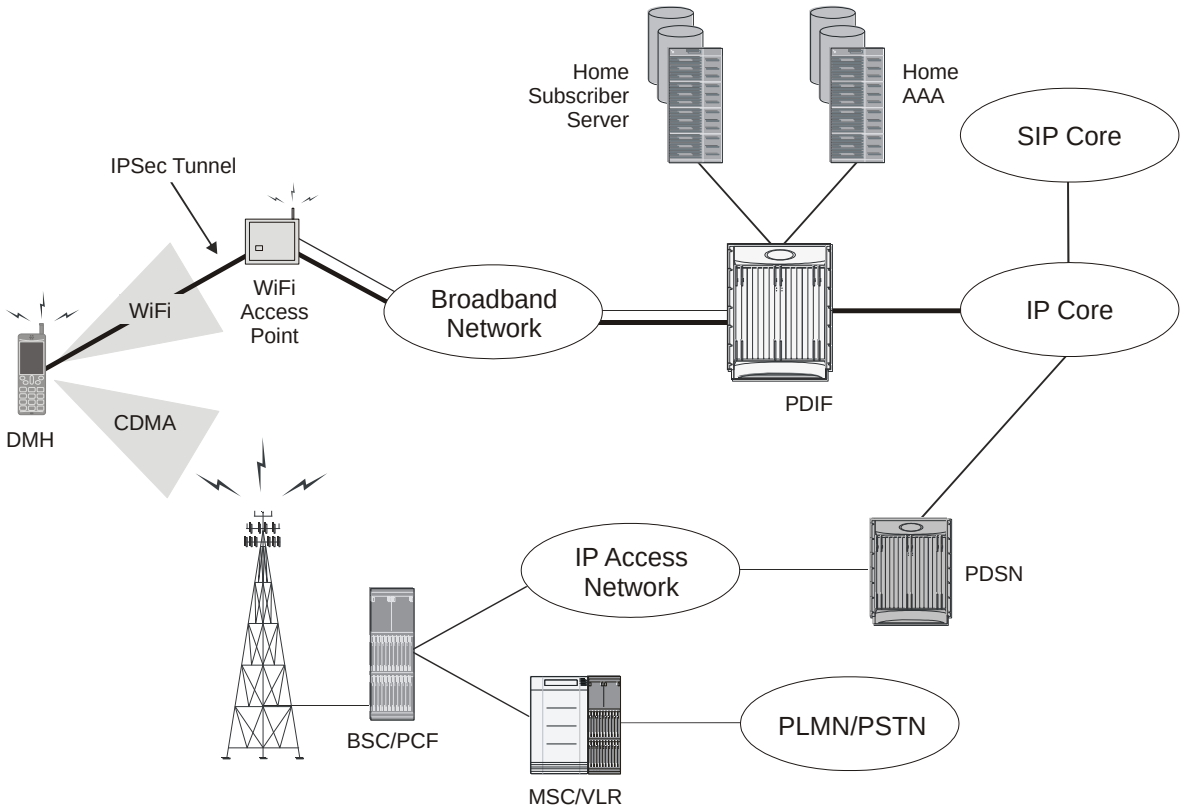
Licenses

The PDIF is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the "Managing License Keys" section of the "Software Management Operations" chapter in the *System Administration Guide*.

Interfaces

The figure below shows how the PDIF/FA acts as a security gateway between the Internet and packet data services. All components are located in the home network.

Figure 186. PDIF/FA Mobile IP Interfaces



1. The IPSec virtual tunnel interface with the MS: The Mode keyword in the IPSec-transform-set configuration mode defaults to Tunnel. In Tunnel mode, the IP datagram is passed to IPSec, where a new IP header is created ahead of the AH and/or ESP IPSec headers. The original IP header is left intact.
2. The Diameter interface: In a mobile IP network, the IMS Sh interface is used for MAC address validation with the HSS as well as HSS subscriber profile updates. In a Proxy-MIP network using multiple authentication, the HSS server is used to authenticate the device during Stage 1 authentication using the EAP-AKA authentication method.
3. The RADIUS authentication and accounting interface: In a mobile IP network, this interface is used for subscriber authentication using the EAP-AKA authentication method. For subscriber accounting, the PDIF/FA sends start, stop and interim messages to the accounting server. When used in a Proxy-MIP network using multiple authentication, RADIUS is used with the AAA servers to authenticate the subscriber using the GTC/MD5 authentication methods.
4. The home agent interface: This interface is used for Proxy mobile IP and mobile IP subscribers. All mobile station packets are tunneled to the HA through this interface. This interface is not used for simple IP subscribers.
5. The simple IP interface: This interface provides internet access for simple IP users.

Sample Deployments

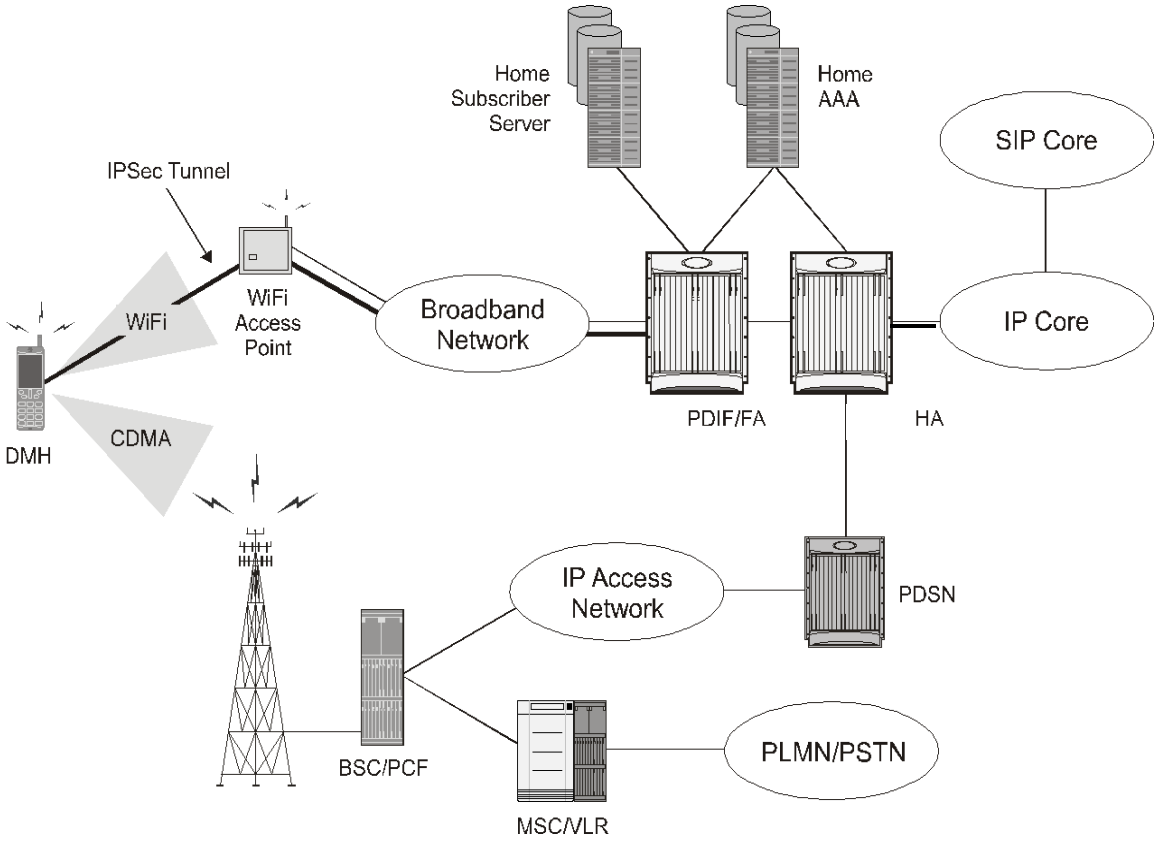
The following are some sample deployments using a PDIF/FA.

Mobile Station using Mobile IP with PDIF/FA

Overview

As shown in the figure below, the PDIF/FA supports the Fixed Mobile Convergence (FMC) application, which employs a Dual Mode Handset (DMH) to provide a VoIP solution over an IP-based WiFi broadband network. The DMH can access the traditional CDMA voice and data networks over the Radio Access Network (RAN). Over the RAN, the DMH implements circuit-switched voice and standard mobile IP (MIP) data over EVDO Rev. A, using the services of a PDSN and an HA.

Figure 187. PDIF/FA Mobile IP Implementation



Alternately, the DMH can send both voice and data over WiFi when a local AP is available. When the DMH connects to the AP, it establishes an IPSec tunnel over the broadband access network. This tunnel terminates at the PDIF/FA.

The DMH initially gets an IP address, also known as a Tunnel Inner Address (TIA), from the PDIF/FA when the DMH establishes the first IPSec tunnel. The PDIF/FA assigns the TIA from its IP address pool. The DMH then starts mobile IP through this initial TIA-based IPSec tunnel.

When the DMH successfully sets up mobile IP, it receives the home address from the HA. The DMH then establishes a second IPSec tunnel using this HA. Once the DMH successfully establishes the second IPSec tunnel with the PDIF/FA, the PDIF/FA tears down the first TIA-based IPSec tunnel to free the TIA, which then returns to the IP address pool. If required, use the **`no release-tia`** command in config-subscriber mode to prevent the TIA from returning to the pool. The DMH sends packetized voice and data through the PDIF/FA to the HA through the second IPSec tunnel.

In this scenario, the PDIF/FA forwards all the packets between the DMH and the HA. From there, voice packets are delivered to the Session Initiation Protocol (SIP) infrastructure, while data is delivered to the Internet or other appropriate destinations.

Mobile IP / Native Simple IP Call Minimum Requirements

The following provides the minimum requirements for each call type:

Mobile IP Calls

The PDIF/FA assumes MIP tunnel establishment over IPSec tunnel as part of the PDIF call flow as so on as any one of the following three possible conditions is met:

1. The default subscriber profile has configured, or:
2. The Radius VSA SN1-PDIF-MIP-Required is returned by AAA during user authentication, or,
3. The MS requests the MIP session type by injecting the IKEv2 configuration attribute 3GPP2_MIP4_MODE.

Native Simple IP Calls

The PDIF/FA assumes a native simple IP session over an IPSec tunnel if:

1. The MS (DMH) does not request 3GPP2_MIP4_MODE in IKEv2 exchange, and:
2. If a subscriber profile is defined, it does not have the pdif mobile-ip required parameter, and:
3. The AAA server does not return the VSA SN1-PDIF-MIP-Required during MS user authentication.

Mobile IP Session Setup over IPSec

The following diagram and table describe the mobile IP session setup over IPSec.

Figure 188. Mobile IP Session Setup over IPSec

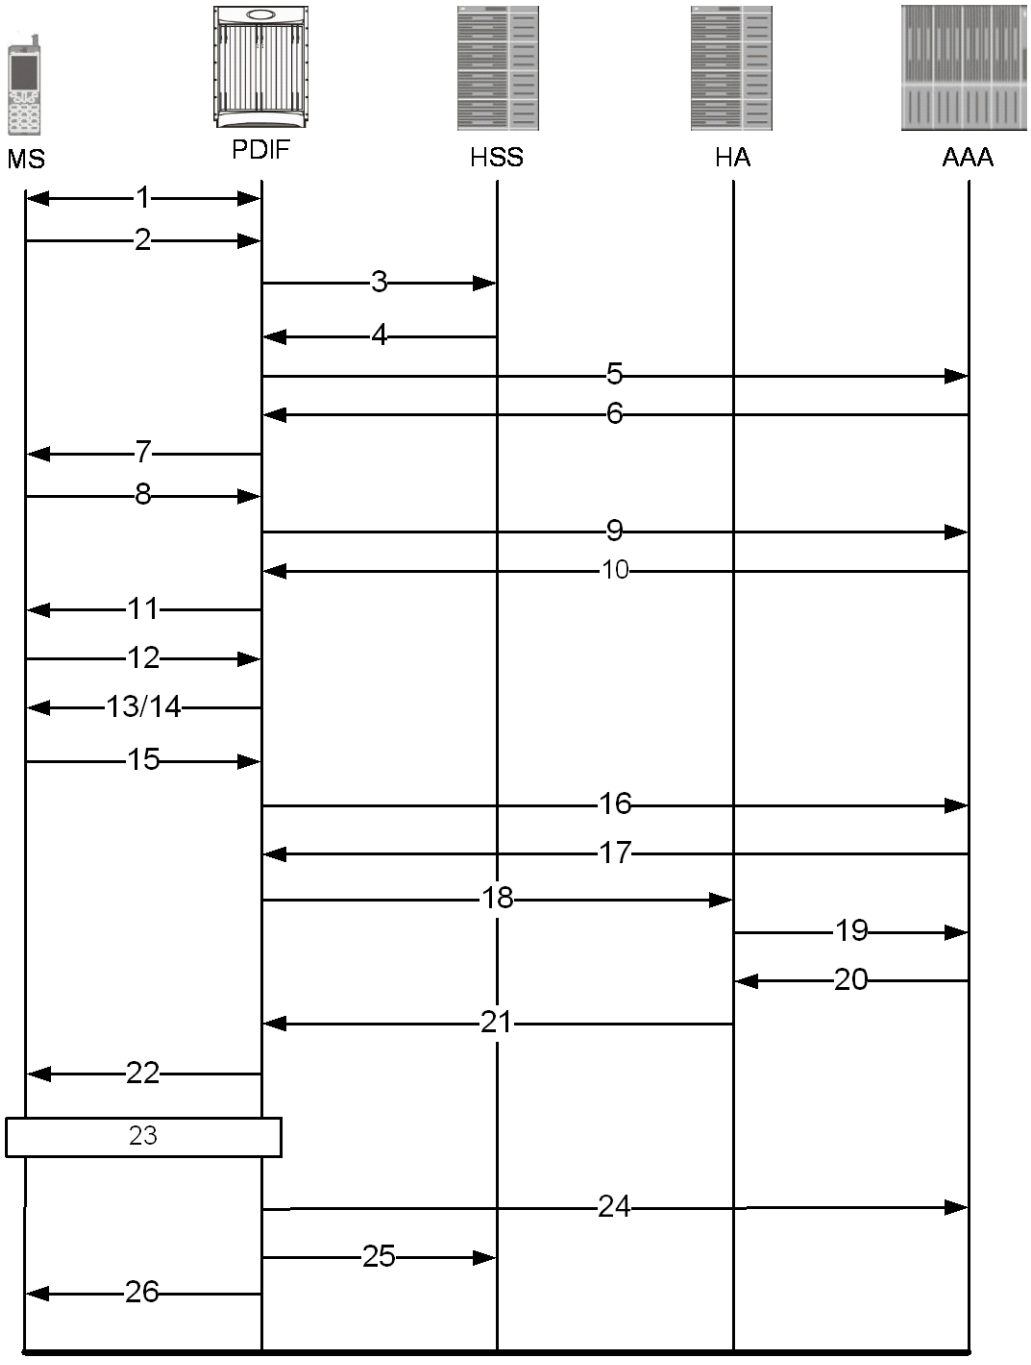


Table 83. Mobile IP over IPSec Call Flow Description

Step	Description
1	After the MS learns the IP address of the PDIF, the MS and the PDIF/FA exchange IKE_SA_INIT messages to negotiate an acceptable cryptographic suite.

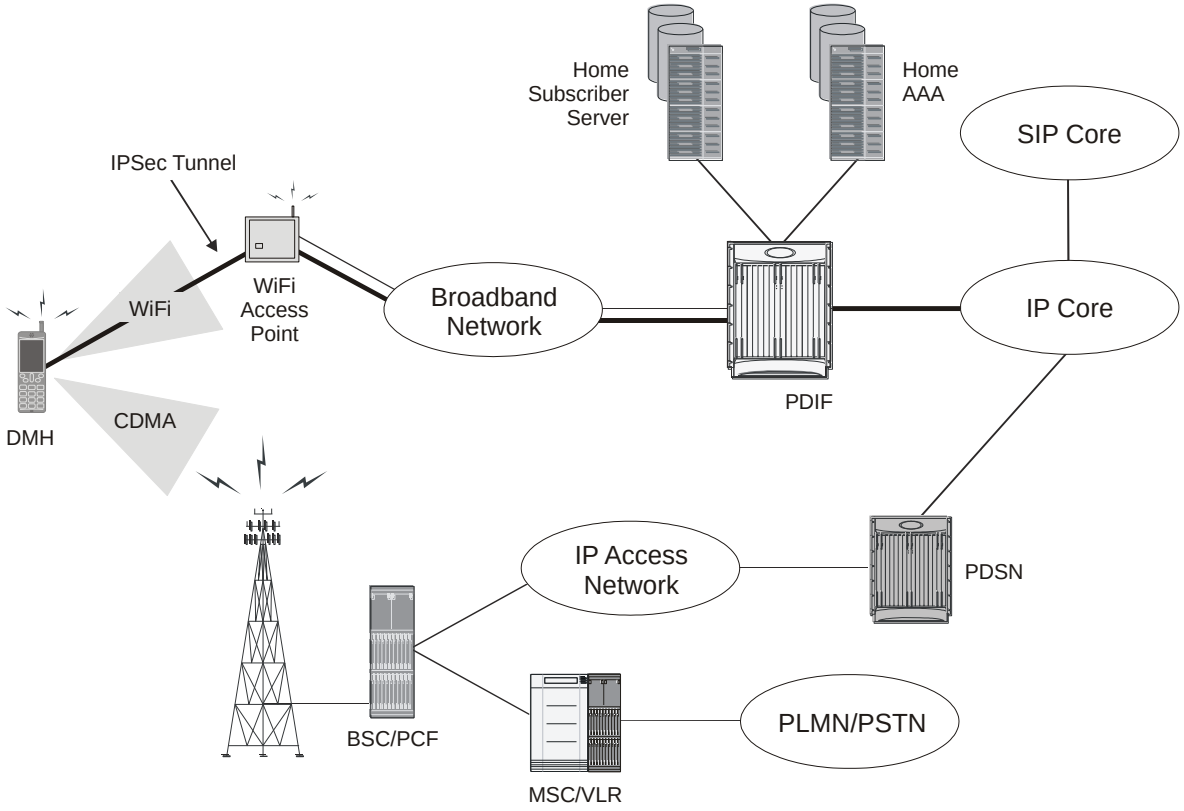
Step	Description
2	The MS initiates IKE_AUTH exchange messages with the PDIF/FA. The MS omits the AUTH parameter to the PDIF/FA, indicating that it wants to use EAP over IKEv2. The MS includes its identity in the IDi payload of the IKE_AUTH request. The IDi is set to be the same as the NAI and the NAI realm is chosen appropriately for M-NAI devices.The MS embeds the MAC address of the WiFi access point (AP) in the NAI and includes the IKEv2 configuration payload. Attributes included in the CFG_REQUEST are at least the INTERNAL_IP4_ADDRESS (with the length set to zero), the INTERNAL_IP4_DNS, and the 3GPP2_MIP_MODE.
3	When the PDIF/FA receives the IKE_AUTH request, it checks if MAC address authorization is enabled. If so, the PDIF/FA uses the ims-sh-service interface to the HSS and requests the list of authorized APs for this user via a User Data Request (UDR).
4	The HSS answers with the list of authorized WiFi APs for the user.
5	After checking that the AP MAC address in the realm portion of the NAI matches with one of the authorized MAC addresses received from the HSS, the PDIF/FA strips the AP MAC address from the realm portion of the NAI and sends the resulting NAI as an EAP response identity to the H-AAA using a RADIUS Access-Request message. This message includes at least the user-name set as the NAI being sent in the EAP response identity, the 3GPP2 correlation ID, the EAP-Message attribute, and the message-authenticator attribute.
6	The H-AAA verifies the identity and checks that WiFi service is allowed for the subscriber. The H-AAA generates a random value RAND and AUTN based on the shared DMU CHAP-key and a sequence number.The H-AAA sends the EAP-Request/AKA Challenge to the PDIF/FA via a RADIUS access-challenge. The EAP-Request/AKA Challenge contains the AT_RAND, AT_AUTN, and the AT_MAC attribute to protect the integrity of the EAP message.
7	The PDIF/FA sends an IKE_AUTH response to the MS with the EAP-Request/AKA-Challenge message received from the H-AAA.
8	The MS verifies the authentication parameters in the EAP-Request/AKA-Challenge message and if the verification is successful, it responds to the challenge with an IKE_AUTH Request message to the PDIF/FA. The main payload of this message is the EAP-Response/AKA-Challenge message.
9	The PDIF/FA forwards the EAP-Response/AKA-Challenge message to the H-AAA via a RADIUS access-request message (RRQ).
10	If authentication succeeds, the H-AAA sends a RADIUS access-accept message with the EAP-message attribute containing EAP Success. The H-AAA sends the EAP-Success and the MSK generated during the EAP-AKA authentication process to the PDIF/FA. The 64-byte MSK is split into two 32-byte parts, with the first 32 bytes sent in the MS-MPPE-REC-KEY and the second 32 bytes sent in the MS-MPEE-SEND-KEY.Both of these attributes (the values of which are encrypted) are needed to construct the 64-byte MSK at the PDIF/FA. If either are missing, the PDIF/FA rejects the session. In addition, the H-AAA sends other attributes equivalent to what it normally sends to the PDSN for a simple IP session. The attributes include at least the following: The Framed-Pool (if required) so that the PDIF/FA can assign a TIA from the right IP address pool, the Session-Timeout, and The Idle-Timeout.
11	The PDIF/FA forwards the EAP Success message to the MS in an IKE_AUTH Response message.
12	The MS calculates the MSK (RFC 4187) and uses it to generate the AUTH payload to authenticate the first IKE_SA_INIT message. The MS sends the AUTH payload in an IKE_AUTH Request message to the PDIF/FA.
13	The PDIF/FA uses the MSK to check the correctness of the AUTH payload received from the MS and calculates its own AUTH payload for the MS to verify [RFC 4306]. The PDIF/FA sends the AUTH payload to the MS together with the Configuration Payload (CP) containing security associations and the rest of the IKEv2 parameters in the IKE_AUTH Response message, and the IKEv2 negotiation terminates.The CP contains the TIA and IP address of the DNS servers that the device had requested earlier. Although the MS requested a DNS address by including only a single payload option for INTERNAL_IP4_DNS, the PDIF/FA may include both a primary DNS address and a secondary DNS address if one is available.

Step	Description
14	After a CHILD_SA is created using the TIA, if the PDIF/FA received 3GPP2_MIP_MODE during the IKEv2 negotiation, or if MIP_Required subscriber configuration is present in the subscriber profiles, the PDIF/FA sends agent advertisements to the MS.
15	The MS sends a MIP RRQ (including the NAI extension), an MN-AAA authentication extension, etc., to the FA. The HA IP address is set to 0 (zero) because the H-AAA assigns the HA. This is the usual NAI without the MAC address of the WiFi AP in the realm.
16	The PDIF/FA sends a RADIUS access-request to the H-AAA to authenticate the MS credential conveyed in the MN-AAA authentication extension and requests the assignment of an HA.
17	The H-AAA authenticates the MS successfully and sends the RADIUS access-accept message with the HA IP address.
18	The PDIF/FA forwards the RRQ to the HA.
19	The HA sends an access-request to the H-AAA to retrieve the MN-HA key in order to authenticate the MN-HA extension.
20	The HA receives the MN-HA key and authenticates the extension.
21	The HA assigns the IP address (HoA) for the MS and sends the RRP back to the PDIF/FA.
22	The PDIF/FA sends the HoA IP address to the MS.
23	After the MS obtains the HoA in the RRP, the MS sends the CREATE_CHILD_SA message with the Traffic Selector payload for Initiator (TSi) set to the HoA. This IKEv2 exchange creates a new IPSec SA.
24	The PDIF/FA sends a RADIUS accounting start message to the H-AAA.
25	The PDIF/FA then updates the subscriber's HSS profile with the indication that the IPSec session is active and the appropriate IP address. In this case, since it is MIP, it is the HoA assigned by the HA. In the case of simple IP Fallback, it would be the TIA assigned by the PDIF/FA. The HSS profile is updated using the Profile Update -Request (PUR) command.
26	PDIF/FA sends Delete payload in the informational message to delete the old IPSec SA associated with the previously assigned TIA.

Simple IP and Simple IP Fallback

For some simple IP deployments, the PDIF/FA authenticates the MS and provides an IP address for packet data services. In addition, the PDIF/FA supports Simple IP fallback if the MS abandons mobile IP operations due to not being able to successfully finish mobile IP registration after the first TIA-based IPSec tunnel is established. These scenarios are described below.

Figure 189. PDIF Simple IP Implementation



As described for mobile IP, during the initial IPSec tunnel establishment the MS gets a publicly routable TIA from a pool specified in the Framed Pool RADIUS attribute. When the IKEv2 negotiation finishes, an IPSec SA with a TIA is established as shown above.

Under normal situations, the MS successfully finishes mobile IP and establishes a new IPSec tunnel. However, if mobile IP fails, and simple IP fallback mode is enabled, the MS can revert to simple IP fallback mode and start using the TIA as the source IP address for all communication.

Important: Simple IP fallback is disabled by default. Use the **`pdif mobile-ip simple-ip-fallback`** command in config-subscriber mode to enable simple IP fallback.

Under these circumstances, the PDIF/FA opens the IPSec tunnel to data traffic and forwards any packets from the MS to the Internet directly. Any received packets from the Internet will be forwarded to the MS. A summary of this process from the point the TIA is assigned is given below:

Figure 190. Simple IP Fallback Message Sequence

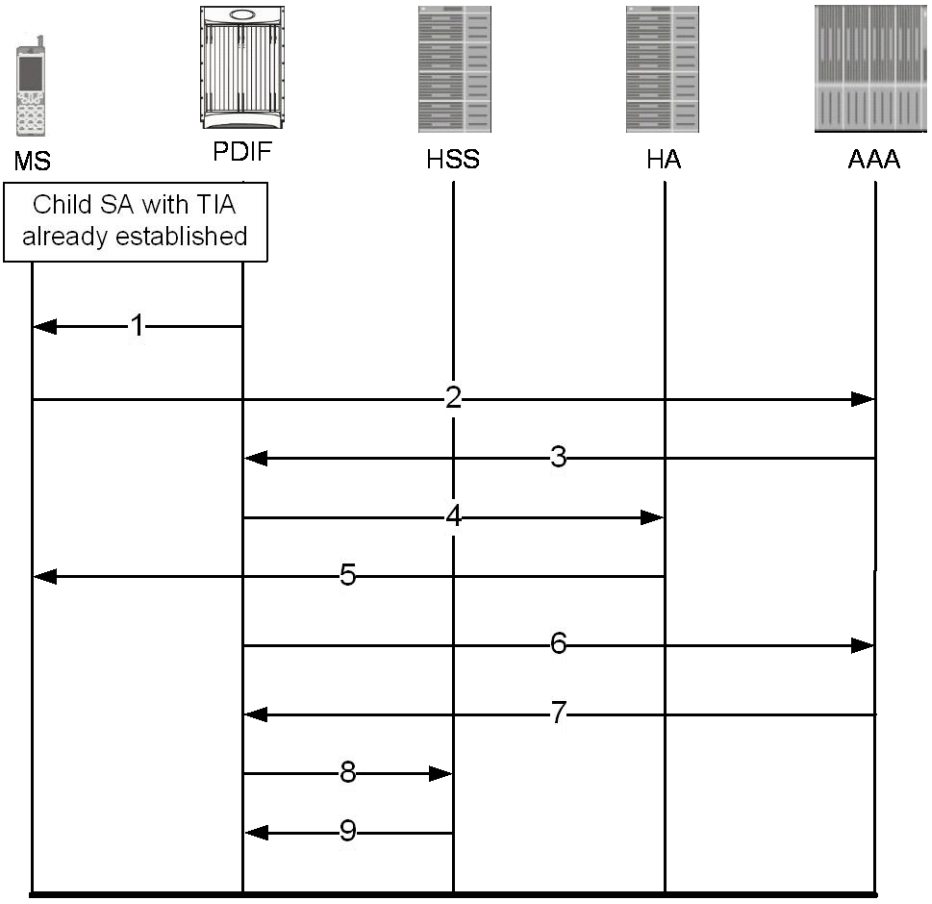


Table 84. Simple IP Fallback Message Sequence

Step	Description
1	With the IPSec Child SA (Security Association) and TIA already in place, the PDIF sends advertisements to the MS.
2	The MS sends a Registration Request (RRQ) message to the PDIF. The PDIF sends an authentication request to the AAA server over the RADIUS interface.
3	The AAA server authenticates successfully and sends the IP address of the HA.
4	The PDIF forwards the RRQ message to the HA.
5	The HA denies the request. The PDIF forwards the denial code to the MS.
6	The session setup timer expires and the PDIF goes into fallback mode. The PDIF sends a RADIUS Accounting Start message.
7	The AAA server sends a RADIUS Accounting Response message.
8	The PDIF updates the HSS with the TIA address of the subscriber.
9	The HSS sends an acknowledgement to the PDIF.

Simple IP Fallback Minimum Requirements

There are certain minimum requirements for simple IP fallback, as follows:

- There must be a context defined in the CLI configuration.
- The default subscriber must be defined in the CLI configuration.
- Mobile IP Simple IP Fallback must be defined in the CLI configuration. For example:

```
configuration
   context <pdif-in>
      subscriber default
      pdif mobile-ip simple-ip-fallback
      exit
```

- The MS has to request MIP by sending an RRQ message to the PDIF/FA. If the MS indicated an intent to use mobile IP (or was configured with the MIP_Required parameter) but failed to send an RRQ message, the IPSec session would be disconnected rather than completing a simple IP fallback call.
- On supported networks, the PDIF/FA only assumes simple IP fallback mode if mobile IP is attempted but fails when the MS tries to use mobile IP as the first choice but encounters a problem such as the HA not responding.

Features and Functionality - Base Software

This section describes the features and functions supported by default in the base PDIF software and the benefits they provide.

Important: All known restrictions are shown in Appendix B.

The following is a list of the features in this section:

- PSC2 Support
- Duplicate Session Detection
- Unsupported Critical Payload Handling
- Registration Revocation
- CHILD SA Rekey Support
- Denial of Service (DoS) Protection: Cookie Challenge
- MAC Address Validation
- RADIUS Accounting
- Special RADIUS Attribute Handling
- IPv6 Support
- IPv6 Neighbor Discovery
- IPv6 Static Routing
- Port-Switch-On-L3-Fail for IPv6
- IKEv2 Keep-Alive (Dead Peer Detection (DPD))
- Congestion Control and Overload Disconnect
- SCTP (Stream Control Transmission Protocol) Support
- X.509 Digital Trusted Certificate Support
- 2048-bit Certificate Key Functionality
- Custom DNS Handling

PSC2 Support

The PDIF supports the Packet Services Card 2 (PSC2). The PSC2 is the next-generation packet forwarding card for the ASR 5000. The PSC2 provides increased aggregate throughput and performance, and a higher number of subscriber sessions.

The PSC2 has been enhanced with a faster network processor unit, featuring two quad-core x86 2.5Ghz CPUs, 32 GB of RAM. These processors run a single copy of the operating system. The operating system running on the PSC2 treats the two dual-core processors as a 4-way multi-processor.

The PSC2 has a 2.5 G/bps-based security processor that provides the highest performance for cryptographic acceleration of next-generation IP Security (IPSec), Secure Sockets Layer (SSL), and wireless LAN/WAN security applications with the latest security algorithms.

Duplicate Session Detection

When an MS sets up a new session, the PDIF automatically checks for any remnants of abandoned calls and if found, clears them.

During a call, the processes of clearing the old session and establishing the new session run in parallel, optimizing processing functions.

With every new session setup, the PDIF supports a mechanism to verify whether there is any old session that is bound with the same International Mobile Subscriber Identity (IMSI) number. This is derived from the Callback-Id AVP in the last DEA message from the HSS after it has verified the subscriber.

For example, if an MS accesses the PDIF and subsequently moves out of the Wi-Fi coverage area, when the MS comes back on line, it could initiate a new session. After authentication, if an old session with the same IMSI is detected, the PDIF starts clearing it by sending a proxy-MIP Deregistration request to the HA. Once a Deregistration request is sent and a Deregistration response is received, the PDIF resumes the new session setup by sending a proxy-MIP Registration request. This setup procedure continues after the PDIF receives a proxy-MIP Deregistration response from the HA.

IMSI-based duplicate session detection is supported per source PDIF context. The PDIF requires only one source context to be configured per PDIF, therefore duplicate session detection across the entire chassis is possible. The feature is designed with the assumption that no more than one call with duplicate identifies are in the setup stage at any time. There is no limit to the number of duplicate session handling iterations.

When an old session is cleared, the PDIF sends Diameter STR messages and Radius Accounting STOP messages to corresponding AAA servers.

The PDIF allows duplicate session detection based on the NAI or IMSI. Note that when detecting based on the NAI, it is the first-phase (Multi-Authentication device authentication phase) NAI that is used.

If NAI-based duplication session handling is enabled, the PDIF sends an INFORMATIONAL (Delete) message to the MS.

Duplicate Session Detection is configured in PDIF-Service mode. The default is NAI-based.

Note that this configuration applies only to calls established after the configuration is made. It is therefore suggested that this selection be made in the boot-time configuration before any calls are established. For example, if NAI-based is used initially and an X number of calls is established, and then the configuration changes to IMSI-based, IMSI-based duplicate session handling does not apply to the calls established before the configuration change.

Unsupported Critical Payload Handling

This feature provides a mechanism whereby the PDIF ignores all unsupported critical payloads and continues processing as if those payloads were never received.

For MOBIKE IKEv2 messages, the PDIF returns UNSUPPORTED_CRITICAL_PAYLOAD in the IKEv2 response messages. The PDIF also drops all NAT-T keep-alive messages.

Registration Revocation

Registration Revocation is a general mechanism whereby the HA providing mobile IP or proxy mobile IP functionality to a mobile node notifies the PDIF/FA of the termination of a binding. This functionality provides the following benefits:

- Timely release of mobile IP resources at the FA and/or HA
- Accurate accounting
- Timely notification to mobile node of change in service

Important: Mobile IP registration revocation is also supported for proxy mobile IP. However, in this implementation, only the HA can initiate the revocation.

CHILD SA Rekey Support

During Child SA (Security Association) rekeying, there exists momentarily (500ms or less) two Child SAs. This is to make sure that transient packets for the old Child SA are still processed and not dropped.

PDIF-initiated rekeying is disabled by default. This is the recommended setting, although rekeying can be enabled through the Crypto Configuration Payload mode commands. By default, rekey request messages from the MS are ignored.

Denial of Service (DoS) Protection: "Cookie Challenge"

There are several known Denial of Service (DoS) attacks associated with IKEv2. Through a configurable option in the **`Config Crypto-Template`** mode, the PDIF can implement the IKEv2 "cookie challenge" payload method as described in [RFC 4306]. This is intended to protect against the PDIF creating too many half-opened sessions or other similar mechanisms. The default is not enabled. If the IKEv2 cookie feature is enabled, when the number of half-opened IPSec sessions exceeds the reasonable limit (or the trigger point with other detection mechanisms), the PDIF invokes the cookie challenge payload mechanism to insure that only legitimate subscribers are initiating the IKEv2 tunnel request, and not a spoofed attack.

If the IKEv2 cookie feature is enabled, and the number of half-opened IPSec sessions exceeds the configured limit of any integer between 0 and 100,000, the call setup is as shown in the figure below.

Figure 191. DoS Cookie-Challenge-Enabled IKEv2 Message Exchange

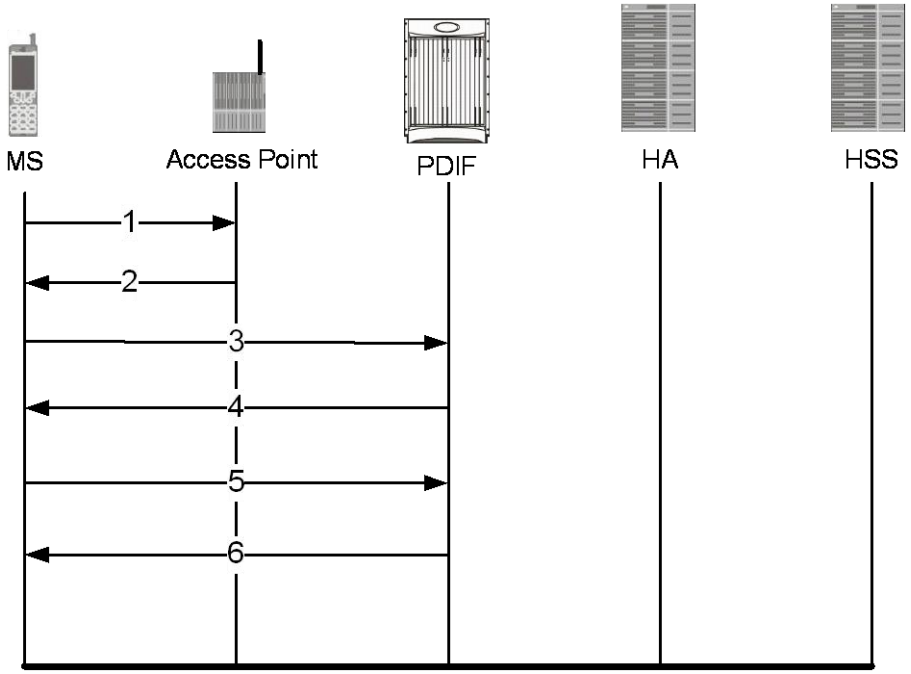


Table 85. DoS Cookie Challenge Enabled IKEv2 Message Exchange

Step	Description
1	The MS places a call to the WiFi AP.
2	The WiFi AP returns the IP address of the PDIF.
3	The MS sends an IKE_SA_INIT request. message.
4	The PDIF sends the Notify (cookie) payload to the MS to request retransmission of the IKE_SA_INIT request message to include the Notify (cookie) payload in the message.
5	Upon receipt of the retransmitted message, the PDIF verifies the cookie payload and ensures it is the same cookie as the one it had sent.
6	If the cookie challenge is met, setup continues as normal with an IKE_SA_INIT response message.

Cookie Challenge Statistics

Cookie challenge statistics appear in the outputs for the following commands:

- **`show crypto managers summary ikev2-stats`**: Shows the total number of invalid cookies per manager instance.
- **`show crypto managers summary npu-stats`**: Shows NPU statistics on each IPSec manager.
- **`show crypto statistics`**: Shows the combined data statistics for the given context name. Includes the number of cookie flows, the number of cookie flow packets, and the total number of cookie errors.

- **show crypto statistics ikev2**: Shows the control statistics for a given context name. Includes the output for **show crypto statistics**, plus Total IKEv2 Cookie Statistics, Cookie Notify Sent, Cookie Notify Received, Cookie Notify Match, Cookie Notify NOT Match, and Invalid Notify Payload Cookie.

MAC Address Validation

The MS embeds the MAC address from the WiFi AP in the NAI when it sends an IKEv2 AUTH request. If MAC address validation is enabled on the PDIF, it sends a Diameter User-Data-Request (UDR) message to the HSS with the NAI from the MS. The HSS returns a User-Data-Answer (UDA) message to the PDIF containing a list of authorized MAC addresses.

If the PDIF finds the MAC address in this list, the MAC address validation succeeds, and the PDIF continues with the IKEv2 call. The MS starts EAP authentication through IKEv2 AUTH procedures. If configured to do so, the PDIF removes the MAC address from the NAI when sending authentication requests to external RADIUS servers. If the embedded MAC address is not removed, the authentication check fails, because the AAA server cannot accommodate embedded MAC addresses.

If the MAC address is not in the list, the MAC address authorization fails, and the IKEv2 session is terminated with a Notify Message Type 16382 - Private User Errors message.

If the HSS interface is not reachable, it is possible that the IKEv2 session setup could continue as if the MAC authorization had succeeded. However, such error behaviors, including various Diameter error codes from the HSS, are configuration options. That means if an HSS returns an error, the action could be either to continue or to terminate the session. This is discussed in Diameter Failure Handling.

Important: See also *Diameter Authentication Failure-Handling* in the *CDMA Command Line Interface Reference*.

RADIUS Accounting

RADIUS Accounting messages are not generated while mobile IP setup is in progress.

- A RADIUS accounting START message is generated when the session is established.
- RADIUS INTERIM accounting messages are generated at configured intervals in a call.
- A RADIUS STOP accounting message is sent to the AAA server when the call ends.

There is no session dormancy in the PDIF. Once the session is active, the session never goes to a dormant state.

Special RADIUS Attribute Handling

Certain attributes require special handling on the PDIF with the attribute values either controlled by a RADIUS dictionary entry or a PDIF-service configurable. No configuration has no behavioral effect.

- 3GPP2-Serving-PCF. The generation of each new custom dictionary requires a new PDIF image. Configured in the pdif-service mode, the command **aaa attribute 3gpp2-serving-pcf** <*ip-address*> specifies the required values for the attribute without building a new software image. If configured, this attribute is sent in RADIUS accounting messages.

The following attributes are in custom dictionaries but have a customer-requested component.

- Calling-Station-ID. Required for PDIF RADIUS messages, there is a "dummy" value of 000000000000000 (fifteen zeros) set in this attribute. For non-PDIF product lines, the configured value may be taken only if no attributes are received through the corresponding access protocols. Configurable in the PDIF-service.
- NAS-Port-Type. The 3GPP2 X.P0028-200 standard requires this value to be set as "5 (= Virtual)." Controlled through the RADIUS dictionary.
- Service-Type. Cisco specifies a Service Type of "framed" for PDIF messages. Controlled through the RADIUS dictionary.
- Framed-Protocol. There is no attribute value defined for IPSec. Cisco specifies a value of "PPP" for PDIF messages. Controlled through the RADIUS dictionary.
- BSID. Base Station ID is used in billing for calculating time-zone offsets. There is a dummy value set in this attribute for RADIUS messages from the PDIF. Configured in the PDIF-service.
- 3GPP2-MEID and 3GPP2-ESN. Since the customer billing system expects these attributes, a null value is set in these attributes for RADIUS messages from the PDIF. Mobile Equipment Identifier (MEID) uniquely identifies the mobile equipment and is the future replacement for Electronic Serial Number (ESN) of the Mobile Station. Controlled through the RADIUS dictionary.
- 3GPP2-Last-Activity. The event timestamp is set in this attribute where applicable in RADIUS messages from PDIF. This attribute is the same as the 3GPP2-Last-User-Activity-Time standard attribute.
- 3GPP2-Service-Option. Set with a default value of 4095. Configurable in the PDIF-service.
- SN-Disconnect-Reason. This is a Cisco VSA that specifies a more detailed reason for session disconnection.
- 3GPP2-Active-Time If required for billing purposes, this VSA could be populated with the session length by generating a new RADIUS dictionary with this attribute. Unless specifically requested, a custom RADIUS dictionary does not include the 3GPP2-Active-Time VSA.

Mobile IP and Proxy Mobile IP Attributes

The SN-Proxy-MIP attribute is required when PDIF supports proxy mobile IP. The PDIF-Mobile-IP-Required attribute is SN1-PDIF-MIP-Required. These attributes need to be returned in a AAA response message or the mobile IP call fails, although there might be an option for simple IP call setup. See the *Sample Deployments* section for more information on attribute messaging.

IPv6 Support

This section describes the level of IPv6 support. All known restrictions are shown in Engineering Restrictions. Configuration examples are shown in Configuration.

Native IPv6 supports configuration of interfaces and routes with IPv6 (128-bit) addressing. PDIF supports IPv6 for communication with Diameter servers over SCTP. Using the Diameter proxy mechanism, each PSC needs a unique IPv6 address. Multiple IPv6 interfaces per context are supported.

Native IPv6 interfaces communicate with the Diameter servers. PDIF supports the configuration of 32 IPv6 Ethernet interfaces and 32 IPv6 loopback interfaces per context:

- One configured (CIDR global or site-local) IPv6 address per interface.
- Support for auto-configuration of link-local address based on an assigned MAC address. If the MAC address changes, the link-local addresses are updated accordingly. If a virtual MAC address is configured, it uses that MAC address for the link-local IFID. Note that this is distinct from the manual configuration of IPv6 addresses described below.

IPv6 Neighbor Discovery

IPv6 Neighbor Discovery protocol is used to dynamically discover the directly attached devices on IPv6 interfaces. It facilitates the mapping of MAC addresses to IPv6 Addresses. PDIF supports a subset of IPv6 Neighbor Discovery as defined by [RFC 2461] as follows:

- Uses IPv6 Neighbor Discovery to learn the Ethernet link-layer addresses of the directly connected next-hop gateway.
- Supports configuration of static IPv6 neighbors.
- Adds link-local addresses to Ethernet type interfaces automatically.
- Performs Unsolicited Neighbor Advertisement on line card switchover.
- Responds to neighbor discovery requests for the PDIF IPv6 addresses.

IPv6 Static Routing

Native IPv6 routing allows the forwarding of IPv6 packets between IPv6 networks. The forwarding lookup is based on destination IPv6 address longest prefix match.

PDIF supports configuration of static routes including a default route. If a default route is configured, all IPv6 traffic is forwarded to the configured next-hop defined by the default route.

Port-Switch-On-L3-Fail for IPv6

IPv4 port failover redundancy if L3 connectivity is lost is extended to support IPv6 addresses.

For more information on configuring port-switch-on-l3-fail, see *Ethernet Interface Configuration Commands* in the *CDMA Command Line Interface Reference* and *Creating and Configuring Ethernet Interfaces and Ports* in the *System Element Configuration Procedures* section of the *System Administration Guide*.

IKEv2 Keep-Alive (Dead Peer Detection (DPD))

PDIF supports DPD protocol messages originating from both the MS and the PDIF/FA. DPD is configured on a per-PDIF-service basis. The administrator can also disable DPD and the PDIF/FA does not initiate DPD exchanges with the MS when disabled. However, the PDIF/FA always responds to DPD availability checks initiated by the MS regardless of the PDIF/FA idle timer configuration.

Important: For a number of failure scenarios involving Dead Peer Detection, refer to the *Troubleshooting* chapter.

Congestion Control and Overload Disconnect

Congestion control is an operator-configurable facility. When the PDIF chassis reaches certain limits (based on CPU utilization, port utilization, and other controls) the system enters a congested state. When in a congested state, existing calls are not impacted but new calls are potentially restricted. There is a separate subscriber-level configuration to

enable/disable the feature on a per-subscriber basis. There is also a subscriber-level configurable for **inactivity-time** and **connect-time** thresholds to remove some old and abandoned calls from the system.

The disconnection scenario is as follows:

- If only **idle-time-threshold** is configured, sessions exceeding this threshold would be selected for disconnection.
- If only **connect-time-threshold** is configured, sessions exceeding this threshold would be selected for disconnection.
- If both **idle-time-threshold** and **connect-time-threshold** are configured, sessions with an idle-time greater than the idle-time threshold and a connect-time greater than the connect-time-threshold would be selected for disconnection.
- If neither **idle-time-threshold** nor **connect-time-threshold** is configured, sessions are sorted based on the idle-timer, and sessions with a longer idle-timer are deleted first.

SCTP (Stream Control Transmission Protocol) Support

PDIF provides support for SCTP (Stream Control Transmission Protocol) for use in communicating with Diameter peers over IPv6.

Diameter/SCTP connections are set up for administratively enabled Diameter peers whenever the system configuration is loaded. In the event of certain card or task-level failures, SCTP connections are torn down and re-established (but note that the Diameter state will still be maintained).

SCTP complies with the description in [RFC 2960 Section 5.1.1] for how to handle the case where the peer is incapable of supporting all of the outbound streams that the endpoint wants to configure. Specifically, PDIF does not abort the session but instead adjusts the association's number of outbound streams to match the number of inbound streams advertised by the peer (in the event that the number sent is less).

X.509 Digital Trusted Certificate Support

A digital certificate is an electronic credit card that establishes one's credentials when doing business or other transactions on the Web. Some digital certificates conform to ITU-T standard X.509 for a Public Key Infrastructure (PKI) and Privilege Management Infrastructure (PMI). X.509 specifies, among other things, standard formats for public key certificates, certificate revocation lists, attribute certificates, and a certification path validation algorithm.

The PDIF generates an SNMP notification when the certificate is within 30 days of expiration and approximately once a day until a new certificate is provided. The operator needs to generate a new certificate and then configure the new certificate using the CLI. The certificate is then used for all new sessions.

Important: For more configuration information, refer to *Global Configuration* in the *CDMA Command Line Interface Reference*.

2048-bit Certificate Key Functionality

This enhancement supports PDIF session setup when a PDIF certificate includes an RSA key size of 2048 bits. This includes:

- 2048-bit private key configuration in Global Configuration mode

- Certificate configuration that includes 2048-bit public keys
- The calculation of the AUTH payload in the first IKE_AUTH response from the PDIF using the 2048-bit private key
- Transporting the PDIF certificates with 2048-bit RSA public key length from the PDIF to the MS
- The validation of the AUTH payload using the 2048-bit RSA public key
- Dynamically applying a new 2048-bit certificate to 1024-bit or 2048-bit certificates with no service interruption to in-process calls
- The PDIF forwarding the first IKE_AUTH response with the CERT payload to the IPMS server (as it currently does), and the IPMS server decoding the message (if the proper IKE SA key is available)
- Collocation of certificates with 1024-bit keys and 2049-bit keys
- Fragmented packets are visible in external system output

Collocation of Certificate with 1024-bit Key and 2048-bit Key

Currently available dual-mode Wi-Fi handsets can only process certificates with 1024-bit key length. Future handsets will support 2048-bit key length only. During the transition period, it is possible that two different types of certificates must be configured in the same PDIF chassis:

- 1024-bit key only handset (abbreviated as 1024-bit MS)
- 2048-bit key only handset (abbreviated as 2048-bit MS)

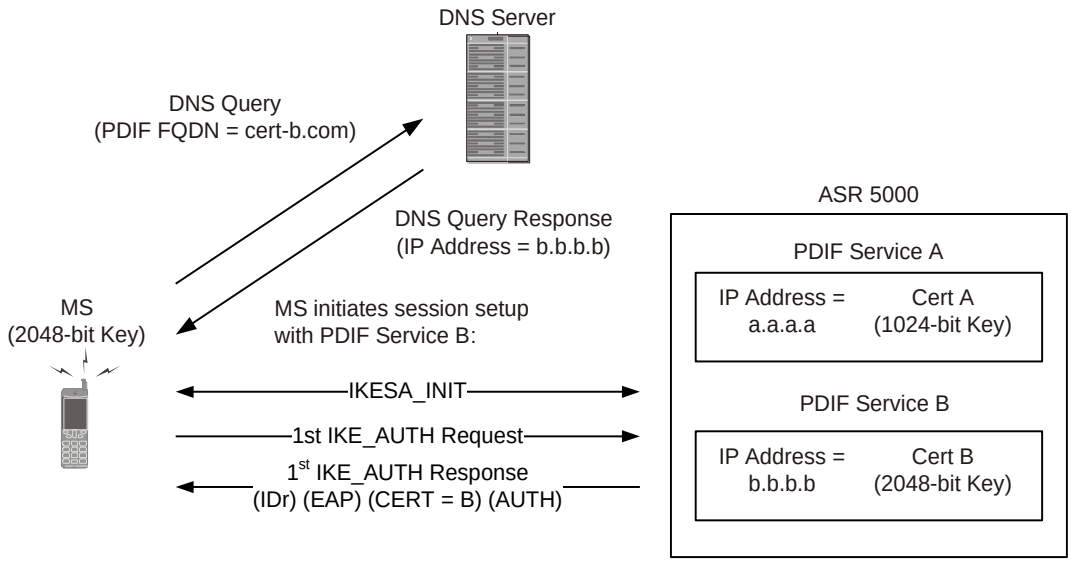
PDIF is informed as to which certificate the MS will receive.

Distinguishing Between 1024-bit and 2048-bit Key Certificates

Certificates are configured in Global Configuration mode and bound to a crypto template. The crypto template is mapped to each PDIF service. The MS uses different PDIF services depending on which key length it can support.

The MS first resolves the IP address of the PDIF service with a pre-configured FQDN. As shown in the following figure, two different FQDNs may be configured to initiate session setup with two different PDIF services.

Figure 192. Distinguishing Between 1024-bit and 2048-bit Key Certificates



Custom DNS Handling

By default, the PDIF always returns a DNS address in the CP payload if one is received from the configuration or the HA. A new CLI has been added defining an alternate series of supported behaviors depending on the number of INTERNAL_IP4_DNS. These include, but are not limited to, the following:

- Provides a mechanism whereby the DNS address present in configurations will be sent to the MS in the CP payload only if the MS requests one.
- The address 0.0.0.0 is treated as invalid and not included.

Gx+ and Gy Interface Support via HA

To provide secure WiFi access between the MS and the PDIF via PMIPv4 (Proxy Mobile IP version 4), this feature enables the PDIF and the FA (Foreign Agent) to run on the same chassis and enables the PDIF/FA to forward the following three attribute-value pairs (AVPs) to the HA in a PMIP RRQ (Proxy MIP Registration Request) message:

- Location information for the WiFi access point in the form of a BSID (Base Station Identifier) per 3GPP2.
- RAT (Radio Access Type) of WiFi per RFC 5563.
- User identity in the form of an NAI (Network Access Identifier) per RFC 2486.

After receiving these AVPs from the PDIF/FA, the HA forwards the AVPs over the Gx+ or Gy interface to the PCRF (Policy Charging and Rules Function) or OCS (Online Charging System). Optionally, this feature allows the co-location of the PDIF/FA and HA on one chassis.

The following call flow shows how this feature supports secure WiFi access between the MS and the PDIF via PMIPv4.

Figure 193. Gx+ and Gy Interface Support

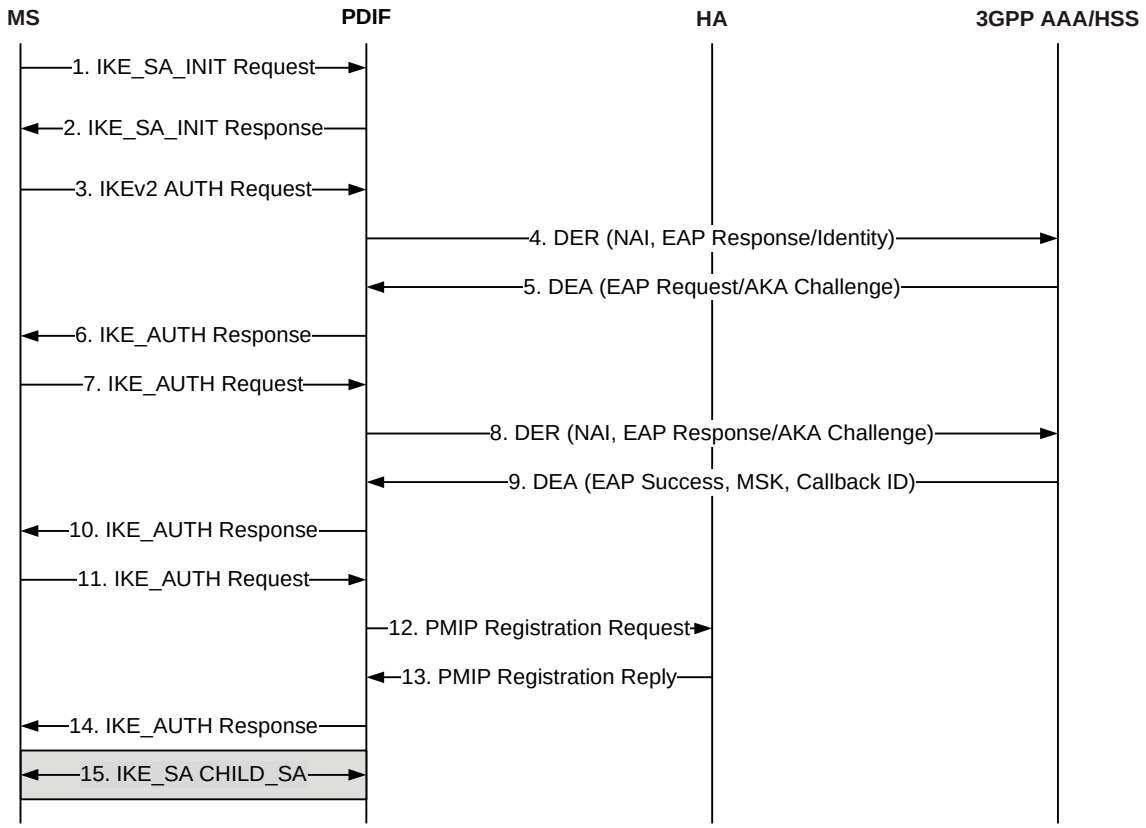


Table 86. Gx+ and Gy Interface Support

Step	Description
1.	The MS sets up an IKEv2/IPSec tunnel by sending an IKE_SA_INIT Request message to the PDIF. The MS includes SA, KE, Ni, and NAT-DETECTION Notify payloads in the IKEv2 exchange.
2.	The PDIF processes the IKE_SA_INIT Request message for the appropriate PDIF service (bound by the destination IP address in the message). The PDIF responds with an IKE_SA_INIT Response message with SA, KE, Nr, and NAT-Detection Notify payloads. The PDIF starts the IKEv2 setup timer upon sending the IKE_SA_INIT Response message.
3.	Upon receiving a successful IKE_SA_INIT Response message from the PDIF, the MS sends an IKE_ AUTH Request message for EAP-AKA authentication. The MS includes an IDi payload, which contains the NAI, SA, TSi, TSr, and CP (requesting an IP address and DNS address) payloads. The MS does not include an AUTH payload to indicate that it will use EAP methods.
4.	Upon receiving an IKE_AUTH Request message from the MS, the PDIF sends a DER (Diameter EAP Request) message to the Diameter AAA server. AAA servers are selected based on domain profile, default subscriber template, or default domain configurations. The PDIF includes a Multiple-Auth-Support AVP and EAP-Payload AVP with EAP-Response/Identity in the DER message. The PDIF starts the session setup timer upon receiving an IKE_AUTH Request message from the MS.
5.	The PDIF receives a DEA (Diameter EAP Accept) message with a Result-Code AVP indicating that EAP authentication should continue.

Step	Description
6.	The PDIF takes the contents of the EAP-Payload AVP and sends an IKE_ AUTH Response message back to the MS in the EAP payload. The PDIF will optionally include IDr and CERT payloads (depending upon the configuration). The PDIF allows IDr and CERT configurations in the PDIF service. The PDIF optionally includes an AUTH payload in the IKE_AUTH Response message (if the PDIF service is configured to do so). The MS receives the IKE_AUTH Response message from the PDIF.
7.	Upon receiving the IKE_AUTH Response message from the PDIF, the MS processes the exchange and sends a new IKE_AUTH Request message with an EAP payload.
8.	The PDIF receives the new IKE_AUTH Request message from the MS and sends a DER message to the AAA server. This DER message contains the EAP Payload AVP with an EAP AKA Challenge Response and Challenge Received from the MS.
9.	The AAA server sends the DEA message back to the PDIF with a Result-Code AVP of Success. The EAP Payload AVP message also contains an EAP result code of Success. The DEA message also contains the IMSI for the user, which is included in the Callback-Id AVP. The PDIF uses this IMSI for all subsequent session management functions, such as duplicate session detection. The PDIF also receives the MSK from the AAA server, which is used for further key computation.
10.	The PDIF sends the IKE_AUTH Response message back to the MS with the EAP payload.
11.	The MS sends the final IKE_AUTH Request message with the AUTH payload computed from the keys. The PDIF processes the AUTH Request message and responds with an IKE_AUTH Response message with the AUTH payload computed from the MSK.
12.	The PDIF does not assign any IP address for the MS and initiates a Proxy-MIP Registration Request (RRQ) message to the HA if “proxy-mip-required” is enabled. And if "mobile-ip send bsid" and “mobile-ip send access-technology” are enabled, the PDIF includes the additional attributes RAT and BSID in the PMIP RRQ message. The user ID is already include in the PMIP RRQ as the NAI. The HA address may be received from the AAA server in the Access Accept or may be locally configured. The PDIF may also remember the HA address from the authentication received in the final DEA message. If proxy-mip-required is disabled, PDIF will assign the IP address from the local pool. Detailed changes in the RRQ NAI attribute are sent as before under the attribute Username, but no changes are made to this attribute.
13.	The HA returns a Proxy-MIP Registration Reply (RRP) message to the PDIF.
14.	The PDIF sends a final IKE_AUTH Response message to the MS.
15.	An IKE_SA CHILD_SA gets established between the MS and the PDIF.

To enable this feature on the PDIF, from the system CLI, in Subscriber Configuration Mode, you issue the following two commands:

- **mobile-ip send bsid**
- **mobile-ip send access-technology**

For command descriptions, see the *CDMA Command Line Interface Reference*.

Features and Functionality - Licensed Enhanced Feature Support

This section covers any feature not covered by the base PDIF software and is licensed either separately or in a customized bundle of feature licenses.


Important: Contact your local Sales or Support representative for information on how to obtain a license.

This section describes the following features:

- PDIF Service
- Multiple PDIF Services
- Lawful Intercept
- Diameter Authentication Failure Handling
- Online Upgrade
- Operation Over a Common IPv4 Network
- Operation Over a Common IPv6 Network
- Session Recovery Support
- IPSec/IKEv2
- Simple IP Fallback
- Simple IP
- Proxy Mobile IP
- Multiple Authentication in a Proxy Mobile IP Network
- RADIUS Authentication
- Termination
- Session Recovery
- Intelligent Packet Monitoring System (IPMS)
- Multiple Traffic Selectors
- Selective Diameter Profile Update Request Control
- Support of SHA2 Algorithms

PDIF Service

The PDIF service and the processes associated with it define the PDIF itself. The PDIF service enables mobile stations to interface with the PDIF.

The PDIF service configuration includes the following:

- **The IPv4 address for the service:** This is the PDIF IP address to which the MS tries to connect. The MS sends IKEv2 messages to this IP address and this address must be a valid address in the context. PDIF service will not be up and running if this IP address is not configured.
- **The name of the crypto template for IKEv2:** A crypto template is used to configure an IKEv2 PDIF IPSec policy. It includes most of the IPSec parameters and IKEv2 parameters for keep-alive, lifetime, NAT-T and cryptographic and authentication algorithms. There must be one crypto template per PDIF service. The PDIF service will not be up and running without a crypto-template configuration.
- **The EAP profile name:** This profile defines the EAP authentication methods.
- **Multiple authentication support:** The multiple authentication configuration is a part of the crypto template.
- **IKEv2 and IPSec transform sets:** These define the negotiable algorithms for IKE SA and CHILD SA setup to connect calls to the PDIF/FA.
- **Configure the setup timeout value:** The MS connection attempt is terminated if the MS does not establish a successful connection within the configured value.
- **Mobile IP foreign agent context and foreign agent service:** This defines the system context where mobile IP foreign agent functionalities are configured.
- **Max-sessions:** The maximum number of subscriber sessions allowed by this PDIF service.
- **PDIF supports a domain template for storing domain related configuration:** The domain name is taken from the received NAI and searched in the domain template database.
- **3GPP2 serving PCF address:** This configurable specifies what value in the RADIUS attribute when sending authentication and accounting messages.
- **Duplicate session detection parameters:** PDIF supports either NAI (first phase authentication) or IMSI to be used for duplicate session detection. This configuration specifies whether duplicate session detection is based on IMSI or NAI. The default is NAI.

When the PDIF service is configured in the system with the IP address, crypto template, etc., the PDIF is ready to accept IKEv2 control packets for establishing IKEv2 PDIF sessions.

There is a limit to the number of CHILD SAs supported by each PDIF service. Traditionally, other Cisco services limit this to the number of subscriber sessions. The PDIF treats this as the number of CHILD SAs. This means that if each subscriber establishes only a single CHILD SA, the limit will be equal to the number of subscriber sessions. During CHILD SA rekeying, for a small duration of time, there are two CHILD SAs in the system. This is to make sure that transient packets for the old CHILD SA are still processed (not dropped).

Multiple PDIF Services

The PDIF supports multiple PDIF services running simultaneously on the same chassis. This feature enables operators to configure PDIF services with different crypto templates to support multiple subscriber handsets and to set per-service maximum session limits. The total number of sessions for all PDIF services running simultaneously on the same chassis must fall under the PDIF session counting license limit.

Lawful Intercept

The Cisco Lawful Intercept feature is supported on the PDIF. Lawful Intercept is a license-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your Cisco account representative.

Diameter Authentication Failure Handling

Diameter EAP failure handling defines error handling for both Session Termination Requests and for EAP Requests.

Specific actions (continue, retry-and-terminate, or terminate) can be associated with each possible result-code. EAP failure handling is flexible enough that wide ranges of result codes can be defined with the same action, or actions can be bound on a per-result-code basis.

A failure does not necessarily mean a summary termination of a call.

```
diameter authentication <failure-handling> session-termination-request

diameter result-code 5001-5005 action continue
```

configures result codes 5001, 5002, 5004 and 5005 to mean the session could continue regardless of the error,

and

```
diameter authentication <failure-handling> session-termination-request

 diameter result-code 5003 action terminate
```

configures result code 5003 to mean terminate the session immediately.

In this scenario, the PDIF receives the DEA from an HSS with the failure code 5003 to terminate the IKE setup for the session. The PDIF sends the IKE_AUTH Response containing a Notify Payload with the type as AUTH_FAILED plus the EAP payload if one was received in the DEA.

When the PDIF received the last DEA message with AVPs that are not in the dictionary, and with the M-bit set to 1, the PDIF disconnects the session.

Important: For more information on Diameter specific configuration, refer to the *Configuring Diameter Authentication Failure Handling* section in the *AAA and GTPP Interface Administration and Reference*.

Online Upgrade

The customer has the benefits of upgrading software from a fully redundant device without the expense of maintaining a fully loaded, fully redundant ASR 5000 in a permanent state of standby.

The PDIF supports online software upgrades with a single software version difference between two chassis. For example, upgrading from Release 8.1 to 8.2 is supported. Support for a chassis running greater differences in software versions would be qualified by Cisco on an as-needed basis.

Important: Refer to the *Maintenance* chapter in this guide for information on how to perform the upgrade.

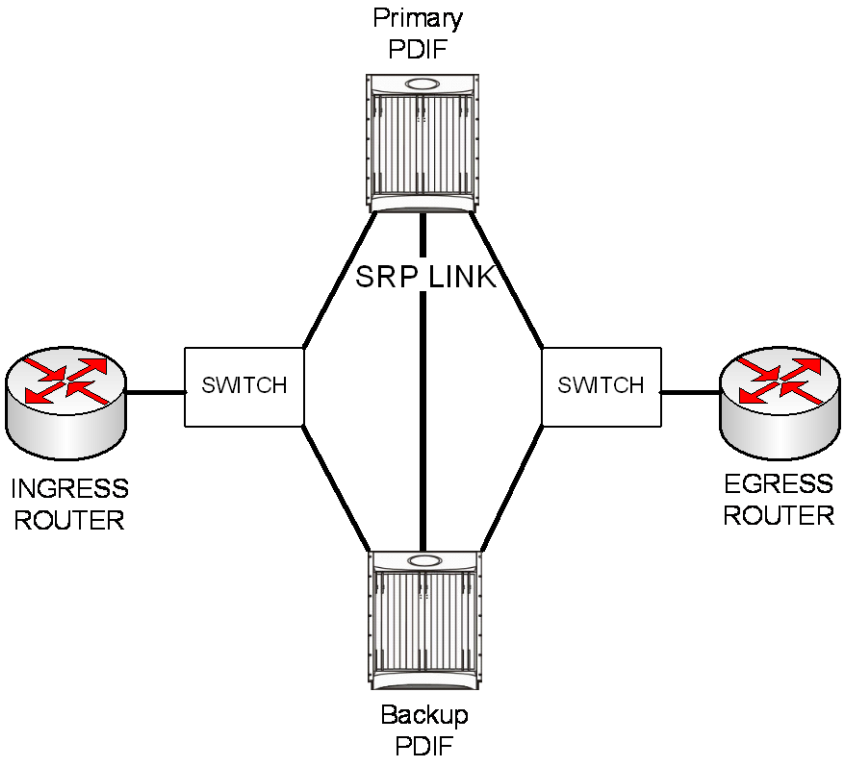
The online upgrade process calls for a spare ASR 5000 to temporarily perform the services currently being provided by a live networked chassis and upgrade the software with minimal service interruption. This model is called Active-Standby, as one chassis is designated as active and the other as standby. The standby chassis does not handle any new, incoming sessions because the DNS allocating new sessions does not know about the backup chassis. The backup is only required to handle sessions that were already on the primary chassis when it was administratively disconnected from the DNS server. Except for the data loss during the brief chassis switch-over, the session information (accounting and timers) are synchronized so that they are accurate when the backup becomes the active PDIF.

Important: Online upgrade requires miscellaneous internal processing that may result in intensive CPU utilization. Up to 50% CPU utilization overhead should be expected during the upgrade.

The Active-Standby Upgrade Model

The Active-Standby model is shown below:

Figure 194. Active-Standby Online Upgrade Model



The active and standby chassis are connected by an SRP redundancy link to monitor and control the chassis state. Both active and standby chassis have SRP-activated resources defined. Resources could mean loopback interfaces, broadcast interfaces, or IP pools, depending on the installation. For this example, use loopback interfaces.

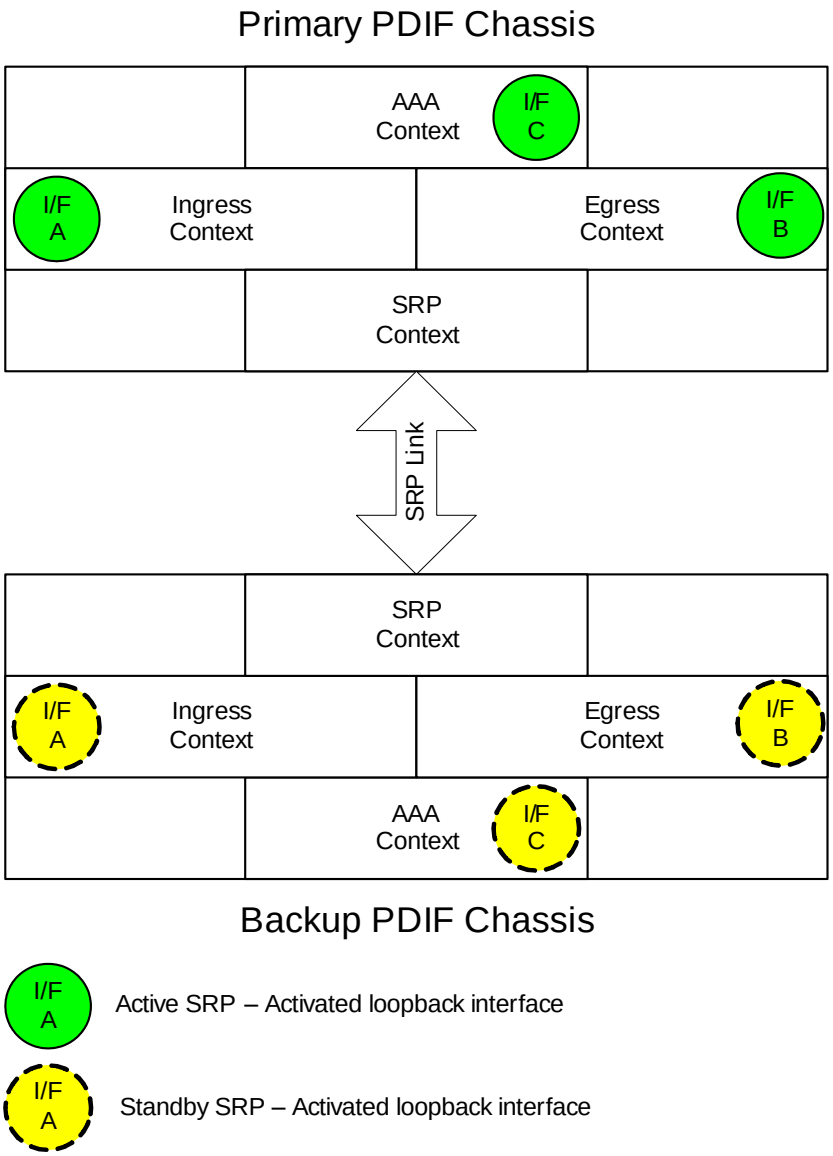
These resources are the same between the active and standby PDIF. Loopback IP addresses in ingress and egress contexts, and IP pools in egress contexts, are usually SRP-activated resources. The result is that only the currently active chassis enables the SRP-activated resources. The activate command is **`srp-activate`**.

Important: Ingress and egress contexts could be the same context. The SRP context must be a separate context.

In the network diagram below, each ingress context has loopback interface A defined, which is SRP-activated. PDIF service A is bound to this interface. The standby chassis has the same interface and PDIF service defined. Both interface and service can only be enabled on the active chassis. Similarly, interface B is defined in the egress context, which can be activated only in the active chassis.

When the active chassis switches over, the standby chassis becomes active and enables all SRP-activated IP interfaces and IP pools so that it can function as a mirror image of the former primary PDIF.

Figure 195. Loopback Interface Configuration



Operation Over a Common IPv4 Network

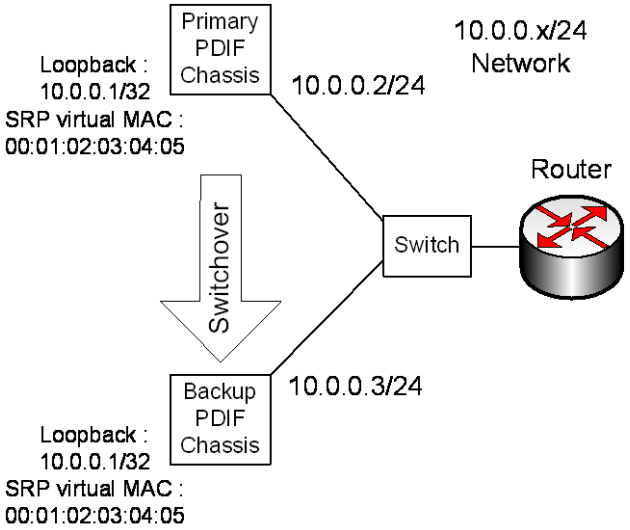
The PDIF supports L2 switching to enable carriers not using dynamic routing between the core nodes to perform an online upgrade.

In the example below, the SRP virtual MAC address is configured for the SRP-activated loopback address for the subnet. This allows the standby chassis to seamlessly assume the active role in the network after a switchover. Attached devices continue to send to the same SRP virtual MAC address and the currently active chassis responds to ARP

requests for the shared loopback IP address. This scheme allows fast standby-to-active transitions, since the SRP virtual MAC address does not change during the switchover.

When the ASR 5000 transitions from backup to primary, the PDIF sends Gratuitous ARPs to update the port-MAC table of the adjacent switch.

Figure 196. Switchover Example for Common IPv4 Subnet

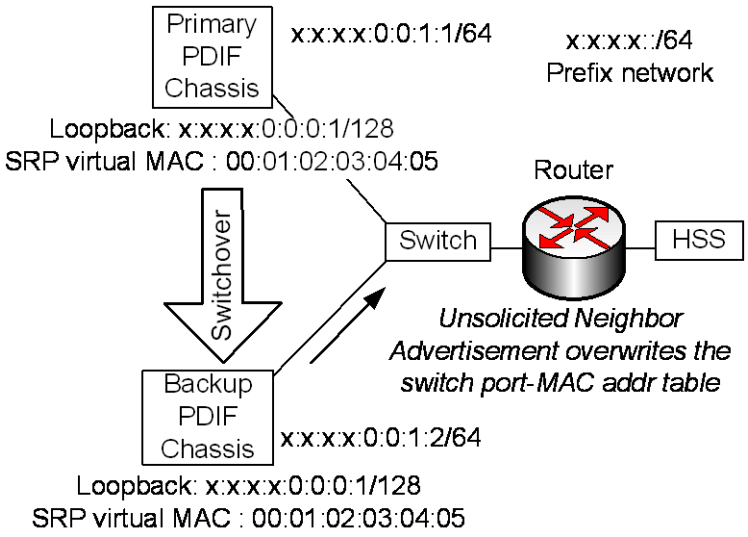


Operation Over a Common IPv6 Network

For AAA context with Diameter/SCTP/IPv6 configuration, multiple loopback IPv6 addresses are configured as Diameter endpoints. The customer can SRP-activate these loopback addresses and, upon SRP switchover, the HSS/SLF still sees the same Diameter peer endpoint. No new Diameter peer configuration to the HSS/SLF is required.

With SRP switchover operation in effect, the PDIF shuts down all the SCTP connections to the HSS/SLF. Then the former backup PDIF immediately creates new SCTP connections with the HSS/SLF. In this reestablishment process, the backup chassis sends an Unsolicited Neighbor Advertisement message to the adjacent switch, which is then used to overwrite its port MAC address table as shown in the diagram below.

Figure 197. Switchover Example for a Common IPv6 Subnet



Other Devices

The following table summarizes how other network devices see two ASR 5000s chassis during online upgrade. The table below assumes that a SRP-activated loopback address is configured in the source (toward the MS), the destination (toward the HA), and the AAA contexts (Diameter and RADIUS).

Table 87. The Chassis as seen from Other Network Devices During Upgrade

Network Entity	Consideration in Two-Chassis Configuration
L3 switch (MS ~ PDIF)	This L3 switch sees two chassis as a single entity. Only the physical port in the switch changes due to the switchover operation by G-ARP. The rest of the ASR 5000 information (IP address and MAC address) remain the same.
L3 switch (PDIF ~ HA)	This L3 switch sees two chassis as a single entity. Only the physical port in the switch changes due to the switchover operation by G-ARP. The rest of the ASR 5000 information (IP address and MAC address) remains the same.
Diameter Server	The MS sees two PDIFs as the same entity. However, upon switchover the SCTP connection is disconnected and then a new SCTP connection with ASR 5000 is established immediately. If an L3 switch exists between the PDIF and Diameter server, it sees two chassis as a single entity. Only the physical port in the switch changes due to the switchover operation by IPv6 Unsolicited Neighbor Advertisement. The rest of the ASR 5000 information (IP address and MAC address) remains the same.
RADIUS Server	This L3 switch sees these two chassis as a single entity. Only the physical port in the switch changes due to the switchover operation by G-ARP. The rest of the chassis information (IP address and MAC address) remains the same. If there should be an L3 switch between the PDIF and a RADIUS server, it sees two chassis as a single entity. Only the physical port in the switch changes due to the switchover operation by G-ARP, and the rest of the ASR 5000 information (IP address and MAC address) remains the same.
IPMS Server	Each chassis is connected to an independent IPMS Server. When a switchover takes place, the new IPMS Server continues to capture and store the call logs (signaling messages and events).

Network Entity	Consideration in Two-Chassis Configuration
O&M Device	Each chassis is connected to an independent O&M Device. When a switchover takes place, the new O&M Device continues to perform the function as the original device was configured.

Session Recovery Support

The session recovery feature provides seamless failover and almost instantaneous reconstruction of subscriber session information in the event of a hardware or software fault within the same chassis, preventing a fully connected user session from being dropped.

Session recovery is performed by mirroring key software processes (the session manager and the AAA manager, for example) within a single PDIF. These mirrored processes remain in an idle state (in standby mode), wherein they perform no processing, until they may be needed in the case of a software failure (a session manager task aborts, for example). The system spawns new instances of standby mode sessions and AAA managers for each active Control Processor (CP) being used.

Additionally, other key system-level software tasks such as VPN manager are performed on a physically separate Packet Services Card (PSC/PSC2) to ensure that a double software fault (the session manager and the VPN manager fail at same time on same card, for example) cannot occur. The PSC used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

The additional hardware resources required for session recovery include a standby System Management Card (SMC) and a standby PSC.

There are two modes for session recovery.

- **Task recovery mode**: Wherein one or more session manager failures occur and are recovered without the need to use resources on a standby PSC. In this mode, recovery is performed by using the mirrored standby-mode session manager tasks running on active PSCs. The standby-mode task is renamed, made active, and is then populated using information from other tasks such as AAA manager.
- **Full PSC recovery mode**: Used when a PSC hardware failure occurs, or when a PSC migration failure happens. In this mode, the standby PSC is made active and the standby-mode session manager and AAA manager tasks on the newly-activated PSC perform session recovery.

Session/call state information is saved in the peer AAA manager task because each AAA manager and session manager task is paired together. To ensure task recovery, these pairs are started on physically different PSCs.

Important: For more information on session recovery support, refer to *Session Recovery* in this guide.

IPSec/IKEv2

IKEv2 and IPSec transform sets configured in the crypto template define the negotiable algorithms for IKE SA and CHILD SA setup to connect calls to the PDIF/FA by creating two secure tunnels. The first, called the Tunnel Inner Address (TIA) is for signaling traffic, but in some cases it can be used for user traffic which can then use the TIA IP address. The second IPSec SA connects the MS to an HA for a mobile IP call.

Refer to *Sample Deployments* for a full description of how a variety of calls are successfully set up (and torn down) in a variety of network scenarios.

At the beginning of IKEv2 session setup, the PDIF and MS exchange capability for multiple authentication. Multiple authentication is configured in the crypto template of the PDIF service. When multiple authentication is enabled in the PDIF service, the PDIF will include MULTIPLE_AUTH_SUPPORTED Notify payload in the initial IKEv2 setup response.

The MS first sends an NAI for the device authentication, in which EAP-AKA is used. After the successful EAP-AKA transaction between the MS and the HSS, the HSS is expected to return the IMSI number for this subscriber. The PDIF uses the authorized IMSI number for session management.

Once the device authentication is successful, the MS notifies the PDIF of its intention to continue subscriber authentication only if the PDIF indicates it has multiple authentication support during the initial IKEv2 exchanges. The MS sends the second NAI that may be different from the first one used during the device authentication. The subscriber authentication is completed either using EAP-MD5 or EAP-GTC. Upon successful authentication, the PDIF continues proxy MIP registration before granting its access to the network.

Even if the PDIF sends the MULTIPLE_AUTH_SUPPORTED capability in the initial IKEv2 setup response, the MS may not support multiple authentication and hence may not include MULTIPLE_AUTH_SUPPORTED Notify payload in the subsequent IKEv2 AUTH exchange. In this case, the MS may only go through the first authentication (which is EAP-AKA authentication). After EAP-AKA authentication, if proxy-mip-required is configured for the session (either through the domain or the default subscriber or the corresponding Diameter AVP), the PDIF will establish a proxy mobile IP session with the HA. The assigned IP address is normally done by the HA and the PDIF receives this address through proxy mobile IP RRP. The PDIF will pass this address back to the MS through the final IKE_AUTH exchange. On the other hand, if proxy-mip-required configuration is not present or disabled, then the PDIF will continue the simple IP session setup by allocating the IP address for the MS from the locally configured pool.

When the MS sends MULTI_AUTH_SUPPORTED Notify payload in subsequent IKE_ AUTH exchanges, the PDIF knows the MS wants to do the second authentication. After the first successful EAP-AKA authentication, the MS will indicate to the PDIF regarding the second authentication (through ANOTHER_AUTH_FOLLOWS Notify payload in the final IKEv2 AUTH request). Please note that the IP address of the MS will not be assigned during the first authentication if the second authentication is to happen. The MS will then initiate the second authentication IKEv2 exchanges. In some networks, this second authentication uses the RADIUS AAA interface. The proxy-mip-required attribute will normally be present in the subscriber profile (or in the domain or default subscriber template) through a RADIUS attribute in the Access Accept message. After successful authentication, if proxy-mip-required is enabled, the PDIF will setup a proxy mobile IP session with the HA, and the HA assigns an IP address to the MS. If proxy-mip-required is disabled (or not present in the subscriber/domain profile), the PDIF establishes a simple IP session and routes traffic using the direct IP interface.

Simple IP Fallback

Network operators with handsets that are mobile IP capable may want the MS to be connected to the network and capable of doing data transfer even though the mobile IP registration process might fail under certain situations. If the mobile IP registration failures are due to HA reachability issues or any authentication problems, the MS should still be able to connect to the network using a simple IP connection, assuming that simple IP fallback is enabled in the PDIF configuration. See *Simple IP* and *Simple IP Fallback* in this chapter for a full description of this type of network configuration.

Simple IP

Simple IP is a solution for network providers whose subscribers fall primarily within a limited set of requirements. It provides the following:

- A mobility solution for subscribers who do not typically roam outside their immediate coverage area.

- An appropriate level of service for users who do not use the network in such as way as to need constant service between coverage areas. For example, subscribers who do not perform large file downloads.
- A mechanism to complete a call even if the proxy-mip-required or mip-required attributes are not configured in the subscriber or domain profile.

Proxy Mobile IP

Proxy mobile IP has the following benefits:

- Allows an MS that does not support mobile IP to have the same roaming benefits of one that does.
- The PDIF communicates with the HA and acts as if the PDIF itself were the handset.
- Proxy mobile IP is configured through the **`proxy-mip-required`** configuration, or the corresponding Diameter AVP or RADIUS Access Accept messages. If neither are present, the PDIF establishes a simple IP session and the PDIF routes the call to the Internet or corporate network.

Proxy mobile IP provides a mobility solution for subscribers whose mobile nodes do not support mobile IP protocol. The PDIF sets up the mobile IP tunnel with the HA and the PDIF proxies or acts on behalf of the handset as if it were the handset. The subscriber receives an IP address from either the service provider or from their home network. As the subscriber roams through the network as if it were using a full mobile IP connection, the IP address is maintained providing the subscriber with the opportunity to use IP applications that require seamless mobility such as transferring files.

Important: Refer to *Proxy Mobile-IP* in the *System Administration Guide* for more information.

Multiple Authentication in a Proxy Mobile IP Network

Multiple authentication requires authenticating both the device and the subscriber.

At the beginning of the IKEv2 session setup, the PDIF and the MS exchange capability for multiple authentication. Multiple authentication is configured in the PDIF service as part of the crypto template where it is associated with an EAP profile. The EAP profile defines the authentication mode and method. If multiple authentication is enabled in the crypto template, the PDIF includes a MULTIPLE_AUTH_SUPPORTED Notify payload in the initial IKEv2 setup response.

Important: Even if the PDIF confirms MULTIPLE_AUTH_SUPPORTED capability in the initial IKEv2 setup response, the MS may not support multiple authentication and hence may not include a MULTIPLE_AUTH_SUPPORTED Notify payload in the subsequent IKEv2 AUTH exchange. In this case, the MS may only go through the first-phase (EAP-AKA) of device authentication.

During initial IKEv2/IPSec security setup exchanges, the MS undergoes both device authentication and subscriber authentication. This is because even if the device is fully authenticated, a PDIF may not be able to tell which service profile is applicable for the MS, nor the correct IP address to assign.

Important: First-phase authentication refers to device authentication, and second-phase authentication refers to subscriber authentication.

AAA Group Selection

A maximum of 64 AAA groups is allowed on the ASR 5000. This could be spread across multiple contexts or all groups can be configured within a single VPN context.

A maximum of 320 RADIUS servers is allowed on the chassis.

When the **aaa-large-configuration** command is issued, this number becomes 800 AAA groups and 1600 RADIUS servers configured within the chassis.

The PDIF service allows you to specify a different AAA group for each authentication phase. A given AAA group supports either Diameter or RADIUS authentication, but not both. In deployments where the NAI used in the first-phase authentication is different from the NAI used in the second-phase authentication, each NAI can point to different domain profiles in the PDIF.

RADIUS Authentication

For more information on AAA, RADIUS, and Diameter groups, refer to the *AAA and GTPP Interface Administration and Reference*.

The second authentication uses RADIUS for subscriber authentication. The PDIF supports EAP termination mode during the second half of multiple authentication. In this mode, EAP exchange takes place between the MS and the PDIF, and the PDIF takes the information exchanged in the EAP payload over IKEv2 into RADIUS attributes to support CHAP/PAP authentication with the RADIUS server, and vice versa.

By default, the PDIF initiates EAP-MD5 authentication and sends an EAP payload with an MD5-Challenge to the MS. The MS returns an MD5-Challenge response in the EAP payload. Upon receipt, the PDIF sends an RADIUS Access Request message which includes an NAI, a CHAP-Password, a CHAP-challenge (derived from the EAP payload), and an IMSI number (which is the calling station ID). Once the AAA server returns an Access-Accept message, optional attributes such as Framed-IP-Address and HA address are expected for the subsequent session setup processing. The PDIF translates this Access-Accept message into an EAP Success message, and returns this in an IKE_AUTH Response message.

It is possible that some MSs may not support CHAP authentication. In this case, the MS is expected to return the EAP payload with a legacy-Nak message when the PDIF sends an MD5-Challenge message. Upon receipt of the legacy-Nak message, the PDIF initiates an EAP-GTC procedure. When the MS returns EAP-GTC including its own password, the PDIF sends a RADIUS Access Request message which includes an NAI, a password, and an IMSI number. Once the AAA server returns an Access-Accept message, attributes such as Framed-IP-Address and HA address are expected for the subsequent session setup processing. The PDIF translates the Access-Accept message as EAP success, and returns this in an IKE_AUTH Response message.

If EAP-GTC is configured, then the EAP-GTC method is used instead of the EAP-MD5 method.

The PDIF does the following for IKEv2 and RADIUS authentication:

The PDIF terminates EAP-MD5/GTC authentication. The PDIF understands the values in the EAP payload, and maps them as RADIUS attributes for CHAP/PAP authentication.

Upon request from the MS, the PDIF performs EAP-GTC authentication instead of EAP-MD5.

Each domain profile may be configured with two AAA groups, one for Diameter and the other for RADIUS.

In deployments where both NAI happen to be the same for both authentications, it will point to the same AAA group and thereafter only one protocol (either RADIUS or Diameter) is used.

There are cases where the domain template may not be associated with a given NAI. In such cases, the default AAA groups are used for authentication. Since authentication happens in two phases, and each using Diameter and RADIUS

AAA groups respectively, there needs to be two default AAA groups (one for Diameter authentication and one for RADIUS authentication) for multiple authentication. The default AAA groups are configured in the PDIF service.

First-Phase Authentication

During first-phase authentication, the HSS authenticates the device. The MS first sends an NAI for device authentication. After the successful EAP-AKA transaction between the MS and the HSS, the HSS is expected to return an IMSI number for this subscriber. The PDIF takes this authorized IMSI number for session management.

This authentication method uses EAP between the MS and the AAA server, and the PDIF acts as a pass-through agent.

Important: First-phase authentication must use the EAP-AKA method.

Depending on the number of HSSs in the network, it is possible that a Subscription Locator Function (SLF) would be introduced into the network as a Diameter proxy or relay agent. If deployed, the SLF would be the first point of contact for the PDIF.

The protocol stack between the PDIF and the HSS/SLF is Diameter over SCTP over IPv6.

Second-Phase Authentication

Second-phase authentication uses EAP-MD5 or EAP-GTC authentication with IKEv2 using a legacy RADIUS server, which does not understand or implement EAP. This could be the same AAA server as those deployed in any existing EV-DO network. In this case, EAP authentication happens between the MS and the PDIF.

The protocol stack between the PDIF and the AAA server is RADIUS over UDP over IPv4.

The two algorithms for second-phase authentication are EAP-MD5 (which is the same as CHAP authentication) and EAP-GTC (which is the same as PAP authentication). When the MS sends the NAI to identify the subscriber, the PDIF initiates the EAP-Request with a challenge. Once the MS returns the challenge response, the PDIF maps it to a RADIUS ACCESS_REQUEST message to complete CHAP authentication. There is an internal mechanism to inform each peer if one method is not supported and to renegotiate to use the other supported method.

In general, session attributes during first-phase authentication are overwritten by those from second-phase authentication, unless specified separately. Exceptions to this include **`session-timeout`** and **`idle-timeout`**, when the lower values are taken.

Termination

During session setup, if there are any configuration mismatches or the PDIF cannot get the required information, the session setup process is terminated and appropriate log messages are generated.

If **`multiple-auth-supported`** is not enabled on the PDIF, and the MS still sends a MULTIPLE_AUTH_SUPPORTED Notify payload marked with the critical bit set, the PDIF returns UNSUPPORTED_PAYLOAD. Otherwise, the PDIF ignores it and processes the IKE packet as if the payload was never received. This is non-standard MS behavior.

Important: The multiple authentication process in a proxy mobile IP network is described in the Proxy-MIP chapter in this guide.

Session Recovery

The session recovery feature provides reconstruction of subscriber session information in the event of a hardware or software fault within the system, providing seamless failover and preventing a fully connected user session from being dropped.

In addition to maintaining call state information, information is retained in order to:

- Recover IPSec manager policies, all template maps, and all subscriber maps.
- Use the policies (including templates) to recover CHILD SA tunnels, flow IDs, and statistics.
- Recover or reconfigure NPU flow IDs and data path handles.
- Recover and restore the IKEv2 stack state for all tunnels.
- Supply the IKEv2 stack with needed data statistics to determine rekey and DPD states.
- Recover Diameter session information.

Recovery requires a complex interaction between IPSec and session subsystems. The IPSec subsystem also interacts with a Datapath that includes daughter cards, daughter card managers, and the NPU. The session recovery feature is disabled by default on the system, even when the feature use key is present.

The IPSec controller does not send an IPSec manager death notification to any subsystem. This allows the daughter card to continue to receive and decrypt IPSec tunnel data. It also allows both the session manager and daughter card to continue carrying subscriber traffic using NPU flows and IPSec SAs to transmit the data.

A session manager is created on a PSC and a corresponding AAA manager is created on a different PSC but is created with the same instance number. A session manager saves (check-points) its Call Recovery Record (CRR) on the AAA manager with an instance ID the same as its own. This pairs up the session manager and the AAA manager and at the same time guarantees session recovery in the event of a single PSC failure.

IPSec manager is also created on a PSC. When a PDIF call request arrives, the IPSec manager picks a session manager for this particular call using a demux library on the same PSC. This means the IPSec manager is associated with the session managers on the PSC.

The session subsystem continues to use the AAA manager as its storage system for the PDIF because AAA needs to provide other subscriber-related information to the session manager. Now that the session manager and the IPSec manager are paired on the same PSC, the IPSec manager is assured of data recovery in case of PSC failure. This is because the session manager saves its data on the AAA manager on a backup PSC.

Important: For more information, refer to the *PDIF Session Recovery* chapter in this guide.

Intelligent Packet Monitoring System (IPMS)

The IPMS provides a control-packet capture, database, and query facility. It provides the functions to assist operators to analyze and investigate call-related events at a later time.

Important: IPMS is described in the *IPMS System Administration Guide*.

Multiple Traffic Selectors

The PDIF can be configured with multiple IPSec traffic classes, each containing up to 128 traffic selectors, which are used during traffic selector negotiation with UEs. Multiple traffic selectors allow the PDIF to direct outbound traffic to selected IP addresses based on the following protocols: IP, TCP, UDP, and ICMP. The PDIF can also direct TCP and UDP traffic to selected IP addresses and port ranges.

Important: In this software release, the PDIF supports IPv4 traffic selectors only.

Per RFC 4306, when a packet arrives at an IPSec subsystem and matches a 'protect' selector in its Security Policy Database (SPD), the subsystem must protect the packet via IPSec tunneling. Traffic selectors enable an IPSec subsystem to accomplish this by allowing two endpoints to share information from their SPDs. Traffic selectors can be used to assure that both endpoint SPDs are consistent and can aid in the dynamic update of an SPD. Traffic selector payloads contain the selection criteria for packets being sent over IPSec security associations (SAs).

During traffic selector negotiation, each endpoint sends two traffic selector payloads in the messages exchanged during the creation of an IPSec SA. The first traffic selector payload is known as the TSi (Traffic Selector-initiator) and the second is known as the TSr (Traffic Selector-responder). Each traffic selector payload contains one or more traffic selectors, and each traffic selector can contain an IP address range, a port range, and an IP protocol ID. During traffic selector negotiation between the UE and the PDIF, the UE assumes the role of the initiator as it initiates an IPSec SA for its traffic, and the PDIF assumes the role of the responder. The PDIF can use multiple traffic selectors in its role as the responder.

Traffic selectors are applied to calls via an AAA attribute. During call setup, the PDIF's AAA manager selects the traffic class to use for a call based on the Radius vendor-specific attribute (VSA) TrafficSelector-class, which is received from the AAA server. The PDIF's Session Manager passes the selected traffic class configuration from its AAA Manager to its IPSec Manager, which then sends the traffic selectors to the UE in the TSr for all CHILD SAs in the call. If no matching traffic selector classes or traffic selectors have been configured on the PDIF, or if the PDIF does not receive the TrafficSelector-class attribute from the AAA server, or if the value of the received TrafficSelector-class attribute is 0, the PDIF returns the default traffic selector to the UE in the TSr, which allows all inbound traffic.

The PDIF saves the traffic class configuration in each call during call setup. Configuration changes made to the existing traffic class configuration will apply to new calls only. There is no hard limit to the maximum number of allowed traffic classes, but the recommended limit is 50.

When incoming traffic from a UE does not match any of the configured traffic selectors, the PDIF does not reject the traffic. Instead, the PDIF keeps a per-call counter to record the number of packets that do not match the configured traffic selectors. Outgoing traffic from the PDIF to the UE is not subject to traffic selection or checking.

Selective Diameter Profile Update Request Control

For mobile IP calls, the Selective Diameter Profile Update Request Control feature allows WiFi data-only sessions to co-exist with VoIP sessions on the PDIF platform.

When the PDIF is accessed by voice-enabled devices, it needs to interact with the HSS in order for a subscriber session to access the IP core network. When the PDIF is accessed by data-only devices, there is no need to interact with the HSS.

This feature is used to identify which subscriber sessions need to have the PDIF and the HSS exchange Diameter Profile Update Request (PUR) and Profile Update Answer (PUA) messages, and allows the PDIF to handle the call setup for a data-only client without having to interact with the HSS.

Selective PUR profiles on the AAA server are mapped to subscribers during AAA authentication via the Radius vendor-specific attribute (VSA) FMC-Type. FMC-Type has these possible values: voice or data. When the AAA server sets the FMC-Type value to voice, the PDIF and the HSS exchange PUR and PUA messages. When the AAA server sets the FMC-Type value to data, the PDIF and the HSS do not exchange PUR and PUA messages.

This feature is enabled by default and requires no configuration.

Support of SHA2 Algorithms

In addition to SHA1, the following algorithm variations, collectively known as SHA2, are now supported in the Cisco ASR5000 IKEv2 stack: SHA-256, SHA-384, and SHA-512. These algorthims provide an additional level of security as authentication mechanisms for IKE and IKEv2 protocols.

The truncated variants, known as Hash-based Message Authentication Mode (HMAC), are used in conjunction with the with SHA-256, SHA-384, and SHA-512 algorithms in IKE and IKEv2. The untruncated variants are known as Pseudo-Random Functions (PRFs), also used with IKE and IKEv2.

The authentication variants ensure that packets are authentic and cannot be modified in transit. PRF variants provide secure pseudo-random functions for generating keyed material and protocol-specific numeric quantities.

The following table summarizes the sizes associated with the algorithms.

Table 88. Sizes Associated with SHA2 Algorithms

Algorithm ID	Block Size	Output Length	Truncated Length	Key Length	Algorithm Type
For Authentication (with truncated output					
HMAC-SHA-256-128	512	256	128	256	Auth/integ
HMAC-SHA-384-192	1024	384	192	384	Auth/inte
HMAC-SHA-512-256	1024	512	256	512	Auth/inte
For PRFs (with untruncated output					
PRF-HMAC-SHA-256	512	256	None	Variable	PRF
PRF-HMAC-SHA-384	1024	384	None	Variable	PRF
PRF-HMAC-SHA-512	1024	512	None	Variable	PRF

Supported Standards and RFCs

3GPP2 References

- P.S0001-B Version 2.0 cdma2000 Wireless IP Network Standard
- X.S0011-001-C v3.0 cdma2000 Wireless IP Network Standard; Introduction
- X.S0011-002-C v3.0 cdma2000 Wireless IP Network Standard: Simple IP and Mobile IP Services
- X-S0013-000-A v1.0 All-IP Core Network Multimedia Domain - Overview
- X.S0013-010-0 v2.0 All-IP Core Network Multimedia Domain - IP Multimedia Subsystem Sh Interface; Signaling Flows and Message Contents – Stage 2
- X.S0013-010-A v1.0 All-IP Core Network Multimedia Domain - IP Multimedia Subsystem Sh Interface; Signaling Flows and Message Contents – Stage 2
- X.S0013-011-A v1.0 All-IP Core Network Multimedia Domain - Sh Interface Based on Diameter Protocol; Protocol Details – Stage 3
- X.S0016-000-B v1.0 3GPP2 MMS Specification Overview Multimedia Messaging System Specification
- X.S0016-000-C v1.0 Multimedia Messaging Service - Overview
- X.S0028-000-0 v1.0 cdma2000 Packet Data Services: Wireless Local Area Network (WLAN) Interworking - List of Parts
- X.S0028-100-0 v1.0 cdma2000 Packet Data Services: Wireless Local Area Network (WLAN) Interworking - Access to Internet
- X.S0028-200-0 v1.0 cdma2000 Packet Data Services: Wireless Local Area Network (WLAN) Interworking - Access to Operator Service and Mobility

IETF References

- RFC 1594 (March 1994): “FYI on Questions and Answers to Commonly asked “New Internet User” Questions”
- RFC 2104 (February 1997): “HMAC: Keyed-Hashing for Message Authentication”
- RFC 2401 (November 1998): “Security Architecture for the Internet Protocol”
- RFC 2403 (November 1998): “The Use of HMAC-MD5-96 within ESP and AH”
- RFC 2404 (November 1998): “The Use of HMAC-SHA-1-96 within ESP and AH”
- RFC 2405 (November 1998): “The ESP DES-CBC Cipher Algorithm With Explicit IV”
- RFC 2451 (November 1998): “The ESP CBC-Mode Cipher Algorithms”
- RFC 3526 (May 2003): “More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)”
- RFC 3539: (June 2003): “Authentication, Authorization and Accounting (AAA) Transport Profile”
- RFC 3588 (September 2003): “Diameter Base Protocol”

- RFC 3602 (September 2003): "The AES-CBC Cipher Algorithm and Its Use with IPSec"
- RFC 3706 (February 2004): "A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers"
- RFC 3775 (June 2004): "Mobility Support in IPv6"
- RFC 4187 (January 2006): "Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement"
- RFC 4301 (December 2005): "Security Architecture for the Internet Protocol"
- RFC 4302 (December 2005): "IP Authentication Header"
- RFC 4303 (December 2005): "IP Encapsulating Security Payload (ESP)"
- RFC 4305 (December 2005): "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)"
- RFC 4306 (December 2005): "Internet Key Exchange (IKEv2) Protocol"
- RFC 4307 (December 2005): "Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)"
- RFC 4308 (December 2005): "Cryptographic Suites for IPSec"
- RFC 4718 (October 2006): "IKEv2 Clarifications and Implementation Guidelines"
- RFC 4835 (April 2007): "Cryptographic Algorithm Implementation RFC Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)"

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

Chapter 26
PDG/TTG Overview

This chapter contains general overview information about the Packet Data Gateway/Tunnel Termination Gateway (PDG/TTG), including:

- Product Description
- Network Deployment(s) and Interfaces
- Features and Functionality
- Features Not Supported in This Release
- How the PDG/TTG Works
- Supported Standards

Product Description

The Cisco® Packet Data Gateway/Tunnel Termination Gateway (PDG/TTG) enables mobile operators with 3G UMTS wireless data networks to provide Fixed Mobile Convergence (FMC) services to subscribers with dual-mode handsets and dual-mode access cards via WiFi access points. The PDG/TTG makes it possible for operators to provide secure access to the operator's 3G network from a non-secure network, reduce the load on the macro wireless network, enhance in-building wireless coverage, and make use of existing backhaul infrastructure to reduce the cost of carrying wireless calls.

The PDG is a network element that provides WLAN-to-3GPP network interworking. This interworking is achieved by the subscriber UEs in the WLAN initiating an IKEv2/IPSec tunnel toward the PDG, which functions as a security gateway that provides the UEs a secure connection to the Packet Data Network (PDN).

The TTG is a network element that enables PDG functionality for existing GGSN deployments. The TTG and a subset of existing GGSN functions work together to provide PDG functionality to the subscriber UEs in the WLAN.

Platform Requirements

The PDG/TTG software runs on a Cisco® ASR 5000 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the installation guide for the chassis and/or contact your Cisco account representative.

Licenses

The PDG/TTG is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the "Managing License Keys" section of the "Software Management Operations" chapter in the *System Administration Guide*.

Summary of PDG/TTG Features and Functions

The TTG features and functions include:

- PDG service
- PDG mode
- TTG mode
- IKEv2 and IP Security (IPSec) encryption
- Multiple digital certificate selection based on APN
- Subscriber traffic policing for IPSec access
- DSCP marking for IPSec access
- WLAN access control
- RADIUS and Diameter support
- EAP fast re-authentication
- Pseudonym NAI support

- Multiple APN support for IPSec access
- Multiple authenctication support
- Lawful intercept
- IMS emergency call handling
- Session recovery support
- Congestion control
- Bulk statistics
- Threshold crossing alerts (TCAs)
- DIAMETER authentication/authorization support
- QCI-level QoS configuration support
- AAA mediation accounting and offline charging

Network Deployment(s) and Interfaces

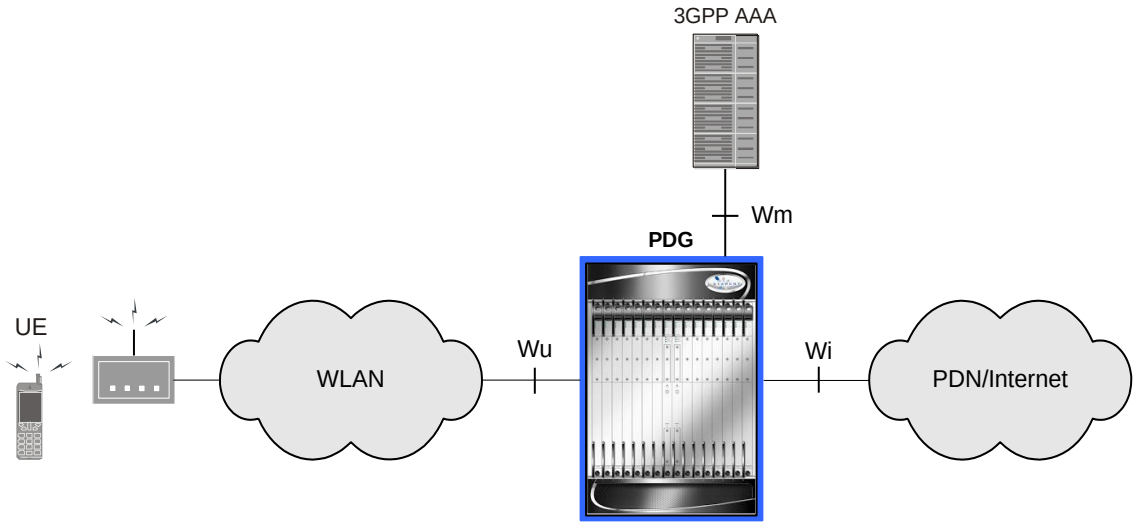
This section describes the PDG/TTG as it functions in a GPRS/UMTS data network.

The PDG in a GPRS/UMTS Data Network

The PDG is a network element that provides WLAN-to-3GPP network interworking. This interworking is achieved by the subscriber UEs in the WLAN access network initiating an IKEv2/IPSec tunnel to the 3GPP PDG, which functions as a security gateway that provides the UEs a secure connection to the Packet Data Network (PDN) or Internet.

The following figure shows a sample PDG implementation.

Figure 198. Sample PDG Implementation



In the implementation above, the subscriber UEs first gain access to the WLAN via a Wi-Fi access point. The UEs get an IP address from the WLAN and use this IP address for communication over the access IP network. To gain access to the PDN/Internet, the UEs use IKEv2 protocol to request a secure IPSec tunnel from the PDG. The UEs may get the IP address of the PDG either from a statically configured IP address, or they may use DNS resolution. During IKEv2 negotiation, the PDG allocates a remote IP address for each subscriber UE to use to access the PDN/Internet. The PDG binds each UE's local IP address with its allocated remote IP address.

The authorization decision for tunnel establishment for access to a 3GPP W-APN belongs to the 3GPP AAA server. Upon authorization, the PDG terminates a secure IPSec tunnel for each WLAN UE subscriber session established over the Wu reference point. The UE establishes a corresponding connection over the Wi reference point toward the PDN/Internet after the call is set up by the PDG.

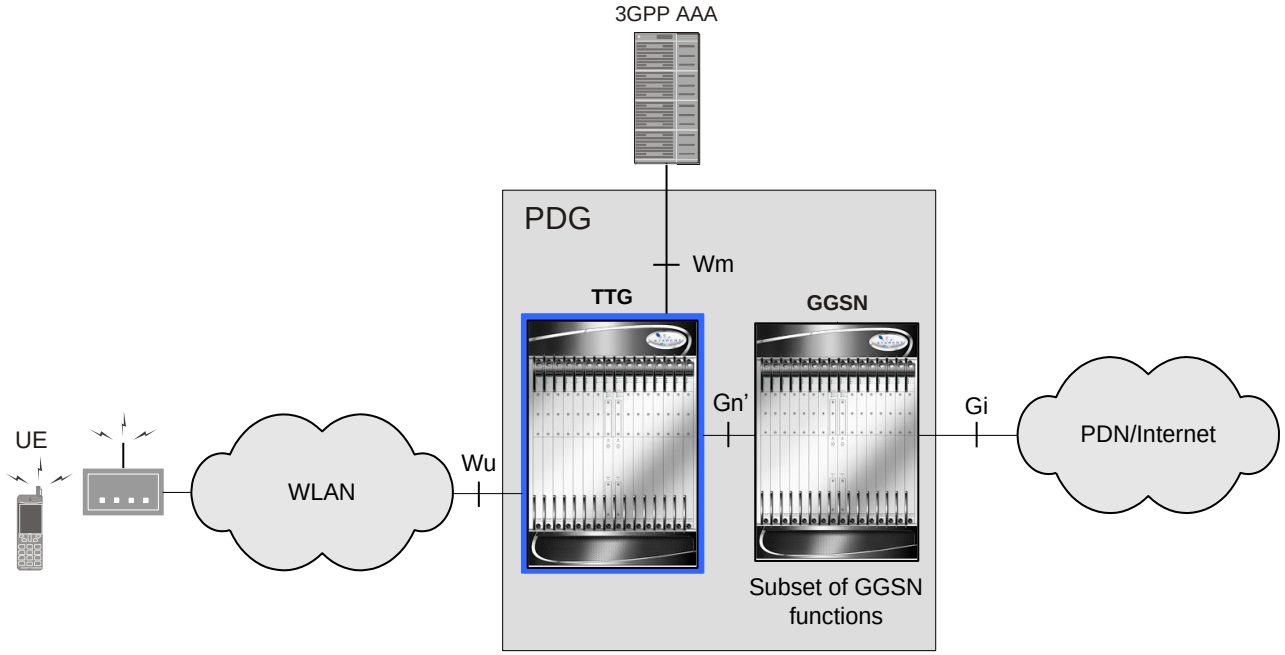
The TTG in a GPRS/UMTS Data Network

The TTG is a GPRS/UMTS network element that enables the implementation of PDG functionality in existing GGSN deployments. It achieves this by using a subset of the Gn reference point called the Gn' (Gn prime) reference point to

connect to currently-deployed GGSNs. This provides the means by which GPRS mobile operators can offer new services to current subscribers by implementing PDG functionality using existing infrastructure.

The following figure shows a PDG implementation that consists of the TTG working together with a currently-deployed GGSN. In this implementation, only a subset of the GGSN functionality is used.

Figure 199. The TTG and GGSN in a PDG Implementation



In the implementation above, the TTG terminates a secure IPSec tunnel for each WLAN UE subscriber session established over the Wu reference point. The TTG also establishes a corresponding GTP (GPRS Tunneling Protocol) tunnel over the Gn' reference point to the GGSN. The TTG and the subset of GGSN functions work together to provide PDG functionality to the UEs in the WLAN.

The TTG functions as an SGSN in the GPRS/UMTS network to provide an SGTP (SGSN GPRS Tunneling Protocol) service. The SGTP service enables the TTG to use GTP over the Gn' interface to carry packet data between itself and the GGSN. The GGSN establishes a corresponding connection over the Gi reference point toward the PDN/Internet.

PDG/TTG Logical Network Interfaces (Reference Points)

The following table provides descriptions of the logical network interfaces supported by the PDG/TTG in a GPRS/UMTS data network.

Table 89. PDG/TTG Logical Network Interfaces

Interface	Description
Wu	The Wu reference point is located between the WLAN UE and the PDG/TTG. The Wu interface carries the secure IPSec tunnels between the UEs in the WLAN and the PDG/TTG. The IPSec tunnels carry the ESP (Encapsulating Security Payload) packets between the UEs and the PDG/TTG.

Interface	Description
Wm	The Wm reference point is located between the 3GPP AAA server and the PDG/TTG. The PDG/TTG uses this reference point to retrieve tunneling attributes and UE IP configuration parameters.
Wi (PDG mode only)	The Wi reference point is located between the PDG and the Packet Data Network (PDN) for WLAN IP access.
Gn' (TTG mode only)	The Gn' reference point is located between the TTG and the GGSN. To provide PDG functionality in existing GGSN deployments, the TTG functions as an SGSN. For every IPSec tunnel that is established between the TTG and a WLAN UE, the TTG initiates a PDP context and a corresponding GTP tunnel over the Gn' interface to the GGSN. The TTG forwards the W-APN and IMSI of the WLAN UE to the GGSN in the Create-PDP-Context-Request message. The following messages are supported over the Gn' reference point: • Create PDP Context Request / Response • Update PDP Context Request / Response • Delete PDP Context Request / Response • Error Indication • Version Not Supported • GTP Payload Forwarding • GTP Echo
Gi (TTG mode only)	The Gi reference point is located between the GGSN and the Packet Data Network (PDN) for WLAN IP access when the PDG/TTG is in TTG mode.

Features and Functionality

This section describes the features and functions supported by the PDG/TTG software.

The following features are supported and described in this section:

- PDG Service
- PDG Mode
- TTG Mode
- IKEv2 and IP Security (IPSec) Encryption
- Multiple Digital Certificate Selection Based on APN
- Subscriber Traffic Policing for IPSec Access
- DSCP Marking for IPSec Access
- WLAN Access Control
- RADIUS and Diameter Support
- EAP Fast Re-authentication Support
- Pseudonym NAI Support
- Multiple APN Support for IPSec Access
- Multiple Authentication Support
- Lawful Intercept
- IMS Emergency Call Handling
- Session Recovery Support
- Congestion Control
- Bulk Statistics
- Threshold Crossing Alerts
- AAA Mediation Accounting and Offline Charging

PDG Service

The PDG service provides both PDG and TTG functionality, operating in either PDG mode or TTG mode . The PDG service enables the UEs in the WLAN to connect with the core network elements via a secure IPSec interface.

During configuration, you create the PDG service in a PDG context, which is a routing domain on the ASR 5000. PDG context and service configuration includes the following main steps:

- **Configure the IPv4 address for the service:** This is the IP address of the TTG to which the UEs in the WLAN attempt to connect, sending IKEv2 messages to this address to establish IPSec tunnels.
- **Configure the name of the crypto template for IKEv2/IPSec:** A crypto template is used to define an IKEv2/IPSec policy. It includes IKEv2 and IPSec parameters for keepalive, lifetime, NAT-T, and cryptographic and authentication algorithms. There must be one crypto template per PDG service.

- **The name of the EAP profile:** The EAP profile defines the EAP authentication method and associated parameters.
- **Multiple authentication support:** Multiple authentication is specified as a part of crypto template configuration.
- **IKEv2 and IPSec transform sets:** Transform set defines the negotiable algorithms for IKE SAs and Child SAs to enable calls to connect to the PDG/TTG.
- **The setup timeout value:** This parameter specifies the session setup timeout timer value. The PDG/TTG terminates a UE connection attempt if the UE does not establish a successful connection within the specified timeout period.
- **Max-sessions:** This parameter sets the maximum number of subscriber sessions allowed by this PDG service.

PDG Mode

The PDG mode of operation uses IKEv2/IPsec tunnels to deliver packet data services over untrusted Wireless Local Area Networks (WLANs) with connectivity to the Internet or managed networks.

Important: PDG mode is not fully qualified and is not supported for field deployment. It is available for lab use only.

In PDG mode, the system terminates an IPSec tunnel for each WLAN UE subscriber session established over the Wu reference point. The UE establishes a corresponding connection over the Wi reference point toward the PDN/Internet after the call is set up by the PDG.

TTG Mode

The TTG mode of operation uses IKEv2/IPsec tunnels to deliver packet data services over untrusted WLANs with connectivity to the Internet or managed networks.

In TTG mode, the system terminates an IPSec tunnel for each WLAN UE subscriber session established over the Wu reference point. The TTG also establishes a corresponding GTP tunnel over the Gn' reference point to the GGSN. The TTG and a subset of GGSN functions work together to provide PDG functionality to the WLAN UEs. In this configuration, the GGSN sees the TTG as an SGSN, and no additional changes are required at the GGSN to support this functionality.

IKEv2 and IP Security (IPSec) Encryption

The PDG/TTG supports IKEv2 and IPSec encryption using IPv4 addressing. IKEv2 and IPSec encryption enables network domain security for all IP packet-switched networks in order to provide confidentiality, integrity, authentication, and anti-replay protection. These capabilities are insured through use of cryptographic techniques.

IKEv2 and IP Security (IPSec) encryption includes the following options:

- **IKEv2 encryption protocols:** AES-CBC with 128 bits, AES-CBC with 256 bits, 3DES-CBC, and DES-CBC
- **IKEv2 pseudo-random functions:** PRF-HMAC-SHA1, PRF-HMAC-MD5
- **IKEv2 integrity:** HMAC-SHA1-96, HMAC-MD5
- **IKEv2 Diffie-Hellman groups:** 1, 2, 5, and 14

- **IPSec ESP (Encapsulating Security Payload) encryption:** AES-CBC with 128 bits, AES-CBC with 256 bits, 3DES-CBC, and DES-CBC
- **IPSec integrity:** HMAC-SHA1-96, HMAC-MD5
- **IKEv2 and IPSec rekeying**

Multiple Digital Certificate Selection Based on APN

Selecting digital certificates based on Access Point Name (APN) allows you to apply digital certificates per the requirements of each APN and associated packet data network. A digital certificate is an electronic credit card that establishes a subscriber's credentials when doing business or other transactions on the Internet. Some digital certificates conform to ITU-T standard X.509 for a Public Key Infrastructure (PKI) and Privilege Management Infrastructure (PMI). X.509 specifies standard formats for public key certificates, certificate revocation lists, attribute certificates, and a certification path validation algorithm.

During session establishment, the PDG/TTG can select a digital certificate from multiple certificates based on the APN. The selected certificate is associated with the APN that the WLAN UE includes in the IDr payload of the first IKE_AUTH_REQ message.

When configuring APN-based certificate selection, ensure that the certificate names match the associated APNs exactly. The PDG/TTG can then examine each APN received in the IDr payload and select the correct certificate.

The PDG/TTG generates an SNMP notification when the certificate is within 30 days of expiration and approximately once a day until a new certificate is provided. Operators need to generate a new certificate and then configure the new certificate using the system's CLI. The certificate is then used for all new sessions.

Subscriber Traffic Policing for IPSec Access

Traffic policing allows you to manage bandwidth usage on the network and limit bandwidth allowances to subscribers. QoS configurations are stored per QCI level (Quality of Service class identities.

Traffic policing enables the configuration and enforcement of bandwidth limitations on individual subscribers of a particular traffic class in a 3GPP service. Bandwidth enforcement is configured and enforced independently in the downlink and uplink directions.

When configured in the Subscriber Configuration Mode of the system's CLI, the PDG/TTG performs traffic policing. However, if the GGSN changes the QoS via an Update PDP Context Request, the PDG/TTG uses the QoS values from the GGSN.

Per RFC 2698, a Token Bucket Algorithm is used to implement the traffic policing feature on the PDG/TTG. The following criteria is used when determining how to mark a packet:

- **Committed Data Rate (CDR):** The guaranteed rate (in bits per second) at which packets can be transmitted/received for the subscriber during the sampling interval. Note that the committed (or guaranteed) data rate does not apply to the Interactive and Background traffic classes.
- **Peak Data Rate (PDR):** The maximum rate (in bits per second) that subscriber packets can be transmitted/received for the subscriber during the sampling interval.

Using negotiated QoS data rates, the system calculates the burst size, which is the maximum number of bytes that can be transmitted/received for the subscriber during the sampling interval for both committed and peak rate conditions. The committed burst size (CBS) and peak burst size (PBS) for each subscriber depends on the guaranteed bit rate (GBR) and maximum bit rate (MBR) respectively. This represents the maximum number of tokens that can be placed in the subscriber's "bucket". The burst size is the bucket size used by the Token Bucket Algorithm.

Tokens are removed from the subscriber's bucket based on the size of the packets being transmitted/received. Every time a packet arrives, the system determines how many tokens need to be added (returned) to a subscriber's CBS (and PBS) bucket. This value is derived by computing the product of the time difference between incoming packets and the CDR (or PDR). The computed value is then added to the tokens remaining in the subscriber's CBS (or PBS) bucket. The total number of tokens can not be greater than the burst size. If the total number of tokens is greater than the burst size, the number is set to equal the burst size.

After passing through the Token Bucket Algorithm, the packet is internally classified with a color, as follows:

- If there are not enough tokens in the PBS bucket to allow a packet to pass, the packet is considered to be in violation and is marked "red" and the violation counter is incremented by one.
- If there are enough tokens in the PBS bucket to allow a packet to pass, but not in the CBS "bucket", then the packet is considered to be in excess and is marked "yellow", the PBS bucket is decremented by the packet size, and the exceed counter is incremented by one.
- If there are more tokens present in the CBS bucket than the size of the packet, then the packet is considered as conforming and is marked "green" and the CBS and PBS buckets are decremented by the packet size.

The system can be configured with actions to take for red and yellow packets. Any of the following actions may be specified:

- **Drop:** The offending packet is discarded.
- **Transmit:** The offending packet is passed.
- **Lower the IP Precedence:** The packet's ToS octet is set to "0", thus downgrading it to Best Effort, prior to passing the packet.

Different actions can be specified for red and yellow, as well as for uplink and downlink directions and for different 3GPP traffic classes.

DSCP Marking for IPSec Access

The Differentiated Service Code Point (DSCP) marking feature on the PDG/TTG provides support for more granular configuration of DSCP marking.

Important: Support for storing the QoS configuration on a per-QCI level instead of the class level is not fully qualified and is not supported for field deployment. It is available for lab use only.

The PDG/TTG functioning as a TTG can perform DSCP marking of packets sent over the Wu interface in the downlink direction to the WLAN UEs and over the Gn' interface in the uplink direction to the GGSN.

In the PDG Service Configuration Mode of the system's CLI, you use the **`ip qos-dscp`** command to control DSCP markings for downlink packets sent over the Wu interface in IPSec tunnels, and use the **`ip gnp-qos-dscp`** command to control DSCP markings for uplink packets sent over the Gn' interface in GTP tunnels.

The DSCP markings are applied to the IP header of every transmitted subscriber data packet. DSCP levels can be assigned to specific traffic patterns in order to ensure that the data packets are delivered according to the precedence with which they are tagged. The four traffic patterns have the following order of precedence: background (lowest), interactive, streaming, and conversational (highest).

For the interactive traffic class, the PDG/TTG supports per-gateway service and per-APN configurable DSCP marking for the uplink and downlink directions based on Allocation/Retention Priority in addition to the current priorities.

The following matrix can be used to determine the DSCP markings used based on the configured traffic class and Allocation/Retention Priority:

Table 90. Default DSCP Value Matrix

Allocation Priority	1	2	3
Traffic Handling Priority			
1	ef	ef	ef
2	af21	af21	af21
3	af21	af21	af21

You can assign DSCP to specific traffic patterns to ensure that data packets are delivered according to the precedence with which they are tagged. The diffserv markings are applied to the outer IP header of every GTP data packet. The diffserv marking of the inner IP header is not modified.

The traffic patterns are defined by QCI (1 to 9). Data packets falling under the category of each of the traffic patterns are tagged with a DSCP that further indicate their precedence as shown in the following tables:

Table 91. Class structure for assured forwarding (af) levels

Drop Precedence	Class			
	Class 1	Class 2	Class 3	Class 4
Low	af11	af21	af31	af41
Medium	af12	af22	af32	af41
High	af13	af23	af33	af43

Table 92. DSCP Precedence

Precedence (low to high)	DSCP
0	Best Effort (be)
1	Class 1
2	Class 2
3	Class 3
4	Class 4
5	Express Forwarding (ef)

WLAN Access Control

The PDG/TTG in TTG mode enables WLAN access control by enabling you to limit the number of IKEv2/IPSec tunnels per subscriber session.

Important: WLAN access control is supported in TTG mode only.

In the PDG Service Configuration Mode of the system's CLI, the **max-tunnels-per-ue** command can be used to specify the maximum number of IKEv2/IPSec tunnels per subscriber session.

The number of tunnels per UE is limited by the Network Service Access Point Identifier (NSAPI) range, which is 5 to 15. Hence, the configurable maximum number of tunnels is 11, within the range of 1 to 11, with a default value of 11.

RADIUS and Diameter Support

RADIUS and Diameter support on the PDG/TTG provides a mechanism for performing authentication, authorization, and accounting (AAA) for subscribers.

Important: Diameter is not fully qualified and is not supported for field deployment. It is available for lab use only.

The benefits of using AAA are:

- Higher flexibility for subscriber access control
- Better accounting, charging, and reporting options
- Industry-standard RADIUS and Diameter authentication

The Remote Authentication Dial-In User Service (RADIUS) and Diameter protocols can be used to provide AAA functionality for subscribers. The PDG/TTG supports EAP authentication based on both RADIUS and Diameter protocols.

The AAA functionality on the PDG/TTG provides a wide range of configuration options via AAA server groups, which allow a number of RADIUS/Diameter parameters to be configured in support of the PDG service.

Currently, two types of authentication load-balancing methods are supported: first-server and round-robin. The first-server method sends requests to the highest priority active server. A request will be sent to a different server only if the highest priority server is not reachable. With the round-robin method, requests are sent to all active servers in a round-robin fashion.

The PDG/TTG can detect the status of the AAA servers. Status checking is enabled by configuration in the AAA Server Group Configuration Mode of the system's CLI. Once an AAA server is detected to be down, it is kept in the down state up to a configurable duration of time called the dead-time period. After the dead-time period expires, the AAA server is eligible to be retried. If a subsequent request is directed to that server and the server properly responds to the request, the system makes the server active again.

The PDG/TTG generates accounting messages on successful session establishment. For a TTG session, the system creates an IPSec SA for a subscriber session after it creates the GTP tunnel to the GGSN over the Gn' interface. The TTG sends an accounting START message to the AAA server after successful completion of both GTP tunnel creation on the Gn' interface and IPsec SA creation on the Wu interface.

Additional Command Option for APN Configuration for Diameter AAA

For APN configuration for Diameter AAA on the PDG, the APN profile can specify whether the PDG combines authentication and authorization in the same access request to the AAA server, or sends an authentication access request followed by a separate authorization access request. The command syntax for this authentication command option in the APN Configuration Mode of the system CLI is:

authentication eap initial-access-request { authenticate-authorize | authenticate-only }

When set to authenticate-authorize, the PDG performs EAP authentication and authorization using a Diameter DER-DEA message exchange for all subscribers requesting access to the specified APN. When set to authenticate-only, the PDG performs authentication only using a Diameter DER-DEA message exchange. A successful authentication will be followed by a separate authorization access request via an AAR-AAA message exchange.

This command option applies to the PDG/TTG in PDG mode only.

Important: For more information on AAA configuration, refer to the *AAA and GTPP Interface Administration and Reference*.

EAP Fast Re-authentication Support

When subscriber authentication is performed frequently, it can lead to a high network load, especially when the number of currently connected subscribers is high. To address this issue, the PDG/TTG can employ fast re-authentication, which is a more efficient method than full authentication.

Fast re-authentication is an EAP (Extensible Authentication Protocol) exchange that is based on keys derived from a preceding full authentication exchange. The fast re-authentication mechanism can be used during both EAP-AKA and EAP-SIM authentication.

When fast re-authentication is enabled, the PDG/TTG receives a fast re-auth ID from the UE in the IDi payload of the IKE_AUTH_REQ message. The PDG/TTG sends the fast re-auth ID to the AAA server in an Authentication Request message to initiate fast re-authentication.

During fast re-authentication, the PDG/TTG handles two separate IKE/IPSec SAs, one for the original session and one for re-authentication. The re-authentication SA remains for a very short period until the fast re-authentication is successful. After the successful fast re-authentication, the PDG/TTG assigns the UE with the same IP address. The SGTP service running on the PDG/TTG identifies the original session and replicates the same session using the same IP address assignment. The PDG/TTG then deletes the original session SA.

The AAA server falls back to full authentication in the following scenarios:

- When the AAA server does not support fast re-authentication.
- When the number of times a fast re-authentication is allowed after a successful full authentication exceeds the limit configured on the AAA server.
- When the EAP server does not have the permanent subscriber identity to perform a fast re-authentication.

Pseudonym NAI Support

The PDG/TTG supports the use of pseudonym Network Access Identifiers (NAIs) to protect the identity of subscribers against tracing from unauthorized access networks.

Pseudonym NAIs are allocated to the WLAN UEs by the EAP server along with the last successful full authentication. The EAP server maintains the mapping of pseudonym-to-permanent identity for each subscriber. The UEs store this mapping in non-volatile memory to save it across reboots, and then use the pseudonym NAI instead of the permanent one in responses to identity requests from the EAP server.

Multiple APN Support for IPSec Access

The PDG/TTG supports multiple wireless APNs for the same UE (the same IMSI) for use during subscriber authorization.

To support subscribers while they attempt to access multiple services, the PDG/TTG enables multiple subscriber authorizations via multiple wireless APNs. Each time a UE attempts to access a service, the PDG/TTG receives a new APN from the UE in the IDr payload of its first IKE_AUTH_REQ message, and the PDG/TTG initiates a new authorization as a distinct session.

Multiple Authentication Support

The PDG/TTG operating in PDG mode supports multiple authentication phases per TS 23.234 and RFC 4739. For EAP protocol, the PDG service operates in authenticator pass-through mode. For PAP and CHAP protocols, the PDG service operates in EAP authenticator terminator mode. The PDG exchanges authentication credentials on the access side using EAP-GTC for PAP payloads and EAP-MD5 for CHAP payloads, and on the PDN side, it exchanges the same parameters inside RADIUS or Diameter AAA protocol messages.

Important: Multiple authentication is not fully qualified and is not supported for field deployment. It is available for lab use only.

When an APN access request received from a WLAN UE requires authentication and authorization, the PDG negotiates with the UE to determine whether it supports multiple authentication exchanges in IKEv2 protocol. If the UE supports multiple authentication exchanges, and if the UE requests additional authentication with the 3GPP AAA server after successful initial EAP authentication, the PDG performs the next phase of authentication with the external AAA server.

The PDG requests additional authentication of the UE from the external AAA server as follows:

- If an EAP procedure is required (EAP profile = AKA/SIM), the PSG performs a second phase authentication similar to the first phase authentication using EAP-AKA or EAP-SIM. If the PDG receives a Legacy-Nak response from the UE, the PDG sends an EAP-Failure message to the UE.
- If a PAP procedure is required (EAP profile = GTC), the PDG sends an EAP-GTC request to the UE. Upon receiving an EAP-GTC response from the UE within an IKE_AUTH request, the PDG performs the next phase authentication with the AAA server. If the PDG receives a Legacy-Nak response containing the type EAP-MD5, and if the specified W-APN allows it, the PDG changes the authentication and authorization procedure to CHAP. If the specified W-APN does not allow CHAP procedures or if the PDG receives a Legacy-Nak response that does not contain EAP-MD5, the PDG sends an EAP-Failure message to the UE.
- If a CHAP procedure is required (EAP profile = MD5), the PDG sends an MD5-Challenge request to the UE. Upon receiving an MD5-Challenge response from the UE within an IKE_AUTH request, the PDG performs the next phase authentication with the AAA server. If the PDG receives a Legacy-Nak response containing the type EAP-GTC, and if the specified W-APN allows it, the PDG changes the authentication and authorization procedure to PAP. If the specified W-APN does not allow PAP procedures or if the PDG receives a Legacy-Nak response that does not contain EAP-GTC, the PDG sends an EAP-Failure message to the UE.

Fast Re-authentication and Pseudonym Re-authentication

After a successful call set-up, if the UE sends a re-authentication request to the PDG using fast re-authentication or pseudonym re-authentication (received in the first phase EAP authentication in the IDi payload of the IKE_AUTH message), the PDG performs EAP-based re-authentication with the 3GPP AAA server. If the UE sends an ANOTHER_AUTH_FOLLOWS Notify payload and IDi for a second phase authentication after that, the PDG performs

a second phase authentication with the external AAA server. If the UE sends a re-authentication request to the PDG using fast re-authentication or pseudonym re-authentication (received in the second phase EAP authentication in the IDi payload of the IKE_AUTH message), the PDG drops the call. In the case of PAP and CHAP, neither the UE nor the PDG initiate re-authentication for a second phase authentication directly.

Re-authorization and RADIUS CoA Handling

An AAA-initiated change-of-authorization / re-authorization for the external AAA server is handled in the same way as is handled for the 3GPP AAA server.

Lawful Intercept

The Cisco Lawful Intercept feature is supported on the PDG/TTG. Lawful Intercept is a license-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your Cisco account representative.

IMS Emergency Call Handling

The PDG/TTG supports IMS emergency call handling per 3GPP TS 33.234. This feature is enabled by configuring a special WLAN access point name (W-APN), which includes a W-APN network identifier for emergency calls (sos, for example), and can be configured with no authentication.

The DNSs in the network are configured to resolve the special W-APN to the IP address of the PDG/TTG. When a WLAN UE initiates an IMS emergency call, the UE sends a W-APN that includes the same W-APN network identifier (sos) as the one that is configured on the PDG/TTG. This W-APN network identifier is prefixed to the W-APN operator identifier per 3GPP TS 23.003. The W-APN operator identifier sent by the UE must match the PLMN ID (MCC and MNC) that is configured on the PDG/TTG (visited network). When the PDG/TTG receives the W-APN from the UE in the IDr, the PDG/TTG marks the call as an emergency call and proceeds with call establishment, even in the event of an authentication or EAP failure from the AAA/EAP server.

If the PDG/TTG detects that an old IKE SA for the special W-APN already exists, it deletes the IKE SA and sends an INFORMATIONAL message with a Delete payload to the WLAN UE to delete the old IKE SA on the UE.

Session Recovery Support

The session recovery feature provides seamless failover and nearly instantaneous reconstruction of subscriber session information in the event of a hardware or software fault within the same chassis, preventing a fully-connected user session from being dropped.

Important: Use of the session recovery feature requires that a valid license key be installed. Contact your Cisco account representative for detailed information on specific licensing requirements.

Session recovery is performed by mirroring key software processes (the IPSec manager, session manager, and AAA manager, for example) on the PDG/TTG. These mirrored processes remain in an idle state (in standby mode), where they perform no processing until they may be needed in the case of a software failure (a session manager task aborts, for example). The system spawns new instances of standby mode sessions and AAA managers for each active control processor being used.

Additionally, other key system-level software tasks such as VPN manager are performed on a physically separate packet processing card to ensure that a double software fault (the session manager and the VPN manager fail at same time on same card, for example) cannot occur. The packet processing card used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled. At a minimum, four packet processing cards (3 active and 1 standby) are required on the chassis to support the session recovery feature.

Congestion Control

Congestion control allows you to set policies and thresholds and specify how the system reacts when faced with a heavy load condition.

Congestion control monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are quickly resolved. However, continuous or large numbers of these conditions within a specific time interval may have an impact on the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establishes limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operation thresholds that are configured for the system as described in the *Thresholding Configuration Guide*. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated.

A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.

- **Port Utilization Thresholds:** If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
- **Port-specific Thresholds:** If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
- **Service Congestion Policies:** Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

Bulk Statistics

Bulk statistics allow operators to choose to view not only statistics that are of importance to them, but to also configure the format in which they are presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. The following is a partial list of supported schemas:

- **System**: Provides system-level statistics
- **Card**: Provides card-level statistics
- **Port**: Provides port-level statistics
- **PDG**: Provides PDG service statistics
- **APN**: Provides Access Point Name statistics

The system supports the configuration of up to four sets (primary/secondary) of receivers. Each set can be configured to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, system host name, system uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics Server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Important: For more information on bulk statistic configuration, refer to the *Configuring and Maintaining Bulk Statistics* chapter of the *System Administration Guide*.

Threshold Crossing Alerts

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outages. Typically, these conditions are temporary (i.e., high CPU utilization or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

Important: The threshold crossing alerts feature for the PDG/TTG is not fully qualified and is not supported for field deployment. It is available for lab use only.

The system supports threshold crossing alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure a threshold on these resources whereby, should the resource depletion cross the configured threshold, an SNMP trap would be sent.

The following thresholding models are supported by the system:

- **Alert:** A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated, then generated and/or sent again at the end of the polling interval.
- **Alarm:** Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated, then generated and/or sent again at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps:** SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values. Generation of specific traps can be enabled or disabled on the chassis, ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.

- **Logs:** The system provides a facility for which active and event logs can be generated. As with other system facilities, logs are generated messages pertaining to the condition of a monitored value and are generated with a severity level of WARNING. Logs are supported in both the Alert and the Alarm models.
- **Alarm System:** High threshold alarms generated within the specified polling interval are considered outstanding until a condition no longer exists or a condition clear alarm is generated. Outstanding alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

AAA Mediation Accounting and Offline Charging

Offline charging is a process where charging information is collected at the same time as resource use. The charging information is then passed through a chain of charging functions. At the end of the process, CDR files are generated by the network, which are then transferred to the network operator's billing domain.

Important: AAA mediation accounting and offline charging is not fully qualified and is not supported for field deployment. It is available for lab use only.

The charging trigger function (CTF) generates charging events and forwards them to the charging data function (CDF). The CDF, in turn, generates CDRs which are transferred to the charging gateway function (CGF). Finally, the CGF creates CDR files and forwards them to the billing domain. The CTF and CDF are integrated in the PDG, however, the CGF may exist as a separate entity or be integrated with the PDG. If the CGF is external to the PDG then the CDF forwards the CDRs to the CGF using the GTPP protocol.

In the ASR 5000, the PDG is integrated with the CTF and CDF and generates WLAN-CDRs based on triggered events over the Wz interface. The PDG offline charging involves the following functionalities for WLAN 3GPP IP access:

- Charging Trigger Function
- Charging Data Function
- Wz Reference Point

Accounting Session Management

The WLAN-CDR is generated as part of AAAmgr accounting management in PDG. When the Accounting-Mode is set to GTPP, it indicates that the charging is enabled and the Wz reference point is to be used for passing charging records to the CGF.

Triggers for WLAN CDRs Charging Information

The PDG/TTG uses the charging characteristics to determine whether to activate or deactivate CDR generation. The charging characteristics al set the chargeable event conditions, such as the time/volume limits that trigger CDR generation or the addition of information. Multiple charging characteristics profiles may be configured on the PDG to allow different sets of trigger values.

Triggers for WLAN-CDR Closure

The following events trigger closure and sending of a partial WLAN-CDR:

- **Time trigger**: every x seconds configured using **`interval x.`**
- **Volume trigger**: every x octets configured using **`volume x`**
- The command **`gtpp interim now`**

A WLAN-CDR is closed as the final record of a session for the following events:

- normalRelease: UE terminates the IPSec tunnel. The call is handed from one PDG to another PDG.
- **abnormalRelease**: Failure due to multiple software failures.
- **volumeLimit**: The configured volume threshold has been exceeded.
- **timeLimit**: the configured interval has been reached.
- **maxChangeCondition**: the limit for the LOTV containers has been exceeded.
- **managementIntervention**: The command **`gtpp interim now`** has been issued.
- **managementIntervention**: The command **`clear sub all`** has been issued.

Triggers for WLAN-CDRs Charging Information Addition

The List of Traffic Volumes attribute of the WLAN-CDR consists of a set of containers, which are added when specific trigger conditions are met. They identify the volume count per PDP context, and are separated for uplink and downlink traffic:

- **QoS Change**: A change in the QoS results in closing the List of Traffic Data Volumes that were open. The volumes are added to the CDR and a new bearer-specific container is opened.
- **tariffTime**: On reaching the Tariff Time Change, a List of Traffic Data Volumes container is added to the CDR.
- **recordClosure**: A list of List of Traffic Data Volumes containers is added to the WLAN CDR.

Features Not Supported in This Release

The following features are not supported in this PDG/TTG software release:

- Link aggregation
- IPv6
- MPLS
- NAT
- Firewall
- Peer-to-Peer

How the PDG/TTG Works

This section describes the PDG/TTG during connection establishment.

PDG Connection Establishment

The figure below shows the message flow during PDG connection establishment. The table that follows the figure describes each step in the message flow.

Figure 200. PDG Connection Establishment

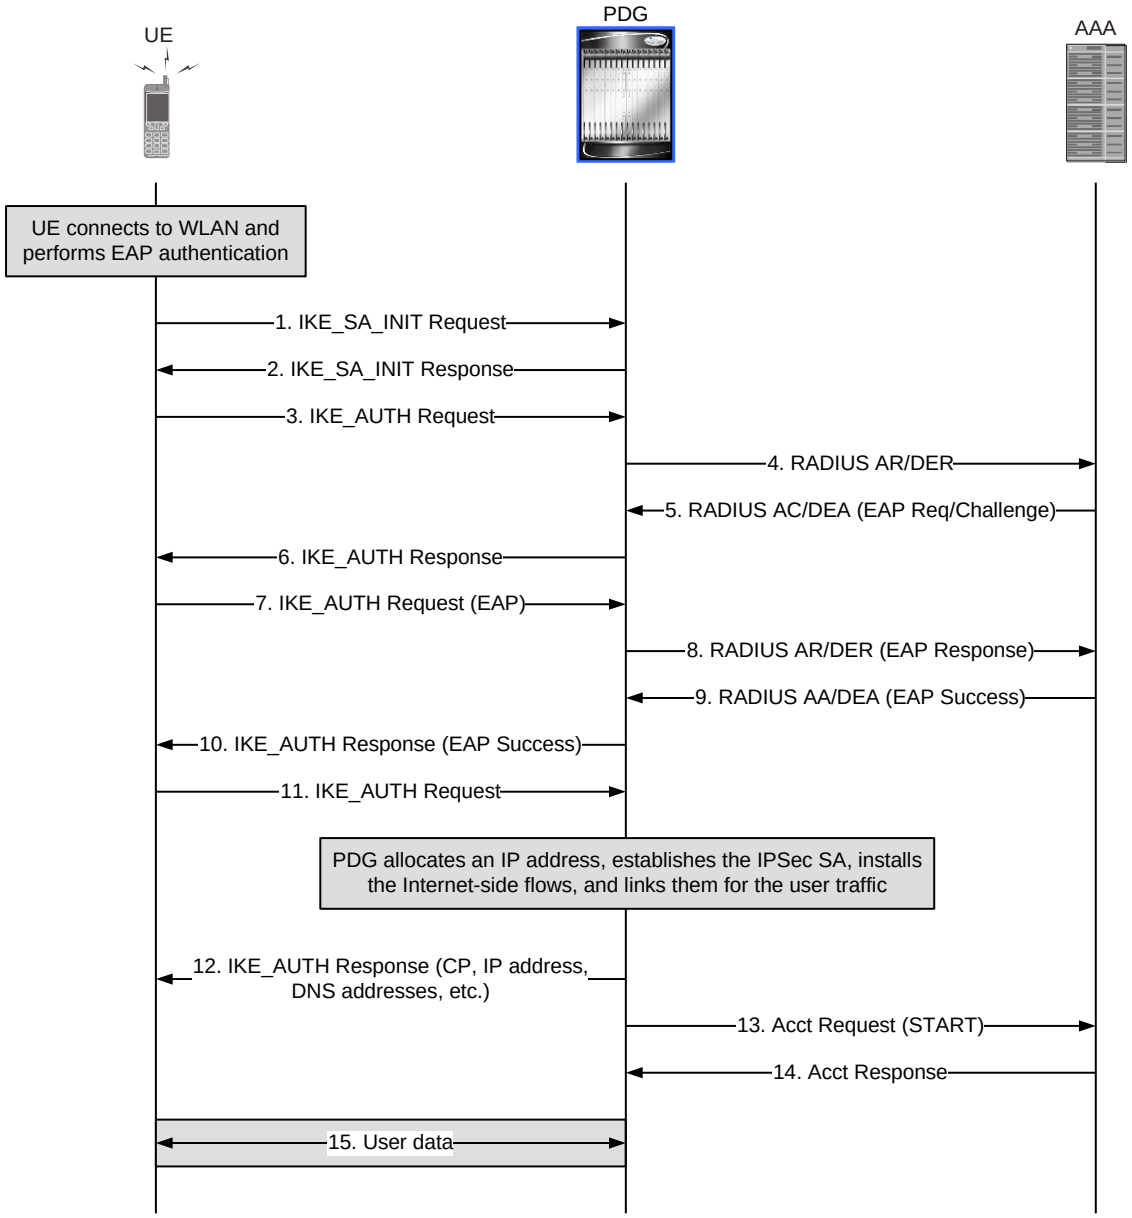


Table 93. PDG Connection Establishment

Step	Description
1.	After the UE connects to the WLAN and performs EAP authentication, the UE initiates an IKEv2/IPSec tunnel by sending an IKE_SA_INIT Request to the PDG. The UE includes the SA, KE, Ni, and NAT-Detection Notify payloads in the IKEv2 exchange.

Step	Description
2.	The PDG processes the IKE_SA_INIT Request for the appropriate PDG service (bound by the destination IP address in the IKEv2 INIT Request). The PDG responds with an IKE_SA_INIT Response with the SA, KE, and Nr payloads, and NAT-Detection Notify payloads. The PDG will start the IKEv2 setup timer when sending the IKE_SA_INIT Response. With the IKEv2 SA INIT exchanges, the WLAN UE negotiates cryptographic algorithms, exchanges the nonce, and performs a Diffie-Hellman exchange.
3.	Upon receiving a successful IKE_SA_INIT Response from the PDG, the UE sends an IKE_ AUTH Request for the first EAP-AKA authentication. The UE also includes an IDi payload, which contains the NAI, SA, TSi, TSr, CP (requesting an IP address and DNS address) payloads. The IDr payload is the requested W-APN. The UE does not include AUTH payload to indicate that it will use the EAP method. The NAI can either be the IMSI or a pseudonym.
4.	Upon receiving the IKE_AUTH Request from UE, the PDG sends an Authentication Request (RADIUS Access Request or DER) message to the AAA server. The PDG sends the Authentication Request message with an EAP (Identity Response) AVP to the AAA server, including the user identity and W-APN. The W-APN information is included in the called-station-id RADIUS attribute in all Access-Request messages towards the AAA server. The PDG includes a parameter indicating that the authentication is being performed for tunnel establishment. This helps the AAA server to distinguish between authentications for WLAN access and authentications for tunnel setup. The PDG starts the session setup timer upon receiving the IKE_AUTH Request from the UE. Note that the PDG sends the W-APN received in the IDr payload in IKEv2 messages as is to the AAA server. This helps the AAA server to look up the authorization database based on the W-APN name. When sending messages to the HLR (or HSS), the AAA server maps the W-APN name into the real APN configured in the HLR (or HSS).
5.	The AAA server initiates the authentication challenge. The user identity is not requested again, as in a normal authentication process, because there is the certainty that the user identity received in the EAP Identity Response message has not been modified or replaced by any intermediate node. This is because the user identity is received via an IKEv2 secure tunnel which can only be decrypted and authenticated by the end points (the PDG and the WLAN UE). The PDG receives a DEA with a Result-Code AVP specifying to continue EAP authentication. For RADIUS, this is an access challenge message. The PDG accepts the EAP-Payload AVP contents.
6.	The PDG sends an IKE_ AUTH Response back to the UE in the EAP payload. Depending upon the configuration, the PDG can include IDr (PDG-ID) and CERT payloads. The PDG allows IDr and CERT configurations in the PDG service. If the PDG service is configured to do so, the PDG can also include an AUTH payload in the IKE_AUTH Response. The UE receives the IKE_AUTH Response from PDG.
7.	Upon receiving the IKE_AUTH Response from the PDG, the UE processes the exchange and sends a new IKE_AUTH Request with an EAP payload. The PDG receives the new IKE_AUTH Request from the UE.
8.	The PDG sends a DER (or RADIUS AR) message to the AAA server. This DER message contains the EAP-Payload AVP with an EAP-AKA challenge or EAP-SIM challenge response and challenge received from the UE.
9.	The AAA server sends the DEA back to the PDG with Result-Code AVP as Success. The EAP-Payload AVP message also contains an EAP result code as Success. The PDG also receives the MSK (keying materials) from the AAA server, which is used for further key computation. When using Diameter, the MSK is encapsulated in the EAP-Master-Session-Key parameter. The AAA server also includes several authorization AVPs. When the checks for an IMS emergency call fail, the AAA Server also sends an Authentication Answer that includes an EAP Failure to the PDG.
10.	The PDG sends the IKE_AUTH Response back to UE with the EAP payload.
11.	The UE sends the final IKE_AUTH Request with the AUTH payload computed from the keys. The PDG uses the MSK to generate the AUTH parameters in order to authenticate the IKE_SA_INIT phase messages. These first two messages had not been authenticated before as there was no key material available yet. When used over IKEv2, the shared secret generated in the EAP exchange (the MSK) is used to generate the AUTH parameters. The PDG processes the IKE_AUTH Request, checks the validity of AUTH payload, allocates an IP address for the UE, establishes an IPSec SA, install the Internet-side flows, and links them for the user traffic.

Step	Description
12.	The PDG sends a Diameter Accounting-Request (START) message to the AAA server. For RADIUS, it sends an Accounting-Start message.
13.	Upon receiving the Accounting-Request message, the AAA server sends an Accounting-Response reply to the PDG. For RADIUS, it sends an Accounting-Stop message.
14.	The PDG sends an IKE_AUTH Response with the AUTH payload computed from the MSK. The PDG assigns the IP address to the UE in the configuration payload along with DNS addresses and other parameters.
15.	The PDG session/IPSec SA is fully established and ready for data transfer.

TTG Connection Establishment

The figure below shows the message flow during TTG connection establishment. The table that follows the figure describes each step in the message flow.

Figure 201. TTG Connection Establishment

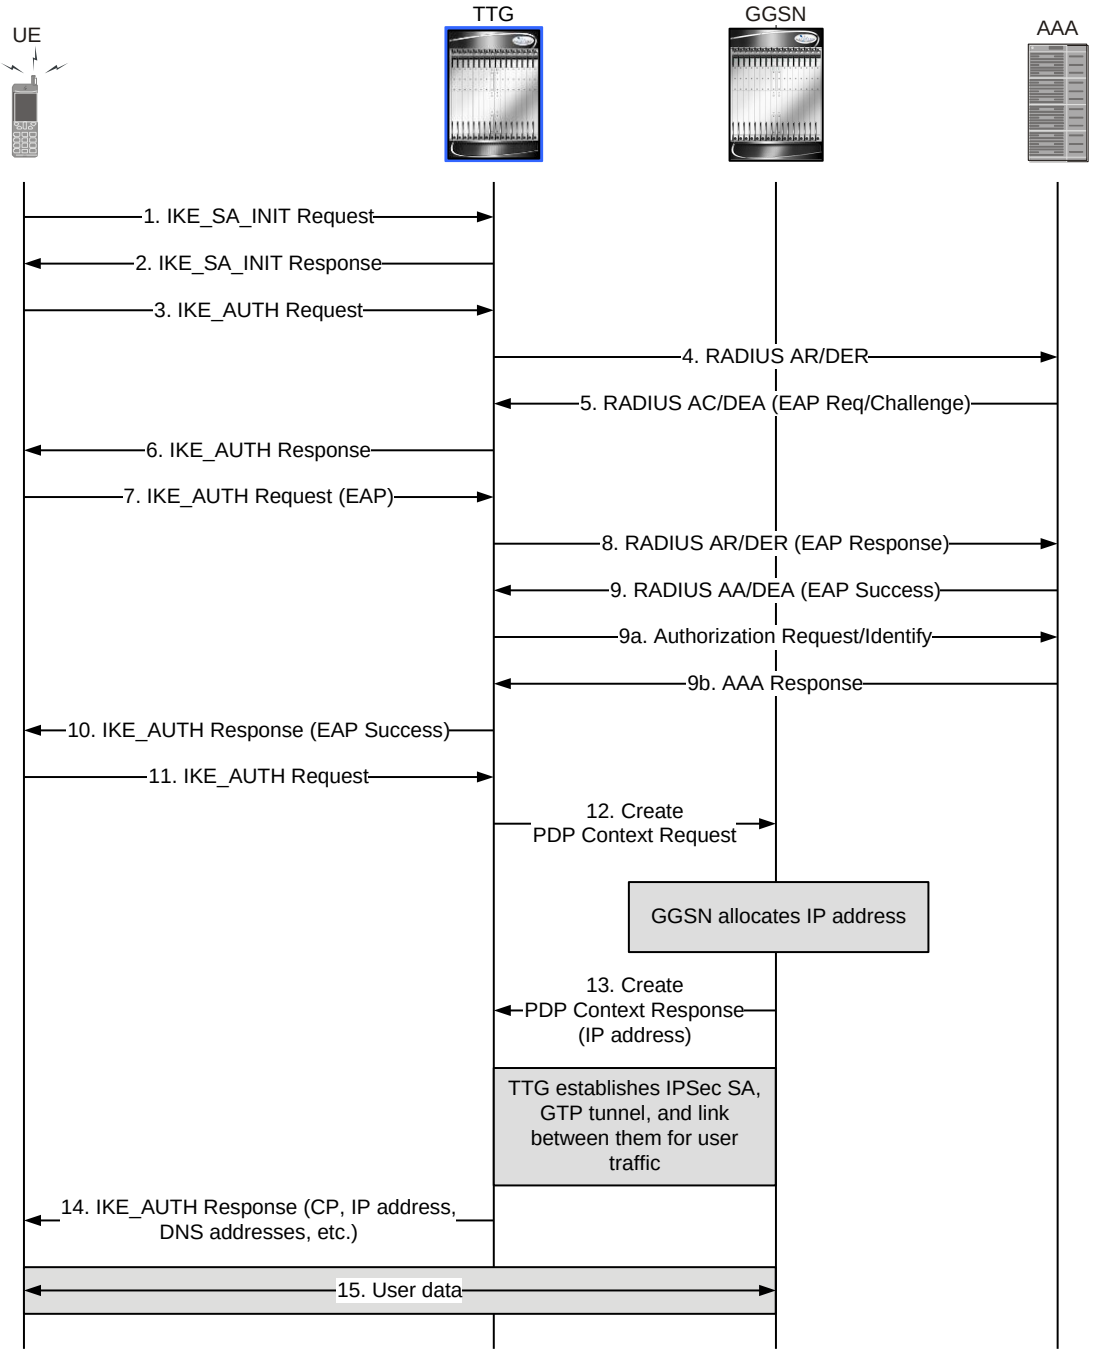


Table 94. TTG Connection Establishment

Step	Description

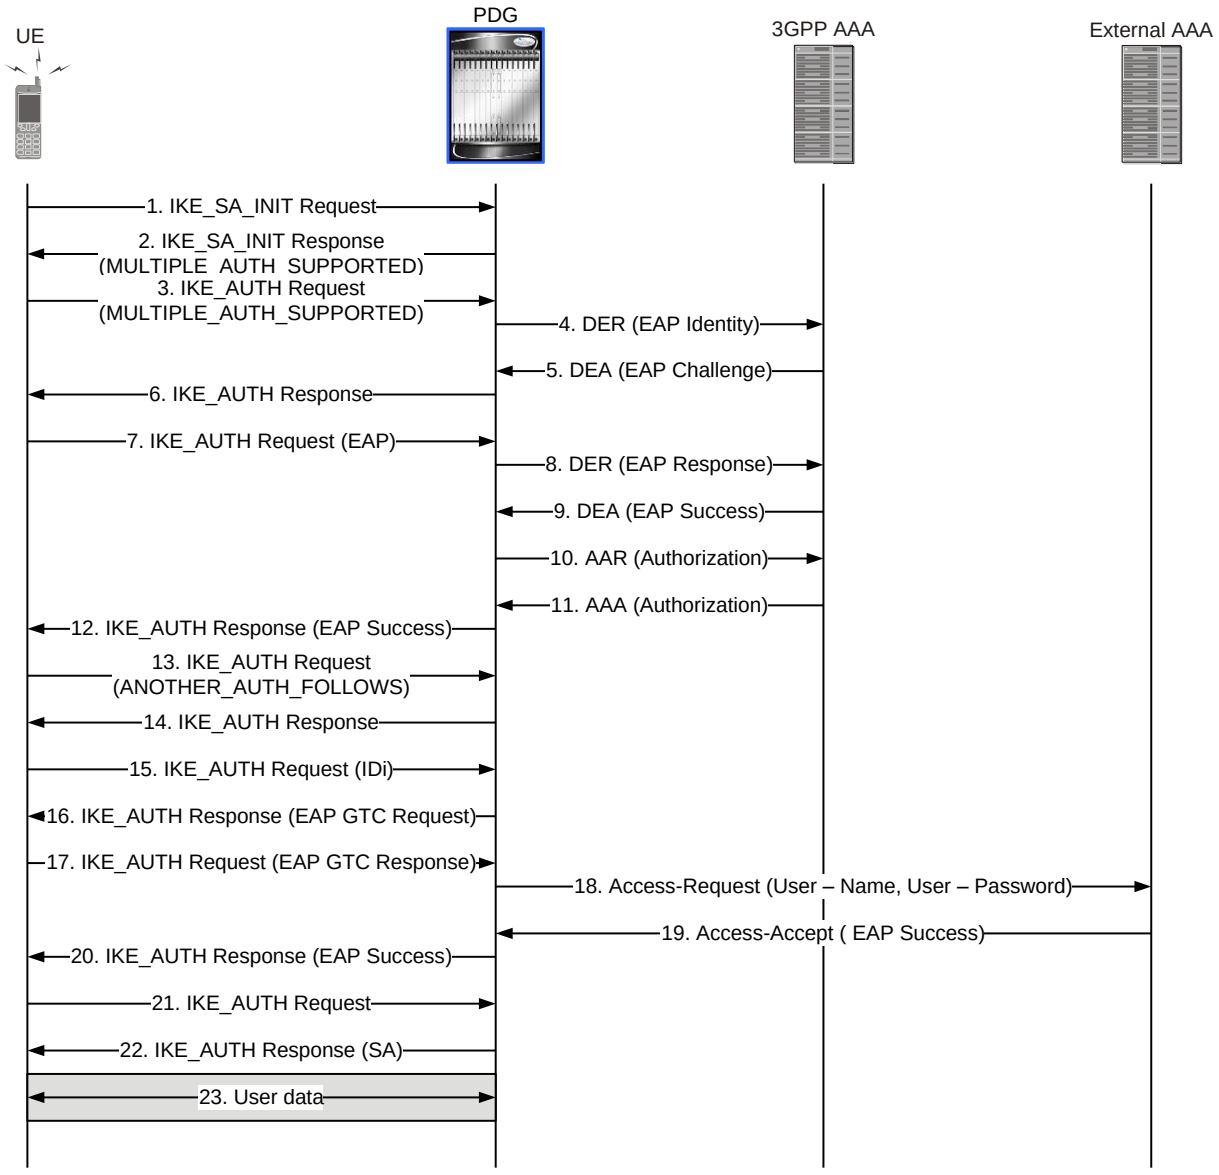
Step	Description
1.	After receiving the IP address of the TTG from the WiFi access point, the UE initiates an IKEv2/IPSec tunnel by sending an IKE_SA_INIT Request to the TTG. The UE includes the SA, KE, Ni, and NAT-Detection Notify payloads in the IKEv2 exchange.
2.	The TTG processes the IKE_SA_INIT Request for the appropriate PDG service (bound by the destination IP address in the IKEv2 INIT Request). The TTG responds with an IKE_SA_INIT Response with the SA, KE, and Nr payloads, and NAT-Detection Notify payloads. The TTG will start the IKEv2 setup timer when sending the IKE_SA_INIT Response. With the IKEv2 SA INIT exchanges, the WLAN UE negotiates cryptographic algorithms, exchanges the nonce, and performs a Diffie-Hellman exchange.
3.	Upon receiving a successful IKE_SA_INIT Response from the TTG, the UE sends an IKE_ AUTH Request for the first EAP-AKA authentication. The UE also includes an IDi payload, which contains the NAI, SA, TSi, TSr, CP (requesting an IP address and DNS address) payloads. The IDr payload is the requested W-APN. The UE does not include AUTH payload to indicate that it will use the EAP method. The NAI can either be the IMSI or a pseudonym.
4.	Upon receiving the IKE_AUTH Request from UE, the TTG sends an Authentication Request (RADIUS Access Request or DER) message to the AAA server. The TTG sends the Authentication Request message with an EAP (Identity Response) AVP to the AAA server, including the user identity and W-APN. The W-APN information is included in the called-station-id RADIUS attribute in all Access-Request messages towards the AAA server. The TTG includes a parameter indicating that the authentication is being performed for tunnel establishment. This helps the AAA server to distinguish between authentications for WLAN access and authentications for tunnel setup. The TTG starts the session setup timer upon receiving the IKE_AUTH Request from the UE. Note that the TTG sends the W-APN received in the IDr payload in IKEv2 messages as is to the AAA server. This helps the AAA server to look up the authorization database based on the W-APN name. When sending messages to the HLR (or HSS), the AAA server maps the W-APN name into the real APN configured in the HLR (or HSS).
5.	The AAA server initiates the authentication challenge. The user identity is not requested again, as in a normal authentication process, because there is the certainty that the user identity received in the EAP Identity Response message has not been modified or replaced by any intermediate node. This is because the user identity is received via an IKEv2 secure tunnel which can only be decrypted and authenticated by the end points (the TTG and the WLAN UE). The TTG receives a DEA with a Result-Code AVP specifying to continue EAP authentication. For RADIUS, this is an access challenge message. The TTG accepts the EAP-Payload AVP contents.
6.	The TTG sends an IKE_ AUTH Response back to the UE in the EAP payload. Depending upon the configuration, the TTG can include IDr (TTG-ID) and CERT payloads. The TTG allows IDr and CERT configurations in the PDG service. If the PDG service is configured to do so, the TTG can also include an AUTH payload in the IKE_AUTH Response. The UE receives the IKE_AUTH Response from TTG.
7.	Upon receiving the IKE_AUTH Response from the TTG, the UE processes the exchange and sends a new IKE_AUTH Request with an EAP payload. The TTG receives the new IKE_AUTH Request from the UE.
8.	The TTG sends a DER (or RADIUS AR) message to the AAA server. This DER message contains the EAP-Payload AVP with an EAP-AKA challenge or EAP-SIM challenge response and challenge received from the UE.

Step	Description
9.	The AAA server sends the DEA back to the TTG with Result-Code AVP as Success. The EAP-Payload AVP message also contains an EAP result code as Success. The TTG also receives the MSK (keying materials) from the AAA server, which is used for further key computation. When using Diameter, the MSK is encapsulated in the EAP-Master-Session-Key parameter. The AAA server also includes several authorization AVPs. When the checks for an IMS emergency call fail, the AAA Server also sends an Authentication Answer that includes an EAP Failure to the TTG. Note that steps 9a. and 9b. (described below) may not be required if authorization attributes or AVPs are present in the Access-Accept message containing the EAP-Success. As explained in step 5 above, if the W-APN is present in all the Access-Request messages from the TTG to the AAA server, the AAA server can use the W-APN to look up the authorization database to retrieve the parameters. If the TTG has done the W-APN-to-real-APN mapping and includes the mapped APN in the AAA messages, the TTG performs steps 9a. and 9b., and includes the W-APN in a separate message after successful EAP-authentication. 9a. The TTG sends an Authorization Request message with an empty EAP AVP, but containing the W-APN, to the AAA server. The AAA server checks the user's subscription information whether the user is authorized to establish a tunnel. The IKE SA counter for that W-APN is incremented. If the maximum number of IKE SAs for that W-APN is exceeded, the AAA server sends an indication to the TTG that established the oldest active IKE SA (it could be the same TTG or a different one) to delete the oldest established IKE SA. The AAA server then updates the counters tracking the active IKE SAs for the W-APN accordingly. 9b. The AAA server sends the AA-Answer to the TTG. The AAA server sends the IMSI within the AA-Answer.
10.	The TTG sends the IKE_AUTH Response back to UE with the EAP payload.
11.	The UE sends the final IKE_AUTH Request with the AUTH payload computed from the keys. The TTG uses the MSK to generate the AUTH parameters in order to authenticate the IKE_SA_INIT phase messages. These first two messages had not been authenticated before as there was no key material available yet. When used over IKEv2, the shared secret generated in the EAP exchange (the MSK) is used to generate the AUTH parameters. The TTG processes the IKE_AUTH Request, checks the validity of AUTH payload, and initiates PDP context activation with the GGSN.
12.	The TTG sends a Create PDP Context Request to the GGSN. The GGSN processes the request and assigns an IP address to the UE.
13.	The GGSN sends a Create PDP Context Response to the TTG. The TTG sets up an IPSec SA.
14.	The TTG sends an IKE_AUTH Response with the AUTH payload computed from the MSK. The TTG assigns the IP address received from the GGSN to the UE in the configuration payload along with DNS addresses and other parameters.
15.	The TTG session/IPSec SA is fully established and ready for data transfer.

Multiple Authentication Using EAP and PAP

The figure below shows the message flow during PDG connection establishment with multiple authentication. This message flow shows the PDG performing first-phase authentication using EAP (EAP-AKA or EAP-SIM) over Diameter protocol and second-phase authentication using PAP (EAP-GTC) over RADIUS protocol.

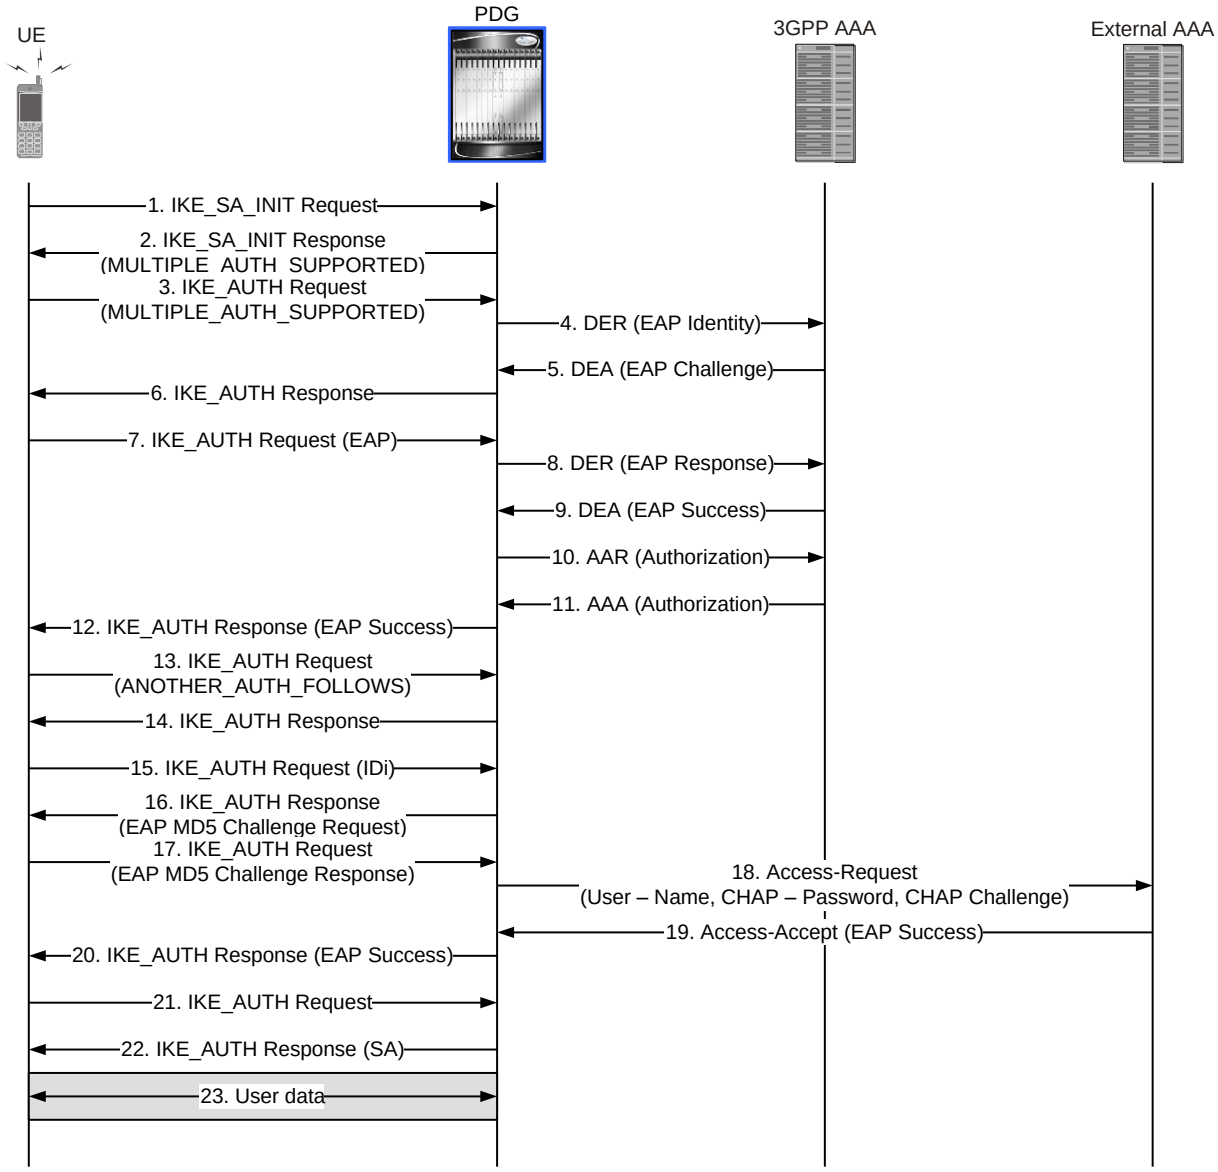
Figure 202. Multiple Authentication Using EAP and PAP



Multiple Authentication Using EAP and CHAP

The figure below shows the message flow during PDG connection establishment with multiple authentication. This message flow shows the PDG performing first-phase authentication using EAP (EAP-AKA or EAP-SIM) over Diameter protocol and second-phase authentication using CHAP (EAP-MD5) over RADIUS protocol.

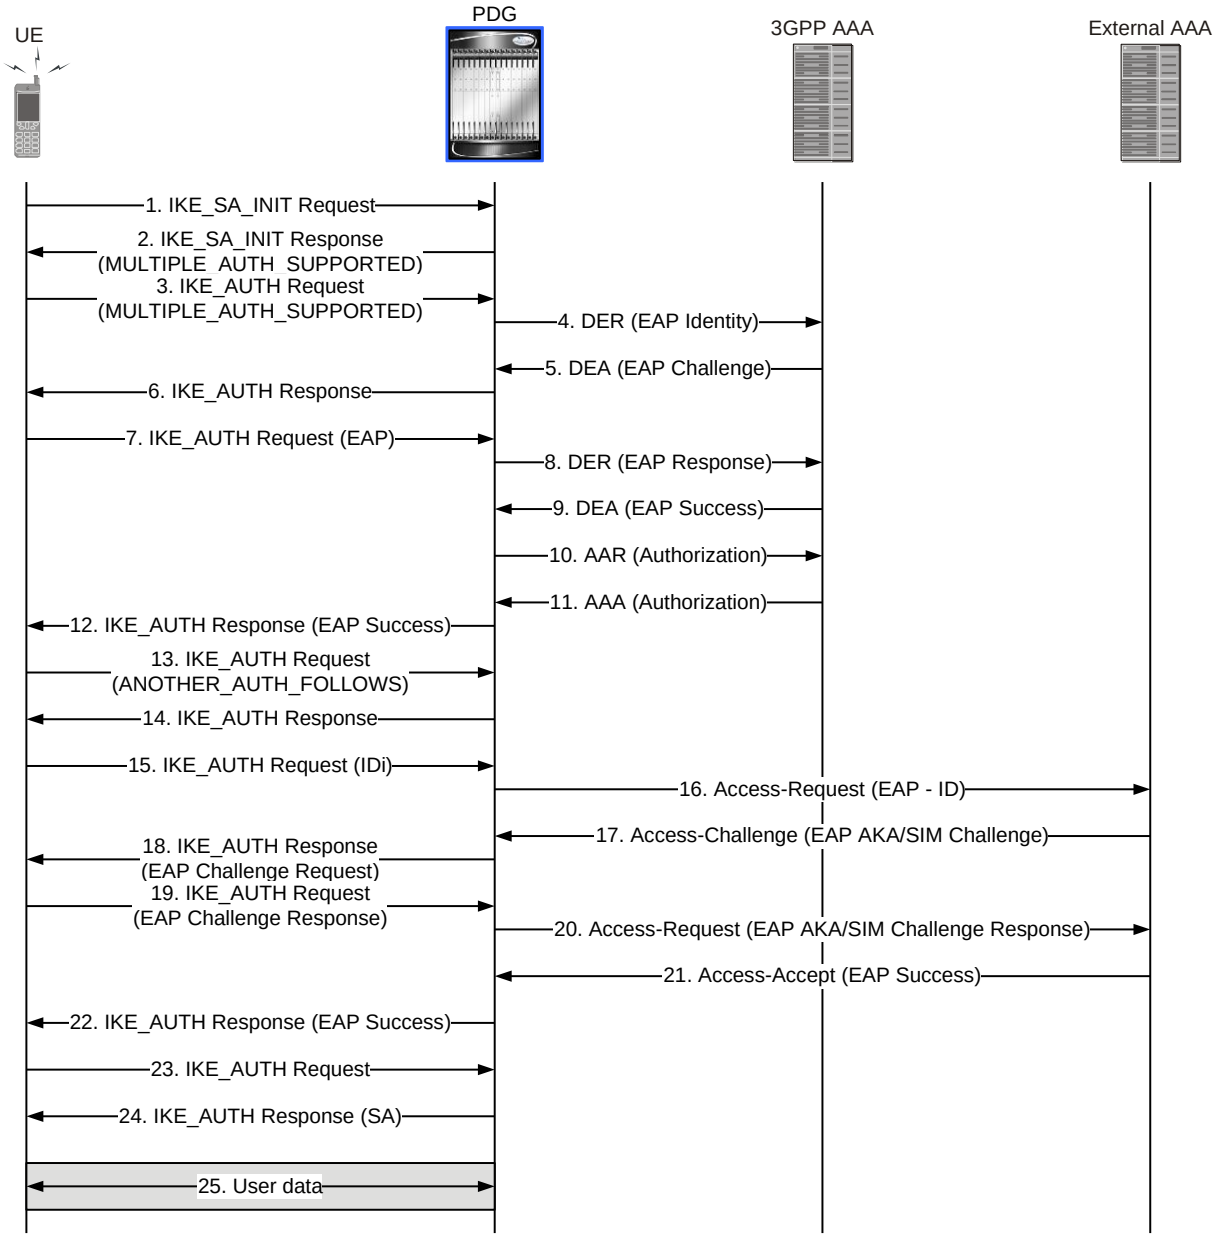
Figure 203. Multiple Authentication Using EAP and CHAP



Multiple Authentication Using EAP and EAP

The figure below shows the message flow during PDG connection establishment with multiple authentication. This message flow shows the PDG performing first-phase authentication using EAP (EAP-AKA or EAP-SIM) over Diameter protocol and second-phase authentication using EAP (EAP-AKA or EAP-SIM) over RADIUS protocol.

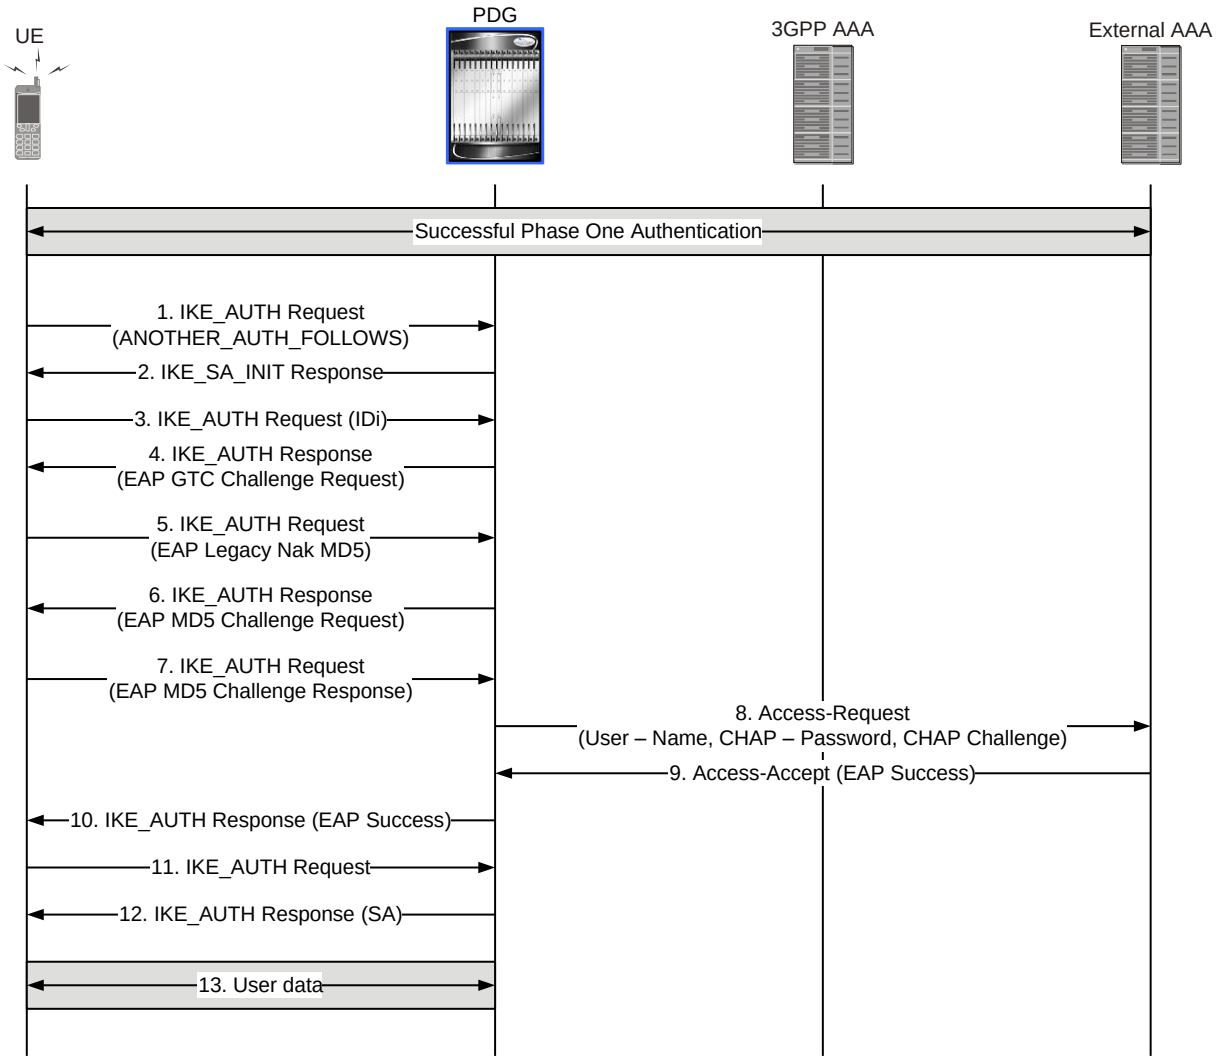
Figure 204. Multiple Authentication Using EAP and EAP



Multiple Authentication with Request from UE for Change of Second Phase Protocol

The figure below shows the message flow during PDG connection establishment with multiple authentication. This message flow shows a scenario in which the UE does not support EAP-MD5. When the PDG sends an EAP-MD5 challenge, the UE sends an EAP Legacy-Nak requesting that the PDG use EAP-GTC instead. The first phase authentication is over Diameter protocol and the second phase authentication is over RADIUS protocol.

Figure 205. Multiple Authentication with Request from UE for Change of Second Phase Protocol



Supported Standards

The PDG/TTG complies with the following standards.

- 3GPP References
- IETF References

3GPP References

- 3GPP TS 22.234 (V8.1.0): "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Requirements on 3GPP system to Wireless Local Area Network (WLAN) interworking (Release 7)".
- 3GPP TS 23.003 (V7.9.0): "3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Numbering, addressing and identification (Release 7)".
- 3GPP TS 23.234 (V6.10.0 and V7.5.0): "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3GPP system to Wireless Local Area Network (WLAN) interworking; System description (Release 6)".
- 3GPP TS 23.327 (V8.4.0): "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Mobility between 3GPP-Wireless Local Area Network (WLAN) interworking and 3GPP systems (Release 8)".
- 3GPP TS 24.234 (V8.3.0): "Group Core Network and Terminals; 3GPP System to Wireless Local Area Network (WLAN) interworking; WLAN User Equipment (WLAN UE) to network protocols; Stage 3 (Release 8)".
- 3GPP TS 29.060 (V7.9.0): "3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface (Release 7)".
- 3GPP TS 29.234 (V8.1.0): "3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; 3GPP system to Wireless Local Area Network (WLAN) interworking; Stage 3 (Release 8)".
- 3GPP TS 32.252 (V7.0.0): "3rd Generation Partnership Project; Technical Specification Group Service and System Aspects; Telecommunication management; Charging management; Wireless Local Area Network (WLAN) charging (Release 7)".
- 3GPP TS 33.234 (V6.9.0): "3rd Generation Partnership Project; Technical Specification Group Service and System Aspects; 3G Security; Wireless Local Area Network (WLAN) interworking security (Release 6)".

IETF References

- RFC 2104 (February 1997): "HMAC: Keyed-Hashing for Message Authentication".
- RFC 2246 (January 1999): "The TLS Protocol, Version 1.0".
- RFC 2401 (November 1998): "Security Architecture for the Internet Protocol".
- RFC 2403 (November 1998): "The Use of HMAC-MD5-96 within ESP and AH".
- RFC 2404 (November 1998): "The Use of HMAC-SHA-1-96 within ESP and AH".
- RFC 2405 (November 1998): "The ESP DES-CBC Cipher Algorithm With Explicit IV".

- RFC 2451 (November 1998): "The ESP CBC-Mode Cipher Algorithms".
- RFC 3526 (May 2003): "More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)".
- RFC 3539: (June 2003): "Authentication, Authorization and Accounting (AAA) Transport Profile".
- RFC 3602 (September 2003): "The AES-CBC Cipher Algorithm and Its Use with IPSec".
- RFC 3706 (February 2004): "A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers".
- RFC 4186 (January 2006): "Extensible Authentication Protocol Method for Global System for Mobile Communications (GSM) Subscriber Identity Modules (EAP-SIM)".
- RFC 4187 (January 2006): "Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA)".
- RFC 4301 (December 2005): "Security Architecture for the Internet Protocol".
- RFC 4302 (December 2005): "IP Authentication Header".
- RFC 4303 (December 2005): "IP Encapsulating Security Payload (ESP)".
- RFC 4305 (December 2005): "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)".
- RFC 4306 (December 2005): "Internet Key Exchange (IKEv2) Protocol".
- RFC 4307 (December 2005): "Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)".
- RFC 4308 (December 2005): "Cryptographic Suites for IPSec".
- RFC 4478 (April 2006): "Repeated Authentication in Internet Key Exchange (IKEv2) Protocol".
- RFC 4718 (October 2006): "IKEv2 Clarifications and Implementation Guidelines".
- RFC 4835 (April 2007): "Cryptographic Algorithm Implementation RFC Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)".

Chapter 27
PDN Gateway Overview

The Cisco® ASR 5x00 provides wireless carriers with a flexible solution that functions as a Packet Data Network (PDN) Gateway (P-GW) in 3GPP2 Long Term Evolution-System Architecture Evolution (LTE-SAE) and evolved High Rate Packet Data (eHRPD) wireless data networks.

This overview provides general information about the P-GW including:

- Product Description
- Network Deployment(s)
- Features and Functionality - Base Software
- Features and Functionality - Inline Service Support
- Features and Functionality - External Application Support
- Features and Functionality - Optional Enhanced Feature Software
- How the PDN Gateway Works
- Supported Standards

Product Description

The P-GW is the node that terminates the SGi interface towards the PDN. If a UE is accessing multiple PDNs, there may be more than one P-GW for that UE. The P-GW provides connectivity to the UE to external packet data networks by being the point of exit and entry of traffic for the UE. A UE may have simultaneous connectivity with more than one P-GW for accessing multiple PDNs. The P-GW performs policy enforcement, packet filtering for each user, charging support, lawful interception and packet screening.

Figure 206. P-GW in the Basic E-UTRAN/EPC Network

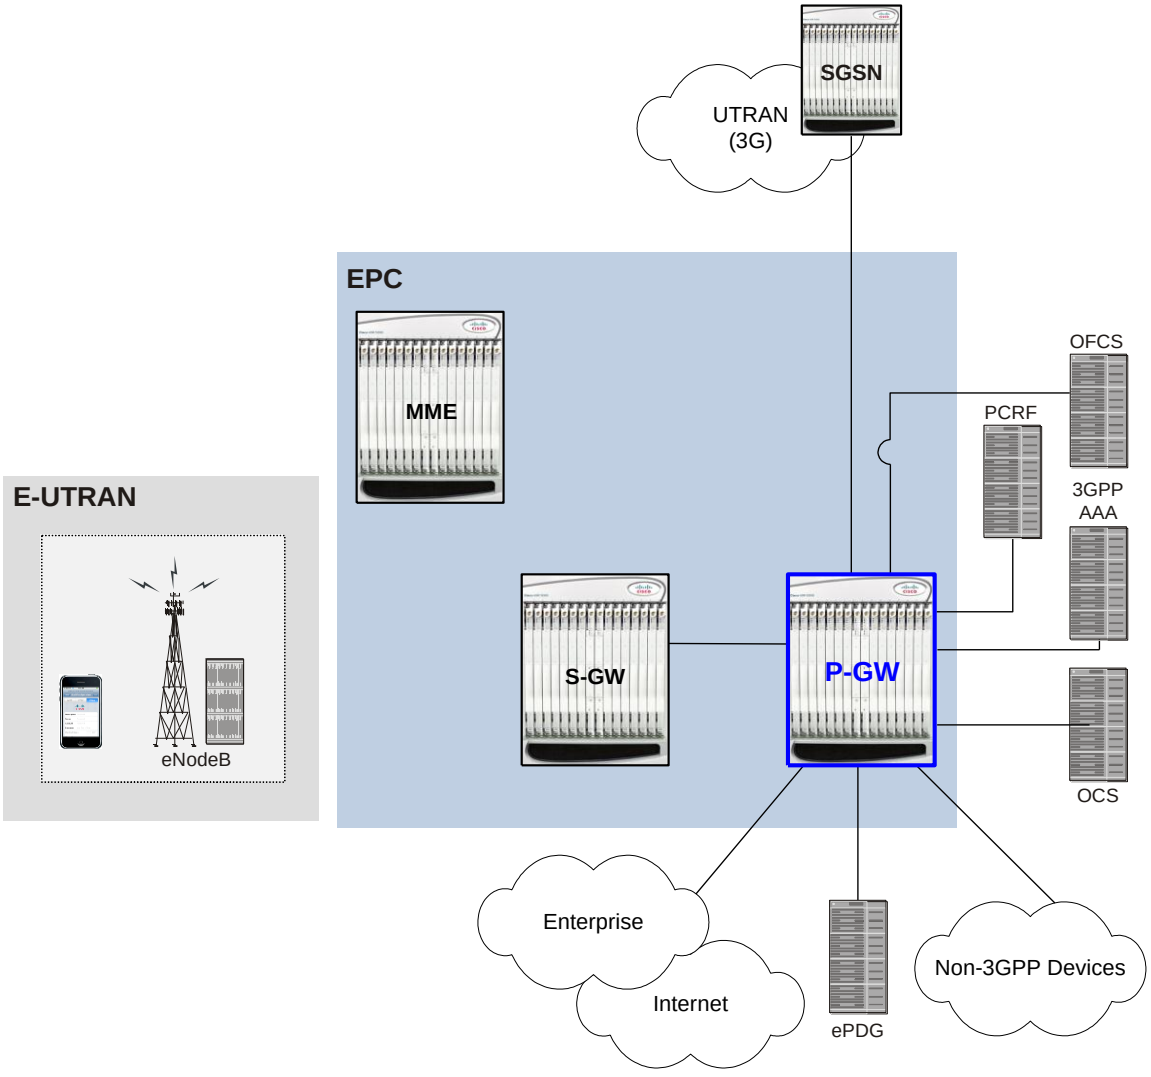
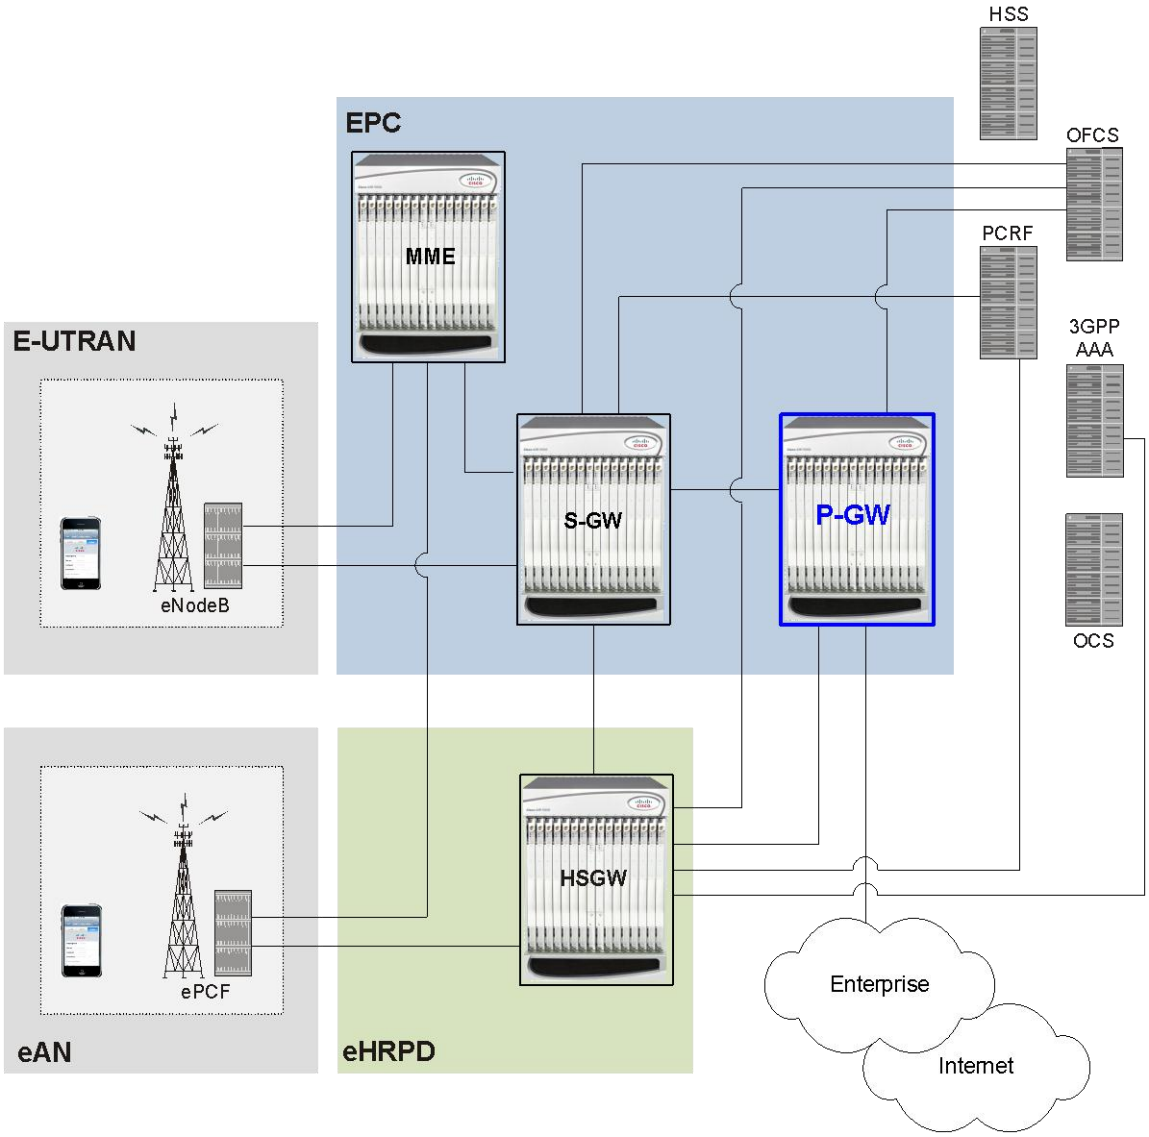


Figure 207. P-GW in the Basic E-UTRAN/EPC and eHRPD Network



Another key role of the P-GW is to act as the anchor for mobility between 3GPP and non-3GPP technologies such as WiMAX and 3GPP2 (CDMA 1X and EvDO).

P-GW functions include:

- Mobility anchor for mobility between 3GPP access systems and non-3GPP access systems. This is sometimes referred to as the SAE Anchor function.
- Policy enforcement (gating and rate enforcement)
- Per-user based packet filtering (deep packet inspection)
- Charging support
- Lawful Interception
- UE IP address allocation

- Packet screening
- Transport level packet marking in the downlink;
- Down link rate enforcement based on Aggregate Maximum Bit Rate (AMBR)

The following are additional P-GW functions when supporting non-3GPP access (eHRPD):

- P-GW includes the function of a Local Mobility Anchor (LMA) according to draft-ietf-netlmm-proxymip6, if PMIP-based S5 or S8 is used.
- The P-GW includes the function of a DSMIPv6 Home Agent, as described in draft-ietf-mip6-nemo-v4traversal, if S2c is used.

Platform Requirements

The P-GW service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the Installation Guide for the chassis and/or contact your Cisco account representative.

Licenses

The P-GW is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

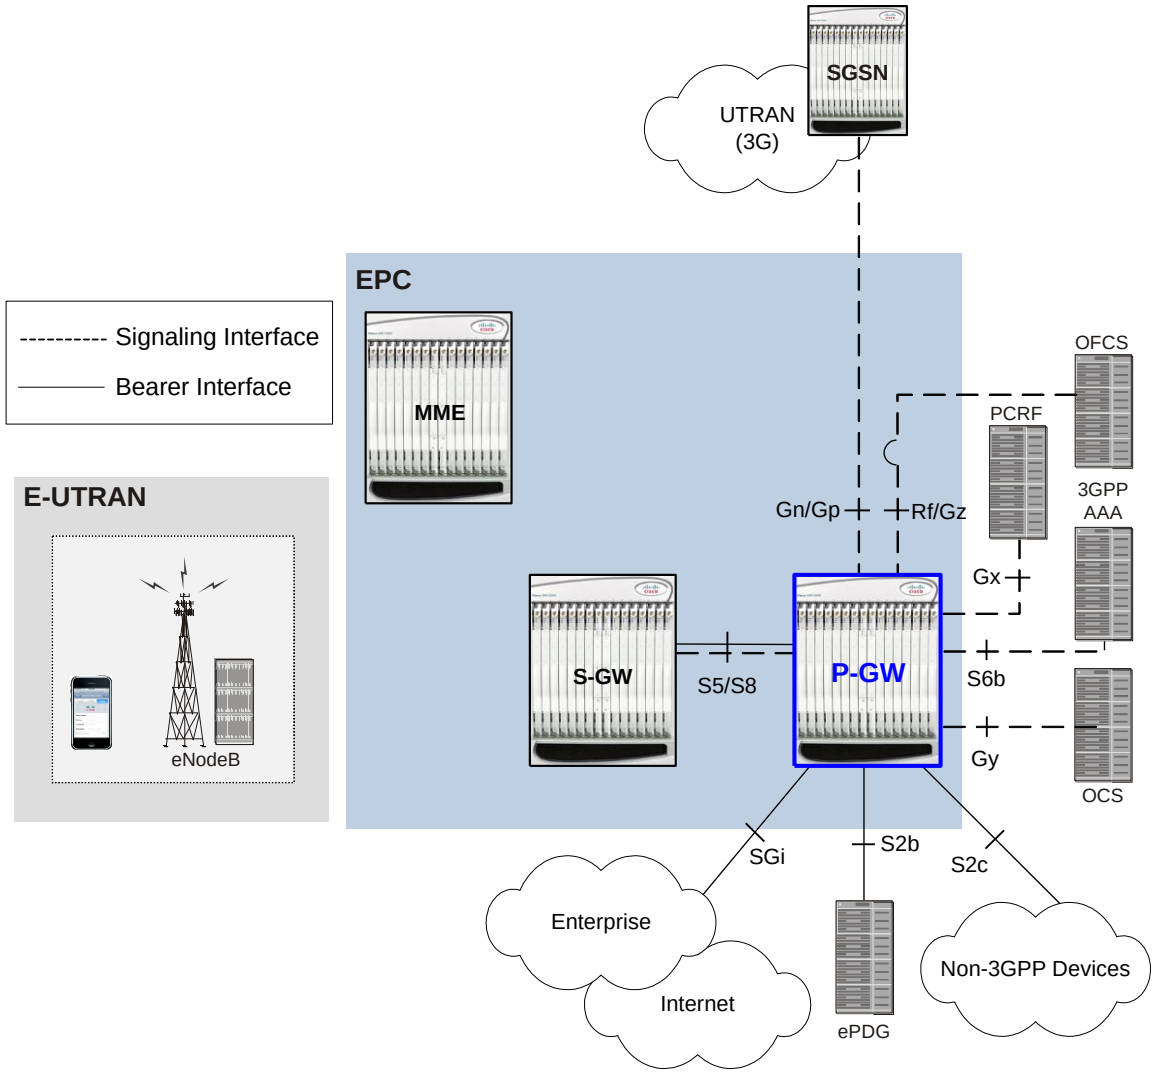
Network Deployment(s)

This section describes the supported interfaces and the deployment scenarios of a PDN Gateway.

PDN Gateway in the E-UTRAN/EPC Network

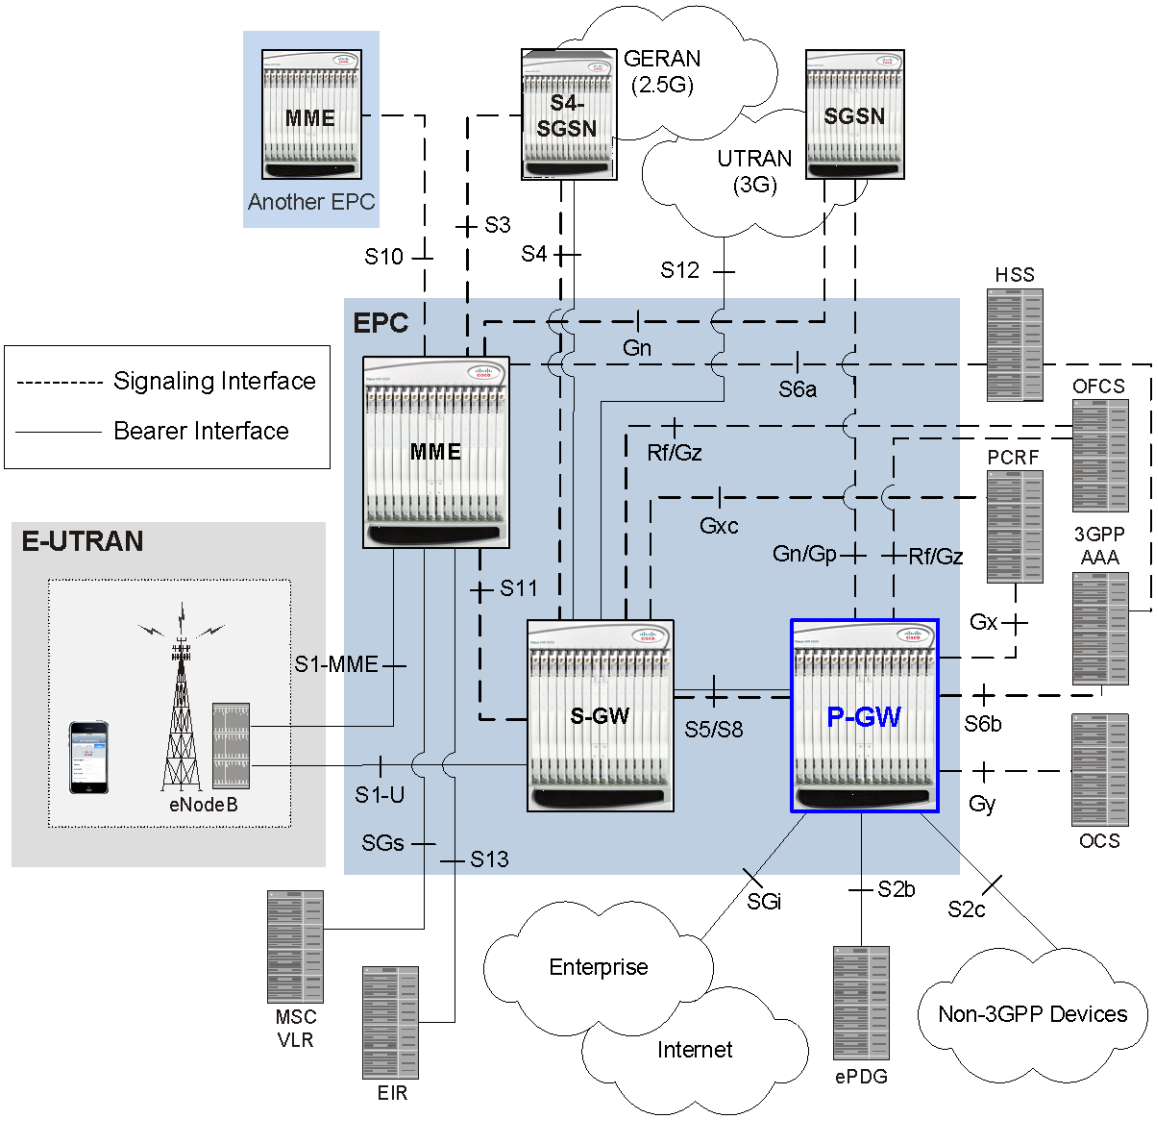
The following figure displays the specific network interfaces supported by the P-GW. Refer to Supported Logical Network Interfaces (Reference Points) for detailed information about each interface.

Figure 208. Supported P-GW Interfaces in the E-UTRAN/EPC Network



The following figure displays a sample network deployment of a P-GW, including all of the interface connections with other 3GPP Evolved-UTRAN/Evolved Packet Core network devices.

Figure 209. P-GW in the E-UTRAN/EPC Network



Supported Logical Network Interfaces (Reference Points)

The P-GW provides the following logical network interfaces in support of E-UTRAN/EPC network:

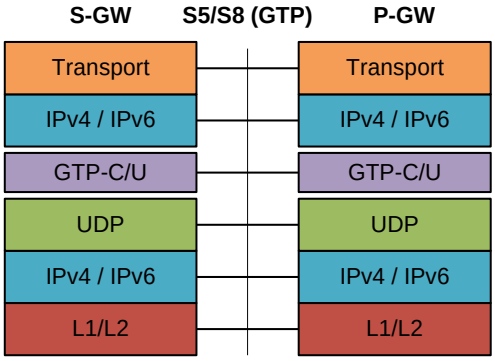
S5/S8 Interface

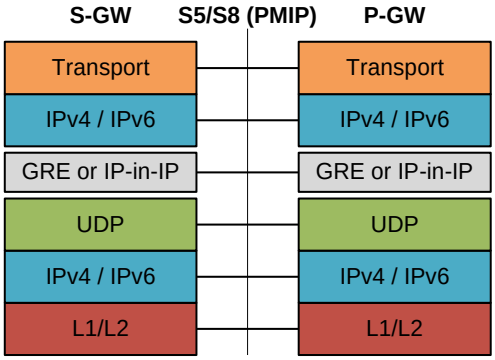
This reference point provides tunneling and management between the S-GW and the P-GW, as defined in 3GPP TS 23.401 and TS 23.402. The S8 interface is an inter-PLMN reference point between the S-GW and the P-GW used during roaming scenarios. The S5 interface is used between an S-GW and P-GW located within the same administrative domain (non-roaming). It is used for S-GW relocation due to UE mobility and if the S-GW needs to connect to a non-collocated P-GW for the required PDN connectivity.

Supported protocols

- Transport Layer: UDP, TCP

- Tunneling:
  - GTP: GTPv2-C (signaling channel), GTPv1-U (bearer channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



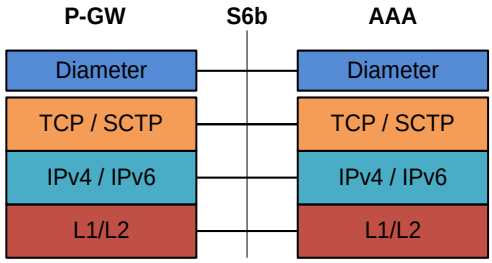


S6b Interface

This reference point, between a P-GW and a 3GPP AAA server/proxy, is used for mobility-related authentication. It may also be used to retrieve and request parameters related to mobility and to retrieve static QoS profiles for UEs (for non-3GPP access) in the event that dynamic PCC is not supported.

Supported protocols:

- Transport Layer: TCP, SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

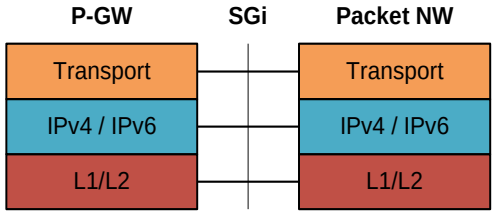


SGi Interface

This reference point provides connectivity between the P-GW and a packet data network (3GPP TS 23.401). This interface can provide access to a variety of network types including an external public or private PDN and/or an internal IMS service provisioning network.

Supported protocols:

- Transport Layer: TCP, UDP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

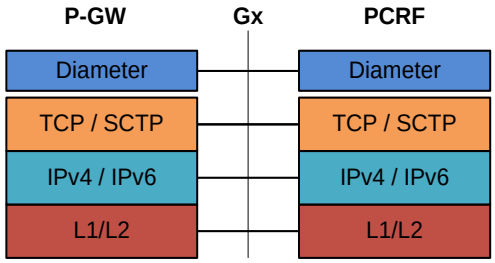


Gx Interface

This signalling interface supports the transfer of policy control and charging rules information (QoS) between the Policy and Charging Enforcement Function (PCEF) on the P-GW and a Policy and Charging Rules Function (PCRF) server (3GPP TS 23.401).

Supported protocols:

- Transport Layer: TCP, SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



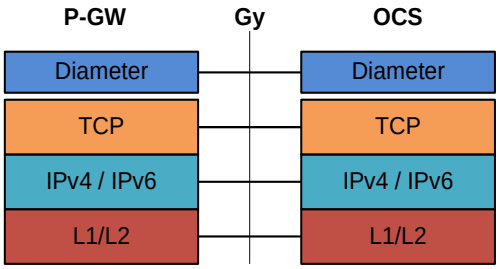
For more information on the Gx interface, refer to Dynamic Policy Charging Control (Gx Reference Interface) in the *Features and Functionality - Base Software* section of this chapter.

Gy Interface

The Gy reference interface enables online accounting functions on the P-GW in accordance with 3GPP Release 8 specifications.

Supported protocols:

- Transport Layer: TCP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



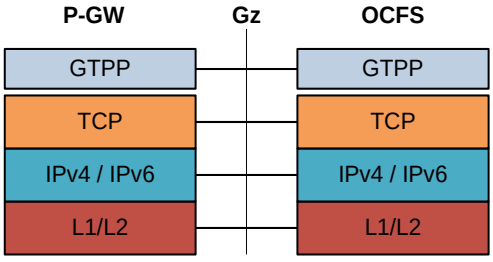
For more information on the Gy interface and online accounting, refer to Gy Interface Support in the *Features and Functionality - Base Software* section of this chapter.

Gz Interface

The Gz reference interface enables offline accounting functions on the P-GW. The P-GW collects charging information for each mobile subscriber UE pertaining to the radio network usage.

Supported protocols:

- Transport Layer: TCP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

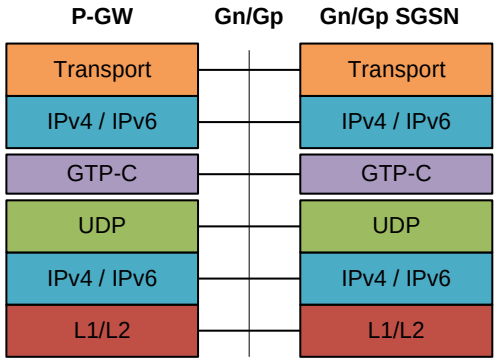


Gn/Gp Interface

This reference point provides tunneling and management between the P-GW and the SGSN during handovers between the EPS and 3GPP 2G and/or 3G networks (3GPP TS 29.060). For more information on the Gn/Gp interface, refer to Gn/Gp Handoff Support in the *Features and Functionality - Base Software* section of this chapter.

Supported protocols

- Transport Layer: UDP, TCP
- Tunneling: GTP: GTP-C (signaling channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

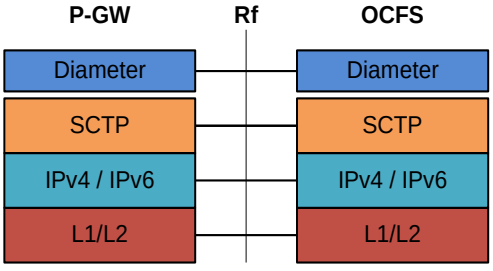


Rf Interface

The Rf interface enables offline accounting functions on the P-GW in accordance with 3GPP Release 8 specifications. The P-GW collects charging information for each mobile subscriber UE pertaining to the radio network usage.

Supported protocols:

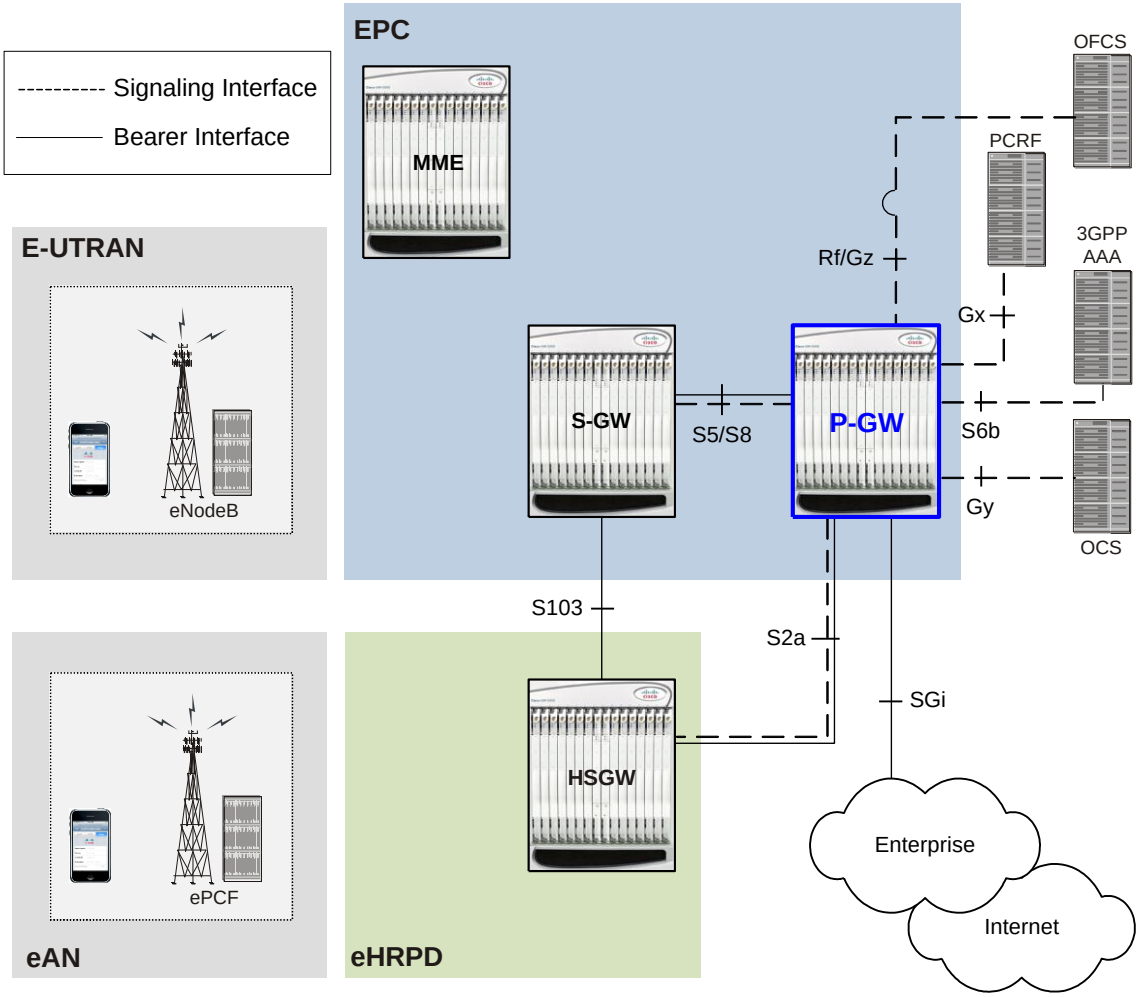
- Transport Layer: TCP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



PDN Gateway Supporting eHRPD to E-UTRAN/EPC Connectivity

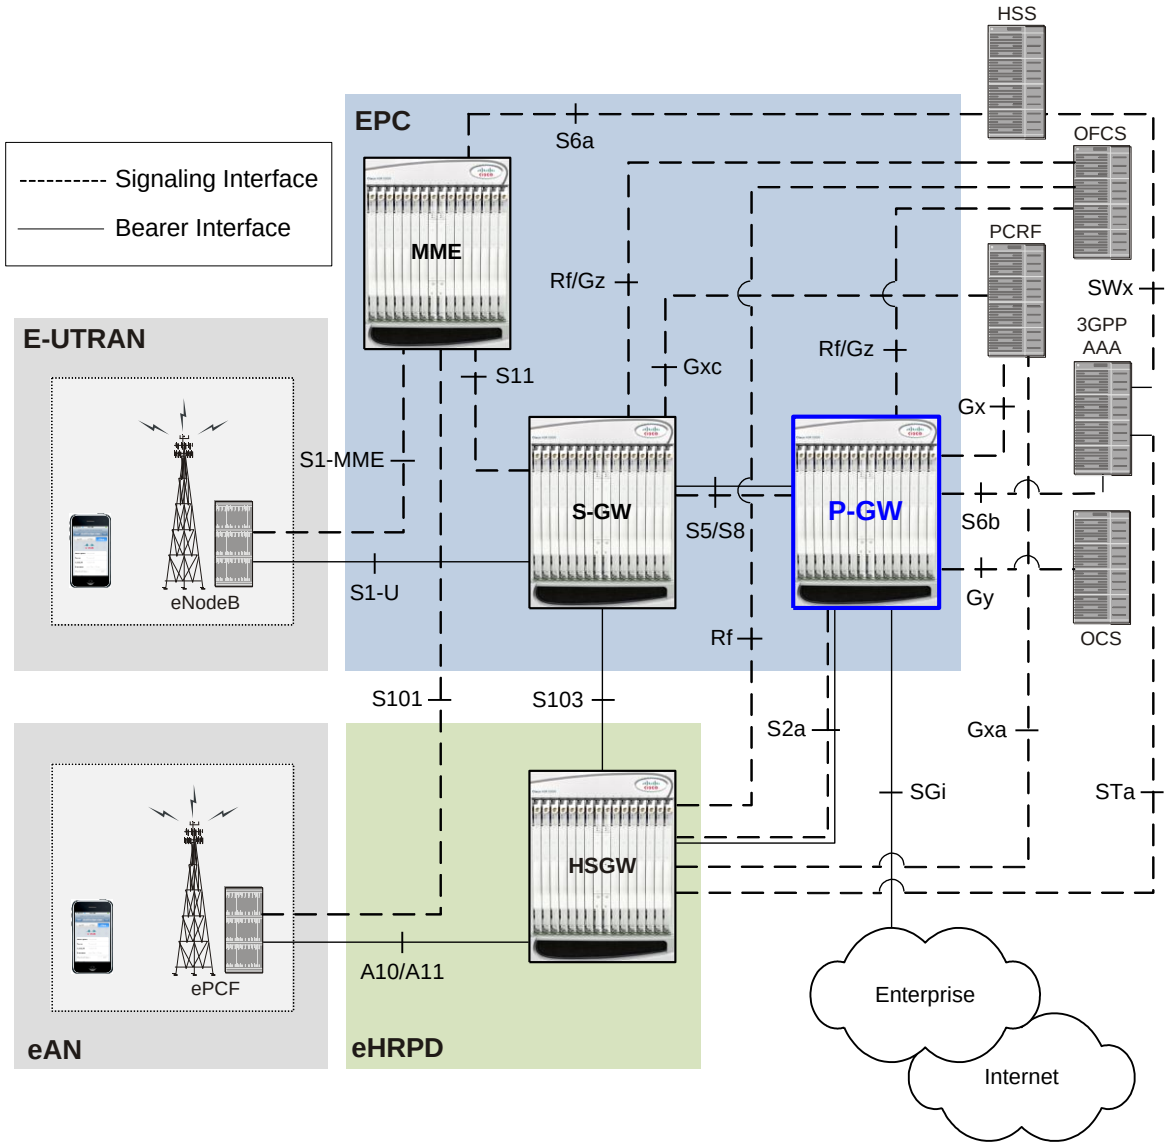
The following figure displays the specific network interfaces supported by the P-GW in an eHRPD network. Refer to Supported Logical Network Interfaces (Reference Points) for detailed information about each interface.

Figure 210. P-GW Interfaces Supporting eHRPD to E-UTRAN/EPC Connectivity



The following figure displays a sample network deployment of a P-GW in an eHRPD Network, including all of the interface connections with other 3GPP Evolved-UTRAN/Evolved Packet Core network devices.

Figure 211. P-GW in the E-UTRAN/EPC Network Supporting the eHRPD Network



Supported Logical Network Interfaces (Reference Points)

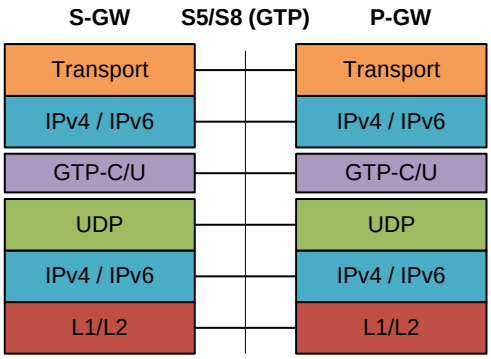
The P-GW provides the following logical network interfaces in support of eHRPD to E-UTRAN/EPC connectivity:

S5/S8 Interface

This reference point provides tunneling and management between the S-GW and the P-GW, as defined in 3GPP TS 23.401. The S8 interface is an inter-PLMN reference point between the S-GW and the P-GW used during roaming scenarios. The S5 interface is used between an S-GW and P-GW located within the same administrative domain (non-roaming). It is used for S-GW relocation due to UE mobility and if the S-GW needs to connect to a non-collocated P-GW for the required PDN connectivity.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling:
    - GTP: IPv4 or IPv6 GTP-C (signaling channel) and GTP-U (bearer channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

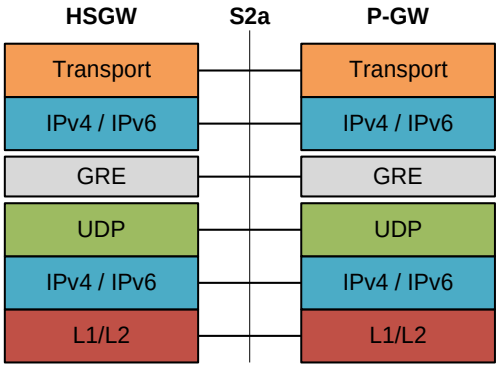


S2a Interface

This reference point supports the bearer interface by providing signaling and mobility support between a trusted non-3GPP access point (HSGW) and the PDN Gateway. It is based on Proxy Mobile IP but also supports Client Mobile IPv4 FA mode which allows connectivity to trusted non-3GPP IP access points that do not support PMIP.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: GRE IPv6
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

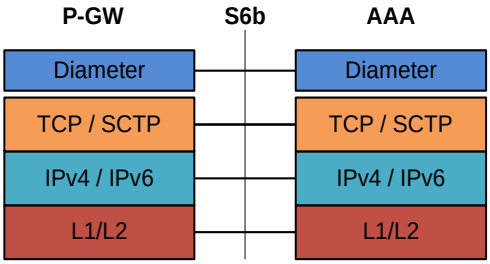


S6b Interface

This reference point, between a P-GW and a 3GPP AAA server/proxy, is used for mobility-related authentication. It may also be used to retrieve and request parameters related to mobility and to retrieve static QoS profiles for UEs (for non-3GPP access) in the event that dynamic PCC is not supported.

Supported protocols:

- Transport Layer: TCP, SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

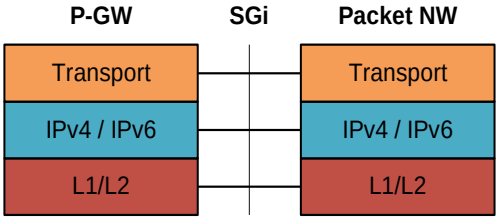


SGi Interface

This reference point provides connectivity between the P-GW and a packet data network. This interface can provide access to a variety of network types including an external public or private PDN and/or an internal IMS service provisioning network.

Supported protocols:

- Transport Layer: TCP, UDP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

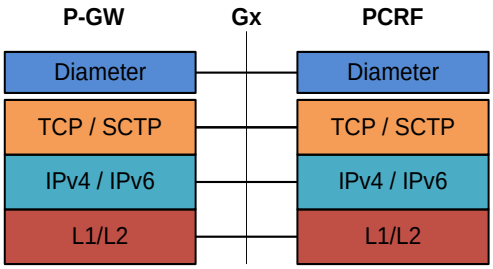


Gx Interface

This signalling interface supports the transfer of policy control and charging rules information (QoS) between the Policy and Charging Enforcement Function (PCEF) on the P-GW and a Policy and Charging Rules Function (PCRF) server (3GPP TS 23.401).

Supported protocols:

- Transport Layer: TCP, SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



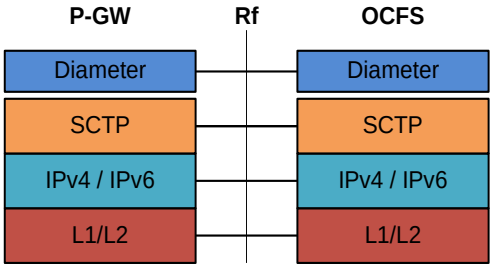
For more information on the Gx interface, refer to Dynamic Policy Charging Control (Gx Reference Interface) in the *Features and Functionality - Base Software* section of this chapter.

Rf Interface

The Rf reference interface enables offline accounting functions on the P-GW in accordance with 3GPP Release 8 specifications. The P-GW collects charging information for each mobile subscriber UE pertaining to the radio network usage.

Supported protocols:

- Transport Layer: TCP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



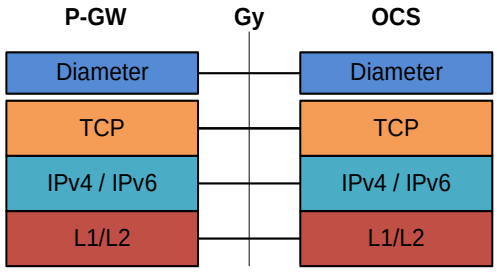
For more information on Rf accounting, refer to the section in the *Features and Functionality - Base Software* section of this chapter.

Gy Interface

The Gy reference interface enables online accounting functions on the P-GW in accordance with 3GPP Release 8 specifications.

Supported protocols:

- Transport Layer: TCP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



For more information on the Gy interface and online accounting, refer to Gy Interface Support in the *Features and Functionality - Base Software* section of this chapter.

Features and Functionality - Base Software

This section describes the features and functions supported by default in the base software for the P-GW service and do not require any additional licenses to implement the functionality.

This section describes the following features:

- 3GPP R9 Volume Charging Over Gx
- AAA Server Groups
- ANSI T1.276 Compliance
- APN Support
- Bulk Statistics Support
- Congestion Control
- Default and Dedicated EPC Bearers
- DHCP Support
- Direct Tunnel Support
- DSCP Marking
- Dynamic Policy Charging Control (Gx Reference Interface)
- Enhanced Charging Service (ECS)
- GnGp Handoff Support
- IP Access Control Lists
- IP Address Hold Timers
- IPv6 Capabilities
- Local Break-Out
- Management System Overview
- Mobile IP Registration Revocation
- Non-Optimized e-HRPD to Native LTE (E-UTRAN) Mobility Handover
- Multiple PDN Support
- Online/Offline Charging
- Proxy Mobile IPv6 (S2a)
- QoS Bearer Management
- RADIUS Support
- Source IP Address Validation
- Subscriber Level Trace
- Threshold Crossing Alerts (TCA) Support
- Virtual APN Support

3GPP R9 Volume Charging Over Gx

Also known as accumulated usage tracking over Gx, this 3GPP R9 enhancement provides a subset of the volume and charging control functions defined in TS 29.212 based on usage quotas between a P-GW and PCRF. The quotas can be assigned to the default bearer or any of the dedicated bearers for the PDN connection.

This feature enables volume reporting over Gx, which entails usage monitoring and reporting of the accumulated usage of network resources on an IP-CAN session or service data flow basis. PCRF subscribes to the usage monitoring at session level or at flow level by providing the necessary information to PCEF. PCEF in turn reports the usage to the PCRF when the conditions are met. Based on the total network usage in real-time, the PCRF will have the information to enforce dynamic policy decisions.

When usage monitoring is enabled, the PCEF can monitor the usage volume for the IP-CAN session, or applicable service data flows, and report accumulated usage to the PCRF based on any of the following conditions:

- When a usage threshold is reached,
- When all PCC rules for which usage monitoring is enabled for a particular usage monitoring key are removed or deactivated,
- When usage monitoring is explicitly disabled by the PCRF,
- When an IP CAN session is terminated or,
- When requested by the PCRF.

Accumulated volume reporting can be measured by total volume, the uplink volume, or the downlink volume as requested by the PCRF. When receiving the reported usage from the PCEF, the PCRF deducts the value of the usage report from the total allowed usage for that IP-CAN session, usage monitoring key, or both as applicable.

AAA Server Groups

Value-added feature to enable VPN service provisioning for enterprise or MVNO customers. Enables each corporate customer to maintain its own AAA servers with its own unique configurable parameters and custom dictionaries.

This feature provides support for up to 800 AAA server groups and 800 NAS IP addresses that can be provisioned within a single context or across the entire chassis. A total of 128 servers can be assigned to an individual server group. Up to 1,600 accounting, authentication and/or mediation servers are supported per chassis.

ANSI T1.276 Compliance

ANSI T1.276 specifies security measures for Network Elements (NE). In particular it specifies guidelines for password strength, storage, and maintenance security measures.

ANSI T1.276 specifies several measures for password security. These measures include:

- Password strength guidelines
- Password storage guidelines for network elements
- Password maintenance, e.g. periodic forced password changes

These measures are applicable to the ASR 5x00 and the Web Element Manager since both require password authentication. A subset of these guidelines where applicable to each platform will be implemented. A known subset of guidelines, such as certificate authentication, are not applicable to either product. Furthermore, the platforms support a variety of authentication methods such as RADIUS and SSH which are dependent on external elements. ANSI T1.276 compliance in such cases will be the domain of the external element. ANSI T1.276 guidelines will only be implemented for locally configured operators.

APN Support

The P-GW's Access Point Name (APN) support offers several benefits:

- Extensive parameter configuration flexibility for the APN.
- Creation of subscriber tiers for individual subscribers or sets of subscribers within the APN.
- Virtual APNs to allow differentiated services within a single APN.

In StarOS v12.x and earlier, up to 1024 APNs can be configured in the P-GW. An APN may be configured for any type of PDP context, i.e., PPP, IPv4, IPv6 or both IPv4 and IPv6. Many dozens of parameters may be configured independently for each APN.

Here are a few highlights of what may be configured:

- **Accounting**: RADIUS, GTPP or none. Server group to use. Charging characteristics. Interface with mediation servers.
- **Authentication**: Protocol, such as, CHAP or PAP or none. Default username/password. Server group to use. Limit for number of PDP contexts.
- **Enhanced Charging**: Name of rulebase to use, which holds the enhanced charging configuration (e.g., eG-CDR variations, charging rules, prepaid/postpaid options, etc.).
- **IP**: Method for IP address allocation (e.g., local allocation by P-GW, Mobile IP, DHCP, etc.). IP address ranges, with or without overlapping ranges across APNs.
- **Tunneling**: PPP may be tunneled with L2TP. IPv4 may be tunneled with GRE, IP-in-IP or L2TP. Load-balancing across multiple tunnels. IPv6 is tunneled in IPv4. Additional tunneling techniques, such as, IPsec and VLAN tagging may be selected by the APN, but are configured in the P-GW independently from the APN.
- **QoS**: IPv4 header ToS handling. Traffic rate limits for different 3GPP traffic classes. Mapping of R98 QoS attributes to work around particular handset defections. Dynamic QoS renegotiation (described elsewhere).

After an APN is determined by the P-GW, the subscriber may be authenticated/authorized with an AAA server. The P-GW allows the AAA server to return VSAs (Vendor Specific Attributes) that override any/all of the APN configuration. This allows different subscriber tier profiles to be configured in the AAA server, and passed to the P-GW during subscriber authentication/authorization.

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Following is a list of supported schemas for P-GW:

- **APN**: Provides Access Point Name statistics
- **Card**: Provides card-level statistics
- **Context**: Provides context service statistics
- **Diameter-acct**: Provides Diameter Accounting statistics
- **Diameter-auth**: Provides Diameter Authentication statistics
- **ECS**: Provides Enhanced Charging Service statistics
- **EGTPC**: Provides Evolved GPRS Tunneling Protocol - Control message statistics
- **FA**: Provides FA service statistics
- **GTPC**: Provides GPRS Tunneling Protocol - Control message statistics
- **GTPP**: Provides GPRS Tunneling Protocol - Prime message statistics
- **GTPU**: Provides GPRS Tunneling Protocol - User message statistics
- **HA**: Provides HA service statistics
- **IMSA**: Provides IMS Authorization service statistics
- **IP Pool**: Provides IP pool statistics
- **LMA**: Provides Local Mobility Anchor service statistics
- **P-GW**: Provides P-GW node-level service statistics
- **Port**: Provides port-level statistics
- **PPP**: Provides Point-to-Point Protocol statistics
- **RADIUS**: Provides per-RADIUS server statistics
- **System**: Provides system-level statistics

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, system host name, system uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Important: For more information on bulk statistic configuration, refer to the *Configuring and Maintaining Bulk Statistics* chapter in the *System Administration Guide*.

Congestion Control

The congestion control feature allows you to set policies and thresholds and specify how the system reacts when faced with a heavy load condition.

Congestion control monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are quickly resolved. However, continuous or large numbers of these conditions within a specific time interval may have an impact the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establishes limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operation thresholds that are configured for the system as described in the Thresholding Configuration Guide. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated.

  A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.

  - **Port Utilization Thresholds**: If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
    - **Port-specific Thresholds**: If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
  - **Service Congestion Policies**: Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

Important: For more information on congestion control, refer to the *Congestion Control* chapter in the *System Administration Guide*.

Default and Dedicated EPC Bearers

Provides a foundation for contributing towards improved Quality of User Experience (QoE) by enabling deterministic end-to-end forwarding and scheduling treatments for different services or classes of applications pursuant to their requirements for committed bandwidth resources, jitter and delay. In this way, each application receives the service treatment that users expect.

In the StarOS 9.0 release and later, the Cisco EPC core platforms support one or more EPS bearers (default plus dedicated). An EPS bearer is a logical aggregate of one or more Service Data Flows (SDFs), running between a UE and a P-GW in the case of a GTP-based S5/S8 interface, and between a UE and HSGW (HRPD Serving Gateway) in case of a PMIP-based S2a interface. In networks where GTP is used as the S5/S8 protocol, the EPS bearer constitutes a concatenation of a radio bearer, S1-U bearer and an S5/S8 bearer anchored on the P-GW. In cases where PMIPv6 is used the EPS bearer is concatenated between the UE and HSGW with IP connectivity between the HSGW and P-GW.

Note: The P-GW supports GTP-based S5/S8 and PMIPv6 S2a capabilities, with no commercial support for PMIPv6 S5/S8.

An EPS bearer uniquely identifies traffic flows that receive a common QoS treatment between a UE and P-GW in the GTP-based S5/S8 design, and between a UE and HSGW in the PMIPv6 S2a approach. If different QoS scheduling

priorities are required between Service Data Flows, they should be assigned to separate EPS bearers. Packet filters are signalled in the NAS procedures and associated with a unique packet filter identifier on a per-PDN connection basis.

One EPS bearer is established when the UE connects to a PDN, and that remains established throughout the lifetime of the PDN connection to provide the UE with always-on IP connectivity to that PDN. That bearer is referred to as the default bearer. A PDN connection represents a traffic flow aggregate between a mobile access terminal and an external Packet Data Network (PDN) such as an IMS network, a walled garden application cloud or a back-end enterprise network. Any additional EPS bearer that is established to the same PDN is referred to as a dedicated bearer. The EPS bearer Traffic Flow Template (TFT) is the set of all 5-tuple packet filters associated with a given EPS bearer. The EPC core elements assign a separate bearer ID for each established EPS bearer. At a given time a UE may have multiple PDN connections on one or more P-GWs.

DHCP Support

The P-GW supports dynamic IP address assignment to subscriber IP PDN contexts using the Dynamic Host Control Protocol (DHCP), as defined by the following standards:

- RFC 2131, Dynamic Host Configuration Protocol
- RFC 2132, DHCP Options and BOOTP Vendor Extensions

The method by which IP addresses are assigned to a PDN context is configured on an APN-by-APN basis. Each APN template dictates whether it will support static or dynamic addresses. Dynamically assigned IP addresses for subscriber PDN contexts can be assigned through the use of DHCP.

The P-GW acts as a DHCP server toward the UE and a DHCP client toward the external DHCP server. The DHCP server function and DHCP client function on the P-GW are completely independent of each other; one can exist without the other.

The P-GW does not support DHCP-relay.

Important: Currently, the P-GW only supports DHCP with IPv4 addresses. IPv6 address support is planned at a later date.

Deferred IPv4 Address Allocation

Apart from obtaining IP addresses during initial access signalling, a UE can indicate via PCO options that it prefers to obtain IP address and related configuration via DHCP after default bearer has been established. This is also know as Deferred Address Allocation.

IPv4 addresses are becoming an increasingly scarce resource. Since 4G networks like LTE are always on, scarce resources such as IPv4 addresses cannot/should not be monopolized by UEs when they are in an ECM-IDLE state.

PDN-type IPv4v6 allows a dual stack implementing. The P-GW allocates an IPv6 address only by default for an IPv4v6 PDN type. The UE defers the allocation of IPv4 addresses based upon its needs, and relinquishes any IPv4 addresses to the global pool once it is done. The P-GW may employ any IPv4 address scheme (local pool or external DHCP server) when providing an IPv4 address on demand.

Direct Tunnel Support

When Gn/Gp interworking with pre-release SGSNs is enabled, the GGSN service on the P-GW supports direct tunnel functionality.

Direct tunnel improves the user experience (e.g. expedited web page delivery, reduced round trip delay for conversational services, etc.) by eliminating SGSN tunnel "switching" latency from the user plane. An additional advantage of direct tunnel from an operational and capital expenditure perspective is that direct tunnel optimizes the usage of user plane resources by removing the requirement for user plane processing on the SGSN.

The direct tunnel architecture allows the establishment of a direct user plane tunnel between the RAN and the GGSN, bypassing the SGSN. The SGSN continues to handle the control plane signalling and typically makes the decision to establish direct tunnel at PDP Context Activation. A direct tunnel is achieved at PDP context activation by the SGSN establishing a user plane (GTP-U) tunnel directly between RNC and GGSN (using an Update PDP Context Request toward the GGSN).

A major consequence of deploying direct tunnel is that it produces a significant increase in control plane load on both the SGSN and GGSN components of the packet core. It is therefore of paramount importance to a wireless operator to ensure that the deployed GGSNs are capable of handling the additional control plane loads introduced of part of direct tunnel deployment. The Cisco GGSN and SGSN offer massive control plane transaction capabilities, ensuring system control plane capacity will not be a capacity limiting factor once direct tunnel is deployed.

DSCP Marking

Provides support for more granular configuration of DSCP marking.

For Interactive Traffic class, the P-GW supports per-gateway service and per-APN configurable DSCP marking for Uplink and Downlink direction based on Allocation/Retention Priority in addition to the current priorities.

The following matrix may be used to determine the Diffserv markings used based on the configured traffic class and Allocation/Retention Priority:

Table 95. Default DSCP Value Matrix

Allocation Priority	1	2	3
Traffic Handling Priority			
1	ef	ef	ef
2	af21	af21	af21
3	af21	af21	af21

Dynamic Policy Charging Control (Gx Reference Interface)

Dynamic policy and charging control provides a primary building block toward the realization of IMS multimedia applications. In contrast to statically provisioned architectures, the dynamic policy framework provides a centralized service control layer with global awareness of all access-side network elements. The centralized policy decision elements simplify the process of provisioning global policies to multiple access gateways. Dynamic policy is especially useful in an Always-On deployment model as the usage paradigm transitions from a short lived to a lengthier online session in which the volume of data consumed can be extensive. Under these conditions dynamic policy management enables dynamic just in-time resource allocation to more efficiently protect the capacity and resources of the network.

Dynamic Policy Control represents the ability to dynamically authorize and control services and application flows between a Policy Charging Enforcement Function (PCEF) on the P-GW and the PCRF. Policy control enables a

centralized and decoupled service control architecture to regulate the way in which services are provisioned and allocated at the bearer resource layer.

The StarOS 9.0 release included enhancements to conform with 3GPP TS 29.212 and 29.230 functions. The Gx reference interface uses Diameter transport and IPv6 addressing. The subscriber is identified to the PCRF at session establishment using IMSI based NAIs within the Subscription-ID AVP. Additionally the IMEI within the Equipment-Info AVP is used to identify the subscriber access terminal to the policy server. The Gx reference interface supports the following capabilities:

- Authorize the bearer establishment for a packet flow
- Dynamic L3/L4 transfer of service data flow filters within PCC rules for selection and policy enforcement of downlink/uplink IP CAN bearers
- Support static pre-provisioned L7 rulebase name attribute as trigger for activating Inline Services such as Peer-to-Peer Detection
- Authorize the modification of a service data flow
- Revoke the authorization of a packet flow
- Provision PCC rules for service data flows mapped to default or dedicated EPS bearers
- Support P-GW initiated event triggers based on change of access network gateway or IP CAN
- Provide the ability to set or modify APN-AMBR for a default EPS bearer
- Create or modify QoS service priority by including QCI values in PCC rules transmitted from PCRF to PCEF functions

Enhanced Charging Service (ECS)

The Enhanced Charging Service provides an integrated in-line service for inspecting subscriber data packets and generating detail records to enable billing based on usage and traffic patterns. Other features include:

- Content Analysis Support
- Content Service Steering
- Support for Multiple Detail Record Types
- Diameter Credit Control Application
- Accept TCP Connections from DCCA Server
- Gy Interface Support

The Enhanced Charging Service (ECS) is an in-line service feature that is integrated within the system. ECS enhances the mobile carrier's ability to provide flexible, differentiated, and detailed billing to subscribers by using Layer 3 through Layer 7 deep packet inspection with the ability to integrate with back-end billing mediation systems.

ECS interacts with active mediation systems to provide full real-time prepaid and active charging capabilities. Here the active mediation system provides the rating and charging function for different applications.

In addition, ECS also includes extensive record generation capabilities for post-paid charging with in-depth understanding of the user session. Refer to the Support for Multiple Detail Record Types section for more information.

The major components include:

- **Service Steering**: Directs subscriber traffic into the ECS subsystem. Service Steering is used to direct selective subscriber traffic flows via an Access Control List (ACL). It is used for other redirection applications as well for both internal and external services and servers.

- **Protocol Analyzer**: The software stack responsible for analyzing the individual protocol fields and states during packet inspection. It performs two types of packet inspection:
  - **Shallow Packet Inspection**: inspection of the layer 3 (IP header) and layer 4 (e.g. UDP or TCP header) information.
  - **Deep Packet Inspection**: inspection of layer 7 and 7+ information. Deep packet inspection functionality includes:
    - Detection of URI (Uniform Resource Identifier) information at level 7 (e.g., HTTP, WTP, RTSP Uniform Resource Locators (URLs)).
    - Identification of true destination in the case of terminating proxies, where shallow packet inspection would only reveal the destination IP address / port number of a terminating proxy.
    - De-encapsulation of upper layer protocol headers, such as MMS-over-WTP, WSP-over-UDP, and IP-over GPRS.
    - Verification that traffic actually conforms to the protocol the layer 4 port number suggests.
- **Rule Definitions**: User-defined expressions, based on protocol fields and/or protocol-states, which define what actions to take when specific field values are true. Expressions may contain a number of operator types (string, =, >, etc.) based on the data type of the operand. Each Ruledef configuration is consisting of multiple expressions applicable to any of the fields or states supported by the respective analyzers.
- **Rule Bases**: a collection of rule definitions and their associated billing policy. The rule base determines the action to be taken when a rule is matched. It is possible to define a rule definition with different actions.

Mediation and Charging Methods

To provide maximum flexibility when integrating with billing mediation systems, ECS supports a full range of charging and authorization interfaces.

- **Pre-paid**: In a pre-paid environment, the subscribers pay for service prior to use. While the subscriber is using the service, credit is deducted from subscriber's account until it is exhausted or call ends. The pre-paid accounting server is responsible for authorizing network nodes (GGSNs) to grant access to the user, as well as grant quotas for either time connected or volume used. It is up to the network node to track the quota use, and when these use quotas run low, the network node sends a request to the pre-paid server for more quota.

  If the user has not used up the purchased credit, the server grants quota and if no credit is available to the subscriber the call will be disconnected. ECS and DCCA manage this functionality by providing the ability to setup quotas for different services.

  Pre-paid quota in ECS is implemented using DIAMETER Credit Control Application (DCCA). DCCA supports the implementation of real-time credit control for a variety of services, such as networks access, messaging services, and download services.

  In addition to being a general solution for real-time cost and credit control, DCCA includes these features:

  - **Real-time Rate Service Information** - DCCA can verify when end subscribers' accounts are exhausted or expired; or deny additional chargeable events.
  - **Support for Multiple Services** - DCCA supports the usage of multiple services within one subscriber session. Multiple Service support includes; 1) ability to identify and process the service or group of services that are subject to different cost structures 2) independent credit control of multiple services in a single credit control sub-session.

    Refer to the Diameter Credit Control Application section for more information.
- **Post-paid**: In a post-paid environment, the subscribers pay after use of the service. A AAA server is responsible for authorizing network nodes (GGSNs) to grant access to the user and a CDR system generates G-CDRs/eG-

CDRs/EDRs/UDRs or Comma Separated Values (CSVs) for billing information on pre-defined intervals of volume or per time.

Important: Support for the Enhanced Charging Service requires a service license; the ECS license is included in the P-GW session use license. For more information on ECS, refer to the *Enhanced Charging Service Administration Guide*.

Content Analysis Support

The Enhanced Charging Service is capable of performing content analysis on packets of many different protocols at different layers of the OSI model.

The ECS content analyzers are able to inspect and maintain state across various protocols at all layers of the OSI stack. ECS system supports, inspects, and analyzes the following protocols:

- IP
- TCP
- UDP
- DNS
- FTP
- TFTP
- SMTP
- POP3
- HTTP
- ICMP
- WAP: WTP and WSP
- Real-Time Streaming: RTP and RTSP
- MMS
- SIP and SDP
- File analysis: examination of downloaded file characteristics (e.g. file size, chunks transferred, etc.) from file transfer protocols such as HTTP and FTP.

Traffic analyzers in enhanced charging subsystem are based on configured rules. Rules used for Traffic analysis analyze packet flows and form usage records. Usage records are created per content type and forwarded to a pre-paid server or to a mediation/billing system. A traffic analyzer performs shallow (Layer 3 and Layer 4) and deep (above Layer 4) packet inspection of the IP packet flows.

The Traffic Analyzer function is able to do a shallow (layer 3 and layer 4) and deep (above layer 4) packet inspection of IP Packet Flows.

It is able to correlate all layer 3 packets (and bytes) with higher layer trigger criteria (e.g. URL detected in a HTTP header) and it is also perform stateful packet inspection to complex protocols like FTP, RTSP, SIP that dynamically open ports for the data path and by this way, user plane payload is differentiated into “categories”.

The Traffic Analyzer works on the application level as well and performs event based charging without the interference of the service platforms.

Important: This functionality is available for use with the Enhanced Charging Service which requires a session-use license. For more information on ECS, refer to the *Enhanced Charging Service Administration Guide*.

Content Service Steering

Content Service Steering (CSS) directs selective subscriber traffic into the ECS subsystem (In-line services internal to the system) based on the content of the data presented by mobile subscribers.

CSS uses Access Control Lists (ACLs) to redirect selective subscriber traffic flows. ACLs control the flow of packets into and out of the system. ACLs consist of "rules" (ACL rules) or filters that control the action taken on packets matching the filter criteria.

ACLs are configurable on a per-context basis and applies to a subscriber through either a subscriber profile or an APN profile in the destination context.

Important: For more information on CSS, refer to the *Content Service Steering* chapter of the *System Administration Guide*.

Important: For more information on ACLs, refer to the *IP Access Control Lists* chapter of the *System Administration Guide*.

Support for Multiple Detail Record Types

To meet the requirements of standard solutions and at the same time, provide flexible and detailed information on service usage, the Enhanced Charging Service (ECS) provides the following type of usage records:

- Event Detail Records (EDRs)
- Usage Detail Records (UDRs)

ECS provides for the generation of charging data files, which can be periodically retrieved from the system and used as input to a billing mediation system for post-processing. These files are provided in a standard format, so that the impact on the existing billing/mediation system is minimal and at the same time, these records contain all the information required for billing based on the content.

GTPP accounting in ECS allows the collection of counters for different types of data traffic into detail records. The following types of detail records are supported:

- **Event Detail Records (EDRs)**: An alternative to standard G-CDRs when the information provided by the G-CDRs is not sufficient to do the content billing. EDRs are generated according to explicit action statements in rule commands that are user-configurable. The EDRs are generated in comma separated values (CSV) format, generated as defined in traffic analysis rules.
- **User Detail Records (UDRs)**: Contain accounting information related to a specific mobile subscriber. The fields to be reported in them are user-configurable and are generated on any trigger of time threshold, volume threshold, handoffs, and call termination. The UDRs are generated in comma separated values (CSV) format, generated as defined in traffic analysis rules.

Important: This functionality is available for use with the Enhanced Charging Service which requires a session-use license. For more information on ECS, refer to the *Enhanced Charging Service Administration Guide*.

Diameter Credit Control Application

Provides a pre-paid billing mechanism for real-time cost and credit control based on the following **standards**:

- RFC 3588, Diameter Base Protocol, September 2003
- RFC 4006, Diameter Credit-Control Application, August 2005

The Diameter Credit Control Application (DCCA) is used to implement real-time credit-control for a variety of end user services such as network access, Session Initiation Protocol (SIP) services, messaging services, download services etc.

Used in conjunction with ECS, the DCCA interface uses a mechanism to allow the user to be informed of the charges to be levied for a requested service. In addition, there are services such as gaming and advertising that may credit as well as debit from a user account.

DCCA also supports the following:

- **Real-time Rate Service Information:** The ability to verify when end subscribers' accounts are exhausted or expired; or deny additional chargeable events.
- **Support for Multiple Services**: The usage of multiple services within one subscriber session is supported. Multiple Service support includes:
  - The ability to identify and process the service or group of services that are subject to different cost structures.
  - Independent credit control of multiple services in a single credit control sub-session.

Important: This functionality is available for use with the Enhanced Charging Service, which requires a session-use license. For more information on ECS, refer to the *Enhanced Charging Service Administration Guide*.

Accept TCP Connections from DCCA Server

This feature allows for peer Diameter Credit Control Application servers to initiate a connection the NGME.

This feature allows peer diameter nodes to connect to the NGME on TCP port 3868 when the diameter server is incapable of receiving diameter incoming diameter requests.

Important: For more information on Diameter support, refer to the *AAA and GTPP Interface Administration and Reference*.

Gy Interface Support

The Gy interface enables the wireless operator to implement a standardized interface for real time content based charging with differentiated rates for time based and volume based charging.

As it is based on a quota mechanism, the Gy interface enables the wireless operator to spare expensive Prepaid System resources.

As it enables time-, volume-, and event-based charging models, the Gy interface flexibly enables the operator to implement charging models tailored to their service strategies.

The Gy interface provides a standardized Diameter interface for real time content based charging of data services. It is based on the 3GPP standards and relies on quota allocation.

It provides an online charging interface that works with the ECS deep packet inspection feature. With Gy, customer traffic can be gated and billed in an “online” or “prepaid” style. Both time- and volume-based charging models are

supported. In all of these models, differentiated rates can be applied to different services based on shallow or deep packet inspection.

Gy is a Diameter interface. As such, it is implemented atop, and inherits features from, the Diameter Base Protocol. The system supports the applicable Base network and application features, including directly connected, relayed or proxied DCCA servers using TLS or plaintext TCP.

In the simplest possible installation, the system exchanges Gy Diameter messages over Diameter TCP links between itself and one "prepay" server. For a more robust installation, multiple servers would be used. These servers may optionally share or mirror a single quota database so as to support Gy session failover from one server to the other. For a more scalable installation, a layer of proxies or other Diameter agents can be introduced to provide features such as multi-path message routing or message and session redirection features.

The Cisco implementation is based on the following standards:

- RFC 4006 generic DCCA, including:
  - CCR Initial, Update, and Final signaling
  - ASR and RAR asynchronous DCCA server messages
  - Time, Total-Octets, and Service-Specific-Units quota management
  - Multiple independent quotas using Multiple-Services-Credit-Control
  - Rating-Group for quota-to-traffic association
  - CC-Failure-Handling and CC-Session-Failover features
  - Final-Unit-Action TERMINATE behavior
  - Tariff-Time-Change feature.
- 3GPP TS 32.299 online mode "Gy" DCCA, including:
  - Final-Unit-Action REDIRECT behavior
  - Quota-Holding-Time: This defines a user traffic idle time, on a per category basis, after which the usage is returned and no new quota is explicitly requested
  - Quota-Thresholds: These AVPs define a low value watermark at which new quota will be sought before the quota is entirely gone; the intent is to limit interruption of user traffic.

    These AVPs exist for all quota flavors, for example "Time-Quota-Threshold".
  - Trigger-Type: This AVP defines a set of events which will induce a re-authentication of the current session and its quota categories.

Gn/Gp Handoff Support

In LTE deployments, smooth handover support is required between 3G/2G and LTE networks, and Evolved Packet Core (EPC) is designed to be a common packet core for different access technologies. P-GW supports handovers as user equipment (UE) moves across different access technologies.

Cisco's P-GW supports inter-technology mobility handover between 4G and 3G/2G access. Interworking is supported between the 4G and 2G/3G SGSNs, which provide only Gn and Gp interfaces but no S3, S4 or S5/S8 interfaces. These Gn/Gp SGSNs provide no functionality introduced specifically for the evolved packet system (EPS) or for interoperation with the E-UTRAN. These handovers are supported only with a GTP-based S5/S8 and P-GW supports handoffs between GTPv2 based S5/S8 and GTPv1 based Gn/Gp tunneled connections. In this scenario, the P-GW works as an IP anchor for the EPC.

Important: To support the seamless handover of a session between GGSN and P-GW, the two independent services must be co-located on the same node and configured within the same context for optimum interoperation.

Important: For more information on Gn/GP handoffs, refer to *Gn/Gp GGSN/SGSN (GERAN/UTRAN)* in the *Supported Logical Network Interfaces (Reference Points)* section in this chapter.

IP Access Control Lists

IP access control lists allow you to set up rules that control the flow of packets into and out of the system based on a variety of IP packet parameters.

IP access lists, or access control lists (ACLs) as they are commonly referred to, are used to control the flow of packets into and out of the system. They are configured on a per-context basis and consist of "rules" (ACL rules) or filters that control the action taken on packets that match the filter criteria. Once configured, an ACL can be applied to any of the following:

- An individual interface
- All traffic facilitated by a context (known as a policy ACL)
- An individual subscriber
- All subscriber sessions facilitated by a specific context

Important: For more information on IP access control lists, refer to the *IP Access Control Lists* chapter in the *System Administration Guide*.

IP Address Hold Timers

Also known as address quarantining, this subscriber-level CLI introduces an address hold timer to temporarily buffer a previously assigned IP address from an IP address pool to prevent it from being recycled and reassigned to a new subscriber session. It is especially useful during inter-RAT handovers that sometimes lead to temporary loss of the mobile data session.

This feature provides a higher quality user experience for location-based services where the remote host server needs to reach the mobile device.

IPv6 Capabilities

Enables increased address efficiency and relieves pressures caused by rapidly approaching IPv4 address exhaustion problem.

The P-GW offers the following IPv6 capabilities:

Native IPv6 and IPv6 transport

- Support for any combination of IPv4, IPv6 or dual stack IPv4/v6 address assignment from dynamic or static address pools on the P-GW.

- Support for native IPv6 transport and service addresses on PMIPv6 S2a interface. Note that transport on GTP S5/S8 connections in this release is IPv4 based.
- Support for IPv6 transport for outbound traffic over the SGi reference interface to external Packet Data Networks.

IPv6 Connections to Attached Elements

IPv6 transport and interfaces are supported on all of the following connections:

- Diameter Gx policy signaling interface
- Diameter Gy online charging reference interface
- S6b authentication interface to external 3GPP AAA server
- Diameter Rf offline charging interface
- Lawful Intercept (X1, X2 interfaces)

Routing and Miscellaneous Features

- OSPFv3
- MP-BGP v6 extensions
- IPv6 flows (Supported on all Diameter QoS and Charging interfaces as well as Inline Services (e.g. ECS)

Local Break-Out

Provides a standards-based procedure to enable LTE operators to generate additional revenues by accepting traffic from visited subscribers based on roaming agreements with other mobile operators.

Local Breakout is a policy-based forwarding function that plays an important role in inter-provider roaming between LTE service provider networks. Local Breakout is determined by the SLAs for handling roaming calls between visited and home networks. In some cases, it is more beneficial to locally breakout a roaming call on a foreign network to the visited P-W rather than incur the additional transport costs to backhaul the traffic to the Home network.

If two mobile operators have a roaming agreement in place, Local Break-Out enables the visited user to attach to the V-PLMN network and be anchored by the local P-GW in the visited network. The roaming architecture relies on the HSS in the home network and also introduces the concept of the S9 policy signaling interface between the H-PCRF in the H-PLMN and the V-PCRF in the V-PLMN. When the user attaches to the EUTRAN cell and MME (Mobility Management Entity) in the visited network, the requested APN name in the S6a NAS signaling is used by the HSS in the H-PLMN to select the local S-GW (Serving Gateway) and P-GWs in the visited EPC network.

Management System Overview

The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management - focusing on providing superior quality network element (NE) and element management system (Cisco Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems - giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

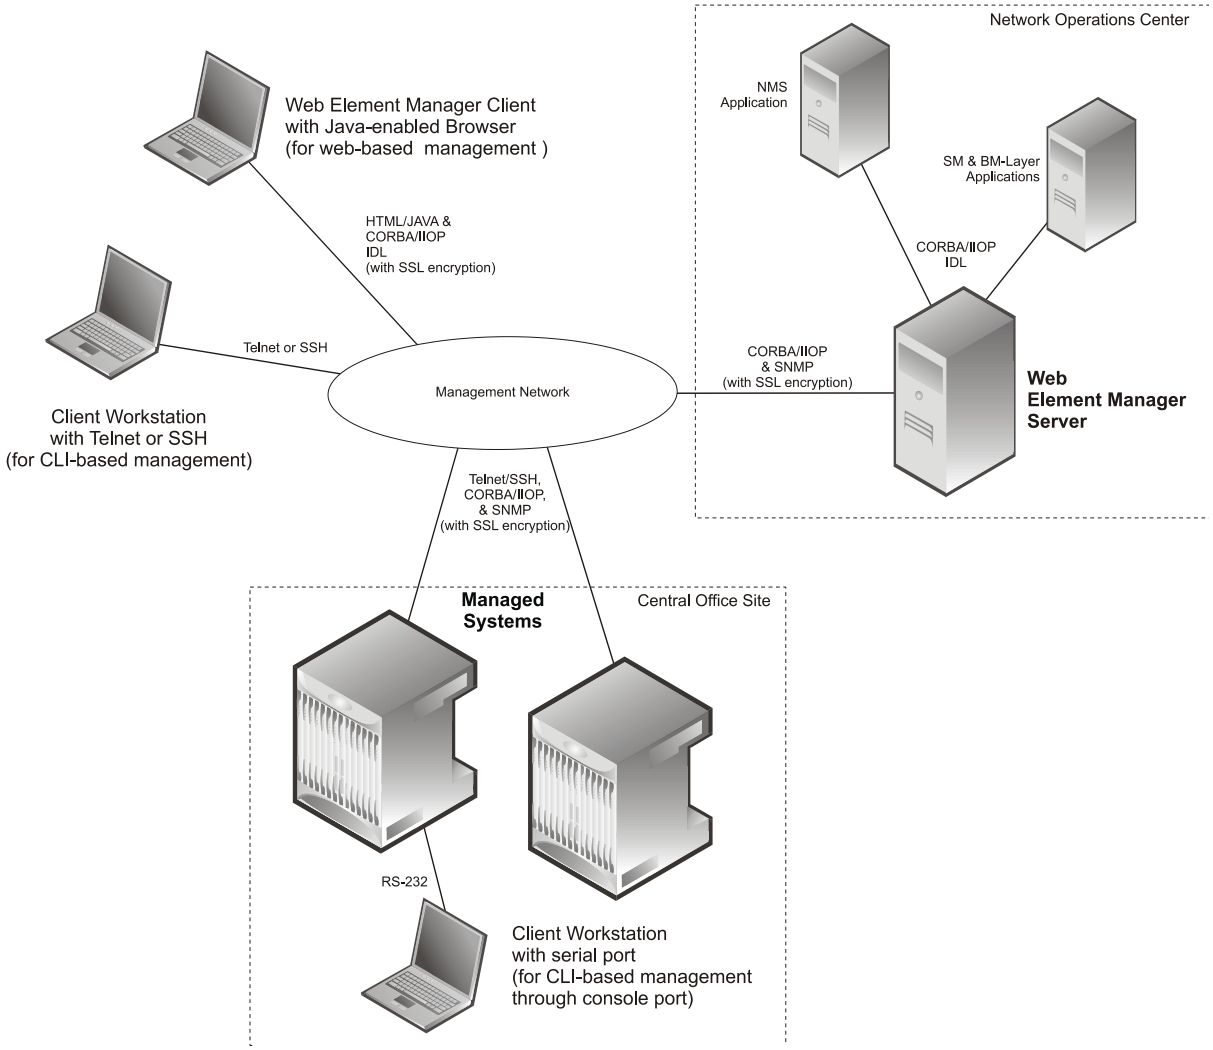
Cisco's O&M module offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces.

These include:

- Using the command line interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager application
- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000
- Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (i.e. Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network.

Figure 212. Element Management Methods



Important: P-GW management functionality is enabled by default for console-based access. For GUI-based management support, refer to the Web Element Management System section in this chapter.

Important: For more information on command line interface based management, refer to the *Command Line Interface Reference*.

Mobile IP Registration Revocation

Mobile IP registration revocation functionality provides the following benefits:

- Timely release of Mobile IP resources at the HSGW and/or P-GW
- Accurate accounting
- Timely notification to mobile node of change in service

Registration Revocation is a general mechanism whereby either the P-GW or the HSGW providing Mobile IP functionality to the same mobile node can notify the other mobility agent of the termination of a binding. Mobile IP Registration Revocation can be triggered at the HSGW by any of the following:

- Session terminated with mobile node for whatever reason
- Session renegotiation
- Administrative clearing of calls
- Session Manager software task outage resulting in the loss of HSGW sessions (sessions that could not be recovered)

Important: Registration Revocation functionality is also supported for Proxy Mobile IP. However, only the P-GW can initiate the revocation for Proxy-MIP calls.

Multiple PDN Support

Enables an APN-based user experience that enables separate connections to be allocated for different services including IMS, Internet, walled garden services, or off-deck content services.

The MAG function on the S-GW can maintain multiple PDN or APN connections for the same user session. The MAG runs a single node level Proxy Mobile IPv6 tunnel for all user sessions toward the LMA function of the P-GW. When a user wants to establish multiple PDN connections, the MAG brings up the multiple PDN connections over the same PMIPv6 session to one or more P-GW LMAs. The P-GW in turn allocates separate IP addresses (Home Network Prefixes) for each PDN connection and each one can run one or multiple EPC default & dedicated bearers. To request the various PDN connections, the MAG includes a common MN-ID and separate Home Network Prefixes, APNs and a Handover Indication Value equal to one in the PMIPv6 Binding Updates.

Important: Up to 11 multiple PDN connections are supported.

Non-Optimized e-HRPD to Native LTE (E-UTRAN) Mobility Handover

This feature enables a seamless inter-technology roaming capability in support of dual mode e-HRPD/e-UTRAN access terminals.

The non-optimized inter-technology mobility procedure is rooted at the P-GW as the mobility anchor point for supporting handovers for dual radio technology e-HRPD/E-UTRAN access terminals. To support this type of call handover, the P-GW supports handoffs between the GTP-based S5/S8 (GTPv2-C / GTPv1-U) and PMIPv6 S2a tunneled connections. It also provisions IPv4, IPv6, or dual stack IPv4/IPv6 PDN connections from a common address pool and preserves IP addresses assigned to the UE during inter-technology handover. In the current release, the native LTE (GTP-based) P-GW service address is IPv4-based, while the e-HRPD (PMIP) address is an IPv6 service address.

During the initial network attachment for each APN that the UE connects to, the HSS returns the FQDN of the P-GW for the APN. The MME uses DNS to resolve the P-GW address. When the PDN connection is established in the P-GW, the P-GW updates the HSS with the IP address of the P-GW on PDN establishment through the S6b authentication process. When the mobile user roams to the e-HRPD network, the HSS returns the IP address of the P-GW in the P-GW Identifier through the STa interface and the call ends up in the same P-GW. The P-GW is also responsible for initiating the session termination on the serving access connection after the call handover to the target network.

During the handover procedure, all dedicated EPS bearers must be re-established. On LTE- handovers to a target e-HRPD access network, the dedicated bearers are initiated by the mobile access terminal. In contrast, on handovers in the

opposite direction from e-HRPD to LTE access networks, the dedicated bearers are network initiated through Gx policy interactions with the PCRF server.

Finally, in order to support the inter-technology handovers, the P-GW uses common interfaces and Diameter endpoint addresses for the various reference points:

- S6b: Non-3GPP authentication
- Gx: QoS Policy and Charging
- Rf: Offline Charging

All three types of sessions are maintained during call handovers. The bearer binding will be performed by the HSGW during e-HRPD access and by the P-GW during LTE access. Thus, the Bearer Binding Event Reporting (BBERF) function needs to migrate between the P-GW and the HSGW during the handover. The HSGW establishes a Gxa session during e-HRPD access for bearer binding and releases the session during LTE access. The HSGW also maintains a limited context during the e-HRPD <->LTE handover to reduce latency in the event of a quick handover from the LTE RAN back to the e-HRPD network.

Important: For more information on handoff interfaces, refer to the *Supported Logical Network Interfaces (Reference Points)* section in this chapter.

Online/Offline Charging

The Cisco EPC platform offers support for online and offline charging interactions with external OCS and CGF/CDF servers.

Online Charging

Gy/Ro Reference Interfaces

The StarOS 9.0 online prepaid reference interface provides compatibility with the 3GPP TS 23.203, TS 32.240, TS 32.251 and TS 32.299 specifications. The Gy/Ro reference interface uses Diameter transport and IPv6 addressing. Online charging is a process whereby charging information for network resource usage must be obtained by the network in order for resource usage to occur. This authorization is granted by the Online Charging System (OCS) upon request from the network. The P-GW uses a charging characteristics profile to determine whether to activate or deactivate online charging. Establishment, modification or termination of EPS bearers is generally used as the event trigger on the PCRF to activate online charging PCC rules on the P-GW.

When receiving a network resource usage request, the network assembles the relevant charging information and generates a charging event towards the OCS in real-time. The OCS then returns an appropriate resource usage authorization that may be limited in its scope (e.g. volume of data or duration based). The OCS assigns quotas for rating groups and instructs the P-GW whether to continue or terminate service data flows or IP CAN bearers.

The following Online Charging models and functions are supported:

- Time based charging
- Volume based charging
- Volume and time based charging
- Final Unit Indication and termination or redirection of service data flows when quota is consumed

- Reauthorization triggers to rearm quotas for one or more rating groups using multi-service credit control (MSCC) instances
- Event based charging
- Billing cycle bandwidth rate limiting: Charging policy is enforced through interactions between the PDN GW and Online Charging Server. The charging enforcement point periodically conveys accounting information for subscriber sessions to the OCS and it is debited against the threshold that is established for the charging policy. Subscribers can be assigned a max usage for their tier (gold, silver, bronze for example), the usage can be tracked over a month, week, day, or peak time within a day. When the subscriber exceeds the usage limit, bandwidth is either restricted for a specific time period, or dropped depending on their tier of service.
- Fair usage controls

Offline Charging

Ga/Gz Reference Interfaces

The Cisco P-GW supports 3GPP-compliant offline charging as defined in TS 32.251, TS 32.297 and 32.298. Whereas the S-GW generates SGW-CDRs to record subscriber level access to PLMN resources, the P-GW creates PGW-CDRs to record user access to external networks. Additionally, when Gn/Gp interworking with pre-release SGSNs is enabled, the GGSN service on the P-GW records G-CDRs to record user access to external networks.

To provide subscriber level accounting, the Cisco S-GW and P-GWs support integrated Charging Transfer Functions (CTF) and Charging Data Functions (CDF). Each gateway uses Charging-ID's to distinguish between default and dedicated bearers within subscriber sessions. The Ga/Gz reference interface between the CDF and CGF is used to transfer charging records via the GTPP protocol. In a standards based implementation, the CGF consolidates the charging records and transfers them via an FTP/S-FTP connection over the Bm reference interface to a back-end billing mediation server. The Cisco EPC gateways also offer the ability to FTP/S-FTP charging records between the CDF and CGF server. CDR records include information such as Record Type, Served IMSI, ChargingID, APN Name, TimeStamp, Call Duration, Served MSISDN, PLMN-ID, etc. The ASR 5x00 platform offers a local directory to enable temporary file storage and buffer charging records in persistent memory located on a pair of dual redundant RAID hard disks. Each drive includes 147GB of storage and up to 100GB of capacity is dedicated to storing charging records. For increased efficiency it also possible to enable file compression using protocols such as GZIP. The Offline Charging implementation offers built-in heart beat monitoring of adjacent CGFs. If the Cisco P-GW has not heard from the neighbor CGF within the configurable polling interval, they will automatically buffer the charging records on the local drives until the CGF reactivates itself and is able to begin pulling the cached charging records.

The P-GW supports a Policy Charging Enforcement Function (PCEF) to enable Flow Based Bearer Charging (FBC) via the Gy reference interface to adjunct OCS servers (See Online Charging description above).

Rf Reference Interface

The Cisco EPC platforms also support the Rf reference interface to enable direct transfer of charging files from the CTF function of the P-GW to external CDF/CGF servers. This interface uses Diameter Accounting Requests (Start, Stop, Interim, and Event) to transfer charging records to the CDF/CGF. Each gateway relies on triggering conditions for reporting chargeable events to the CDF/CGF. Typically as EPS bearers are activated, modified or deleted, charging records are generated. The EPC platforms include information such as Subscription-ID (IMSI), Charging-ID (EPS bearer identifier) and separate volume counts for the uplink and downlink traffic.

Proxy Mobile IPv6 (S2a)

Provides a mobility management protocol to enable a single LTE-EPC core network to provide the call anchor point for user sessions as the subscriber roams between native EUTRAN and non-native e-HRPD access networks

S2a represents the trusted non-3GPP interface between the LTE-EPC core network and the evolved HRPD network anchored on the HSGW. In the e-HRPD network, network-based mobility provides mobility for IPv6 nodes without host involvement. Proxy Mobile IPv6 extends Mobile IPv6 signaling messages and reuses the HA function (now known as LMA) on the P-GW. This approach does not require the mobile node to be involved in the exchange of signaling messages between itself and the Home Agent. A proxy mobility agent (e.g. MAG function on HSGW) in the network performs the signaling with the home agent and does the mobility management on behalf of the mobile node attached to the network.

The S2a interface uses IPv6 for both control and data. During the PDN connection establishment procedures the P-GW allocates the IPv6 Home Network Prefix (HNP) via Proxy Mobile IPv6 signaling to the HSGW. The HSGW returns the HNP in router advertisement or based on a router solicitation request from the UE. PDN connection release events can be triggered by either the UE, the HSGW or the P-GW.

In Proxy Mobile IPv6 applications the HSGW (MAG function) and P-GW (LMA function) maintain a single shared tunnel and separate GRE keys are allocated in the PMIP Binding Update and Acknowledgement messages to distinguish between individual subscriber sessions. If the Proxy Mobile IP signaling contains Protocol Configuration Options (PCOs) it can also be used to transfer P-CSCF or DNS server addresses

QoS Bearer Management

Provides a foundation for contributing towards improved Quality of User Experience (QoE) by enabling deterministic end-to-end forwarding and scheduling treatments for different services or classes of applications pursuant to their requirements for committed bandwidth resources, jitter and delay. In this way, each application receives the service treatment that users expect.

An EPS bearer is a logical aggregate of one or more Service Data Flows (SDFs), running between a UE and a P-GW in case of GTP-based S5/S8, and between a UE and HSGW in case of PMIP-based S2a connection. An EPS bearer is the level of granularity for bearer level QoS control in the EPC/E-UTRAN. The Cisco P-GW maintains one or more Traffic Flow Templates (TFT's) in the downlink direction for mapping inbound Service Data Flows (SDFs) to EPS bearers. The P-GW maps the traffic based on the downlink TFT to the S5/S8 bearer. The Cisco PDN GW offers all of the following bearer-level aggregate constructs:

QoS Class Identifier (QCI): An operator provisioned value that controls bearer level packet forwarding treatments (e.g. scheduling weights, admission thresholds, queue management thresholds, link layer protocol configuration, etc). The Cisco EPC gateways also support the ability to map the QCI values to DiffServ code points in the outer GTP tunnel header of the S5/S8 connection. Additionally, the platform also provides configurable parameters to copy the DSCP marking from the encapsulated payload to the outer GTP tunnel header.

Guaranteed Bit Rate (GBR): A GBR bearer is associated with a dedicated EPS bearer and provides a guaranteed minimum transmission rate in order to offer constant bit rate services for applications such as interactive voice that require deterministic low delay service treatment.

Maximum Bit Rate (MBR): The MBR attribute provides a configurable burst rate that limits the bit rate that can be expected to be provided by a GBR bearer (e.g. excess traffic may get discarded by a rate shaping function). The MBR may be greater than or equal to the GBR for a given Dedicated EPS bearer.

Aggregate Maximum Bit Rate (AMBR): AMBR denotes a bit rate of traffic for a group of bearers destined for a particular PDN. The Aggregate Maximum Bit Rate is typically assigned to a group of Best Effort service data flows over the Default EPS bearer. That is, each of those EPS bearers could potentially utilize the entire AMBR, e.g. when the other EPS bearers do not carry any traffic. The AMBR limits the aggregate bit rate that can be expected to be provided

by the EPS bearers sharing the AMBR (e.g. excess traffic may get discarded by a rate shaping function). AMBR applies to all Non-GBR bearers belonging to the same PDN connection. GBR bearers are outside the scope of AMBR.

Policing and Shaping: The Cisco P-GW offers a variety of traffic conditioning and bandwidth management capabilities. These tools enable usage controls to be applied on a per-subscriber, per-EPS bearer or per-PDN/APN basis. It is also possible to apply bandwidth controls on a per-APN AMBR capacity. These applications provide the ability to inspect and maintain state for user sessions or Service Data Flows (SDFs) within them using shallow L3/L4 analysis or high touch deep packet inspection at L7. Metering of out-of-profile flows or sessions can result in packet discards or reducing the DSCP marking to Best Effort priority. When traffic shaping is enabled the P-GW enqueues the non-conforming session to the provisioned memory limit for the user session. When the allocated memory is exhausted, the inbound/outbound traffic for the user can be transmitted or policed in accordance with operator provisioned policy.

RADIUS Support

Provides a mechanism for performing authorization, authentication, and accounting (AAA) for subscriber PDP contexts based on the following standards:

- RFC-2618, RADIUS Authentication Client MIB, June 1999
- RFC-2620, RADIUS Accounting Client MIB, June 1999
- RFC-2865, Remote Authentication Dial In User Service (RADIUS), June 2000
- RFC-2866, RADIUS Accounting, June 2000
- RFC-2867, RADIUS Accounting Modifications for Tunnel Protocol Support, June 2000
- RFC-2868, RADIUS Attributes for Tunnel Protocol Support, June 2000
- RFC-2869, RADIUS Extensions, June 2000

The Remote Authentication Dial-In User Service (RADIUS) protocol is used to provide AAA functionality for subscriber PDP contexts. (RADIUS accounting is optional since GTPP can also be used.)

Within contexts configured on the system, there are AAA and RADIUS protocol-specific parameters that can be configured. The RADIUS protocol-specific parameters are further differentiated between RADIUS Authentication server RADIUS Accounting server interaction.

Among the RADIUS parameters that can be configured are:

- **Priority**: Dictates the order in which the servers are used allowing for multiple servers to be configured in a single context.
- **Routing Algorithm**: Dictate the method for selecting among configured servers. The specified algorithm dictates how the system distributes AAA messages across the configured AAA servers for new sessions. Once a session is established and an AAA server has been selected, all subsequent AAA messages for the session will be delivered to the same server.

In the event that a single server becomes unreachable, the system attempts to communicate with the other servers that are configured. The system also provides configurable parameters that specify how it should behave should all of the RADIUS AAA servers become unreachable.

The system provides an additional level of flexibility by supporting the configuration RADIUS server groups. This functionality allows operators to differentiate AAA services for subscribers based on the APN used to facilitate their PDP context.

In general, 128 AAA Server IP address/port per context can be configured on the system and it selects servers from this list depending on the server selection algorithm (round robin, first server). Instead of having a single list of servers per context, this feature provides the ability to configure multiple server groups. Each server group, in turn, consists of a list of servers.

This feature works in following way:

- All RADIUS authentication/accounting servers configured at the context-level are treated as part of a server group named "default". This default server group is available to all subscribers in that context through the realm (domain) without any configuration.
- It provides a facility to create "user defined" RADIUS server groups, as many as 399 (excluding "default" server group), within a context. Any of the user defined RADIUS server groups are available for assignment to a subscriber through the APN configuration within that context.

Since the configuration of the APN can specify the RADIUS server group to use as well as IP address pools from which to assign addresses, the system implements a mechanism to support some in-band RADIUS server implementations (i.e. RADIUS servers which are located in the corporate network, and not in the operator's network) where the NAS-IP address is part of the subscriber pool. In these scenarios, the P-GW supports the configuration of the first IP address of the subscriber pool for use as the RADIUS NAS-IP address.

Important: For more information on RADIUS AAA configuration, refer to the *AAA and GTPP Interface Administration and Reference*.

Source IP Address Validation

Insures integrity between the attached subscriber terminal and the PDN GW by mitigating the potential for unwanted spoofing or man-in-the-middle attacks.

The P-GW includes local IPv4/IPv6 address pools for assigning IP addresses to UEs on a per-PDN basis. The P-GW defends its provisioned address bindings by insuring that traffic is received from the host address that it has awareness of. In the event that traffic is received from a non-authorized host, the P-GW includes the ability to block the non-authorized traffic. The P-GW uses the IPv4 source address to verify the sender and the IPv6 source prefix in the case of IPv6.

Subscriber Level Trace

Provides a 3GPP standards-based session level trace function for call debugging and testing new functions and access terminals in an LTE environment.

As a complement to Cisco's protocol monitoring function, the P-GW supports 3GPP standards based session level trace capabilities to monitor all call control events on the respective monitored interfaces including S5/S8, S2a, SGi, and Gx. The trace can be initiated using multiple methods:

- Management initiation via direct CLI configuration
- Management initiation at HSS with trace activation via authentication response messages over S6a reference interface
- Signaling based activation through signaling from subscriber access terminal

Important: Once the trace is provisioned, it can be provisioned through the access cloud via various signaling interfaces.

The session level trace function consists of trace activation followed by triggers. The time between the two events is treated much like Lawful Intercept where the EPC network element buffers the trace activation instructions for the provisioned subscriber in memory using camp-on monitoring. Trace files for active calls are buffered as XML files

using non-volatile memory on the local dual redundant hard drives on the ASR 5x00 platform. The Trace Depth defines the granularity of data to be traced. Six levels are defined including Maximum, Minimum and Medium with ability to configure additional levels based on vendor extensions.

All call control activity for active and recorded sessions is sent to an off-line Trace Collection Entity (TCE) using a standards-based XML format over a FTP or secure FTP (SFTP) connection. In the current release the IPv4 interfaces are used to provide connectivity to the TCE. Trace activation is based on IMSI or IMEI. Once a subscriber level trace request is activated it can be propagated via the S5/S8 signaling to provision the corresponding trace for the same subscriber call on the P-GW. The trace configuration will only be propagated if the P-GW is specified in the list of configured Network Element types received by the S-GW. Trace configuration can be specified or transferred in any of the following message types:

- S5/S8: Create Session Request
- S5/S8: Modify Bearer Request
- S5/S8: Trace Session Activation (New message defined in TS 32.422)

Performance Goals: As subscriber level trace is a CPU intensive activity the max number of concurrently monitored trace sessions per Cisco P-GW is 32. Use in a production network should be restricted to minimize the impact on existing services.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions with in a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert**: A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm**: Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps**: SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values.

  Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.

- **Logs**: The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

  Logs are supported in both the Alert and the Alarm models.

- **Alarm System**: High threshold alarms generated within the specified polling interval are considered "outstanding" until a the condition no longer exists or a condition clear alarm is generated. "Outstanding"

alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

Virtual APN Support

Virtual APNs allow differentiated services within a single APN.

The Virtual APN feature allows a carrier to use a single APN to configure differentiated services. The APN that is supplied by the MME is evaluated by the P-GW in conjunction with multiple configurable parameters. Then, the P-GW selects an APN configuration based on the supplied APN and those configurable parameters.

APN configuration dictates all aspects of a session at the P-GW. Different policies imply different APNS. After basic APN selection, however, internal re-selection can occur based on the following parameters:

- Service name
- Subscriber type
- MCC-MNC of IMSI
- Domain name part of username (user@domain)
- S-GW address

Features and Functionality - Inline Service Support

This section describes the features and functions of inline services supported on the P-GW. These services require additional licenses to implement the functionality.

This section describes the following features:

- Content Filtering
- Header Enrichment: Header Insertion and Encryption
- Mobile Video Gateway
- Network Address Translation (NAT)
- Peer-to-Peer Detection
- Personal Stateful Firewall
- Traffic Performance Optimization (TPO)

Content Filtering

The Cisco P-GW offers two variants of network-controlled content filtering / parental control services. Each approach leverages the native DPI capabilities of the platform to detect and filter events of interest from mobile subscribers based on HTTP URL or WAP/MMS URI requests:

- Integrated Content Filtering: A turnkey solution featuring a policy enforcement point and category based rating database on the Cisco P-GW. An offboard AAA or PCRF provides the per-subscriber content filtering information as subscriber sessions are established. The content filtering service uses DPI to extract URLs or URIs in HTTP request messages and compares them against a static rating database to determine the category match. The provisioned policy determines whether individual subscribers are entitled to view the content.
- Content Filtering ICAP Interface: This solution is appropriate for mobile operators with existing installations of Active Content Filtering external servers. The service continues to harness the DPI functions of the ASR 5x00 platform to extract events of interest. However in this case, the extracted requests are transferred via the Integrated Content Adaptation Protocol (ICAP) with subscriber identification information to the external ACF server which provides the category rating database and content decision functions.

Integrated Adult Content Filter

Provides a value-added service to prevent unintended viewing of objectionable content that exploits underage children. Content Filtering offers mobile operators a way to increase data ARPU and subscriber retention through a network-based solution for parental controls and content filtering. The integrated solution enables a single policy decision and enforcement point thereby streamlining the number of signaling interactions with external AAA/Policy Manager servers. When used in parallel with other services such as Enhanced Content Charging (ECS) it increases billing accuracy of charging records by insuring that mobile subscribers are only charged for visited sites they are allowed to access.

The Integrated Adult Content Filter is a subscriber-aware inline service provisioned on an ASR 5x00 running P-GW services. Integrated Content Filtering utilizes the local DPI engine and harnesses a distributed software architecture that scales with the number of active P-GW sessions on the system.

Content Filtering policy enforcement is the process of deciding if a subscriber should be able to receive some content. Typical options are to allow, block, or replace/redirect the content based on the rating of the content and the policy defined for that content and subscriber. The policy definition is transferred in an authentication response from a AAA server or Diameter policy message via the Gx reference interface from an adjunct PCRF. The policy is applied to subscribers through rulebase or APN/Subscriber configuration. The policy determines the action to be taken on the content request on the basis of its category. A maximum of one policy can be associated with a rulebase.

ICAP Interface

Provides a value-added service to prevent unintended viewing of objectionable content that exploits underage children. Content Filtering offers mobile operators a way to increase data ARPU and subscriber retention through a network-based solution for parental controls and content filtering. The Content Filtering ICAP solution is appropriate for operators with existing installations of Active Content Filtering servers in their networks.

The Enhanced Charging Service (ECS) for the P-GW provides a streamlined Internet Content Adaptation Protocol (ICAP) interface to leverage the Deep Packet Inspection (DPI) to enable external Application Servers to provide their services without performing the DPI functionality and without being inserted in the data flow. The ICAP interface may be attractive to mobile operators that prefer to use an external Active Content Filtering (ACF) Platform. If a subscriber initiates a WAP (WAP1.x or WAP2.0) or Web session, the subsequent GET/POST request is detected by the deep packet inspection function. The URL of the GET/POST request is extracted by the local DPI engine on the ASR 5x00 platform and passed, along with subscriber identification information and the subscriber request, in an ICAP message to the Application Server (AS). The AS checks the URL on the basis of its category and other classifications like, type, access level, content category and decides if the request should be authorized, blocked or redirected by answering the GET/POST message. Depending upon the response received from the ACF server, the P-GW either passes the request unmodified or discards the message and responds to the subscriber with the appropriate redirection or block message.

Header Enrichment: Header Insertion and Encryption

Header enrichment provides a value-added capability for mobile operators to monetize subscriber intelligence to include subscriber-specific information in the HTTP requests to application servers.

Extension header fields (x-header) are the fields that can be added to headers of a protocol for a specific purpose. The enriched header allows additional entity-header fields to be defined without changing the protocol, but these fields cannot be assumed to be recognizable by the recipient. Unrecognized fields should be ignored by the recipient and must be forwarded by transparent proxies.

Extension headers can be supported in HTTP/WSP GET and POST request packets. The Enhanced Charging Service (ECS) for the P-GW offers APN-based configuration and rules to insert x-headers in HTTP/WSP GET and POST request packets. The charging action associated with the rules will contain the list of x-headers to be inserted in the packets. Protocols supported are HTTP, WAP 1.0 and WAP 2.0 GET, and POST messages.

Important: For more information on ECS, see the *Enhanced Charging Service Administration Guide*.

The data passed in the inserted HTTP header attributes is used by end application servers (also known as Upsell Servers) to identify subscribers and session information. These servers provide information customized to that specific subscriber.

The Cisco P-GW can include the following information in the http header:

- User-customizable, arbitrary text string
- Subscriber's MSISDN (the RADIUS calling-station-id, in clear text)

- Subscriber‘s IMSI
- Subscriber‘s IP address
- S-GW IP address (in clear text)

X-Header encryption enhances the header enrichment feature by increasing the number of fields that can be supported and through encryption of the fields before inserting them.

The following limitations to insertion of x-header fields in WSP headers apply:

- x-header fields are not inserted in IP fragmented packets.
- In case of concatenated request, x-header fields are only inserted in first GET or POST request (if rule matches for the same). X-header fields are not inserted in the second or later GET/POST requests in the concatenated requests. For example, if there is ACK+GET in packet, x-header is inserted in the GET packet. However, if GET1+GET2 is present in the packet and rule matches for GET2 and not GET1 x-header is still inserted in GET2. In case of GET+POST also, x-header is not inserted in POST.
- In case of CO, x-header fields are not inserted if the WTP packets are received out of order (even after proper reordering).
- If route to MMS is present, x-headers are not inserted.
- x-headers are not inserted in WSP POST packet when header is segmented. This is because POST contains header length field which needs to be modified after addition of x-headers. In segmented WSP headers, header length field may be present in one packet and header may complete in another packet.

Mobile Video Gateway

The Cisco ASR 5x00 chassis provides mobile operators with a flexible solution that functions as a Mobile Video Gateway in 2.5G, 3G, and 4G wireless data networks.

The Cisco Mobile Video Gateway consists of new software for the ASR 5x00. The Mobile Video Gateway is the central component of the Cisco Mobile Videoscape. It employs a number of video optimization techniques that enable mobile operators to enhance the video experience for their subscribers while optimizing the performance of video content transmission through the mobile network.

The Mobile Video Gateway features and functions include:

- DPI (Deep Packet Inspection) to identify subscriber requests for video vs. non-video content
- Transparent video re-addressing to the Cisco CAE (Content Adaptation Engine) for retrieval of optimized video content
- CAE load balancing of HTTP video requests among the CAEs in the server cluster
- Video optimization policy control for tiered subscriber services
- Video white-listing, which excludes certain video clips from video optimization
- Video pacing for “just in time” video downloading
- TCP link monitoring
- Dynamic inline transrating
- Dynamically-enabled TCP proxy
- Traffic performance optimization
- N+1 redundancy support
- SNMP traps and alarms (threshold crossing alerts)

- Mobile video statistics
- Bulk statistics for mobile video

The Cisco CAE is an optional component of the Cisco Mobile Videoscape. It runs on the Cisco UCS (Unified Computing System) platform and functions in a UCS server cluster to bring additional video optimization capabilities to the Mobile Videoscape. For information about the features and functions of the Cisco CAE, see the CAE product documentation.

Important: For more information on the Mobile Video Gateway, refer to the *Mobile Video Gateway Administration Guide*.

Network Address Translation (NAT)

NAT translates non-routable private IP address(es) to routable public IP address(es) from a pool of public IP addresses that have been designated for NAT. This enables to conserve on the number of public IP addresses required to communicate with external networks, and ensures security as the IP address scheme for the internal network is masked from external hosts, and each outgoing and incoming packet goes through the translation process.

NAT works by inspecting both incoming and outgoing IP datagrams and, as needed, modifying the source IP address and port number in the IP header to reflect the configured NAT address mapping for outgoing datagrams. The reverse NAT translation is applied to incoming datagrams.

NAT can be used to perform address translation for simple IP and mobile IP. NAT can be selectively applied/denied to different flows (5-tuple connections) originating from subscribers based on the flows' L3/L4 characteristics—Source-IP, Source-Port, Destination-IP, Destination-Port, and Protocol.

NAT supports the following mappings:

- One-to-One
- Many-to-One

Important: For more information on NAT, refer to the *Network Address Translation Administration Guide*.

Peer-to-Peer Detection

Allows operators to identify P2P traffic in the network and applying appropriate controlling functions to ensure fair distribution of bandwidth to all subscribers.

Peer-to-Peer (P2P) is a term used in two slightly different contexts. At a functional level, it means protocols that interact in a peering manner, in contrast to client-server manner. There is no clear differentiation between the function of one node or another. Any node can function as a client, a server, or both—a protocol may not clearly differentiate between the two. For example, peering exchanges may simultaneously include client and server functionality, sending and receiving information.

Detecting peer-to-peer protocols requires recognizing, in real time, some uniquely identifying characteristic of the protocols. Typical packet classification only requires information uniquely typed in the packet header of packets of the stream(s) running the particular protocol to be identified. In fact, many peer-to-peer protocols can be detected by simple packet header inspection. However, some P2P protocols are different, preventing detection in the traditional manner. This is designed into some P2P protocols to purposely avoid detection. The creators of these protocols purposely do not publish specifications. A small class of P2P protocols is stealthier and more challenging to detect. For some protocols no set of fixed markers can be identified with confidence as unique to the protocol.

Operators care about P2P traffic because of the behavior of some P2P applications (for example, Bittorrent, Skype, and eDonkey). Most P2P applications can hog the network bandwidth such that 20% P2P users can generate as much as traffic generated by the rest 80% non-P2P users. This can result into a situation where non-P2P users may not get enough network bandwidth for their legitimate use because of excess usage of bandwidth by the P2P users. Network operators need to have dynamic network bandwidth / traffic management functions in place to ensure fair distributions of the network bandwidth among all the users. And this would include identifying P2P traffic in the network and applying appropriate controlling functions to the same (for example, content-based premium billing, QoS modifications, and other similar treatments).

Cisco's P2P detection technology makes use of innovative and highly accurate protocol behavioral detection techniques.

Important: For more information on peer-to-peer detection, refer to the *Application Detection and Control Administration Guide*.

Personal Stateful Firewall

The Personal Stateful Firewall is an in-line service feature that inspects subscriber traffic and performs IP session-based access control of individual subscriber sessions to protect the subscribers from malicious security attacks.

The Personal Stateful Firewall supports stateless and stateful inspection and filtering based on the configuration.

In stateless inspection, the firewall inspects a packet to determine the 5-tuple—source and destination IP addresses and ports, and protocol—information contained in the packet. This static information is then compared against configurable rules to determine whether to allow or drop the packet. In stateless inspection the firewall examines each packet individually, it is unaware of the packets that have passed through before it, and has no way of knowing if any given packet is part of an existing connection, is trying to establish a new connection, or is a rogue packet.

In stateful inspection, the firewall not only inspects packets up through the application layer / layer 7 determining a packet's header information and data content, but also monitors and keeps track of the connection's state. For all active connections traversing the firewall, the state information, which may include IP addresses and ports involved, the sequence numbers and acknowledgement numbers of the packets traversing the connection, TCP packet flags, etc. is maintained in a state table. Filtering decisions are based not only on rules but also on the connection state established by prior packets on that connection. This enables to prevent a variety of DoS, DDoS, and other security violations. Once a connection is torn down, or is timed out, its entry in the state table is discarded.

The Enhanced Charging Service (ECS) / Active Charging Service (ACS) in-line service is the primary vehicle that performs packet inspection and charging. For more information on ECS, see the *Enhanced Charging Service Administration Guide*.

Important: For more information on Personal Stateful Firewall, refer to the *Personal Stateful Firewall Administration Guide*.

Traffic Performance Optimization (TPO)

Though TCP is a widely accepted protocol in use today, it is optimized only for wired networks. Due to inherent reliability of wired networks, TCP implicitly assumes that any packet loss is due to network congestion and consequently invokes congestion control measures. However, wireless links are known to experience sporadic and usually temporary losses due to several reasons, including the following, which also trigger TCP congestion control measures resulting in poor TCP performance.

Reasons for delay variability over wireless links include:

- Channel fading effect, subscriber mobility, and other transient conditions
- Link-layer retransmissions
- Handoffs between neighboring cells
- Intermediate nodes, such as SGSN and e-NodeB, implementing scheduling polices tuned to deliver better QoS for select services; resulting is variable delay in packet delivery for other services

The TPO inline service uses a combination of TCP and HTTP optimization techniques to improve TCP performance over wireless links.

Important: For more information on TPO, refer to the *Traffic Performance Optimization Administration Guide*.

Features and Functionality - External Application Support

This section describes the features and functions of external applications supported on the P-GW. These services require additional licenses to implement the functionality.

This section describes the following feature(s):

- Web Element Management System

Web Element Management System

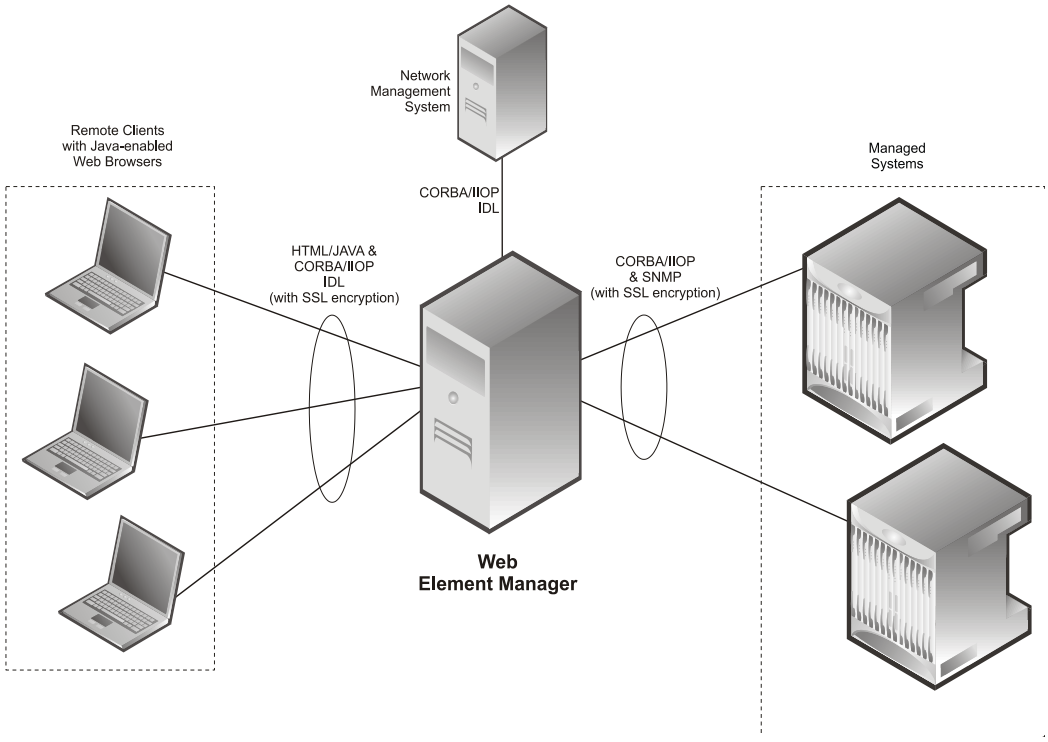
Provides a graphical user interface (GUI) for performing fault, configuration, accounting, performance, and security (FCAPS) management of the ASR 5x00.

The Web Element Manager is a Common Object Request Broker Architecture (CORBA)-based application that provides complete fault, configuration, accounting, performance, and security (FCAPS) management capability for the system.

For maximum flexibility and scalability, the Web Element Manager application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the fault-tolerant Sun® Solaris® operating system.

The following figure demonstrates various interfaces between the Cisco Web Element Manager and other network components.

Figure 213. Web Element Manager Network Interfaces



Important: For more information on WEM support, refer to the *WEM Installation and Administration Guide*.

Features and Functionality - Optional Enhanced Feature Software

This section describes the optional enhanced features and functions for the P-GW service.

Each of the following features requires the purchase of an additional license to implement the functionality with the P-GW service.


Important: For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

This section describes the following features:

- Dynamic RADIUS Extensions (Change of Authorization)
- GRE Protocol Interface Support
- Inter-Chassis Session Recovery
- IP Security (IPSec) Encryption
- L2TP LAC Support
- Lawful Intercept
- Layer 2 Traffic Management (VLANs)
- Local Policy Decision Engine
- NEMO Service Supported
- Session Recovery Support
- Traffic Policing and Shaping
- User Location Information Reporting

Dynamic RADIUS Extensions (Change of Authorization)

Use of Dynamic RADIUS Extensions (CoA and PoD) requires that a valid license key be installed. Contact your Cisco account representative for information on how to obtain a license.

Dynamic RADIUS extension support provide operators with greater control over subscriber PDP contexts by providing the ability to dynamically redirect data traffic, and or disconnect the PDP context.

This functionality is based on the RFC 3576, Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS), July 2003 standard.

The system supports the configuration and use of the following dynamic RADIUS extensions:

- **Change of Authorization**: The system supports CoA messages from the AAA server to change data filters associated with a subscriber session. The CoA request message from the AAA server must contain attributes to identify NAS and the subscriber session and a data filter ID for the data filter to apply to the subscriber session.
- **Disconnect Message**: The DM message is used to disconnect subscriber sessions in the system from a RADIUS server. The DM request message should contain necessary attributes to identify the subscriber session.

The above extensions can be used to dynamically re-direct subscriber PDP contexts to an alternate address for performing functions such as provisioning and/or account set up. This functionality is referred to as Session Redirection, or Hotlining.

Session redirection provides a means to redirect subscriber traffic to an external server by applying ACL rules to the traffic of an existing or a new subscriber session. The destination address and optionally the destination port of TCP/IP or UDP/IP packets from the subscriber are rewritten so the packet is forwarded to the designated redirected address.

Return traffic to the subscriber has the source address and port rewritten to the original values. The redirect ACL may be applied dynamically by means of the Radius Change of Authorization (CoA) extension.

GRE Protocol Interface Support

Use of GRE Interface Tunneling requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

The P-GW supports GRE generic tunnel interfaces in accordance with RFC 2784, Generic Routing Encapsulation (GRE). The GRE protocol allows mobile users to connect to their enterprise networks through GRE tunnels.

GRE tunnels can be used by the enterprise customers of a carrier 1) To transport AAA packets corresponding to an APN over a GRE tunnel to the corporate AAA servers and, 2) To transport the enterprise subscriber packets over the GRE tunnel to the corporation gateway.

The corporate servers may have private IP addresses and hence the addresses belonging to different enterprises may be overlapping. Each enterprise needs to be in a unique virtual routing domain, known as VRF. To differentiate the tunnels between same set of local and remote ends, GRE Key will be used as a differentiation.

GRE tunneling is a common technique to enable multi-protocol local networks over a single-protocol backbone, to connect non-contiguous networks and allow virtual private networks across WANs. This mechanism encapsulates data packets from one protocol inside a different protocol and transports the data packets unchanged across a foreign network. It is important to note that GRE tunneling does not provide security to the encapsulated protocol, as there is no encryption involved (like IPSec offers, for example).

GRE tunneling consists of three main components:

- Passenger protocol-protocol being encapsulated. For example: CLNS, IPv4 and IPv6.
- Carrier protocol-protocol that does the encapsulating. For example: GRE, IP-in-IP, L2TP, MPLS and IPSec.
- Transport protocol-protocol used to carry the encapsulated protocol. The main transport protocol is IP.

Inter-Chassis Session Recovery

Use of Interchassis Session Recovery requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

The ASR 5x00 provides industry leading carrier class redundancy. The systems protects against all single points of failure (hardware and software) and attempts to recover to an operational state when multiple simultaneous failures occur.

The system provides several levels of system redundancy:

- Under normal N+1 PSC/PSC2 hardware redundancy, if a catastrophic packet processing card failure occurs all affected calls are migrated to the standby packet processing card if possible. Calls which cannot be migrated are gracefully terminated with proper call-termination signaling and accounting records are generated with statistics accurate to the last internal checkpoint

- If the Session Recovery feature is enabled, any total PSC/PSC2 failure will cause a PSC switchover and all established sessions for supported call-types are recovered without any loss of session.

Even though Cisco provides excellent intra-chassis redundancy with these two schemes, certain catastrophic failures which can cause total chassis outages, such as IP routing failures, line-cuts, loss of power, or physical destruction of the chassis, cannot be protected by this scheme. In such cases, the MME Inter-Chassis Session Recovery feature provides geographic redundancy between sites. This has the benefit of not only providing enhanced subscriber experience even during catastrophic outages, but can also protect other systems such as the RAN from subscriber re-activation storms.

The Interchassis Session Recovery feature allows for continuous call processing without interrupting subscriber services. This is accomplished through the use of redundant chassis. The chassis are configured as primary and backup with one being active and one in recovery mode. A checkpoint duration timer is used to control when subscriber data is sent from the active chassis to the inactive chassis. If the active chassis handling the call traffic goes out of service, the inactive chassis transitions to the active state and continues processing the call traffic without interrupting the subscriber session. The chassis determines which is active through a propriety TCP-based connection called a redundancy link. This link is used to exchange Hello messages between the primary and backup chassis and must be maintained for proper system operation.

- **Interchassis Communication**

  Chassis configured to support Interchassis Session Recovery communicate using periodic Hello messages. These messages are sent by each chassis to notify the peer of its current state. The Hello message contains information about the chassis such as its configuration and priority. A dead interval is used to set a time limit for a Hello message to be received from the chassis' peer. If the standby chassis does not receive a Hello message from the active chassis within the dead interval, the standby chassis transitions to the active state. In situations where the redundancy link goes out of service, a priority scheme is used to determine which chassis processes the session. The following priority scheme is used:

  - router identifier
  - chassis priority
  - SPIO MAC address

- **Checkpoint Messages**

  Checkpoint messages are sent from the active chassis to the inactive chassis. Checkpoint messages are sent at specific intervals and contain all the information needed to recreate the sessions on the standby chassis, if that chassis were to become active. Once a session exceeds the checkpoint duration, checkpoint data is collected on the session. The checkpoint parameter determines the amount of time a session must be active before it is included in the checkpoint message.

Important: For more information on inter-chassis session recovery support, refer to the *Interchassis Session Recovery* chapter in the *System Administration Guide*.

IP Security (IPSec) Encryption

Use of Network Domain Security requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

IPSec encryption enables network domain security for all IP packet switched LTE-EPC networks in order to provide confidentiality, integrity, authentication, and anti-replay protection. These capabilities are insured through use of cryptographic techniques.

The Cisco P-GW supports IKEv1 and IPSec encryption using IPv4 addressing. IPSec enables the following two use cases:

- Encryption of S8 sessions and EPS bearers in roaming applications where the P-GW is located in a separate administrative domain from the S-GW
- IPSec ESP security in accordance with 3GPP TS 33.210 is provided for S1 control plane, S1 bearer plane and S1 management plane traffic. Encryption of traffic over the S1 reference interface is desirable in cases where the EPC core operator leases radio capacity from a roaming partner's network.

L2TP LAC Support

Use of L2TP LAC requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

The system configured as a Layer 2 Tunneling Protocol Access Concentrator (LAC) enables communication with L2TP Network Servers (LNSs) for the establishment of secure Virtual Private Network (VPN) tunnels between the operator and a subscriber's corporate or home network.

The use of L2TP in VPN networks is often used as it allows the corporation to have more control over authentication and IP address assignment. An operator may do a first level of authentication, however use PPP to exchange user name and password, and use IPCP to request an address. To support PPP negotiation between the P-GW and the corporation, an L2TP tunnel must be setup in the P-GW running a LAC service.

L2TP establishes L2TP control tunnels between LAC and LNS before tunneling the subscriber PPP connections as L2TP sessions. The LAC service is based on the same architecture as the P-GW and benefits from dynamic resource allocation and distributed message and data processing.

The LAC sessions can also be configured to be redundant, thereby mitigating any impact of hardware or software issues. Tunnel state is preserved by copying the information across processor cards.

Lawful Intercept

The feature use license for Lawful Intercept on the P-GW is included in the P-GW session use license.

The Cisco Lawful Intercept feature is supported on the P-GW. Lawful Intercept is a licensed-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your Cisco account representative.

Layer 2 Traffic Management (VLANs)

Use of Layer 2 Traffic Management requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

Virtual LANs (VLANs) provide greater flexibility in the configuration and use of contexts and services.

VLANs are configured as “tags” on a per-port basis and allow more complex configurations to be implemented. The VLAN tag allows a single physical port to be bound to multiple logical interfaces that can be configured in different contexts; therefore, each Ethernet port can be viewed as containing many logical ports when VLAN tags are employed.

Important: For more information on VLAN support, refer to the *VLANs* chapter in the *System Administration Guide*.

Local Policy Decision Engine

Use of the Local Policy Decision Engine requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

The Local Policy Engine is an event-driven rules engine that offers Gx-like QoS and policy controls to enable user or application entitlements. As the name suggests, it is designed to provide a subset of a PCRF in cases where an operator elects not to use a PCRF or scenarios where connections to an external PCRF are disrupted. Local policies are used to control different aspects of a session like QoS, data usage, subscription profiles, and server usage by means of locally defined policies. A maximum of 1,024 local policies can be provisioned on a P-GW system.

Local policies are triggered when certain events occur and the associated conditions are satisfied. For example, when a new call is initiated, the QoS to be applied for the call could be decided based on the IMSI, MSISDN, and APN.

Potential uses cases for the Local Policy Decision Engine include:

- Disaster recovery data backup solution in the event of a loss of PCRF in a mobile core network.
- Dedicated bearer establishment for emergency voice calls.
- Network-initiated bearer establishment on LTE to non-3GPP inter-domain handovers.

NEMO Service Supported

Use of NEMO requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

The P-GW may be configured to enable or disable Network Mobility (NEMO) service.

When enabled, the system includes NEMO support for a Mobile IPv4 Network Mobility (NEMO-HA) on the P-GW platform to terminate Mobile IPv4 based NEMO connections from Mobile Routers (MRs) that attach to an Enterprise PDN. The NEMO functionality allows bi-directional communication that is application-agnostic between users behind the MR and users or resources on Fixed Network sites.

The same NEMO4G-HA service and its bound Loopback IP address supports NEMO connections whose underlying PDN connection comes through GTP S5 (4G access) or PMIPv6 S2a (eHRPD access).

Session Recovery Support

The feature use license for Session Recovery on the P-GW is included in the P-GW session use license.

Session recovery provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

In the telecommunications industry, over 90 percent of all equipment failures are software-related. With robust hardware failover and redundancy protection, any card-level hardware failures on the system can quickly be corrected. However, software failures can occur for numerous reasons, many times without prior indication. StarOS Release 9.0 adds the ability to support stateful intra-chassis session recovery for P-GW sessions.

When session recovery occurs, the system reconstructs the following subscriber information:

- Data and control state information required to maintain correct call behavior
- Subscriber data statistics that are required to ensure that accounting information is maintained
- A best-effort attempt to recover various timer values such as call duration, absolute time, and others

Session recovery is also useful for in-service software patch upgrade activities. If session recovery is enabled during the software patch upgrade, it helps to preserve existing sessions on the active PSC/PSC2 during the upgrade process.

Important: For more information on session recovery support, refer to the *Session Recovery* chapter in the *System Administration Guide*.

Traffic Policing and Shaping

Use of Per-Subscriber Traffic Policing/Shaping requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

Traffic policing and shaping allows you to manage bandwidth usage on the network and limit bandwidth allowances to subscribers. Shaping allows you to buffer excesses to be delivered at a later time.

Traffic Policing

Traffic policing enables the configuring and enforcing of bandwidth limitations on individual subscribers and/or APNs of a particular traffic class in 3GPP/3GPP2 service.

Bandwidth enforcement is configured and enforced independently on the downlink and the uplink directions.

A Token Bucket Algorithm (a modified trTCM) [RFC2698] is used to implement the Traffic-Policing feature. The algorithm used measures the following criteria when determining how to mark a packet:

- Committed Data Rate (CDR): The guaranteed rate (in bits per second) at which packets can be transmitted/received for the subscriber during the sampling interval.
- Peak Data Rate (PDR): The maximum rate (in bits per second) that subscriber packets can be transmitted/received for the subscriber during the sampling interval.
- Burst-size: The maximum number of bytes that can be transmitted/received for the subscriber during the sampling interval for both committed (CBS) and peak (PBS) rate conditions. This represents the maximum number of tokens that can be placed in the subscriber's "bucket". Note that the committed burst size (CBS) equals the peak burst size (PBS) for each subscriber.

The system can be configured to take any of the following actions on packets that are determined to be in excess or in violation:

- Drop: The offending packet is discarded.
- Transmit: The offending packet is passed.
- Lower the IP Precedence: The packet's ToS bit is set to "0", thus downgrading it to Best Effort, prior to passing the packet. Note that if the packet's ToS bit was already set to "0", this action is equivalent to "Transmit".

Traffic Shaping

Traffic Shaping is a rate limiting method similar to the Traffic Policing, but provides a buffer facility for packets exceeded the configured limit. Once the packet exceeds the data-rate, the packet queued inside the buffer to be delivered at a later time.

The bandwidth enforcement can be done in the downlink and the uplink direction independently. If there is no more buffer space available for subscriber data system can be configured to either drop the packets or kept for the next scheduled traffic session.

User Location Information Reporting

Use of User Location Information (ULI) Reporting requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

ULI Reporting allows the eNodeB to report the location of a UE to the MME, when requested by a P-GW.

The following procedures are used over the S1-MME interface to initiate and stop location reporting between the MME and eNodeB:

- **Location Reporting Control**: The purpose of Location Reporting Control procedure is to allow the MME to request that the eNodeB report where the UE is currently located. This procedure uses UE-associated signaling.
- **Location Report Failure Indication**: The Location Report Failure Indication procedure is initiated by an eNodeB in order to inform the MME that a Location Reporting Control procedure has failed. This procedure uses UE-associated signalling.
- **Location Report**: The purpose of Location Report procedure is to provide the UE's current location to the MME. This procedure uses UE-associated signalling.

The start/stop trigger for location reporting for a UE is reported to the MME by the S-GW over the S11 interface. The Change Reporting Action (CRA) Information Element (IE) is used for this purpose. The MME updates the location to the S-GW using the User Location Information (ULI) IE.

The following S11 messages are used to transfer CRA and ULI information between the MME and S-GW:

- **Create Session Request**: The ULI IE is included for E-UTRAN Initial Attach and UE-requested PDN Connectivity procedures. It includes ECGI and TAI. The MME includes the ULI IE for TAU/ X2-Handover procedure if the P-GW has requested location information change reporting and the MME support location information change reporting. The S-GW includes the ULI IE on S5/S8 exchanges if it receives the ULI from the MME. If the MME supports change reporting, it sets the corresponding indication flag in the Create Session Request message.
- **Create Session Response**: The CRA IE in the Create Session Response message can be populated by the S-GW to indicate the type of reporting required.
- **Create Bearer Request**: The CRA IE is included with the appropriate Action field if the Location Change Reporting mechanism is to be started or stopped for the subscriber in the MME.
- **Modify Bearer Request**: The MME includes the ULI IE for TAU/Handover procedures and UE-initiated Service Request procedures if the P-GW has requested location information change reporting and the MME supports location information change reporting. The S-GW includes this IE on S5/S8 exchanges if it receives the ULI from the MME.
- **Modify Bearer Response**: The CRA IE is included with the appropriate Action field if the Location Change Reporting mechanism is to be started or stopped for the subscriber in the MME.
- **Delete Session Request**: The MME includes the ULI IE for the Detach procedure if the P-GW has requested location information change reporting and MME supports location information change reporting. The S-GW includes this IE on S5/S8 exchanges if it receives the ULI from the MME.
- **Update Bearer Request**: The CRA IE is included with the appropriate Action field if the Location Change Reporting mechanism is to be started or stopped for the subscriber in the MME.
- **Change Notification Request**: If no existing procedure is running for a UE, a Change Notification Request is sent upon receipt of an S1-AP location report message. If an existing procedure is running, one of the following messages reports the ULI:
    - Create Session Request
    - Create Bearer Response

- Modify Bearer Request
- Update Bearer Response
- Delete Bearer Response
- Delete Session Request

If an existing Change Notification Request is pending, it is aborted and a new one is sent.

Important: Information on configuring User Location Information (ULI) Reporting support is located in the *Configuring Optional Features on the MME* section of the *Mobility Management Entity Configuration* chapter in the *Mobility Management Entity Administration Guide*.

How the PDN Gateway Works

This section provides information on the function of the P-GW in an EPC E-UTRAN network and presents call procedure flows for different stages of session setup and disconnect.

The P-GW supports the following network flows:

- PMIPv6 PDN Gateway Call Session Procedures in an eHRPD Network
- GTP PDN Gateway Call Session Procedures in an LTE-SAE Network

PMIPv6 PDN Gateway Call/Session Procedures in an eHRPD Network

The following topics and procedure flows are included:

- Initial Attach with IPv6IPv4 Access
- PMIPv6 Lifetime Extension without Handover
- PDN Connection Release Initiated by UE
- PDN Connection Release Initiated by HSGW
- PDN Connection Release Initiated by P-GW

Initial Attach with IPv6/IPv4 Access

This section describes the procedure of initial attach and session establishment for a subscriber (UE).

Figure 214. Initial Attach with IPv6/IPv4 Access Call Flow

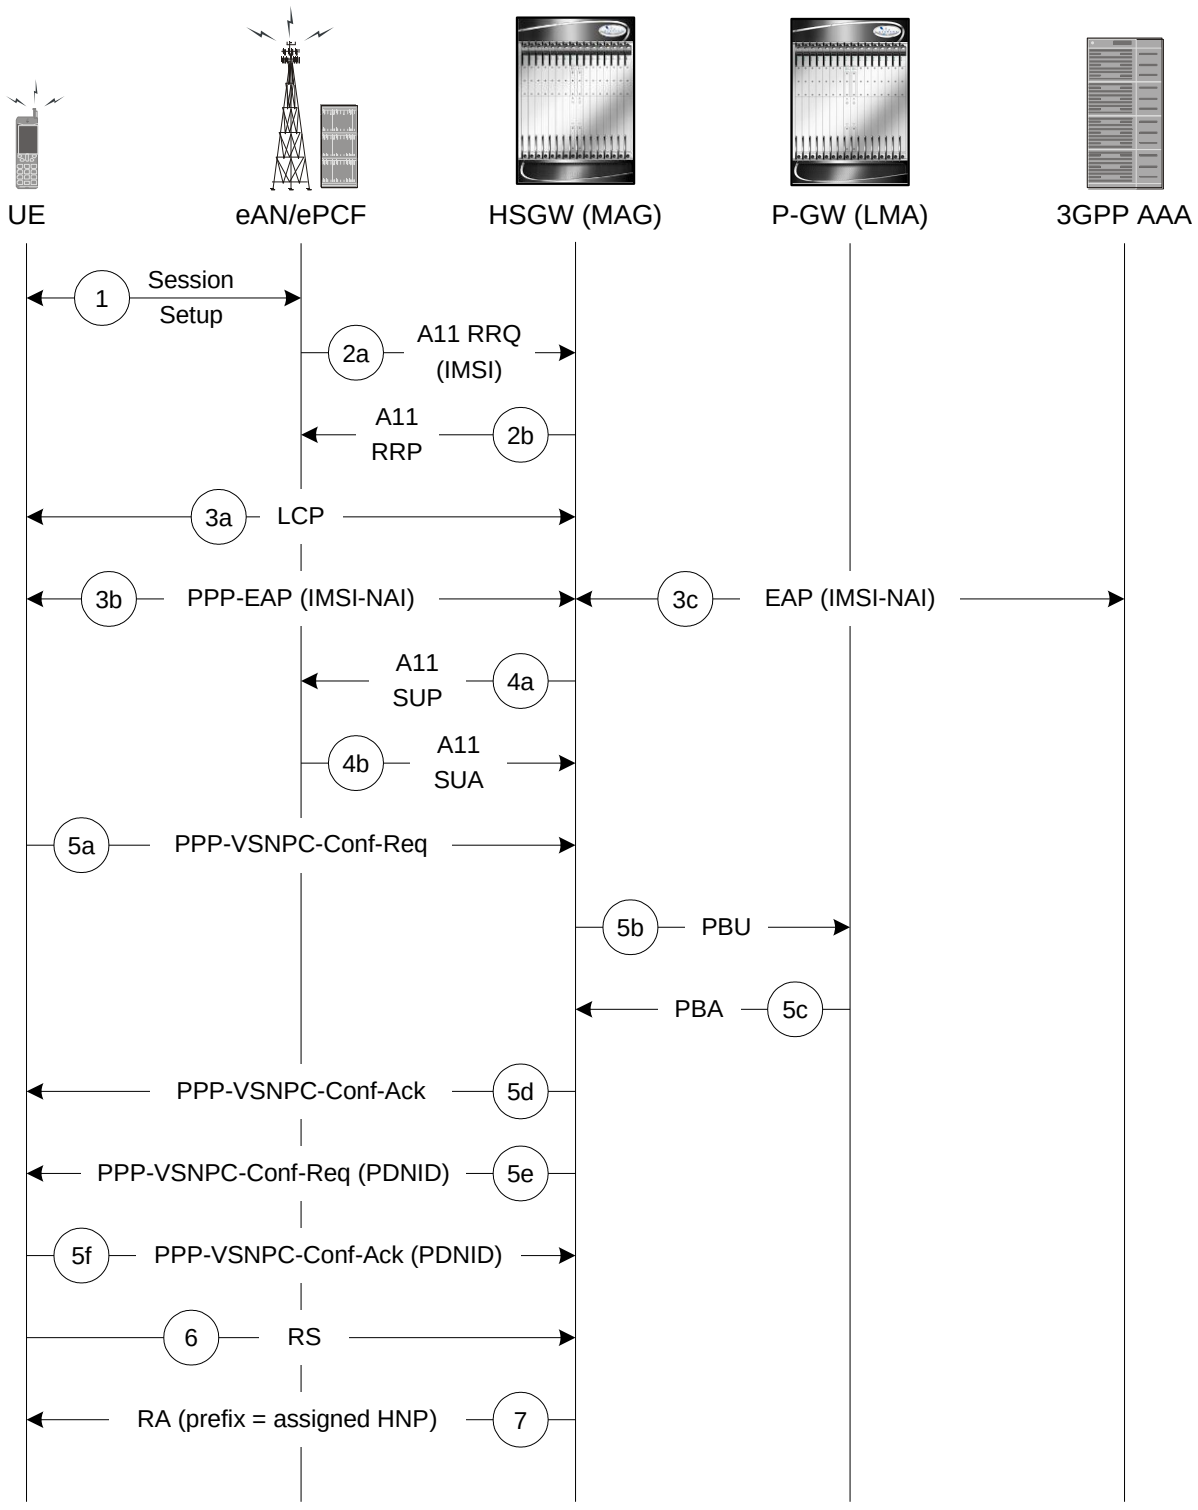


Table 96. Initial Attach with IPv6/IPv4 Access Call Flow Description

Step	Description
1	The subscriber (UE) attaches to the eHRPD network.
2a	The eAN/PCF sends an A11 RRQ to the HSGW. The eAN/PCF includes the true IMSI of the UE in the A11 RRQ.
2b	The HSGW establishes A10s and respond back to the eAN/PCF with an A11 RRP.
3a	The UE performs LCP negotiation with the HSGW over the established main A10.
3b	The UE performs EAP over PPP.
3c	EAP authentication is completed between the UE and the 3GPP AAA. During this transaction, the HSGW receives the subscriber profile from the AAA server.
4a	After receiving the subscriber profile, the HSGW sends the QoS profile in A11 Session Update Message to the eAN/PCF.
4b	The eAN/PCF responds with an A11 Session Update Acknowledgement (SUA).
5a	The UE initiates a PDN connection by sending a PPP-VSNCP-Conf-Req message to the HSGW. The message includes the PDNID of the PDN, APN, PDN-Type=IPv6/[IPv4], PDSN-Address and, optionally, PCO options the UE is expecting from the network.
5b	The HSGW sends a PBU to the P-GW.
5c	The P-GW processes the PBU from the HSGW, assigns an HNP for the connection and responds back to the HSGW with PBA.
5d	The HSGW responds to the VSNCP Conf Req with a VSNCP Conf Ack.
5e	The HSGW sends a PPP-VSNCP-Conf-Req to the UE to complete PPP VSNCP negotiation.
5f	The UE completes VSNCP negotiation by returning a PPP-VSNCP-Conf-Ack.
6	The UE optionally sends a Router Solicitation (RS) message.
7	The HSGW sends a Router Advertisement (RA) message with the assigned Prefix.

PMIPv6 Lifetime Extension without Handover

This section describes the procedure of a session registration lifetime extension by the P-GW without the occurrence of a handover.

Figure 215. PMIPv6 Lifetime Extension (without handover) Call Flow

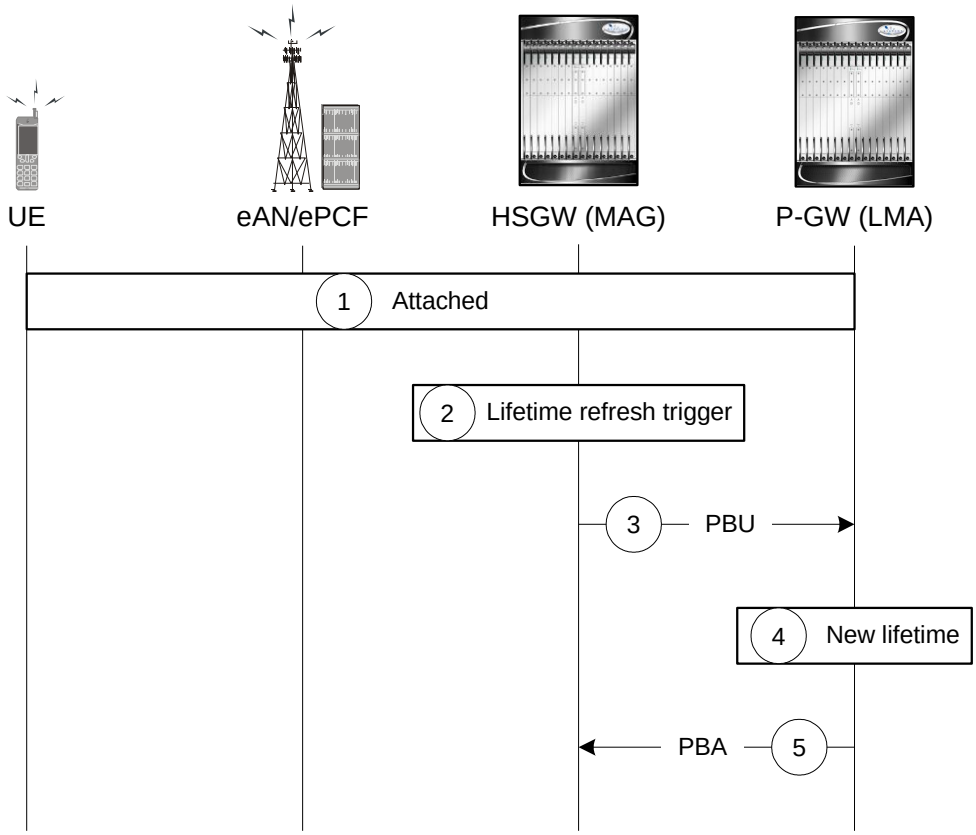


Table 97. PMIPv6 Lifetime Extension (without handover) Call Flow Description

Step	Description
1	The UE is attached to the EPC and has a PDN connection with the P-GW where PDNID=x and an APN with assigned HNP.
2	The HSGW MAG service registration lifetime nears expiration and triggers a renewal request for the LMA.
3	The MAG service sends a Proxy Binding Update (PBU) to the P-GW LMA service with the following attributes: Lifetime, MNID, APN, ATT=HRPD, HNP.
4	The P-GW LMA service updates the Binding Cache Entry (BCE) with the new granted lifetime.
5	The P-GW responds with a Proxy Binding Acknowledgement (PBA) with the following attributes: Lifetime, MNID, APN.

PDN Connection Release Initiated by UE

This section describes the procedure of a session release by the UE.

Figure 216. PDN Connection Release by the UE Call Flow

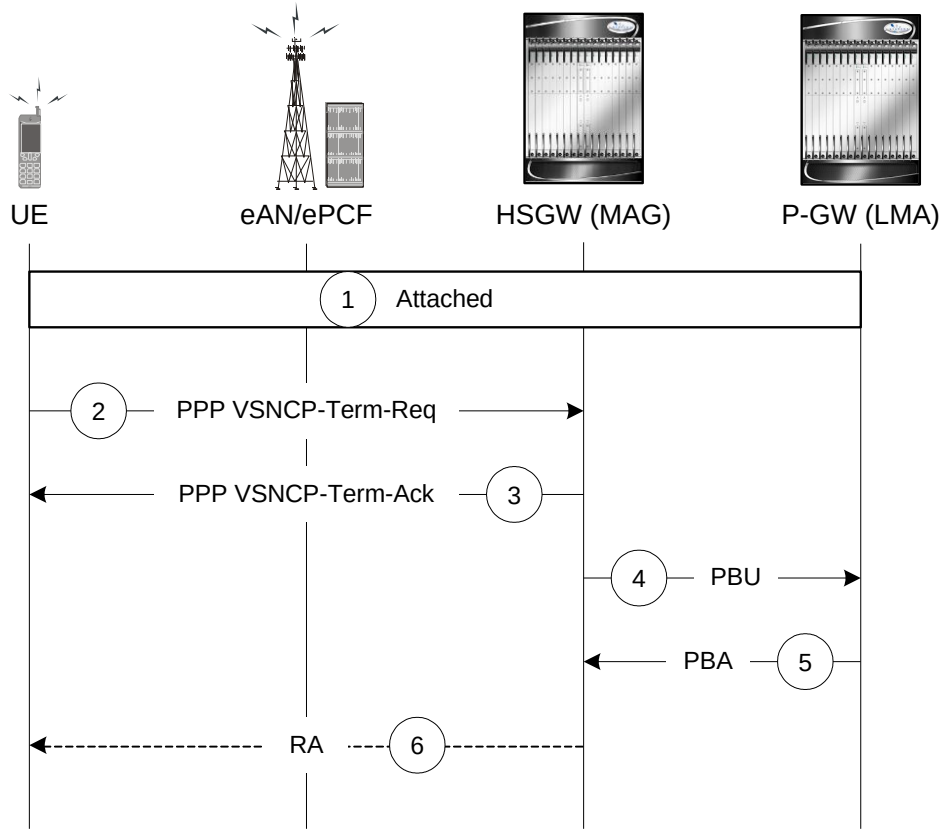


Table 98. PDN Connection Release by the UE Call Flow Description

Step	Description
1	The UE is attached to the EPC and has a PDN connection with the P-GW for PDN-ID=x and APN with assigned HNP.
2	The UE decides to disconnect from the PDN and sends a PPP VSNCP-Term-Req with PDNID=x.
3	The HSGW starts disconnecting the PDN connection and sends a PPP-VSNCP-Term-Ack to the UE (also with PDNID=x).
4	The HSGW begins the tear down of the PMIP session by sending a PBU Deregistration to the P-GW with the following attributes: Lifetime=0, MNID, APN, ATT=HRPD, HNP. The PBU Deregistration message should contain all the mobility options that were present in the initial PBU that created the binding.
5	The P-GW looks up the Binding Cache Entry (BCE) based on the HNP, deletes the binding, and responds to the HSGW with a Deregistration PBA with the same attributes (Lifetime=0, MNID, APN, ATT=HRPD, HNP).
6	The HSGW optionally sends a Router Advertisement (RA) with assigned HNP and prefix lifetime=0.

PDN Connection Release Initiated by HSGW

This section describes the procedure of a session release by the HSGW.

Figure 217. PDN Connection Release by the HSGW Call Flow

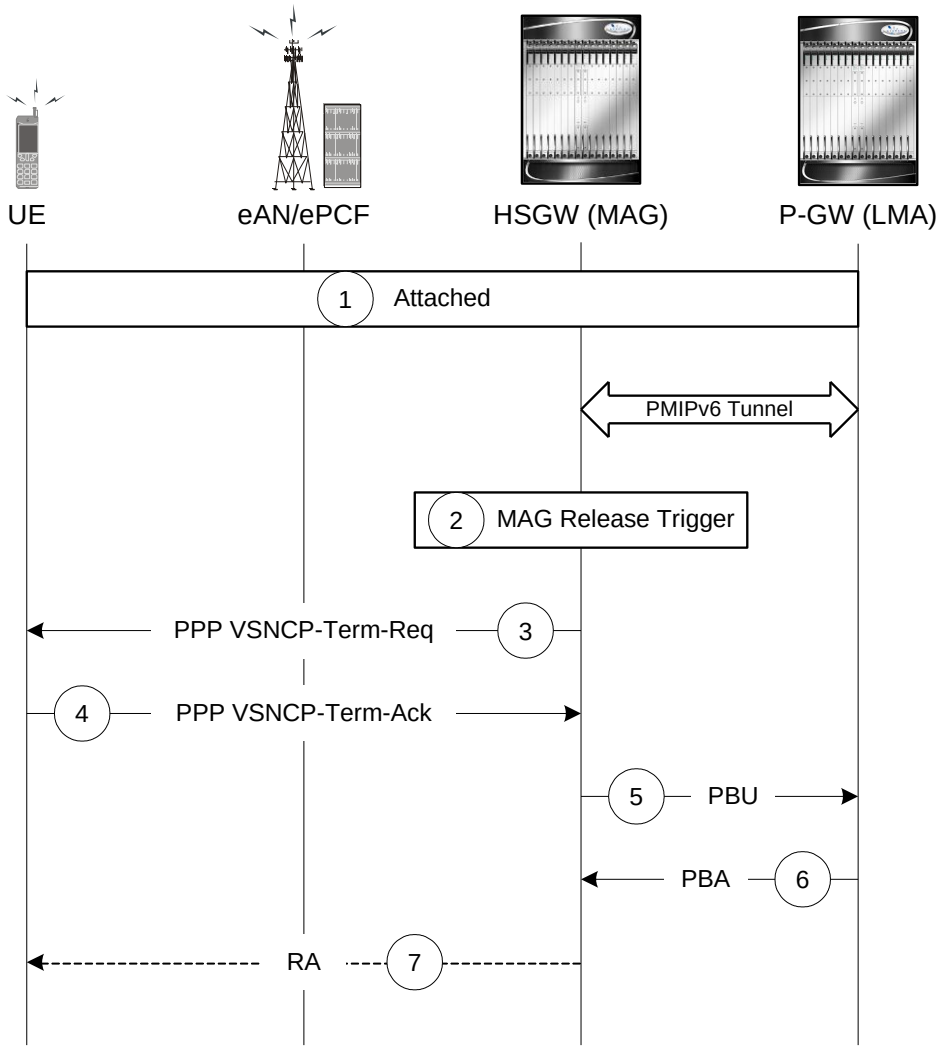


Table 99. PDN Connection Release by the HSGW Call Flow Description

Step	Description
1	The UE is attached to the EPC and has a PDN connection with the P-GW for PDN-ID=x and APN with assigned HNP.
2	The HSGW MAG service triggers a disconnect of the PDN connection for PDNID=x.
3	The HSGW sends a PPP VSNCP-Term-Req with PDNID=x to the UE.
4	The UE acknowledges the receipt of the request with a VSNCP-Term-Ack (PDNID=x).
5	The HSGW begins the tear down of the PMIP session by sending a PBU Deregistration to the P-GW with the following attributes: Lifetime=0, MNID, APN, HNP. The PBU Deregistration message should contain all the mobility options that were present in the initial PBU that created the binding.
6	The P-GW looks up the BCE based on the HNP, deletes the binding, and responds to the HSGW with a Deregistration PBA with the same attributes (Lifetime=0, MNID, APN, ATT=HRPD, HNP).

Step	Description
7	The HSGW optionally sends a Router Advertisement (RA) with assigned HNP and prefix lifetime=0.

PDN Connection Release Initiated by P-GW

This section describes the procedure of a session release by the P-GW.

Figure 218. PDN Connection Release by the P-GW Call Flow

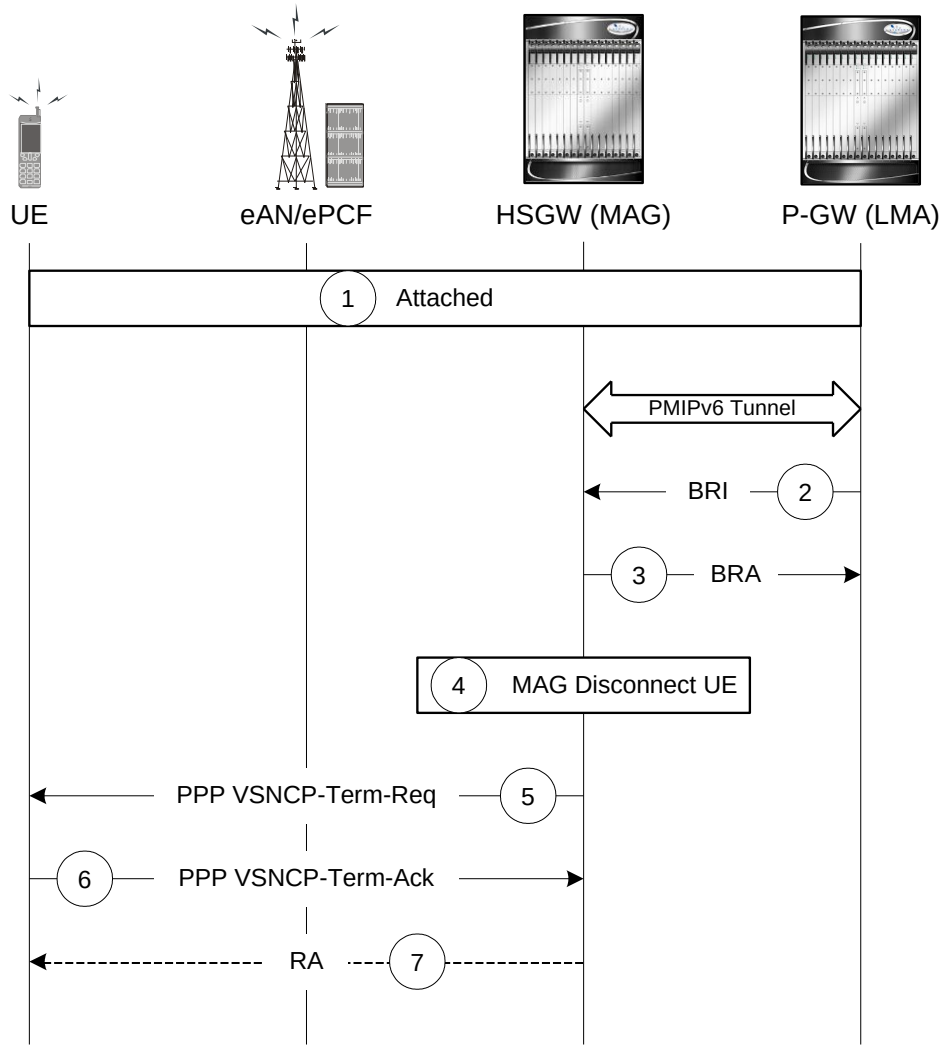


Table 100. PDN Connection Release by the P-GW Call Flow Description

Step	Description
1	The UE is attached to the EPC and has a PDN connection with the P-GW for PDN-ID=x and APN with assigned HNP.

Step	Description
2	A PGW trigger causes a disconnect of the PDN connection for PDNID=x and the PGW sends a Binding Revocation Indication (BRI) message to the HSGW with the following attributes: MNID, APN, HNP.
3	The HSGW responds to the BRI message with a Binding Revocation Acknowledgement (BRA) message with the sane attributes (MNID, APN, HNP).
4	The HSGW MAG service triggers a disconnect of the UE PDN connection for PDNID=x.
5	The HSGW sends a PPP VSNCP-Term-Req with PDNID=x to the UE.
6	The UE acknowledges the receipt of the request with a VSNCP-Term-Ack (PDNID=x).
7	The HSGW optionally sends a Router Advertisement (RA) with assigned HNP and prefix lifetime=0.

GTP PDN Gateway Call/Session Procedures in an LTE-SAE Network

The following topics and procedure flows are included:

- Subscriber-initiated Attach (initial)
- Subscriber-initiated Detach

Subscriber-initiated Attach (initial)

This section describes the procedure of an initial attach to the EPC network by a subscriber.

Figure 219. Subscriber-initiated Attach (initial) Call Flow

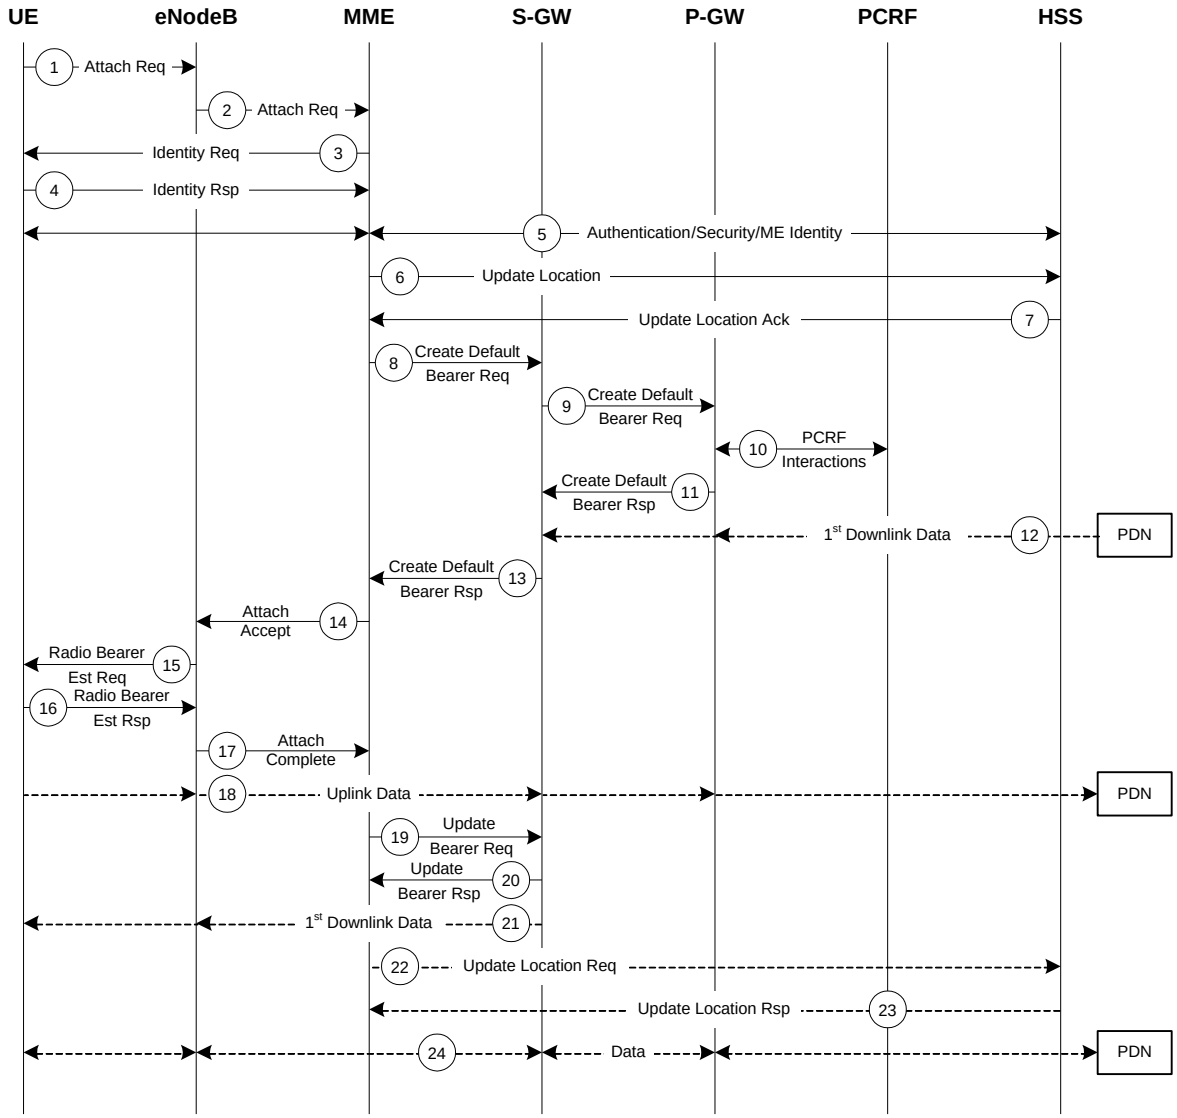


Table 101. Subscriber-initiated Attach (initial) Call Flow Description

Step	Description
1	The UE initiates the Attach procedure by the transmission of an Attach Request (IMSI or old GUTI, last visited TAI (if available), UE Network Capability, PDN Address Allocation, Protocol Configuration Options, Attach Type) message together with an indication of the Selected Network to the eNodeB. IMSI is included if the UE does not have a valid GUTI available. If the UE has a valid GUTI, it is included.
2	The eNodeB derives the MME from the GUTI and from the indicated Selected Network. If that MME is not associated with the eNodeB, the eNodeB selects an MME using an "MME selection function". The eNodeB forwards the Attach Request message to the new MME contained in a S1-MME control message (Initial UE message) together with the Selected Network and an indication of the E-UTRAN Area identity, a globally unique E-UTRAN ID of the cell from where it received the message to the new MME.

Step	Description
3	If the UE is unknown in the MME, the MME sends an Identity Request to the UE to request the IMSI.
4	The UE responds with Identity Response (IMSI).
5	If no UE context for the UE exists anywhere in the network, authentication is mandatory. Otherwise this step is optional. However, at least integrity checking is started and the ME Identity is retrieved from the UE at Initial Attach. The authentication functions, if performed this step, involves AKA authentication and establishment of a NAS level security association with the UE in order to protect further NAS protocol messages.
6	The MME sends an Update Location (MME Identity, IMSI, ME Identity) to the HSS.
7	The HSS acknowledges the Update Location message by sending an Update Location Ack to the MME. This message also contains the Insert Subscriber Data (IMSI, Subscription Data) Request. The Subscription Data contains the list of all APNs that the UE is permitted to access, an indication about which of those APNs is the Default APN, and the 'EPS subscribed QoS profile' for each permitted APN. If the Update Location is rejected by the HSS, the MME rejects the Attach Request from the UE with an appropriate cause.
8	The MME selects an S-GW using "Serving GW selection function" and allocates an EPS Bearer Identity for the Default Bearer associated with the UE. If the PDN subscription context contains no P-GW address the MME selects a P-GW as described in clause "PDN GW selection function". Then it sends a Create Default Bearer Request (IMSI, MME Context ID, APN, RAT type, Default Bearer QoS, PDN Address Allocation, AMBR, EPS Bearer Identity, Protocol Configuration Options, ME Identity, User Location Information) message to the selected S-GW.
9	The S-GW creates a new entry in its EPS Bearer table and sends a Create Default Bearer Request (IMSI, APN, S-GW Address for the user plane, S-GW TEID of the user plane, S-GW TEID of the control plane, RAT type, Default Bearer QoS, PDN Address Allocation, AMBR, EPS Bearer Identity, Protocol Configuration Options, ME Identity, User Location Information) message to the P-GW.
10	If dynamic PCC is deployed, the P-GW interacts with the PCRF to get the default PCC rules for the UE. The IMSI, UE IP address, User Location Information, RAT type, AMBR are provided to the PCRF by the P-GW if received by the previous message.
11	The P-GW returns a Create Default Bearer Response (P-GW Address for the user plane, P-GW TEID of the user plane, P-GW TEID of the control plane, PDN Address Information, EPS Bearer Identity, Protocol Configuration Options) message to the S-GW. PDN Address Information is included if the P-GW allocated a PDN address Based on PDN Address Allocation received in the Create Default Bearer Request. PDN Address Information contains an IPv4 address for IPv4 and/or an IPv6 prefix and an Interface Identifier for IPv6. The P-GW takes into account the UE IP version capability indicated in the PDN Address Allocation and the policies of operator when the P-GW allocates the PDN Address Information. Whether the IP address is negotiated by the UE after completion of the Attach procedure, this is indicated in the Create Default Bearer Response.
12	The Downlink (DL) Data can start flowing towards S-GW. The S-GW buffers the data.
13	The S-GW returns a Create Default Bearer Response (PDN Address Information, S-GW address for User Plane, S-GW TEID for User Plane, S-GW Context ID, EPS Bearer Identity, Protocol Configuration Options) message to the new MME. PDN Address Information is included if it was provided by the P-GW.
14	The new MME sends an Attach Accept (APN, GUTI, PDN Address Information, TAI List, EPS Bearer Identity, Session Management Configuration IE, Protocol Configuration Options) message to the eNodeB.
15	The eNodeB sends Radio Bearer Establishment Request including the EPS Radio Bearer Identity to the UE. The Attach Accept message is also sent along to the UE.
16	The UE sends the Radio Bearer Establishment Response to the eNodeB. In this message, the Attach Complete message (EPS Bearer Identity) is included.
17	The eNodeB forwards the Attach Complete (EPS Bearer Identity) message to the MME.

Step	Description
18	The Attach is complete and UE sends data over the default bearer. At this time the UE can send uplink packets towards the eNodeB which are then tunnelled to the S-GW and P-GW.
19	The MME sends an Update Bearer Request (eNodeB address, eNodeB TEID) message to the S-GW.
20	The S-GW acknowledges by sending Update Bearer Response (EPS Bearer Identity) message to the MME.
21	The S-GW sends its buffered downlink packets.
22	After the MME receives Update Bearer Response (EPS Bearer Identity) message, if an EPS bearer was established and the subscription data indicates that the user is allowed to perform handover to non-3GPP accesses, and if the MME selected a P-GW that is different from the P-GW address which was indicated by the HSS in the PDN subscription context, the MME sends an Update Location Request including the APN and P-GW address to the HSS for mobility with non-3GPP accesses.
23	The HSS stores the APN and P-GW address pair and sends an Update Location Response to the MME.
24	Bidirectional data is passed between the UE and PDN.

Subscriber-initiated Detach

This section describes the procedure of detachment from the EPC network by a subscriber.

Figure 220. Subscriber-initiated Detach Call Flow

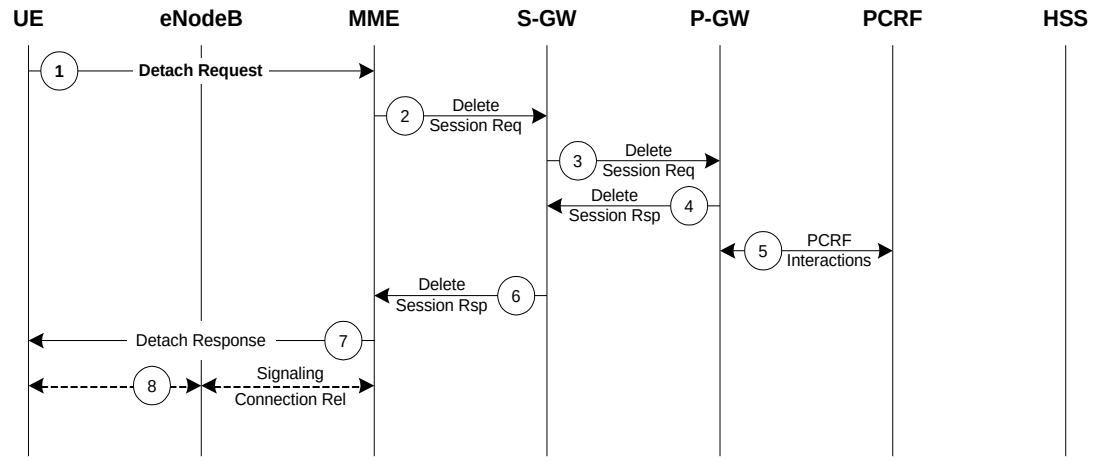


Table 102. Subscriber-initiated Detach Call Flow Description

Step	Description
1	The UE sends NAS message Detach Request (GUTI, Switch Off) to the MME. Switch Off indicates whether detach is due to a switch off situation or not.
2	The active EPS Bearers in the S-GW regarding this particular UE are deactivated by the MME sending a Delete Bearer Request (TEID) message to the S-GW.
3	The S-GW sends a Delete Bearer Request (TEID) message to the P-GW.

Step	Description
4	The P-GW acknowledges with a Delete Bearer Response (TEID) message.
5	The P-GW may interact with the PCRF to indicate to the PCRF that EPS Bearer is released if PCRF is applied in the network.
6	The S-GW acknowledges with a Delete Bearer Response (TEID) message.
7	If Switch Off indicates that the detach is not due to a switch off situation, the MME sends a Detach Accept message to the UE.
8	The MME releases the S1-MME signalling connection for the UE by sending an S1 Release command to the eNodeB with Cause = Detach.

Supported Standards

The P-GW service complies with the following standards.

- Release 8 3GPP References
- 3GPP2 References
- IETF References
- Object Management Group (OMG) Standards

Release 8 3GPP References



Important: The P-GW currently supports the following Release 8 3GPP specifications. Most 3GPP specifications are also used for 3GPP2 support; any specifications that are unique to 3GPP2 are listed under *3GPP2 References*.

- 3GPP TR 21.905: Vocabulary for 3GPP Specifications
- 3GPP TS 23.003: Numbering, addressing and identification
- 3GPP TS 23.007: Restoration procedures
- 3GPP TS 23.060. General Packet Radio Service (GPRS); Service description; Stage 2
- 3GPP TS 23.107: Quality of Service (QoS) concept and architecture
- 3GPP TS 23.203: Policy and charging control architecture
- 3GPP TS 23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
- 3GPP TS 23.402. Architecture enhancements for non-3GPP accesses
- 3GPP TS 24.008: Mobile radio interface Layer 3 specification; Core network protocols
- 3GPP TS 24.229: IP Multimedia Call Control Protocol based on SIP and SDP; Stage 3
- 3GPP TS 27.060: Mobile Station (MS) supporting Packet Switched Services
- 3GPP TS 29.061: Interworking between the Public Land Mobile Network (PLMN) supporting packet based services and Packet Data Networks (PDN)
- 3GPP TS 29.210. Charging rule provisioning over Gx interface
- 3GPP TS 29.212: Policy and Charging Control over Gx reference point
- 3GPP TS 29.213: Policy and Charging Control signaling flows and QoS
- 3GPP TS 29.274: Evolved GPRS Tunnelling Protocol for Control plane (GTPv2-C)
- 3GPP TS 29.275: Proxy Mobile IPv6 (PMIPv6) based Mobility and Tunnelling protocols
- 3GPP TS 29.281: GPRS Tunnelling Protocol User Plane (GTPv1-U)
- 3GPP TS 32.295: Charging management; Charging Data Record (CDR) transfer
- 3GPP TS 32.298: Telecommunication management; Charging management; Charging Data Record (CDR) encoding rules description

- 3GPP TS 32.299: Charging management; Diameter charging applications
- 3GPP TS 36.300. EUTRA and EUTRAN; Overall description Stage 2
- 3GPP TS 36.412. EUTRAN S1 signaling transport
- 3GPP TS 36.413. EUTRAN S1 Application Protocol (S1AP)

3GPP2 References

- X.S0057-0 v3.0 E-UTRAN - eHRPD Connectivity and Interworking: Core Network Aspects

IETF References

- RFC 768: User Datagram Protocol (STD 6).
- RFC 791: Internet Protocol (STD 5).
- RFC 1701, Generic Routing Encapsulation (GRE)
- RFC 1702, Generic Routing Encapsulation over IPv4 networks
- RFC 2131: Dynamic Host Configuration Protocol
- RFC 2460: Internet Protocol, Version 6 (IPv6) Specification
- RFC 2698: A Two Rate Three Color Marker
- RFC 2784: Generic Routing Encapsulation (GRE)
- RFC 2890: Key and Sequence Number Extensions to GRE
- RFC 3319: Dynamic Host Configuration Protocol (DHCPv6) Options for Session Initiation Protocol (SIP) Servers
- RFC 3588: Diameter Base Protocol
- RFC 3646: DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
- RFC 3775: Mobility Support in IPv6
- RFC 4006: Diameter Credit-Control Application
- RFC 4282: The Network Access Identifier
- RFC 4283: Mobile Node Identifier Option for Mobile IPv6 (MIPv6)
- RFC 4861: Neighbor Discovery for IP Version 6 (IPv6)
- RFC 4862: IPv6 Stateless Address Autoconfiguration
- RFC 5094: Mobile IPv6 Vendor Specific Option
- RFC 5149: Service Selection for Mobile IPv6
- RFC 5213: Proxy Mobile IPv6
- Internet-Draft (draft-ietf-netlmm-grekey-option-01.txt): GRE Key Option for Proxy Mobile IPv6, work in progress
- Internet-Draft (draft-ietf-netlmm-pmip6-ipv4-support-02.txt) IPv4 Support for Proxy Mobile IPv6
- Internet-Draft (draft-ietf-netlmm-proxymip6-07.txt): Proxy Mobile IPv6

- Internet-Draft (draft-ietf-mext-binding-revocation-02.txt): Binding Revocation for IPv6 Mobility, work in progress
- Internet-Draft (draft-meghana-netlmm-pmipv6-mipv4-00.txt) Proxy Mobile IPv6 and Mobile IPv4 interworking

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

Chapter 28
Personal Stateful Firewall Overview

This chapter provides an overview of the Personal Stateful Firewall In-line Service.

This chapter covers the following topics:

- Firewall Overview
- Supported Features
- How Personal Stateful Firewall Works
- Understanding Firewall Rules with Stateful Inspection

Firewall Overview

The Personal Stateful Firewall is an in-line service feature that inspects subscriber traffic and performs IP session-based access control of individual subscriber sessions to protect the subscribers from malicious security attacks.

The Personal Stateful Firewall in-line service works in conjuction with the following products:

- GGSN
- HA
- IPSG
- PDSN
- P-GW

The Personal Stateful Firewall supports stateless and stateful inspection and filtering based on the configuration.

In stateless inspection, the firewall inspects a packet to determine the 5-tuple—source and destination IP addresses and ports, and protocol—information contained in the packet. This static information is then compared against configurable rules to determine whether to allow or drop the packet. In stateless inspection the firewall examines each packet individually, it is unaware of the packets that have passed through before it, and has no way of knowing if any given packet is part of an existing connection, is trying to establish a new connection, or is a rogue packet.

In stateful inspection, the firewall not only inspects packets up through the application layer / layer 7 determining a packet's header information and data content, but also monitors and keeps track of the connection's state. For all active connections traversing the firewall, the state information, which may include IP addresses and ports involved, the sequence numbers and acknowledgement numbers of the packets traversing the connection, TCP packet flags, etc. is maintained in a state table. Filtering decisions are based not only on rules but also on the connection state established by prior packets on that connection. This enables to prevent a variety of DoS, DDoS, and other security violations. Once a connection is torn down, or is timed out, its entry in the state table is discarded. For more information see the Connection State and State Table in Personal Stateful Firewall section.

The Enhanced Charging Service (ECS) / Active Charging Service (ACS) in-line service is the primary vehicle that performs packet inspection and charging. For more information on ECS, see the *Enhanced Charging Service Administration Guide*.

Platform Requirements

The Personal Stateful Firewall in-line service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the Installation Guide for the chassis and/or contact your Cisco account representative.

License Requirements

The Personal Stateful Firewall is a licensed Cisco feature. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Supported Features

The Personal Stateful Firewall supports the following features:

- Protection against DoS Attacks
- Application-level Gateway (ALG) Support
- Stateful Packet Filtering and Inspection Support
- Stateless Packet Filtering and Inspection Support
- Host Pool, IMSI Pool, and Port Map Support
- Flow Recovery Support
- SNMP Thresholding Support
- Logging Support

Protection against Denial-of-Service Attacks

Denial-of-Service (DoS) and Distributed Denial-of-Service (DDoS) attacks can deprive network resources/services unavailable to its intended users.

DoS attacks can result in:

- A host consuming excessive resources — memory, disk space, CPU time, etc. — eventually leading to a system crash or providing very sluggish response.
- Flooding of the network to the extent that no valid traffic is able to reach the intended destination.
- Confusing target TCP/IP stack on destination hosts by sending crafted, malformed packets eventually resulting in system crash.

  In this release, malformity check is enhanced for IPv6 and ICMPv6 packets. Port-scan and Flooding attacks are also enhanced to support IPv6. Protection against other L4 attacks are similar to IPv4. The Attacking server feature is also enhanced to store IPv6 servers.

DoS attacks can destroy data in affected mobile nodes. Stateful Firewall is designed to defend subscribers and prevent the abuse of network bandwidth from DoS attacks originating from both the Internet and the internal network.

Types of Denial-of-Service Attacks

Personal Stateful Firewall can detect the following DoS attacks.

The DoS attacks are listed based on the protocol layer that they work on.

- IP-based Attacks:
  - Land attacks
  - Jolt attacks
  - Teardrop attacks — Detected only in downlink direction, i.e. traffic coming from the external network towards the mobile subscribers
  - Invalid IP option length

- IP-unaligned-timestamp attack — Detected only in downlink direction
- Short IP header length
- IP checksum errors
- IP reassembly failure (downlink)
- IP reassembly failure (uplink)
- Source router — Detected only in downlink direction
- IPv6 header checks

- TCP-based Attacks:
  - Data packets received after RST/FIN
  - Invalid SEQ number received with RST
  - Data without connection established
  - Invalid TCP connection requests
  - Invalid TCP pre-connection requests
  - Invalid ACK value (cookie enabled)
  - Invalid TCP packet length
  - Short TCP header length
  - TCP checksum errors
  - SEQ/ACK out-of-range
  - TCP null scan attacks
  - Post connection SYN
  - No TCP flags set
  - All TCP flags set
  - Invalid TCP packets
  - Flows closed by RST before 3-Way handshake
  - Flows timed-out in SYN_RCVD1 state
  - Flows timed-out in SYN_RCVD2 state
  - TCP-SYN flood attacks — Detected only in downlink direction
  - FTP bounce attack — Detected only in downlink direction
  - MIME flood attacks — Detected only in downlink direction
  - Exceeding reset message threshold
  - Source port zero
  - WinNuke attack — Detected only in downlink direction
  - TCP-window-containment — Detected only in downlink direction
- UDP-based Attacks:
  - Invalid UDP echo response
  - Invalid UDP packet length
  - UDP checksum errors

        - Short UDP header length
        - UDP flood attack — Detected only in downlink direction
    - ICMP-based Attacks:
        - Invalid ICMP response
        - ICMP reply error
        - Invalid ICMP type packet
        - ICMP error message replay attacks
        - ICMP packets with duplicate sequence number
        - Short ICMP header length
        - Invalid ICMP packet length
        - ICMP flood attack — Detected only in downlink direction
        - Ping of death attacks
        - ICMP checksum errors
        - ICMP packets with destination unreachable message
        - ICMP echo packets with ID zero
    - Other DoS Attacks:
        - Port-scan attacks — Detected only in downlink direction

Various header integrity checks are performed for IPv6 to ensure the integrity of an IPv6 packet. IPv6 packets with unknown extension headers will not be dropped by Firewall; such packets will be allowed by Firewall. Firewall performs the following header checks:

- Limiting extension headers
- Hop-by-hop Options filtering
- Destination Options filtering
- Router Header filtering
- Fragment Header filtering

Protection against Port Scanning

Port scanning is a technique used to determine the states of TCP/UDP ports on a network host, and to map out hosts on a network. Essentially, a port scan consists of sending a message to each port on the host, one at a time. The kind of response received indicates whether the port is used, and can therefore be probed further for weakness. This way hackers find potential weaknesses that can be exploited.

Stateful Firewall provides protection against port scanning by implementing port scan detection algorithms. Port-scan attacks are only detected in the downlink direction—traffic from external network towards mobile subscribers.

Application-level Gateway Support

A stateful firewall while ensuring that only legitimate connections are allowed, also maintains the state of an allowed connection. Some network applications require additional connections to be opened up in either direction and information regarding such connections is sent in the application payload. For these applications to work properly, a

stateful firewall must inspect, analyze, and parse these application payloads to get the additional connection information, and open partial connections/pinholes in the firewall to allow the connections.

To parse application payloads, firewall employs ALGs. ALGs also check for application-level attacks. Stateful Firewall provides ALG functionality for the following protocols. ALG support for Simple Mail Transfer Protocol (SMTP) and HTTP is ECS functionality.

ALG Type	Support in IPv4 Firewall	Support in IPv6 Firewall
File Transfer Protocol (FTP)	Yes	Yes
H323	Yes	No
Point-to-Point Tunneling Protocol (PPTP)	Yes	Yes
Real Time Streaming Protocol (RTSP)	Yes	Yes
Session Initiation Protocol (SIP TCP)	Yes	Yes
Session Initiation Protocol (SIP UDP)	Yes	Yes
Trivial File Transfer Protocol (TFTP)	Yes	Yes

PPTP ALG Support

PPTP exchanges IP or port specific information over its control connection and that information will be used to transfer the data over tunnel. If a PPTP client resides behind NAT and uses private IP to communicate with the outside world, it is possible that the information exchange over PPTP control flow consists of private IPs. So NAT translates the private IP specific information to public IP (NATed IP) for good communication. To achieve this, PPTP ALG is supported.

To establish a GRE session, PPTP exchanges call IDs from both peers to form a unique triple value, that is, client IP, server IP and Call ID. For Many-to-One NAT, PPTP analyzer is implemented to analyze the PPTP Control Flow traffic. It can be configured to send all the PPTP Control Flow packets to PPTP analyzer. PPTP analyzer analyzes the packet and allocates a new unique Call ID. Packet payload will be modified for the new Call ID and the binding between the two Call IDs will be maintained. Similarly, the PPTP first packet will be NAT-ed, Call ID translated and sent to the PPTP Server. This Call ID translation happens for all the downlink packets after the first packet. For GRE Data Tunnel Flow translation, it can be configured to send all the GRE downlink packets to PPTP analyzer. PPTP analyzer then analyzes the GRE header and translates the GRE Call ID if a Call ID binding exists.

Important: In 9.0 and earlier releases, GRE flow was detected and identified as a 3-tuple. In later releases, GRE flow is identified as a 5-tuple as part of the PPTP analyzer implementation. For the PPTP traffic to work after upgrade, a PPTP analyzer need to be configured in the rulebase or add a downlink allow rule for GRE traffic in Firewall-and-NAT policy.

TFTP ALG Support

Trivial File Transfer Protocol (TFTP) ALG enables Firewall or NAT enabled users to seamlessly use applications using TFTP Protocol. TFTP ALG feature analyzes the TFTP packets and selectively allows the downlink data flow by creating pin holes. This feature also ensures NAT/PAT IP/Port translation for NAT enabled users.

TFTP ALG analyzes the packets for basic TFTP signatures. A TFTP analyzer is implemented for this purpose. A routing rule is created for routing the packets to TFTP analyzer. Potential TFTP packets are parsed and information like

query type and mode are stored. After confirming that the packet is TFTP, a dynamic route is created for MS IP, MS Port, Server IP and Protocol. When the data flow starts, dynamic route is matched and data is sent to the TFTP analyzer. For NAT enabled calls, same Client port used for the control connection will be used for Data flow.

Stateful Packet Inspection and Filtering Support

As described in the Overview section, stateful packet inspection and filtering uses Layer-4 information as well as the application-level commands up to Layer-7 to provide good definition of the individual connection states to defend from malicious security attacks.

Personal Stateful Firewall overcomes the disadvantages of static packet filters by disallowing any incoming packets that have the TCP SYN flag set (which means a host is trying to initiate a new connection). If configured, stateful packet filtering allows only packets for new connections initiated from internal hosts to external hosts and disallows packets for new connections initiated from external hosts to internal hosts.

TCP stateful processing is enhanced for processing IPv6 packets. The functionality is similar to IPv4 packets.

Stateless Packet Inspection and Filtering Support

Stateful Firewall service can be configured for stateless processing. In stateless processing, packets are inspected and processed individually.

Stateless processing is only applicable for TCP and ICMP protocols. By nature UDP is a stateless protocol without any kind of acking or request and reply mechanism at transport level.

When TCP FSM is disabled, flows can start with any kind of packet and need not respect the TCP FSM. Such flows are marked as dummy (equivalent to flows established during flow recovery timer running). For these flows only packet header check is done; there will be no FSM checks, sequence number validations, or port scan checks done.

When ICMP FSM is disabled, ICMP reply without corresponding requests, ICMP error message without inner packet data session, and duplicate ICMP requests are allowed by firewall.

Host Pool, IMSI Pool, and Port Map Support

This section describes the Host Pool, IMSI Pool, and Port Map features that can be used while configuring access ruledefs.

Host Pool Support

Host pools allow operators to group a set of host or IP addresses that share similar characteristics together. Access rule definitions (ruledefs) can be configured with host pools. Up to 10 sets of IP addresses can be configured in each host pool. Host pools are configured in the ACS Host Pool Configuration Mode.

Host pools are enhanced to support IPv6 addresses and address ranges. It can also be a combination of IPv4 and IPv6 addresses.

IMSI Pool Support

IMSI pools allow the operator to group a set of International Mobile Station Identifier (IMSI) numbers together. Up to 10 sets of IMSI numbers can be configured in each IMSI pool. IMSI pools are configured in the ACS IMSI Pool Configuration Mode.

Port Map Support

Port maps allow the operator to group a set of port numbers together. Access ruledefs can be configured with port maps. Up to 10 sets of ports can be configured in each port map. Port maps are configured in the ACS Port Map Configuration Mode.

The Personal Stateful Firewall uses standard application ports to trigger ALG functionality. The operator can modify the existing set to remove/add new port numbers.

Flow Recovery Support

Stateful Firewall supports call recovery during session failover. Flows associated with the calls are recovered.

A recovery-timeout parameter is configurable for uplink and downlink directions. If the value is set to zero, firewall flow recovery is disabled. If the value is non-zero, then firewall will be bypassed for packets from MS/Internet until the time configured (uplink/downlink). Once the manager recovers, the recovery-timeout timer is started. During this time:

- If any ongoing traffic arrives from the subscriber and no association is found, and flow recovery is enabled, basic checks like header processing, attacks, etc. are done (stateful checks of packet is not done), and if all is okay, an association is created and the packet is allowed to pass through.
- If any ongoing traffic arrives from the Internet to MS and no association is found, and flow recovery is not enabled, it is dropped. No RESET is sent. Else, basic checks like header processing, flooding attack check are done (stateful checks are not done), and if all is okay, an association is created and the packet is allowed to pass through.
- In case flow recovered from ongoing traffic arrives from Internet to MS, and MS sends a NACK, the Unwanted Traffic Suppression feature is triggered, i.e. upon repeatedly receiving NACK from MS for a 5-tuple, further traffic to the 5-tuple is blocked for some duration and not sent to MS.
- If any new traffic (3-way handshake) comes, whether it is a new flow or a new flow due to pin-hole, based on the direction of packet and flow-recovery is enabled, basic checks like header processing, attacks, etc. are done (stateful checks are not done) and if all is okay, an association is created and the packet is allowed to pass through.

For any traffic coming after the recovery-timeout:

- If any ongoing traffic arrives, it is allowed only if an association was created earlier. Else, it is dropped and reset is sent.
- If any new traffic (3-way handshake) arrives, the usual Stateful Firewall processing is done.

If recovery-timeout value is set to zero, Stateful Firewall flow recovery is not done.

Stateful Firewall now supports IPv6 flows recovery similar to IPv4 flows.

SNMP Thresholding Support

Personal Stateful Firewall allows to configure thresholds to receive notifications for various events that are happening in the system. Whenever a measured value crosses the specified threshold value at the given time, an alarm is generated. And, whenever a measured value falls below the specified threshold clear value at the given time, a clear alarm is generated. The following events are supported for generating and clearing alarms:

- Dos-Attacks: When the number of DoS attacks crosses a given value, a threshold is raised, and it is cleared when the number of DoS attacks falls below a value in a given period of time.

- Drop-Packets: When the number of dropped packets crosses a given value, a threshold is raised, and it is cleared when the number of dropped packets falls below a value in a given period of time.
- Deny-Rule: When the number of Deny Rules cross a given value, a threshold is raised, and it is cleared when the number of Deny Rules falls below a value in a given period of time.
- No-Rule: When the number of No Rules cross a given value, a threshold is raised, and it is cleared when the number of No Rules falls below a value in a given period of time.

Logging Support

Stateful Firewall supports logging of various messages on screen if logging is enabled for firewall. These logs provide detailed messages at various levels, like critical, error, warning, and debug. All the logs displaying IP addresses are enhanced to display IPv6 addresses.

Logging is also supported at rule level, when enabled through rule a message will be logging whenever a packet hits the rule. This can be turned on/off in a rule.

These logs are also sent to a syslog server if configured in the system.

How Personal Stateful Firewall Works

This section describes how Personal Stateful Firewall works.

Important: In release 8.x, Stateful Firewall for CDMA and early UMTS releases used rulebase-based configurations, whereas later UMTS releases used policy-based configurations. In release 9.0, Stateful Firewall for UMTS and CDMA releases, both use policy-based configurations. For more information, please contact your local service representative.

Firewall-and-NAT policies are configured in the Firewall-and-NAT Policy Configuration Mode. Each policy contains a set of access ruledefs and the firewall configurations. Multiple such policies can be configured, however, only one policy is applied to a subscriber at any point of time.

The policy used for a subscriber can be changed either from the CLI, or by dynamic update of policy name in Diameter and RADIUS messages.

The Firewall-and-NAT policy to be used for a subscriber can be configured in:

- ACS Rulebase: The default Firewall-and-NAT policy configured in the ACS rulebase has the least priority. If there is no policy configured in the APN/subscriber template, and/or no policy to use is received from the AAA/OCS, only then the default policy configured in the ACS rulebase is used.
- APN/Subscriber Template: The Firewall-and-NAT policy configured in the APN/subscriber template overrides the default policy configured in the ACS rulebase. To use the default policy configured in the ACS rulebase, in the APN/subscriber configuration, the command to use the default rulebase policy must be configured.
- AAA/OCS: The Firewall-and-NAT policy to be used can come from the AAA server or the OCS. If the policy comes from the AAA/OCS, it will override the policy configured in the APN/subscriber template and/or the ACS rulebase.

Important: The Firewall-and-NAT policy received from the AAA and OCS have the same priority. Whichever comes latest, either from AAA/OCS, is applied.

The Firewall-and-NAT policy to use can be received from RADIUS during authentication.

Disabling Firewall Policy

Important: By default, Stateful Firewall processing for subscribers is disabled.

Stateful Firewall processing is disabled for subscribers in the following cases:

- If Stateful Firewall is explicitly disabled in the APN/subscriber template configuration.
- If the AAA/OCS sends the SN-Firewall-Policy AVP with the string “disable”, the locally configured firewall policy does not get applied.
- If the SN-Firewall-Policy AVP is received with the string “NULL”, the existing policy will continue.
- If the SN-Firewall-Policy AVP is received with a name that is not configured locally, the subscriber session is terminated.

Mid-session Firewall Policy Update

The Firewall-and-NAT policy can be updated mid-session provided firewall policy was enabled during call setup.

Important: When the SN-Firewall-Policy AVP contains "disable" during mid-session firewall policy change, there will be no action taken as the Firewall-and-NAT policy cannot be disabled dynamically. The policy currently applied will continue.

Important: When a Firewall-and-NAT policy is deleted, for all subscribers using the policy, Firewall processing is disabled, also ECS sessions for the subscribers are dropped. In case of session recovery, the calls are recovered but with Stateful Firewall disabled.

How it Works

The following figures illustrate packet flow in Stateful Firewall processing for a subscriber.

Figure 221. Stateful Firewall Processing

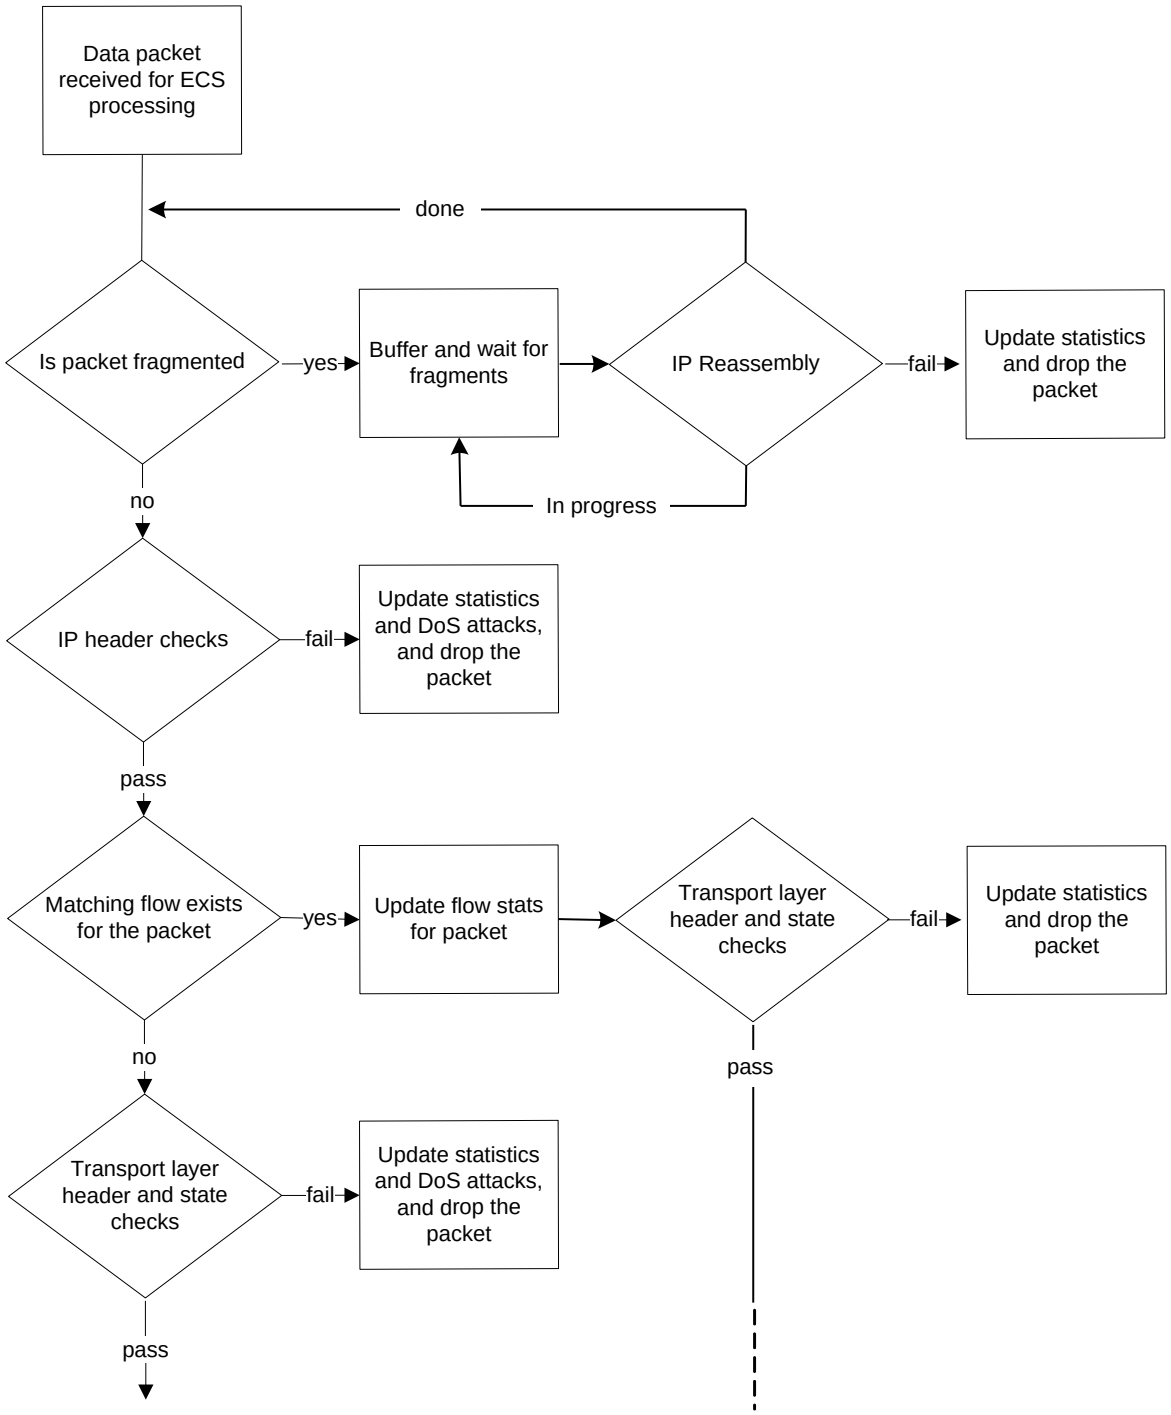


Figure 222. Continued... Stateful Firewall Processing

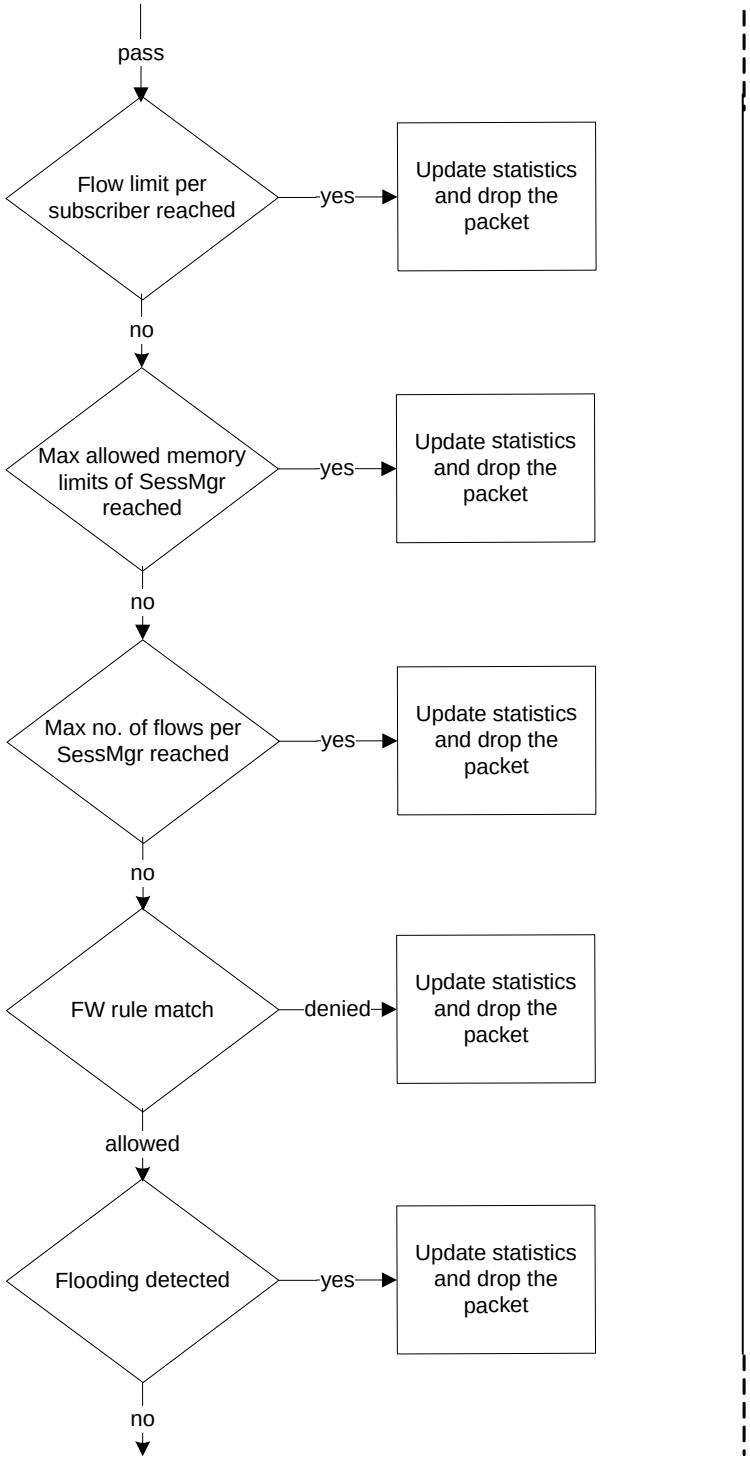
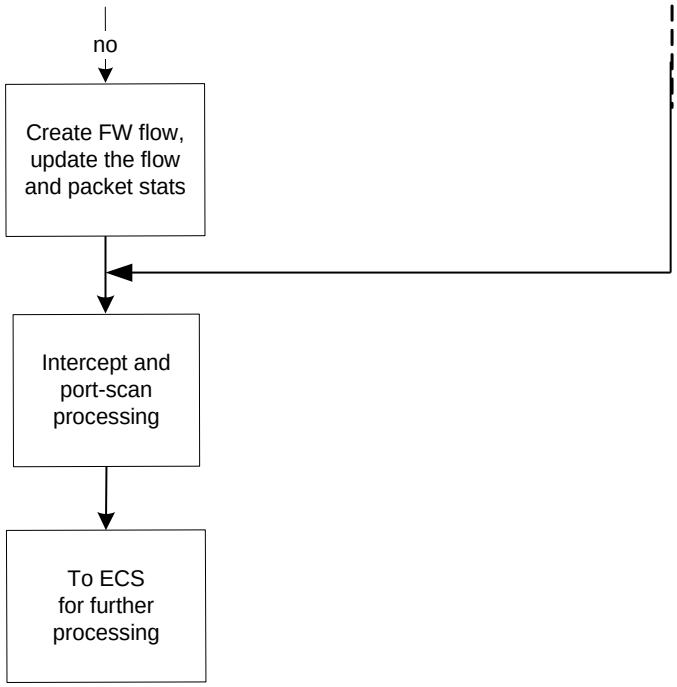


Figure 223. Continued... Stateful Firewall Processing



Understanding Rules with Stateful Inspection

This section describes terms used in the Personal Stateful Firewall context.

- **Access Ruledefs**: The Personal Stateful Firewall's stateful packet inspection feature allows operators to configure rule definitions (ruledefs) that take active session information into consideration to permit or deny incoming or outgoing packets.

  An access ruledef contains the criteria for multiple actions that could be taken on packets matching the rules. These rules specify the protocols, source and destination hosts, source and destination ports, direction of traffic parameters for a subscriber session to allow or reject the traffic flow.

  An access ruledef consists of the following fields:

  - Ruledef name
  - Source IP address
  - Source port number — not required if the protocol is other than TCP or UDP
  - Destination IP address
  - Destination port number — not required if the protocol is other than TCP or UDP
  - Transport protocol (TCP/UDP/ICMP/ICMPv6/AH/ESP)
  - Direction of connection (Uplink/Downlink)
  - Bearer (IMSI-pool and APN)
  - Logging action (enable/disable)
  - IP version - IPv4 or IPv6

  An access ruledef can be added to multiple Firewall-and-NAT policies.

  A combined maximum of 4096 rules (host pools + IMSI pools + port maps + charging ruledefs + firewall/access ruledefs + routing ruledefs) can be created in a system. Access ruledefs are different from ACS ruledefs.

  Firewall access ruledefs are enhanced to support IPv6 addresses and parameters like IP version and ICMPv6 protocol. The existing rule lines "ip src-address" and "ip dst-address" are capable of accepting both IPv4 and IPv6 addresses and there is no CLI change for them.
- **Firewall-and-NAT Policy**: Firewall policies can be created for individual subscribers, domains, or all callers within a referenced context. Each policy contains a set of access ruledefs with priorities defined for each rule and the firewall configurations. Firewall-and-NAT policies are configured in the Firewall-and-NAT Policy Configuration Mode.
- **Service Definition**: User-defined firewall service for defining Stateful Firewall policy for initiating an outgoing connection on a primary port and allowing opening of auxiliary ports for that association in the reverse direction.
- **Maximum Association**: The maximum number of Stateful Firewall associations for a subscriber.

Connection State and State Table in Personal Stateful Firewall

This section describes the state table and different connection states for transport and network protocols.

After packet inspection, the Personal Stateful Firewall stores session state and other information into a table. This state table contains entries of all the communication sessions of which the firewall subsystem is aware of. Every entry in this

table holds a list of information that identifies the subscriber session it represents. Generally this information includes the source and destination IP address, flags, sequence, acknowledgement numbers, etc.

When a connection is permitted through the Personal Stateful Firewall enabled chassis, a state entry is created. If a session connection with same information (source address, source port, destination address, destination port, protocol) is requested the firewall subsystem compares the packet's information to the state table entry to determine the validity of session. If the packet is currently in a table entry, it allows it to pass, otherwise it is dropped.

Transport and Network Protocols and States

Transport protocols have their connection's state tracked in various ways. Many attributes, including IP address and port combination, sequence numbers, and flags are used to track the individual connection. The combination of this information is kept as a hash in the state table.

TCP Protocol and Connection State

TCP is considered as a stateful connection-oriented protocol that has well defined session connection states. TCP tracks the state of its connections with flags as defined for TCP protocol. The following table describes different TCP connection states.

Table 103. TCP Connection States

State Flag	Description
TCP (Establishing Connection)	
CLOSED	A "non-state" that exists before a connection actually begins.
LISTEN	The state a host is in waiting for a request to start a connection. This is the starting state of a TCP connection.
SYN-SENT	The time after a host has sent out a SYN packet and is waiting for the proper SYN-ACK reply.
SYN-RCVD	The state a host is in after receiving a SYN packet and replying with its SYN-ACK reply.
ESTABLISHED	The state a host is in after its necessary ACK packet has been received. The initiating host goes into this state after receiving a SYN-ACK.
TCP (Closing Connection)	
FIN-WAIT-1	The state a connection is in after it has sent an initial FIN packet asking for a graceful termination of the TCP connection.
CLOSE-WAIT	The state a host's connection is in after it receives an initial FIN and sends back an ACK to acknowledge the FIN.
FIN-WAIT-2	The connection state of the host that has received the ACK response to its initial FIN, as it waits for a final FIN from its connection peer.
LAST-ACK	The state of the host that just sent the second FIN needed to gracefully close the TCP connection back to the initiating host while it waits for an acknowledgement.
TIME-WAIT	The state of the initiating host that received the final FIN and has sent an ACK to close the connection and waiting for an acknowledgement of ACK from the connection peer. Note that the amount of time the TIME-STATE is defined to pause is equal to the twice of the Maximum Segment Lifetime (MSL), as defined for the TCP implementation.
CLOSING	A state that is employed when a connection uses the unexpected simultaneous close.

UDP Protocol and Connection State

UDP is a connection-less transport protocol. Due to its connection-less nature, tracking of its state is a more complicated process than TCP. The Personal Stateful Firewall tracks a UDP connection in a different manner than TCP. A UDP packet has no sequence number or flag field in it. The port numbers used in UDP packet flow change randomly for any given session connection. So the Personal Stateful Firewall keeps the status of IP addresses.

UDP traffic cannot correct communication issues on its own and it relies entirely on ICMP as its error handler. This method makes ICMP an important part of a UDP session for tracking its overall state.

UDP has no set method of connection teardown that announces the session's end. Because of the lack of a defined ending, the Personal Stateful Firewall clears a UDP session's state table entries after a preconfigured timeout value reached.

ICMP Protocol and Connection State

ICMP is also a connection-less network protocol. The ICMP protocol is often used to return error messages when a host or protocol cannot do so on its own. ICMP response-type messages are precipitated by requests using other protocols like TCP or UDP. This way of messaging and its connection-less and one-way communication make the tracking of its state a much more complicated process than UDP. The Personal Stateful Firewall tracks an ICMP connection based on IP address and request message type information in a state table.

Like UDP, the ICMP connection lacks a defined session ending process, the Personal Stateful Firewall clears a state table entry on a predetermined timeout.

Firewall now supports ICMP Traceroute to handle ICMP packets with type value 30 that were being dropped. ICMP packets with ICMP type value 30 are called ICMP Traceroute packets.

It is now possible to allow/deny the ICMP echo packets having identifier value zero. By default, these packets are allowed. This feature will be effective only if Firewall is enabled (Firewall or Firewall+NAT) for a call. For only NAT enabled calls, there is no change in the behavior. Configuration is available only if Firewall license is present.

Application-Level Traffic and States

The Personal Stateful Firewall uses Deep Packet Inspection (DPI) functionality to manage application-level traffic and its state. With the help of DPI functionality, the Personal Stateful Firewall inspects packets up to Layer-7. It takes application behaviors into account to verify that all session-related traffic is properly handled and then decides which traffic to allow into the network.

Different applications follow different rules for communication exchange so the Personal Stateful Firewall manages the different communication sessions with different rules through DPI functionality.

The Personal Stateful Firewall also provides inspection and filtering functionality on application content with DPI. Personal Stateful Firewall is responsible for performing many simultaneous functions and it detect, allow, or drop packets at the ingress point of the network.

HTTP Application and State

HTTP is the one of the main protocols used on the Internet today. It uses TCP as its transport protocol, and its session initialization follows the standard TCP connection method.

Due to the TCP flow, the HTTP allows an easier definition of the overall session's state. It uses a single established connection from the client to the server and all its requests are outbound and responses are inbound. The state of the connection matches with the TCP state tracking.

For content verification and validation on the HTTP application session, the Personal Stateful Firewall uses DPI functionality in the chassis.

PPTP Application and State

Point-to-Point Tunneling Protocol (PPTP) is one of the protocols widely used to achieve Virtual Private Networks (VPN). PPTP allows the Point-to-Point Protocol (PPP) to be tunneled through an IP network. PPTP uses an enhanced GRE (Generic Routing Encapsulation) to carry PPP packets.

PPTP protocol has 2 connection states - Control connection (TCP) and Data connection (GREv1). PPTP exchanges IP or port specific information over its control connection and that information will be used to transfer the data over tunnel. If a PPTP client resides behind NAT and uses private IP to communicate with the outside world, it is possible that the information exchange over PPTP control flow has private IPs.

TFTP Application and State

Trivial File Transfer Protocol (TFTP) is an application layer protocol which is used by File Transfer applications. TFTP uses UDP (User Datagram Protocol) as its transport protocol and has only basic functionalities. TFTP file operations include sending a file and receiving a file. TFTP supports different modes for File Transfer which are netascii, ascii, octet, and binary.

TFTP has two connection states - Control connection and Data connection that operate on UDP. Initially, TFTP starts the control flow (uses UDP Port 69) for communicating the type of file operation to be performed. The Client initiates the connection towards Server on port 69 (UDP). Server replies to the Client from a port other than 69 and data is transferred in this flow. Negative reply is sent using different error codes supported by TFTP.

File Transfer Protocol and State

FTP is an application to move files between systems across the network. This is a two way connection and uses TCP as its transport protocol.

Due to TCP flow, FTP allows an easier definition of the overall session's state. As it uses a single established connection from the client to the server, the state of the connection matches with the TCP state tracking.

Personal Stateful Firewall uses application-port mapping along with FTP application-level content verification and validation with DPI functionality in the chassis. It also supports Pinhole data structure and Initialization, wherein FTP ALG parses FTP Port command to identify the initiation and termination end points of future FTP DATA sessions. The source/destination IP and destination Port of FTP DATA session is stored.

When a new session is to be created for a call, a check is made to see if the source/destination IP and Destination Port of this new session matches with the values stored. Upon match, a new ACS data session is created.

This lookup in the pinhole list is made before port trigger check and stateful firewall ruledef match. If the look up returns a valid pinhole then a particular session is allowed. Whenever a new FTP data session is allowed because of a pinhole match the associated pinhole is deleted. Pinholes are also expired if the associated FTP Control session is deleted in, or when the subscriber call goes down.

Chapter 29
Serving Gateway Overview

The Cisco® ASR 5x00 core platform provides wireless carriers with a flexible solution that functions as a Serving Gateway (S-GW) in Long Term Evolution-System Architecture Evolution (LTE-SAE) wireless data networks.

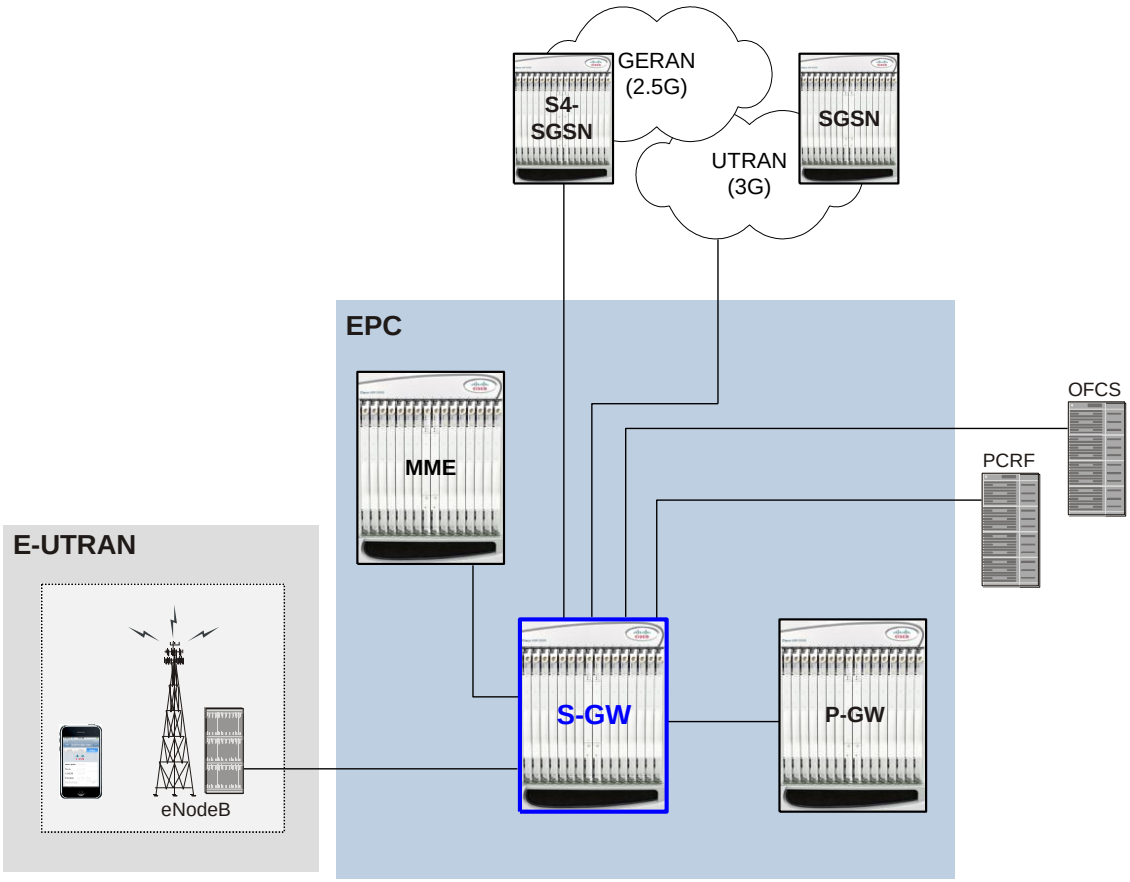
This overview provides general information about the S-GW including:

- Product Description
- Network Deployment(s)
- Features and Functionality - Base Software
- Features and Functionality - External Application Support
- Features and Functionality - Optional Enhanced Feature Software
- How the Serving Gateway Works
- Supported Standards

Product Description

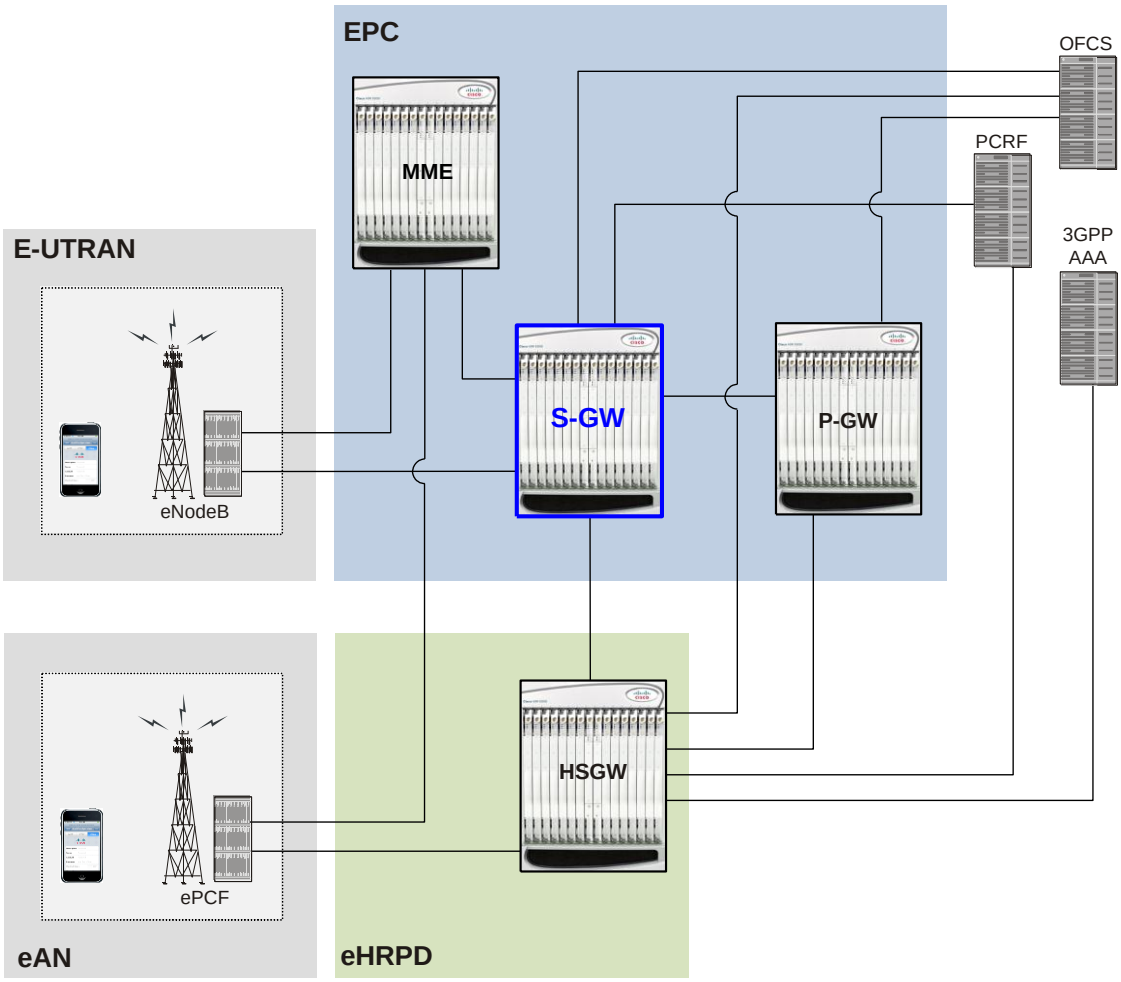
The Serving Gateway routes and forwards data packets from the UE and acts as the mobility anchor during inter-eNodeB handovers. Signals controlling the data traffic are received on the S-GW from the MME which determines the S-GW that will best serve the UE for the session. Every UE accessing the EPC is associated with a single S-GW.

Figure 224. S-GW in the Basic E-UTRAN/EPC Network



The S-GW is also involved in mobility by forwarding down link data during a handover from the E-UTRAN to the eHRPD network. An interface from the eAN/ePCF to an MME provides signaling that creates a GRE tunnel between the S-GW and the eHRPD Serving Gateway.

Figure 225. S-GW in the Basic E-UTRAN/EPC and eHRPD Network



The functions of the S-GW include:

- packet routing and forwarding.
- providing the local mobility anchor (LMA) point for inter-eNodeB handover and assisting the eNodeB reordering function by sending one or more "end marker" packets to the source eNodeB immediately after switching the path.
- mobility anchoring for inter-3GPP mobility (terminating the S4 interface from an SGSN and relaying the traffic between 2G/3G system and a PDN gateway.
- packet buffering for ECM-IDLE mode downlink and initiation of network triggered service request procedure.
- replicating user traffic in the event that Lawful Interception (LI) is required.
- transport level packet marking.
- user accounting and QoS class indicator (QCI) granularity for charging.
- uplink and downlink charging per UE, PDN, and QCI.
- support of circuit switched fallback (CSFB) for re-using deployed CS domain access for voice and other CS domain services.

Platform Requirements

The S-GW service runs on a Cisco® ASR 5x00 Series chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

Licenses

The S-GW is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

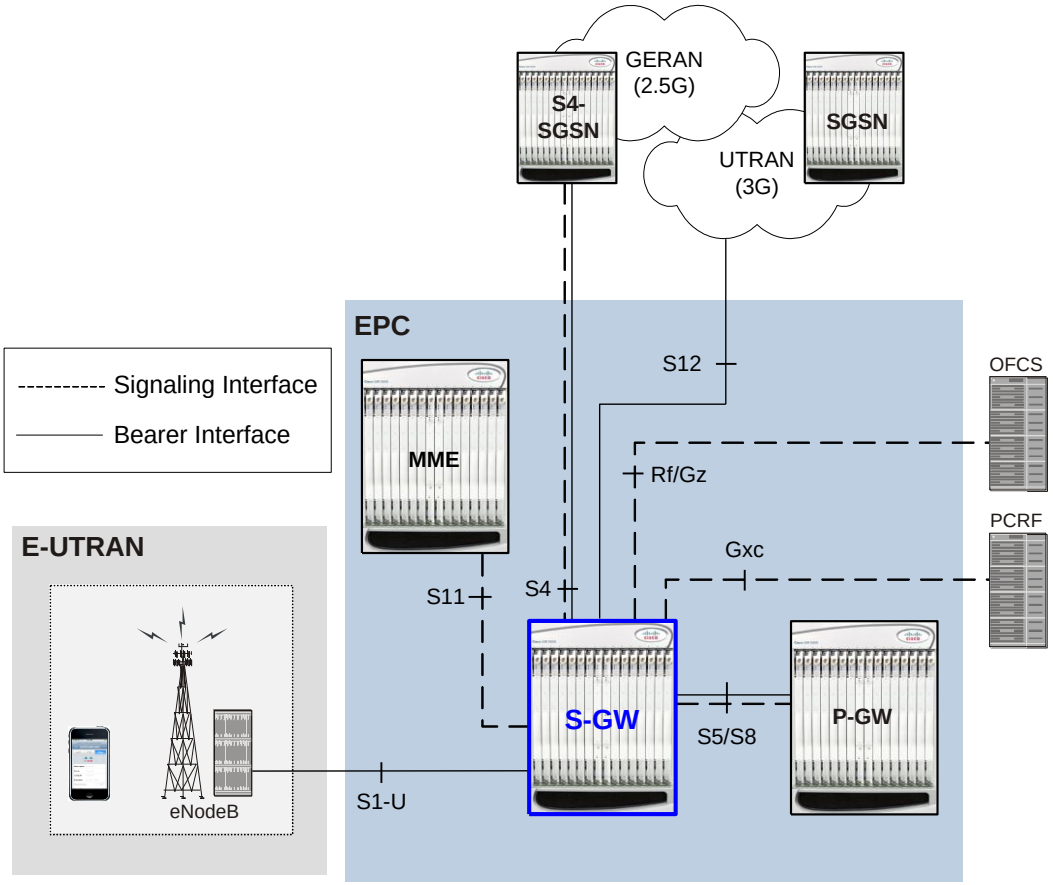
Network Deployment(s)

This section describes the supported interfaces and the deployment scenarios of a Serving Gateway.

Serving Gateway in the E-UTRAN/EPC Network

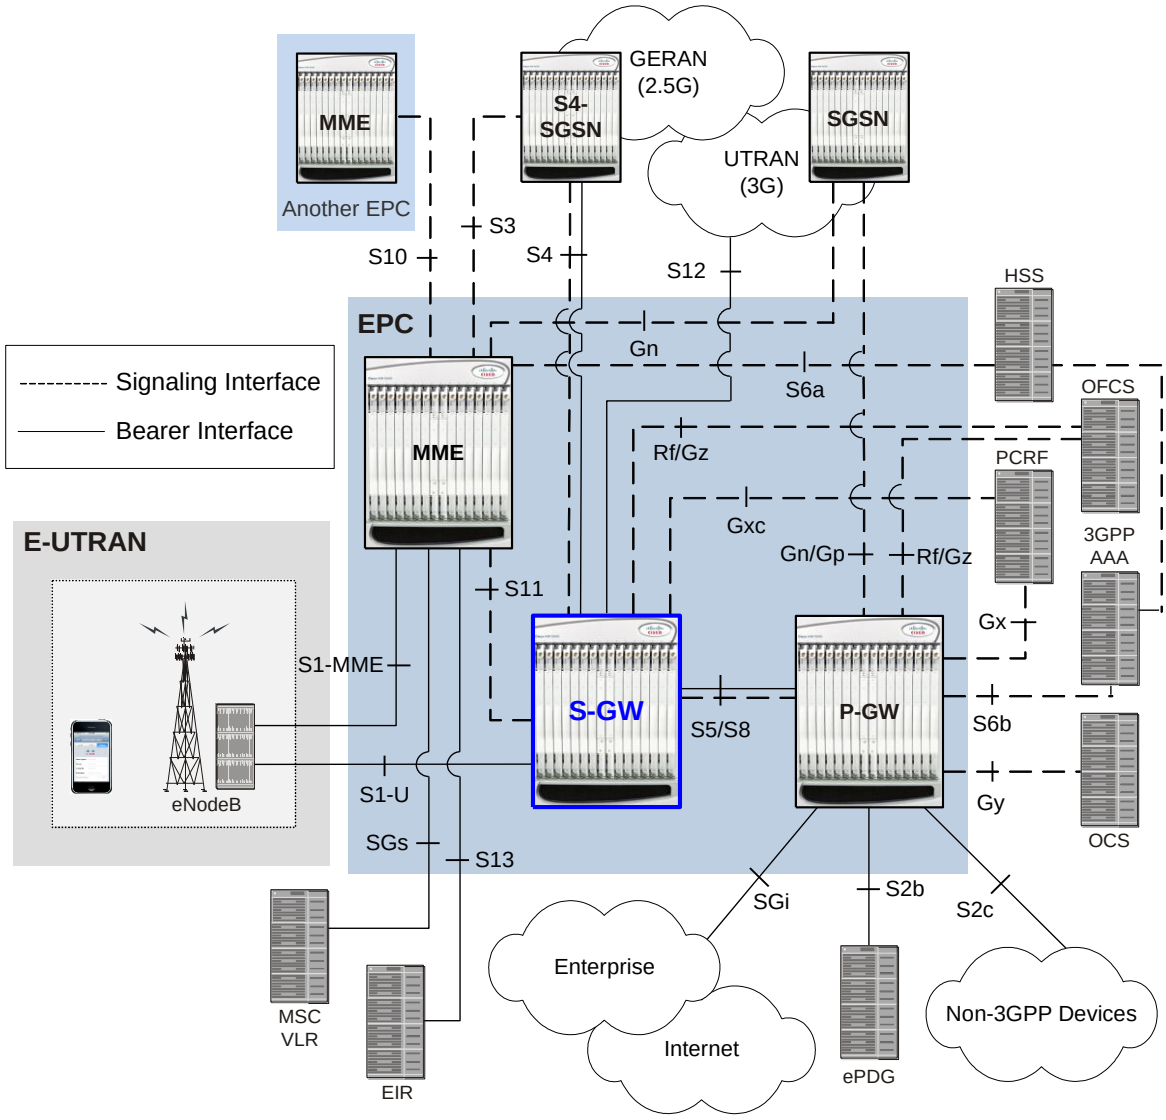
The following figure displays the specific network interfaces supported by the S-GW. Refer to Supported Logical Network Interfaces (Reference Points) for detailed information about each interface.

Figure 226. Supported S-GW Interfaces in the E-UTRAN/EPC Network



The following figure displays a sample network deployment of an S-GW, including all of the interface connections with other 3GPP Evolved-UTRAN/Evolved Packet Core network devices.

Figure 227. S-GW in the E-UTRAN/EPC Network



Supported Logical Network Interfaces (Reference Points)

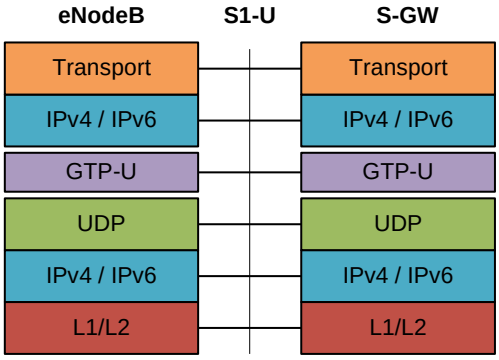
The S-GW provides the following logical network interfaces in support of the E-UTRAN/EPC network:

S1-U Interface

This reference point provides bearer channel tunneling between the eNodeB and the S-GW. It also supports eNodeB path switching during handovers. The S-GW provides the local mobility anchor point for inter-eNodeB hand-overs. It provides inter-eNodeB path switching during hand-overs when the X2 handover interface between base stations cannot be used. The S1-U interface uses GPRS tunneling protocol for user plane (GTP-Uv1). GTP encapsulates all end user IP packets and it relies on UDP/IP transport. The S1-U interface also supports IPSec IKEv2. This interface is defined in 3GPP TS 23.401.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: IPv4 or IPv6 GTPv1-U (bearer channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

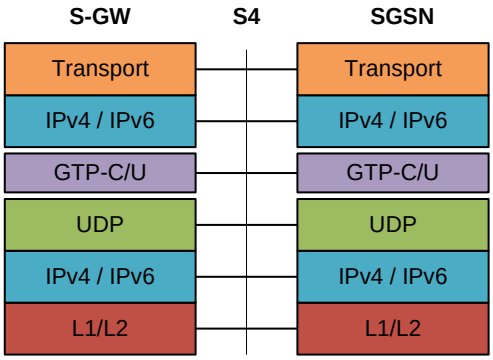


S4 Interface

This reference point provides tunneling and management between the S-GW and a 3GPP S4 SGSN. The interface facilitates soft hand-offs with the EPC network by providing control and mobility support between the inter-3GPP anchor function of the S-GW. This interface is defined in 3GPP TS 23.401.

Supported protocols:

- Transport Layer: UDP
- Tunneling:
  - GTP: IPv4 or IPv6 GTP-C (GTPv2 control/signaling channel) and GTP-U (GTPv1 user/bearer channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



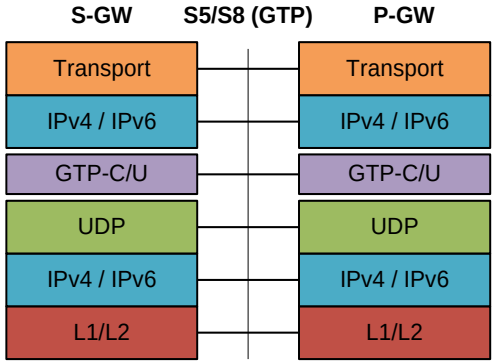
S5/S8 Interface

This reference point provides tunneling and management between the S-GW and the P-GW, as defined in 3GPP TS 23.401. The S8 interface is an inter-PLMN reference point between the S-GW and the P-GW used during roaming

scenarios. The S5 interface is used between an S-GW and P-GW located within the same administrative domain (non-roaming). It is used for S-GW relocation due to UE mobility and if the S-GW needs to connect to a non-collocated P-GW for the required PDN connectivity.

Supported protocols:

- Transport Layer: UDP, TCP
- Tunneling: GTP: GTPv2-C (signaling channel), GTPv1-U (bearer channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

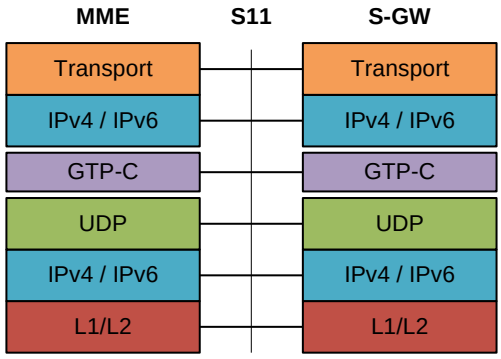


S11 Interface

This reference point provides GTP-C control signal tunneling between the MME and the S-GW. One GTP-C tunnel is created for each mobile terminal between the MME and S-GW. This interface is defined in 3GPP TS 23.401.

Supported protocols:

- Transport Layer: UDP
- Tunneling: IPv4 or IPv6 GTPv2-C (signalling channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

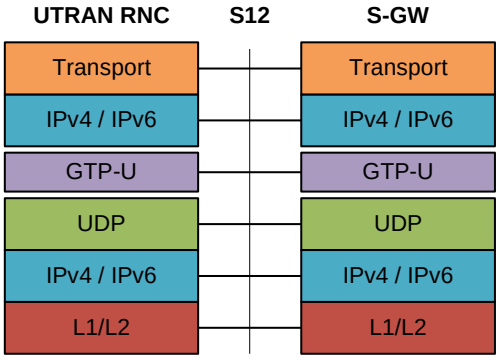


S12 Interface

This reference point provides GTP-U bearer/user direct tunneling between the S-GW and a UTRAN Radio Network Controller (RNC), as defined in 3GPP TS 23.401. This interface provides support for inter-RAT handovers between the 3G RAN and EPC allowing a direct tunnel to be initiated between the RNC and S-GW, thus bypassing the S4 SGSN and reducing latency.

Supported protocols:

- Transport Layer: UDP
- Tunneling: IPv4 or IPv6 GTP-U (GTPv1 bearer/user channel)
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

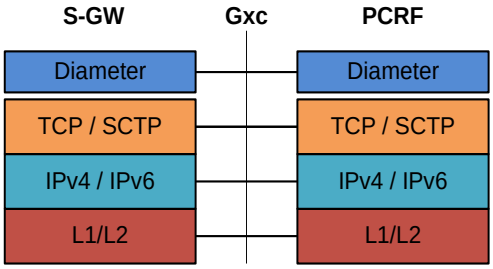


Gxc Interface

This signaling interface supports the transfer of policy control and charging rules information (QoS) between the Bearer Binding and Event Reporting Function (BBERF) on the S-GW and a Policy and Charging Rules Function (PCRF) server.

Supported protocols:

- Transport Layer: TCP or SCTP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet

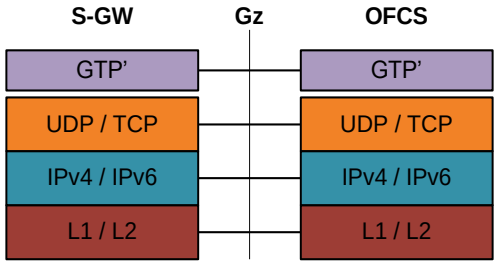


Gz Interface

The Gz reference interface enables offline accounting functions on the S-GW. The S-GW collects charging information for each mobile subscriber UE pertaining to the radio network usage. The Gz interface and offline accounting functions are used primarily in roaming scenarios where the foreign P-GW does not support offline charging.

Supported protocols:

- Transport Layer: TCP
- Network Layer: IPv4, IPv6
- Data Link Layer: ARP
- Physical Layer: Ethernet



Features and Functionality - Base Software

This section describes the features and functions supported by default in the base software for the S-GW service and do not require any additional licenses to implement the functionality.

Important: To configure the basic service and functionality on the system for the S-GW service, refer to the configuration examples provided in the *Serving Gateway Administration Guide*.

The following features are supported and brief descriptions are provided in this section:

- ANSI T1.276 Compliance
- Bulk Statistics Support
- Circuit Switched Fall Back (CSFB) Support
- Congestion Control
- IP Access Control Lists
- IPv6 Capabilities
- Management System Overview
- Multiple PDN Support
- Online/Offline Charging
- Operator Policy Support
- QoS Bearer Management
- Subscriber Level Trace
- Threshold Crossing Alerts (TCA) Support

ANSI T1.276 Compliance

ANSI T1.276 specifies security measures for Network Elements (NE). In particular it specifies guidelines for password strength, storage, and maintenance security measures.

ANSI T1.276 specifies several measures for password security. These measures include:

- Password strength guidelines
- Password storage guidelines for network elements
- Password maintenance, e.g. periodic forced password changes

These measures are applicable to the ASR 5x00 Platform and the Web Element Manager since both require password authentication. A subset of these guidelines where applicable to each platform will be implemented. A known subset of guidelines, such as certificate authentication, are not applicable to either product. Furthermore, the platforms support a variety of authentication methods such as RADIUS and SSH which are dependent on external elements. ANSI T1.276 compliance in such cases will be the domain of the external element. ANSI T1.276 guidelines will only be implemented for locally configured operators.

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Following is a partial list of supported schemas:

- **System**: Provides system-level statistics
- **Card**: Provides card-level statistics
- **Port**: Provides port-level statistics
- **MAG**: Provides MAG service statistics
- **S-GW**: Provides S-GW node-level service statistics
- **IP Pool**: Provides IP pool statistics
- **APN**: Provides Access Point Name statistics

The system supports the configuration of up to four sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, system host name, system uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

The Cisco Web Element Manager is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in its PostgreSQL database. It can also generate XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, the Bulk Statistics server can archive files to an alternative directory on the server. The directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Important: For more information on bulk statistic configuration, refer to the *Configuring and Maintaining Bulk Statistics* chapter in the *System Administration Guide*.

Circuit Switched Fall Back (CSFB) Support

Circuit Switched Fall Back (CSFB) enables the UE to camp on an EUTRAN cell and originate or terminate voice calls through a forced switchover to the circuit switched (CS) domain or other CS-domain services (for example, Location Services (LCS) or supplementary services). Additionally, SMS delivery via the CS core network is realized without CSFB. Since LTE EPC networks were not meant to directly anchor CS connections, when any CS voice services are initiated, any PS based data activities on the EUTRAN network will be temporarily suspended (either the data transfer is suspended or the packet switched connection is handed over to the 2G/3G network).

CSFB provides an interim solution for enabling telephony and SMS services for LTE operators that do not plan to deploy IMS packet switched services at initial service launch.

The S-GW supports CSFB messaging over the S11 interface over GTP-C. Supported messages are:

- Suspend Notification
- Suspend Acknowledge
- Resume Notification
- Resume Acknowledgement

Congestion Control

The congestion control feature allows you to set policies and thresholds and specify how the system reacts when faced with a heavy load condition.

Congestion control monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are quickly resolved. However, continuous or large numbers of these conditions within a specific time interval may have an impact the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establish limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operational thresholds that are configured for the system as described in the *Thresholding Configuration Guide*. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated.

  A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.

  - **Port Utilization Thresholds**: If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
  - **Port-specific Thresholds**: If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
- **Service Congestion Policies:** Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

Important: For more information on congestion control, refer to the *Congestion Control* appendix in the *System Administration Guide*.

IP Access Control Lists

IP access control lists allow you to set up rules that control the flow of packets into and out of the system based on a variety of IP packet parameters.

IP access lists, or Access Control Lists (ACLs) as they are commonly referred to, control the flow of packets into and out of the system. They are configured on a per-context basis and consist of "rules" (ACL rules) or filters that control the action taken on packets that match the filter criteria. Once configured, an ACL can be applied to any of the following:

- An individual interface
- All traffic facilitated by a context (known as a policy ACL)
- An individual subscriber
- All subscriber sessions facilitated by a specific context

Important: The S-GW supports interface-based ACLs only. For more information on IP access control lists, refer to the *IP Access Control Lists* appendix in the *System Administration Guide*.

IPv6 Capabilities

IPv6 enables increased address efficiency and relieves pressures caused by rapidly approaching IPv4 address exhaustion problem.

The S-GW platform offers the following IPv6 capabilities:

IPv6 Connections to Attached Elements

IPv6 transport and interfaces are supported on all of the following connections:

- Diameter Gxc policy signaling interface
- Diameter Rf offline charging interface
- Lawful Intercept (X1, X2 interfaces)

Routing and Miscellaneous Features

- OSPFv3
- MP-BGP v6 extensions
- IPv6 flows (Supported on all Diameter QoS and Charging interfaces as well as Inline Services (for example, ECS, P2P detection, Stateful Firewall, etc.)

Management System Overview

The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management - focusing on providing superior quality Network Element (NE) and element management system (Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems - giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

Cisco's O+M module offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces.

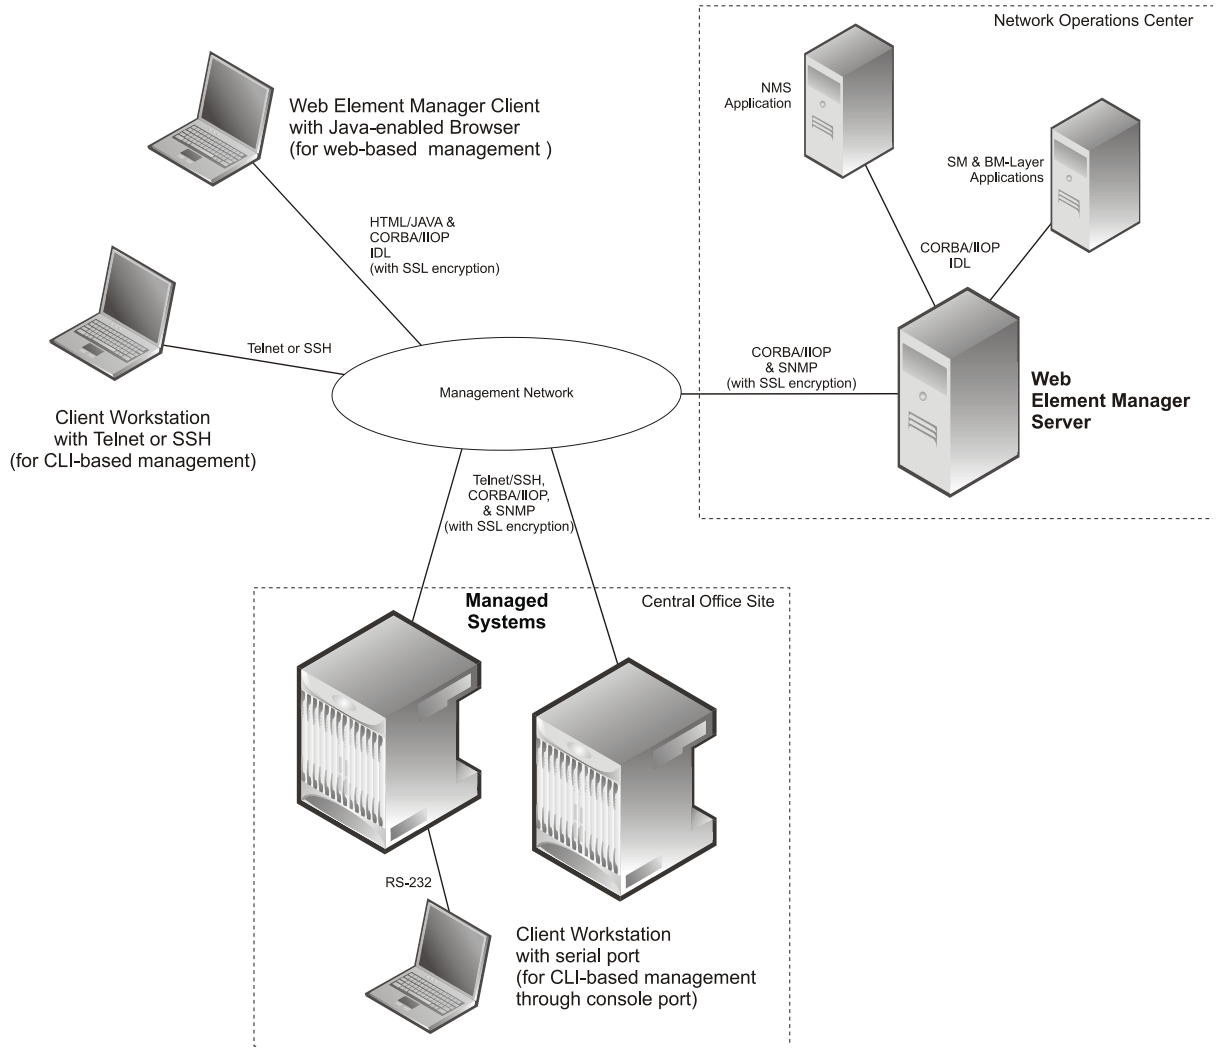
These include:

- Using the Command Line Interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the console port on the SPIO card via an RS-232 serial connection
- Using the Web Element Manager application

- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or
- 1000Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (such as, Microsoft Internet Explorer v6.0 and above or others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network.

Figure 228. Element Management Methods



Important: P-GW management functionality is enabled by default for console-based access. For GUI-based management support, refer to the Web Element Manager section in this chapter.

Important: For more information on command line interface based management, refer to the *Command Line Interface Reference* and *P-GW Administration Guide*.

Multiple PDN Support

Enables an APN-based user experience that enables separate connections to be allocated for different services including IMS, Internet, walled garden services, or offdeck content services.

The Mobile Access Gateway (MAG) function on the S-GW can maintain multiple PDN or APN connections for the same user session. The MAG runs a single node level Proxy Mobile IPv6 (PMIP) tunnel for all user sessions toward the Local Mobility Anchor (LMA) function of the P-GW.

When a user wants to establish multiple PDN connections, the MAG brings up the multiple PDN connections over the same PMIPv6 session to one or more P-GW LMAs. The P-GW in turn allocates separate IP addresses (Home Network Prefixes) for each PDN connection and each one can run one or multiple EPC default and dedicated bearers. To request the various PDN connections, the MAG includes a common MN-ID and separate Home Network Prefixes, APNs and a Handover Indication Value equal to one in the PMIPv6 Binding Updates.

Important: Up to 11 multiple PDN connections are supported.

Online/Offline Charging

The Cisco EPC platforms support offline charging interactions with external OCS and CGF/CDF servers. To provide subscriber level accounting, the Cisco EPC platform supports integrated Charging Transfer Function (CTF) and Charging Data Function (CDF) / Charging Gateway Function (CGF). Each gateway uses Charging-IDs to distinguish between default and dedicated bearers within subscriber sessions.

The ASR 5x00 platform offers a local directory to enable temporary file storage and buffer charging records in persistent memory located on a pair of dual redundant RAID hard disks. Each drive includes 147GB of storage and up to 100GB of capacity is dedicated to storing charging records. For increased efficiency it also possible to enable file compression using protocols such as GZIP.

The offline charging implementation offers built-in heart beat monitoring of adjacent CGFs. If the Cisco P-GW has not heard from the neighboring CGF within the configurable polling interval, it will automatically buffer the charging records on the local drives until the CGF reactivates itself and is able to begin pulling the cached charging records.

Online: Gy Reference Interface

The P-GW supports a Policy Charging Enforcement Function (PCEF) to enable Flow Based Bearer Charging (FBC) via the Gy reference interface to adjunct Online Charging System (OCS) servers. The Gy interface provides a standardized Diameter interface for real-time content-based charging of data services. It is based on the 3GPP standards and relies on quota allocation. The Gy interface provides an online charging interface that works with the ECS Deep Packet Inspection feature. With Gy, customer traffic can be gated and billed. Both time- and volume-based charging models are supported.

Offline: Gz Reference Interface

The Cisco P-GW and S-GW support 3GPP Release 8 compliant offline charging as defined in TS 32.251, TS 32.297 and 32.298. Whereas the S-GW generates SGW-CDRs to record subscriber level access to PLMN resources, the P-GW creates PGW-CDRs to record user access to external networks. Additionally when Gn/Gp interworking with SGSNs is enabled, the GGSN service on the P-GW records G-CDRs to record user access to external networks.

To provide subscriber level accounting, the Cisco S-GW supports integrated Charging Transfer Function (CTF) and Charging Data Function (CDF). Each gateway uses Charging-IDs to distinguish between default and dedicated bearers within subscriber sessions.

The Gz reference interface between the CDF and CGF is used to transfer charging records via the GTPP protocol. In a standards based implementation, the CGF consolidates the charging records and transfers them via an FTP or SFTP connection over the Bm reference interface to a back-end billing mediation server. The Cisco EPC gateways also offer the ability to transfer charging records between the CDF and CGF serve via FTP or SFTP. CDR records include information such as Record Type, Served IMSI, ChargingID, APN Name, TimeStamp, Call Duration, Served MSISDN, PLMN-ID, etc.

Operator Policy Support

The operator policy provides mechanisms to fine tune the behavior of subsets of subscribers above and beyond the behaviors described in the user profile. It also can be used to control the behavior of visiting subscribers in roaming scenarios, enforcing roaming agreements and providing a measure of local protection against foreign subscribers.

An operator policy associates APNs, APN profiles, an APN remap table, and a call-control profile to ranges of IMSIs. These profiles and tables are created and defined within their own configuration modes to generate sets of rules and instructions that can be reused and assigned to multiple policies. In this manner, an operator policy manages the application of rules governing the services, facilities, and privileges available to subscribers. These policies can override standard behaviors and provide mechanisms for an operator to get around the limitations of other infrastructure elements, such as DNS servers and HSSs.

The operator policy configuration to be applied to a subscriber is selected on the basis of the selection criteria in the subscriber mapping at attach time. A maximum of 1,024 operator policies can be configured. If a UE was associated with a specific operator policy and that policy is deleted, the next time the UE attempts to access the policy, it will attempt to find another policy with which to be associated.

A default operator policy can be configured and applied to all subscribers that do not match any of the per-PLMN or IMSI range policies.

Changes to the operator policy take effect when the subscriber re-attaches and subsequent EPS Bearer activations.

QoS Bearer Management

Provides a foundation for contributing towards improved Quality of User Experience (QoE) by enabling deterministic end-to-end forwarding and scheduling treatments for different services or classes of applications pursuant to their requirements for committed bandwidth resources, jitter and delay. In this way, each application receives the service treatment that users expect.

An EPS bearer is a logical aggregate of one or more Service Data Flows (SDFs), running between a UE and a P-GW in case of GTP-based S5/S8, and between a UE and HSGW in case of PMIP-based S2a connection. An EPS bearer is the level of granularity for bearer level QoS control in the EPC/E-UTRAN. The Cisco P-GW maintains one or more Traffic Flow Templates (TFTs) in the downlink direction for mapping inbound Service Data Flows (SDFs) to EPS bearers. The P-GW maps the traffic based on the downlink TFT to the S5/S8 bearer. The Cisco P-GW offers all of the following bearer-level aggregate constructs:

QoS Class Identifier (QCI): An operator provisioned value that controls bearer level packet forwarding treatments (for example, scheduling weights, admission thresholds, queue management thresholds, link layer protocol configuration, etc). Cisco EPC gateways also support the ability to map the QCI values to DiffServ codepoints in the outer GTP tunnel header of the S5/S8 connection. Additionally, the platform also provides configurable parameters to copy the DSCP marking from the encapsulated payload to the outer GTP tunnel header.

Guaranteed Bit Rate (GBR): A GBR bearer is associated with a dedicated EPS bearer and provides a guaranteed minimum transmission rate in order to offer constant bit rate services for applications such as interactive voice that require deterministic low delay service treatment.

Maximum Bit Rate (MBR): The MBR attribute provides a configurable burst rate that limits the bit rate that can be expected to be provided by a GBR bearer (e.g. excess traffic may get discarded by a rate shaping function). The MBR may be greater than or equal to the GBR for a given dedicated EPS bearer.

Aggregate Maximum Bit Rate (AMBR): AMBR denotes a bit rate of traffic for a group of bearers destined for a particular PDN. The Aggregate Maximum Bit Rate is typically assigned to a group of Best Effort service data flows over the Default EPS bearer. That is, each of those EPS bearers could potentially utilize the entire AMBR, e.g. when the other EPS bearers do not carry any traffic. The AMBR limits the aggregate bit rate that can be expected to be provided by the EPS bearers sharing the AMBR (e.g. excess traffic may get discarded by a rate shaping function). AMBR applies to all Non-GBR bearers belonging to the same PDN connection. GBR bearers are outside the scope of AMBR.

Policing and Shaping: The Cisco P-GW offers a variety of traffic conditioning and bandwidth management capabilities. These tools enable usage controls to be applied on a per-subscriber, per-EPS bearer or per-PDN/APN basis. It is also possible to apply bandwidth controls on a per-APN AMBR capacity. These applications provide the ability to inspect and maintain state for user sessions or Service Data Flows (SDF's) within them using shallow L3/L4 analysis or high touch deep packet inspection at L7. Metering of out-of-profile flows or sessions can result in packet discards or reducing the DSCP marking to Best Effort priority. When traffic shaping is enabled the P-GW enqueues the non-conforming session to the provisioned memory limit for the user session. When the allocated memory is exhausted, the inbound/outbound traffic for the user can be transmitted or policed in accordance with operator provisioned policy.

Subscriber Level Trace

Provides a 3GPP standards-based session level trace function for call debugging and testing new functions and access terminals in an LTE environment.

As a complement to Cisco's protocol monitoring function, the S-GW supports 3GPP standards based session level trace capabilities to monitor all call control events on the respective monitored interfaces including S1-U, S11, S5/S8, and Gxc. The trace can be initiated using multiple methods:

- Management initiation via direct CLI configuration
- Management initiation at HSS with trace activation via authentication response messages over S6a reference interface
- Signaling based activation through signaling from subscriber access terminal

Note: Once the trace is provisioned it can be provisioned through the access cloud via various signaling interfaces.

The session level trace function consists of trace activation followed by triggers. The EPC network element buffers the trace activation instructions for the provisioned subscriber in memory using camp-on monitoring. Trace files for active calls are buffered as XML files using non-volatile memory on the local dual redundant hard drives on the ASR 5x00 platform. The Trace Depth defines the granularity of data to be traced. Six levels are defined including Maximum, Minimum and Medium with ability to configure additional levels based on vendor extensions.

All call control activity for active and recorded sessions is sent to an off-line Trace Collection Entity (TCE) using a standards-based XML format over an FTP or secure FTP (SFTP) connection. In the current release the IPv4 interfaces are used to provide connectivity to the TCE. Trace activation is based on IMSI or IMEI.

Once a subscriber level trace request is activated it can be propagated via the S5/S8 signaling to provision the corresponding trace for the same subscriber call on the P-GW. The trace configuration will only be propagated if the P-GW is specified in the list of configured Network Element types received by the S-GW. Trace configuration can be specified or transferred in any of the following message types:

- S11: Create Session Request
- S11: Trace Session Activation
- S11: Modify Bearer Request

As subscriber level trace is a CPU intensive activity the maximum number of concurrently monitored trace sessions per Cisco P-GW or S-GW is 32. Use in a production network should be restricted to minimize the impact on existing services.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert**: A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm**: Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps**: SNMP traps have been created that indicate the condition (high threshold crossing and clear) of each of the monitored values.

  Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.
- **Logs**: The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

  Logs are supported in both the Alert and the Alarm models.
- **Alarm System**: High threshold alarms generated within the specified polling interval are considered outstanding until a the condition no longer exists or a condition clear alarm is generated. Outstanding alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

  The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

Features and Functionality - External Application Support

This section describes the features and functions of external applications supported on the S-GW. These services require additional licenses to implement the functionality.

- Web Element Management System

Web Element Manager

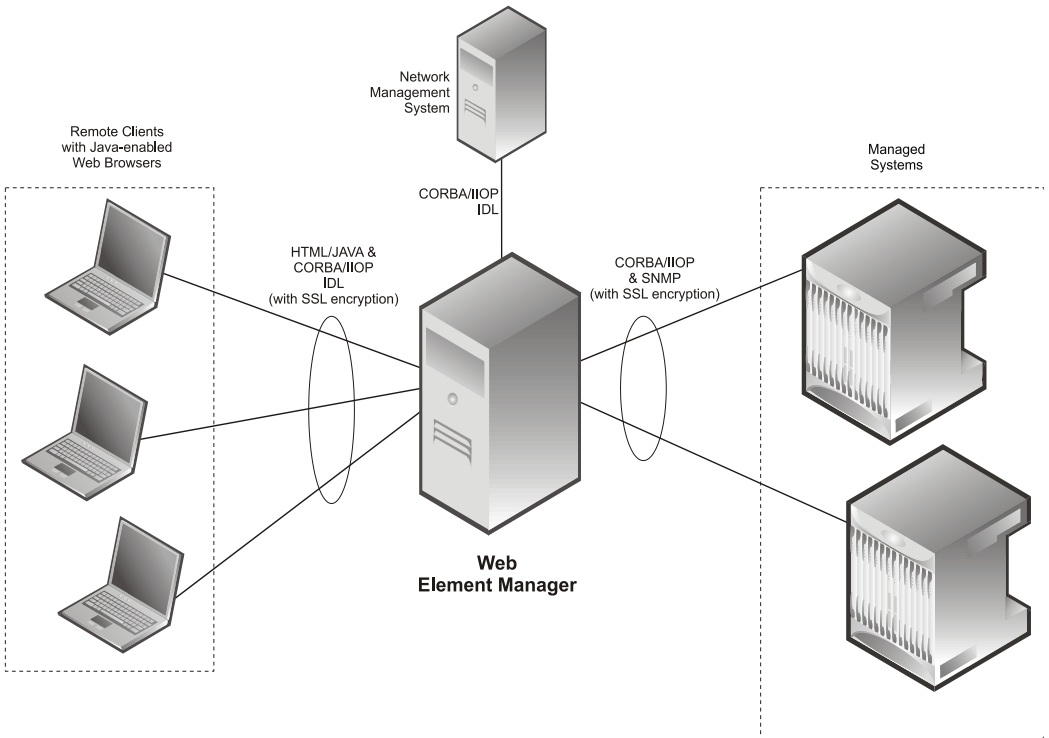
The Web Element Manager (WEM) provides a graphical user interface (GUI) for performing fault, configuration, accounting, performance, and security (FCAPS) management of the ASR 5x00 platform.

The Web Element Manager is a Common Object Request Broker Architecture (CORBA)-based application that provides complete fault, configuration, accounting, performance, and security (FCAPS) management capability for the system.

For maximum flexibility and scalability, the Web Element Manager application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces.

The following figure demonstrates various interfaces between the Cisco Web Element Manager and other network components.

Figure 229. Web Element Manager Network Interfaces



Important: For more information on WEM support, refer to the *WEM Installation and Administration Guide*.

Features and Functionality - Optional Enhanced Feature Software

This section describes the optional enhanced features and functions for the S-GW service.

Each of the following features require the purchase of an additional license to implement the functionality with the S-GW service.

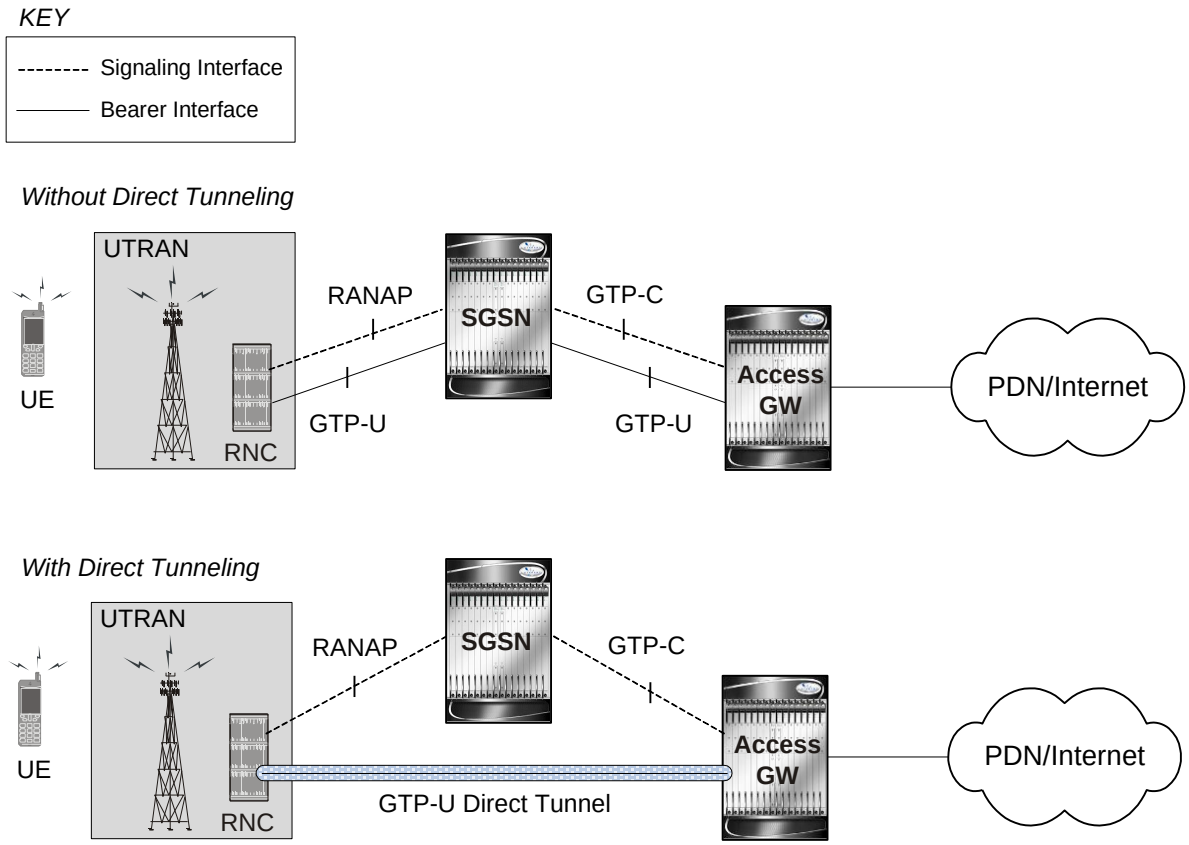
This section describes following features:

- Direct Tunnel
- Inter-Chassis Session Recovery
- IP Security (IPSec) Encryption
- Lawful Intercept
- Layer 2 Traffic Management (VLANs)
- Session Recovery Support

Direct Tunnel

In accordance with standards, one tunnel functionality enables the SGSN to establish a direct tunnel at the user plane level - a GTP-U tunnel, directly between the RAN and the S-GW.

Figure 230. GTP-U with Direct Tunnel



In effect, a direct tunnel reduces data plane latency as the tunnel functionality acts to remove the SGSN from the data plane and limit the SGSN to the control plane for processing. This improves the user experience (for example, expedites web page delivery, reduces round trip delay for conversational services). Additionally, direct tunnel functionality implements the standard SGSN optimization to improve the usage of user plane resources (and hardware) by removing the requirement from the SGSN to handle the user plane processing.

Typically, the SGSN establishes a direct tunnel at PDP context activation using an Update PDP Context Request towards the S-GW. This means a significant increase in control plane load on both the SGSN and S-GW components of the packet core. Hence, deployment requires highly scalable S-GWs since the volume and frequency of Update PDP Context messages to the S-GW will increase substantially. The ASR 5x00 platform capabilities ensure control plane capacity will not be a limiting factor with direct tunnel deployment.

For more information on Direct Tunnel configuration, refer to the *Direct Tunnel Configuration* appendix in this guide.

Inter-Chassis Session Recovery

The ASR 5x00 platform provide industry leading carrier class redundancy. The systems protects against all single points of failure (hardware and software) and attempts to recover to an operational state when multiple simultaneous failures occur.

The system provides several levels of system redundancy:

- Under normal N+1 PSC hardware redundancy, if a catastrophic PSC failure occurs all affected calls are migrated to the standby PSC if possible. Calls which cannot be migrated are gracefully terminated with proper call-termination signaling and accounting records are generated with statistics accurate to the last internal checkpoint
- If the Session Recovery feature is enabled, any total PSC failure will cause a PSC switchover and all established sessions for supported call-types are recovered without any loss of session.

Even though Cisco provides excellent intra-chassis redundancy with these two schemes, certain catastrophic failures which can cause total chassis outages, such as IP routing failures, line-cuts, loss of power, or physical destruction of the chassis, cannot be protected by this scheme. In such cases, the MME Inter-Chassis Session Recovery (ICSR) feature provides geographic redundancy between sites. This has the benefit of not only providing enhanced subscriber experience even during catastrophic outages, but can also protect other systems such as the RAN from subscriber re-activation storms.

ICSR allows for continuous call processing without interrupting subscriber services. This is accomplished through the use of redundant chassis. The chassis are configured as primary and backup with one being active and one in recovery mode. A checkpoint duration timer is used to control when subscriber data is sent from the active chassis to the inactive chassis. If the active chassis handling the call traffic goes out of service, the inactive chassis transitions to the active state and continues processing the call traffic without interrupting the subscriber session. The chassis determines which is active through a propriety TCP-based connection called a redundancy link. This link is used to exchange Hello messages between the primary and backup chassis and must be maintained for proper system operation.

Interchassis Communication

Chassis configured to support ICSR communicate using periodic Hello messages. These messages are sent by each chassis to notify the peer of its current state. The Hello message contains information about the chassis such as its configuration and priority. A dead interval is used to set a time limit for a Hello message to be received from the chassis' peer. If the standby chassis does not receive a Hello message from the active chassis within the dead interval, the standby chassis transitions to the active state. In situations where the redundancy link goes out of service, a priority scheme is used to determine which chassis processes the session. The following priority scheme is used:

- router identifier
- chassis priority
- SPIO MAC address

Checkpoint Messages

Checkpoint messages are sent from the active chassis to the inactive chassis. Checkpoint messages are sent at specific intervals and contain all the information needed to recreate the sessions on the standby chassis, if that chassis were to become active. Once a session exceeds the checkpoint duration, checkpoint data is collected on the session. The checkpoint parameter determines the amount of time a session must be active before it is included in the checkpoint message.

Important: For more information on inter-chassis session recovery support, refer to the *Interchassis Session Recovery* appendix in *System Administration Guide*.

IP Security (IPSec) Encryption

Enables network domain security for all IP packet switched LTE-EPC networks in order to provide confidentiality, integrity, authentication, and anti-replay protection. These capabilities are insured through use of cryptographic techniques.

The Cisco S-GW supports IKEv1 and IPSec encryption using IPv4 addressing. IPSec enables the following two use cases:

- Encryption of S8 sessions and EPS bearers in roaming applications where the P-GW is located in a separate administrative domain from the S-GW
- IPSec ESP security in accordance with 3GPP TS 33.210 is provided for S1 control plane, S1 bearer plane and S1 management plane traffic. Encryption of traffic over the S1 reference interface is desirable in cases where the EPC core operator leases radio capacity from a roaming partner's network.

Important: You must purchase an IPSec license to enable IPSec. For more information on IPSec support, refer to the *IP Security* appendix in this guide.

Lawful Intercept

The Cisco Lawful Intercept feature is supported on the S-GW. Lawful Intercept is a licensed-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your Cisco account representative.

Layer 2 Traffic Management (VLANs)

Virtual LANs (VLANs) provide greater flexibility in the configuration and use of contexts and services.

VLANs are configured as tags on a per-port basis and allow more complex configurations to be implemented. The VLAN tag allows a single physical port to be bound to multiple logical interfaces that can be configured in different contexts. Therefore, each Ethernet port can be viewed as containing many logical ports when VLAN tags are employed.

Important: For more information on VLAN support, refer to the VLANs appendix in the *System Administration Guide*.

Session Recovery Support

Provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

In the telecommunications industry, over 90 percent of all equipment failures are software-related. With robust hardware failover and redundancy protection, any card-level hardware failures on the system can quickly be corrected. However, software failures can occur for numerous reasons, many times without prior indication. StarOS has the ability to support stateful intra-chassis session recovery (ICSR) for S-GW sessions.

When session recovery occurs, the system reconstructs the following subscriber information:

- Data and control state information required to maintain correct call behavior
- Subscriber data statistics that are required to ensure that accounting information is maintained
- A best-effort attempt to recover various timer values such as call duration, absolute time, and others

Session recovery is also useful for in-service software patch upgrade activities. If session recovery is enabled during the software patch upgrade, it helps to preserve existing sessions on the active packet services card during the upgrade process.

Important: For more information on session recovery support, refer to the *Session Recovery* appendix in the *System Administration Guide*.

How the Serving Gateway Works

This section provides information on the function of the S-GW in an EPC E-UTRAN network and presents call procedure flows for different stages of session setup and disconnect.

The S-GW supports the following network flows:

- GTP Serving Gateway Call/Session Procedures in an LTE-SAE Network

GTP Serving Gateway Call/Session Procedures in an LTE-SAE Network

The following topics and procedure flows are included:

- Subscriber-initiated Attach (initial)
- Subscriber-initiated Detach

Subscriber-initiated Attach (initial)

This section describes the procedure of an initial attach to the EPC network by a subscriber.

Figure 231. Subscriber-initiated Attach (initial) Call Flow

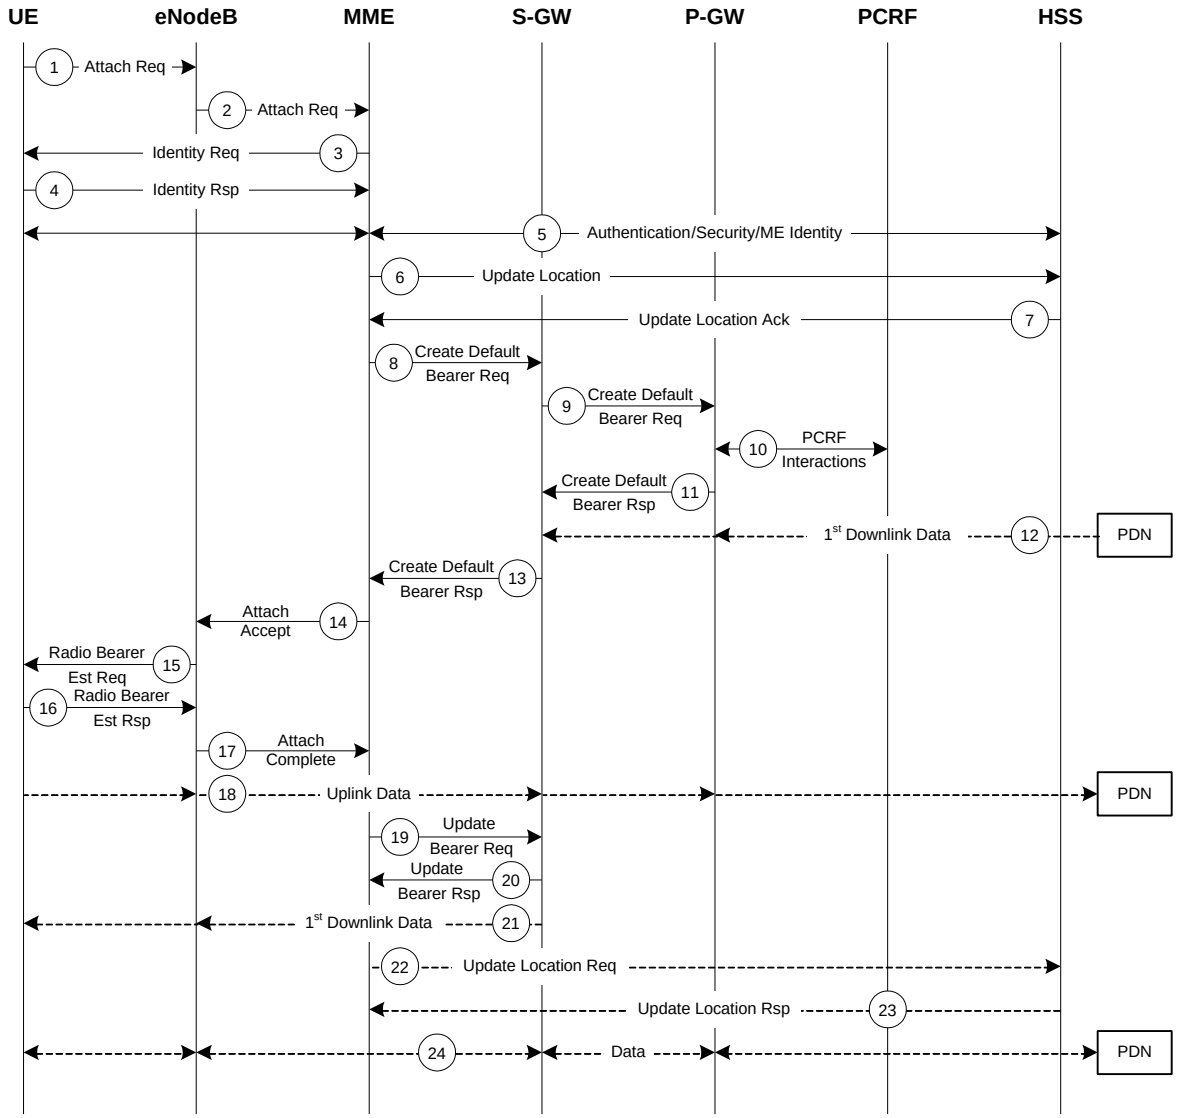


Table 104. Subscriber-initiated Attach (initial) Call Flow Description

Step	Description
1	The UE initiates the Attach procedure by the transmission of an Attach Request (IMSI or old GUTI, last visited TAI (if available), UE Network Capability, PDN Address Allocation, Protocol Configuration Options, Attach Type) message together with an indication of the Selected Network to the eNodeB. IMSI is included if the UE does not have a valid GUTI available. If the UE has a valid GUTI, it is included.
2	The eNodeB derives the MME from the GUTI and from the indicated Selected Network. If that MME is not associated with the eNodeB, the eNodeB selects an MME using an MME selection function. The eNodeB forwards the Attach Request message to the new MME contained in a S1-MME control message (Initial UE message) together with the Selected Network and an indication of the E-UTRAN Area identity, a globally unique E-UTRAN ID of the cell from where it received the message to the new MME.

Step	Description
3	If the UE is unknown in the MME, the MME sends an Identity Request to the UE to request the IMSI.
4	The UE responds with Identity Response (IMSI).
5	If no UE context for the UE exists anywhere in the network, authentication is mandatory. Otherwise this step is optional. However, at least integrity checking is started and the ME Identity is retrieved from the UE at Initial Attach. The authentication functions, if performed this step, involves AKA authentication and establishment of a NAS level security association with the UE in order to protect further NAS protocol messages.
6	The MME sends an Update Location (MME Identity, IMSI, ME Identity) to the HSS.
7	The HSS acknowledges the Update Location message by sending an Update Location Ack to the MME. This message also contains the Insert Subscriber Data (IMSI, Subscription Data) Request. The Subscription Data contains the list of all APNs that the UE is permitted to access, an indication about which of those APNs is the Default APN, and the EPS subscribed QoS profile for each permitted APN. If the Update Location is rejected by the HSS, the MME rejects the Attach Request from the UE with an appropriate cause.
8	The MME selects an S-GW using the Serving GW selection function and allocates an EPS Bearer Identity for the Default Bearer associated with the UE. If the PDN subscription context contains no P-GW address the MME selects a P-GW as described in clause PDN GW selection function. Then it sends a Create Default Bearer Request (IMSI, MME Context ID, APN, RAT type, Default Bearer QoS, PDN Address Allocation, AMBR, EPS Bearer Identity, Protocol Configuration Options, ME Identity, User Location Information) message to the selected S-GW.
9	The S-GW creates a new entry in its EPS Bearer table and sends a Create Default Bearer Request (IMSI, APN, S-GW Address for the user plane, S-GW TEID of the user plane, S-GW TEID of the control plane, RAT type, Default Bearer QoS, PDN Address Allocation, AMBR, EPS Bearer Identity, Protocol Configuration Options, ME Identity, User Location Information) message to the P-GW.
10	If dynamic PCC is deployed, the P-GW interacts with the PCRF to get the default PCC rules for the UE. The IMSI, UE IP address, User Location Information, RAT type, AMBR are provided to the PCRF by the P-GW if received by the previous message.
11	The P-GW returns a Create Default Bearer Response (P-GW Address for the user plane, P-GW TEID of the user plane, P-GW TEID of the control plane, PDN Address Information, EPS Bearer Identity, Protocol Configuration Options) message to the S-GW. PDN Address Information is included if the P-GW allocated a PDN address Based on PDN Address Allocation received in the Create Default Bearer Request. PDN Address Information contains an IPv4 address for IPv4 and/or an IPv6 prefix and an Interface Identifier for IPv6. The P-GW takes into account the UE IP version capability indicated in the PDN Address Allocation and the policies of operator when the P-GW allocates the PDN Address Information. Whether the IP address is negotiated by the UE after completion of the Attach procedure, this is indicated in the Create Default Bearer Response.
12	The Downlink (DL) Data can start flowing towards S-GW. The S-GW buffers the data.
13	The S-GW returns a Create Default Bearer Response (PDN Address Information, S-GW address for User Plane, S-GW TEID for User Plane, S-GW Context ID, EPS Bearer Identity, Protocol Configuration Options) message to the new MME. PDN Address Information is included if it was provided by the P-GW.
14	The new MME sends an Attach Accept (APN, GUTI, PDN Address Information, TAI List, EPS Bearer Identity, Session Management Configuration IE, Protocol Configuration Options) message to the eNodeB.
15	The eNodeB sends Radio Bearer Establishment Request including the EPS Radio Bearer Identity to the UE. The Attach Accept message is also sent along to the UE.
16	The UE sends the Radio Bearer Establishment Response to the eNodeB. In this message, the Attach Complete message (EPS Bearer Identity) is included.
17	The eNodeB forwards the Attach Complete (EPS Bearer Identity) message to the MME.

Step	Description
18	The Attach is complete and UE sends data over the default bearer. At this time the UE can send uplink packets towards the eNodeB which are then tunnelled to the S-GW and P-GW.
19	The MME sends an Update Bearer Request (eNodeB address, eNodeB TEID) message to the S-GW.
20	The S-GW acknowledges by sending Update Bearer Response (EPS Bearer Identity) message to the MME.
21	The S-GW sends its buffered downlink packets.
22	After the MME receives Update Bearer Response (EPS Bearer Identity) message, if an EPS bearer was established and the subscription data indicates that the user is allowed to perform handover to non-3GPP accesses, and if the MME selected a P-GW that is different from the P-GW address which was indicated by the HSS in the PDN subscription context, the MME sends an Update Location Request including the APN and P-GW address to the HSS for mobility with non-3GPP accesses.
23	The HSS stores the APN and P-GW address pair and sends an Update Location Response to the MME.
24	Bidirectional data is passed between the UE and PDN.

Subscriber-initiated Detach

This section describes the procedure of detachment from the EPC network by a subscriber.

Figure 232. Subscriber-initiated Detach Call Flow

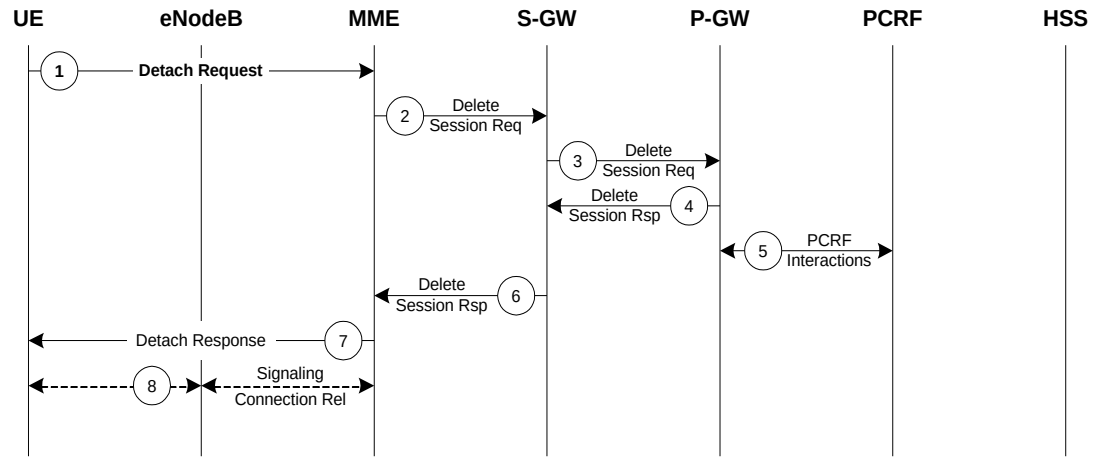


Table 105. Subscriber-initiated Detach Call Flow Description

Step	Description
1	The UE sends NAS message Detach Request (GUTI, Switch Off) to the MME. Switch Off indicates whether detach is due to a switch off situation or not.
2	The active EPS Bearers in the S-GW regarding this particular UE are deactivated by the MME sending a Delete Bearer Request (TEID) message to the S-GW.
3	The S-GW sends a Delete Bearer Request (TEID) message to the P-GW.

Step	Description
4	The P-GW acknowledges with a Delete Bearer Response (TEID) message.
5	The P-GW may interact with the PCRF to indicate to the PCRF that EPS Bearer is released if PCRF is applied in the network.
6	The S-GW acknowledges with a Delete Bearer Response (TEID) message.
7	If Switch Off indicates that the detach is not due to a switch off situation, the MME sends a Detach Accept message to the UE.
8	The MME releases the S1-MME signalling connection for the UE by sending an S1 Release command to the eNodeB with Cause = Detach.

Supported Standards

The S-GW service complies with the following standards.

- 3GPP References
- 3GPP2 References
- IETF References
- Object Management Group (OMG) Standards

3GPP References

Release 8 Supported Standards

- 3GPP TR 21.905: Vocabulary for 3GPP Specifications
- 3GPP TS 23.003: Numbering, addressing and identification
- 3GPP TS 23.007: Restoration procedures
- 3GPP TS 23.107: Quality of Service (QoS) concept and architecture
- 3GPP TS 23.203: Policy and charging control architecture
- 3GPP TS 23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access
- 3GPP TS 23.402: Architecture Enhancements for non-3GPP accesses
- 3GPP TS 23.060. General Packet Radio Service (GPRS); Service description; Stage 2
- 3GPP TS 24.008: Mobile radio interface Layer 3 specification; Core network protocols
- 3GPP TS 24.229: IP Multimedia Call Control Protocol based on SIP and SDP; Stage 3
- 3GPP TS 29.210. Gx application
- 3GPP TS 29.212: Policy and Charging Control over Gx reference point
- 3GPP TS 29.213: Policy and Charging Control signaling flows and QoS
- 3GPP TS 29.214: Policy and Charging Control over Rx reference point
- 3GPP TS 29.274 V8.1.1 (2009-03): 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C); Stage 3 (Release 8)
- 3GPP TS 29.274: Evolved GPRS Tunnelling Protocol for Control plane (GTPv2-C), version 8.2.0 (both versions are intentional)
- 3GPP TS 29.275: Proxy Mobile IPv6 (PMIPv6) based Mobility and Tunnelling protocols, version 8.1.0
- 3GPP TS 29.281: GPRS Tunnelling Protocol User Plane (GTPv1-U)
- 3GPP TS 32.251: Telecommunication management; Charging management; Packet Switched (PS) domain charging
- 3GPP TS 32.295: Charging management; Charging Data Record (CDR) transfer

- 3GPP TS 32.298: Telecommunication management; Charging management; Charging Data Record (CDR) encoding rules description
- 3GPP TS 32.299: Charging management; Diameter charging applications
- 3GPP TS 33.106: 3G Security; Lawful Interception Requirements
- 3GPP TS 36.107: 3G security; Lawful interception architecture and functions
- 3GPP TS 36.300: Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description
- 3GPP TS 36.412. EUTRAN S1 signaling transport
- 3GPP TS 36.413: Evolved Universal Terrestrial Radio Access (E-UTRA); S1 Application Protocol (S1AP)
- 3GPP TS 36.414: Evolved Universal Terrestrial Radio Access Network (E-UTRAN); S1 data transport

3GPP2 References

- X.P0057-0 v0.11.0 E-UTRAN - eHRPD Connectivity and Interworking: Core Network Aspects

IETF References

- RFC 768: User Datagram Protocol (STD 6).
- RFC 791: Internet Protocol (STD 5).
- RFC 2131: Dynamic Host Configuration Protocol
- RFC 2460: Internet Protocol, Version 6 (IPv6) Specification
- RFC 2698: A Two Rate Three Color Marker
- RFC 2784: Generic Routing Encapsulation (GRE)
- RFC 2890: Key and Sequence Number Extensions to GRE
- RFC 3319: Dynamic Host Configuration Protocol (DHCPv6) Options for Session Initiation Protocol (SIP) Servers
- RFC 3588: Diameter Base Protocol
- RFC 3775: Mobility Support in IPv6
- RFC 3646: DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
- RFC 4006: Diameter Credit-Control Application
- RFC 4282: The Network Access Identifier
- RFC 4283: Mobile Node Identifier Option for Mobile IPv6 (MIPv6)
- RFC 4861: Neighbor Discovery for IP Version 6 (IPv6)
- RFC 4862: IPv6 Stateless Address Autoconfiguration
- RFC 5094: Mobile IPv6 Vendor Specific Option
- RFC 5213: Proxy Mobile IPv6
- Internet-Draft: Proxy Mobile IPv6
- Internet-Draft: GRE Key Option for Proxy Mobile IPv6, work in progress

- Internet-Draft: Binding Revocation for IPv6 Mobility, work in progress

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

Chapter 30
Session Control Manager Overview

This chapter contains general overview information about the Session Control Manager (SCM) including:

- Product Description
- Network Deployments and Interfaces
- Features and Functionality - Base Software
- Features and Functionality - Licensed Enhanced Feature Support
- How the SCM Works
- Supported Standards

Product Description

The Session Control Manager (SCM) delivers and controls a robust multimedia environment today, while preparing for the networks of tomorrow. SCM provides an easy on-ramp to deploying Session Initiation Protocol (SIP)-based services and a future-proof migration path to the IP Multimedia Subsystem/Multimedia Domain (IMS/MMD) architectures.

The SCM performs the following functions:

- SIP routing
- Translation and mobility
- Admission control
- Authentication
- Registration
- Emergency Registration
- Packet network access based on pre-established policies and procedures
- Localized policy selection and enforcement
- Multimedia Call Detail Records (CDRs)
- Per-subscriber service facilitation
- SIP Application-level Gateway (ALG)
- Media relay
- Mitigate SIP Denial of Service (DoS)
- Prevent registration hijacking
- Prevent theft of service

The SCM consists of multiple IMS components that can be integrated into a single ASR 5000 platform or distributed as standalone network elements:

- IETF-compliant SIP Proxy/Registrar
- 3GPP/3GPP2-compliant Proxy Call/Session Control Function (P-CSCF)
- 3GPP/3GPP2-compliant Serving Call/Session Control Function (S-CSCF)
    - 3GPP/3GPP2-compliant Interrogating Call/Session Control Function (I-CSCF)
    - 3GPP/3GPP2 Breakout Gateway Control Function (BGCF)
- 3GPP/3GPP2-compliant Emergency Call/Session Control Function (E-CSCF)
- 3GPP/IETF-compliant Access Border Gateway (A-BG)

As standards-based network elements, SCM components can be integrated with each other or with third-party IMS components.

IMS Architecture

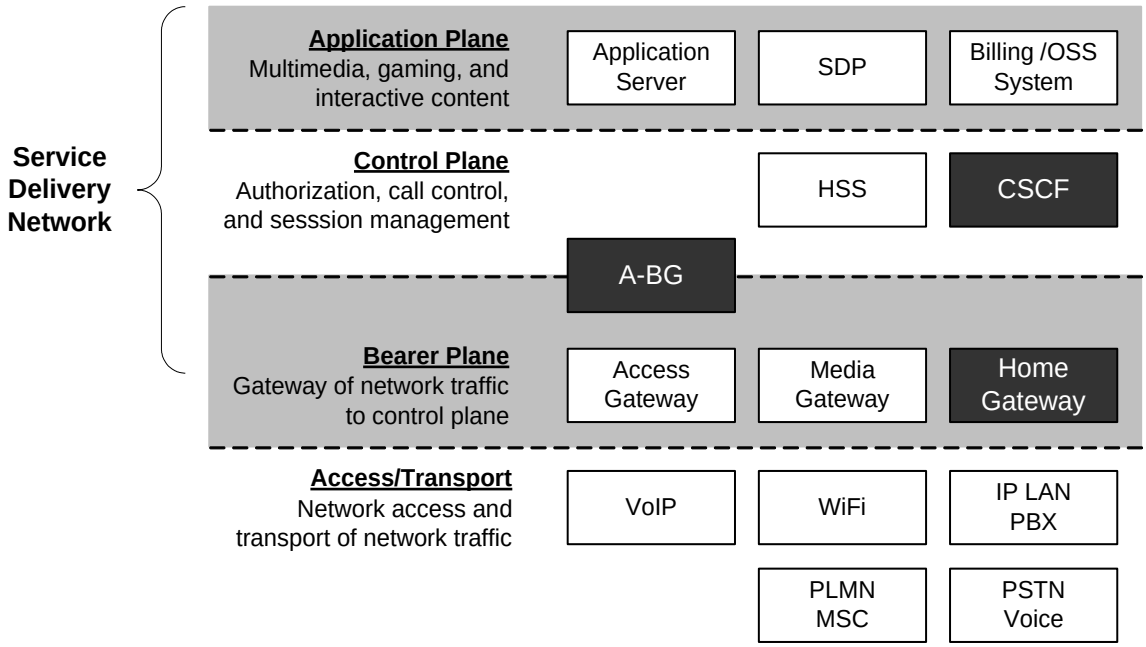
IP Multimedia Subsystem (IMS) specifies a standard architecture for providing combined IP services (voice, data, multimedia) over the existing public switched domain. IMS is an integral part of the 3GPP, 3GPP2, ETSI, and TISPAN network model standards that define circuit switched, packet switched, and IP multimedia domain environments. IMS

also supports multiple access methods such as CDMA2000, DOCSIS, EPS, Ethernet, Fiber, GPRS, WCDMA, WLAN, XDSL, and wireless broadband access.

The call signaling protocol used in IMS is the Session Initiation Protocol (SIP). The primary component in the network for resolving and forwarding SIP messages is the Call/Session Control Function (CSCF). The CSCF provides the control and routing function for all IP sessions accessing the network. CSCFs are located in the control plane or layer of the Service Delivery Network as shown in the figure below.

When the SCM acts as an Access Border Gateway (A-BG), it uses the RFC3261/P-CSCF to provide a SIP/IMS control plane access border, as well as a bearer access border control function. Therefore, the A-BG provides all session border control functions for all SIP UEs attempting to access the mobile network from a network outside of the operator's control and operations.

Figure 233. IMS Service Delivery Networks Components

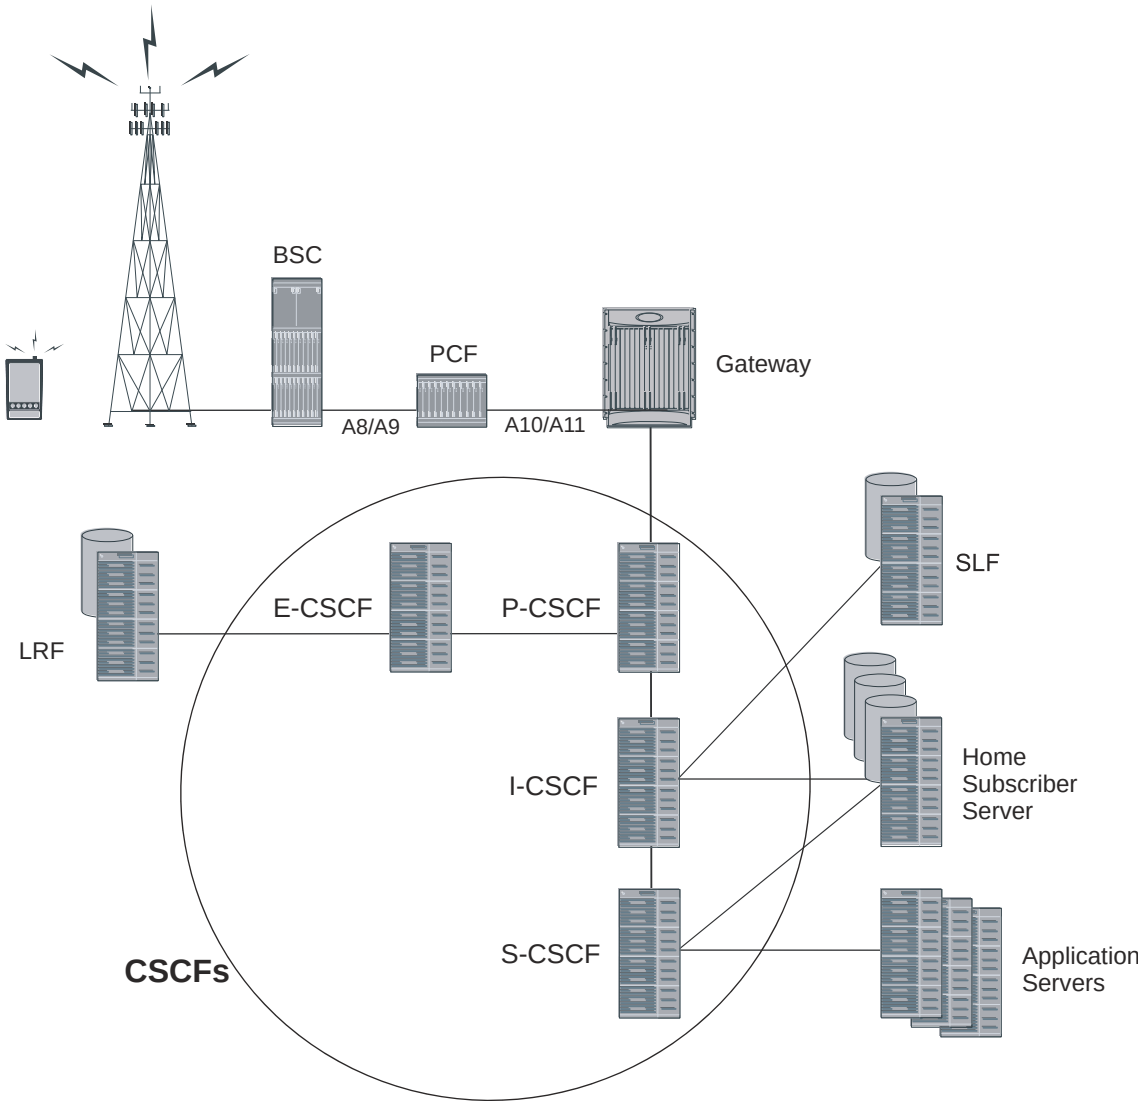


Collectively, CSCFs are responsible for managing an IMS session, including generating Call Detail Records (CDRs). Four functional behaviors are defined for the CSCF:

- Proxy
- Interrogating
- Serving
- Emergency

The following figure shows the general interaction between the CSCF components and the supporting servers.

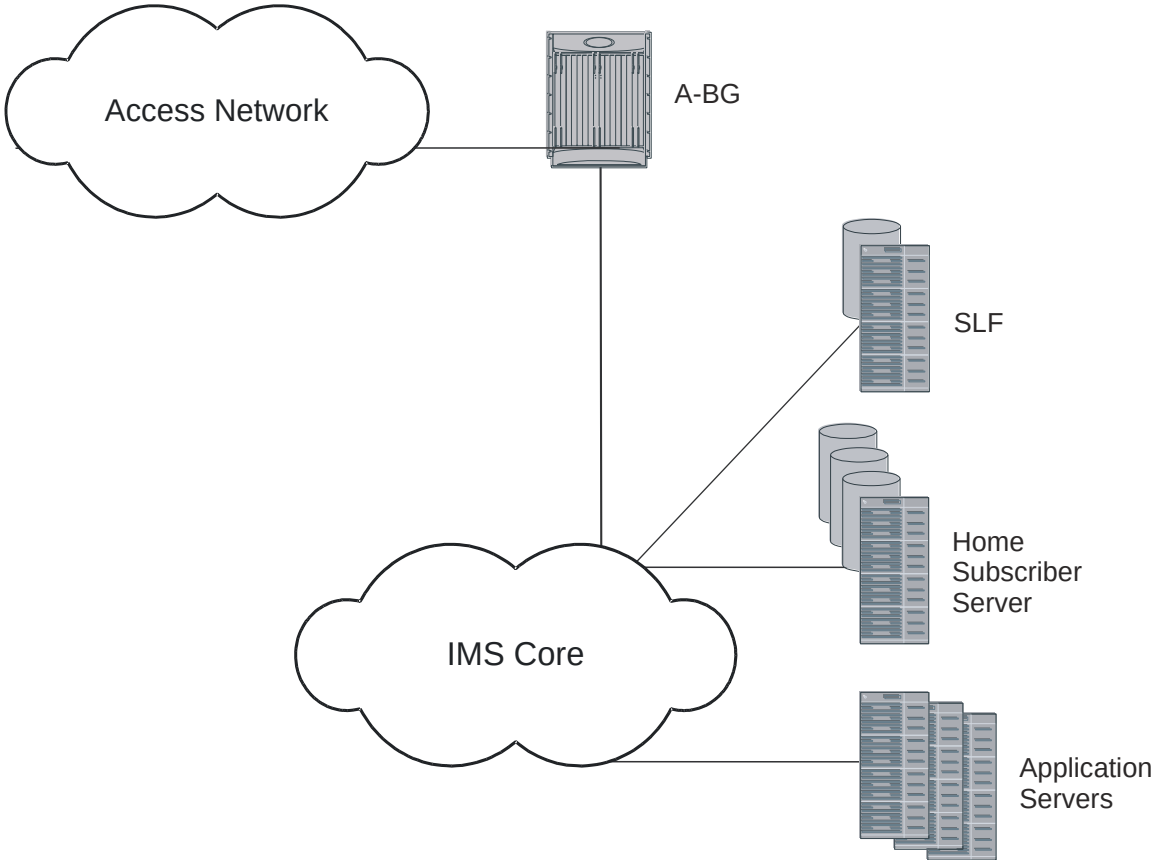
Figure 234. IMS CSCF Components



In addition, the SCM may act as an Access Border Gateway (A-BG).

The following figure shows the general interaction between the A-BG and the supporting servers.

Figure 235. Access Border Gateway



Proxy-CSCF

The primary point of entry into the IMS network is the Proxy-CSCF (P-CSCF). The P-CSCF is responsible for:

- providing message manipulation to allow for localized services (traffic/weather reports, news, directory services, etc.)
- initiating the breakout of emergency service calls
- Topology Hiding Inter-network Gateway (THIG)
- Quality of Service (QoS) authorization
- number conversions for local dialing plans
- terminating IPSec tunnels
- Transport Layer Security (TLS)
- interworking
- Signaling Compression/Decompression (SIGCOMP)
- charging

The P-CSCF is the handset's first point of entry into the IMS and is also the outbound proxy for SIP. Once the P-CSCF has completed all of the functions for which it is responsible, the call setup is handed off to the Interrogating-CSCF (I-CSCF).

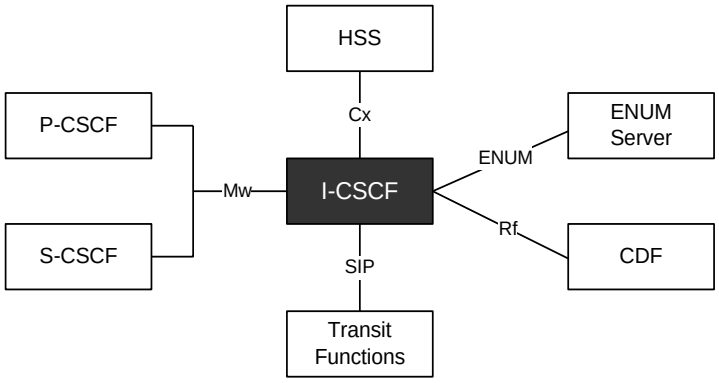
Interrogating-CSCF

The I-CSCF performs mostly as a load distribution device. The I-CSCF queries the Home Subscriber Server (HSS) to identify the appropriate Serving-CSCF (S-CSCF) to which the call is sent. Since the HSS maintains user profile information (much like the Home Location Register (HLR) in the Public Land Mobile Network (PLMN)), the I-CSCF can identify the proper S-CSCF for the call. The I-CSCF may also query a AAA server to determine subscriber profile information using DIAMETER.

Important: The I-CSCF is incorporated into the S-CSCF.

I-CSCF Interfaces

The following diagram shows the interfaces/reference points associated with the I-CSCF:



Serving-CSCF

The Serving-CSCF (S-CSCF) is the access point to services provided to the subscriber. Service examples include session control services, such as call features.

Other services include:

- VPN
- Centralized speed dialing lists
- Charging

The S-CSCF also interacts with the HSS for:

- User authentication
- Subscriber profile download and provisioning filter rules for services
- Network authentication key
- Emergency registration
- Location management
- User data handling

A Breakout Gateway Control Function is integrated into the SCM's S-CSCF to support PSTN calls.

Telephony Application Server (TAS) Basic Supported

The following describe the local basic call features implemented on the S-CSCF:

- Abbreviated Dialing (AD)
- Call Forward Busy Line (CFBL)
- Call Forward No Answer (CFNA)
- Call Forward Not Registered (CFNR)
- Call Forward Unconditional (CFU)
- Call Transfer
- Call Waiting
- Caller ID Display (CID)
- Caller ID Display Blocked (CIDB)
- Feature Code Activation/De-activation
- Follow Me/Find Me
- Locally Allowed Abbreviated Dialing
- Outbound Call Restrictions/Dialing Permissions
- Short Code Dialing

Integrated S/I-CSCF

The following Interrogating-CSCF features are supported for the integrated S/I-CSCF:

- **Assign an S-CSCF to a User Performing SIP Registration** - On a UE registration, the I-CSCF carries out a first step authorization and S-CSCF discovery. For this, the I-CSCF sends a Cx User-Authentication-Request (UAR) to the HSS by transferring the Public and Private User Identities and the visited network identifier (all extracted from the UE REGISTER message). The HSS answers with a Cx User-Authentication-Answer (UAA). The UAA includes the URI of the S-CSCF already allocated to the user. If there is no previously allocated S-CSCF, the HSS returns a set of S-CSCF capabilities that the I-CSCF uses to select the S-CSCF.
- **E.164 Address Translation** - Translates the E.164 address contained in all Request-URIs having the SIP URI with user=phone parameter format into the Tel: URI format before performing the HSS Location Query. In the event the user does not exist, and if configured by operator policy, the I-CSCF may invoke the portion of the transit functionality that translates the E.164 address contained in the Request-URI of the Tel: URI format to a routable SIP URI.
- **Obtain the S-CSCF Address from the HSS** - When the I-CSCF receives a SIP request from another network, it has to route the request to the called party. For this it obtains the S-CSCF address associated with the called party from the HSS by querying with a Cx Location-Information-Request (LIR) message. The Public-Identity AVP in the LIR is the Request-URI of the SIP request. The Location-Information-Answer (LIA) message contains the S-CSCF address in the Server-Name AVP. The request is then routed to the S-CSCF.
- **Route a SIP Request or Forward Response from Another Network** - When the I-CSCF receives a request from another network, it obtains the address of the S-CSCF from the HSS using the procedure detailed above and routes the request to the S-CSCF. Responses are also routed to the S-CSCF.
- **Perform Transit Routing Functions** - The I-CSCF may need to perform transit routing if, based on the HSS query, the destination of the session is not within the IMS. The IMS Transit Functions perform an analysis of the destination address and determine where to route the session. The session may be routed directly to an MGCF, BGCF, or to another IMS entity in the same network, to another IMS network, or to a CS domain or PSTN.

- **Generate CDRs** - The I-CSCF generates CDRs for its interactions. Upon completing a Cx query, the I-CSCF sends an Accounting Request with the Accounting-Record-Type set to EVENT. The CDF acknowledges the data received and creates an I-CSCF CDR.

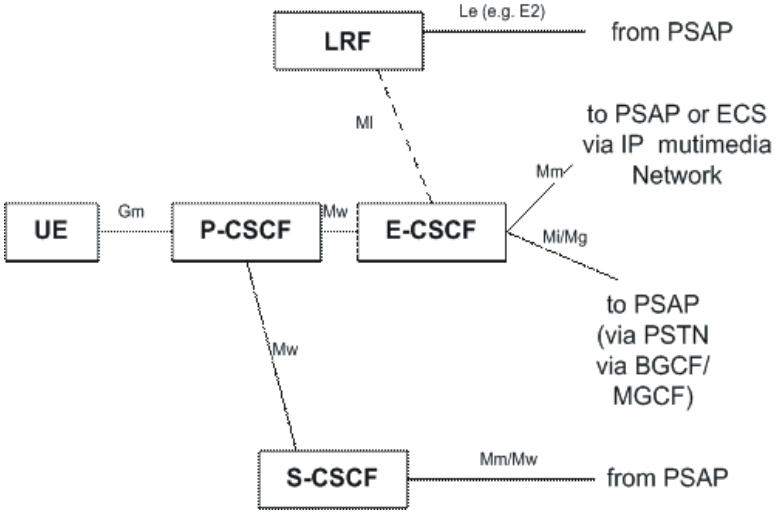
Emergency-CSCF

The Emergency-CSCF (E-CSCF) is a network element in IMS which is responsible for routing an emergency call to a Public Safety Answering Point (PSAP).

To identify the next hop PSAP, E-CSCF interacts with the Location Retrieval Function (LRF). LRF provides the necessary routing information so that E-CSCF can route the request to the appropriate PSAP.

E-CSCF Interfaces

The following diagram shows the interfaces/reference points associated with the E-CSCF:



A-BG

The A-BG is responsible for:

- Border Control for both Signaling and Bearer
- CALEA Support
    - SIP and media taps
- Call Admission and Access Control
    - Access Control based on IP, URL, SIP Identity, and Session Limits
- Intelligent Routing
    - Least Cost, Congestion Based, Call Type, Domain Based
    - As a SIP ALG, supports signaling and media routing with overlapping address ranges
- SIP Application-level Gateway (SIP-ALG)
    - SIP NAT Traversal
    - SIP NAT (IPv4 <–> IPv6 translation)

- Media Relay (Header Manipulation): RTP, MSRP
- SIP Security
  - Prevent Theft of Service
    - Prevent CSCF bypass
    - Robust authentication procedures
    - SIP message checking
  - Prevent Registration Hijacking
    - Authenticate Re-Register (S-CSCF)
    - Early IMS Security: DoS attack prevention, impersonating a server
    - UA authentication (prevent server impersonation)
    - AKA authentication mechanism (further protection)
  - Prevent Message Tampering (IPSec)
  - Prevent Early Session Tear Down
    - Early IMS Security prevents a different user releasing existing session
  - Mitigate SIP Denial of Service (DoS)
    - P-CSCF DoS Attack Prevention
    - Blocking of user/IP address
      - after repeated authentication and bad request failure in Register/INVITE
    - Dropping of Register
      - containing Contact header pointing to CSCF service ip:port
    - Limited number of contacts on which Forking is allowed
    - Dropping of Requests
      - coming from source address other than the Register request's source address
- Topology Hiding Inter-network Gateway (THIG)
- Transport Layer Security (TLS)

Technical Specifications

The following table provides product specifications for the SCM.

Table 106. Session Control Manager Technical Specifications

	Description
Service Instances	Dual-mode proxy: simultaneously supports IETF & 3GPP/3GPP2 Proxies

	Description
SIP	• IETF SIP Proxy/Registrar • 3GPP/3GPP2 Proxy Call Session Control Function (P-CSCF) • Stateful session and subscriber aware control • Signaling Compression/Decompression (SIGCOMP) • Auto discovery, subscriber privacy, network security, call fraud prevention, thwarting network overload conditions
SIP Message Handling	Forking, error handling and discard, header stripping and insertion, multiple public user identities
Logical Interfaces	• IETF: SIP Proxy/Registrar • 3GPP: Mw, Gm, Rx, Rf, Cx, Sh, Dx, MI • 3GPP2: Mw, Gm, Tx, Rf, Cx, Sh, Dx, MI

Platform Requirements

The SCM service runs on a Cisco® ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the Installation Guide for the chassis and/or contact your Cisco account representative.

Licenses

The SCM is a licensed Cisco product. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

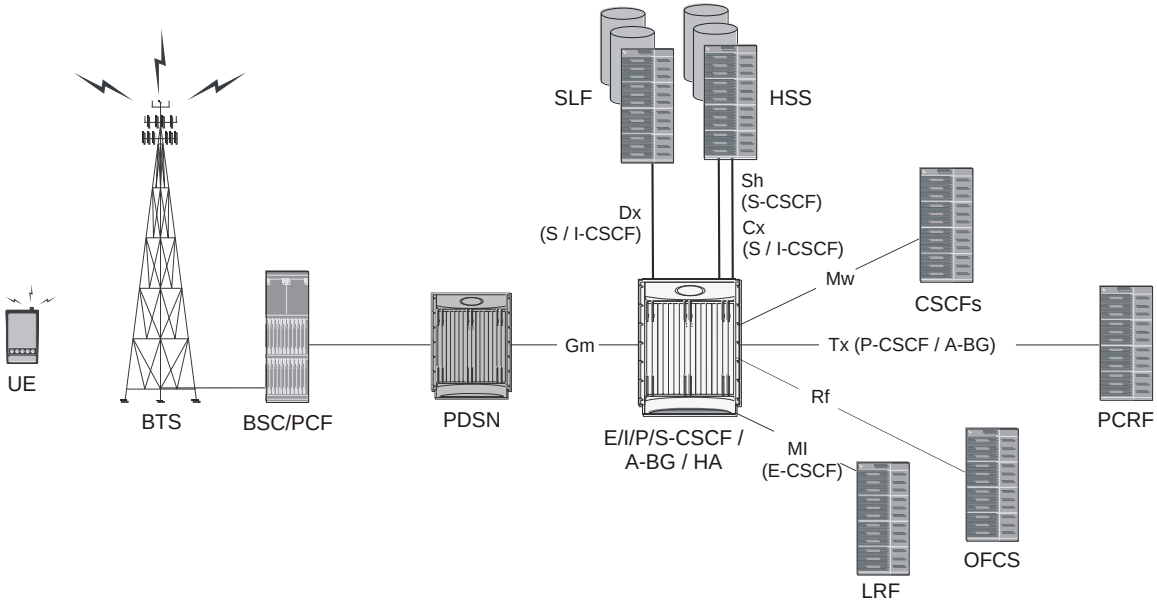
Network Deployments and Interfaces

SCM in a CDMA2000 Data Network Deployment

Integrated CSCF / A-BG / HA

The SCM is designed to function within a CDMA2000 PDSN network. By combining the SCM with a carrier-class Home Agent, a number of advantages emerge such as increased performance, distributed architecture, and high availability. As shown in the figure below, the SCM supports a number of interfaces used to communicate with other components in an IMS environment and supports the interface used to bridge the CDMA network.

Figure 236. CDMA2000 CSCF/A-BG/HA SCM Deployment Example



Logical Network Interfaces (Reference Points)

Interfaces, used to support IMS in a CDMA network, can be defined within two categories: SIP and DIAMETER. The SCM incorporates standards-based interfaces for both SIP and DIAMETER network architectures.

SIP Interfaces

The following table provides descriptions of SIP interfaces supported by the SCM in a CDMA2000 network deployment.

Table 107.SIP Interfaces in a CDMA Network

Interface	Description

Interface	Description
Gm	The reference point between the P-CSCF and the User Equipment (UE).The Gm interface provides SIP signaling between the PDSN and the P-CSCF.
MI	The reference point between the E-CSCF and Location Retrieval Function (LRF). The MI interface is used for routing an emergency call to a Public Safety Answering Point (PSAP). The E-CSCF interacts with the Location Retrieval Function (LRF) to identify the next hop PSAP.
Mw	The reference point between the P/S-CSCF and other CSCFs.The Mw interface provides SIP signaling between two CSCFs.

DIAMETER Interfaces

The following table provides descriptions of DIAMETER interfaces supported by the SCM in a CDMA2000 network deployment.

Table 108.DIAMETER Interfaces in a CDMA Network

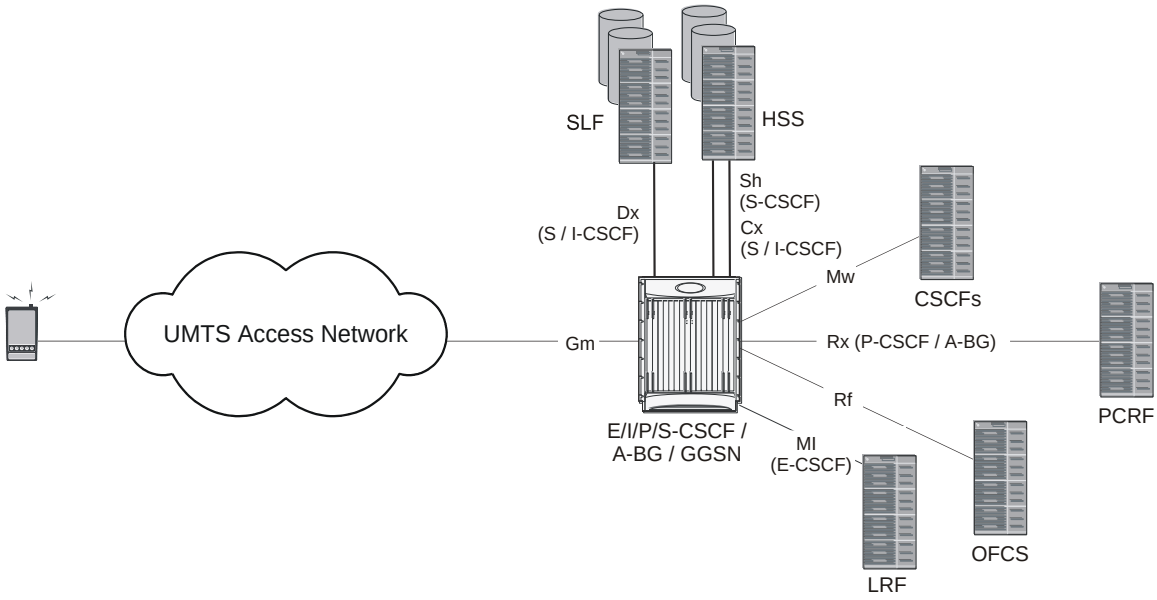
Interface	Description
Cx	The reference point between the S/I-CSCF and the Home Subscriber Server (HSS). The Cx interface is used to authenticate subscribers, provides server assignments, push user profile information from the HSS to the S-CSCF, and, when necessary, transmit a network initiated de-registration.
Dx	The reference point between the S/I-CSCF and Subscriber Location Function (SLF). The Dx interface is used to proxy queries to a subscriber data server (such as an HSS) in which subscription data for a user can be found. The SLF receives a query for the subscriber data server, looks up the address of appropriate subscriber data server, and proxies the query to the appropriate subscriber data server.
Rf	The reference point between the P-CSCF and the Offline Charging System (OFCS). The Rf interface is used to transfer charging information that will not affect, in real-time, the service being rendered. For more information, refer to the 3GPP2 specification X.S0013-007-A v1.0.
Sh	The reference point between the S-CSCF and Home Subscriber Server (HSS). The Sh interface is used for retrieval and update of call feature data parameters.
Tx	The reference point between the P-CSCF/A-BG and the Charging Rule Function (CRF)/Policy Decision Point (PDP) (PCRF) used for Service Based Bearer Control (SBBC). It identifies any P-CSCF/A-BG restrictions to be applied to the identified packet flows.

SCM in a GSM/UMTS Data Network Deployment

CSCF / A-BG / GGSN Deployment

The SCM is designed to function within a UMTS GGSN network. As shown in following figure, the SCM supports a number of interfaces used to communicate with other components in an IMS environment and supports the interface used to bridge the GGSN network.

Figure 237. GSM/UMTS CSCF/A-BG/GGSN SCM Deployment Example



Logical Network Interfaces (Reference Points)

Interfaces, used to support IMS in a UMTS network, can be defined within two categories: SIP and DIAMETER. The SCM incorporates standards-based interfaces for both SIP and DIAMETER network architectures.

SIP Interfaces

The following table provides descriptions of SIP interfaces supported by the SCM in a GSM/UMTS network deployment.

Table 109. SIP Interfaces in a GSM/UMTS Network

Interface	Description
Gm	The reference point between the P-CSCF and the User Equipment (UE). The Gm interface provides SIP signaling between the GGSN and the P-CSCF.
MI	The reference point between the E-CSCF and Location Retrieval Function (LRF). The MI interface is used for routing an emergency call to a Public Safety Answering Point (PSAP). The E-CSCF interacts with the Location Retrieval Function (LRF) to identify the next hop PSAP.
Mw	The reference point between the P/S-CSCF and other CSCFs. The Mw interface provides SIP signaling between two CSCFs.

DIAMETER Interfaces

The following table provides descriptions of DIAMETER interfaces supported by the SCM in a GSM/UMTS network deployment.

Table 110.DIAMETER Interfaces in a GSM/UMTS Network

Interface	Description
Cx	The reference point between the S/I-CSCF and the Home Subscriber Server (HSS). The Cx interface is used to authenticate subscribers, provides server assignments, push user profile information from the HSS to the S-CSCF, and, when necessary, transmit a network initiated de-registration.
Dx	The reference point between the S/I-CSCF and Subscriber Location Function (SLF). The Dx interface is used to proxy queries to a subscriber data server (such as an HSS) in which subscription data for a user can be found. The SLF receives a query for the subscriber data server, looks up the address of appropriate subscriber data server, and proxies the query to the appropriate subscriber data server.
Rf	The reference point between the P-CSCF and the Offline Charging System (OFCS). The Rf interface is used to transfer charging information that will not affect, in real-time, the service being rendered. For more information, refer to the 3GPP2 specification X.S0013-007-A v1.0.
Rx	The reference point between the P-CSCF/A-BG and the Charging Rule Function (CRF)/Policy Decision Point (PDP) (PCRF). The Rx interface (3GPP 29.211) is used to exchange Flow Based Charging (FBC) control information between the PCRF and the P-CSCF/A-BG. The CRF uses the information to make FBC decisions that are then exchanged with the Traffic Plane Function (TPF). This interface is used in a 3GPP2 Release 7 implementation.
Sh	The reference point between the S-CSCF and Home Subscriber Server (HSS). The Sh interface is used for retrieval and update of call feature data parameters.

Features and Functionality - Base Software

The following is a list containing a variety of features found in the SCM and the benefits they provide.

This section describes the following features:

- Bulk Statistics Support
- Call Abort Handling
- Call Forking
- Call Types Supported
- Congestion Control
- DSCP Marking
- Early IMS Security
- Emergency Call Support
- Error Handling
- Future-proof Solution
- Intelligent Integration
- Interworking Function
- IPv6 Support
- Management System Overview
- MSRP Support
- PCRF Policy Control
- Presence Enabled
- Redirection
- Redundancy and Session Recovery
- Registration Event Package
- Signaling Compression (SigComp)
- SIP Denial of Service (DoS) Attack Prevention
- SIP Intelligence at the Core
- SIP Large Message Support
- SIP Routing Engine
- Shared Initial Filter Criteria (SiFC)
- Telephony Application Server (TAS) Basic Supported
- Threshold Crossing Alerts (TCA) Support
- Trust Domain

Bulk Statistics Support

The system's support for bulk statistics allows operators to choose to view not only statistics that are of importance to them, but also to configure the format in which it is presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Following is a list of supported schemas for SCM:

- **Card**: Provides card-level statistics
- **Context**: Provides context-level statistics
- **CSCF**: Provides CSCF service statistics
- **CSCFINTF**: Provides CSCF interface statistics
- **Diameter-acct**: Provides Diameter Accounting statistics
- **Diameter-auth**: Provides Diameter Authentication statistics
- **Map**: Provides Map service statistics
- **Nat-realm**: Provides NAT realm statistics
- **Port**: Provides port-level statistics
- **System**: Provides system-level statistics

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the system or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, system host name, system uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

Important: For more information on bulk statistic configuration, refer to the *Configuring and Maintaining Bulk Statistics* chapter in the *System Administration Guide*.

Call Abort Handling

Call abort handling provides resource cleanup in error scenarios and makes sure resources that are not being used can be used for new calls. This feature is managed gracefully for a P-CSCF failure and CLI-initiated subscriber and session clean up.

Call Forking

Call forking allows subscribers to receive calls wherever they are by enabling multi-location UE registration.

Call Types Supported

In the IMS architecture, telephony features are normally provided by an external application server. Providing these features with the S-CSCF:

- Reduces the need for an additional SIP AS
- Simplifies the network architecture
- Improves latency for call setup and feature invocation

The following call types are supported:

- **Directory service, toll-free, long distance, international, and operator-assisted calls** - are supported through translation lists.
- **Emergency calls -** are managed through the addition of an Emergency Call/Session Control Function (E-CSCF) that routes emergency calls to a Public Safety Answering Point (PSAP).
- **Mobile-to-Mobile SIP calls** - supports SIP-based VoIP calls between mobile data users.
- **Public Switched Telephone Network (PSTN) calls** - can be routed through a 3GPP/2 compliant BGCF located in the S-CSCF.

Congestion Control

The congestion control feature allows you to set policies and thresholds and specify how the system reacts when faced with a heavy load condition.

Congestion control monitors the system for conditions that could potentially degrade performance when the system is under heavy load. Typically, these conditions are temporary (for example, high CPU or memory utilization) and are quickly resolved. However, continuous or large numbers of these conditions within a specific time interval may have an impact the system's ability to service subscriber sessions. Congestion control helps identify such conditions and invokes policies for addressing the situation.

Congestion control operation is based on configuring the following:

- **Congestion Condition Thresholds:** Thresholds dictate the conditions for which congestion control is enabled and establishes limits for defining the state of the system (congested or clear). These thresholds function in a way similar to operation thresholds that are configured for the system as described in the Thresholding Configuration Guide. The primary difference is that when congestion thresholds are reached, a service congestion policy and an SNMP trap, starCongestion, are generated.

  A threshold tolerance dictates the percentage under the configured threshold that must be reached in order for the condition to be cleared. An SNMP trap, starCongestionClear, is then triggered.

- **Port Utilization Thresholds**: If you set a port utilization threshold, when the average utilization of all ports in the system reaches the specified threshold, congestion control is enabled.
    - **Port-specific Thresholds**: If you set port-specific thresholds, when any individual port-specific threshold is reached, congestion control is enabled system-wide.
- **Service Congestion Policies**: Congestion policies are configurable for each service. These policies dictate how services respond when the system detects that a congestion condition threshold has been crossed.

CSCF performs congestion control based on the memory usage inside every sessmgr at two levels.

Level 1: For every new call/event received, the system checks if sessmgr memory-usage is above a threshold value (such as 95 percent). If it is, memory-congestion is triggered and new call messages are rejected with 500 SIP response. Memory congestion is disabled when memory usage drops by a tolerance value (default is 10 percent).

Level 2: If the sessmgr usage reaches 100 percent, all newly received SIP messages are dropped at the socket layer in that sessmgr except for the BYE message. The new SIP messages are not processed until the memory reaches the threshold value (95 percent).

A trap is also generated whenever sessmgr is in congestion state

Important: For more information on congestion control, refer to the *Congestion Control* chapter in the *System Administration Guide*.

DSCP Marking

Provides support for more granular configuration of DSCP marking.

For Interactive Traffic class, the P-CSCF/A-BG supports per-service configurable DSCP marking for Uplink and Downlink direction based on Allocation/Retention Priority in addition to the current priorities.

The following matrix may be used to determine the Diffserv markings used based on the configured traffic class and Allocation/Retention Priority:

Table 111. Default DSCP Value Matrix

Allocation Priority	1	2	3
Traffic Handling Priority			
1	ef	ef	ef
2	af21	af21	af21
3	af21	af21	af21

Early IMS Security

Early IMS security allows authenticating the UE without IMS protocols and clients. Based on the 3GPP TR 33.978 specification, the SCM supports security inter-operation with 2G and non-IPSec user devices.

Emergency Call Support

P-CSCF gives priority to emergency calls, especially in a congested network. In addition, P-CSCF rejects new calls to any user who is in an emergency call.

Error Handling

The SCM supports consistent management of errors in a framework that considers existing and future standards and specifications.

Future-proof Solution

The SCM eliminates the capital and operational barriers associated with deploying traditional, server-based SIP proxies that lack carrier-class characteristics, occupy valuable rack space, and require numerous network interfaces, while also introducing additional control hops in the network that add call setup latency.

When operators deploy IMS/MMD, profitability will improve because a seamless on-ramp will be provided by simultaneously supporting 3GPP/3GPP2-based standards, P-CSCF functionality, and IETF SIP standards.

Intelligent Integration

For deployed platforms, no new hardware is necessary to install or manage. Functionality is enabled with a simple software download.

Intelligent integration lowers operational expenditure and reduces the number of network elements, network interfaces, and call setup latency.

Interworking Function

The SCM allows non-IMS UEs (pre IMS or RFC3261-compliant UEs) to work with the IMS core. When UEs are not IMS compliant, having this protocol interworking function at the edge allows the IMS core to be IMS compliant. After the interworking function inserts all necessary IMS headers toward the IMs core, the call appears to the IMS core network elements as if it is coming from an IMS-compliant UE.

The feature allows simultaneous support of IETF SIP and 3GPP/3GPP2 IMS/MMD clients.

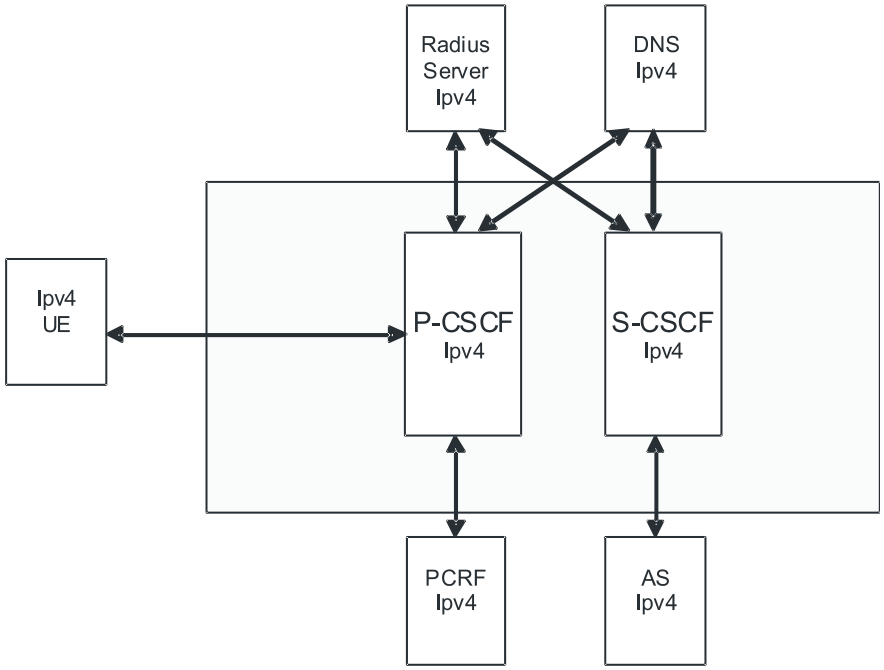
IPv6 Support

In addition to supporting IPv4, the SCM supports IPv6 addressing. A CSCF service can be configured with v6 addresses to support an all v6 network.

Important: For this feature, you may bind a CSCF service to either an IPv4 address or to an IPv6 address, but not both simultaneously.

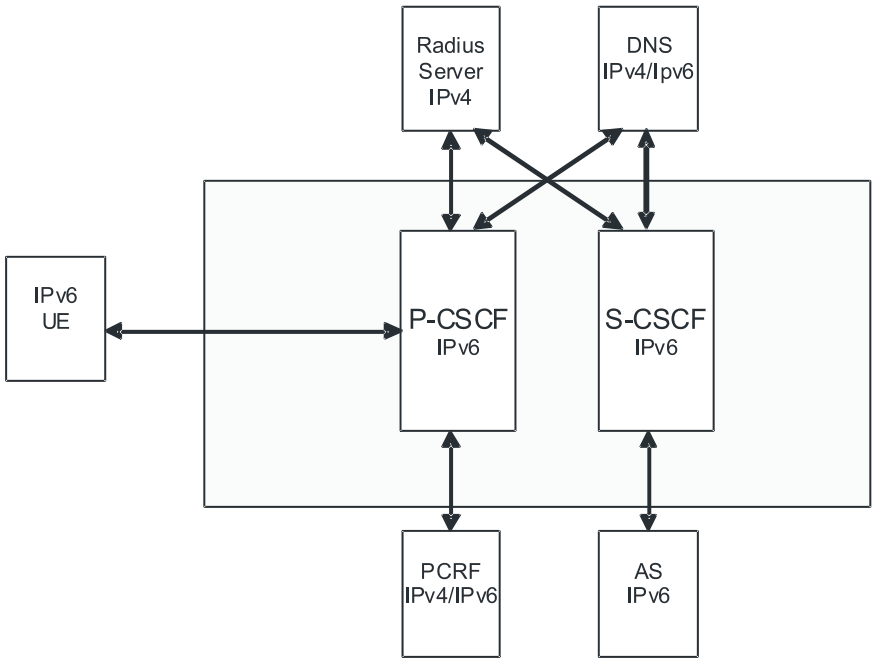
The following diagram shows the implementation where CSCF supports only IPv4.

Figure 238. IPv4 Configuration



With IPv6 support, the configuration supported would look like the following diagram. The DNS server could be either IPv4 or IPv6.

Figure 239. IPv6 Configuration



Important: The policy interface to PCRF will be IPv6 based when DIAMETER supports IPv6.

Management System Overview

The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management - focusing on providing superior quality network element (NE) and element management system (Cisco Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems - giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

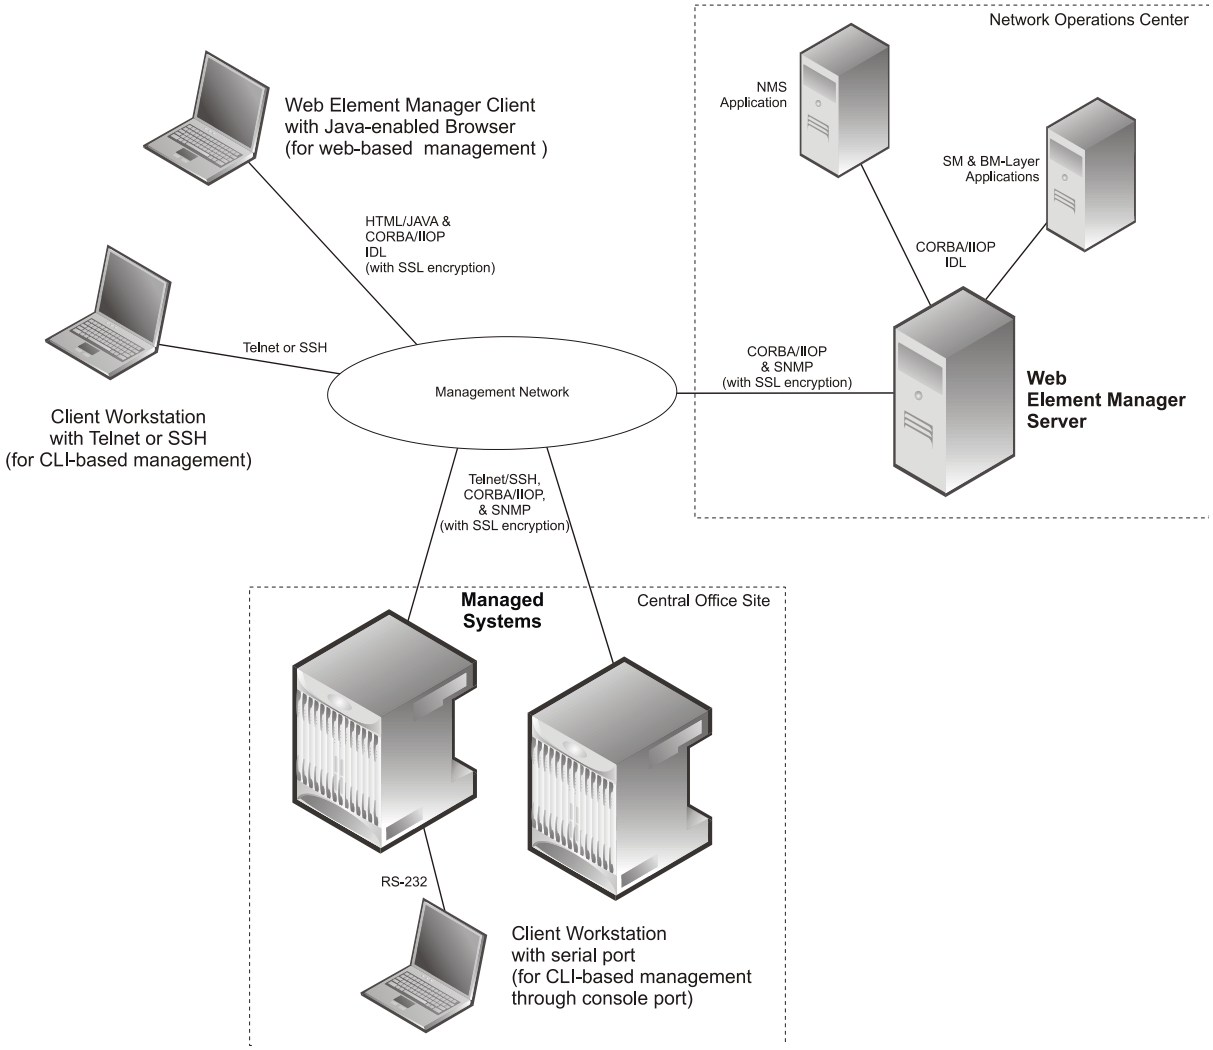
Cisco's O&M module offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces.

These include:

- Using the command line interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager application
- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000
- Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (i.e. Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network.

Figure 240. Element Management Methods



Important: SCM management functionality is enabled by default for console-based access. For GUI-based management support, refer to the *Web Element Management System* section in this chapter.

Important: For more information on command line interface based management, refer to the *Command Line Interface Reference*.

MSRP Support

The SCM supports Message Session Relay Protocol (MSRP) session and page modes.

PCRF Policy Control

P-CSCF can use the IMS Rx interface to communicate with the PCRF during call initiation and renegotiation to provide session information to the PCRF and ensure that a call conforms to policy. P-CSCF uses the IMS Rx interface during registration to subscribe to learn access network information and signaling path status.

PCRF policy control includes:

- IP flow/media authorization
- QoS resource reservation
- Early media/bandwidth authorization
- Notification of media path changes/states
- Notification of signaling path changes/states

Presence Enabled

With its high transaction setup rate, this is an ideal solution to handle a large number of messages generated by presence signaling. CSCF supports all the presence RFC extensions and signaling and interoperates with several presence servers.

Redirection

The SCM supports response to 3xx redirect messages. In addition to supporting redirection as per 3GPP, it supports call redirection to other chassis in the network (based on configuration) in case of system overload.

Redundancy and Session Recovery

When enabled, provides automatic failover of existing CSCF sessions due to hardware or software faults.

The system recovers from a single hardware or software fault with minimal interruption to the subscriber's service and maintains session information to rebuild sessions if multiple faults occur.

Registration Event Package

A set of event notifications used to inform SIP node of changes made to a registration.

Signaling Compression (SigComp)

SigComp compresses SIP call setup messages and is supported on the P-CSCF component. This reduces bandwidth demands on the RAN and reduces setup times.

SIP Denial of Service (DoS) Attack Prevention

The A-BG provides a scalable proxy network and a distributed Network Address Translation (NAT) network which effectively mitigates DoS attacks.

The SCM prevents a variety of DoS attacks specific to CSCF and SIP technology.

SIP Intelligence at the Core

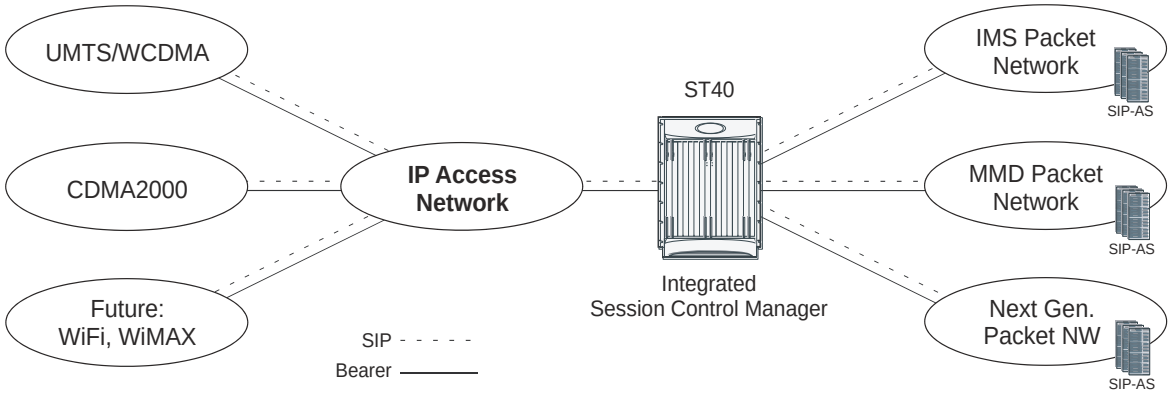
The SCM provides operators with an easy on-ramp for deploying SIP-based subscriber services while supporting various network control operations that provide the necessary intelligent control to insure a robust, carrier-class subscriber experience is achieved in this always changing multimedia environment.

When integrated into Cisco's session-aware Home Agent or GGSN platform, the SCM becomes the first SIP hop in the network, allowing operators to monitor and control all SIP-based sessions and execute additional value-added functions.

As the logical anchor point within the packet core, the SCM improves the user experience with device and location independence, and enhances subscriber control and policy enforcement with faster, more intelligent decisions for multimedia services.

Furthermore, as Fixed Mobile Convergence takes hold, it will be especially important to incorporate the SCM in the packet core in order to achieve mobility and voice continuity between multiple access networks (3G, WiFi, WiMAX, etc.).

Figure 241. Cisco Integrated Session Control Manager



SIP Large Message Support

Large notify contains information about multiple users in one message, which reduces the number of SIP messages in the network. Large SIP messages can be sent on UDP if the endpoint can support fragmentation; otherwise, UDP to TCP switching can be used to transport large messages intact.

If a request is within 200 bytes of the path MTU, or if it is larger than 1300 bytes and the path MTU is unknown, the request MUST be sent using TCP. This prevents fragmentation of messages over UDP and provides congestion control for larger messages. P-CSCF/A-BG is also able to handle messages up to the maximum datagram packet size. For UDP, this size is 65,535 bytes, including IP and UDP headers.

Large message support is needed for handling presence signaling traffic as the size of messages could be as large as 50K.

SIP Routing Engine

The SIP routing engine deploys SIP in a secure and controlled fashion.

Provides auto discovery of SIP elements, subscriber privacy, call fraud prevention, network security, and thwarting of network overload conditions.

Shared Initial Filter Criteria (SiFC)

If both the HSS and the S-CSCF support this feature, subsets of iFC may be shared by several service profiles. The HSS downloads the unique identifiers of the shared iFC sets to the S-CSCF. The S-CSCF uses a locally administered database to map the downloaded identifiers onto the shared iFC sets.

If the S-CSCF does not support this feature, the HSS will not download identifiers of shared iFC sets.

Telephony Application Server (TAS) Basic Supported

The following describe the local basic call features implemented on the S-CSCF:

- Abbreviated Dialing (AD)
- Call Forward Busy Line (CFBL)
- Call Forward No Answer (CFNA)
- Call Forward Not Registered (CFNR)
- Call Forward Unconditional (CFU)
- Call Transfer
- Call Waiting
- Caller ID Display (CID)
- Caller ID Display Blocked (CIDB)
- Feature Code Activation/De-activation
- Follow Me/Find Me
- Locally Allowed Abbreviated Dialing
- Outbound Call Restrictions/Dialing Permissions
- Short Code Dialing

TAS Basic provides basic voice call feature support in the SCM. In the IMS architecture, these telephony features are normally provided by an external application server. Providing these features with the S-CSCF:

- Reduces the need for an additional SIP AS
- Simplifies the network architecture
- Improves latency for call setup and feature invocation

The following describe the local basic call features implemented on the S-CSCF:

- **Abbreviated Dialing (AD)** - This feature allows the subscriber to call a Directory Number by entering less than the usual ten digits.Usually, the subscriber has four digit dialing to mimic PBX dialing privileges but these must be set up prior to use. When the SCM receives these numbers, it translates them and routes the call.
- **Call Forward Busy Line (CFBL)** - This feature forwards the call if busy line indication is received from the UE. If CFBL is enabled on both the AS and the S-CSCF, the call is forwarded by the S-CSCF on Busy Line indication. The feature detects and eliminates call forward loops if the History-Info header is present. It also

terminates forwarding if forwarding causes the forward attempts to be more than the number specified in the Max-Forwards header.

- **Call Forward No Answer (CFNA)** - This feature forwards the call if no answer is received from the UE. If CFNA is enabled on both the AS and the S-CSCF, the call is forwarded by the S-CSCF on No Answer indication. The feature detects and eliminates call forward loops if the History-Info header is present. It also terminates forwarding if forwarding causes the forward attempts to be more than the number specified in the Max-Forwards header.
- **Call Forward Not Registered (CFNR)** - This feature forwards the call if the subscriber is not registered. If CFNR is enabled on both the AS and the S-CSCF, the call is forwarded by the S-CSCF on Not Registered indication. The feature detects and eliminates call forward loops if the History-Info header is present. It also terminates forwarding if forwarding causes the forward attempts to be more than the number specified in the Max-Forwards header.
- **Call Forward Unconditional (CFU)** - This feature unconditionally forwards the call. The check for local CFU is done prior to the filter criteria and before any AS interaction. Thus CFU is enabled on both the S-CSCF and the destination AS, the local CFU occurs and there is no AS interaction. The feature eliminates basic loop detection (A calls B which is forwarded to A) and if the History-Info header is present, enhanced loop detection is performed based on the contents of this header. It also terminates forwarding if forwarding causes the forward attempts to be more than the number specified in the Max-Forwards header.
- **Call Transfer** - This feature allows the subscriber to transfer a call.
- **Call Waiting** - This feature allows the subscriber to receive a second call while on the first call.
- **Caller ID Display (CID)** - This feature inserts P-Preferred-Identity which communicates the identity of the user within the trust domain. If this header is already present, the feature may not do anything different.
- **Caller ID Display Blocked (CIDB)** - This feature removes P-Preferred-Identity and P-Preferred-Asserted-Identity headers and inserts a Privacy header with the privacy value set to “id”.
- **Feature Code Activation/De-activation** - This feature allows for activating and de-activating certain features using a star (*) - number sequence (star code). Registered subscribers have the option of activating or deactivated call features using specified star codes. The SCM translates these codes and routes the call.
- **Follow Me/Find Me** - This feature invokes the incoming call to several configured destinations in parallel and connects the call to the first destination that responds, “tearing down” all the other calls. There are two possible implementations of this feature; one a sequential implementation in which each destination is attempted in sequence till a successful connection. The other is a parallel approach in which several destinations are tried simultaneously. The advantage of the parallel approach is a faster set up.
- **Locally Allowed Abbreviated Dialing** - This feature allows the subscriber to dial a local-only, legacy, short code such as *CG or *POL. The SCM translates these codes to a ten-digit directory number and routes the call.
- **Outbound Call Restrictions/Dialing Permissions** - This feature restricts subscribers from initiating certain outbound calls. For example, if a subscriber attempts to make an international call and is not permitted to, the S-CSCF rejects the call.
- **Short Code Dialing** - This feature allows the subscriber to dial a short code such as #PAY or #MIN. The SCM translates these codes and routes the call.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be

indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, IP pool addresses, etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert**: A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm**: Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps**: SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values.

  Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.
- **Logs**: The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

  Logs are supported in both the Alert and the Alarm models.
- **Alarm System**: High threshold alarms generated within the specified polling interval are considered "outstanding" until a the condition no longer exists or a condition clear alarm is generated. "Outstanding" alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

  The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

Trust Domain

Enables the identification of trusted network entities. This keeps subscriber information confidential when it is received.

Features and Functionality - External Application Support

This section describes the features and functions of external applications supported on the SCM. These services require additional licenses to implement the functionality.

- Web Element Management System

Web Element Management System

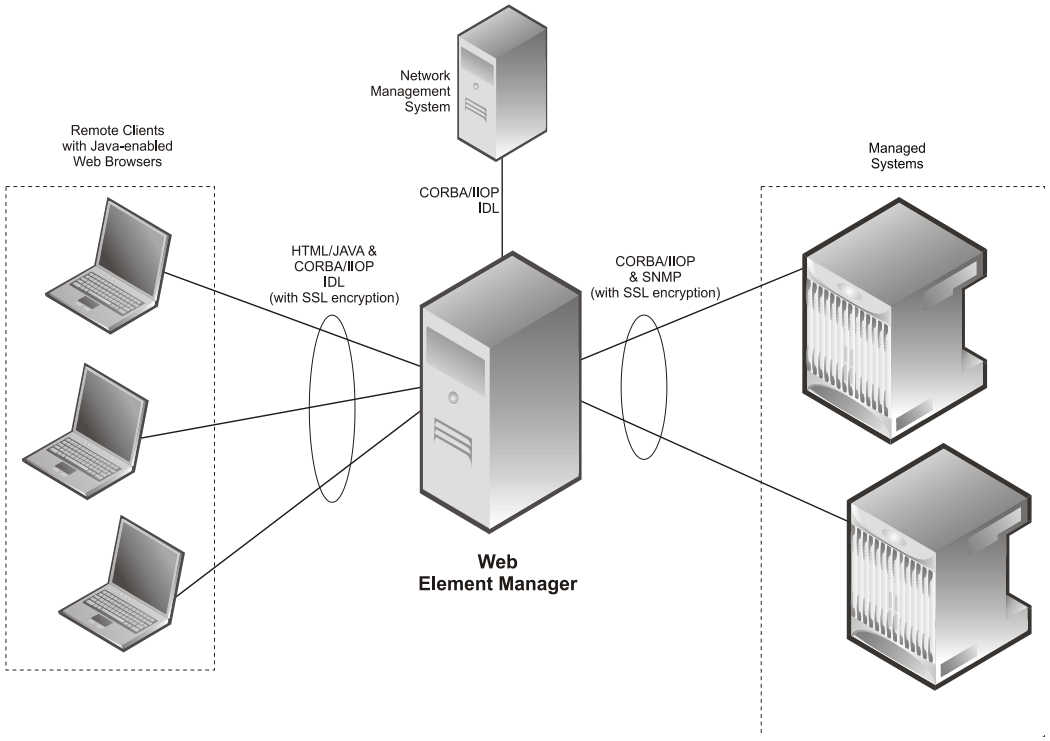
Provides a graphical user interface (GUI) for performing fault, configuration, accounting, performance, and security (FCAPS) management of the ASR 5000.

The Web Element Manager is a Common Object Request Broker Architecture (CORBA)-based application that provides complete fault, configuration, accounting, performance, and security (FCAPS) management capability for the system.

For maximum flexibility and scalability, the Web Element Manager application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the fault-tolerant Sun® Solaris® operating system.

The following figure demonstrates various interfaces between the Cisco Web Element Manager and other network components.

Figure 242. Web Element Manager Network Interfaces





Important: For more information on WEM support, refer to the *WEM Installation and Administration Guide*.

Features and Functionality - Licensed Enhanced Feature Support

This section describes optional enhanced features and functions.

Each of the following optional enhanced features require the purchase of an additional license to implement the functionality with the SCM.

This section describes the following features:

- Interchassis Session Recovery
- IPSec Support
- IPv4-IPv6 Interworking
- Lawful Intercept
- Session Recovery Support
- TLS Support in P-CSCF

Interchassis Session Recovery

Use of Interchassis Session Recovery requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

The ASR 5000 provides industry leading carrier class redundancy. The systems protects against all single points of failure (hardware and software) and attempts to recover to an operational state when multiple simultaneous failures occur.

The system provides several levels of system redundancy:

- Under normal N+1 PSC/PSC2/PSC3 hardware redundancy, if a catastrophic packet processing card failure occurs all affected calls are migrated to the standby packet processing card if possible. Calls which cannot be migrated are gracefully terminated with proper call-termination signaling and accounting records are generated with statistics accurate to the last internal checkpoint
- If the Session Recovery feature is enabled, any total packet processing card failure will cause a packet processing card switchover and all established sessions for supported call-types are recovered without any loss of session.

Even though Cisco Systems provides excellent intra-chassis redundancy with these two schemes, certain catastrophic failures which can cause total chassis outages, such as IP routing failures, line-cuts, loss of power, or physical destruction of the chassis, cannot be protected by this scheme. In such cases, the Interchassis Session Recovery feature provides geographic redundancy between sites. This has the benefit of not only providing enhanced subscriber experience even during catastrophic outages, but can also protect other systems such as the RAN from subscriber re-activation storms.

The Interchassis Session Recovery feature allows for continuous call processing without interrupting subscriber services. This is accomplished through the use of redundant chassis. The chassis are configured as primary and backup with one being active and one in recovery mode. A checkpoint duration timer is used to control when subscriber data is sent from the active chassis to the inactive chassis. If the active chassis handling the call traffic goes out of service, the inactive chassis transitions to the active state and continues processing the call traffic without interrupting the subscriber session. The chassis determines which is active through a propriety TCP-based connection called a redundancy link.

This link is used to exchange Hello messages between the primary and backup chassis and must be maintained for proper system operation.

- **Interchassis Communication**

  Chassis configured to support Interchassis Session Recovery communicate using periodic Hello messages. These messages are sent by each chassis to notify the peer of its current state. The Hello message contains information about the chassis such as its configuration and priority. A dead interval is used to set a time limit for a Hello message to be received from the chassis' peer. If the standby chassis does not receive a Hello message from the active chassis within the dead interval, the standby chassis transitions to the active state. In situations where the redundancy link goes out of service, a priority scheme is used to determine which chassis processes the session. The following priority scheme is used:

  - router identifier
  - chassis priority
  - SPIO MAC address

- **Checkpoint Messages**

  Checkpoint messages are sent from the active chassis to the inactive chassis. Checkpoint messages are sent at specific intervals and contain all the information needed to recreate the sessions on the standby chassis, if that chassis were to become active. Once a session exceeds the checkpoint duration, checkpoint data is collected on the session. The checkpoint parameter determines the amount of time a session must be active before it is included in the checkpoint message.

Important: For more information on interchassis session recovery support, refer to the *Interchassis Session Recovery* chapter in the *System Administration Guide*.

IPSec Support

Use of IPSec requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

Encrypted IPSec tunnels are terminated and decrypted so that traffic coming from untrusted networks are secured before entering the secure operator network. This prevents eavesdropping, hijacking, and other intrusive behavior from occurring.

IP Security (IPSec) is a suite of protocols that interact with one another to provide secure private communications across IP networks. These protocols allow the system to establish and maintain secure tunnels with peer security gateways.

Important: IPSec implementation is a mandatory part of IPv6, but it is optional to secure IPv4 traffic.

IPv4-IPv6 Interworking

Use of IPv4-IPv6 interworking requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

This feature allows the P-CSCF to provide IPv4-IPv6 interworking in the following scenarios:

- When UEs are IPv6-only and the IMS core network is IPv4-only
- When UEs are IPv4-only and the IMS core network is IPv6-only

In addition, IPv4-IPv6 interworking helps an IPv4 IMS network transition to an all-IPv6 IMS network.

The following interworking requirements are currently supported:

- MSRP support when IPv4-IPv6 interworking is enabled
- IPv4 TCP and IPv6 TCP
- Transport switching allowed based on size for both v4 and v6 network
- UDP fragmentation allowed for both v4 and v6 networks
- P-CSCF supports Mw and Gm interfaces on both v4 and v6
- KPIs for Mw and Gm interfaces are supported on both v4 and v6
- DNS supported for v4 and v6 networks
- Interworking supported for IM and presence
- Both v4 and v6 handsets are supported simultaneously on the same P-CSCF node

P-CSCF will provide IPv4-IPv6 interworking functionality between IPv6-only UEs and IPv4-only core network elements (I/S-CSCF) by acting as a dual stack. To achieve the dual-stack behavior, P-CSCF will be configured in two services with the first service (V6-SVC) listening on an IPv6 address and the second service (V4-SVC) listening on an IPv4 address. SIP messages coming from IPv6 UEs will come to V6-SVC and will be forwarded to the IPv4 core network through V4-SVC. Similarly, messages from the IPv4 core network come to V4-SVC and will be forwarded to IPv6 UEs via V6-SVC. P-CSCF also provides interworking functionality between IPV4-only UEs and IPv6-only core network elements.

P-CSCF handling different v4-v6 interworking scenarios is shown below.

Figure 243. Interworking Between IPv6 UE and IPv4 IMS Core Network

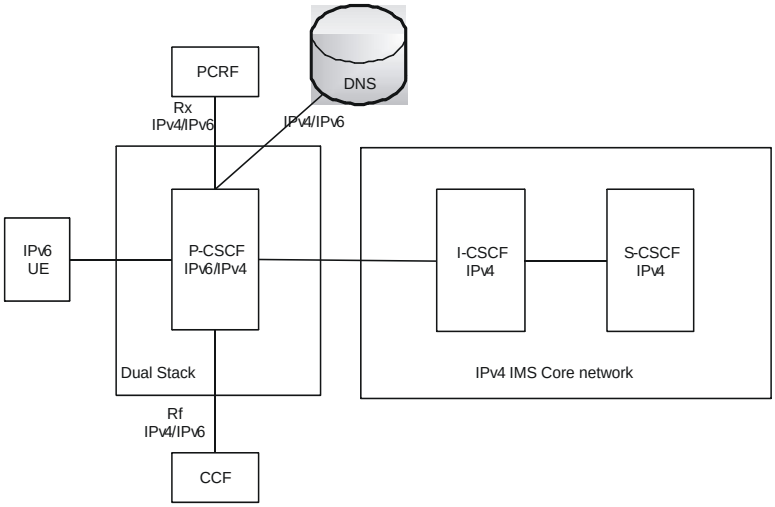
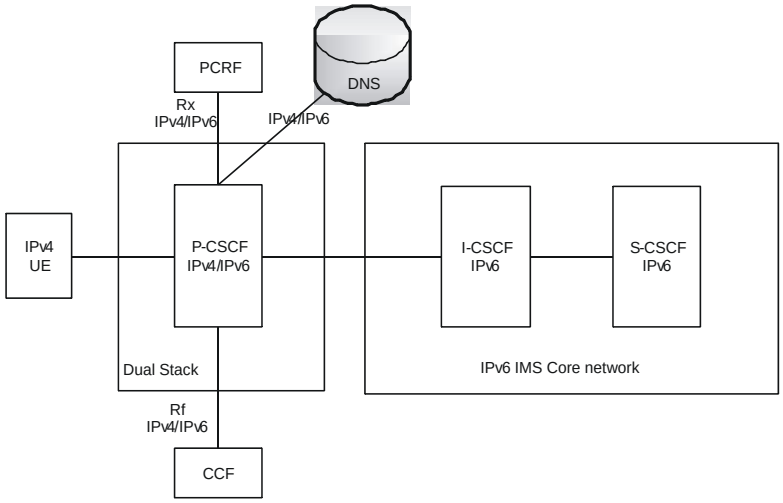


Figure 244. Interworking Between IPv4 UE and IPv6 IMS Core Network



To identify the need for IPv4-IPv6 interworking for a new incoming IPv6 REGISTER arriving at V6-SVC, a route lookup is performed based on the request-uri, first in V4-SVC context and then in V6-SVC context if the first lookup does not return any matching route entry. If a matching IPv4 next-hop route entry is found, then this indicates that interworking needs to be done. If no route entry is found, then a DNS query on request-uri domain is done for both A and AAAA type records. If DNS response yields only an IPv4 address, then this is also the case for performing IPv4-IPv6 interworking.

Headers (such as Via, Path, etc.) are automatically set to IPv4 bind address of P-CSCF V4-SVC. Remaining headers will be not be altered and sent as is toward the S-CSCF. The IPv4 address in a Path header received from S-CSCF in 200Ok of REGISTER will be replaced with V6-SVC's IPv6 address before forwarding to UE.

Lawful Intercept

Use of Lawful Intercept requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

The Cisco Lawful Intercept feature is supported on the SCM. Lawful Intercept is a licensed-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your Cisco account representative.

Session Recovery Support

Use of Session Recovery requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

Session recovery is performed by mirroring key software processes (e.g. session manager and AAA manager) within the system. These mirrored processes remain in an idle state (in standby-mode), wherein they perform no processing, until

they may be needed in the case of a software failure (e.g. a session manager task aborts). The system spawns new instances of "standby mode" session and AAA managers for each active Control Processor (CP) being used.

Additionally, other key system-level software tasks, such as VPN manager, are performed on a physically separate Packet Services Card (PSC/PSC2/PSC3) to ensure that a double software fault (e.g. session manager and VPN manager fails at same time on same card) cannot occur. The PSC/PSC2/PSC3 used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

The additional hardware resources required for session recovery include a standby System Management Card (SMC) and a standby PSC/PSC2/PSC3.

There are two modes for Session Recovery.

- **Task recovery mode**: Wherein one or more session manager failures occur and are recovered without the need to use resources on a standby PSC. In this mode, recovery is performed by using the mirrored "standby-mode" session manager task(s) running on active PSCs. The "standby-mode" task is renamed, made active, and is then populated using information from other tasks such as AAA manager.
- **Full PSC/PSC2/PSC3 recovery mode**: Used when a PSC hardware failure occurs, or when a PSC migration failure happens. In this mode, the standby PSC is made active and the "standby-mode" session manager and AAA manager tasks on the newly activated PSC perform session recovery.

Session/Call state information is saved in the peer AAA manager task because each AAA manager and session manager task is paired together. These pairs are started on physically different PSCs/PSC2s/PSC3s to ensure task recovery.

Important: Session Recovery is supported for either IPv4 or IPv6 traffic.

Important: For more information on session recovery support, refer to the *Session Recovery* chapter in the *System Administration Guide*.

TLS Support in P-CSCF

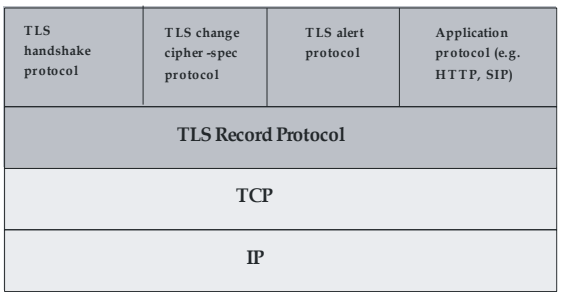
Use of SSL requires that a valid license key be installed. Contact your local Sales or Support representative for information on how to obtain a license.

Transport Layer Security (TLS) provides confidentiality and integrity protection for SIP signaling messages between the UE and P-CSCF/A-BG. TLS is a layered protocol that runs upon reliable transport protocols like TCP and SCTP.

The SCM supports the following two scenarios:

- TLS as a transport between UE and P-CSCF/A-BG, as per RFC 3261
- Use of TLS by Security Mechanism agreement between UE and P-CSCF/A-BG, as per RC 3329 and TS 33.203

The following figure shows the TLS protocol layers.



How the SCM Works

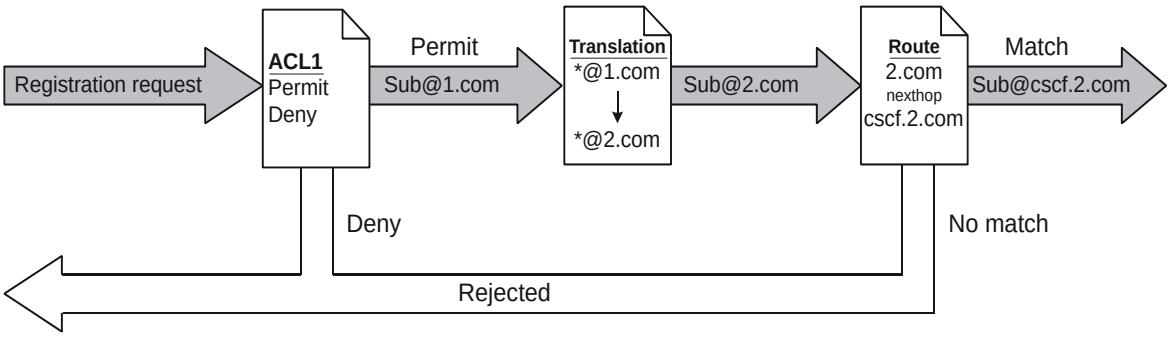
This section provides information on the function of the SCM in a CDMA2000 PDSN or UMTS GGSN network and presents call procedure flows for different stages of session setup.

Admission and Routing

Admission and routing of subscriber URIs is performed through a number of configurable lists in the SCM.

The following sections describe the main admission and routing techniques used in the SCM. The following figure presents the method and order for admitting and routing sessions within the SCM.

Figure 245. Admission and Routing Method



CSCF Access Control Lists

Access Control Lists (ACLs) are a set of rules that are applied during CSCF session establishment. A typical use of these rules is to accept or deny registration or session establishment requests. ACLs may be tied to subscribers and/or the whole service. Subscriber based ACLs can also be imported from an external ACL/policy server. In that event, the external policy server address would be configured with the service.

Translation Lists

Translation lists help modify request-uri (i.e. addressing of a CSCF session). One example is that E.164 numbers could be altered by adding prefixes and suffixes or the request-uri could be modified based on the registration database.

Route Lists

Route lists are service level lists that assist in finding the next CSCF/UA hop. These are static routes and will override any dynamic routes (based on DNS queries for FQDNs).

Signaling Compression

The Session Initiation Protocol (SIP) is a text-based protocol designed for higher bandwidth networks. As such, it is inherently less suited for lower bandwidth environments such as wireless networks. If a wireless handset uses SIP to set up a call, the setup time is significantly increased due to the high overhead of text-based signaling messages.

Signaling Compression (SigComp) is a solution for compressing/decompressing messages generated by application protocols such as SIP. The P-CSCF component of the SCM uses SigComp to reduce call setup times on the access network, typically between the P-CSCF and the UE. The following features are supported:

- **SigComp Detection** - P-CSCF detects if the UE supports SigComp and compresses messages it sends to the UE. The P-CSCF also detects if messages it receives are compressed and decompresses them.
- **SigComp Parameter Configuration** - P-CSCF allows the configuration of Decompression Memory Size (DMS), State Memory Size (SMS), and Cycles Per Bit (CPB).
- **Failure Acknowledgement** - P-CSCF replies with NACK on decompression failure.
- **SIP/SDP Static Dictionaries** - P-CSCF supports the Session Initiation Protocol/Session Description Protocol Static Dictionary for Signaling Compression.

Supported Standards

The SCM service complies with the following standards for CDMA2000 PDSN, UMTS GGSN, and LTE network wireless data services.

Release 9 3GPP References

Important: The SCM currently supports the following Release 9 3GPP specifications. Most 3GPP specifications are also used for 3GPP2 support; any specifications that are unique to 3GPP2 would be listed under 3GPP2 References.

- TS 23.167 IP Multimedia Subsystem (IMS) emergency sessions
- TS 23.204 Support of Short Message Service (SMS) over generic 3GPP Internet Protocol (IP) access; Stage 2
- TS 23.207 End-to-end Quality of Service (QoS) concept and architecture
- TS 23.228 IP Multimedia Subsystem (IMS); Stage 2
- TS 23.981 Interworking aspects and migration scenarios for IPv4-based IP Multimedia Subsystem (IMS) implementations
- TS 24.229 Internet Protocol (IP) multimedia call control protocol based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP); Stage 3
- TS 24.341 Support of SMS over IP networks; Stage 3
- TS 29.208 End-to-end Quality of Service (QoS) signalling flows
- TS 29.214 Policy and charging control over Rx reference point
- TS 29.228 IP Multimedia (IM) Subsystem Cx and Dx interfaces; Signalling flows and message contents
- TS 29.229 Cx and Dx interfaces based on the Diameter protocol; Protocol details
- TS 32.240 Telecommunication management; Charging management; Charging architecture and principles
- TS 32.260 Telecommunication management; Charging management; IP Multimedia Subsystem (IMS) charging
- TS 33.203 3G security; Access security for IP-based services
- TS 33.978 Security aspects of early IP Multimedia Subsystem (IMS)

Release 8 3GPP References

Important: The SCM currently supports the following Release 8 3GPP specifications. Most 3GPP specifications are also used for 3GPP2 support; any specifications that are unique to 3GPP2 are listed under 3GPP2 References.

- TR 23.806 Voice call continuity between Circuit Switched (CS) and IP Multimedia Subsystem (IMS) Study
- TR 23.808 Supporting Globally Routable User Agent URI (GRUU) in IMS; Report and conclusions
- TR 23.816 Identification of Communication Services in IMS
- TS 24.229 IP multimedia call control protocol based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP); Stage 3

- TR 24.930 IP Multimedia core network Subsystem (IMS) based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP); Stage 3
- TR 29.847 Conferencing based on SIP, SDP, and other protocols; Functional models, information flows and protocol details
- TR 33.978 Security aspects of early IP Multimedia Subsystem (IMS)
- TS 22.101 Service principles
- TS 23.003 Numbering, addressing and identification
- TS 23.107 Quality of Service (QoS) concept and architecture
- TS 23.125 Overall high level functionality and architecture impacts of flow based charging; Stage 2
- TS 23.141 Presence service; Architecture and functional description; Stage 2
- TS 23.167 IP Multimedia Subsystem (IMS) emergency sessions
- TS 23.203 Policy and charging control architecture
- TS 23.204 Support of Short Message Service (SMS) over generic 3GPP Internet Protocol (IP) access; Stage 2
- TS 23.207 End-to-end Quality of Service (QoS) concept and architecture
- TS 23.218 IP Multimedia (IM) session handling; IM call model; Stage 2
- TS 23.221 Architectural Requirements
- TS 23.228 IP Multimedia Subsystem (IMS); Stage 2
- TS 23.271 Functional description of Location Services (LCS)
- TS 23.981 Interworking aspects and migration scenarios for IPv4-based IP Multimedia Subsystem (IMS) implementations
- TS 24.141 Presence service using the IP Multimedia (IM) Core Network (CN) subsystem; Stage 3
- TS 24.228 Signalling flows for the IP multimedia call control based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP); Stage 3
- TS 24.229 Internet Protocol (IP) multimedia call control protocol based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP); Stage 3
- TS 24.341 Support of SMS over IP networks; Stage 3
- TS 26.114 IP Multimedia Subsystem (IMS); Multimedia telephony; Media handling and interaction
- TS 26.141 IP Multimedia System (IMS) Messaging and Presence; Media formats and codecs
- TS 26.234 Transparent end-to-end Packet-switched Streaming Service (PSS); Protocols and codecs
- TS 26.235 Packet switched conversational multimedia applications; Default codecs
- TS 26.236 Packet switched conversational multimedia applications; Transport protocols
- TS 29.207 Policy control over Go interface
- TS 29.208 End-to-end Quality of Service (QoS) signalling flows
- TS 29.209 Policy control over Gq interface
- TS 29.213 Policy and charging control signalling flows and Quality of Service (QoS) parameter mapping
- TS 29.214 Policy and charging control over Rx reference point
- TS 29.228 IP Multimedia (IM) Subsystem Cx and Dx interfaces; Signalling flows and message contents
- TS 29.229 Cx and Dx interfaces based on the Diameter protocol; Protocol details

- TS 29.328 IMS Sh interface: signalling flows and message content
- TS 29.329 IMS Sh interface based on the Diameter protocol; Protocol details
- TS 31.103 Characteristics of the IMS Identity Module (ISIM) application
- TS 32.225 Telecommunication management; Charging management; Charging data description for the IP Multimedia Subsystem (IMS)
- TS 32.240 Telecommunication management; Charging management; Charging architecture and principles
- TS 32.260 Telecommunication management; Charging management; IP Multimedia Subsystem (IMS) charging
- TS 32.299 Telecommunication management; Charging management; Diameter charging applications
- TS 33.102 3G security; Security architecture
- TS 33.178 Security aspects of early IP Multimedia Subsystem (IMS)
- TS 33.203 3G security; Access security for IP-based services
- TS 33.978 Security aspects of early IP Multimedia Subsystem (IMS)

3GPP2 References

- S.R0079-A v1.0 Support for End-to-End QoS - Stage 1 Requirements
- S.R0086-A v1.0 IMS Security Framework
- X.S0013-000-A v1.0 All-IP Core Network Multimedia Domain - Overview
- X.S0013-002-A v1.0 All-IP Core Network Multimedia Domain - IP Multimedia Subsystem Stage 2
- X.S0013-003-0 v2.0 All-IP Core Network Multimedia Domain - IP Multimedia (IMS) Session Handling; IP Multimedia (IM) Call Model - Stage 2
- X.S0013-004-A v1.0 All-IP Core Network Multimedia Domain - IP Multimedia Call Control Protocol Based on SIP and SDP Stage 3
- X.S0013-005-0 All-IP Core Network Multimedia Domain: IP Multimedia Subsystem Cx Interface Signaling Flows and Message Contents
- X.S0013-006-0 All-IP Core Network Multimedia Domain - Cx Interface Based on the Diameter Protocol; Protocol Details
- X.S0013-007-0 All-IP Core Network Multimedia Domain: IP Multimedia Subsystem - Charging Architecture
- X.S0013-007-A v1.0 All-IP Core Network Multimedia Domain - IP Multimedia Subsystem - Charging Architecture
- X.S0013-008-0 All-IP Core Network Multimedia Domain: IP Multimedia Subsystem - Accounting Information Flows and Protocol
- X.S0013-008-A All-IP Core Network Multimedia Domain - IP Multimedia Subsystem - Offline Accounting Information Flows and Protocol
- X.S0013-010-0 v1.0 All-IP Core Network Multimedia Domain: IP Multimedia Subsystem Sh Interface; Signaling Flows and Message Contents - Stage 2
- X.S0013-011-0 v1.0 All-IP Core Network Multimedia Domain: Sh Interface Based on Diameter Protocols Protocol Details - Stage 3
- X.S0013-012-0 v1.0 All-IP Core Network Multimedia Domain - Service Based Bearer Control - Stage 2

- X.S0013-014-0 v1.0 All-IP Core Network Multimedia Domain - Service Based Bearer Control - Tx Interface Stage 3
- X.S0016-000-A v1.0 3GPP2 Multimedia Messaging System MMS Specification Overview, Revision A
- X.S0027-002-0 v1.0 Presence Security
- X.S0027-003-0 v1.0 Presence Stage 3
- X.S0029-0 v1.0 Conferencing Using the IP Multimedia (IM) Core Network (CN) Subsystem
- X.S0049-0 v1.0 All-IP Network Emergency Call Support

IETF References

- RFC 1594 (March 1994): "FYI on Questions and Answers to Commonly Asked "New Internet User" Questions"
- RFC 1889 (January 1996): "RTP: A Transport Protocol for Real-Time Applications"
- RFC 2246 (January 1999): "TLS protocol version 1.0"
- RFC 2327 (April 1998): SDP: Session Description Protocol
- RFC 2401 (November 1998): "Security Architecture for the Internet Protocol (IPSec)"
- RFC 2403 (November 1998): "The Use of HMAC-MD5-96 within ESP and AH"
- RFC 2404 (November 1998): "The Use of HMAC-SHA-1-96 within ESP and AH"
- RFC 2462 (December 1998): "IPv6 Address Autoconfiguration"
- RFC 2617 (June 1999): "HTTP Authentication: Basic and Digest Access Authentication"
- RFC 2753 (January 2000): "A Framework for Policy-based Admission Control"
- RFC 2833 (May 2000): "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals"
- RFC 2915 (September 2000): The Naming Authority Pointer (NAPTR) DNS Resource Record
- RFC 2976 (October 2000): "The SIP INFO Method"
- RFC 3041 (January 2001): "Privacy Extensions for Stateless Address Autoconfiguration in IPv6"
- RFC 3261 (June 2002): "SIP: Session Initiation Protocol"
- RFC 3262 (June 2002): "Reliability of provisional responses in Session Initiation Protocol (SIP)"
- RFC 3263 (June 2002): "Session Initiation Protocol (SIP): Locating SIP Servers"
- RFC 3264 (June 2002): "An Offer/Answer Model with Session Description Protocol (SDP)"
- RFC 3265 (June 2002): "Session Initiation Protocol (SIP) - Specific Event Notification"
- RFC 3280 (April 2002): "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile"
- RFC 3310 (September 2002): "Hypertext Transfer Protocol (HTTP) Digest Authentication Using Authentication and Key Agreement (AKA)"
- RFC 3311 (September 2002): "The Session Initiation Protocol (SIP) UPDATE Method".
- RFC 3312 (October 2002): "Integration of Resource Management and Session Initiation Protocol (SIP)"
- RFC 3313 (January 2003): "Private Session Initiation Protocol (SIP) Extensions for Media Authorization"
- RFC 3315 (July 2003): "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)"
- RFC 3320 (January 2003): "Signaling Compression (SigComp)"

- RFC 3321 (January 2003): "Signaling Compression (SigComp) - Extended Operations"
- RFC 3323 (November 2002): "A Privacy Mechanism for the Session Initiation Protocol (SIP)"
- RFC 3325 (November 2002): "Private Extensions to the Session Initiation Protocol (SIP) for Network Asserted Identity within Trusted Networks"
- RFC 3326 (December 2002): "The Reason Header Field for the Session Initiation Protocol (SIP)"
- RFC 3327 (December 2002): "Session Initiation Protocol (SIP) Extension Header Field for Registering Non-Adjacent Contacts"
- RFC 3329 (January 2003): "Security Mechanism Agreement for the Session Initiation Protocol (SIP)"
- RFC 3388 (December 2002): "Grouping of Media Lines in the Session Description Protocol (SDP)"
- RFC 3428 (December 2002): "Session Initiation Protocol (SIP) Extension for Instant Messaging"
- RFC 3455 (January 2003): "Private Header (P-Header) Extensions to the Session Initiation Protocol (SIP) for the 3rd-Generation Partnership Project (3GPP)"
- RFC 3485 (February 2003): "The Session Initiation Protocol (SIP) and Session Description Protocol (SDP) Static Dictionary for Signaling Compression (SigComp)"
- RFC 3486 (February 2003): "Compressing the Session Initiation Protocol (SIP)"
- RFC 3515 (April 2003): "The Session Initiation Protocol (SIP) Refer method"
- RFC 3556 (July 2003): "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth"
- RFC 3581 (August 2003): "An Extension to the Session Initiation Protocol (SIP) for Symmetric Response Routing"
- RFC 3588 (September 2003): "Diameter Base Protocol"
- RFC 3608 (October 2003): "Session Initiation Protocol (SIP) Extension Header Field for Service Route Discovery During Registration"
- RFC 3665 (December 2003): "Session Initiation Protocol (SIP) Basic Call Flow Examples"
- RFC 3680 (March 2004): "A Session Initiation Protocol (SIP) Event Package for Registrations"
- RFC 3761 (April 2004): "The E.164 to Uniform Resource Identifiers (URI) Dynamic Delegation Discovery System (DDDS) Application (ENUM)"
- RFC 3824 (June 2004): "Using E.164 numbers with the Session Initiation Protocol (SIP)"
- RFC 3840 (August 2004): "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)"
- RFC 3841 (August 2004): "Caller Preferences for the Session Initiation Protocol (SIP)"
- RFC 3842 (August 2004): "A Message Summary and Message Waiting Indication Event Package for the Session Initiation Protocol (SIP)"
- RFC 3856 (August 2004): "A Presence Event Package for the Session Initiation Protocol (SIP)"
- RFC 3857 (August 2004): "A Watcher Information Event Template-Package for the Session Initiation Protocol (SIP)"
- RFC 3858 (August 2004): "An Extensible Markup Language (XML) Based Format for Watcher Information"
- RFC 3861 (August 2004): "Address Resolution for Instant Messaging and Presence"
- RFC 3891 (September 2004): "The Session Initiation Protocol (SIP) "Replaces" Header"
- RFC 3892 (September 2004): "The Session Initiation Protocol (SIP) Referred-By Mechanism"
- RFC 3903 (October 2004): "Session Initiation Protocol (SIP) Extension for Event State Publication"

- RFC 3911 (October 2004): "The Session Initiation Protocol (SIP) "Join" Header"
- RFC 3966 (December 2004): "The tel URI for Telephone Numbers"
- RFC 3986 (January 2005): "Uniform Resource Identifier (URI): Generic Syntax"
- RFC 4028 (April 2005): "Session Timers in the Session Initiation Protocol (SIP)"
- RFC 4032 (March 2005): "Update to the Session Initiation Protocol (SIP) Preconditions Framework"
- RFC 4077 (May 2005): "A Negative Acknowledgement Mechanism for Signaling Compression"
- RFC 4244 (November 2005): "An Extension to the Session Initiation Protocol (SIP) for Request History Information"
- RFC 4317 (December 2005): "Session Description Protocol (SDP) Offer/Answer Examples"
- RFC 4353 (February 2006): "A Framework for Conferencing with the Session Initiation Protocol (SIP)"
- RFC 4475 (May 2006): "Session Initiation Protocol (SIP) Torture Test Messages"
- RFC 4566 (July 2006): "SDP: Session Description Protocol"
- RFC 4975 (September 2007): "Message Session Relay Protocol (MSRP)"
- RFC 5031 (January 2008): "A Uniform Resource Name (URN) for Emergency and Other Well-Known Services"
- RFC 5049 (December 2007): "Applying Signaling Compression (SigComp) to the Session Initiation Protocol (SIP)"
- RFC 5112 (January 2008): "The Presence-Specific Static Dictionary for Signaling Compression (Sigcomp)"
- RFC 5491 (March 2009): "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations"
- RFC 5626 (October 2009): "Managing Client Initiated Connections in the Session Initiation Protocol (SIP)"

Other

- Packet-Cable spec (PKT-TR-SEC-V02-061013)

Chapter 31
Serving GPRS Support Node (SGSN) Overview

This chapter contains general overview information about the Serving GPRS Support Node (SGSN), including sections for:

- Product Description
- Network Deployments and Interfaces
- SGSN Core Functionality
- Features and Functionality
- How the SGSN Works
- Supported Standards

Product Description

The ASR 5000 provides a highly flexible and efficient Serving GPRS Support Node (SGSN) service to the wireless carriers. Functioning as an SGSN, the system readily handles wireless data services within 2.5G General Packet Radio Service (GPRS) and 3G Universal Mobile Telecommunications System (UMTS) data networks.

Important: Throughout this chapter the designation for the subscriber equipment is referred to in various ways: UE for user equipment (common to 3G/4G scenarios), MS or mobile station (common to 2G/2.5G scenarios), and MN or mobile node (common to 2G/2.5G scenarios involving IP-level functions). Unless noted, these terms are equivalent and the term used usually complies with usage in the relevant standards.

In a GPRS/UMTS network, the SGSN works in conjunction with radio access networks (RANs) and Gateway GPRS Support Nodes (GGSNs) to:

- Communicate with home location registers (HLR) via a Gr interface and mobile visitor location registers (VLRs) via a Gs interface to register a subscriber's user equipment (UE), or to authenticate, retrieve or update subscriber profile information.
- Support Gd interface to provide short message service (SMS) and other text-based network services for attached subscribers.
- Activate and manage IPv4, IPv6, or point-to-point protocol (PPP) -type packet data protocol (PDP) contexts for a subscriber session.
- Setup and manage the data plane between the RAN and the GGSN providing high -speed data transfer with configurable GEA0-3 ciphering.
- Provide mobility management, location management, and session management for the duration of a call to ensure smooth handover.
- Provide various types of charging data records (CDRs) to attached accounting/billing storage mechanisms such as our SMC-based hard drive or a GTPP Storage Server (GSS) or a charging gateway function (CGF).
- Provide CALEA support for lawful intercepts.

This chapter catalogs many of the SGSN key components and features for data services within the GPRS/UMTS environment. Also, a range of SGSN operational and compliance information is summarized with pointers to other information sources.

Platform Requirements

The SGSN service runs on a Cisco® ASR 5000 Series chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

Licenses

The SGSN is a licensed Cisco product and requires the purchase and installation of the SGSN Software License. Separate feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements.

For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Network Deployments and Interfaces

The following logical connection maps illustrate the SGSN's ability to connect to various radio access network types, core network types, and network components:

- GSM edge radio access network (GERAN) provides access to the 2.5G general packet radio service (GPRS) network
- UMTS terrestrial radio access network (UTRAN) provides access to the 3G universal mobile telecommunications system (UMTS) network
- another SGSN
- standalone gateway GPRS support node (GGSN)
- co-located GGSN
- mobile service center (MSC)
- visitor location register (VLR)
- home location register (HLR)
- charging gateway (CF - sometimes referred to as a charging gateway function (CGF))
- GTPP storage server (GSS)
- equipment identity registry (EIR)
- CAMEL service's GSM service control function (gsmSCF)
- short message service server center (SMS-C)
- network devices in another PLMN

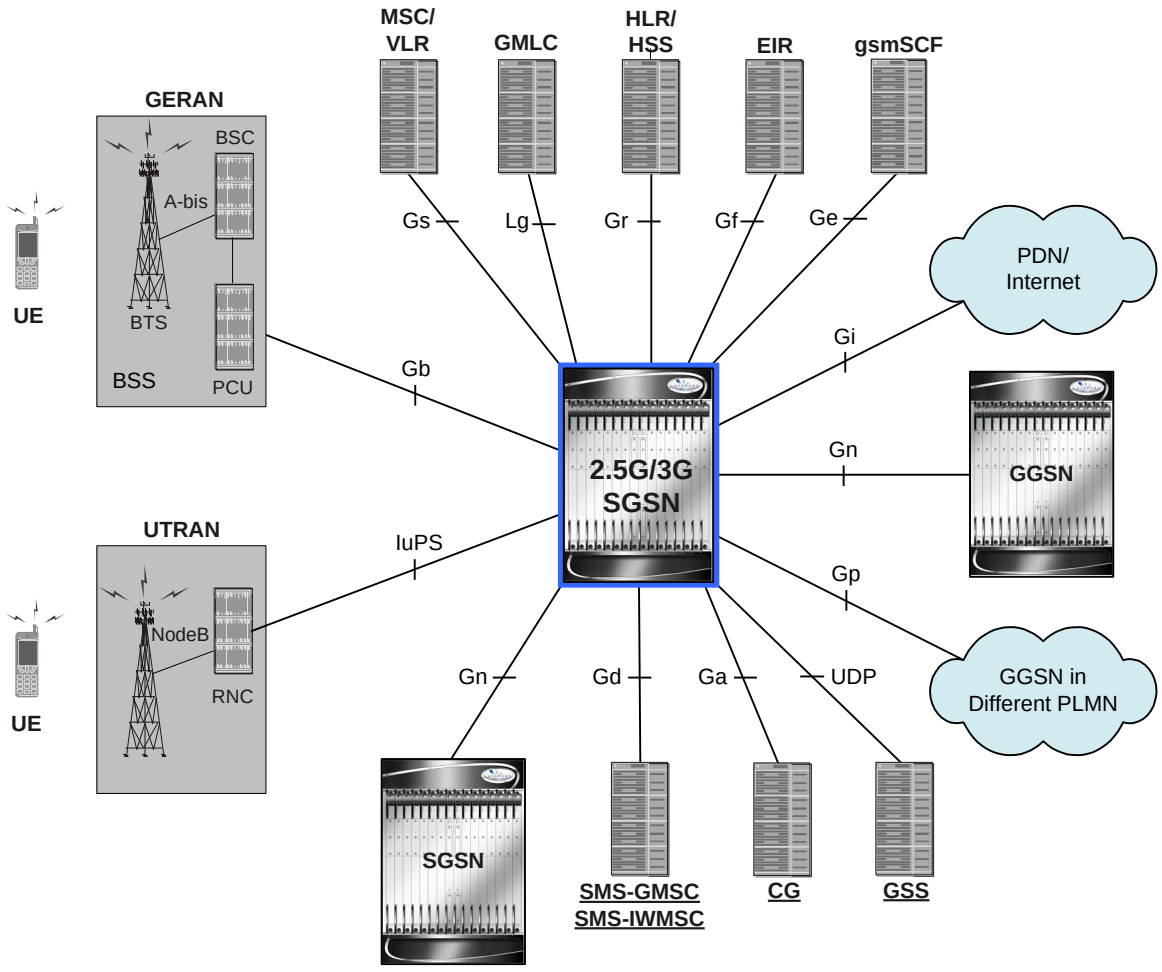
SGSN and Dual Access SGSN Deployments

SGSNs and GGSNs work in conjunction within the GPRS/UMTS network. As indicated earlier in the section on *System Configuration Options*, the flexible architecture of the ASR 5000 enables a single chassis to reduce hardware requirements by supporting integrated co-location of a variety of the SGSN services.

A chassis can be devoted solely to SGSN services or the SGSN system can include any co-location combination, such as multiple instances of 2.5G SGSNs (configured as GPRS services); or multiple instances of 3G SGSNs (configured as SGSN services); or a combination of 2.5G and 3G SGSN to comprise a dual access SGSN.

Important: The following illustrates the GPRS/UMTS Dual Access architecture with a display of all the interfaces supported as of Release 14.0. The *SGSN Logical Network Interfaces* section below lists the interfaces available for the release applicable to the version of this manual.

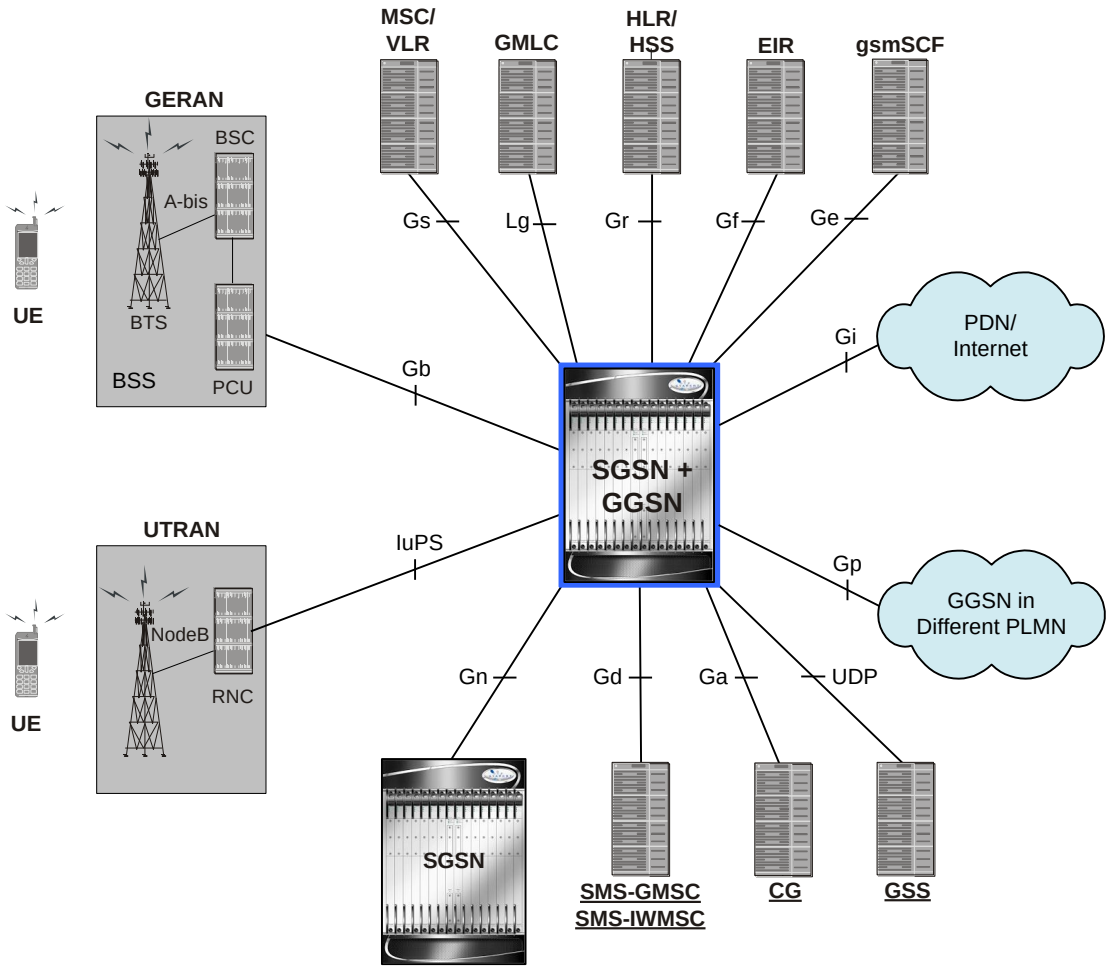
Figure 246. 2.5G & 3G Dual Access Architecture



SGSN/GGSN Deployments

The co-location of the SGSN and the GGSN in the same chassis facilitates handover. A variety of GSN combos is possible, 2.5G or 3G SGSN with the GGSN.

Figure 247. GSN Combo Architecture



SGSN Logical Network Interfaces

The SGSN provides IP-based transport on all RAN and core network interfaces, in addition to the standard IP-based interfaces (Ga, Gn, Gp, Iu-PS). This means enhanced performance, future-proof scaling and reduction of inter-connectivity complexity. The all-IP functionality is key to facilitating evolution to the next generation technology requirements.

The SGSN provides the following functions over the logical network interfaces illustrated above:

- **Ga:** The SGSN uses the Ga interface with GPRS Transport Protocol Prime (GTPP) to communicate with the charging gateway (CG, also known as CGF) and/or the GTPP storage server (GSS). The interface transport layer is typically UDP over IP but can be configured as TCP over IP for:
  - One or more Ga interfaces per system context, and
  - An interface over Ethernet 10/100 or Ethernet 1000 interfaces

  The charging gateway handles buffering and pre-processing of billing records and the GSS provides storage for Charging Data Records (CDRs). For additional information regarding SGSN charging, refer to the Charging section.

- **IuPS:** The SGSN provides an IP over ATM (IP over AAL5 over ATM) interface between the SGSN and the RNCs in the 3G UMTS radio access network (UTRAN). RANAP is the control protocol that sets up the data plane (GTP-U) between these nodes. SIGTRAN (M3UA/SCTP) or QSAAL (MTP3B/QSAAL) handle IuPS-C (control) for the RNCs.

  Some of the procedures supported across this interface are:

  - Control plane based on M3UA/SCTP
  - Up to 128 Peer RNCs per virtual SGSN. Up to 256 peers per physical chassis
  - SCTP Multi-Homing supported to facilitate network resiliency
  - M3UA operates in and IPSP client/server and single/double-ended modes
  - Multiple load shared M3UA instances for high-performance and redundancy
  - Works over Ethernet and ATM (IPoA) interfaces
  - Facilitates SGSN Pooling
  - RAB (Radio Access Bearer) Assignment Request
  - RAB Release Request
  - Iu Release Procedure
  - SGSN-initiated Paging
  - Common ID
  - Security Mode Procedures
  - Initial MN Message
  - Direct Transfer
  - Reset Procedure
  - Error Indication
  - SRNS relocation
- **Gb:** This is the SGSN's interface to the base station system (BSS) in a 2G radio access network (RAN). It connects the SGSN via UDP/IP (via an Ethernet interface) or Frame Relay (via a Channelized SDH or SONET interface). Gb-IP is the preferred interface as it improves control plane scaling as well as facilitates the deployment of SGSN Pools.

  Some of the procedures supported across this interface are:

  - BSS GSM at 900/1800/1900 MHz
  - BSS Edge
  - Frame Relay congestion handling
  - Traffic management per Frame Relay VC
  - NS load sharing
  - NS control procedures
  - BVC management procedures
  - Paging for circuit-switched services
  - Suspend/Resume

- Flow control
- Unacknowledged mode
- Acknowledged mode

- **Gn/Gp:** The Gn/Gp interfaces, comprised of GTP/UDP/IP-based protocol stacks, connect the SGSNs and GGSNs to other SGSNs and GGSNs within the same public land mobile network (PLMN) - the Gn - or to GGSNs in other PLMNs - the Gp.

  This implementation supports:

  - GTPv0 and GTPv1, with the capability to auto-negotiate the version to be used with any particular peer
  - GTP-C (control plane) and GTP-U (user plane)
  - Transport over ATM/STM-1Optical, Fast Ethernet, and Ethernet 1000 line cards/QGLCs)
  - One or more Gn/Gp interfaces configured per system context

  As well, the SGSN can support the following IEs from later version standards:

  - IMEI-SV
  - RAT TYPE
  - User Location Information

- **Ge:** This is the interface between the SGSN and the SCP that supports the CAMEL service. It supports both SS7 and SIGTRAN and uses the CAP protocol.

- **Gr:** This is the interface to the HLR. It supports SIGTRAN (M3UA/SCTP/IP) over Ethernet.

  Some of the procedures supported by the SGSN on this interface are:

  - Send Authentication Info
  - Update Location
  - Insert Subscriber Data
  - Delete Subscriber Data
  - Cancel Location
  - Purge
  - Reset
  - Ready for SM Notification
  - SIGTRAN based interfaces M3UA/SCTP
  - Peer connectivity can be through an intermediate SGP or directly depending on whether the peer (HLR, EIR, SMSC, GMLC) is SIGTRAN enabled or not
  - SCTP Multi-Homing supported to facilitate network resiliency
  - M3UA operates in IPSP client/server and single/double-ended modes
  - Multiple load shared M3UA instances for high-performance and redundancy
  - Works over Ethernet (IPoA) interface

- **Gs:** This is the interface used by the SGSN to communicate with the visitor location register (VLR) or mobile switching center (MSC) to support circuit switching (CS) paging initiated by the MSC. This interface uses Signaling Connection Control Part (SCCP) connectionless service and BSSAP+ application protocols.

- **Gd:** This is the interface between the SGSN and the SMS Gateway (SMS-GMSC / SMS-IWMSC) for both 2G and 3G technologies through multiple interface mediums. Implementation of the Gd interface requires purchase of an additional license.
- **Gf:** Interface is used by the SGSN to communicate with the equipment identity register (EIR) which keeps a listing of UE (specifically mobile phones) being monitored. The SGSN's Gf interface implementation supports functions such as:
  - International Mobile Equipment Identifier-Software Version (IMEI-SV) retrieval
  - IMEI-SV status confirmation

SGSN Core Functionality

The SGSN core functionality is comprised of:

- All-IP Network (AIPN)
- SS7 Support
- PDP Context Support
- Mobility Management
- Location Management
- Session Management
- Charging

All-IP Network (AIPN)

AIPN provides enhanced performance, future-proof scaling and reduction of inter-connectivity complexity.

In accordance with 3GPP, the SGSN provides IP-based transport on all RAN and core network interfaces, in addition to the standard IP-based interfaces (Ga, Gn, Gp, Iu-Data). The all-IP functionality is key to facilitating Iu and Gb Flex (SGSN pooling) functionality as well as evolution to the next generation technology requirements.

The following IP-based protocols are supported on the SGSN:

SS7 Support

The ASR 5000 SGSN implements SS7 functionality to communicate with the various SS7 network elements, such as HLRs and VLRs.

The SGSN employs standard SS7 addressing (point codes) and global title translation. SS7 feature support includes:

- Transport layer support includes:
  - Broadband SS7 (MTP3B/SSCF/SSCOP/AAL5)
  - Narrowband SS7 (high speed and low speed)
  - SIGTRAN (M3UA/SCTP/IP)
- SS7 variants supported:
  - ITU-T (International Telecommunication Union - Telecommunications - Europe)
  - ANSI (American National Standards Institute - U.S.)
  - B-ICI (B-ISDN Inter-Carrier Interface)
  - China
  - TTC (Telecommunication Technology Committee - Japan)
  - NTT (Japan)
- SS7 protocol stack components supported:

- MTP2
- MTP3
- SCCP with BSSAP+ and RANAP
- ISUP
- TCAP and MAP

PDP Context Support

Support for subscriber primary and secondary Packet Data Protocol (PDP) contexts in compliance with 3GPP standards ensure complete end-to-end GPRS connectivity.

The SGSN supports a total of 11 PDP contexts per subscriber. Of the 11 PDP context, all can be primaries, or 1 primary and 10 secondaries or any combination of primary and secondary. Note that there must be at least one primary PDP context in order for secondaries to establish.

PDP context processing supports the following types and functions:

- Types: IPv4, IPv6, and/or PPP
- GTPP accounting support
- PDP context timers
- Quality of Service (QoS)

Mobility Management

The SGSN supports mobility management (MM) in compliance with applicable 3GPP standards and procedures to deliver the full range of services to the mobile device. Some of the procedures are highlighted below:

GPRS Attach

The SGSN is designed to accommodate a very high rate of simultaneous attaches. The actual attach rate depends on the latencies introduced by the network and scaling of peers. In order to optimize the entire signaling chain, the SGSN eliminates or minimizes bottlenecks caused by large scale control signaling. For this purpose, the SGSN implements features such as an in-memory data-VLR and SuperCharger. Both IMSI and P-TMSI based attaches are supported.

The SGSN provides the following mechanisms to control MN attaches:

- **Attached Idle Timeout** - When enabled, if an MN has not attempted to setup a PDP context since attaching, this timer forces the MN to detach with a cause indicating that the MN need not re-attach. This timer is particularly useful for reducing the number of attached subscribers, especially those that automatically attach at power-on.
- **Detach Prohibit** - When enabled, this mechanism disables the Attached Idle Timeout functionality for selected MNs which aggressively re-attach when detached by the network.
- **Prohibit Reattach Timer** - When enabled, this timer mechanism prevents MNs, that were detached due to inactivity, from re-attaching for a configured period of time. Such MNs are remembered by the in-memory data-VLR until the record needs to be purged.
- **Attach Rate Throttle** - It is unlikely that the SGSN would become a bottleneck because of the SGSN's high signaling rates. However, other nodes in the network may not scale commensurately. To provide network

overload protection, the SGSN provides a mechanism to control the number of attaches occurring through it on a per second basis.

Beside configuring the rate, it is possible to configure the action to be taken when the overload limit is reached. Note, this is a soft control and the actual attach rate may not match exactly the configured value depending on the load conditions.

GPRS Detach

The SGSN is designed to accommodate a very high rate of simultaneous detaches. However, the actual detach rate is dependent on the latencies introduced by the network and scaling of peers. A GPRS detach results in the deactivation of all established PDP contexts.

There are a variety of detaches defined in the standards and the SGSN supports the following detaches:

- **MN Initiated Detach** - The MN requests to be detached.
- **SGSN Initiated Detach** - The SGSN requests the MN to detach due to expiry of a timer or due to administrative action.
- **HLR Initiated Detach** - The detach initiated by the receipt of a cancel location from the HLR.

Mass detaches triggered by administrative commands are paced in order to avoid flooding the network and peer nodes with control traffic.

Paging

CS-Paging is initiated by a peer node - such as the MSC - when there is data to be sent to an idle or unavailable UE. CS-paging requires the Gs interface. This type of paging is intended to trigger a service request from the UE. If necessary, the SGSN can use PS-Paging to notify the UE to switch channels. Once the UE reaches the connected state, the data is forwarded to it.

Paging frequency can be controlled by configuring a paging-timer.

Service Request

The Service Request procedure is used by the MN in the PMM Idle state to establish a secure connection to the SGSN as well as request resource reservation for active contexts.

The SGSN allows configuration of the following restrictions:

- Prohibition of services
- Enforce identity check
- PLMN restriction
- Roaming restrictions

Authentication

The SGSN authenticates the subscriber via the authentication procedure. This procedure is invoked on attaches, PDP activations, inter-SGSN routing Area Updates (RAUs), and optionally on configurable periodic RAUs. The procedure requires the SGSN to retrieve authentication quintets/triplets from the HLR (AuC) and issuing an authentication and ciphering request to the MN. The SGSN implements an in-memory data-VLR functionality to pre-fetch and store authentication vectors from the HLR. This decreases latency of the control procedures.

Additional configuration at the SGSN allows for the following:

- Enforcing ciphering
- Retrieval of the IMEI-SV

P-TMSI Reallocation

The SGSN supports standard Packet-Temporary Mobile Identity (P-TMSI) Reallocation procedures to provide identity confidentiality for the subscriber.

The SGSN can be configured to allow or prohibit P-TMSI reallocation on the following events:

- Routing Area Updates
- Attaches
- Detaches
- Service Requests

The SGSN reallocates P-TMSI only when necessary.

P-TMSI Signature Reallocation

The SGSN supports operator definition of frequency and interval for Packet Temporary Mobile Subscriber Identity (P-TMSI) signature reallocation for all types of routing area update (RAU) events.

Identity Request

This procedure is used to retrieve IMSI and IMEI-SV from the MN. The SGSN executes this procedure only when the MN does not provide the IMSI and the MM context for the subscriber is not present in the SGSN's data-VLR.

Location Management

The SGSN's 3GPP compliance for location management ensures efficient call handling for mobile users.

The SGSN supports routing area updates (RAU) for location management. The SGSN implements standards based support for:

- Periodic RAUs
- Intra-SGSN RAUs
- Inter-SGSN RAUs.

The design of the SGSN allows for very high scalability of RAUs. In addition, the high capacity of the SGSN and Flex functionality provides a great opportunity to convert high impact Inter-SGSN RAUs to lower impact Intra-SGSN RAUs. The SGSN provides functionality to enforce the following RAU restrictions:

- Prohibition of GPRS services
- Enforce identity request
- Enforce IMEI check
- PLMN restriction

- Roaming restrictions

The SGSN also provides functionality to optionally supply the following information to the MN:

- P-TMSI Signature and Allocated P-TMSI
- List of received N-PDU numbers for loss less relocation
- Negotiated READY timer value
- Equivalent PLMNs
- PDP context status
- Network features supported

Session Management

Session management ensures proper PDP context setup and handling.

For session management, the SGSN supports four 3GPP-compliant procedures for processing PDP contexts:

- Activation
- Modification
- Deactivation
- Preservation

PDP Context Activation

The PDP context activation procedure establishes a PDP context with the required QoS from the MN to the GGSN. These can be either primary or secondary contexts. The SGSN supports a minimum of 1 PDP primary context per attached subscriber, and up to a maximum of 11 PDP contexts per attached subscriber.

The PDP context types supported are:

- PDP type IPv4
- PDP type IPv6
- PDP type PPP

Both dynamic and static addresses for the PDP contexts are supported.

The SGSN provides configuration to control the duration of active and inactive PDP contexts.

When activating a PDP context the SGSN can establish the GTP-U data plane from the RNC through the SGSN to the GGSN or directly between the RNC and the GGSN (one tunnel).

The SGSN is capable of interrogating the DNS infrastructure to resolve the specified APN to the appropriate GGSN. The SGSN also provides default and override configuration of QoS and APN.

PDP Context Modification

This procedure is used to update the MN and the GGSN. The SGSN is capable of initiating the context modification or negotiating a PDP context modification initiated by either the MN or the GGSN.

PDP Context Deactivation

This procedure is used to deactivate PDP contexts. The procedure can be initiated by the MN or the SGSN. The SGSN provides configurable timers to initiate PDP deactivation of idle contexts as well as active contexts.

PDP Context Preservation

The SGSN provides this functionality to facilitate efficient radio resource utilization. This functionality comes into play on the following triggers:

- **RAB (Radio Access Bearer) Release Request**

  This is issued by the RAN to request the release of RABs associated with specific PDP contexts. The SGSN responds with a RAB assignment request, waits for the RAB assignment response and marks the RAB as having been released. The retention of the PDP contexts is controlled by configuration at the SGSN. If the PDP contexts are retained the SGSN is capable of receiving downlink packets on them.

- **Iu Release Request**

  The RAN issues an Iu release request to release all RABs of an MN and the Iu connection. The retention of the PDP contexts is controlled by configuration at the SGSN. When PDP contexts are retained the SGSN is capable of receiving downlink packets on them.

  When PDP contexts are preserved, the RABs can be restored on a service request from the MN without having to go through the PDP context establishment process again. The service request is issued by the MN either when it has some data to send or in response to a paging request, on downlink data, from the SGSN.

Charging

Charging functionality for the SGSN varies depending upon the type of network in which it is deployed.

SGSN in GPRS/UMTS Network

The SGSN provides an efficient and accurate billing system for all calls and SMSs passing through the SGSN. The charging-specific interfaces and 3GPP standards supported by the SGSN deployments are listed below:

- Allows the configuration of multiple CGFs and a single GSS in a single GTPP group along with their relative priorities.
- Implements the standardized Ga interface.
- Fully supports the GPRS Tunneling Protocol Prime (GTPP) over UDP/TCP.
- Supports the relevant charging information as defined in:
  - **3GPP TS 29.060 v7.9.0 (2008-09):** Technical Specification; 3rd Generation Partnership Project; Technical Specification Group Core Network; General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface (Release 6)
  - **3GPP TS 32.215 v5.9.0 (2005-06):** 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Charging management; Charging data description for the Packet Switched (PS) domain (Release 4)
  - **3GPP TS.32.251 V8.8.0 (2009-12):** 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Charging management; Packet Switched (PS) domain charging (Release 8).

- **3GPP TS 32.298 V8.7.0 (2009-12):** 3rd Generation Partnership Project; Technical Specification Group Service and System Aspects; Telecommunication management; Charging management; Charging Data Record (CDR) parameter description (Release 8)

Charging Data Records (CDRs)

The SGSN generates CDRs with the charging information. The following sections outline the types of CDRs generated by the SGSN.

SGSN Call Detail Records (S-CDRs)

These charging records are generated for PDP contexts established by the SGSN. They contain attributes as defined in TS 32.251 v7.2.0.

Mobility Call Detail Records (M-CDRs)

These charging records are generated by the SGSN's mobility management (MM) component and correspond to the mobility states. They contain attributes as defined in 3GPP TS 32.251 v7.2.0.

Short Message Service CDRs

SGSN supports following CDRs for SMS related charging:

- SMS-Mobile Originated CDRs (SMS-MO-CDRs)
- SMS Mobile Terminated CDRs (SMS-MT-CDRs)

These charging records are generated by the SGSN's Short Message Service component. They contain attributes as defined in 3GPP TS 32.215 v5.9.0.

Features and Functionality

It is impossible to list all of the features supported by the ASR 5000 2.5G and/or 3G SGSN.

Those features listed below are only a few of the features that enable the operator to control the SGSN and their network. All of these features are either proprietary or comply with relevant 3GPP specifications.

Some of the proprietary features may require a separate license. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

The following is an alphabetical list of the features described in this overview:

- APN Aliasing
- APN Resolution with SCHAR or RNC-ID
- Automatic Protection Switching (APS)
- Authentications and Reallocations -- Selective
- Avoiding PDP Context Deactivations
- Bulk Statistics Support
- CAMEL Service Phase 3, Ge Interface
- Configurable RAB Asymmetry Indicator in RAB Assignment Request
- Direct Tunnel
- DSCP Templates for Control and Data Packets - Iu or Gb over IP
- ECMP over ATM
- Equivalent PLMN
- First Vector Configurable Start for MS Authentication
- GMM-SM Event Logging
- Gn/Gp Delay Monitoring
- GTP-C Path Failure Detection and Management
- Handling Multiple MS Attaches All with the Same Random TLLI
- Ignore Context-ID During 4G3G Handovers
- Intra- or Inter-SGSN Serving Radio Network Subsystem (SRNS) Relocation (3G only)
- Lawful Intercept
- Link Aggregation - Horizontal
- Local DNS
- Local Mapping of MBR
- Local QoS Capping
- Management System Overview
- Multiple PLMN Support

- Network Sharing
- NPU FastPath
- NRPCA - 3G
- Operator Policy
- Overcharging Protection
- QoS Traffic Policing per Subscriber
- Reordering of SNDCP N-PDU Segments
- Session Recovery
- SGSN Pooling and Iu-Flex Gb-Flex
- Short Message Service (SMS over Gd)
- SMS Authentication Repetition Rate
- SMSC Address Denial
- Threshold Crossing Alerts (TCA) Support
- Tracking Usage of GEA Encryption Algorithms
- VLR Pooling via the Gs Interface

APN Aliasing

In many situations, the APN provided in the Activation Request is unacceptable – perhaps it does not match with any of the subscribed APNs or it is misspelled – and would result in the SGSN rejecting the Activation Request. The APN Aliasing feature enables the operator to override an incoming APN – specified by a subscriber or provided during the APN selection procedure (TS 23.060) – or replace a missing APN with an operator-preferred APN.

The APN Aliasing feature provides a set of override functions: Default APN, Blank APN, APN Remapping, and Wildcard APN to facilitate such actions as:

- overriding a mismatched APN with a default APN.
- overriding a missing APN (blank APN) with a default or preferred APN.
- overriding an APN on the basis of charging characteristics.
- overriding an APN by replacing part or all of the network or operator identifier with information defined by the operator, for example, MNC123.MCC456.GPRS could be replaced by MNC222.MCC333.GPRS.
- overriding an APN for specific subscribers (based on IMSI) or for specific devices (based on IMEI).

Default APN

Operators can configure a “default APN” for subscribers not provisioned in the HLR. The default APN feature will be used in error situations when the SGSN cannot select a valid APN via the normal APN selection process. Within an APN remap table, a default APN can be configured for the SGSN to:

- override a requested APN when the HLR does not have the requested APN in the subscription profile.
- provide a viable APN if APN selection fails because there was no “requested APN” and wildcard subscription was not an option.

In either of these instances, the SGSN can provide the default APN as an alternate behavior to ensure that PDP context activation is successful.

Recently, the SGSN's default APN functionality was enhanced so that if a required subscription APN is not present in the subscriber profile, then the SGSN will now continue the activation with another configured 'dummy' APN. The call will be redirected, via the GGSN, to a webpage informing the user of the error and prompting to subscribe for services.

APN Resolution with SCHAR or RNC-ID

It is now possible to append charging characteristic information to the DNS string. The SGSN includes the profile index value portion of the CC as binary/decimal/hexadecimal digits (type based on the configuration) after the APN network identification. The charging characteristic value is taken from the subscription record selected for the subscriber during APN selection. This enables the SGSN to select a GGSN based on the charging characteristics information.

After appending the charging characteristic the DNS string will take the following form: `<apn_network_id>.<profile_index>.<apn_operator_id >`. The profile index in the following example has a value 10: `quicknet.com.uk.1010.mnc234.mcc027.gprs.`

If the RNC_ID information is configured to be a part of the APN name, and if inclusion of the profile index of the charging characteristics information is enabled before the DNS query is sent, then the profile index is included after the included RNC_ID and the DNS APN name will appear in the following form: `<apn_network_id>.<rnc_id>.<profile_index>.<apn_operator_id>`. In the following example, the DNS query for a subscriber using RNC 0321 with the profile index of value 8 would appear as: `quicknet.com.uk.0321.1000.mnc234.mcc027.gprs.`

Automatic Protection Switching (APS)

Automatic protection switching (APS is now available on an inter-card basis for SONET configured CLC2 (Frame Relay) and OLC2 (ATM) optical line cards. Multiple switching protection (MSP) version of is also available for SDH configured for the CLC2 and OLC2 (ATM) line cards.

APS/MSP offers superior redundancy for SONET/SDH equipment and supports recovery from card failures and fiber cuts. APS allows an operator to configure a pair of SONET/SDH lines for line redundancy. In the event of a line problem, the active line switches automatically to the standby line within 60 milliseconds (10 millisecond initiation and 50 millisecond switchover).

At this time, the ASR 5000 APS/MSP supports the following parameters:

- 1+1 - Each redundant line pair consists of a working line and a protection line.
- uni-directional - Protection on one end of the connection.
- non-revertive - Upon restoration of service, this parameter prevents the network from automatically reverting to the original working line.

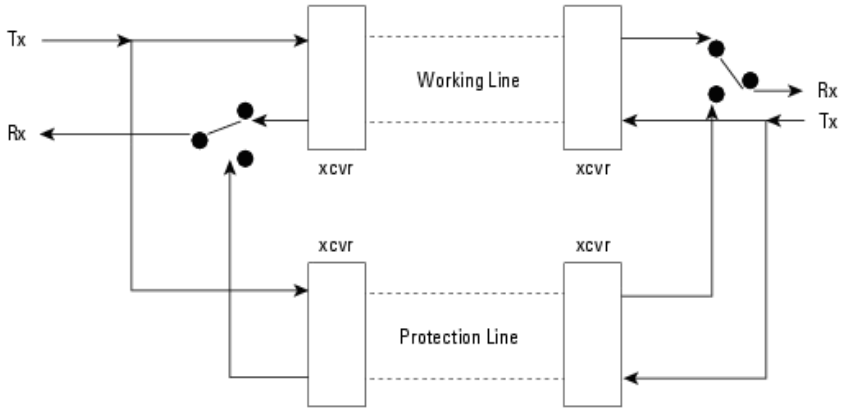
The protection mechanism used for the APS/MSP uses a linear 1+1 architecture, as described in the ITU-T G.841 standard and the Bellcore publication GR-253-CORE, SONET Transport Systems; Common Generic Criteria, Section 5.3. The connection is unidirectional.

With APS/MSP 1+1, each redundant line pair consists of a working line and a protection line. Once a signal fail condition or a signal degrade condition is detected, the hardware switches from the working line to the protection line.

With the non-revertive option, if a signal fail condition is detected, the hardware switches to the protection line and does not automatically revert back to the working line.

Since traffic is carried simultaneously by the working and protection lines, the receiver that terminates the APS/MSP 1+1 must select cells from either line and continue to forward one consistent traffic stream. The receiving ends can switch from working to protection line without coordinating at the transmit end since both lines transmit the same information.

Figure 248. SONET APS 1+1



Refer to the section on *Configuring APS/MSP Redundancy* in the *SGSN Service Configuration Procedures* chapter for configuration details.

Authentications and Reallocations -- Selective

Subscriber event authentication, P-TMSI reallocation, and P-TMSI signature reallocation are now selective rather than enabled by default.

The operator can enable and configure them to occur according to network requirements:

- every instance or every nth instance;
- on the basis of UMTS, GPRS or both;
- on the basis of elapsed time intervals between events.

There are situations in which authentication will be performed unconditionally:

- IMSI Attach – all IMSI attaches will be authenticated
- When the subscriber has not been authenticated before and the SGSN does not have a vector
- When there is a P-TMSI signature mismatch
- When there is a CKSN mismatch

There are situation in which P-TMSI will be reallocated unconditionally:

- Inter SGSN Attach/RAU
- Inter-RAT Attach/RAU in 2G
- IMSI Attach

Avoiding PDP Context Deactivations

The SGSN can be configured to avoid increased network traffic resulting from bursts of service deactivations/activations resulting from erroneous restart counter change values in received messages (Create PDP Context Response or Update PDP Context Response or Update PDP Context Request). Be default, the SGSN has the responsibility to verify possible GTP-C path failure by issuing an Echo Request/Echo Response to the GGSN. Path failure will only be confirmed if the Echo Response contains a new restart counter value. Only after this confirmation of the path failure does the SGSN begin deactivation of PDP contexts.

Bulk Statistics Support

System support for bulk statistics allows operators to choose which statistics to view and to configure the format in which the statistics are presented. This simplifies the post-processing of statistical data since it can be formatted to be parsed by external, back-end processors.

When used in conjunction with the Web Element Manager, the data can be parsed, archived, and graphed.

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. The following is the list of schemas supported for use by the SGSN:

- **System**: Provides system-level statistics
- **Card**: Provides card-level statistics
- **Port**: Provides port-level statistics
- **DLCI-Util**: Provides statistics specific to DLCIs utilization for CLC-type line cards
- **GPRS**: Provides statistics for LLC, BSSGP, SNDCP, and NS layers
- **SCCP**: Provides SCCP network layer statistics
- **SGTP**: Provides SGSN-specific GPRS Tunneling Protocol (GTP) statistics
- **SGSN**: Provides statistics for: mobility management (MM) and session management (SM) procedures; as well, MAP, TCAP, and SMS counters are captured in this schema. SGSN Schema statistic availability is per service (one of: SGSN, GPRS, MAP) and per routing area (RA)
- **SS7Link**: Provides SS7 link and linkset statistics
- **SS7RD**: Provides statistics specific to the proprietary SS7 routing domains

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the chassis or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files.

The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, chassis host name, chassis uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

When the Web Element Manager is used as the receiver, it is capable of further processing the statistics data through XML parsing, archiving, and graphing.

The Bulk Statistics Server component of the Web Element Manager parses collected statistics and stores the in formation in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local file system or on an NFS-mounted file system on the Web Element Manager server.

CAMEL Service Phase 3, Ge Interface

The ASR 5000 SGSN provides PDP session support as defined by Customized Applications for Mobile network Enhanced Logic (CAMEL) phase 3.

CAMEL Service

CAMEL service enables operators of 2.5G/3G networks to provide operator-specific services (such as prepaid GPRS service and prepaid SMS service) to subscribers, even when the subscribers are roaming outside their HPLMN.

CAMEL Support

ASR 5000 SGSN support for CAMEL phase 3 services expands with each SGSN application release. Current support enables operators of 2.5G/3G networks to provide operator-specific services (such as prepaid GPRS service and prepaid SMS service) to subscribers, even when the subscribers are roaming outside their HPLMN.

For this release the SGSN has expanded its support for CAMEL Scenario 1 adding:

- Implementation of Scenario 1 triggers (TDP-Attach, TDP-Attach-ChangeofPosition)
- Implementation of Scenario 1 Dynamic triggers (DP-Detach, DP-ChangeofPosition)
- Expanded conformance to 3GPP spec 23.078 (Release 4)

The ASR 5000 SGSN supports the following GPRS-related functionality in CAMEL phase 3:

- Control of GPRS PDP contexts

Functional support for CAMEL interaction includes:

- PDP Context procedures per 3GPP TS 29.002
    - GPRS TDP (trigger detection point) functions
    - Default handling codes, if no response received from SCP
    - GPRS EDP (event detection points) associated with SCP
    - Charging Procedures: Handle Apply Charging GPRS & Handle Apply Charging Report GPRS
- “GPRS Dialogue scenario 2" for CAMEL control with SCP
- CAMEL-related data items in an S-CDR:
    - SCF Address
    - Service Key
    - Default Transaction Handling
    - Level of CAMEL service (phase 3)
- Session Recovery for all calls have an ESTABLISHED CAMEL association.

Ge Interface

The ASR 5000 implementation of CAMEL uses standard CAP protocol over a Ge interface between the SGSN and the SCP. This interface can be deployed over SS7 or SIGTAN.

The SGSN's Ge support includes use of the gprsSSF CAMEL component with the SGSN and the gsmSCF component with the SCP.

CAMEL Configuration

To provide the CAMEL interface on the SGSN, a new service configuration mode, called "CAMEL Service", has been introduced on the SGSN.

1. An SCCP Network configuration must be created or exist already.
2. A CAMEL Service instance must be created.
3. The CAMEL Service instance must be associated with either the SGSN Service configuration or the GPRS Service configuration in order to enable use of the CAMEL interface.
4. The CAMEL Service must be associated with the SCCP Network configuration.

Until a CAMEL Service is properly configured, the SGSN will not process any TDP for pdp-context or mo-sms.

Configurable RAB Asymmetry Indicator in RAB Assignment Request

The SGSN sets the value for the RAB Asymmetry Indicator that is included in the RAB Assignment Request.

In releases prior to R12.0, the SGSN set the RAB asymmetry indicator to "Symmetric-Bidirectional" when downlink and uplink bit rates were equal. Now, the SGSN selects the value based on the symmetry of negotiated maximum bitrates as follows:

- If the uplink and downlink bitrates are equal then it is set to "Symmetric-Bidirectional",
- If uplink bitrate is set to 0 kbps, then it is set to "Asymmetric-Unidirectional-Downlink",
- If downlink bitrate is set to 0 kbps, then it is set to "Asymmetric-Unidirectional-Uplink",
- If the uplink and downlink bitrates are non-zero and different, then it is set to "Asymmetric-Bidirectional".

A change in CLI configuration allows the SGSN to override the above functionality and set the RAB Asymmetry Indicator to "Asymmetric-Bidirectional" when uplink and downlink bitrates are equal. As a result, two sets of bitrates - one for downlink and one for uplink - will be included in the RAB Assignment Requests as mandated in 3GPP TS 25.413.

Direct Tunnel

In accordance with standards, one tunnel functionality enables the SGSN to establish a direct tunnel at the user plane level - a GTP-U tunnel, directly between the RAN and the GGSN. Feature details and configuration procedures are provided in the *Direct Tunnel* feature chapter in this guide.

DSCP Templates for Control and Data Packets - Iu or Gb over IP

The SGSN supports a mechanism for differentiated services code point (DSCP) marking of control packets and signaling messages for the SGSN's LLC messages for the Gb interface.

This DSCP marking feature enables the SGSN to perform classifying and managing of network traffic and to determine quality of service (QoS) for the interfaces to an IP network.

ECMP over ATM

Iu Redundancy is the ASR 5000's implementation of equal-cost multi-path routing (ECMP) over ATM.

Iu Redundancy is based on the standard ECMP multi-path principle of providing multiple next-hop-routes of equal cost to a single destination for packet transmission. ECMP works with most routing protocols and can provide increased bandwidth when traffic load-balancing is implemented over multiple paths.

ECMP over ATM will create an ATM ECMP group when multiple routes with different destination ATM interfaces are defined for the same destination IP address. When transmitting a packet with ECMP, the NPU performs a hash on the packet header being transmitted and uses the result of the hash to index into a table of next hops. The NPU looks up the ARP index in the ARP table (the ARP table contains the next-hop and egress interfaces) to determine the next-hop and interface for sending packets.

Equivalent PLMN

This feature is useful when an operator deploys both GPRS and UMTS access in the same radio area and each radio system broadcasts different PLMN codes. It is also useful when operators have different PLMN codes in different geographical areas, and the operators' networks in the various geographical areas need to be treated as a single HPLMN.

This feature allows the operator to consider multiple PLMN codes for a single subscriber belonging to a single home PLMN (HPLMN). This feature also allows operators to share infrastructure and it enables a UE with a subscription with one operator to access the network of another operator.

First Vector Configurable Start for MS Authentication

Previously, the SGSN would begin authentication towards the MS only after the SGSN received all requested vectors. This could result in a radio network traffic problem when the end devices timed out and needed to re-send attach requests.

Now, the SGSN can be configured to start MS authentication as soon as it receives the first vector from the AuC/HLR while the SAI continues in parallel. After an initial attach request, some end devices restart themselves after waiting for the PDP to be established. In such cases, the SGSN restarts and a large number of end devices repeat their attempts to attach. The attach requests flood the radio network, and if the devices timeout before the PDP is established then they continue to retry, thus even more traffic is generated. This feature reduces the time needed to retrieve vectors over the GR interface to avoid the high traffic levels during PDP establishment and to facilitate increased attach rates.

GMM-SM Event Logging

To facilitate troubleshooting, the SGSN will capture procedure-level information per 2G or 3G subscriber (IMSI-based) in CSV formatted event data records (EDRs) that are stored on an external server.

This feature logs the following events:

- Attaches
- Activation of PDP Context

- RAU
- ISRAU
- Deactivation of PDP Context
- Detaches
- Authentications
- PDP Modifications

The new SGSN event logging feature is enabled/disabled per service with via the CLI commands. For more information on this feature, refer to the chapter *GMM/SM Event Logging* in this guide.

Gn/Gp Delay Monitoring

The SGSN measures the control plane packet delay for GTP-C signaling messages on the SGSN's Gn/Gp interface towards the GGSN.

If the delay crosses a configurable threshold, an alarm will be generated to prompt the operator.

A delay trap is generated when the GGSN response to an ECHO message request is delayed more than a configured amount of time and for a configured number of consecutive responses. When this occurs, the GGSN will be flagged as experiencing delay.

A clear delay trap is generated when successive ECHO Response (number of successive responses to detect a delay clearance is configurable), are received from a GGSN previously flagged as experiencing delay.

This functionality can assist with network maintenance, troubleshooting, and early fault discovery.

GTP-C Path Failure Detection and Management

The SGSN now provides the ability to manage GTP-C path failures detected as a result of spurious restart counter change messages received from the GGSN.

Previous Behavior: The old default behavior was to have the Session Manager (SessMgr) detect GTP-C path failure based upon receiving restart counter changes in messages (Create PDP Context Response or Update PDP Context Response or Update PDP Context Request) from the GGSN and immediately inform the SGTPC Manager (SGTPCMgr) to pass the path failure detection to all other SessMgrs so that PDP deactivation would begin.

New Behavior: The new default behavior has the SessMgr inform the SGTPCMgr of the changed restart counter value. The SGTPCMgr now has the responsibility to verify a possible GTP-C path failure by issuing an Echo Request/Echo Response to the GGSN. Path failure will only be confirmed if the Echo Response contains a new restart counter value. Only after this confirmation of the path failure does the SGTPCMgr inform all SessMgrs so that deactivation of PDP contexts begins.

Handling Multiple MS Attaches All with the Same Random TLLI

Some machine-to-machine (M2M) devices from the same manufacturer will all attempt PS Attaches using the same fixed random Temporary Logical Link Identifier (TLLI).

The SGSN cannot distinguish between multiple M2M devices trying to attach simultaneously using the same random TLLI and routing area ID (RAI). As a result, during the attach process of an M2M device, if a second device tries to

attach with the same random TLLI, the SGSN interprets that as an indication that the original subscriber moved during the Attach process and the SGSN starts communicating with the second device and drops the first device.

The SGSN can be configured to allow only one subscriber at a time to attach using a fixed random TLLI. While an Attach procedure with a fixed random TLLI is ongoing (that is, until a new P-TMSI is accepted by the MS), all other attaches sent to the SGSN with the same random TLLI using a different IMSI will be dropped by the SGSN's Linkmgr.

To limit the wait-time functionality to only the fixed random TLLI subscribers, the TLLI list can be configured to control which subscribers will be provided this functionality.

HSPA Fallback

Besides enabling configurable support for either 3GPP Release 6 (HSPA) and 3GPP Release 7 (HSPA+) to match whatever the RNCs support, this feature enables configurable control of data rates on a per RNC basis. This means that operators can allow subscribers to roam in and out of coverages areas with different QoS levels.

The SGSN can now limit data rates (via QoS) on a per-RNC basis. Some RNCs support HSPA rates (up to 16 Mbps in the downlink and 8 Mbps in the uplink) and cannot support higher data rates - such as those enabled by HSPA+ (theoretically, up to 256 Mbps both downlink and uplink). Being able to specify the QoS individually for each RNC makes it possible for operators to allow their subscribers to move in-and-out of coverage areas with different QoS levels, such as those based on 3GPP Release 6 (HSPA) and 3GPP Release 7 (HSPA+).

For example, when a PDP context established from an RNC with 21 Mbps is handed off to an RNC supporting only 16 Mbps, the end-to-end QoS will be re-negotiated to 16 Mbps. Note that an MS/UE may choose to drop the PDP context during the QoS renegotiation to a lower value.

This data rate management per RNC functionality is enabled, in the radio network controller (RNC) configuration mode, by specifying the type of 3GPP release specific compliance, either release 7 for HSPA+ rates or pre-release 7 for HSPA rates.

Ignore Context-ID During 4G/3G Handovers

HSS and HLR, when operating as separate network nodes, are required to use the same context-ID for a given APN-configuration of a subscriber. During inter-RAT cell reselections and handovers between 2G/3G and 4G, if the SGSN does not find a matching APN-configuration for the given context-ID learnt from the the peer node, then the PDP does not get established. This could result in SRNS relocation failures when none of the PDP's learnt from the SGSN has a matching context-ID in the HLR.

New commands have been added to enable the operator to configure the SGSN to ignore the context-ID provided by the peer and to use the PDP-type and address information to search through HLR subscription and to update the context-ID information within the PDP. For details, refer to the description for the **`rau-inter`** command under the *Call-Control Profile Configuration Mode Commands* section of the Cisco ASR 5x00 GPRS / UMTS Command Line Interface Reference.

Intra- or Inter-SGSN Serving Radio Network Subsystem (SRNS) Relocation (3G only)

Implemented according to 3GPP standard, the SGSN supports both inter- and intra-SGSN RNS relocation (SRNS) to enable handover of an MS from one RNC to another RNC.

The relocation feature is triggered by subscribers (MS/UE) moving from one RNS to another. If the originating RNS and destination RNS are connected to the same SGSN but are in different routing areas, the behavior triggers an intra-

SGSN Routing Area Update (RAU). If the RNS are connected to different SGSNs, the relocation is followed by an inter-SGSN RAU. This feature is configured through the Call-Control Profile Configuration Mode which is part of the feature set.

Lawful Intercept

The Cisco Lawful Intercept feature is supported on the SGSN. Lawful Intercept is a license-enabled, standards-based feature that provides telecommunications service providers with a mechanism to assist law enforcement agencies in monitoring suspicious individuals for potential illegal activity. For additional information and documentation on the Lawful Intercept feature, contact your Cisco account representative.

Link Aggregation - Horizontal

The SGSN supports enhanced link aggregation (LAG) within ports on different XGLCs. Ports can be from multiple XGLCs. LAG works by exchanging control packets (Link Aggregation Control Marker Protocol) over configured physical ports with peers to reach agreement on an aggregation of links. LAG sends and receives the control packets directly on physical ports attached to different XGLCs. The link aggregation feature provides higher aggregated bandwidth, auto-negotiation, and recovery when a member port link goes down.

Local DNS

Previously, the SGSN supported GGSN selection for an APN only through operator policy, and supported a single pool of up to 16 GGSN addresses which were selected in round robin fashion.

The SGSN now supports configuration of multiple pools of GGSNs; a primary pool and a secondary. As part of DNS resolution, the operator can use operator policies to prioritize local GGSNs versus remote ones. This function is built upon existing load balancing algorithms in which weight and priority are configured per GGSN, with the primary GGSN pool used first and the secondary used if no primary GGSNs are available.

The SGSN first selects a primary pool and then GGSNs within that primary pool; employing a round robin mechanism for selection. If none of the GGSNs in a pool are available for activation, then the SGSN proceeds with activation selecting a GGSN from a secondary pool on the basis of assigned weight. A GGSN is considered unavailable when it does not respond to GTP Requests after a configurable number of retries over a configurable time period. Path failure is detected via GTP-echo.

Local Mapping of MBR

The SGSN provides the ability to map a maximum bit rate (MBR) value (provided by the HLR) to an HSPA MBR value.

The mapped value is selected based on the matching MBR value obtained from the HLR subscription. QoS negotiation then occurs based on the converted value.

This feature is available within the operator policy framework. MBR mapping is configured via new keywords added to the `qos class` command in the APN Profile configuration mode. A maximum of four values can be mapped per QoS per APN.

Important: To enable this feature the `qos prefer-as-cap`, also a command in the APN Profile configuration mode, must be set to either `both-hlr-and-local` or to `hlr subscription`.

Local QoS Capping

The operator can configure a cap or limit for the QoS bit rate.

The SGSN can now be configured to cap the QoS bit rate parameter when the subscribed QoS provided by the HLR is lower than the locally configured value.

Depending upon the keywords included in the command, the SGSN can:

- take the QoS parameter configuration from the HLR configuration.
- take the QoS parameter configuration from the local settings for use in the APN profile.
- during session establishment, apply the lower of either the HLR subscription or the locally configured values.

Management System Overview

The system's management capabilities are designed around the Telecommunications Management Network (TMN) model for management - focusing on providing superior quality network element (NE) and element management system (Web Element Manager) functions. The system provides element management applications that can easily be integrated, using standards-based protocols (CORBA and SNMPv1, v2), into higher-level management systems - giving wireless operators the ability to integrate the system into their overall network, service, and business management systems. In addition, all management is performed out-of-band for security and to maintain system performance.

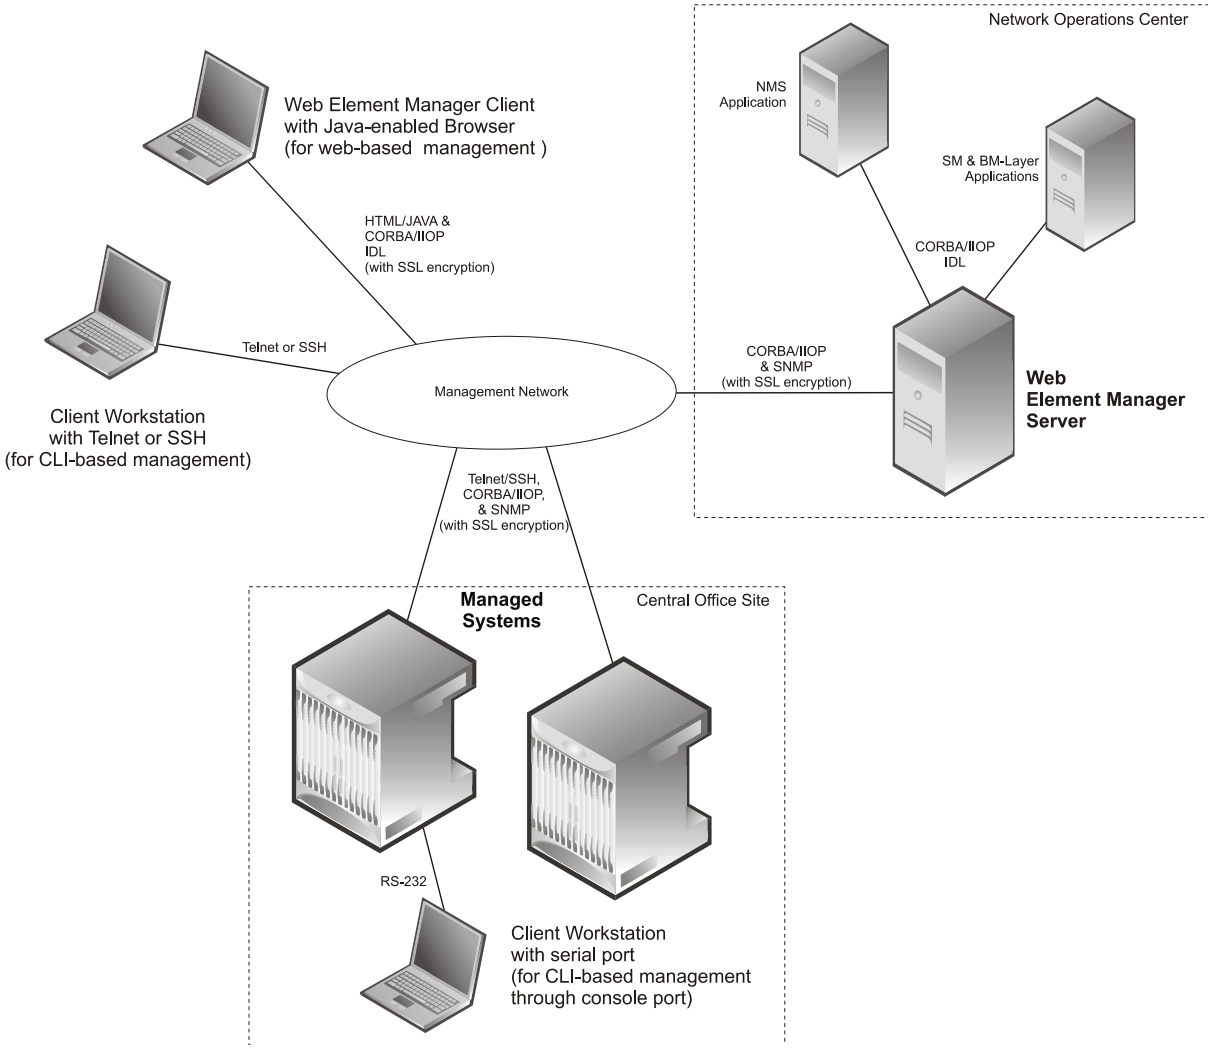
The Operation and Maintenance module of the system offers comprehensive management capabilities to the operators and enables them to operate the system more efficiently. There are multiple ways to manage the system either locally or remotely using its out-of-band management interfaces.

These include:

- Using the command line interface (CLI)
- Remote login using Telnet, and Secure Shell (SSH) access to CLI through SPIO card's Ethernet management interfaces
- Local login through the Console port on SPIO card using an RS-232 serial connection
- Using the Web Element Manager (WEM) application (requires a separate license)
- Supports communications through 10 Base-T, 100 Base-TX, 1000 Base-TX, or 1000
- Base-SX (optical gigabit Ethernet) Ethernet management interfaces on the SPIO
- Client-Server model supports any browser (i.e. Microsoft Internet Explorer v5.0 and above or Netscape v4.7 or above, and others)
- Supports Common Object Request Broker Architecture (CORBA) protocol and Simple Network Management Protocol version 1 (SNMPv1) for fault management
- Provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capabilities
- Can be easily integrated with higher-level network, service, and business layer applications using the Object Management Group's (OMG's) Interface Definition Language (IDL)

The following figure demonstrates these various element management options and how they can be utilized within the wireless carrier network.

Figure 249. Element Management Methods



Important: By default, SGSN management functionality is enabled for console-based access.

Multiple PLMN Support

With this feature, the 2.5G and 3G SGSNs now support more than one PLMN ID per SGSN. Multiple PLMN support facilitates MS handover from one PLMN to another PLMN.

Multiple PLMN support also means an operator can 'hire out' their infrastructure to other operators who may wish to use their own PLMN IDs. As well, multiple PLMN support enables an operator to assign more than one PLMN ID to a cell-site or an operator can assign each cell-site a single PLMN ID in a multi-cell network (typically, there are no more than 3 or 4 PLMN IDs in a single network).

This feature is enabled by configuring, within a single context, multiple instances of either an IuPS service for a single 3G SGSN service or multiple GPRS services for a 2.G SGSN. Each IuPS service or GPRS service is configured with a unique PLMN ID. Each of the SGSN and/or GPRS services must use the same MAP, SGTPU and GS services so these only need to be defined one-time per context.

Network Sharing

In accordance with 3GPP TS 23.251, the SGSN provides an operator the ability to share the RAN and/or the core network with other operators. Depending upon the resources to be shared, there are 2 network sharing modes of operation: the Gateway Core Network (GWCN) and the Multi-Operator Core Network (MOCN).

Benefits of Network Sharing

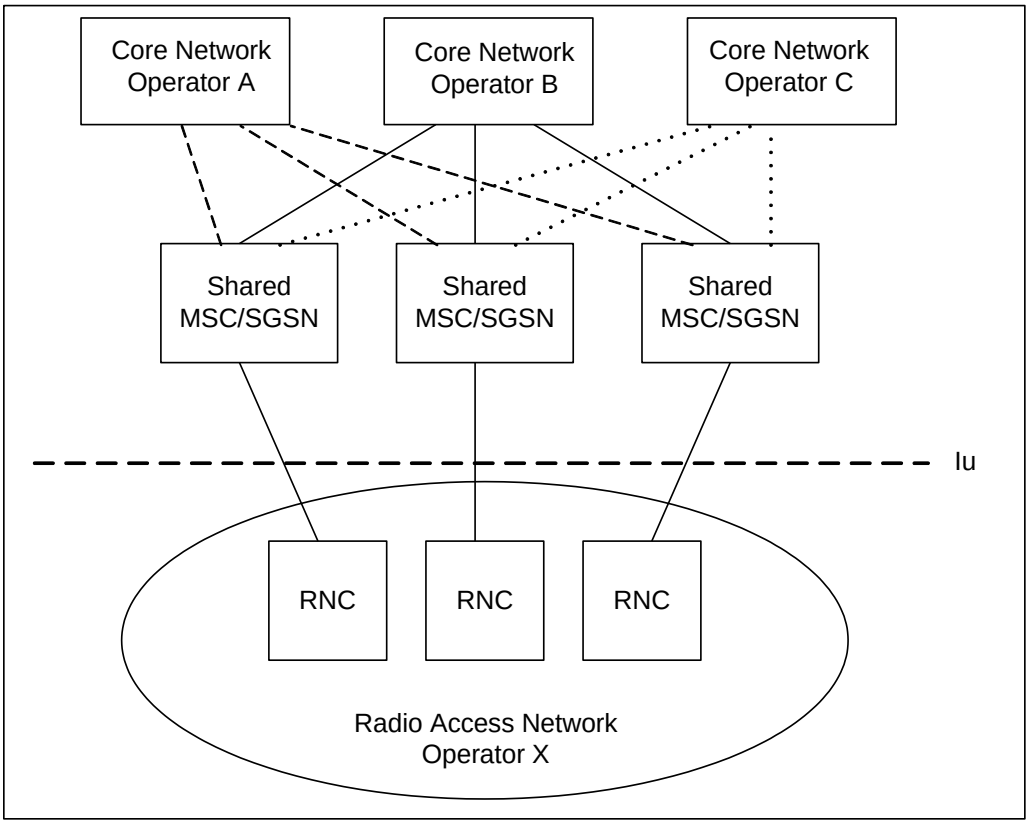
Network sharing provides operators with a range of logistical and operational benefits:

- Enables two or more network operators to share expensive common network infrastructure.
- A single operator with multiple MCC-MNC Ids can utilize a single physical access infrastructure and provide a single HPLMN view to the UEs.
- Facilitates implementation of MVNOs.

GWCN Configuration

With a gateway core network configuration, the complete radio access network and part of the core network are shared (for example, MSC/SGSN) among different operators, while each operator maintains its own separate network nodes (for example, GGSN/HLR).

Figure 250. GWCN-type Network Sharing



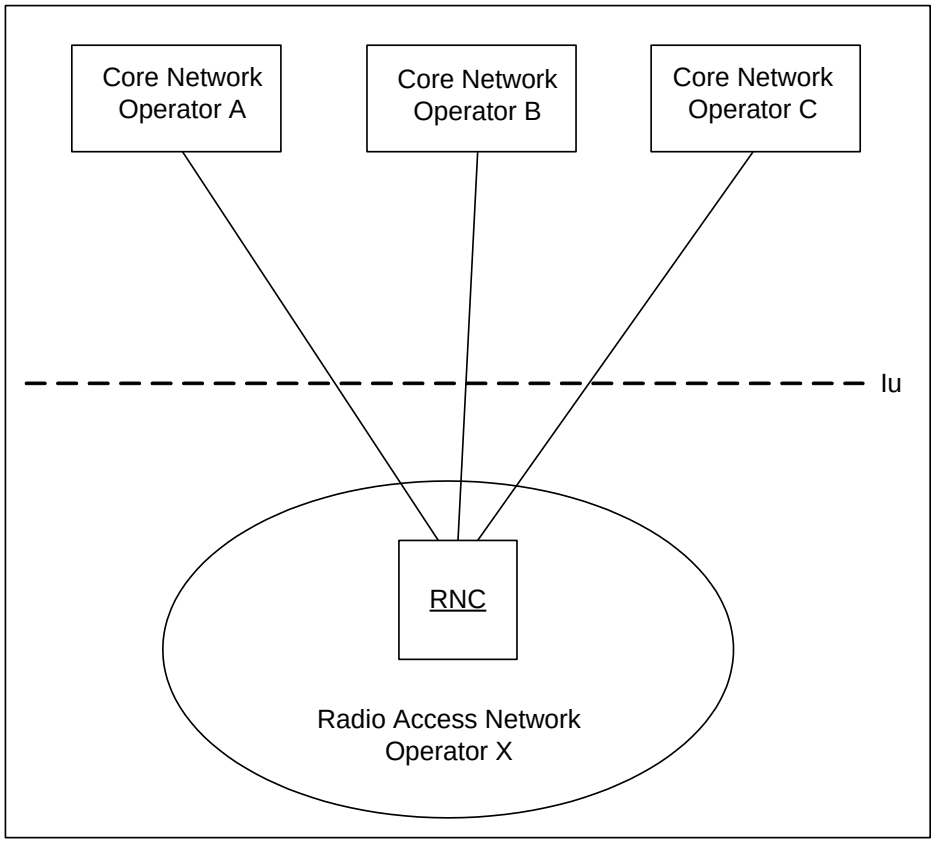
With the GWCN configuration, the SGSN supports two scenarios:

- GWCN with non-supporting UE
- GWCN with supporting UE

MOCN Configuration

In the multi-operator core network configuration, the complete radio network is shared among different operators, while each operators maintains its own separate core network.

Figure 251. MOCN-type Network Sharing



With the MOCN configuration, the SGSN supports the following scenarios:

- MOCN with non-supporting UE
- MOCN with supporting UE

Implementation

To facilitate network sharing, the SGSN implements the following key features:

- Multiple virtual SGSN services in a single physical node.
- Sharing operators can implement independent policies, such as roaming agreements.
- Equivalent PLMN configuration.
- RNC identity configuration allows RNC-ID + MCC-MNC instead of just RNC-ID.

Configuration for network sharing is accomplished by defining:

- NRI in the SGSN service configuration mode
- PLMN IDs and RNC IDs in the IuPS configuration mode
- Equivalent PLMN IDs and configured in the Call-Control Profile configuration mode.
- IMSI ranges are defined in the SGSN-Global configuration mode

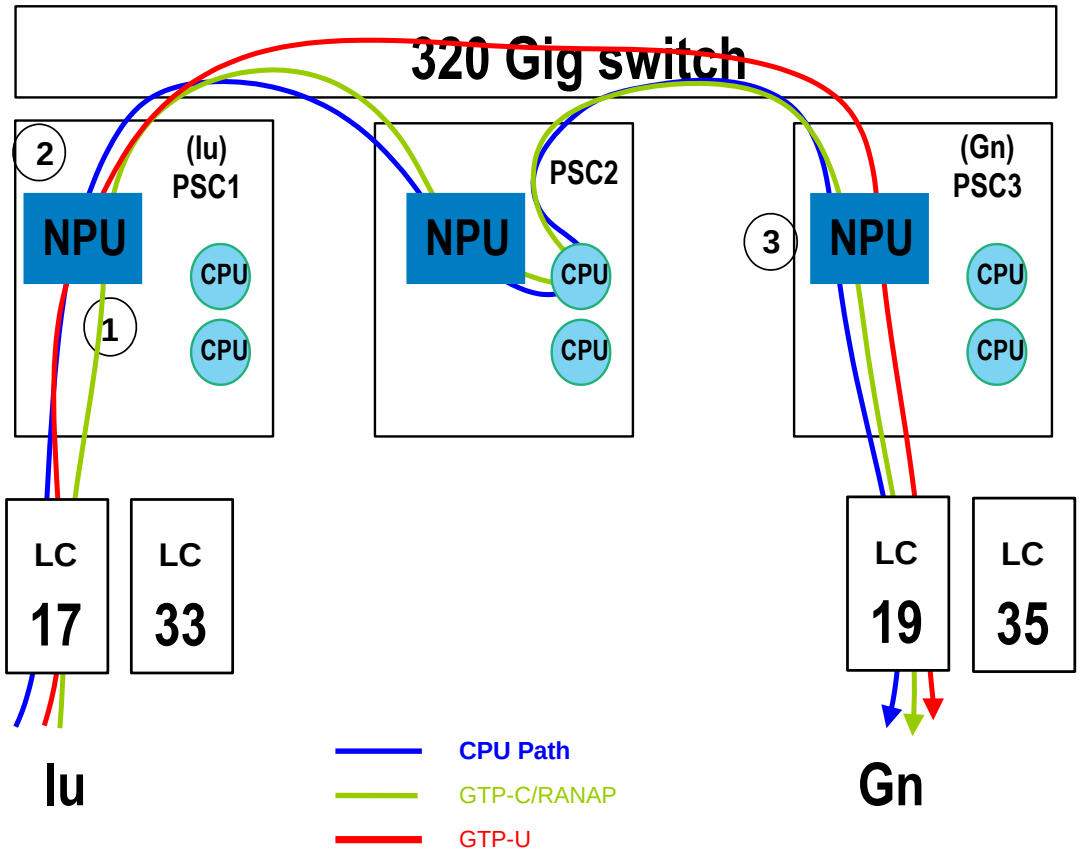
- The Call-Control Profile and IMSI ranges are associated in the configuration mode.

NPU FastPath

NPU FastPath's proprietary internal direct tunnel optimizes resource usage and reduces latency when processing GTP-U packets. This proprietary feature is only available on the ASR 5000 SGSN.

Incoming traffic passes through the switch fabric and the routing headers are changed to re-route traffic from the incoming network processing unit (NPU) of the ingress packet processing card directly to the outgoing NPU of the egress packet processing card. This means that intervening NPUs and CPUs are by-passed. This provides the SGSN with router-like latency and increased node signaling capacity.

Figure 252. SGSN NPU FastPath



FastPath is established when both ends of a tunnel are available. Two FastPath flows are established, one for the uplink and one for the downlink direction for a given PDP context. FastPath will temporarily go down or be disengaged so that packets temporarily do not move through FastPath when either an Intra-SGSN RAU or an Iu-Connection Release occurs.

If FastPath cannot be established, the NPU forwards the GTP-U packets to a CPU for processing and they are processed like all other packets.

FastPath can not be established for subscriber PDP sessions if:

- Traffic Policing and Shaping is enabled.

- Subscriber Monitoring is enabled.
- Lawful Intercept (LI) is enabled,
- IP Source Violation Checks are enabled.
- GTP-v0 tunnel is established with an GGSN.

For NPU fast path configuration, refer to Enabling NPU Fast Path for GTP-U Processing section of "Service Configuration Procedures" chapter of *Cisco ASR 5000 Series Serving GPRS Support Node Administration Guide*.

NRPCA - 3G

The SGSN supports the Network Requested Primary PDP Context Activation (NRPCA) procedure for 3G attachments.

There are no interface changes to support this feature. Support is configured with existing CLI commands (network-initiated-pdp-activation, location-area-list) in the call control profile configuration mode and timers (`T3385-timeout` and `max-actv-retransmission`) are set in the SGSN service configuration mode.

Operator Policy

The non-standard feature is unique to the ASR 5000 SGSN. This feature empowers the carrier with unusual and flexible control to manage functions that aren't typically used in all applications and to determine the granularity of the implementation of any: to groups of incoming calls or to simply one single incoming call. For details about the feature, its components, and how to configure it, refer to the *Operator Policy* chapter in this guide.

Important: **SGSN configurations created prior to Release 11.0 are not forward compatible.** All configurations for SGSNs, with -related configurations that were generated with software releases prior to Release 11.0, must be converted to enable them to operate with an SGSN running Release 11.0 or higher. Your Cisco Representative can accomplish this conversion for you.

Some Features Managed by Operator Policies

The following is a list of some of the features and functions that can be controlled via configuration of Operator Policies:

- APN Aliasing
- Authentication
- Direct Tunnel - for feature description and configuration details, refer to the *Direct Tunnel* chapter in this guide
- Equivalent PLMN
- IMEI Override
- Intra- or Inter-SGSN Serving Radio Network Subsystem (SRNS) Relocation (3G only)
- Network Sharing
- QoS Traffic Policing per Subscriber
- SGSN Pooling - Gb/Iu Flex
- SuperCharger

- Subscriber Overcharging Protection - for feature description and configuration details, refer to the *Subscriber Overcharging Protection* chapter in this guide.

Overcharging Protection

Overcharging Protection enables the SGSN to avoid overcharging the subscriber if/when a loss of radio coverage (LORC) occurs in a UMTS network. For details and configuration information, refer to the *Subscriber Overcharging Protection* chapter in this book.

QoS Traffic Policing per Subscriber

Traffic policing enables the operator to configure and enforce bandwidth limitations on individual PDP contexts for a particular traffic class.

Traffic policing typically deals with eliminating bursts of traffic and managing traffic flows in order to comply with a traffic contract.

The SGSN conforms to the DiffServ model for QoS by handling the 3GPP defined classes of traffic, QoS negotiation, DSCP marking, traffic policing, and support for HSDPA/HSUPA.

QoS Classes

The 3GPP QoS classes supported by the SGSN are:

- Conversational
- Streaming
- Interactive
- Background

The SGSN is capable of translating between R99 and R97/98 QoS attributes.

QoS Negotiation

On PDP context activation, the SGSN calculates the QoS allowed, based upon:

- **Subscribed QoS -** This is a per-APN configuration, obtained from the HLR on an Attach. It specifies the highest QoS allowed to the subscriber for that APN.
- **Configured QoS -** The SGSN can be configured with default and highest QoS profiles in the configuration.
- **MS requested QoS -** The QoS requested by the UE on pdp-context activation.

DSCP Marking

The SGSN performs diffserv code point (DSCP) marking of the GTP-U packets according to allowed-QoS to PHB mapping. The default mapping matches that of the UMTS to IP QoS mapping defined in 3GPP TS 29.208.

The SGSN also supports DSCP marking of the GTP control plane messages on the Gn/Gp interface. This allows QoS to be set on GTP-C messages, and is useful if Gn/Gp is on a less than ideal link. DSCP marking is configurable via the CLI, with default = Best Effort Forwarding.

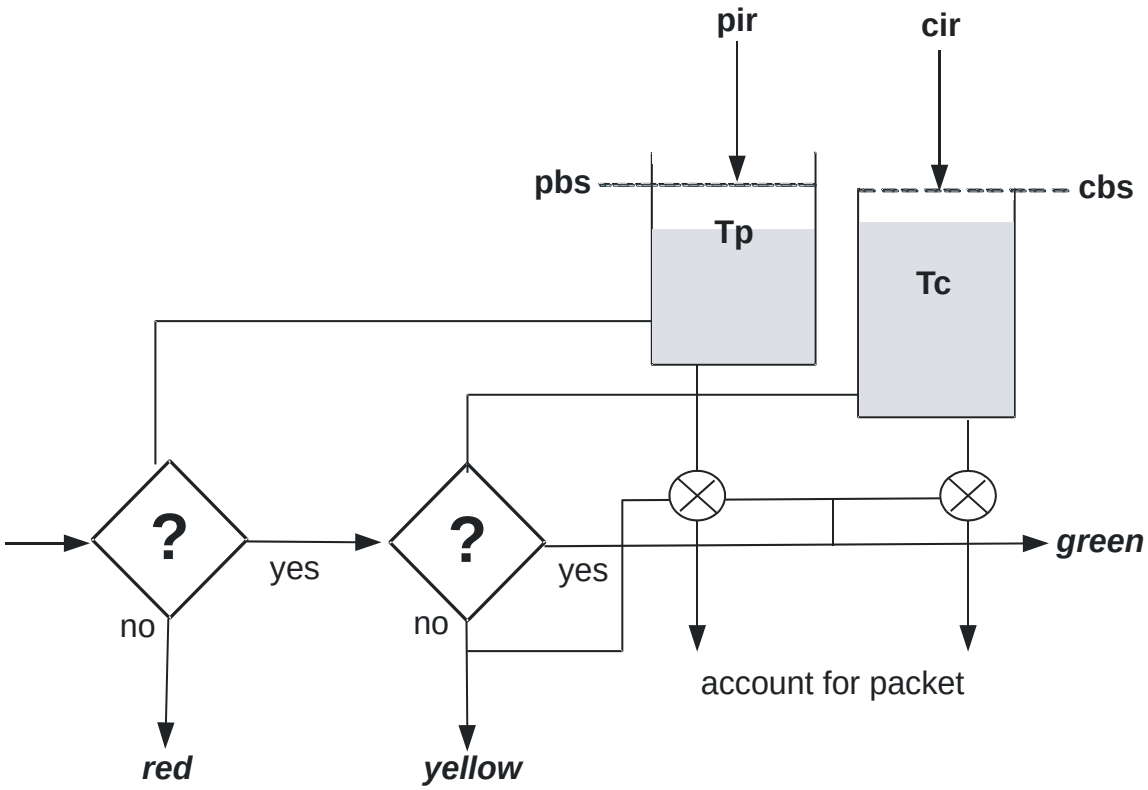
Traffic Policing

The SGSN can police uplink and downlink traffic according to predefined QoS negotiated limits fixed on the basis of individual contexts - either primary or secondary. The SGSN employs the Two Rate Three Color Marker (RFC2698) algorithm for traffic policing. The algorithm meters an IP packet stream and marks its packets either green, yellow, or red depending upon the following variables:

- **PIR -** Peak Information Rate (measured in bytes/second)
- **CIR -** Committed Information Rate (measured in bytes/second)
- **PBS -** Peak Burst Size (measured in bytes)
- **CBS -** Committed Burst Size (measured in bytes)

The following figure depicts the working of the TCM algorithm:

Figure 253. TCM Algorithm Logic for Traffic Policing



Reordering of SNDCP N-PDU Segments

The SGSN fully supports reordering of out-of-order segments coming from the same SNDCP N-PDU. The SGSN waits the configured amount of time for all segments of the N-PDU to arrive. If all the segments are not received before the timer expiries, then all queued segments are dropped.

Session Recovery

Session recovery provides a seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault that prevents a fully attached user session from having the PDP contexts removed or the attachments torn down.

Session recovery is performed by mirroring key software processes (e.g., session manager and AAA manager) within the system. These mirrored processes remain in an idle state (in standby-mode) until they may be needed in the case of a software failure (e.g., a session manager task aborts). The system spawns new instances of "standby mode" session and AAA managers for each active control processor (CP) being used.

As well, other key system-level software tasks, such as VPN manager, are performed on a physically separate packet processing card to ensure that a double software fault (e.g., session manager and VPN manager fail at the same time on the same card) cannot occur. The packet processing card used to host the VPN manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

The additional hardware resources required for session recovery include a standby System Management Card and a standby packet processing card.

There are two modes for Session Recovery.

- **Task recovery mode:** One or more session manager failures occur and are recovered without the need to use resources on a standby packet processor card. In this mode, recovery is performed by using the mirrored "standby-mode" session manager task(s) running on active packet processor cards. The "standby-mode" task is renamed, made active, and is then populated using information from other tasks such as AAA manager.
- **Full packet processing card recovery mode:** Used when a packet processing card hardware failure occurs, or when a packet processor card migration failure happens. In this mode, the standby packet processor card is made active and the "standby-mode" session manager and AAA manager tasks on the newly activated packet processor card perform session recovery.

Session/Call state information is saved in the peer AAA manager task because each AAA manager and session manager task is paired together. These pairs are started on physically different packet processor cards to ensure task recovery.

When session recovery occurs, the system reconstructs the following subscriber information:

- Data and control state information required to maintain correct call behavior
- Subscriber data statistics that are required to ensure that accounting information is maintained
- A best-effort attempt to recover various timer values such as call duration, absolute time, and others

For more information on session recovery use and session recovery configuration, refer to the *Session Recovery* chapter in the *Cisco ASR 5000 Series System Administration Guide*.

SGSN Pooling and Iu-Flex / Gb-Flex

This implementation allows carriers to load balance sessions among pooled SGSNs, to improve reliability and efficiency of call handling, and to use Iu-Flex / Gb-Flex to provide carriers with deterministic failure recovery.

The SGSN, with its high capacity, signaling performance, and peering capabilities, combined with its level of fault tolerance, delivers many of the benefits of Flex functionality even without deploying SGSN pooling.

As defined by 3GPP TS 23.236, the SGSN implements Iu-Flex and Gb-Flex functionality to facilitate network sharing and to ensure SGSN pooling for 2.5G and 3G accesses as both separate pools and as dual-access pools.

SGSN pooling enables the following:

- Eliminates the single point of failure between an RNC and an SGSN or between a BSS and an SGSN.

- Ensures geographical redundancy, as a pool can be distributed across sites.
- Minimizes subscriber impact during service, maintenance, or node additions or replacements.
- Increases overall capacity via load sharing across the SGSNs in a pool.
- Reduces the need/frequency for inter-SGSN RAUs. This substantially reduces signaling load and data transfer delays.
- Supports load redistribution with the SGSN offloading procedure.

Gb/Iu Flex Offloading

The SGSN supports Gb/Iu Flex subscriber offloading from one SGSN to another specific SGSN in a 2G/3G pool.

In addition, the operator can configure the offloading Target NRI in P-TMSI, and the quantity to offload to the Target. This can be used to provide load balancing, or to offload a single node in pool, take it out of service for whatever reason (e.g., maintenance).

Short Message Service (SMS over Gd)

The SGSN implements a configurable Short Message Service (SMS) to support sending and receiving text messages up to 140 octets in length. The SGSN handles multiple, simultaneous messages of both types: those sent from the MS/UE (SMS-MO: mobile originating) and those sent to the MS/UE (SMS-MT: mobile terminating). Short Message Service is disabled by default.

After verifying a subscription for the PLMN's SMS service, the SGSN connects with the SMSC (short message service center), via a Gd interface, to relay received messages (from a mobile) using MAP-MO-FORWARD-REQUESTs for store-and-forward.

In the reverse, the SGSN awaits messages from the SMSC via MAP-MT-FORWARD-REQUESTs and checks the subscriber state before relaying them to the target MS/UE.

The SGSN will employ both the Page procedure and MNRG (mobile not reachable for GPRS) flags in an attempt to deliver messages to subscribers that are absent.

The SGSN supports

- charging for SMS messages, and
- lawful intercept of SMS messages

SMS Authentication Repetition Rate

The SGSN provides an authentication procedures for standard GMM events like Attach, Detach, RAU, and Service-Request, and SMS events such as Activate, all with support for 1-in-N Authenticate functionality. The SGSN did not provide the capability to authenticate MO/MT SMS events.

Now, the authentication functionality has been expanded to the Gs interface where the SGSN now supports configuration of the authentication repetition rate for SMS-MO and SMS-MT, for every nth event. This functionality is built on existing SMS CLI, with configurable MO and/or MT. The default is not to authenticate.

SMSC Address Denial

Previously, the SGSN supported restricting MO-SMS and MT-SMS only through SGSN operator policy configuration.

Now, the SGSN can restrict forwarding of SMS messages to specific SMSC addresses, in order to allow operators to block SMS traffic that cannot be charged for. This functionality supports multiple SMSCs and is configurable per SMSC address with a maximum of 10 addresses. It is also configurable for MO-SMS and/or MT-SMS messages.

Threshold Crossing Alerts (TCA) Support

Thresholding on the system is used to monitor the system for conditions that could potentially cause errors or outage. Typically, these conditions are temporary (i.e high CPU utilization, or packet collisions on a network) and are quickly resolved. However, continuous or large numbers of these error conditions within a specific time interval may be indicative of larger, more severe issues. The purpose of thresholding is to help identify potentially severe conditions so that immediate action can be taken to minimize and/or avoid system downtime.

The system supports Threshold Crossing Alerts for certain key resources such as CPU, memory, number of sessions etc. With this capability, the operator can configure threshold on these resources whereby, should the resource depletion cross the configured threshold, a SNMP Trap would be sent.

The following thresholding models are supported by the system:

- **Alert**: A value is monitored and an alert condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.
- **Alarm**: Both high and low threshold are defined for a value. An alarm condition occurs when the value reaches or exceeds the configured high threshold within the specified polling interval. The alert is generated then generated and/or sent at the end of the polling interval.

Thresholding reports conditions using one of the following mechanisms:

- **SNMP traps**: SNMP traps have been created that indicate the condition (high threshold crossing and/or clear) of each of the monitored values.

Generation of specific traps can be enabled or disabled on the chassis. Ensuring that only important faults get displayed. SNMP traps are supported in both Alert and Alarm modes.

- **Logs**: The system provides a facility called threshold for which active and event logs can be generated. As with other system facilities, logs are generated Log messages pertaining to the condition of a monitored value are generated with a severity level of WARNING.

Logs are supported in both the Alert and the Alarm models.

- **Alarm System**: High threshold alarms generated within the specified polling interval are considered "outstanding" until a the condition no longer exists or a condition clear alarm is generated. "Outstanding" alarms are reported to the system's alarm subsystem and are viewable through the Alarm Management menu in the Web Element Manager.

The Alarm System is used only in conjunction with the Alarm model.

Important: For more information on threshold crossing alert configuration, refer to the *Thresholding Configuration Guide*.

Tracking Usage of GEA Encryption Algorithms

GPRS encryption algorithm (GEA) significantly affects the SGSN processing capacity based on the GEAx level used - GEA1, GEA2, or GEA3.

Operators would like to be able to identify the percentages of their customer base that are using the various GEA encryption algorithms. The same tool can also track the migration trend from GEA2 to GEA3 and allow an operator to forecast the need for additional SGSN capacity.

New fields and counters have been added to the output generated by the **`show subscribers gprs-only|sgsn-only summary`** command. This new information enables the operator to track the number of subscribers capable of GEA0-GEO3 and to easily see the number of subscribers with negotiated GEAx levels.

VLR Pooling via the Gs Interface

VLR Pooling, also known as Gs Pooling, helps to reduce call delays and call dropping, when the MS/UE is in motion, by routing a service request to a core network (CN) node with available resources.

VLR pools are configured in the Gs Service, which supports the Gs interface configuration for communication with VLRs and MSCs.

A *pool area* is a geographical area within which an MS/UE can roam without the need to change the serving CN node. A pool area is served by one or more CN nodes in parallel. All the cells, controlled by an RNC or a BSC belong to the same one (or more) pool area(s).

VLR hash is used when a pool of VLRs is serving a particular LAC (or list of LACs). The selection of VLR from this pool is based on the IMSI digits. From the IMSI, the SGSN derives a hash value (V) using the algorithm: [(IMSI div 10) modulo 1000]. Every hash value (V) from the range 0 to 999 corresponds to a single MSC/VLR node. Typically many values of (V) may point to the same MSC/VLR node.

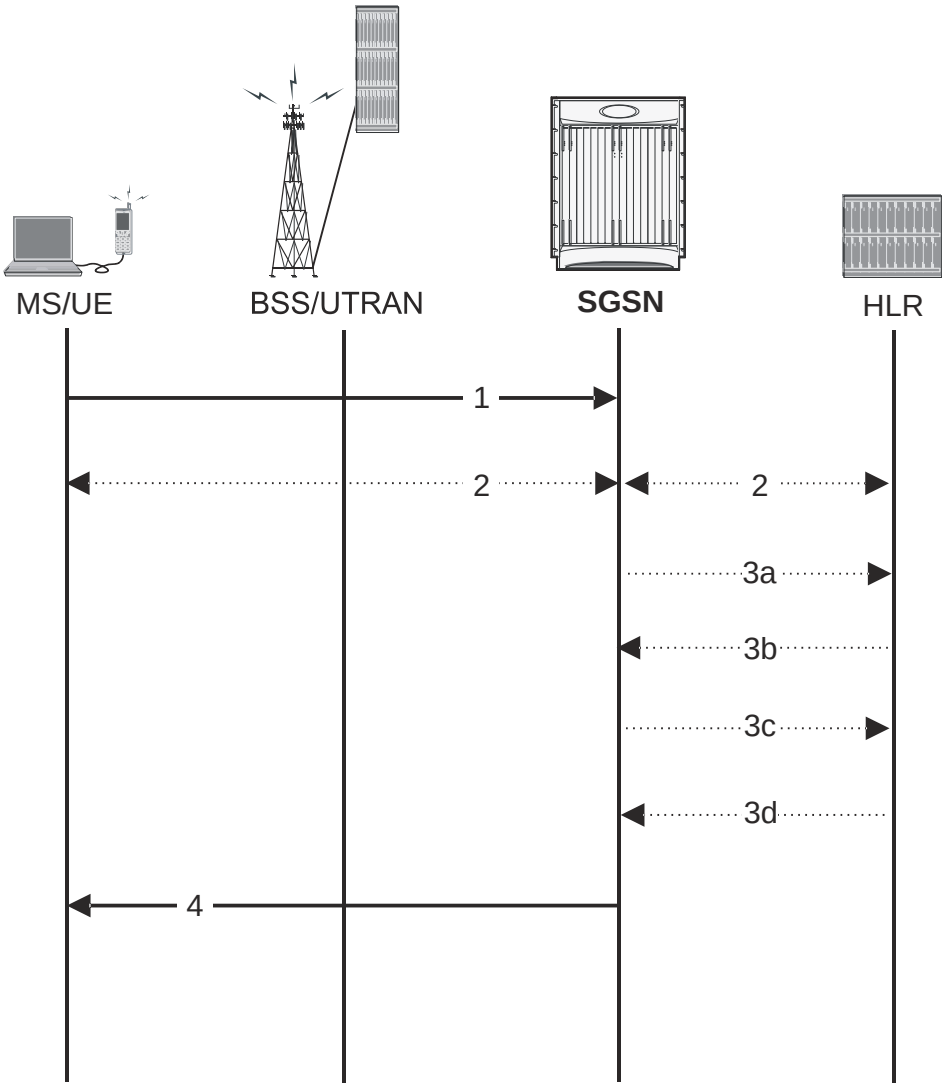
How the SGSN Works

This section illustrates some of the GPRS mobility management (GMM) and session management (SM) procedures the SGSN implements as part of the call handling process. All SGSN call flows are compliant with those defined by 3GPP TS 23.060.

First-Time GPRS Attach

The following outlines the setup procedure for a UE that is making an initial attach.

Figure 254. Simple First-Time GPRS Attach



This simple attach procedure can connect an MS via a BSS through the Gb interface (2.5G setup) or it can connect a UE via a UTRAN through the Iu interface in a 3G network with the following process:

Table 112. First-Time GPRS Attach Procedure

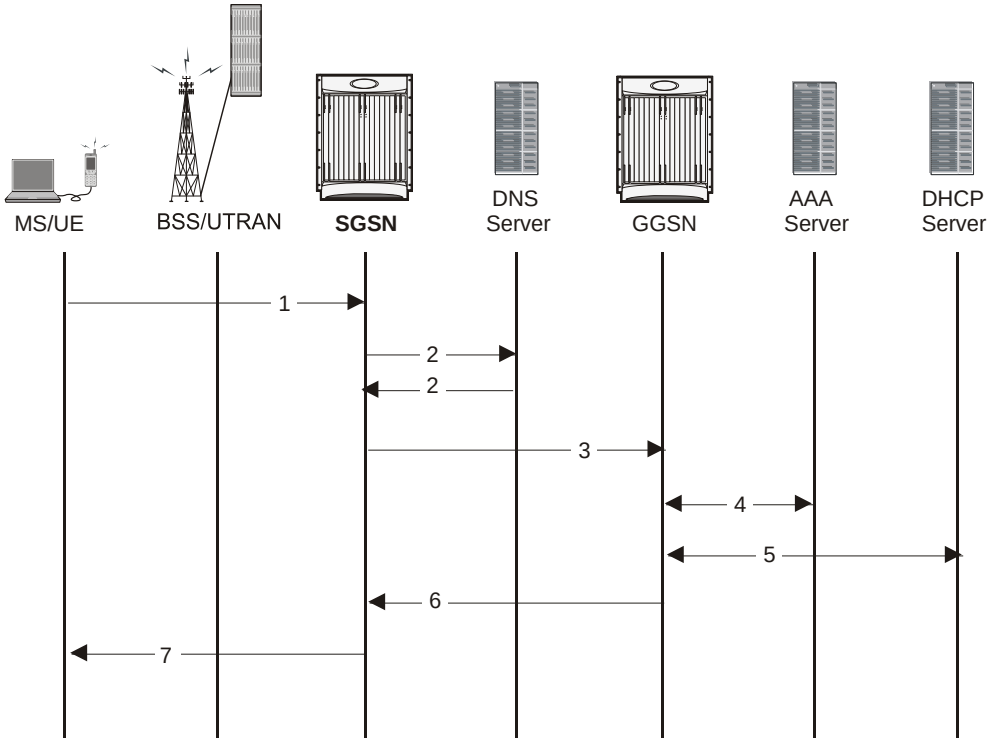
Step	Description
1	The MS/UE sends an Attach Request message to the SGSN. Included in the message is information, such as: • Routing area and location area information • Mobile network identity • Attach type
2	Authentication is mandatory if no MM context exists for the MS/UE: • The SGSN gets a random value (RAND) from the HLR to use as a challenge to the MS/UE. • The SGSN sends a Authentication Request message to the UE containing the random RAND. • The MS/UE contains a SIM that contains a secret key (Ki) shared between it and the HLR called a Individual Subscriber Key. The UE uses an algorithm to process the RAND and Ki to get the session key (Kc) and the signed response (SRES). • The MS/UE sends a Authentication Response to the SGSN containing the SRES.
3	The SGSN updates location information for the MS/UE: a) The SGSN sends an Update Location message, to the HLR, containing the SGSN number, SGSN address, and IMSI. b) The HLR sends an Insert Subscriber Data message to the "new" SGSN. It contains subscriber information such as IMSI and GPRS subscription data. c) The "New" SGSN validates the MS/UE in new routing area: If invalid: The SGSN rejects the Attach Request with the appropriate cause code. If valid: The SGSN creates a new MM context for the MS/UE and sends a Insert Subscriber Data Ack back to the HLR. d) The HLR sends a Update Location Ack to the SGSN after it successfully clears the old MM context and creates new one
4	The SGSN sends an Attach Accept message to the MS/UE containing the P-TMSI (included if it is new), VLR TMSI, P-TMSI Signature, and Radio Priority SMS. At this point the GPRS Attach is complete and the SGSN begins generating M-CDRs.

If the MS/UE initiates a second call, the procedure is more complex and involves information exchanges and validations between "old" and "new" SGSNs and "old" and "new" MSC/VLRs. The details of this combined GPRS/IMSI attach procedure can be found in 3GPP TS23.060.

PDP Context Activation Procedures

The following figure provides a high-level view of the PDP Context Activation procedure performed by the SGSN to establish PDP contexts for the MS with a BSS-Gb interface connection or a UE with a UTRAN-Iu interface connection.

Figure 255. Call Flow for PDP Context Activation



The following table provides detailed explanations for each step indicated in the figure above.

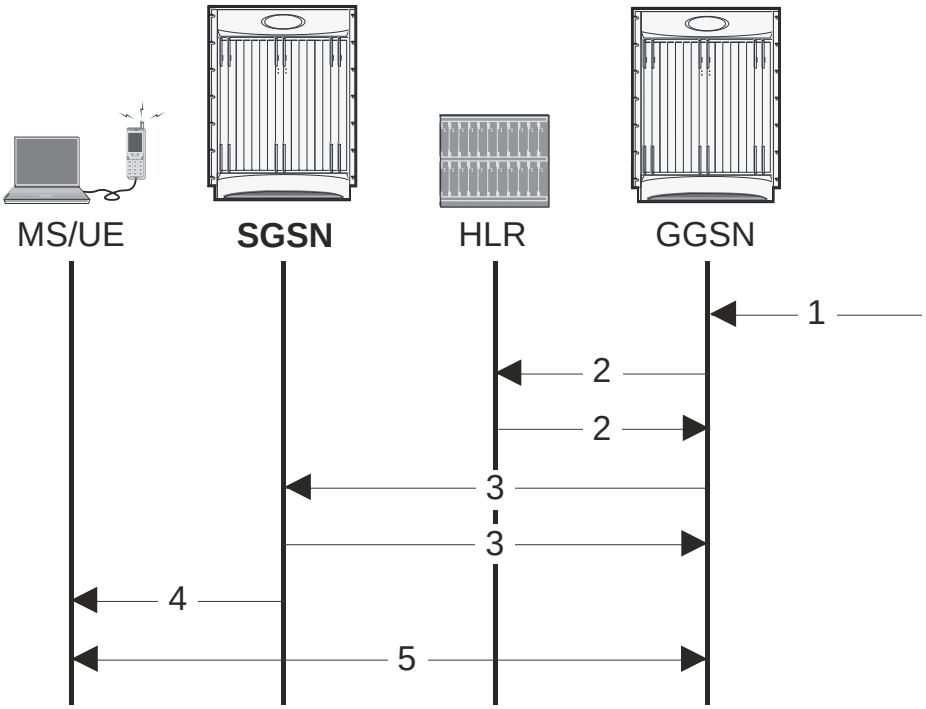
Table 113. PDP Context Activation Procedure

Step	Description
1	The MS/UE sends a PDP Activation Request message to the SGSN containing an Access Point Name (APN).
2	The SGSN sends a DNS query to resolve the APN provided by the MS/UE to a GGSN address. The DNS server provides a response containing the IP address of a GGSN.
3	The SGSN sends a Create PDP Context Request message to the GGSN containing the information needed to authenticate the subscriber and establish a PDP context.
4	If required, the GGSN performs authentication of the subscriber.
5	If the MS/UE requires an IP address, the GGSN may allocate one dynamically via DHCP.
6	The GGSN sends a Create PDP Context Response message back to the SGSN containing the IP Address assigned to the MS/UE.
7	The SGSN sends a Activate PDP Context Accept message to the MS/UE along with the IP Address. Upon PDP Context Activation, the SGSN begins generating S-CDRs. The S-CDRs are updated periodically based on Charging Characteristics and trigger conditions. A GTP-U tunnel is now established and the MS/UE can send and receive data.

Network-Initiated PDP Context Activation Process

In some cases, the GGSN receives information that requires it to request the MS/UE to activate a PDP context. The network, or the GGSN in this case, is not actually initiating the PDP context activation -- it is requesting the MS/UE to activate the PDP context in the following procedure:

Figure 256. Network-Initiated PDP Context Activation



The table below provides details describing the steps indicated in the graphic above.

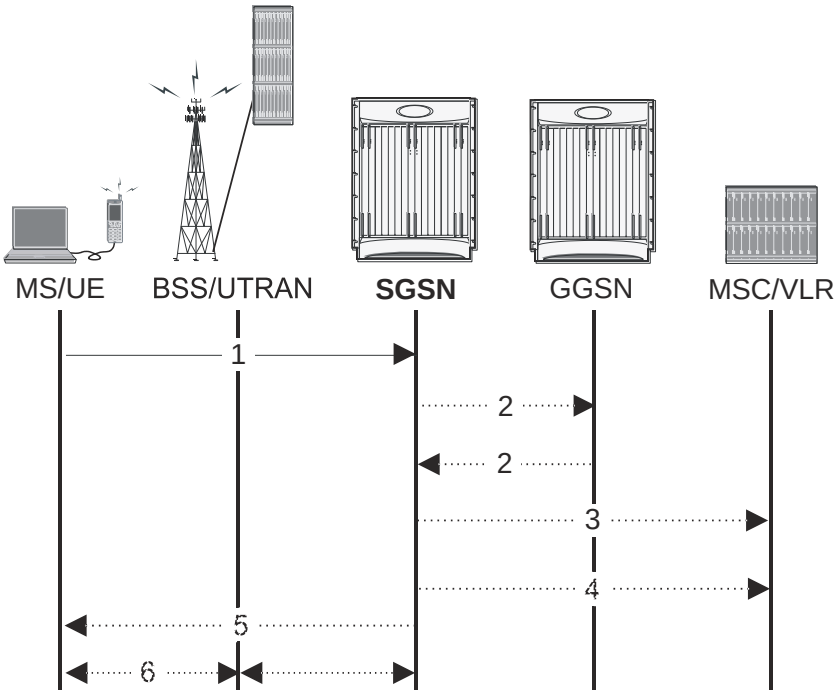
Table 114. Network Invites MS/UE to Activate PDP Context

Step	Description
1	The GGSN receives a PDU with a static PDP address that the GGSN 'knows' is for an MS/UE in its PLMN.
2	The GGSN uses the IMSI in place of the PDP address and sends an SRI (send routing information for GPRS) to the HLR. The HLR sends an SRI response back to the GGSN. The response may include the access of the target SGSN and it may also indicate it the MS/UE is not reachable, in which case it will include the reason in the response message.
3	The GGSN sends a PDU Notification Request to the SGSN (if the address was received). If the address was not received or if the MS/UE continues to be unreachable, the GGSN sets a flag marking that the MS/UE was unreachable. The notified SGSN sends a PDU Notification Response to the GGSN.
4	The SGSN determines the MS/UE's location and sets up a NAS connection with the MS/UE. The SGSN then sends a Request PDP Context Activation message to the MS/UE.
5	If the MS/UE accepts the invitation to setup a PDP context, the MS/UE then begins the PDP context activation process indicated in the preceding procedure.

MS-Initiated Detach Procedure

This process is initiated by the MS/UE for a range of reasons and results in the MS/UE becoming inactive as far as the network is concerned.

Figure 257. MS-Initiated Combined GPRS/IMSI Detach



The following table provides details for the activity involved in each step noted in the diagram above.

Table 115.MS-Initiated Combined GPRS/IMSI Detach Procedure

Step	Description
1	The UE sends a Detach Request message to the SGSN containing the Detach Type, P-TMSI, P-TMSI Signature, and Switch off indicator (i.e. if UE is detaching because of a power off).
2	The SGSN sends Delete PDP Context Request message to the GGSN containing the TEID. The GGSN sends a Delete PDP Context Response back to the SGSN. The SGSN stops generating S-CDR info at the end of the PDP context.
3	The SGSN sends a IMSI Detach Indication message to the MSC/VLR.
4	The SGSN sends a GPRS Detach Indication message to the MSC/VLR. The SGSN stops generating M-CDR upon GPRS Detach.
5	If the detach is not due to a UE switch off, the SGSN sends a Detach Accept message to the UE.

Step	Description
6	Since the UE GPRS Detached, the SGSN releases the Packet Switched Signaling Connection.

Supported Standards

The SGSN services comply with the followng standards for GPRS/UMTS wireless data services.

IETF Requests for Comments (RFCs)

- **RFC-1034,** Domain Names - Concepts and Facilities, November 1987; 3GPP TS 24.008 v7.8.0 (2007-06)
- **RFC-1035,** Domain Names - Implementation and Specification, November 1987; 3GPP TS 23.003 v7.4.0 (2007-06)
- **RFC-2960,** Stream Control Transmission Protocol (SCTP), October 2000; 3GPP TS 29.202 v6.0.0 (2004-12)
- **RFC-3332,** MTP3 User Adaptation Layer (M3UA), September 2002; 3GPP TS 29.202 v6.0.0 (2004-12)
- **RFC-4187,** Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA), January 2006
- **RFC-4666,** Signaling System 7 (SS7) Message Transfer Part 3 (MTP3) - User Adaptation Layer (M3UA), September 2006; 3GPP TS 29.202 v6.0.0 (2004-12)

3GPP Standards

3GPP Release 6 and higher is supported for all specifications unless otherwise noted. Support for higher releases is indicated below, in relation to current and planned development, including support for IEs and messages determined by supported functionality. Product development is aiming ultimately towards full compliance with the releases listed below:

Table 116.3GPP Standards Supported

3GPP Standard	R12.0/ R12.2	R14.0
3GPP TS 9.60, 3rd Generation Partnership Project; Technical Specification Group Core Network; General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp Interface (R98).	v7.10.0 (2002-12)	v7.10.0 (2002-12)
3GPP TS 22.041, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Operator Determined Barring (ODB)	v7.0.0 (2007-2003)	v9.0.0 (2009-12)
3GPP TS 22.042, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Network Identity and Timezone (NITZ); Service description, Stage 1	v7.0.0 (2007-2006)	v9.0.0 (2009-12)
3GPP TS 23.003, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Numbering, addressing and identification	v7.11.0 (2012-2003)	v9.9.0 (2011-12)
3GPP TS 23.007, 3rd Generation Partnership Project; Technical Specification Group Core Network; Restoration procedures	v7.0.0 (2006-06)	v9.10.0 (2011-12)
3GPP TS 23.015, 3rd Generation Partnership Project; Technical Specification Group Core Network; Technical realization of Operator Determined Barring (ODB)	v7.0.0 (2007-03)	v9.0.0 (2009-12)

3GPP Standard	R12.0/ R12.2	R14.0
3GPP TS 23.016, 3rd Generation Partnership Project; Technical Specification Group Core Network; Subscriber data management; Stage 2	v7.0.0 (2007-06)	v9.1.0 (2010-03)
3GPP TS 23.040, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Technical realization of the Short Message Service (SMS)	v7.2.0 (2009-03)	v9.3.0 (2010-09)
3GPP TS 23.060, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; General Packet Radio Service (GPRS); Service description; Stage 2	v7.11.0 (2011-2012)	V9.11.0 (2011-12)
3GPP TS 23.078, 3rd Generation Partnership Project; Technical Specification Group Core Network; Customized Applications for Mobile network Enhanced Logic (CAMEL) Phase 3 - Stage 2 (Release 4)	v4.11.1 (2004-04)	v4.11.1 (2004-04)
3GPP TS 23.107, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Quality of Service (QoS) concept and architecture	v7.1.0 (2007-06)	v9.3.0 (2011-12)
3GPP TS 23.236, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Intra-domain connection of Radio Access Network (RAN) nodes to multiple Core Network (CN) nodes	v7.0.0 (2006-12)	v9.0.0 (2009-12)
3GPP TS 23.251, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Network Sharing; Architecture and functional description	v7.0.0 (2007-06)	v9.4.0 (2011-03)
3GPP TS 23.271, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Functional stage 2 description of Location Services (LCS) (Release 9)	N/A	v9.6.0 (2011-03)
3GPP TS 23.401, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access (Release 9)	N/A	v9.12.0 (2012-03)
3GPP TS 24.007, 3rd Generation Partnership Project; Technical Specification Group Core Network; Mobile radio interface signalling layer 3; General aspects	v7.0.0 (2005-2009)	v10.0.0 (2011-03)
3GPP TS 24.008, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Mobile radio interface Layer 3 specification; Core network protocols; Stage 3	v7.15.0 (2010-2003)	v9.10.0 (2012-03)
3GPP TS 24.011, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Point-to-Point (PP) Short Message Service (SMS) support on mobile radio interface (Release 7)	v7.1.0 (2009-2006)	v7.1.0 (2009-2006)
3GPP TS 24.030, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Location Services (LCS); Supplementary service operations; Stage 3 (Release 9)	N/A	V9.0.0 (2009-12)
3GPP TS 24.080, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Mobile radio interface layer 3 supplementary services specification; Formats and coding (Release 9)	N/A	V9.2.0 (2010-06)
3GPP TS 25.410, 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu Interface: general aspects and principles	v7.0.0 (2006-03)	v9.0.1 (2011-03)
3GPP TS 25.411, 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface layer 1	v7.1.0 (2007-2009)	v9.0.1 (2011-03)

3GPP Standard	R12.0/ R12.2	R14.0
3GPP TS 25.412, 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface signaling transport	v7.1.0 (2006-06)	v9.0.1 (2011-03)
3GPP TS 25.413, 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface RANAP signalling (Release 9)	v9.8.0 (2011 - 12) (sections 8.19.2 and 8.20.2)	v9.8.0 (2011 - 12) (sections 8.19.2 and 8.20.2)
3GPP TS 25.414, 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface data transport and transport signaling	v7.1.0 (2006-06)	v9.0.1 (2011-03)
3GPP TS 25.415, 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; UTRAN Iu interface user plane protocols	v7.3.0 (2006-2012)	v9.0.1 (2011-03)
3GPP TS 29.002, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Mobile Application Part (MAP) specification	v7.15.0 (2010-06)	v9.8.0 (2012-03)
3GPP TS 29.016, 3rd Generation Partnership Project; Technical Specification Group Core Network; General Packet Radio Service (GPRS); Serving GPRS Support Node (SGSN) - Visitors Location Register (VLR); Gs interface network service specification	v7.0.0 (2007-2006)	v8.0.0 (2008-12)
3GPP TS 29.018, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; General Packet Radio Service (GPRS); Serving GPRS Support Node (SGSN) - Visitors Location Register (VLR) Gs interface layer 3 specification	v7.3.0 (2006-2012)	v9.1.0 (2010-12)
3GPP TS 29.060,3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface	v7.17.0 (2011-2009)	v9.9.0 (2011-12)
3GPP TS 29.078, 3rd Generation Partnership Project; Technical Specification Group Core Network; Customized Applications for Mobile network Enhanced Logic (CAMEL) Phase 3; CAMEL Application Part (CAP) specification (Release 4)	v4.9.0 (2009-2009)	v4.9.0 (2009-2009)
3GPP TS 29.202, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; SS7 Signalling Transport in Core Network; Stage 3	v7.0.0 (2007-2006)	v8.0.0 (2007-06)
3GPP TS 29.272, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Evolved Packet System (EPS); Mobility Management Entity (MME) and Serving GPRS Support Node (SGSN) related interfaces based on Diameter protocol (Release 9)	N/A	v9.7.0 (2011-2006)
3GPP TS 29.274, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C); Stage 3 (Release 9)	N/A	v9.9.0 (2011-12)
3GPP TS 29.303, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Domain Name System Procedures; Stage 3 (Release 9)	N/A	v9.4.0 (2011-2006)
3GPP TS 32.215, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Charging management; Charging data description for the Packet Switched (PS) domain	v5.9.0 (2007-10)	v5.9.0 (2007-10)
3GPP TS 32.251, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Charging management; Packet Switched (PS) domain charging	v7.8.0 (2010-10)	v9.8.0

3GPP Standard	R12.0/ R12.2	R14.0
3GPP TS 32.298, 3rd Generation Partnership Project; Technical Specification Group Service and System Aspects; Telecommunication management; Charging management; Charging Data Record (CDR) parameter description	v7.5.0 (2008-2012)	v9.10.0 (2011-12)
3GPP TS 32.406, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Performance Management (PM); Performance measurements Core Network (CN) Packet Switched (PS) domain	v7.1.0 (2006-06)	v9.0.0 (2009-12)
3GPP TS 32.410, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Key Performance Indicators (KPI) for UMTS and GSM	v8.0.0 (2009-2003)	v9.0.0 (2009-09)
3GPP TS 33.102, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G Security; Security architecture	v7.1.0 (2006-2012)	v9.4.0 (2010-12)
3GPP TS 33.106, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G security; Lawful Interception requirements	v7.0.1 (2006-2001)	v9.0.0 (2009-12)
3GPP TS 33.107, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G security; Lawful interception architecture and functions	v7.7.0 (2007-2009)	v9.4.0 (2011-03)
3GPP TS 33.108, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G security; Handover interface for Lawful Interception (LI) (Release 7)	v7.10.0 (2010-2012)	v7.10.0 (2010-2012)
3GPP TS 44.064, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Mobile Station - Serving GPRS Support Node (MS-SGSN); Logical Link Control (LLC) layer specification	v7.4.0 (2011-2012)	v9.1.0 (2011-12)
3GPP TS 44.065, 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Mobile Station (MS) - Serving GPRS Support Node (SGSN); Subnetwork Dependent Convergence Protocol (SNDCP)	v7.0.0 (2006-09)	v9.0.0 (2009-12)
3GPP TS 48.014, 3rd Generation Partnership Project; Technical Specification Group GSM EDGE Radio Access Network; General Packet Radio Service (GPRS); Base Station System (BSS) - Serving GPRS Support Node (SGSN) interface; Gb interface Layer 1	v7.0.0 (2007-2009)	v9.0.0 (2009-12)
3GPP TS 48.016, 3rd Generation Partnership Project; Technical Specification Group GSM EDGE Radio Access Network; General Packet Radio Service (GPRS); Base Station System (BSS) - Serving GPRS Support Node (SGSN) interface; Network Service	v7.4.0 (2008-2003)	v9.0.0 (2009-12)
3GPP TS 48.018, 3rd Generation Partnership Project; Technical Specification Group GSM/EDGE Radio Access Network; General Packet Radio Service (GPRS); Base Station System (BSS) - Serving GPRS Support Node (SGSN); BSS GPRS Protocol (BSSGP) (Release 7)	v7.13.0 (2009-2012)	v7.13.0 (2009-2012)

ITU Standards

- **Q711**; 3GPP TS 29.002 v7.15.0 (2006-2010), 3GPP TS 29.016 v7.0.0 (2007-08), and 3GPP TS 25.410 v7.0.0 (2006-03)

- **Q712**; 3GPP TS 29.002 v7.15.0 (2006-2010), 3GPP TS 29.016 v7.0.0 (2007-08), and 3GPP TS 25.410 v7.0.0 (2006-03)
- **Q713**; 3GPP TS 29.002 v7.15.0 (2006-2010), 3GPP TS 29.016 v7.0.0 (2007-08), and 3GPP TS 25.410 v7.0.0 (2006-03)
- **Q714**; 3GPP TS 29.002 v7.15.0 (2006-2010), 3GPP TS 29.016 v7.0.0 (2007-08), and 3GPP TS 25.410 v7.0.0 (2006-03)
- **Q715**; 3GPP TS 29.002 v7.15.0 (2006-2010), 3GPP TS 29.016 v7.0.0 (2007-08), and 3GPP TS 25.410 v7.0.0 (2006-03)
- **Q716**; 3GPP TS 29.002 v7.15.0 (2006-2010), 3GPP TS 29.016 v7.0.0 (2007-08), and 3GPP TS 25.410 v7.0.0 (2006-03)
- **Q771**; 3GPP TS 29.002 v7.15.0 (2006-2010)
- **Q772**; 3GPP TS 29.002 v7.15.0 (2006-2010)
- **Q773**; 3GPP TS 29.002 v7.15.0 (2006-2010)
- **Q774**; 3GPP TS 29.002 v7.15.0 (2006-2010)
- **Q775**; 3GPP TS 29.002 v7.15.0 (2006-2010)

Object Management Group (OMG) Standards

- CORBA 2.6 Specification 01-09-35, Object Management Group

Chapter 32
Traffic Performance Optimization Overview

This chapter provides an overview of the Traffic Performance Optimization (TPO) in-line service.

The following topics are covered in this chapter:

- Overview
- Feature Specifications
- TPO Administration
- How TPO Works

Overview

TPO is a licensed in-line service supported on the Cisco 5x00 chassis running any of the following products:

- GGSN
- HA
- P-GW

Though TCP is a widely accepted protocol in use today, it is optimized only for wired networks. Due to inherent reliability of wired networks, TCP implicitly assumes that any packet loss is due to network congestion and consequently invokes congestion control measures. However, wireless links are known to experience sporadic and usually temporary losses due to several reasons, including the following, which also trigger TCP congestion control measures resulting in poor TCP performance.

Reasons for delay variability over wireless links include:

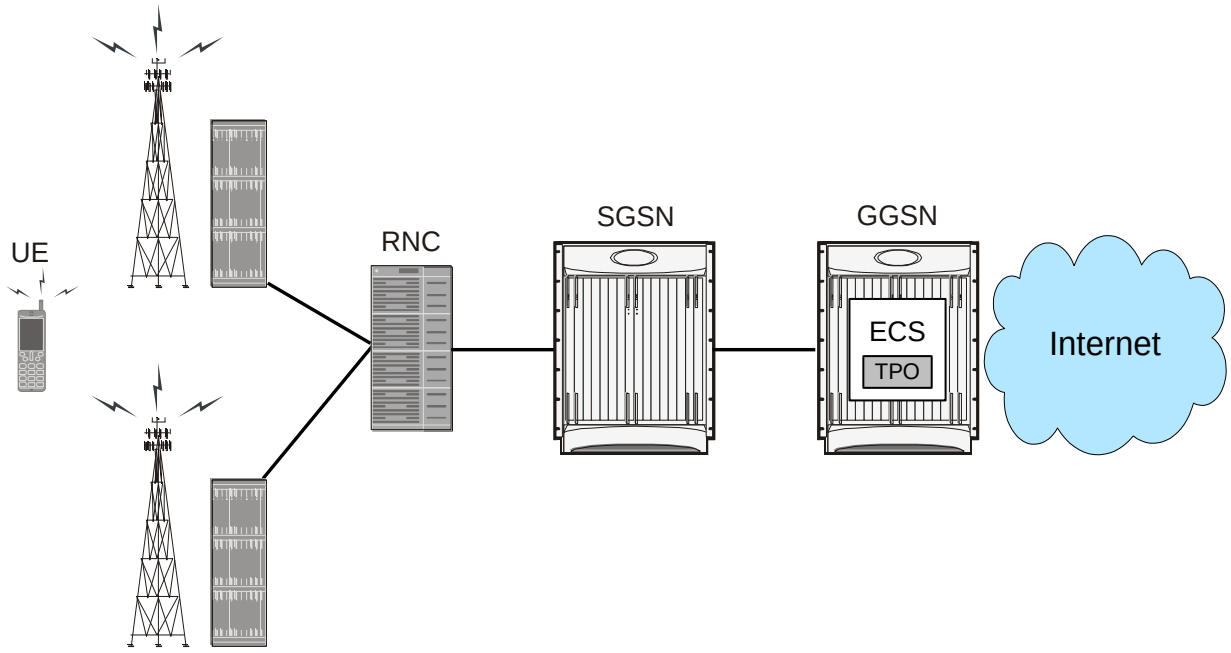
- Channel fading effect, subscriber mobility, and other transient conditions
- Link-layer retransmissions
- Handoffs between neighboring cells
- Intermediate nodes, such as SGSN and e-NodeB, implementing scheduling polices tuned to deliver better QoS for select services — resulting in variable delay in packet delivery for other services

The TPO in-line service uses a combination of TCP and HTTP optimization techniques to improve TCP performance over wireless links.

TPO Deployment

TPO uses the Enhanced Charging Service (ECS) framework for TCP Proxy functionality, and as depicted in the following figure, it is part of the ECS module in the Gateway.

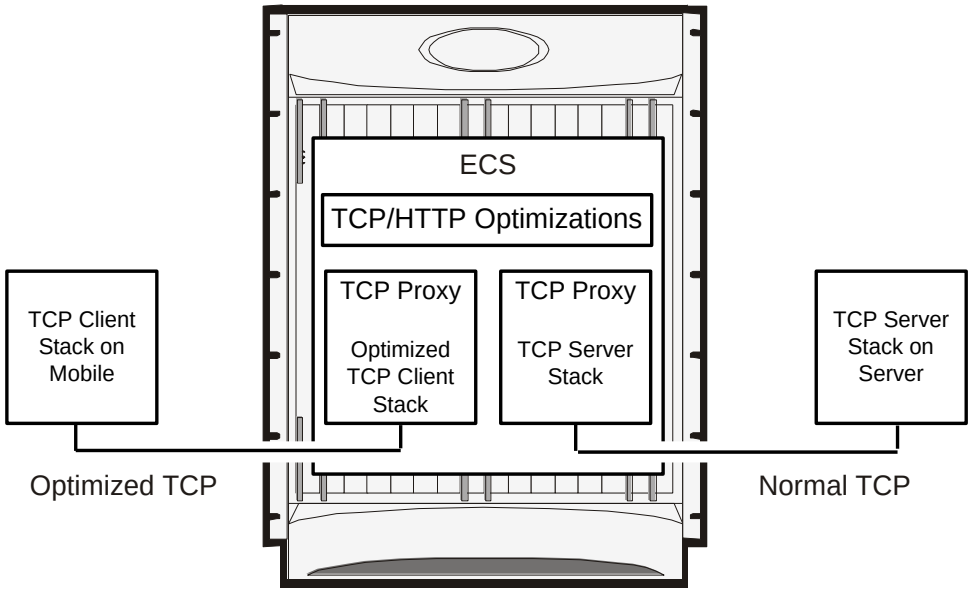
Figure 258. TPO Deployment



TPO uses the TCP Proxy to split end-to-end TCP connections between the client and server into two separate TCP connections, and apply the TCP and HTTP optimization techniques in the TCP stack towards the client. The split TCP connections isolate impacts of packet errors and delay variability for the wireless link from the wired connection, so that TCP congestion control, timeout, and retransmission mechanisms in the wired link do not suffer from the fluctuating quality of the radio channel.

Important: In this release TPO optimizes only downlink radio usage.

Figure 259. TPO Optimization Model



Feature Specifications

This section describes features of the TPO in-line service.

Platform Requirements

The TPO in-line service runs on Cisco ASR 5x00 chassis running StarOS. The chassis can be configured with a variety of components to meet specific network deployment requirements. For additional information, refer to the *Installation Guide* for the chassis and/or contact your Cisco account representative.

License Requirements

The TPO is a licensed Cisco feature. Separate session and feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements.

For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Supported Standards

TPO is based on the following RFCs and standards:

- RFC 793 Transmission Control Protocol; 1981-09
- RFC 2616 Hypertext Transfer Protocol — HTTP/1.1; 1999-06
- RFC 3481 TCP over Second (2.5G) and Third (3G) Generation Wireless Networks; 2003-02

TCP Optimization

To improve TCP performance over wireless links TPO uses TCP optimization techniques that enable:

- Improved downlink throughput during periods of congestion
- Closer to theoretical bandwidth utilization by reducing TCP overheads
- Adoption to wireless network events and conditions optimizing data transmission rate for available bandwidth and RTT
- Minimized TCP retransmissions

TCP Optimization Techniques

This section describes the TCP optimization techniques supported by TPO.

Selection of the optimization techniques to gain TCP performance depends on the network characteristics and the prevalent wireless conditions. To trigger appropriate congestion control techniques for a given situation based on the wireless events, TPO utilizes available information such as subscriber QoS, system load, and so on.

Optimizations in TCP Three-way Handshake

During TCP Three-way Handshake, the TCP client and server negotiate to establish a connection. TCP requires three round-trip time (RTT) measurements between the client and server before data transfer is initiated. Determining the RTT adds extra time for each wireless connection.

TPO supports the following optimizations during TCP Three-way Handshake:

- TCP options such as SACK and timestamp are negotiated.
- RTT is calculated based on Timestamp.

Optimizations in TCP Slow Start Phase

During TCP Slow Start phase, to discover available bandwidth for a connection, TCP calculates the lowest possible bandwidth and increases exponentially until a packet loss is detected. In wireless environments, this phase implies periods during which the link is under-utilized and perceived by the subscribers as slow.

TPO supports the following fast start techniques:

- TPO uses subscriber QoS settings to set the initial congestion window.

  The subscriber QoS information is received from the AAA/OCS.
- TPO uses information from other TCP connections to set the initial congestion window, which is based on the RTT and used bandwidth for the other connection(s).
- TPO allows to configure the initial congestion window to any of the following values:
  - Units of 1 through 255 MSS segments.
  - A dynamically computed value at runtime, which is calculated as a percentage of bandwidth-delay product (BDP). The bandwidth is a CLI-configured value, and the delay is calculated using the SYN-ACK exchange.
  - A value recommended by RFC 5681, which varies from 2 through 4 based on MSS. This the default setting.

Optimizations in TCP Congestion Avoidance Phase

In the TCP Congestion Avoidance phase, TCP after detecting available bandwidth for a new connection linearly adjusts the congestion window to discover incremental bandwidth.

TPO allows to configure any of the following congestion control algorithms:

- TCP Westwood Plus: The TCP Westwood Plus algorithm can significantly increase throughput over wireless links. It relies on end-to-end bandwidth estimation to discriminate the cause of packet loss — congestion or wireless channel effect. The bandwidth estimate is determined by measuring and averaging the rate of returning acknowledgements. This estimate is then used to compute the congestion window and slow start threshold after a congestion episode.
- Vegas: The Vegas algorithm instead of relying on detection of packet loss, uses a bandwidth estimation scheme to proactively gauge network congestion. The Sender watches for signs of congestion is setting in, such as RTT growing and sending rate flattening.
- Basic: This is a custom TCP congestion algorithm based on the TCP Reno algorithm. This is the default setting.

Fast Retransmits

To guarantee reliability of transfers, TCP requires the Receiver to respond to the Sender with an acknowledgment for each segment it receives. The Sender keeps a record of each segment it sends, and waits for an acknowledgment before

sending the next segment. If the Sender does not receive an acknowledgment within the timeout period, under the assumption that the segment was lost in the network, the Sender fast-retransmits that segment (that is, retransmit without waiting for retransmission timeout (RTO)).

The Receiver will generate duplicate acknowledgements for every out-of-order (OOO) packet it receives. If the Sender receives three duplicate acknowledgements with the same acknowledgement number (that is, a total of four acknowledgements with the same acknowledgement number), the Sender decides that the segment with the next higher sequence number was dropped, and will not arrive out of order. The Sender will then fast-retransmit that segment.

TPO supports the following optimization with fast retransmits:

- TPO allows to configure the number of duplicate acknowledgements that will trigger fast retransmits. This can be:
  - Static value: When high amount of re-ordering is present in the network the static threshold of three duplicate acknowledgements does not work well. Under such conditions a higher static value is required as the threshold. This is a CLI-configurable parameter and can be a value 1 through 10.

    However, note that a higher static value will sometimes lead to not reaching the threshold because of less number of in-flight packets (which will roughly determine the number of duplicate ACKs received by the sender).
  - Dynamic value: The duplicate acknowledgement threshold is dynamically computed at runtime based on the number of in-flight packets (one-third of the in-flight packets, subject to a minimum of two). This will enable to adapt in networks where high amount of packet re-ordering is observed.

TCP Handoff Optimization

TPO supports intra-tech and inter-tech handoff events.

When an intra-tech handoff is detected, TPO takes the following actions:

1. Restarts RTT/RTO calculation based on first packet sent after the handoff.
2. Ignores duplicate acknowledgements as congestion event. Considers duplicate acknowledgements as packet drops due to handoff for next one RTT.
3. If RTO is triggered after the handoff event, for the next RTT, ignores congestion window adjustment.
4. If RTO happens and then handoff event is received, undoes congestion window due to RTO.
5. Restarts BWE for Vegas and Westwood congestion control algorithms.
6. Handles one handoff event per X RTT only.

When an inter-tech handoff is detected, TPO takes the following actions:

- Handoff Pre-event: As soon as bearer creation request is received at the gateway:
  1. Stops forwarding packets when handoff pre-event is received.
  2. Stops RTO and retransmission.
- Handoff Post-event: When bearer creation is complete at PGW/GGSN, restarts congestion control algorithm. This will take care of resending old unacknowledged packets.

Enabling/disabling TCP Handoff Optimization is a CLI-configurable parameter.

Retransmission Timeout Optimizations

TPO allows to configure the scaling factor for Round Trip Time Variation (RTTVAR). The configured scaling factor is used as a power of 2, so values of 0 through 4 correspond to 1, 2, 4, 8, and 16. In TCP RTO is calculated using the following formula:

*RTO = SRTT + K*RTTVAR*

where:

- *SRTT* = mean Round Trip Time (RTT)
- *RTTVAR* = Round Trip Time Variation

As wireless networks exhibit high RTT variation, the value of K is configurable. The value of K decides the extent to which Retransmission Timeout (RTO) timer depends on RTT variance. If RTT variance is higher, then K should be higher. The default RTTVAR scaling value, as recommended by RFC 2988, is 2.

TPO also allows to configure the RTO retransmission back-off factor. Once RTO fires for a packet, TCP will retransmit that packet and set the RTO to be a factor X, which is CLI-configurable, of the previous RTO. The default RTO backoff value, as recommended by RFC 5681 is 2.0.

HTTP Optimizations

To optimize incoming HTTP payload over TCP connections TPO uses HTTP optimization techniques that enable:

- Reducing payload by compressing text-based Web pages (packet headers are not compressed), albeit such pages do not account for much traffic and Web servers themselves could do the compression
- Minimizing number of round trips from clients, enabling faster response time
- Reducing payload by filtering advertisements

HTTP Optimization Techniques

This section describes HTTP optimization techniques supported by TPO.

HTTP Compression

The HTTP Compression feature enables to compress HTTP content transferred to the mobile client. HTTP Compression, by reducing the amount of traffic to be transported, enables better use of available bandwidth and improves transmission speeds.

Important: In this release TPO supports only the standards-based gzip compression algorithm.

Note that TPO will not attempt compression in the following cases:

- If the mobile client cannot accept compressed encodings
- If the HTTP version used is earlier than 1.1
- If the response from the Web server is already compressed

Enabling/disabling the HTTP Compression feature is a CLI-configurable parameter. By default HTTP Compression is disabled.

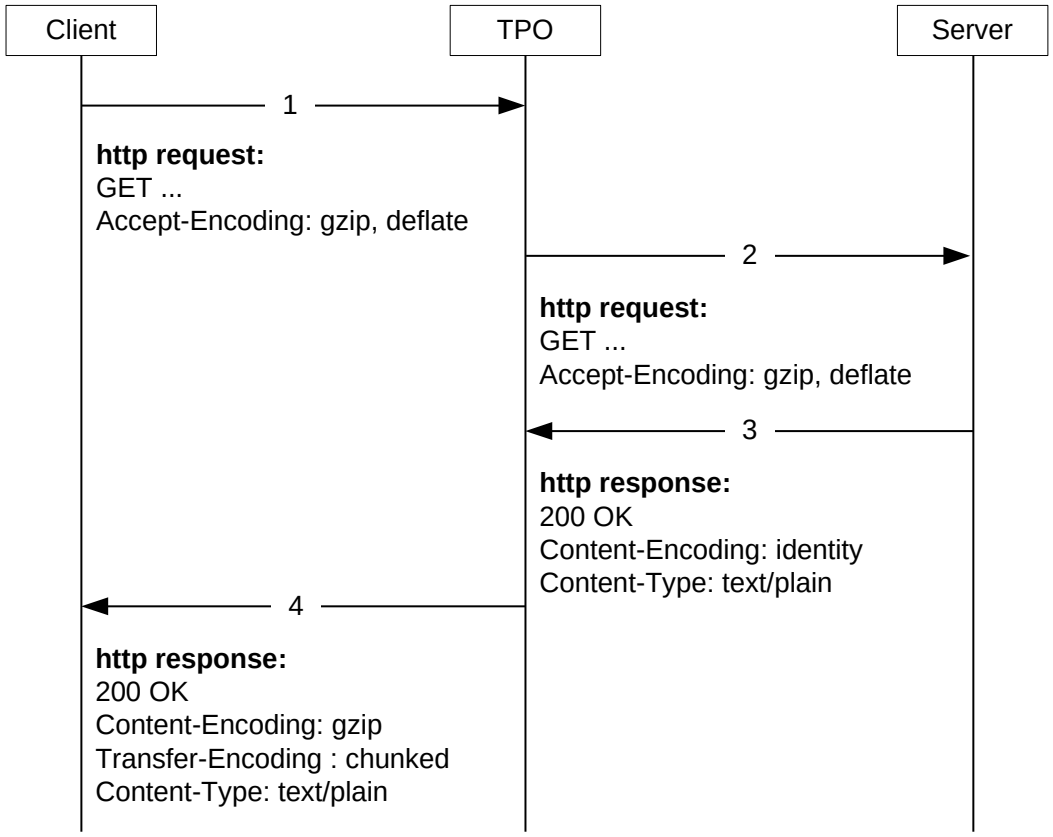
Compression level is a CLI-configurable parameter. Compression levels 1 through 9 are supported. The higher the compression level, the better the compression efficiency but with increased CPU and memory utilization. By default the compression level is set to 6.

TPO supports Prevention of Compression at the Web server. This enables TPO to receive uncompressed data from the Web server, on which it can apply other HTTP optimization techniques, and then compress the optimized data and send

it to the mobile client. TPO achieves this by manipulating the HTTP requests. Enabling/disabling this feature is a CLI-configurable parameter. By default, the Prevention of Server Compression feature is disabled.

The following figure and steps explain how the HTTP Compression feature works:

Figure 260. HTTP Compression



1. Mobile client's browser in its HTTP request for HTML content incudes an Accept-Encoding field with comma separated list of supported compression schema names.
2. TPO forwards the HTTP request to the Web server.
3. If the Web server does not support the compression schema(s) in the Accept-Encoding field or cannot undertake compression, in the HTTP response it sends the uncompressed data.

   If the Web server supports the compression schema(s) in the Accept-Encoding field, in the HTTP response the Web server may send the compressed data, and include the Content-Encoding field with name(s) of the compression schema(s) used. TPO forwards the HTTP response to the mobile client.
4. If the Web server in its response sends uncompressed data, TPO compresses the data and then uses internal thresholds for amounts of internal resources available compared to the amount of packet size reduction achieved.

   If the comparison is favorable, TPO forwards the compressed response to the mobile client, and any subsequent response packets are also compressed. TPO updates the Content-Encoding field with name(s) of the compression schema(s) used. The mobile client's browser parses the requested data and displays it.

   If the comparison is not favorable, TPO forwards the original response to the mobile client and then forwards any subsequent response packets.

URL Rewrite

The URL Rewrite feature enables to preemptively resolve host names in embedded URLs present within HTML content and rewrite them with resolved IP addresses. This rewriting helps to eliminate DNS round trips in high latency mobile networks resulting in faster responses.

Enabling/disabling the URL Rewrite feature is a CLI-configurable parameter. By default the URL Rewrite feature is disabled.

Important: The URL Rewrite feature needs a valid DNS client to be configured in the ISP (destination) context.

When the URL Rewrite feature is enabled, TPO rewrites URLs of the following format

http://*<host_name[:port]>*/*<url_path>*/*<file_name.extension>*

into

http://*<resolved_ip_address[:port]>*/*<url_rewrite_prefix>*/*<host_name[:port]>*/*<url_path>*/*<file_name.extension>*

For example, if the URL Rewrite prefix is *urlrewrite*, TPO rewrites the URL

http://www.google.com/test.img

into

http://209.85.153.103/urlrewrite/www.google.com/test.img

When the mobile client requests for the URL

http://*<resolved_ip_address[:port]>*/*<url_rewrite_prefix>*/*<host_name[:port]>*/*<url_path>*/*<file_name.extension>*

TPO rewrites the URL back to

http://*<host_name[:port]>*/*<url_path>*/*<file_name.extension>*

The URL Rewrite prefix is a CLI-configurable parameter. By default, the prefix is set to “urlrewrite”.

URL Rewrite works only on HTML content, hence it is called only in the following cases:

- Content in response to HTTP GET request where MIME type of the content is HTML/CSS/JavaScript
- Content is not encoded

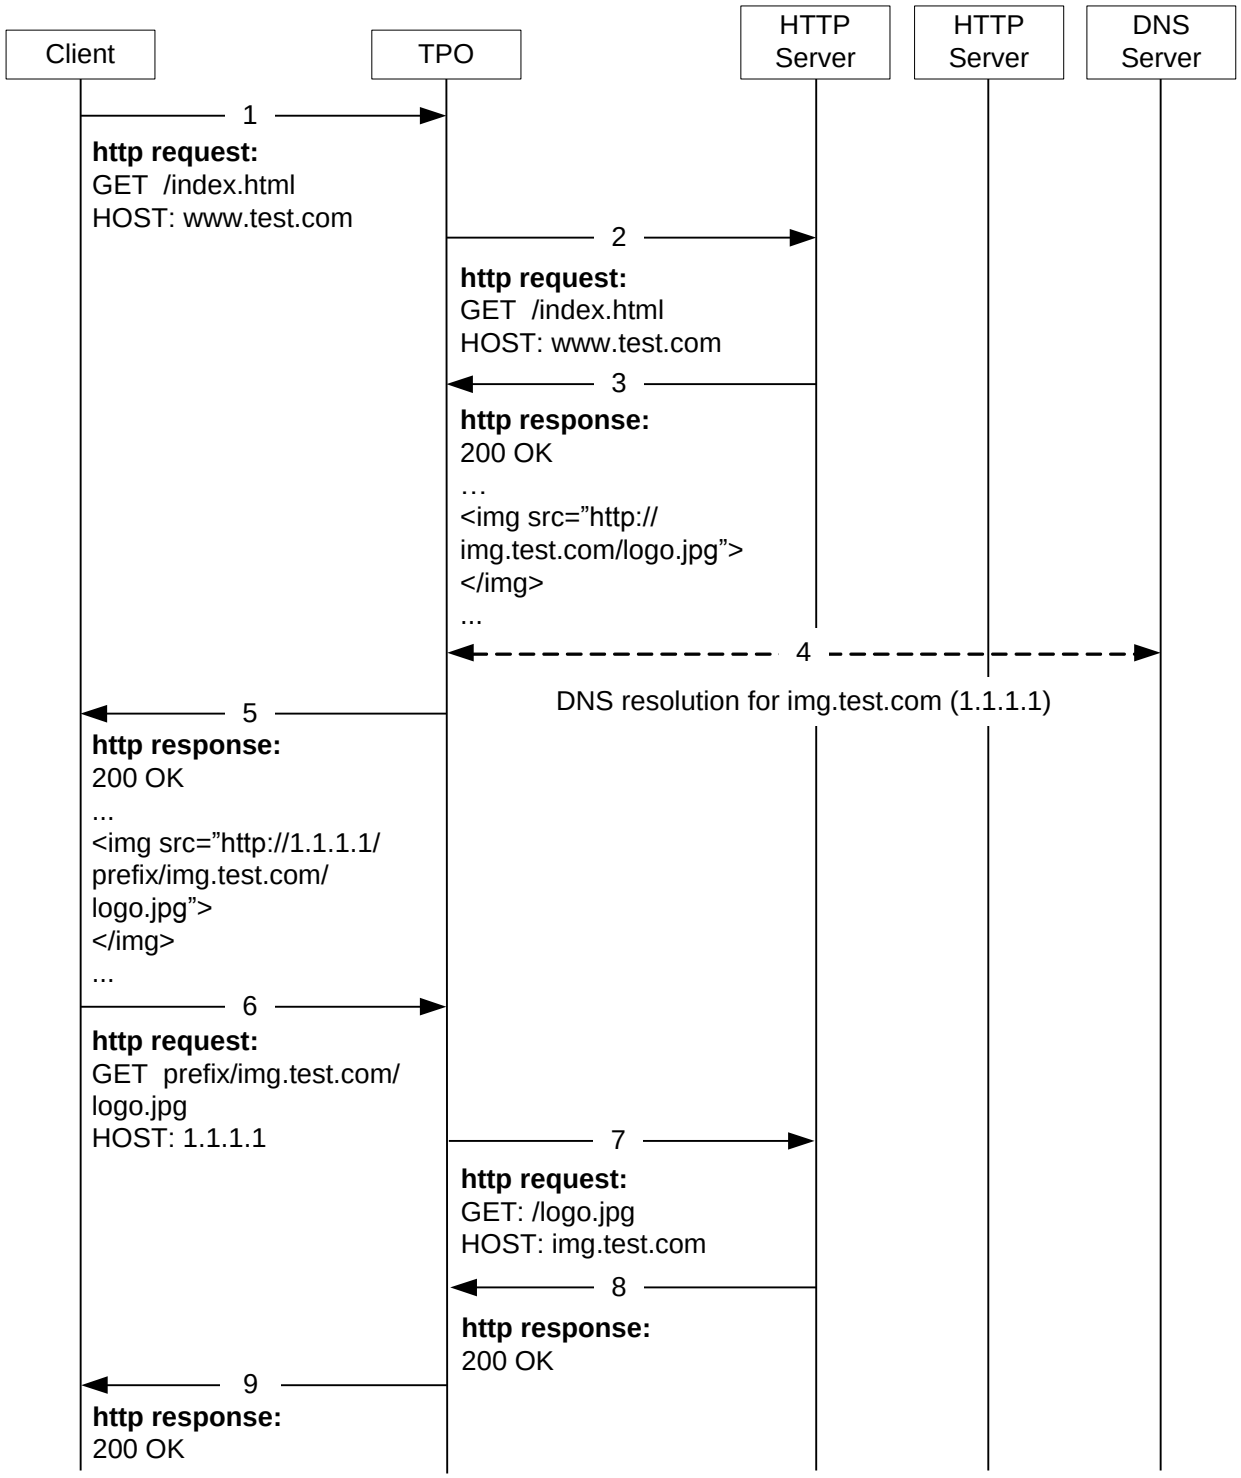
TPO rewrites only those URLs that are present in the following HTML tags:

- image
- imagepath
- img
- input
- link
- script

Important: URLs that are part of JavaScript and VBScript are not rewritten. If an HTML tag spans across packets, TPO will queue only two packets and will rewrite the URL if found.

The following figure and steps explain how the HTTP URL Rewrite feature works:

Figure 261. HTTP URL Rewrite



1. The mobile client sends HTTP request for HTML content.
2. TPO forwards the HTTP request to the Web server.

3. The Web server sends an HTTP response with the requested HTML content.
4. TPO resolves host names in embedded URLs present within the HTML content and rewrites them with corresponding IP addresses.
5. TPO forwards the HTTP response to the mobile client.
6. The mobile client's browser parses the HTML and sends HTTP requests for images and other content to the modified URLs.
7. TPO rewrites the URLs and forwards the HTTP requests to the Web server.
8. The Web server returns HTTP response with the requested content.
9. TPO forwards the HTTP response to the mobile client.

   In case both HTTP Compression and URL Rewrite features are enabled, URL Rewrite processing will happen before HTTP Compression.

Advertisement Filter

The Advertisement Filter feature enables to block advertisement content in Web pages delivered to mobile clients. This filtering reduces over-the-air bandwidth usage as advertisements are not always downloaded, and faster response times as advertisement-related server connections are eliminated.

Important: In this release, TPO considers only images and Flash objects as advertisements.

TPO is configured with URLs of the advertisement sites to be blocked, typically sites such as www.doubleclick.net/*, ad.yahoo.com, and so on. When the mobile client's browser receives HTML content, it parses the HTML and sends out requests for images and other content. If there is a request for an image or Flash object whose URL matches any of the URLs to be blocked, TPO blocks the advertisement as per the configuration.

The Advertisement Filter feature supports the following advertisement blocking methods:

- Advertisement Blocking with NO Text: In the mobile client's browser each blocked advertisement is replaced with the "cannot display image" icon (usually an X mark). Subscribers cannot view the advertisements even if they want to.
- Advertisement Blocking with Text: In the mobile client's browser each blocked advertisement is replaced with a placeholder frame. Each placeholder frame contains standard operator-configured text and the advertisement's URL. Subscribers cannot view the blocked advertisements even if they want to.
- Advertisement Blocking with On-click Function: In the mobile client's browser each blocked advertisement is replaced with a placeholder frame. Each placeholder frame contains standard operator-configured text and the advertisement's URL. To view a blocked advertisement the subscriber must click the placeholder frame.

  To enable retrieving the blocked advertisement, in the HTTP request a bypass string is added to the advertisement's URL, which TPO interprets and forwards to the Web Server allowing the advertisement to be retrieved. The bypass string is a CLI-configurable parameter.

The background color of the placeholder frames and the text displayed in them are CLI-configurable parameters. Operators can use the text to indicate that the advertisement is blocked. Note that different text strings can be configured for "Advertisement Blocking with Text" and "Advertisement Blocking with On-click Function" configurations.

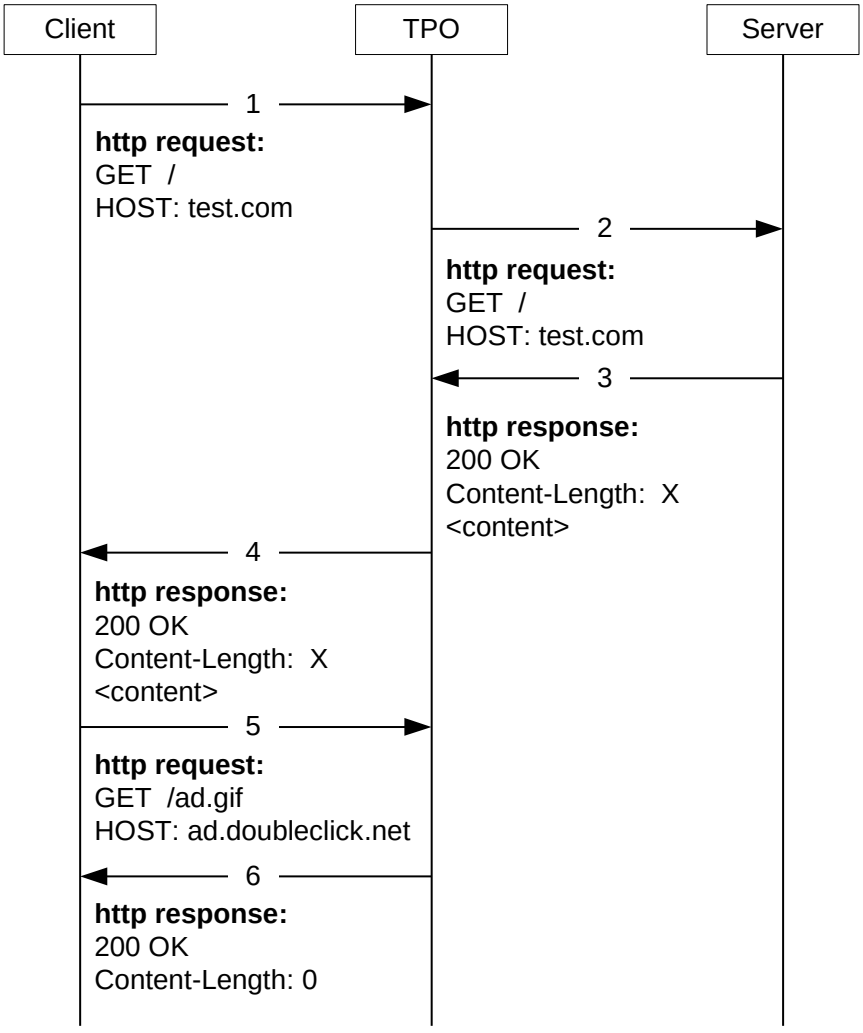
Important: The text string and URL displayed in the placeholder frames may be truncated to fit dimensions of the frames.

Important: The "Advertisement Blocking with Text" and "Advertisement Blocking with On-click Function" Advertisement Filter functionality is achieved using JavaScript code sent from TPO and executed whenever a Web page is loaded in the mobile client's browser. If the mobile client's browser does not support JavaScript or the subscriber has disabled JavaScript, instead of the placeholder frames the subscribers will see the "cannot display image" icons.

Basic Advertisement Blocking

The following figure and steps explain how the basic advertisement blocking feature works.

Figure 262. Basic Advertisement Blocking



1. The mobile client sends HTTP request for HTML content.
2. TPO forwards the HTTP request to the Web server.
3. The Web server sends an HTTP response with the requested HTML content.
4. TPO forwards the HTTP response to the mobile client.
5. The mobile client's browser parses the HTML and sends HTTP requests for images, Flash objects, and other content.

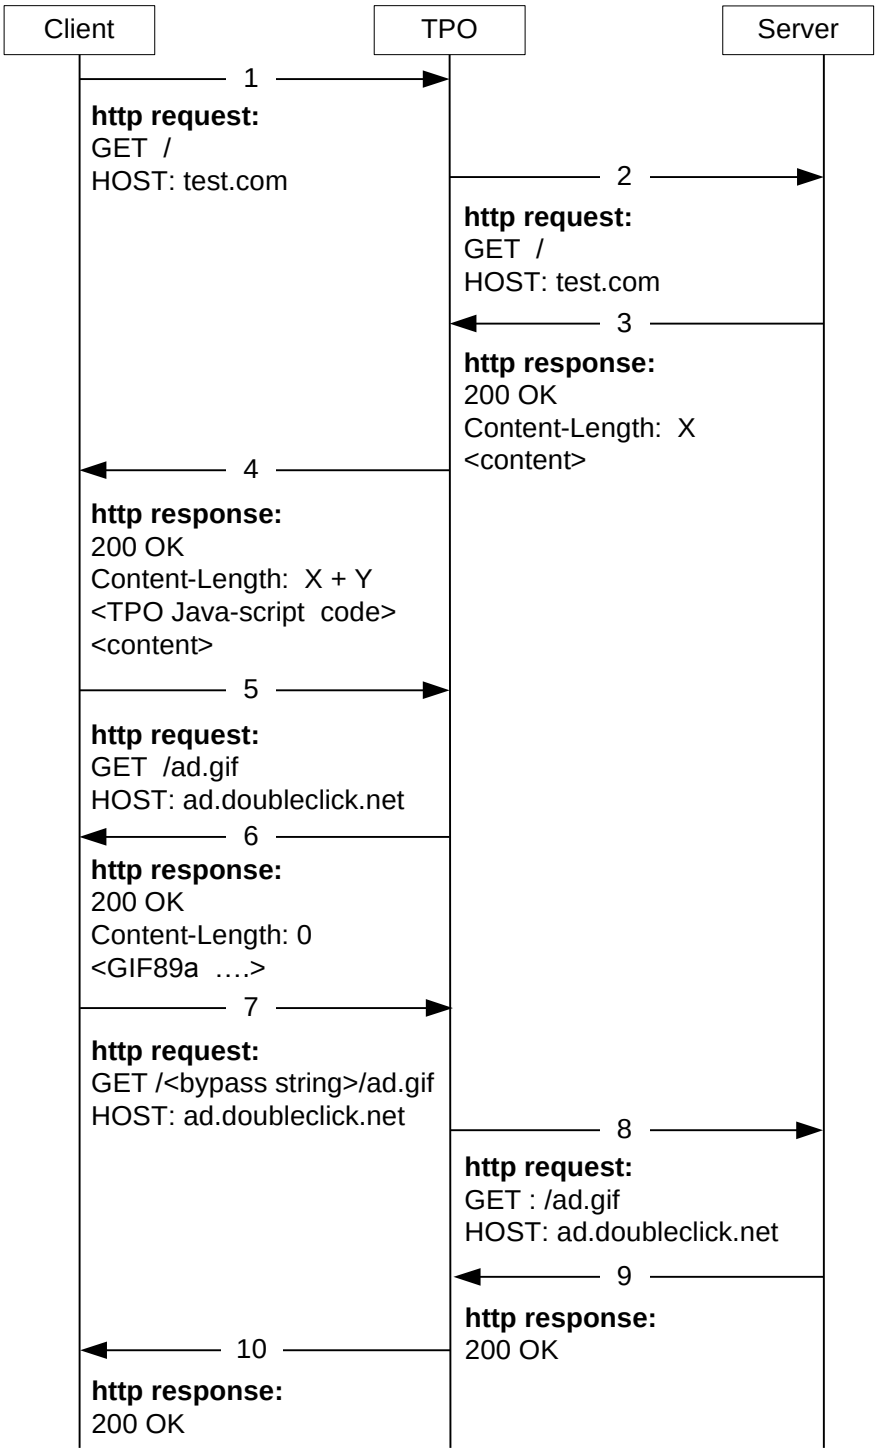
6. If there is an HTTP request for an image or Flash object, TPO matches the requested URL with the list of advertisement URLs to be blocked.

   If there is a match, TPO responds with Content-Length 0 for the request thereby blocking the advertisement. In the mobile client's browser, the image or Flash object is replaced with the "cannot display image" icon.

Advertisement Blocking with On-click Function

The following figure and steps explain how the Advertisement Blocking with On-click feature works.

Figure 263. Advertisement Blocking with On-click Function



1. The mobile client's browser sends an HTTP request for HTML content.
2. TPO forwards the HTTP request to the Web server.
3. The Web server sends an HTTP response with the requested HTML content.

4. TPO forwards the HTTP response to the mobile client's browser along with a JavaScript code containing the list of advertisement site URLs to be blocked.
5. The mobile client's browser parses the HTML and sends HTTP requests for images, Flash objects, and other content.
6. If there is an HTTP request for an image or a Flash object, TPO matches the requested URL with the URLs to be blocked. If there is a match, TPO responds with local content based on the requested file's extension.
7. JavaScript in the mobile client's browser parses the HTML to check for blocked images and Flash objects, and replaces them with placeholder frames. If "Advertisement Blocking with On-click Function" is enabled, on-click functionality is added to the frames.

   When the subscriber clicks the placeholder frame for a blocked image or Flash object, the mobile client's browser sends an HTTP request with a bypass string appended to the requested URL.
8. TPO forwards the HTTP request to the Web server.
9. The Web server sends an HTTP response with the requested data.
10. TPO forwards the HTTP response to the mobile client's browser. The mobile client's browser displays the requested advertisement content.

Session Recovery

The Session Recovery feature provides seamless failover and reconstruction of subscriber session information in the event of a hardware or software fault within the system preventing a fully connected user session from being disconnected.

Session recovery is performed by mirroring key software processes (for example, Session Manager and AAA Manager) within the system. These mirrored processes remain in an idle state (in standby-mode), wherein they perform no processing, until they may be needed in the case of a software failure (for example, a Session Manager task aborts). The system spawns new instances of "standby mode" session and AAA Managers for each active Control Processor (CP) being used.

Additionally, other key system-level software tasks, such as VPN Manager, are performed on a physically separate packet processing card to ensure that a double software fault (for example, Session Manager and VPN Manager fails at same time on same card) cannot occur. The packet processing card used to host the VPN Manager process is in active mode and is reserved by the operating system for this sole use when session recovery is enabled.

For more information on Session Recovery, refer to the *Session Recovery* chapter in the *System Administration Guide*.

TPO Administration

This section describes TPO administration activities and covers the following topics:

- DisablingEnabling TPO Optimizations
- MUR Reporting for TPO
- Switching TPO Policy

Disabling/Enabling TPO Optimizations

The TPO in-line service allows to disable/continue TPO optimizations when a peer-to-peer (P2P) flow is detected. This is a CLI-configurable feature.

Important: In this release, on disabling TPO optimizations only TCP-based optimizations are disabled. Disabling HTTP-based optimizations will be supported in a future release.

MUR Reporting for TPO

This section lists the EDR fields supported for TPO - MUR integration.

The Mobility Unified Reporting (MUR) is a Web-based application providing a unified reporting interface for a variety of data from Cisco Systems in-line service and storage applications. For more information on the MUR, see the *Mobility Unified Reporting System Online Help*.

The following are the EDR generation points:

- End of TCP connection
- End of HTTP transaction

The following EDR fields are supported for TPO - MUR integration:

- TCP flow related fields:
    - TCP data transferred
    - TCP duration
    - TPO enable/disabled
- HTTP transaction related fields:
    - HTTP URL
    - HTTP DNS local resolutions
    - HTTP DNS server resolutions
    - HTTP compression bytes in
    - HTTP compression bytes out
    - HTTP advertisement replaced (advertisement blocked)
    - HTTP advertisement delivered (advertisement accessed)

- TPO enabled/disabled

Switching TPO Policy

TPO allows to switch a TPO policy in use with a different TPO policy.

Important: The switch takes effect only on current calls. For new calls, the RADIUS-returned/APN/subscriber template configured policy is used.

To be able to change the TPO policy in mid session, TPO must have been enabled for the subscriber in the APN/Subscriber template configuration during call setup.

The CLI indicates the number of sessions for which the policy got switched.

How TPO Works

This section describes how TPO works.

Terms and Definitions

The following is a list of terms specific to TPO functionality:

- TPO Policy: A TPO policy specifies the match-rule definitions to select a TPO profile. The match-rule definitions enable to use different sets of optimizations for different TCP/HTTP flows of a subscriber.

  The TPO policy to be used for a subscriber can be from one of the following:

  - AAA/OCS: The TPO policy can come from the AAA server or the OCS. During initial authentication the AAA server returns the TPO policy for the subscriber, which is applied to the corresponding session. For this purpose the system uses the RADIUS AVP SN-TPO-Policy. If the policy comes from the AAA/OCS, it will override the policy configured in the subscriber's APN/subscriber template and/or the ECS rulebase.

    Important: The TPO policy received from the AAA and OCS have the same priority. Whichever comes first, either from AAA or the OCS, is applied.

  - APN/Subscriber Template: If no TPO policy is received from the AAA/OCS, the TPO policy configured in the subscriber's APN/subscriber template is applied to the flows over the sessions using that APN/subscriber profile.
  - ECS Rulebase: The default TPO policy configured in the ECS rulebase has the least priority. If no TPO policy to use is received from the AAA/OCS, and there is no TPO policy configured in the subscriber's APN/subscriber template, only then will the default TPO policy configured in the ECS rulebase be used.

  A maximum of 2048 TPO policies can be configured in the system.
- TPO Profile: A TPO profile specifies the optimizations to be performed for a specific flow. A maximum of 2048 TPO profiles can be configured in the system. A TPO profile can be used in more than one TPO policy.
- Ruledef: A ruledef specifies the criteria to identify a specific flow, such as HTTP flow to google.com. The criteria matches one or more protocol fields and/or protocol-state information supported by the protocol analyzers.
- Rulebase: A rulebase specifies the protocol analyzers that run, and which packets they should analyze. Within a rulebase, ruledefs are specified to select the appropriate analyzers. Also within a rulebase, a collection of ruledefs are associated with charging actions.
- Charging Action: A charging action specifies the action to be taken when a ruledef is matched. Many possible actions are possible, such as an action could indicate whether to count retransmitted packets for a subscriber, or it could be to disconnect the subscriber upon certain protocol specific triggers.
- Subscriber Profile/APN: The subscriber profile/APN specifies the TPO policy to be applied to the flows over the sessions using that subscriber profile/APN.

TPO Processing

TPO can be enabled in either of the following ways:

- Policy-based TPO processing: On all flows of a subscriber based on the TPO policy specified by the AAA/OCS, APN/subscriber template, or the ECS rulebase.
- Charging Action-based TPO processing: On specific flows of a subscriber based on the flow matching a charging action with a TPO profile configured in it.

Policy-based TPO Processing

The following steps describe how policy-based TPO processing works:

1. When a subscriber initiates a data session, the TPO policy from the AAA/OCS, or the TPO policy configured in the subscriber's APN/subscriber profile, or the ECS rulebase is associated with that data session.
2. When a flow is created over the session (as when the subscriber initiates a browsing session), first the ECS routing ruledefs are applied to determine the protocol analyzer (such as HTTP).
3. The optimization ruledefs configured in the TPO policy are applied one after the other, in the specified order. When a match is found, the optimizations configured in the corresponding TPO profile are applied to the session.

   During the course of a flow, first the match-rule logic is applied at SYN time — this mostly results in a default match-rule that selects the default TPO profile. This is because many of the match-rule conditions would not apply at SYN time. The match-rule is generally more useful with deep-packet inspection (DPI). During DPI, when the complete HTTP header information is received, the match-rule is invoked and a new TPO profile (if any) is obtained and applied. This new TPO profile (selected during DPI) will be used to perform HTTP optimization. However, the original TPO profile selected during SYN time will be used for TCP optimization.

   If the first SYN does not match any TPO profile (in the absence of a default match rule), TPO is not applied to that flow. In 12.0 and later releases, with dynamic TCP Proxy this is no longer the case.

Important: The match-rule is also invoked after the HTTP request line is received. At this time, a TPO profile is used to only apply the advertisement block rule (if any). This is required to block any unwanted HTTP request packets for and advertisement site that could be potentially sent to the server. None of the other rules are applied even if present in the profile.

Charging Action Based TPO Processing

The following steps describe how charging action based TPO processing works:

1. When a flow is created (as in when the subscriber initiates a browsing session), ECS routing ruledefs are applied to determine the protocol analyzer (such as HTTP).
2. Based on the ECS rule matching the charging action to apply to the flow is selected.
3. If a TPO profile is configured in that charging action, optimizations configured in that TPO profile are applied on the flow.

Important: In this release, when the TPO profile specified by a charging action is applied on a flow, only TCP optimizations and the decision to cease/continue TPO optimizations for P2P flows will be controlled by that TPO profile. HTTP optimizations will not be affected.

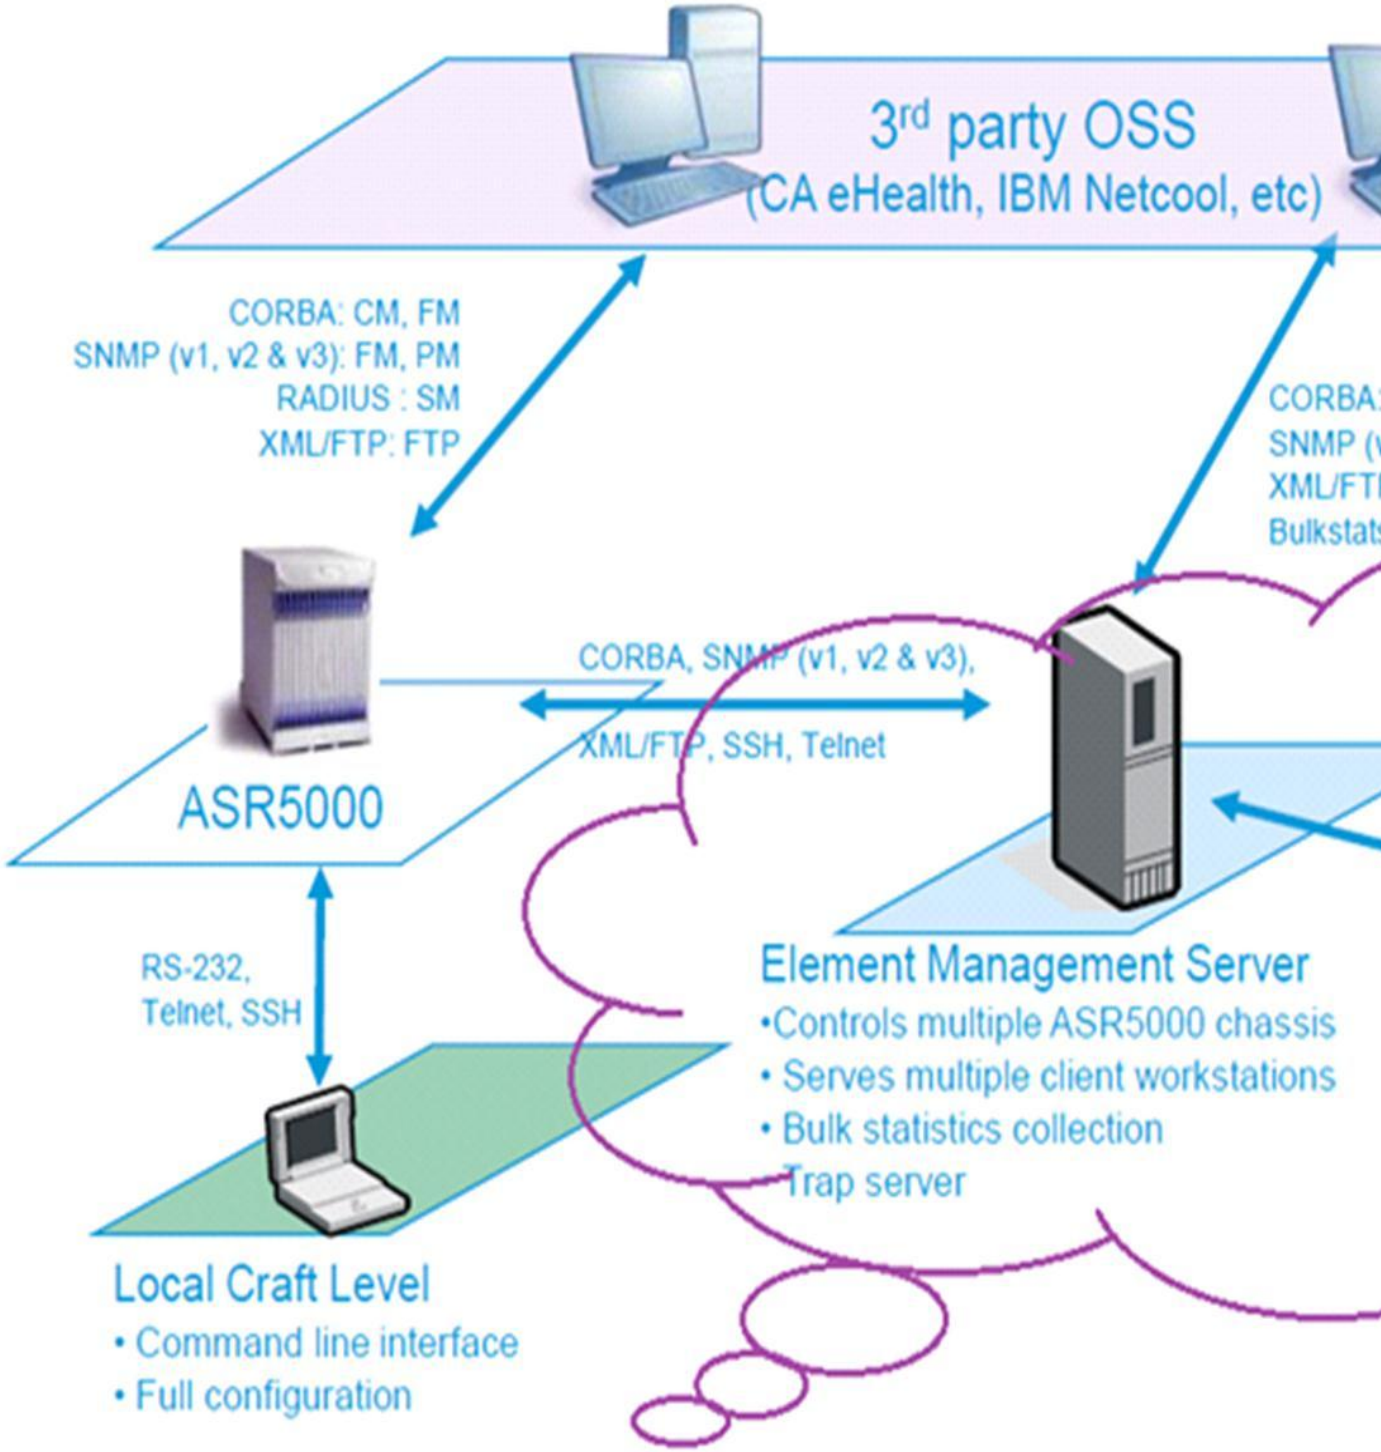
When the TPO profile specified by a charging action is applied on a flow, and subsequently a different charging action that does not have a TPO profile configured in it is applied on the same flow, TPO will not be disabled.

Chapter 33
Web Element Manager Overview

The Web Element Manager (WEM) is a Common Object Request Broker Architecture (CORBA)-based application that provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capability for the system under management.

For maximum flexibility and scalability, the WEM application implements a client-server architecture. This architecture allows remote clients with Java-enabled web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the fault-tolerant Sun® Solaris® and Cisco MITG Red Hat Enterprise Linux operating systems. For added security, management traffic can be encrypted using the Secure Sockets Layer (SSL) protocol, as shown in the following diagram:

Figure 264. Network Topology Example


3rd party OSS
(CA eHealth, IBM Netcool, etc)
CORBA: CM, FM
SNMP (v1, v2 & v3): FM, PM
RADIUS : SM
XML/FTP: FTP
CORBA, SNMP (v1, v2 & v3),
XML/FTP, SSH, Telnet
ASR5000
RS-232,
Telnet, SSH
Element Management Server
•Controls multiple ASR5000 chassis
• Serves multiple client workstations
• Bulk statistics collection
Trap server
Local Craft Level
• Command line interface
• Full configuration

Supported Features

FCAPS Support

The Web Element Manager application providses Fault, Configuration, Accounting, Performance and Security (FCAPS) management functionality for the chassis.

Fault Management

Fault management consists of an event logging function wherein all alarms, warnings, and other faults can be configured, reported, and acknowledged by network operations personnel.

The Simple Network Management Protocol (SNMP) is used by both the Web Element Manager and the chassis to report event notifications. The application's fault management system offers the following support for generated alarms:

- Provide mechanisms for viewing both current and pending alarms for both the chassis and the Web Element Manager server.
- Generate audio and visual alerts for alarms based on their severity (the Web Element Manager also supports the configuration of a severity level for each alarm).
- Maintain statistics for generated alarms.
- Store alarm information in the PostgreSQL® database.
- Execute scripts through the Script Server component of the application.
- Send E-mail notifications and/or forward notifications to Network Management Servers (NMSs) using a CORBA/IIOP-based Northbound Interface.
- Compliancy with the following standards:
    - TS 32.111-3, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Fault Management; Part 3: Alarm Integration Reference Point (IRP): Common Object Request Broker Architecture (CORBA) Solution Set (SS)
    - TS 32.303, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Configuration Management (CM); Notification Integration Reference Point (IRP): Common Object Request Broker Architecture (CORBA) Solution Set (SS)

Configuration Management

The Web Element Manager implements an easy to use, point-and-click GUI for providing configuration management for one or more systems. This GUI provides all the flexibility offered by the system's command Line Interface (CLI), while providing the scalability of performing certain functions across multiple chassis. All configuration information is stored in the PostgreSQL Database.

At the system-level, the Web Element Manager application provides support for the following:

- Adding, modifying, or deleting systems to/from the management system
- Performing configuration of card and port-level parameters
- Adding, modifying, or deleting contexts

- Configuring specific protocols and services within defined contexts such as AAA servers, PDSN services, GGSN services, IP access lists, IP interfaces, IP routes, IP address pools, RADIUS accounting and authentication, PPP, subscribers, and others

At the network level, the application is capable of transferring configuration and/or software images to multiple systems simultaneously in advance to performing software upgrades.

The Web Element Manager supports the configuration of all parameters required to perform software upgrades including:

- Adding, deleting, and sorting system boot stack entries; these entries allow multiple fall-backs in the event the system experiences an error in the loading of a particular image or configuration file
- Configuring network options for bootup
- Transferring configuration and image files to/from a chassis
- Initiating and monitoring upgrade status

The Web Element Manager further simplifies the software upgrade process by providing tools for managing system configuration files:

- **Back-up Tool:** Enables the Web Element Manager to transfer a copy of the configuration file currently being used by a managed system at user-defined intervals. Files are transferred to the host server in a specific directory. The number of files to retain in the directory is also configurable. This tool provides a useful mechanism for testing configurations and/or quickly restoring a last-known-good configuration in the event of an error.
- **Compare Tool:** Provides a powerful tool for comparing the configuration files of two managed systems. Once the two files are specified, a dialog appears displaying the two documents side-by-side. Line numbers are added for convenience. Text additions, modifications, and deletions are displayed in different colors for easy recognition. This tool can be useful on its own to determine variations between multiple iterations of the same configuration file, or, when used in conjunction with the Back-up tool, it can provide an audit trail of configuration changes that occurred during system operation.

Accounting Management

Accounting management operations allow users to examine and perform post-process statistical analysis on systems managed by the Web Element Manager application.

The type of statistics used for element management-based accounting are called bulk statistics. Bulk statistics are grouped into categories called schemas and are polled by the system at fixed polling intervals and then transferred to the Web Element Manager at a different transfer intervals (defined in minutes).

Once the Web Element Manager server application, called the receiver, has received bulk statistics files from the managed system, these files are parsed and added to the PostgreSQL database. This database is updated as new files are received.

The Web Element Manager's accounting management functionality is compliant with *TS 32.401, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Performance Management (PM); Concept and requirements* and allows you to:

- Collect statistics pertaining to the transfer and collection of bulk statistics
- Views statistics stored on the chassis prior to transfer to the receiver
- Graph multiple received bulk statistics over time as either a line or bar graph; these graphs can be printed to network printers accessible by the server
- Generate eXtensible Markup Language (XML) files for transfer to a Northbound NMS or bulk statistics processor.

- Archive collected bulk statistic information to conserve disk space on the server

Performance Management

Performance management operations supported by the Web Element Manager allow users to examine and perform real-time statistical analysis on systems managed by the application as well as on the server on which the application is running.

Information pertaining to various aspects of the Web Element Manager (CPU and memory utilization, disk space, and process status) and its managed systems (hardware, protocols, software subsystems, and subscribers) is collected in real time and is displayed in tabular format. Alternatively, most of the information can be graphed as a function of time in either line or bar-chart format. Multiple statistics can be graphed simultaneously for quick comparison of data.

In addition to collecting and providing mechanisms for the real-time viewing of statistical information, the Web Element Manager provides useful monitoring tools similar to those found in the CLI. These tools can be used to monitor active subscriber sessions, protocol flows, and port information. Data collected during this monitor operation can be saved to the client machine for further analysis.

Security Management

Security management pertains to the operations related to management users. This includes both Web Element Manager application users and local management users who are configured on the chassis. In many cases, management users can be allowed access to both the system (via its CLI) and the application. It is possible for both management user accounts to share the same username and password.

The security management features of the Web Element Manager allow you to:

- Add, modify, or delete administrative users for both the application and the managed system.

  Regardless of the administrative user type, there are four levels of management user privileges:

  - **Inspector**: Inspectors are limited to a small number of read-only Exec Mode commands. The bulk of these are "show" commands giving the inspector the ability to view a variety of statistics and conditions. The Inspector cannot execute **`show configuration`** commands and do not have the privilege to enter the Config Mode.
  - **Operator**: Operators have read-only privileges to a larger subset of the Exec Mode commands. They can execute all commands that are part of the inspector mode, plus some system monitoring, statistic, and fault management functions. Operators do not have the ability to enter the Config Mode.
  - **Administrator**: Administrators have read-write privileges and can execute any command throughout the CLI except for a few security-related commands that can only be configured by Security Administrators. Administrators can configure or modify the system and are able to execute all system commands, including those available to the Operators and Inspectors.
  - **Security Administrator**: Security Administrators have read-write privileges and can execute all CLI commands including those available to Administrators, Operators, and Inspectors.
- Provide authentication and privilege restoration based on the login information entered by administrative users.
- Monitor current system or application-level administrative users in real-time and optionally terminate their management session.
- Perform an audit of all managed system configurations performed through both the application and the CLI as well as other operations performed within the application.

  The audit trail functionality supports the configuration of filters defining the type of operations to audit and also provides a dialog for performing the audit in real-time.

Audit trail results are stored in the PostgreSQL database for later retrieval and analysis.

The new Network Audit Tool functionality in WEM supports the on-demand or periodic auditing of chassis configuration attributes such as PPP MRU, Auth Sequence, Bulkstats Schema Needs Update, etc.

Additional Features

Additional features provided by the Web Element Manager application include:

High Availability Redundant Server Clustering

Beginning with Release 12.0 redundant servers can be configured using Oracle Cluster software. Much of the information in the following chapters applies equally to a cluster installation although at this time we recommend using the GUI installation method described in *Installing the WEM Software*. A separate appendix addresses the differences between a Standalone installation and a Failover installation and it also defines the steps required to create and configure a cluster.

Management Integration Capabilities

Utilizing the Object Management Group's (OMG) standard CORBA Northbound interface, the Web Element Manager application can be integrated with higher-level TMN-modeled applications such as network, business, and service layer applications. The OMG's Interface Definition Language (IDL) can be used to develop custom interfaces to various other third-party components such as Application Servers, etc.

Database Management and Redundancy Support

All databases used for audit trail, performance and statistical information, event management, and device inventory information will be stored on the Web Element Manager server using the UNIX file system.

Multiple Language Support

The Web Element Manager provides the ability for users to select a specific language the information is provided in. The currently supported languages include U.S. English and Korean.

Context-Sensitive Help System

The Web Element Manager has a complete web-based Help system that provides user assistance for every screen and function available within the application. This Help system resides on the Web Element Manager server and is accessible from any supported client workstation.

Stand-alone Offline Help System

A stand-alone version of the WEM's online help file is provided with the WEM. This stand-alone help contains content identical to the context-sensitive help. Users can save the standalone (.chm) help file to a local drive to conduct off-line WEM-related research. The standalone help file can be downloaded from the main WEM browser page by clicking the *Web Element Manager Offline Help* link under **Help Resources**.

Multiple OS Support

Web Element Manager can be installed on servers running the Sun Solaris or the custom Cisco MITG Red Hat Linux Enterprise v5.5 operating systems. For detailed operating system and hardware platform requirements, refer to the *Web Element Manager System Requirements* section of this chapter.

Caution: The Cisco MITG RHEL v5.5 OS is a custom image that contains only those software packages required to support compatible Cisco MITG external software applications. Users must not install any other applications on servers running the Cisco MITG v5.5 OS. For detailed software compatibility information, refer to the *Cisco MITG RHEL v5.5 OS Application Note*.

Web Element Manager System Requirements

Important: The hardware required for the Web Element Manager server may vary, depending on the operating system being used, the number of chassis being managed, the number of clients that require access, and other variables. WEM will support a maximum of 100 chassis for bulkstats reporting and 100 simultaneous web clients.

Server Hardware Requirements

This section describes the WEM server hardware requirements for both the Sun Solaris and Cisco MITG Red Hat Enterprise Linux Operating Systems.

Sun Solaris Server Hardware Requirements

WEM servers running the Sun Solaris operating system must be deployed on the following hardware platform:

- Sun Microsystems Netra™ T5220 server
- 1 x 1.2GHz 8 core UltraSPARC T2 processor with 32GB RAM
- 2 x 146GB SAS hard disk drives
- Quad Gigabit Ethernet interfaces
- Internal DVD-ROM drive
- AC or DC power supplies depending on the application

Red Hat Enterprise Linux Server Hardware Requirements

WEM servers running the Cisco MITG RHEL v5.5 operating system must be deployed on the following hardware platform:

- Cisco UCS C210 M2 Rack Server
- 2x Intel Xeon X5675 processors with 2x 16 GB DDR3 RAM
- 300GB SAS 10K RPM SFF Hard Disk Drive
- Quad Gigabit Ethernet interfaces
- Internal DVD-ROM drive
- AC or DC power supplies depending on the application

Operating System Requirements

This section describes the Sun Solaris and Cisco MITG Red Hat Enterprise Linux (RHEL) v5.5 operating system requirements for WEM servers.

Sun Solaris Operating System Requirements

This section describes the required Sun Solaris OS requirements for WEM servers, including the required OS patches.

Important: Ensure that all recommended patches are installed before performing a new installation or software upgrade, as follows:

Solaris 8 with Recommended Patch Cluster dated on or after April 2006.

Users based in the United States should ensure that the timezone patch 109809-05 (or later) and libc patch 108993-52 (or later) be installed in support of extended daylight savings time (DST).

Important: Solaris 8 Patch Level 23. Following required for **cluster-based configuration only**

- PCI Dual FC 2Gb HBA with SFS
- Optical 5 meter null ethernet cable
- PCI-based video card or Keyboard-Video-Mouse (KVM) card (optional)

Solaris 9 with Recommended Patch Cluster dated on or after April 2006

Users based in the United States should ensure that the timezone patch 113225-07 (or later) and libc patch 112874-33 (or later) be installed in support of extended daylight savings time (DST) support.

In addition, if Solaris 9 is used, it must be installed using the "End User System support 64-bit" software group must be specified during the installation of the operating system. This option installs the libraries required for proper operation of the Web Element Manager.

Solaris 10 with Recommended Patch Cluster dated on or after April 2011.

Users based in the United States using Solaris 10 should ensure that the timezone patch 138856-02 or later is installed in support of extended Daylight Savings Time (DST).

Important: If you plan to install software and maintain the Web Element Manager application and server remotely, it is recommended that you use an X-Windows client.

Red Hat Enterprise Linux Operating System Requirements

This section describes the required Cisco MITG Red Hat Linux (RHEL) operating system (OS) requirements for WEM servers.

- Cisco MITG RHEL v5.5 OS

For hardware platform requirements for the Cisco MITG RHEL v5.5 OS, refer to the *Server Hardware Requirements* section in this chapter. For information related to installation, refer to the *Cisco MITG RHEL OS v5.5 Application Note*.

Caution: The Cisco MITG RHEL v5.5 OS is a custom image that contains only those software packages required to support compatible Cisco MITG external software applications. Users must not install any other applications on servers running the Cisco MITG v5.5 OS. For detailed software compatibility information, refer to the *Cisco MITG RHEL v5.5 OS Application Note*.

Client Access Requirements

- Workstation supporting Solaris/Sun, Linux, UNIX, Microsoft Windows XP, Windows 2000, Windows 7 or Windows NT operating system
- Java Runtime Environment (JRE) version 1.5 or 1.6

Important: It is recommended that users should use JRE 1.4.2_11 (or later) or 1.5 update 6 (or later).

- Java policy file (obtained during initial access to the Web Element Manager server)
- Microsoft Internet Explorer version 5.0 (or higher), Netscape Navigator version 4.72 (or higher), or other Internet browser
- Access to the Web Element Manager server's host network

Important: Web Element Manager clients cannot access the Web Element Manager server if the server is separated by an NAT'd firewall or other device that restricts access between the client workstation and server.

- Configured application user account on Web Element Manager server

WEM Architecture

The WEM architecture consists of the following components:

- Host Filesystem
- Apache Web Server
- WEM Server FCAPS Support
- WEM Process Monitor
- Bulk Statistics Server
- Script Server
- PostgreSQL Database Server
- Northbound Server
- WEM Logger

Host Filesystem

Running on the fault-tolerant Sun Solaris or Red Hat Enterprise Linux operating system, the WEM uses the native filesystem for such things as creating and writing to log files, storing alarm and bulk statistic-related information, and configuration file management.

Apache Web Server

Remote clients interface with the WEM by establishing session with the server using the Hyper Text Transport Protocol (HTTP). The session is hosted by the Apache Web Server which launches a Java applet providing a graphical user interface for managing the system. When HTTPS is mentioned in the URL instead of HTTP, secure connection is established between the WEM client and WEM server. The Apache Web Server is also used to execute Common Gateway Interfaces (CGIs) invoked by the applet using CORBA/Internet Inter-ORB Protocol (IIOP).

WEM Server FCAPS Support

This component provides Fault, Configuration, Accounting, Performance, and Security (FCAPS) functionality.

Important: The Admin can make any menu or submenu item visible or not visible to users as he sees fit. To show or hide a particular menu option, set the flag in the *menu.xml* file as described in the *Configuration File Parameters* chapter of this guide. Please note that it is not possible to actually delete or add a menu or submenu.

Fault Management

Fault management consists of an event logging function wherein all alarms, warnings, and other faults can be configured, reported, and acknowledged by network operations personnel.

The Simple Network Management Protocol (SNMP) is used by both the Web Element Manager and the chassis to report event notifications. The application's fault management system offers the following support for generated alarms:

- Provide mechanisms for viewing both current and pending alarms for both the chassis and the Web Element Manager server.
- Generate audio and visual alerts for alarms based on their severity (the Web Element Manager also supports the configuration of a severity level for each alarm).
- Maintain statistics for generated alarms.
- Store alarm information in the PostgreSQL® database.
- Execute scripts through the Script Server component of the application.
- Send E-mail notifications and/or forward notifications to Network Management Servers (NMSs) using a CORBA/IIOP-based Northbound Interface.
- Compliance with the following standards:
  - TS 32.111-3, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Fault Management; Part 3: Alarm Integration Reference Point (IRP): Common Object Request Broker Architecture (CORBA) Solution Set (SS)
  - TS 32.303, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Configuration Management (CM); Notification Integration Reference Point (IRP): Common Object Request Broker Architecture (CORBA) Solution Set (SS).

Configuration Management

The Web Element Manager implements an easy to use, point-and-click GUI for providing configuration management for one or more systems. This GUI provides all the flexibility offered by the system's command Line Interface (CLI), while providing the scalability of performing certain functions across multiple systems. All configuration information is stored in the PostgreSQL Database.

At the system-level, the Web Element Manager application provides support for the following:

- Adding, modifying, or deleting systems to/from the management system
- Performing configuration of card and port-level parameters
- Adding, modifying, or deleting contexts
- Configuring specific protocols and services within defined contexts such as AAA servers, PDSN services, GGSN services, IP access lists, IP interfaces, IP routes, IP address pools, RADIUS accounting and authentication, PPP, subscribers, and others

At the network level, the application is capable of transferring configuration and/or software images to multiple systems simultaneously in advance to performing software upgrades.

The Web Element Manager supports the configuration of all parameters required to perform software upgrades including:

- Adding, deleting, and sorting system boot stack entries; these entries allow multiple fall-backs in the event the system experiences an error in the loading of a particular image or configuration file.
- Configuring network options for bootup.
- Transferring configuration and image files to/from systems.
- Initiating and monitoring upgrade status.

The Web Element Manager further simplifies the software upgrade process by providing tools for managing system configuration files:

- **Back-up Tool:** Enables the Web Element Manager to transfer a copy of the configuration file currently being used by a managed system at user-defined intervals. Files are transferred to the host server in a specific directory. The number of files to retain in the directory is also configurable. This tool provides a useful mechanism for testing configurations and/or quickly restoring a last-known-good configuration in the event of an error.
- **Compare Tool:** Provides a powerful tool for comparing the configuration files of two managed systems. Once the two files are specified, a dialog appears displaying the two documents side-by-side. Line numbers are added for convenience. Text additions, modifications, and deletions are displayed in different colors for easy recognition. This tool can be useful on its own to determine variations between multiple iterations of the same configuration file, or, when used in conjunction with the Back-up tool, it can provide an audit trail of configuration changes that occurred during system operation.

Accounting Management

Accounting management operations allow users to examine and perform post-process statistical analysis on systems managed by the Web Element Manager application.

The type of statistics used for element management-based accounting are called bulk statistics. Bulk statistics are grouped into categories called schemas and are polled by the system at fixed polling intervals and then transferred to the Web Element Manager at a different transfer intervals (defined in minutes).

Once the Web Element Manager server application, called the receiver, has received bulk statistics files from the managed system, these files are parsed and added to the PostgreSQL database. This database is updated as new files are received.

The Web Element Manager's accounting management functionality is compliant with *TS 32.401, 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Telecommunication management; Performance Management (PM); Concept and requirements* and allows you to:

- Collect statistics pertaining to the transfer and collection of bulk statistics.
- Views statistics stored on the chassis prior to transfer to the receiver.
- Graph multiple received bulk statistics over time as either a line or bar graph; these graphs can be printed to network printers accessible by the server.
- Generate eXtensible Markup Language (XML) files for transfer to a Northbound NMS or bulk statistics processor.
- Archive collected bulk statistic information to conserve disk space on the server.

Performance Management

Performance management operations supported by the Web Element Manager allow users to examine and perform real-time statistical analysis on systems managed by the application as well as on the server on which the application is running.

Information pertaining to various aspects of the Web Element Manager (CPU and memory utilization, disk space, and process status) and its managed systems (hardware, protocols, software subsystems, and subscribers) is collected in real time and is displayed in tabular format. Alternatively, most of the information can be graphed as a function of time in either line or bar-chart format. Multiple statistics can be graphed simultaneously for quick comparison of data.

In addition to collecting and providing mechanisms for the real-time viewing of statistical information, the Web Element Manager provides useful monitoring tools similar to those found in the CLI. These tools can be used to monitor

active subscriber sessions, protocol flows, and port information. Data collected during this monitor operation can be saved to the client machine for further analysis.

Security Management

Security management pertains to the operations related to management users. This includes both Web Element Manager application users and local management users who are configured on the chassis. In many cases, management users can be allowed access to both the system (via its CLI) and the application. It is possible for both management user accounts to share the same username and password.

While it is possible to authenticate users via RADIUS server configuration, each RADIUS server has its own configuration that is outside the scope of this document. However, you can find the mapping information you would need to configure each of the four levels of administrative user in the "CLI Administrative Users" section of the *CLI Overview* chapter in the *Command Line Interface Reference*.

The security management features of the Web Element Manager allow you to:

- Add, modify, or delete administrative users for both the application and the managed system.
- Regardless of the administrative user type, there are four levels of management user privileges:
    - **Inspector**: Inspectors are limited to a small number of read-only Exec Mode commands. The bulk of these are "show" commands giving the inspector the ability to view a variety of statistics and conditions. The Inspector cannot execute **`show configuration`** commands and do not have the privilege to enter the Config Mode.
    - **Operator**: Operators have read-only privileges to a larger subset of the Exec Mode commands. They can execute all commands that are part of the inspector mode, plus some system monitoring, statistic, and fault management functions. Operators do not have the ability to enter the Config Mode.
    - **Administrator**: Administrators have read-write privileges and can execute any command throughout the CLI except for a few security-related commands that can only be configured by Security Administrators. Administrators can configure or modify the system and are able to execute all system commands, including those available to the Operators and Inspectors.
    - **Security Administrator**: Security Administrators have read-write privileges and can execute all CLI commands including those available to Administrators, Operators, and Inspectors.
- Provide authentication and privilege restoration based on the login information entered by administrative users.
- Monitor current system or application-level administrative users in real-time and optionally terminate their management session.
- Perform an audit of all managed system configurations performed through both the application and the CLI as well as other operations performed within the application.

  The audit trail functionality supports the configuration of filters defining the type of operations to audit and also provides a dialog for performing the audit in real-time.

  Audit trail results are stored in the PostgreSQL database for later retrieval and analysis.

  The new Network Audit Tool functionality in WEM supports the on-demand or periodic auditing of chassis configuration attributes such as PPP MRU, Auth Sequence, Bulkstats Schema Needs Update, etc.

ANSI T1.276 Compliance

The WEM supports ANSI standard T1.276, providing a set of baseline security features to help mitigate security risks in the management of telecommunication networks. New users will be sent a randomly generated password automatically, and will be prompted to provide a new password upon first login. New passwords must meet strict requirements to comply with the ANSI standard:

- Passwords must be a minimum of eight characters long.
- Passwords must not be a repeat or the reverse of the associated user ID.
- Passwords must not be more than three of the same characters used consecutively.
- Passwords must contain at least three of the following character types:
    - At least one lower case alpha character
    - At least one upper case alpha character
    - At least one numeric character
    - At least one special character

Users will also be required to change passwords after a configurable number of days, and will be barred from reusing the same password for a configurable number of password change cycles. Too many failed login attempts will result in an account lockout, which may be removed either by an administrator or by waiting for a defined period of time to elapse.

WEM Process Monitor

The Process Monitor (PSMon) is a Perl script that monitors the status of processes pertaining to the WEM application.

The script is a plain text Apache-style configuration file that allows the user to define a set of rules. These rules describe what processes should always be running on the system, any limitations on concurrent instances, Time-To-Live (TTL), and maximum CPU/memory usage of processes. It can be run as a stand alone program or a fully functional background daemon.

PSMon scans the UNIX process table and, using the set of defined rules, will re-spawn any dead processes, and/or slay or "deal with" any aggressive or illegal processes. The number of retries and time interval the PSMon scans the table is configurable meaning that it will never try to start the process if 'number of retries' exceeds in given time interval.

PSMon logs events to syslog and to a log file and is equipped with customizable e-mail notification facilities.

Bulk Statistics Server

The Bulk Statistics Server process is responsible for collecting and processing all bulk statistic-related information from the system as part of the WEM's accounting management functionality.

The Bulk Statistics Server parses collected statistics and stores the information in the PostgreSQL database. If XML file generation and transfer is required, this element generates the XML output and can send it to a Northbound NMS or an alternate bulk statistics server for further processing.

Additionally, if archiving of the collected statistics is desired, the Bulk Statistics server writes the files to an alternative directory on the server. A specific directory can be configured by the administrative user or the default directory can be used. Regardless, the directory can be on a local filesystem or on an NFS-mounted filesystem on the WEM server.

Script Server

The WEM supports the ability to configure the properties for alarms. One of the properties that can be configured is specifying a script that can be executed upon receipt of that alarm. The Script Server process is responsible for executing the specified script.

Upon receipt of the alarm, the WEM Server FCAPS Support function passes the name of the script to execute and the trap logged time to the Script Server. An acknowledgement is sent and the script is executed by the Script Server. In the event, an error is experienced while executing the script, the Script Server generates an SNMP trap.

PostgreSQL Database Server

The PostgreSQL Database consists of multiple databases maintaining information pertaining to the following WEM functions:

- **Configuration:** This database contains tables which maintain configuration information for user details, topology for maps and manageable systems.
- **Trap:** This database contains tables which maintain SNMP trap configuration information and all the received SNMP traps.
- **MIB:** This database contains all the information required to translate SNMP Object identifiers to proper MIB names and their types as given in the MIB file.
- **Audit Trail:** This database contains table that maintains the configuration trail including the following:
  - Configuration performed on each system through the WEM
  - Configuration done through the system's CLI (this is known via the CORBA notification service)
  - Login/out from the WEM and system CLI
  - The addition/deletion of a new system in the managed system list
- **Bulk Statistics:** This database contains various tables containing counter values periodically received from the system via the File Transfer Protocol (FTP).

Northbound Server

The Northbound Server process is responsible for collecting and transmitting information about WEM system management to an NMS.

WEM supports the Northbound Interface as defined in the 3GPP standards for Telecom Management. 3GPP defines a standard interface (Interface-N) between the EMS and the NMS. It also defines Integration Reference Points (IRPs) through which various aspects of system management (FCAPS) are performed by the NMS.

When the Northbound Server process is enabled, WEM will respond to NMS requests by fetching the required information and transmitting it over the CORBA interface.

Northbound Server is a separately licensed feature.

Currently, WEM supports five IRPs defined by the 3GPP standards. The supported IRPs and corresponding 3GPP standards are:

- **Alarm**: 3GPP TS 32.111-3 (V6.6.0)
- **Basic CM**: 3GPP TS 32.603 (V6.4.0)
- **Notification**: 3GPP TS 32.303 (V6.6.0)
- **Communication Surveillance**: 3GPP TS 32.353 (V6.4.0)
- **Entry Point**: 3GPP 32.363 (V6.4.0)

The supported Network Resource Model is **Generic Network Resource (NRM) IRP**: 3GPP TS 32.623 (v6.50)

The following table lists the Northbound Interface operations and notifications that are supported for each of the IRPs:

IRP Category	Notification/Operation Description	
Alarm	*Notifications*	notifyNewAlarm
		notifyAckStateChanged
		notifyClearedAlarm
		notifyAlarmListRebuilt
	Operations	get_alarm_irp_versions
		clear_alarms
		get_alarm_list
		next_alarmInformations
		get_alarm_count
		Basic CM
	Operations	get_basic_cm_irp_version
		find_managed_objects
		next_basicCmInformations
Notification		not applicable
	Operations	get_notification_irp_versions
		attach_push
		change_subscription_filter
		get_subscription_status
		get_subscription_ids
		detach
Communication Surveillance	*Notifications*	notifyHeartbeat
	Operations	get_CS_IRP_versions
		get_heartbeat_period
		set_heartbeat_period
		trigger_heartbeat
Entry Point	*Notifications*	Not Applicable
	Operations	get_EP_IRP_versions
		get_IRP_outline
		get_IRP_reference
		release_IRP_reference

WEM Logger

The WEM application generates and stores logs pertaining to server installation and operation. The logs can be stored locally or to another server. In addition, the WEM provides enhanced logging functionality for customizing log output and log files.

Chapter 34
Technical Specifications

This chapter lists physical dimensions, power specifications, mounting requirements and interface specifications for ASR 5000 system components.

It includes the following sections:

- Physical Dimensions
- Weights
- Power Specifications
- Mounting Requirements
- Interface Specifications

Physical Dimensions

The ASR 5000 can be mounted in any standard (EIA-310-D, IEC 60297) 19-inch (482.6 mm) equipment cabinet or telecommunications rack. The table below lists the dimensions for the chassis and each component that can be placed within the chassis.

Table 117. Physical Dimensions - ASR 5000 Chassis and Components

Component	Height	Width	Depth
Chassis	24.50 in. (62.23 cm)	17.5 in. (44.45 cm)	24.0 in. (60.96 cm)
Application Card	17.05 in. (46.31 cm)	1.01 in. (2.56 cm)	14.10 in. (35.81 cm)
Line Card (half-height)	8.59 in. (21.82 cm)	1.01 in. (2.56 cm)	5.24 in. (13.31 cm)
XGLC (full-height)	17.48 in. (44.40 cm)	1.01 in. (2.56 cm)	5.24 in. (13.31 cm)
Fan Tray (Lower)	2.50 in. (6.35 cm)	16.25 in. (41.27 cm)	17.25 in. (43.82 cm)
Fan Tray (Upper)	2.875 in. (7.30 cm)	16.25 in. (41.27 cm)	19.375 in. (49.21 cm)
Power Filtering Unit (PFU)	3.6 in. (9.14 cm)	8.25 in. (20.96 cm)	5.12 in. (13.00 cm)

Weights

The following table identifies the maximum weights for fully-loaded systems—cards installed in all slots and all other components installed.

Table 118. ASR 5000 Component Weights

Component	Weight
Chassis	
Empty	65 lbs. (29.48 kg)
As Shipped (empty chassis with PFUs, fan trays, bezels and blanking panels)	160 lbs. (72.57 kg)
Shipping (as shipped chassis, shipping container and packing materials)	251 lbs. (113.85 kg)
Fully loaded (as shipped chassis with all slots filled with cards)	307 lbs. (139.25 kg)
Packet Processing Cards	
Packet Services Card (PSC)	11.50 lbs. (5.22 kg)
Packet Services Card 2 (PSC2)	11.50 lbs. (5.22 kg)
Packet Service Card 3 (PSC3)	11.0 lbs. (4.95 kg)
Packet Processing Card (PPC)	11.50 lbs. (5.22 kg)
Switch Process I/O Card (SPIO)	1.25 lbs. (0.57 kg)
System Management Card (SMC)	10.00 lbs. (4.54 kg)
Line Cards	
Channelized Line Card (CLC)	1.25 lbs. (0.57 kg)
Channelized Line Card 2 (CLC2)	1.25 lbs. (0.57 kg)
Fast Ethernet (10/100) Line Card (FELC)	1.00 lbs. (0.45 kg)
Fast Ethernet (10/100) Line Card 2 (FLC2)	1.00 lbs. (0.45 kg)
Gigabit Ethernet Line Card (GELC/GLC2)	1.00 lbs. (0.45 kg)
Optical Line Card (OLC)	1.25 lbs. (0.57 kg)
Optical Line Card 2 (OLC2)	1.25 lbs. (0.57 kg)
Quad Gigabit Ethernet Line Card (QGLC)	1.25 lbs. (0.57 kg)
Redundancy Crossbar Card (RCC)	1.00 lbs. (0.45 kg)
10 Gigabit Ethernet Line Card (XGLC)	2.25 lbs. (1.02 kg)

Power Specifications

The following table provides essential power specifications for the chassis and all associated cards within the system.

Table 119. Chassis Power Requirements

Characteristic	Value
Input Voltage	Maximum range: -40VDC to -60VDC Nominal range: -48VDC to -60 VDC
TUV Rated Peak Current Load	165A @ -48 VDC
Maximum Peak Power Load	5760W
Chassis Maximum Power Load	800W
Line Card (rear-installed) Maximum Power Load	**SPIO:** 15W **FELC/FLC2:** 13.5W **GELC/GLC2:** 10.5W **QGLC:** 15W **XGLC:** 25W **OLC2:** 23W **CLC2:** 23W **RCC:** 20W
Application Card (front-installed) Maximum Power Load	**SMC:** 130W **PPC:** 325W **PSC:** 250W **PSC2:** 325W **PSC3:** 330W

Estimating Power Requirements

Use the following formula to estimate total power consumption for each deployed chassis.

(Total Application Card Maximum Power Load) + (Total Line Card Maximum Power Load) + (Chassis Maximum Power Load)

For example, the calculation for estimating the power required for an ASR 5000 installation with 3 PSCs, 2 SMCs, 2 SPIOs, 2 RCCs, and 4 GELC/GLC2s would be:

(250W x 3) + (130W x 2) + (20W x 2) + (13.5W x 4) + 800W = 1934W

Mounting Requirements

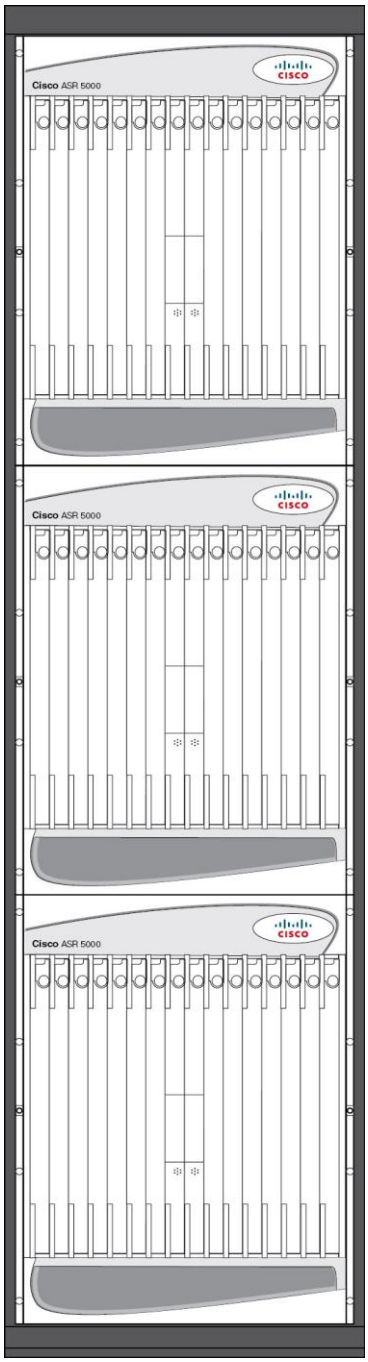
Each 24.5 in. (62.23 cm.) height chassis requires 14 Rack Units (RUs) of space. You can mount the system into any 19-inch (482.6 mm) equipment rack or telco cabinet with the mounting brackets supplied with the chassis. Additional hardware (not supplied), such as extension brackets, may be used to install the chassis in a standard 23-inch (584.2 mm) cabinet or rack. Both front and mid-mount installations are possible, depending on the position of the mounting brackets on the chassis.

You can mount a maximum of three ASR 5000 chassis in a 2- or 4-post equipment rack or telco cabinet, provided that all system cooling and ventilation requirements are met. Three stacked chassis will occupy a minimum of 42 RUs.

Caution: When planning chassis installation, ensure that equipment rack or cabinet hardware does not hinder air flow at any of the intake or exhaust vents. Also, make sure that the rack/cabinet hardware, as well as the ambient environment, allow the system to function within the required limits. For more information, refer to *Environmental Specifications* in this guide.

Rack mounting requires the use of industry-standard (EIA-310-D, IEC 60297) equipment racks and cabinets, as well as supplier-recommended fasteners. The following figure depicts how three chassis can be mounted in a 42 RU equipment rack.

Figure 265. Three ASR 5000 Chassis in a 42 RU Rack



Interface Specifications

Following is a list of interfaces for use within the chassis. Each interface is shown with its specific pin-out.

Important: Some interfaces, such as an RJ-45 interface used for Ethernet connectivity, may have more than one pin-out configuration, depending on the type of cable used.

SPIO Card Interfaces

Each interface on the SPIO card is described below. In each accompanying figure, the interface is shown in the same orientation as the way it appears on the card.

Console Port Interface

The Console port is an RJ-45 RS-232 interface used to access the command line interface. The interface communicates at a baud rate of 9600 to 115,200 bps (115.2 Kbps). The default is 115,200 bps.

The interface's pin out detail is provided in the following figure and table.

Figure 266. SPIO Console Port Pinout

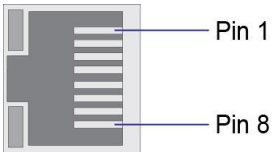


Table 120.SPIO Console Port Pinout

Pin	Signal Description	Signal Type
1	Clear to Send (CTS)	Input
2	Data set Ready (DSR)	Input
3	Receive Data (RX)	Input
4	Signal Ground (SGND)	N/A
5	Ready to Send (RTS)	Output
6	Transmit Data (TX)	Output
7	Data Carrier Detect (DCD)	Input
8	Data Terminal Ready (DTR)	Output

Console Cable Specifications

SPIO cards are shipped with a console cable assembly that includes a 7-foot (2 meter) serial cable with RJ-45 connectors on each end, and an RJ-45-to-DB-9 adapter. Use the RJ-45-to-DB-9 adapter to connect the console cable to a terminal server or terminal emulation device such as a laptop computer. The cable's pin-out is provided in the following figure and table.

Figure 267. SPIO Console Cable Assembly

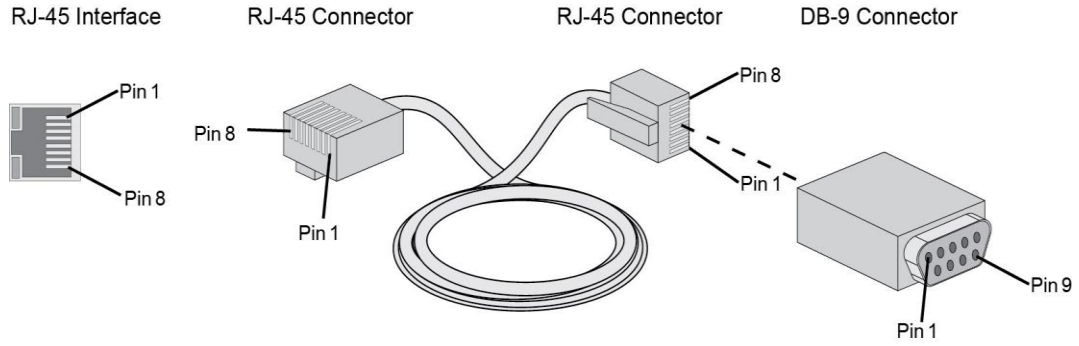


Table 121.RJ-45 to DB-9 Cable

Signal Description	Signal Type	RJ-45 Pin	DB-9 Pin	Signal
Clear to Send (CTS)	Input	1	7	RTS
Data set Ready (DSR)	Input	2	4	DTR
Receive Data (RxD)	Input	3	3	TxD
Signal Ground (SGND)	N/A	4	5	SGND
Ready to Send (RTS)	Output	5	8	CTS
Transmit Data (TxD)	Output	6	2	RxD
Data Carrier Detect (DCD)	Input	7	1	DCD
Data Terminal Ready (DTR)	Output	8	6	DSR

To construct a RJ-45 to DB-25 cable for modem connectivity, refer to the table that follows.

Table 122. RJ-45 to DB-25 Cable

Signal Description	Signal Type	RJ-45 Pin	DB-25 Pin	Signal
Clear to Send (CTS)	Input	1	5	CTS
Data set Ready (DSR)	Input	2	6	DSR
Receive Data (RX)	Input	3	3	RxD
Signal Ground (SGND)	-	4	7	SGND
Ready to Send (RTS)	Output	5	4	RTS

Signal Description	Signal Type	RJ-45 Pin	DB-25 Pin	Signal
Transmit Data (TX)	Input	6	2	TxD
Data Carrier Detect (DCD)	Output	7	8	DCD
Data Terminal Ready (DTR)	Output	8	20	DTR

Fiber SFP Interface

The fiber SFP interface has two host connectors that receive SFP transceivers.

Figure 268. SPIO Gigabit Ethernet Fiber SFP Pinout

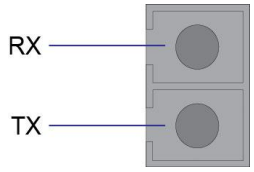


Table 123. Fiber SFP Interface Transmit and Receive Levels

Signal	Level
Max TX:	0 dBm
Min TX:	-9.5 dBm
Max RX:	0 dBm (saturation average power)
Min RX:	-20 (typical) / -17 (max.) dBm (sensitivity average power)

10/100/1000 Mbps RJ-45 Interface

The two RJ-45 interfaces are auto-sensing 10/100/1000 Ethernet (10Base-T/100Base-TX/1000Base-T) that require unshielded twisted pair (UTP) copper cable. Refer to the following figure and table for pin-outs for the RJ-45 Ethernet ports.

Figure 269. SPIO RJ-45 Ethernet Interface Pinouts

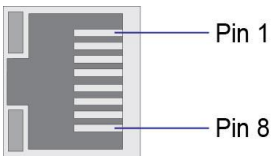


Table 124.SPIO RJ-45 Ethernet Interface Pinouts

Pin	10Base-T 10Mbps Cat3	100Base-Tx 100Mbps Cat5	1000Base-Tx 1Gbps Cat5+
1	TX+	TX+	BI DA+
2	TX-	TX-	BI DA-
3	RX+	RX+	BI DB+
4	Not used	Not used	BI DC+
5	Not used	Not used	BI DC-
6	RX-	RX-	BI DB-
7	Not used	Not used	BI DD+
8	Not used	Not used	BI DD-

Central Office Alarm Interface

The Central Office (CO) alarm interface is a 10-pin Molex connector supporting three dry-contact relay switches. The three normally closed (NC) relays can support normally open (NO) or NC devices. The following two figures show the pin-out details for this interface and the next figure shows an example CO alarm configuration.

Figure 270. SPIO CO Alarms Interface Pinout

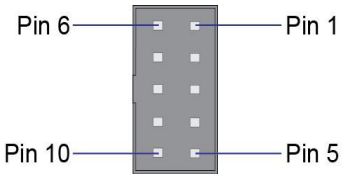


Table 125.SPIO CO Alarms Interface Pinout

Pin	Signal
1	Major Alarm - Normally closed
2	Major Alarm - Common
3	Major Alarm - Normally open
4	Minor Alarm - Normally closed
5	Minor Alarm - Common
6	Minor Alarm - Normally open
7	Critical Alarm - Normally closed
8	Critical Alarm - Common
9	Critical Alarm - Normally open

Pin	Signal
10	Not used

The 8-foot (2.4 meter) CO alarm cable shipped with the chassis supports redundant SPIO card installations. This "Y" cable has two Molex connectors on one end that are keyed to fit into the CO Alarm interfaces in one direction only. Each connector mates with one of the side-by-side SPIO cards. On the opposite end is a 9-pin terminal block that you can mount to the telco cabinet or equipment rack frame.

The following figure and table display this cable assembly and its pinouts.

Figure 271. SPIO CO Alarms Cable Assembly

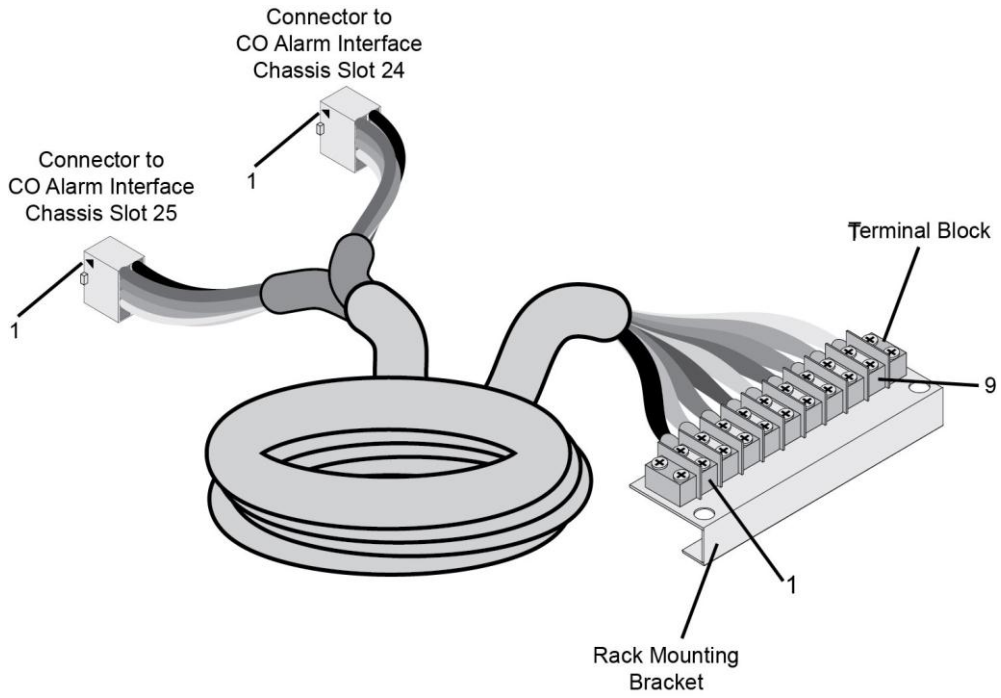


Table 126.CO Alarms Cable Pinout

CO Alarms IF Pin No.	Cable Wire Color	Cable Terminal Block Position No.	Signal
1	Black	1	Major Alarm - Normally closed
2	Orange	2	Major Alarm - Common
3	Red	3	Major Alarm - Normally open
4	Brown	4	Minor Alarm - Normally closed
5	Yellow	5	Minor Alarm - Common
6	Green	6	Minor Alarm - Normally open
7	Blue	7	Critical Alarm - Normally closed

CO Alarms IF Pin No.	Cable Wire Color	Cable Terminal Block Position No.	Signal
8	Violet	8	Critical Alarm - Common
9	Gray	9	Critical Alarm - Normally open
10	Not wired	Not equipped	Not applicable

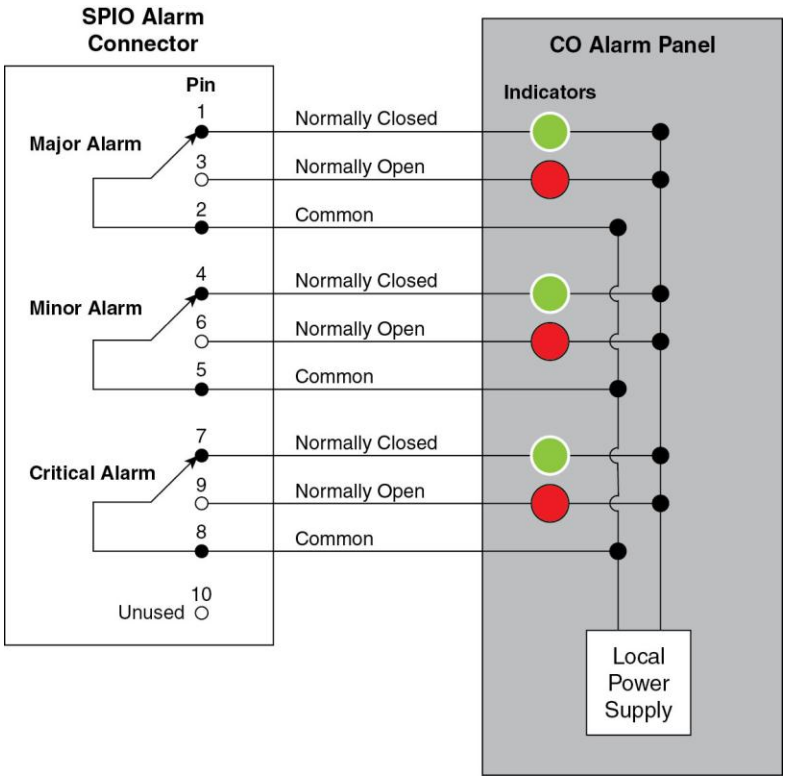
Electrical Characteristics

Each of the three dry-contact relay switches is rated to support a maximum switching current of 1A@30VDC. The relay contacts should not directly connected to high current devices such as sirens and flashing lamps.

Central Office Alarm Wiring Example

The following figure depicts how the dry-contact relays can each control up to two external alarm indicators. In this example, the CO alarm interface is connected to a CO Alarm Panel, where green LEDs are wired to indicate normal operation, and red LEDs are wired to indicate an alarm condition.

Figure 272. CO Alarm Interface Schematic



With all relays de-energized (normally closed), the green LED is illuminated. If an alarm relay is energized, the NO (normally open) contact closes and the red LED is illuminated.

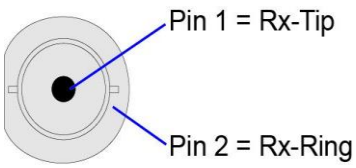
BITS Timing Interface

Important: This interface is not used on SPIOs when the system is configured to perform data services.

BITS BNC Timing Interface

The BNC version of the SPIO interface card uses a BNC connector. The following figure shows the BITS BNC timing interface.

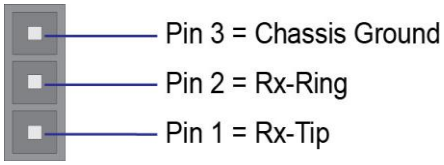
Figure 273. SPIO BITS BNC Timing Interface Pinout



BITS 3-Pin Timing Interface

This 3-pin version of the SPIO interface card uses a 3-pin wire-wrap connector instead of a BNC interface. The following figure shows the BITS 3-wire timing interface wire-wrap pin-out.

Figure 274. SPIO T1 BITS Timing Wire-Wrap Pinout



Fast Ethernet Line Card (FELC/FLC2) Interfaces

Each of the eight RJ-45 interfaces available on the FELC/FLC2 supports auto-sensing 10Base-Tx or 100Base-Tx Ethernet interfaces.

10/100 Mbps RJ-45 Interface

The RJ-45 interfaces on the Fast Ethernet line card support the following cable types and transfer rates. The following figure shows the pin-outs for the RJ-45 Ethernet ports.

Figure 275. FELC/FLC2 RJ-45 Ethernet Interface Pinouts

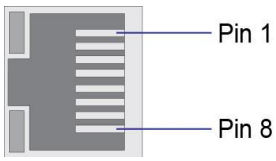


Table 127. FELC/FLC2 RJ-45 Ethernet Interface Pinouts

Pin	10Base-T 10Mbps Cat3	100Base-TX 100Mbps Cat5
1	TX+	TX+
2	TX-	TX-
3	RX+	RX+
4	na	na
5	na	na
6	RX-	RX-
7	na	na
8	na	na

Gigabit Ethernet Card (GELC/GLC2/QGLC) SFP Interfaces

1000Base-SX

The 1000Base-SX fiber SFP interface on the Gigabit Ethernet Card (GELC/GLC2) has one pair of fiber connectors, as shown below. The Quad Gigabit Ethernet Card (QGLC) has four pairs.

Figure 276. GELC/GLC2/QGLC Fiber Connector

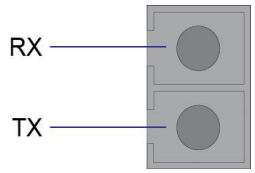


Table 128. 100Base-SX Fiber Interface Transmit and Receive Levels

Signal	Level
Max TX:	0 dBm
Min TX:	-9.5 dBm
Max RX:	0 dBm (saturation average power)
Min RX:	-20 (typical) / -17 (max.) dBm (sensitivity average power)

1000Base-LX Interface

The 1000Base-LX fiber SFP interface on the Ethernet 1000 LX line card has one pair of host connectors. The QGLC has four pairs.

Figure 277. QGLC/1000 Base-LX Fiber Connector

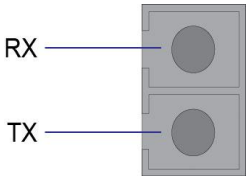


Table 129. 1000Base-LX Fiber Interface Transmit and Receive Levels

Signal	Level
Max TX:	0 dBm
Min TX:	-9.5 dBm
Max RX:	0 dBm (saturation average power)
Min RX:	-20 (typical) / -19 (max.) dBm (sensitivity average power)

1000Base-T

The 1000Base-T SFP copper interfaces on the GELC/GLC2 and QGLC line cards require unshielded twisted pair (UTP) copper CAT-5 cable with a bit error rate (BER) less than 10e-10. Pinouts for the RJ-45 Ethernet ports are shown in the table below.

Figure 278. GELC/GLC2/QGLC RJ-45 Copper Interface Pinouts

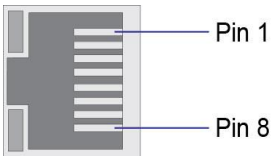


Table 130. 1000Base-T RJ-45 Ethernet Copper Interface Pinouts

Pin	1000Base-Tx 1Gbps Cat5+
1	BI DA+
2	BI DA-
3	BI DB+
4	BI DC+

Pin	1000Base-Tx 1Gbps Cat5+
5	BI DC-
6	BI DB-
7	BI DD+
8	BI DD-
RX = Receive Data TX = Transmit Data BI = BI directional data DA, DB, DC, DD = Data Pair A, B, C, and D	

10 Gigabit Ethernet Line Card (XGLC) SFP+

10GBase-SR

The 10GBase-SR fiber SFP+ interface on the XGLC has one pair of fiber connectors, as shown below.

Figure 279. XGLC 10GBase-SR Fiber Connector

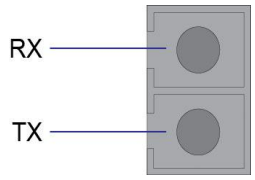


Table 131. 10GBase-SR Fiber Interface Transmit and Receive Levels

Signal	Level
Max TX:	-1.0 dBm
Min TX:	-7.3 dBm
Max RX:	-1.0 dBm (saturation average power)
Min RX:	-11.1 (max.) dBm (sensitivity average power)

10 Base-LR Interface

The 10GBase-LR fiber SFP+ interface on the XGLC has one pair of host connectors.

Figure 280. XGLC 10GBase LR Fiber Connector

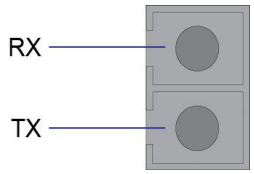


Table 132.10GBase-LR Fiber Interface Transmit and Receive Levels

Signal	Level
Max TX:	0.5 dBm
Min TX:	-8.2 dBm
Max RX:	0.5 dBm (saturation average power)
Min RX:	-12.6 (max.) dBm (sensitivity average power)

Fiber ATM/POS OC-3 (OLC2) Multi-Mode Interface

Fiber ATM/POS OC-3 SM IR-1 Interface

The fiber-optic SFP interface on OLC2 Optical ATM Line Cards with the SM IR-1 interface has one pair of host connectors as shown in The following figure.

Figure 281. OLC2 (ATM) SM IR-1 SFP Pinout

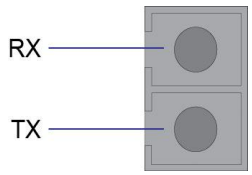


Table 133.OC-3 SM IR-1 Fiber Interface Transmit and Receive Levels

Signal	Level
Max TX:	-8 dBm
Min TX:	-15 dBm
Max RX:	-8 dBm (saturation average power)
Min RX:	-28 (max.) dBm (sensitivity average power)

The fiber-optic SFP interface on OLC2 Optical ATM Line Cards with the multi-mode interface has one pair of host connectors as shown in figure that follows.

Figure 282. OLC2 (ATM) Multi-Mode SFP Pinout

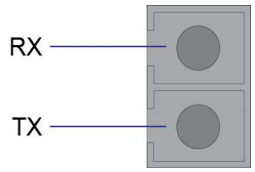


Table 134. Multi-Mode Fiber Interface Transmit and Receive Levels

Signal	Level
Max TX:	-14 dBm
Min TX:	-19 dBm
Max RX:	-12 dBm (saturation average power)
Min RX:	-30 (max.) dBm (sensitivity average power)

Channelized Line Cards

Channelized Line Card (CLC2) with Single-Mode Interface

The optical SFP interface on the 4-port CLC2 with the single-mode interface has four pairs of connectors that receive SFP transceivers, as shown in the following figure.

Figure 283. Channelized Line Cards with Single-Mode SFP Pinout

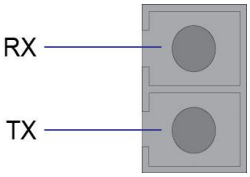


Table 135. Single-Mode Fiber Interface Transmit and Receive Levels

Signal	Level
Max TX:	-8 dBm
Min TX:	-15 dBm
Max RX:	-8 dBm (saturation average power)
Min RX:	-28 (max.) dBm (sensitivity average power)

Channelized Line Cards (CLC2) with Multi-Mode Interface

The fiber SFP interface on the 4-port CLC2 with the multi-mode interface has four pairs of connectors that receive SFP transceivers, as shown in the following figure.

Figure 284. Channelized Line Cards with Multi-Mode SFP Pinout

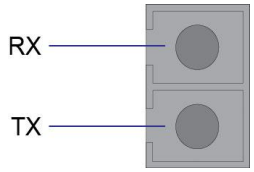


Table 136. Multi-Mode Fiber Interface Transmit and Receive Levels

Signal	Level
Max TX:	-14 dBm
Min TX:	-19 dBm
Max RX:	-12 dBm (saturation average power)
Min RX:	-30 (max.) dBm (sensitivity average power)

Chapter 35
Safety, Electrical and EMC Certifications

This chapter lists FCC warnings, as well as safety, electrical and environmental certifications for the ASR 5000 system.

This chapter includes the following sections:

- Federal Communications Commission Warning
- Safety Certifications
- Electrical Certifications
- Environmental Certifications
- Acoustic Noise
- Electromagnetic Compatibility (EMC) Compliance

Federal Communications Commission Warning

The ASR 5000 complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules and Regulations. Operation is subject to the following two conditions:

- This device must not cause harmful interference.
- This device must withstand any interference received, including interference that may cause undesired operation.

These limits provide reasonable protection against harmful interference when this equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio and television communications. Operation of this equipment in a residential area is likely to cause interference, in which case your organization is responsible for the expenses incurred to correct the interference.

Modifications to this product not authorized by Cisco could void the FCC approval and negate your authority to operate the product.

Important: Shielded cables must be used with this unit to ensure compliance with the FCC Class A limits. Blanking panels must be installed in the chassis.

ICS Notice

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Laser Notice

The lasers in this equipment are Class 1 devices. Class 1 laser devices are not considered to be hazardous.

Safety Certifications

The ASR 5000 complies with all safety certifications listed below.

- UL60950 - Standard for Safety for Information Technology Equipment, 3rd Edition
- European Union EN 60950 (CE Mark)

Electrical Certifications

The ASR 5000 complies with all electrical certifications listed below.

- Telcordia GR-1089-Core, Network Equipment-Building System (NEBS) Requirements: Electromagnetic Compatibility and Electrical Safety Criteria for Network Telecommunication Equipment
- FCC, Part 15 B, Class A Requirements for Non-residential Equipment
- ETSI EN 300 019
- ETSI 300 386
- ETSI/EN 300 386-2 Electrical Fast Transients
- SBC TP76200MP
- Taiwan - BMSI

Environmental Certifications

The ASR 5000 complies with all environmental certifications listed below.

- Telcordia GR-63-Core, Network Equipment-Building System (NEBS) Requirements: Physical Protection
- The chassis equipped with the 165A PFU is compliant to the European Union's RoHS Directive (Directive 2002/95/EC)
- Waste Electrical and Electronic Equipment (WEEE) Directive 2002/96/EC

Acoustic Noise

The maximum acoustic noise level for the ASR 5000 chassis is 76 dBA.

Caution: The maximum acoustic noise level of the ASR 5000 exceeds 70 dBA.

Caution: Maschinenlärminformations-Verordnung - 3. GPSGV, der höchste Schalldruckpegel beträgt 76 dB(A) gemäss EN ISO 7779.

Electromagnetic Compatibility (EMC) Compliance

Electromagnetic compatibility is the ability of electronic devices to operate as intended in proximity to other electronic devices or in the presence of electromagnetic fields. Unintentional radio frequency emissions from an electronic device and immunity of the device to radio frequency interference from other electromagnetic sources are included within electromagnetic compatibility.

Japan VCCI-A

The ASR 5000 has been registered for compliance with the Voluntary Council for Control of Interference, VCCI.

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 VCCI－A

Korean EMC

Class A device (Broadcasting Communication Device for Office Use): This device obtained EMC registration for office use (Class A), and may be used in places other than home. Sellers and/or users need to take note of this.

A급 기기 (업무용 방송통신기기): 이 기기는 업무용(A급)으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다.

Chapter 36
Environmental Specifications

This chapter provides information related to environmental considerations and storage characteristics associated with the ASR 5000.

This chapter includes the following sections:

- Operating and Storage Parameters
- Supported Environmental Standards
- Chassis Air Flow

Operating and Storage Parameters

Use the following information to plan your network installation for the ASR 5000 platform.

Table 137. Temperature, Humidity and Altitude Recommendations

Temperature	
Operating	0 degrees C to +55 degrees C
Storage	-40 degrees C to +70 degrees C
Humidity	
Operating	20 to 80 percent non-condensing
Storage	10 to 95 percent non-condensing
Altitude	
Operating	197 ft. (60m) below to 13,123 ft. (4,000m) above sea level
Non-operating	197 ft. (60m) below to 49,212 ft. (15,000m) above sea level

Supported Environmental Standards

The system has been successfully tested against the following environmental standards:

- Operational Thermal, Operating Conditions - GR-63 Criteria [72, 73]
- Airborne Contaminants, Indoor Levels - GR-63 Criterion [125]
- Operational Thermal, Short-term Conditions - GR-63 Criteria [72, 73]
- Storage Environments, and Transportation and Handling - GR-63 Criteria [69-71, 107-109, 124]
- Earthquake Zone 4 - GR-63 Criteria [110-112, 114, 115, 117, 119]
- Airborne Contaminants, Outdoor Levels - GR-63 Criteria [126, 127]
- Altitude - GR-63 Criteria [74, 76]
- Thermal Heat Dissipation - GR-63 Criteria [77-79]
- Acoustic Noise - GR-63 Criterion [128]
- ESTI 300 019 - Environmental conditions and environmental tests for telecommunications equipment

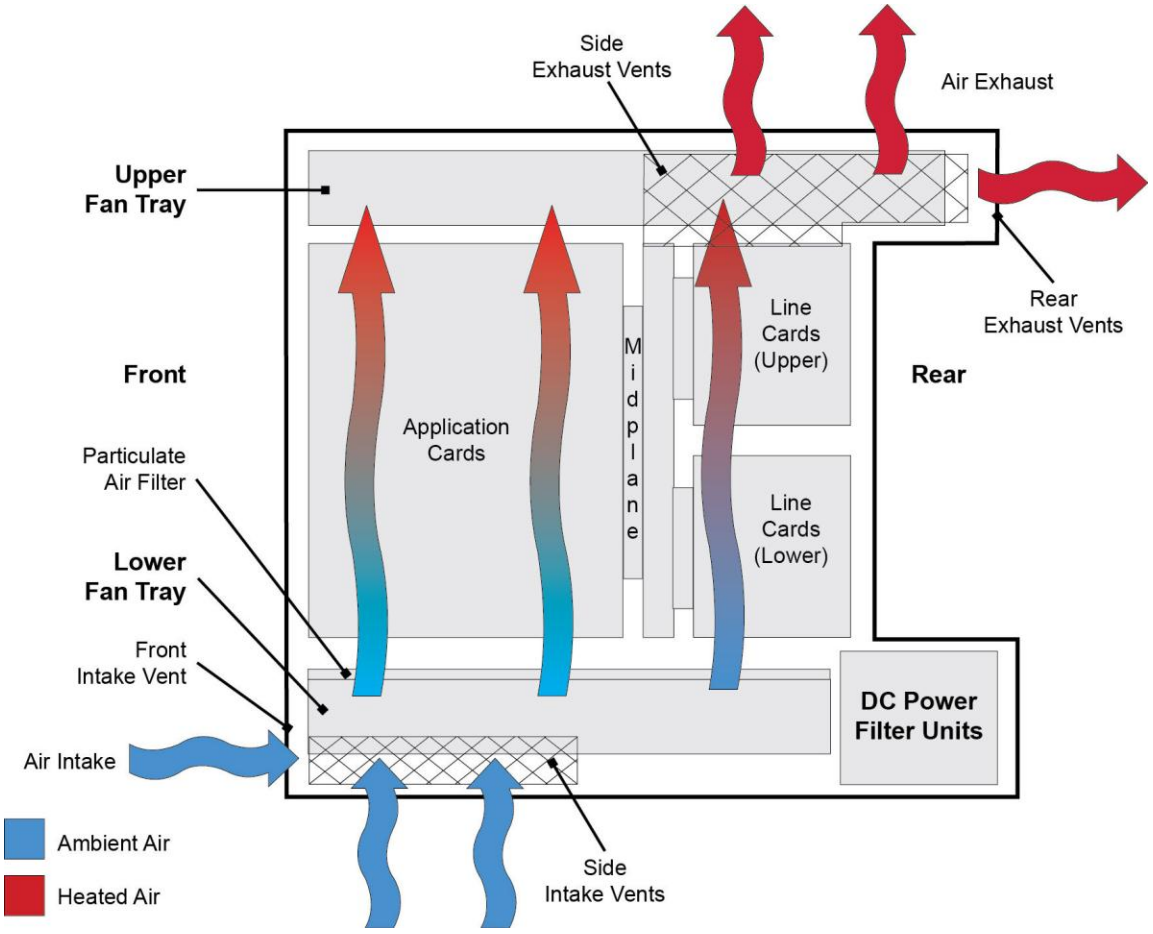
Chassis Air Flow

Airflow within the ASR 5000 is designed per Telcordia recommendations to ensure the proper vertical convection cooling of the system.

As shown in the figure below, the lower fan tray pulls ambient air into the chassis from the front and side intake vents located at the bottom of the chassis. The air is pushed upwards through the system and absorbs heat while passing over system components. The airflow through the chassis, as measured by the speed of the airflow across the various card slots, maintains an average of 290 feet per minute (8.2 meters per minute).

The upper fan tray pulls the heated air up through the chassis. The heated air then exits through the side and rear exhaust vents located at the top of the chassis.

Figure 285. System Airflow and Ventilation



Caution: When planning chassis installation, ensure that equipment rack or cabinet hardware does not hinder air flow at any of the intake or exhaust vents. Additionally, ensure that the rack/cabinet hardware, as well as the ambient environment, allow the system to function within the operating limits specified in this chapter.

Chapter 37
Glossary of Terms

Numerics

1xEV-DO. See **EV-DO**.

1xEV-DV. The third phase of CDMA2000 following 1xEV-DO deployment. 1xEV-DV stands for 1x Evolution - Data Voice, and is characterized by a maximum data rate of 5.2 Mbps and the ability to support wireless Voice over IP (VoIP) services.

1xRTT. The first phase of CDMA2000, characterized by the ability to support a maximum data rate of 1.44 Kbps. 1xRTT stands for 1x, denoting the one radio channel of 1.25 MHz in Radio Transmission Technology.

2G. The second generation of wireless technology that was characterized by its use of digital transmissions rather than analog methods. Radio bandwidth is used for data transmissions. Data transmissions are limited to a maximum rate of 1.44 Kbps for CDMA 2G services. Radio bandwidth is consumed whenever the Mobile Node (MN) is connected to the Internet, regardless of whether it is receiving or transmitting data. This is based on the IS-95A standard for CDMA.

2.5G. An evolutionary step between 2G and 3G wireless services wherein two enhancements were introduced over 2G. The first is that the MN only consumes radio bandwidth when data is being transmitted or received. The second is that the maximum data rate increased to approximately 64 Kbps. Most 2.5G services only support data rates between 1.15 Kbps and 384 Kbps. This is based on the IS-95B standard for CDMA.

3G. The third generation of wireless technology, wherein data services are packetized, with speeds up to 2 Mbps. Based on the CDMA2000 standards.

3GPP. Third Generation Partnership Project. A group of organizational partners from ETSI, TIA/EIA, and other standardization bodies who are working together to define the evolution of GSM-based wireless communication core networks.

3GPP2. Third Generation Partnership Project 2. A second group of organizational partners from ETSI, TIA/EIA, and other standardization bodies who are working together to define the evolution of CDMA-based wireless communication networks

4G. The fourth generation of wireless technology provides mobile ultra-broadband Internet access. Based on the ITUR-R International Mobile Telecommunications Advanced (IMT-Advanced) specification, 4G has a target peak speed requirement of 100 Mbps for high mobility communication (such as from trains and cars) and 1 Gbps for low mobility communication (such as pedestrians and stationary users). A 4G system does not support traditional circuit-switched telephony service, only Internet Protocol (IP) based communication. See **LTE**.

A

A10. The subscriber data portion of the R-P interface (based on GRE as defined in RFC-2784 and IP Encapsulation Within IP as defined in RFC-2003). See **R-P interface**.

A11. The control portion of the R-P interface (based on Mobile IPv4 as defined in RFC-2002).

A11 Manager. A StarOS task that controls the signalling de-multiplexing tasks of the A11 interface used for wireless communications.

AA-A. Authentication/Authorization Answer.

AA-R. Authentication/Authorization Request.

AAA. Authentication, Authorization, and Accounting. The security and billing methodology used by operators to ensure a user's identity and to determine their network usage so that they are properly billed. Often interchanged with RADIUS and Diameter protocols.

AAA Manager. This is a StarOS task that performs all AAA protocol operations and functions for subscribers and context-level administrative users within the system.

AAL. ATM Adaptation Layer supports information transfer protocols that are not based on ATM. This adaptation layer defines how to segment and reassemble higher-layer packets into ATM cells, and how to handle various transmission aspects in the ATM layer.

Acceptable Cell. This is a cell that the MS may camp on to make emergency calls. It must satisfy criteria which are defined for A/Gb mode in 3GPP TS 43.022 and for Iu mode in 3GPP TS 25.304.

Access Technology. The access technology associated with a PLMN. The MS uses this information to determine what type of radio carrier to search for when attempting to select a specific PLMN (for example, GSM, UTRAN, GSM COMPACT or EUTRAN). A PLMN may support more than one access technology.

ACL. Access Control List is a filtering mechanism used by many access IP routers that controls which traffic may be received or transmitted on an interface or port.

ACO. Alarm Cut Off is a toggle switch used to temporarily disable a central office alarm that occurs on a specific network device.

ACR. Accounting Request Message. See **Diameter**.

ACR. Active Charging Record.

ACS. Active Charging System is an in-line service feature that enables operators to reduce billing-related costs and gives the ability to offer tiered, detailed, and itemized billing to their subscribers. Also known as ECS.

ADC. Application Detection and Control is an in-line service that runs on the ASR 5x00 chassis and detects Peer-to-Peer (P2P) protocols via traffic analysis.

Address resolution. The process of determining the link-layer address of a node whose network-layer address is known.

ADMF. Administrative Function manages orders for interception. See **CALEA**.

AE. Application Entity. See **LI**.

AF. Access Function. See **LI**.

AF. Application Function is an element offering applications that use IP bearer resources. The AF is capable of communicating with the CRF to transfer dynamic charging rules related service information. One example of an AF is the P-CSCF of the IM CN subsystem. See **PCC**.

AFI. Address Family Identifier carries the identity of the Network Layer protocol associated with the Network Address that follows. See **BGP** and **SAFI**.

Agent advertisement. The procedure by which a mobility agent becomes known to the mobile node.

Agent discovery. The process by which a mobile node can obtain the IP address of a Home Agent or Foreign Agent, depending upon whether the mobile node is home or away from home. Agent discovery occurs when a mobile node receives an agent advertisement, either as a result of periodic broadcast or in response to a solicitation.

Aggregate Maximum Bit Rate. The maximum bit rate that limits the aggregate bit rate of a set of non-GBR bearers of a UE. The label (EUTRAN only) indicates this subclause or paragraph applies only if EUTRAN is used as current radio access network.

AGW. Access Gateway.

AH. Authentication Header is part of the IPSec specification. Other IPSec header mechanisms include Diffie-Hellman, DES, 3DES, and others.

AKA. Authentication and Key Agreement. This mechanism performs authentication and session key distribution in networks. AKA is a challenge- response based mechanism that uses symmetric cryptography. AKA is typically run in a Services Identity Module.

Allowable PLMN. For an MS operation mode A or B, this is a PLMN which is not in the list of "forbidden PLMNs" in the MS. For an MS operating in MS operation mode C, this is a PLMN which is not in the list of forbidden PLMNs or in the list of "forbidden PLMNs for GPRS service" in the MS.

Allowed CSG List. A list of CSG IDs stored in the UE. A UE is able to access only those CSG cells that have a CSG ID in this list.

APN. Access Point Name is a logical name for a packet data network and/or a service to which the GGSN supports access.

APS. Automatic Protection Switching is a means of achieving network redundancy via automatic switching mechanisms that switch from a primary circuit to a pre-defined secondary circuit.

ARP. Address Resolution Protocol is a standard protocol for performing address resolution between IP addresses and various link-layer addresses.

ARQ. Automatic Repeat Request. A link layer may automatically retransmit packets that were not correctly received by the next hop link layer. This improves the robustness of the packet delivery, but comprises the latency and packet overhead.

ASBR. Autonomous System Border Router. An Inter-AS network deployment can be configured so that the ASBRs exchange IPv4/IPv6 routes and MPLS labels.

ASN.1 Abstract Syntax Notation One is a standard and flexible notation that describes rules and structures for representing, encoding, transmitting, and decoding data in telecommunications and computer networking. The formal rules enable representation of objects that are independent of machine-specific encoding techniques. See **OID**.

ASN-GW. Access Service Network Gateway processes subscriber control and bearer data traffic, and supports connection and mobility management across cell sites and inter-service provider network boundaries. An ASN-GW is a logical entity in the ASN of a WiMAX radio access network and interfaces directly with base transceiver station or base station via an R6 GRE reference interface. ASN-GW performs control plane functions, bearer plane routing or bridging functions, resident functions in the connectivity service network, or a function in another ASN.

AT. Access Terminal.

ATM. Asynchronous Transfer Mode is a connection-oriented data link layer protocol used in cell relay/packet switched networks.

Authorization Token. The authorization token consists of the AF session identifier as well as the PDF identifier. The AF session identifier is assigned by the P-CSCF on successful IMS session establishment. The authorization token is sent to the UE by P-CSCF as part of the session establishment. The UE passes the authorization token in the binding information to the AGW. AGW uses the authorization token to get the PDF to be communicated for policy authorization and the session identifier is used for the authorization request to indicate the session to which authorization event belongs.

Automatic Home Agent discovery. The process by which a mobile node can obtain the IP address of an HA on its home network, involving the transmission of a registration request to the subnet broadcast address of its home network.

AVP. Attribute-Value Pair corresponds to an Information Element in an AAA message.

B

BBERF. Bearer Binding and Event Reporting Function. See **HSGW** or **S-GW**.

BBU. Baseband Unit. See **eNodeB**.

BCE. Binding Cache Entry. See **PBU**.

BE. Best Effort.

BFD. Bidirectional Forwarding Detection is a network protocol used to detect faults between two forwarding engines connected by a link. BFD establishes a session between two endpoints over a particular link. If more than one link exists between two systems, multiple BFD sessions may be established to monitor each one of them.

BGP. Border Gateway Protocol is a routing protocol used for interdomain routing within large networks to maintain integrity of the network. It allows the routers to exchange only pre-specified information with pre-specified routers in other domains.

BHSA. Busy Hour Session Attempts. A measure of dynamic sessions (traffic calls) that can be attempted in an average Busy Hour.

BHSC. Busy Hour Session Completion. A measure of dynamic sessions (traffic calls) that can be completed in an average Busy Hour.

Binding. The triplet of numbers that contains the mobile node's home address, its care-of address, and the registration lifetime-how long the mobility agents may use the binding. Binding within StarOS creates the association of a virtual interface to a physical port on the system. This process allows the flow of traffic from the context through the physical port with which the interface is associated.

Binding Information. The binding information associates a PDP context to the IP flows of a media. The binding information is generated by the P-CSCF and sent to UE during the IMS session establishment. The system receives the binding information from the UE during PDP context activation or modification. The binding information consists of a single authorization token and one or more flow identifiers for the IMS session. See **IMS** and **SCM**.

Binding Mechanism. This mechanism is used to associate a PDP context bearer with the IP flow(s) of an IMS session in the PDF.

Binding update. The message that supplies a new binding to an entity that needs to know the new care-of address for a mobile node. The binding update contains the mobile node's home address, new care-of address, and a new registration lifetime.

BITS. Building Integrated Timing Supply is a method for distributing a precision clock among telecommunications equipment. The physical medium is a T1/E1 (DS1) line.

Blacklisting. Refers to the disabling of the use of a mobile device based on its IMEI appearing in an database of blacklisted handsets. See **EIR**.

BLOB. BLock Of Bits.

BRA. Binding Revocation Acknowledgement.

BRI. Binding Revocation Indication.

BS. Base Station is an entity in the public radio telecommunications system used for bidirectional radio communications with mobile stations or mobile nodes.

BSC. Base Station Controller is a significant device within the 2G/2.5G RAN, the BSC allocates channels and manages BTS handoff. In 2G wireless, the BSC's upstream interfaces (to the MSC) are always TDM. In 2.5G, a BSC supports both TDM and packet upstream interfaces. In 3G, a BSC can support any combination of TDM and packet, TDM only, or packet only interfaces.

BSCS. Business Support and Control System. See **OSS**.

BSS. Base Station Subsystem is the 2G/2.5G Radio Access Network (RAN) technology responsible for connecting the mobile User Equipment (UE) with the Core Network (CN) in a GPRS/UMTS wireless network. The BSS incorporates the BTS, the BSC, and the PCU.

BSS. Business Support System consists of components that a service provider uses to run its business operations towards the customer. See **OSS**.

BTS. Base Transceiver Station is a component of the base station that includes the transmitting and receiving radio equipment. A BTS is sometimes equated with the physical cell site of a wireless network.

Busy Hour. An uninterrupted 60-minute period during which the average volume of traffic is at its maximum.

C

Cached EPS security context. This is a cached security context used in ECS.

CAE. Content Adaptation Engine is an optional component of the Cisco Mobile Video Solution. It runs on the Cisco UCS (Unified Computing System) platform and functions in a UCS cluster to bring video storage and additional video optimization capabilities to the Mobile Video Solution.

CALEA. Communications Assistance for Law Enforcement Act is a United States wiretapping law (Pub. L. No. 103-414, 108 Stat. 4279, codified at 47 USC 1001-1010). CALEA's purpose is to enhance the ability of law enforcement and intelligence agencies to conduct electronic surveillance by requiring that telecommunications carriers and manufacturers of telecommunications equipment modify and design their equipment, facilities, and services to ensure that they have built-in surveillance capabilities, allowing federal agencies to monitor all telephone, broadband internet, and VoIP traffic in real-time. See **LI**.

CAMEL. Customized Applications for Mobile networks Enhanced Logic. CAMEL service enables operators of 2.5G/3G networks to provide operator-specific services (such as prepaid GPRS service and prepaid SMS service) to subscribers, even when the subscribers are roaming outside their home public land mobile network (HPLMN). See **Ge interface**.

Camped on a cell. The MS (ME if there is no SIM) has completed the cell selection/reselection process and has chosen a cell from which it plans to receive all available services. Note that the services may be limited, and that the PLMN may not be aware of the existence of the MS (ME) within the chosen cell.

CAP protocol. CAMEL Application Part protocol is a signalling protocol used in the Intelligent Network (IN) architecture. It is a Remote Operations Service Element (ROSE) user protocol that is layered on top of TCAP in the SS7 protocol suite. CAP is based on a subset of the ETSI Core and allows for the implementation of carrier-grade, value added services such as unified messaging, prepaid and fraud control in both the GSM voice and GPRS data networks.

Care-of address. An IP address at the mobile node's current point of attachment to the Internet, when the mobile node is not attached to the home network. A collocated care-of address is a care-of address assigned to one of the mobile node's network interfaces, instead of one being offered by a foreign agent.

CC. Contents of Communication. See **LI**.

CCA. CC-Answer. See **AVP** and **Diameter**.

CCC. Call Content Channel. See **CALEA**.

CCP. Compression Configuration Protocol.

CCR. CC-Request. See **AVP** and **Diameter**.

CDF. Charging Data Function addresses are used for offline billing and Online Charging Function (OCF) for online billing. See **GTP'**.

CDMA. Code Division Multiple Access is one of three wireless technology classes that encompasses 2G, 2.5G, and 3G communications. The other two are GSM and TDMA.

cdmaOne. Defines the 2G and 2.5G versions of CDMA technology. Based on IS-95A and IS-95B standards respectively.

CDMA2000. Defines the 3G version of CDMA technology.

CDR. Charging Data Record (also known as call detail record) is a GTPP-based subscriber accounting record that consists of formatted information that includes event-based billing information such as call duration. Different systems generate different types of CDRs. The types, content and handling of CDRs is defined in various 3GPP specs within the TS 32.2xx series,

CE. Customer Edge router is an MPLS VPN router that provides customer peering service. See **PE**.

Cell. The unit of a base station having the ability to radiate in a given geographic area; a "sector" or "face" of a physical radio equipment implementation.

CEPT. European Conference of Postal and Telecommunications Administrations is the coordinating body for European state telecommunications and postal organizations.

CF. Collection Function. See **LI**.

CF. Content Filtering is an in-line service available for 3GPP and 3GPP2 networks to filter HTTP and WAP requests from mobile subscribers based on the URLs in the requests. This enables operators to filter and control the content that an individual subscriber can access, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences. CF runs on an external server under Linux OS (RHEL).

CFE. Common Firmware Environment refers the control processor-based firmware running on hardware components within the ASR 5x00.

CG. Charging Gateway is the device on the GSM GPRS or UMTS network that collects and maintains Call Detail Records (CDRs) for subscriber PDP contexts. Also referred to as a Charging Gateway Function (CGF). See **Ga interface**.

CGF. Charging Gateway Function. See **CG**.

CGI. Common Gateway Interface is a standard (see RFC 3875: CGI Version 1.1) method for web server software to delegate the generation of web content to executable files. Such files are known as CGI scripts; they are programs, often stand-alone applications, usually written in a scripting language.

Charging Rule. A set of information including the service data flow filters (IP 5-tuple), the gating status (pass/drop packets matching the rule) and the rating group, for a single service data flow. For an IMS media component a charging rule typically defines a single IP flow associated to a media component (for example, RTP or RTCP).

CLC2. Channelized Line Card (OC-3/STM-1, Frame Relay, ASR 5000). The CLC2 provides frame relay over SONET or SDH. It supports network connectivity through a gigabit interface to connect to the Packet Control Unit (PCU) of the base station subsystem (BSS) in a mobile network. These interfaces are commonly used with the SGSN product to support frame relay. Each CLC2 provides four optical fiber physical interfaces (ports).

CLCI Client. Common Language Circuit Identification DCCA client located in a GGSN.

CLCI Server. Common Language Circuit Identification DCCA server typically located in the Online Charging System.

CLI. Command Line Interface is a Man-machine Interface (MMI) used to configure, monitor, and administer a network device through its Operating System (OS). See **StarOS**.

CMIP. Client Mobile IP is an MIP stack that allows mobile devices to interact with Mobility Agents in the home and foreign network. CMIP is the opposite of PMIP.

CMISE. Common Management Information Service is the service interface specified in ITU-T Recommendation X.710, ISO/IEC International Standard 9595 that is employed by OSI network elements for network management. It defines the service interface that is implemented by the Common Management Information Protocol (CMIP) as specified in ITU-T Recommendation X.711, ISO/IEC International Standard 9596-1. See **SNMP**.

CN. Core Network.

CO. Central Office. The telecommunications facility where calls are switched.

Context. A specific group of StarOS configuration parameters that apply to the ports, interfaces, protocols, and services supported by a system. Each system can support multiple contexts and each context can reside as a separate, logically independent instance. Multiple context support allows numerous like or disparate services to exist on the same physical hardware.

CORBA. Common Object Request Broker Architecture is the Object Management Group's (MAG's) core specification for distributed object interoperability. CORBA enables separate pieces of software written in different languages and running on different computers to work with each other like a single application or set of services. More specifically,

CORBA is a mechanism in software for normalizing the method-call semantics between application objects residing either in the same address space (application) or remote address space (same host, or remote host on a network). See **WEM**.

Correspondent node. A node that sends or receives a packet to an MN; the correspondent node may be another mobile node or a non-mobile Internet node.

CP. Control Processor, a high-speed state-of-the-art CPU used by the system.

CPE. Customer Premises Equipment is any terminal and associated equipment located at a subscriber's premises and connected with a carrier's telecommunication channel(s) at the demarcation point (demarc). The demarc is a point established in a building or complex to separate customer equipment from the equipment located in either the distribution infrastructure or central office of the Communications Service Provider.

CRF. Charging Rule Function. See **PCRF**.

CS. Circuit Switched is a methodology of implementing a telecommunications network in which two network nodes establish a dedicated communications channel (circuit) through the network before the nodes may communicate. The circuit guarantees the full bandwidth of the channel and remains connected for the duration of the communication session. The circuit functions as if the nodes were physically connected as with an electrical circuit.

CSCF. Call Session Control Function consists of a collection of SIP servers or proxies that are used to process SIP signalling packets in IMS. See **E-CSCF** and **P-CSCF**.

CSFB. Circuit Switched Fallback. In MME deployments CSFB enables the UE to camp on an EUTRAN cell and originate or terminate voice calls through a forced switchover to the circuit switched (CS) domain or other CS-domain services.

CSG. Closed Subscriber Group identifies subscribers of an operator who are permitted to access one or more cells of the PLMN but which have restricted access (CSG cells).

CSG Cell. A CSG cell identifies a specific CSG that is accessible by the members of the closed subscriber group for that CSG. All CSG cells sharing the same CSG identity use the same radio access technology.

CSG ID. A CSG ID is an identifier associated with a cell or group of cells with access restricted to a defined group of users.

CSP. Card Slot Port is a StarOS subsystem that manages all cards, slots, and physical ports installed in a system.

CTF. Charging Trigger Function receives charging events from the CTF via the Rf reference point. It then uses the information contained in the charging events to construct CDRs. See **OFCS**.

Current EPS security context. the EPS security context which has been taken into use by the network most recently.

Current serving cell. This is the cell on which the MS is camped.

D

Data Radio Bearer. Data Radio Bearer transports the packets of an E-RAB between a UE and an eNodeB. There is an one-to-one mapping between the E-RAB and the Data Radio Bearer.

DCCA. Diameter Credit Control Application is a networking protocol for implementing real-time credit-control for a variety of end user services (IETF RFC 4006).

Dedicated Bearer. An EPS bearer that is associated with uplink packet filters in the UE and downlink packet filters in the P-GW where the filters only match certain packets.

Dedicated PDP Context. A PDP context with associated TFT filters, this may be a secondary or a primary PDP context (updated after its activation). There can be several such PDP contexts for a UE IP address.

Default APN. A Default APN is defined as the APN which is marked as default in the subscription data and used during the Attach procedure for PDN connection.

Default Bearer. The EPS bearer which is first established for a new PDN connection and remains established throughout the lifetime of the PDN connection.

Destination Context. The virtual context, or location, where a particular service configuration resides that mobile subscriber is directed to upon successful authentication through the system.

DF. Delivery Function. See **LI**.

DHCP. Dynamic Host Configuration Protocol is a protocol by which a host obtains from a server certain information it needs to communicate, such as an IP address, prefix length, and Domain Name System (DNS) server address.

Diameter. A second-generation AAA protocol. Diameter Applications extend the base protocol by adding new commands and/or attributes, such as those for use of EAP.

DiffServ. See **DSCP Marking**.

DL. Downlink is the transmission path from a cell site to the cell phone. See **UL**.

DNS. Domain Naming System. A system within the network that maps host-names into IP addresses.

DoS. Denial-of-Service is an attempt to make a machine or network resource unavailable to its intended users. A DoS attack generally consists of the efforts of one or more people to temporarily or indefinitely interrupt or suspend services of a host connected to the Internet.

DPC. Data Processing Card (ASR 5500). DPCs manage subscriber sessions and control traffic.

DPCA. Diameter Policy Control Application. See **PCRF**.

DPD. Dead Peer Detection, also known as Keepalive, is a timer that starts after the last IKE_AUTH message is sent to the MS and resets when traffic is received from the MS. If no valid messages are received when the timer expires the session is disconnected.

DPI. Deep Packet Inspection is a form of computer network packet filtering that examines data within a packet as it passes an inspection point. DPI searches for protocol non-compliance (viruses, spam, intrusions) or defined criteria to decide whether: (a) the packet may pass or if it needs to be routed to a different destination; (b) to collect statistical information.

DSCP Marking. DiffServ Code Point (IP Differentiated Services). DiffServ is a computer networking architecture that specifies a simple, scalable and coarse-grained mechanism for classifying and managing network traffic and providing quality of service (QoS) on modern IP networks. DiffServ uses the 6-bit Differentiated Services Field (DS field) in the IP header for packet classification purposes. The DS field and ECN field replace the outdated IPv4 TOS field.

Dynamic Charging Rule. Charging rule where some or all of the data within the charging rule (example, service data flow filter information) is assigned via real-time analysis using for example dynamic application derived criteria. An example of a dynamic charging rule is a rule determined by the E-PDF by means of real-time SDP derived information analysis.

E

E1. An ITU-T compliant E1 link operates at a line data rate of 2.048 Mbps, full duplex, which is split into 32 timeslots of 8 bits. Time slot 0 is devoted to transmission management and time slot 16 for signaling; the rest are available for voice/data transport. See **T1**.

E-CSCF. Emergency-Call Session Control Function handles certain aspects of emergency sessions, for example, routing of emergency requests to the correct emergency center or Public Safety Answering Point (PSAP). See **IMS**.

E-PDF. Enhanced Policy Decision Function. See **PCRF**.

E-RAB. Evolved Radio Access Bearer uniquely identifies the concatenation of an S1 Bearer and the corresponding Data Radio Bearer. When an E-RAB exists, there is a one-to-one mapping between this E-RAB and an EPS bearer of the Non Access Stratum.

E-RAB identity. Uniquely identifies an E-RAB for one UE. Note. The E-RAB identity remains unique for the UE even if the UE-associated logical S1-connection is released during periods of user inactivity.

eAN/ePCF. Evolved Access Network/Evolved Packet Control Function.

EAP. Extensible Authentication Protocol is an authentication protocol which provides an infrastructure that enables clients to authenticate with a central authentication server.

EAP-AKA. An extension to the EAP enabling authentication and session key distribution using the UMTS AKA (Authentication and Key Agreement) mechanism.

ECM. EPS Connection Management.

ECMP. Equal Cost Multiple Path routing distributes traffic across multiple routes that have the same cost to lessen the burden on any one route.

ECN. Explicit Congestion Notification is an extension to the Internet Protocol and to TCP (RFC 3168). ECN allows end-to-end notification of network congestion without dropping packets. ECN is an optional feature that is only used when both endpoints support it and are willing to use it.

ECS. Enhanced Charging Services (also known as Active Charging Service [ACS]). This is an in-line service feature that enables operators to reduce billing-related costs and offer tiered, detailed, and itemized billing to their subscribers. Using shallow and deep packet inspection (DPI), ECS allows operators to charge subscribers based on actual usage, number of bytes, premium services, location, and so on. ECS also generates CDRs for postpaid and prepaid billing systems.

EDGE. Enhanced Data rates for GSM Evolution is a digital mobile phone technology that allows improved data transmission rates as a backward-compatible extension of GSM.

EDR. Event Data Records are generated by the ECS and sent to an external billing system. An EDR contains information about different flow types (HTTP, SMTP, MMS) and are used for content billing purposes. EDR content is configurable by an operator.

EHPLMN. Equivalent Home PLMN. Any of the PLMN entries contained in the Equivalent HPLMN list. To allow provision for multiple HPLMN codes, PLMN codes that are present within this list shall replace the HPLMN code derived from the IMSI for PLMN selection purposes. This list is stored on the USIM and is known as the EHPLMN list. The EHPLMN list may also contain the HPLMN code derived from the IMSI. If the HPLMN code derived from the IMSI is not present in the EHPLMN list then it shall be treated as a Visited PLMN for PLMN selection purposes.

eHRPD. Evolved High Rate Packet Data (3GPP2).

EIR. Equipment Identity Register is a security database that enables network operators to track mobile phones in a wireless network and to disable stolen equipment. See **IMEI** and **S13 interface**.

EMACS. This is a standard UNIX text editor that employs macros to manipulate command lines in the CLI.

EMM. EPS Mobility Manager. Supports UE mobility, such as informing the network of its present location and providing user identity confidentiality. It also provides connection management services to the session management (SM) sublayer. See **EPS**.

EMM context. An EMM context is established in the UE and the MME when an attach procedure is successfully completed.

EMM-CONNECTED mode. A UE is in EMM-CONNECTED mode when a NAS signalling connection between UE and network is established. The term EMM-CONNECTED mode used in the present document corresponds to the term ECM-CONNECTED state used in 3GPP TS 23.401.

EMM-IDLE mode. A UE is in EMM-IDLE mode when no NAS signalling connection between UE and network exists.

EMS. Element Management System defines the system or application used to manage a network device, or groups of like network devices. See **WEM**.

Encapsulation. The process of incorporating an original IP packet (less any preceding fields such as a MAC header) inside another IP packet, making the fields within the original IP header temporarily lose their effect.

eNodeB. EUTRAN Node B or Evolved Node B is the hardware that communicates directly with mobile handsets (UEs), like a base transceiver station (BTS) in GSM networks.

EPC. Evolved Packet Core. The successor to the 3GPP Release 7 packet-switched core network, developed by 3GPP within the framework of the 3GPP System Architecture Evolution (SAE).

ePDG. Evolved Packet Data Gateway. The ePDG allows interworking between the EPC and fixed non-3GPP networks that require secure access, such as Wireless LAN. The ePDG provides a secure access to EPC and supports what LTE standards refers to as un-trusted non-3GPP access through the SWu and S2b or S2c interfaces.

EPS. Evolved Packet System or evolved 3GPP packet-switched domain consists of the evolved packet core network and EUTRAN.

ESD. Electrostatic discharge is the sudden flow of electricity between two objects caused by contact, an electrical short, or dielectric breakdown. ESD can be caused by a buildup of static electricity or by electrostatic induction.

ESM. EPS Session Management.

ESN. Electronic Serial Number is a unique 32-bit binary number that identifies each cellular device. This information is passed as part of the call setup.

ESS. External Storage System. Collects, stores, and reports billing information from the ECS running on the ASR 5000 chassis. The ESS can simultaneously fetch any type of files from one or more chassis, including CDR, EDR, NBR, and UDR files.

ETSI. European Telecommunications Standards Institute (ETSI) is an independent, non-profit, standardization organization in the telecommunications industry (equipment makers and network operators) in Europe, with worldwide projection.

EUTRAN. Enhanced UMTS Terrestrial Radio Access Network (3GPP). This is the air interface of the 3GPP Long Term Evolution (LTE) upgrade path for mobile networks. It replaces UMTS, HSDPA and HSUPA technologies specified in 3GPP releases 5 and beyond. EUTRAN is incompatible with W-CDMA. It provides higher data rates, lower latency and is optimized for packet data.

EV-DO. The second phase of CDMA2000 following 1xRTT deployment. 1xEV-DO stands for 1x Evolution - Data Only, and is characterized by a maximum data rate of 2.4 Mbps.

eWAG. Enhanced Wireless Access Gateway enables Wi-Fi integration into 3G mobile packet core (MPC), allowing clientless UE attaches to trusted Wireless Local Area Networks (WLANs) for access to 3G services. The UE does not require a client, has no dependencies on the Wi-Fi architecture, and does not realize that it is connecting to a 3G network.

F

FA. Foreign Agent is a mobility agent on the foreign network that can assist the mobile node in receiving datagrams delivered to the care-of address.

FDMA. Frequency Division Multiple Access is a method of allocating a discrete amount of frequency bandwidth to individual users to allow multiple conversations across many users.

FDR. Flow Data Record. See **EDR**.

FEC. Forward Error Correction. The physical link layer may add many extra bits to the data before transmitting it. The receiving physical link layer uses those bits to automatically correct errors in the received data, without needing the data to be retransmitted. The transmitter and receiver must use the same FEC algorithm.

FEC. Forwarding Equivalency Class is a term used in MPLS to describe a set of packets with similar and/or identical characteristics which may be forwarded the same way; that is, they may be bound to the same MPLS label. An FEC tends to correspond to a label switched path (LSP). See **FTN** and **NHLFE**.

FELC/FLC2. Fast Ethernet Line Card (Ethernet 10/100, ASR 5000).

femtocell. This is a small, low-power cellular base station, typically designed for use in a home or small business. It connects to the service provider's network via broadband (such as DSL or cable). See **HNB**.

Firewall. A device that protects a private network against intrusion from nodes that are using the public network.

FITS. Failure in Time Statistics is a statistical method of determining the number of failures that are expected to occur over a specific time period. The telecommunications industry generally assumes this number to represent the number of failures per million hours (Fpmh).

Flow Identifier. An IP flow is indicated uniquely in an IMS session by means of a flow identifier. The flow identifier is created based on the ordinal number of the media stream and of the IP flow in the media where the IP flows are arranged based on the ports used.

FNG. Femto Network Gateway enables mobile operators to provide 3G network services to subscribers with wireless handsets via Femtocell Access Points (FAPs). The FNG makes it possible for operators to provide secure access to the operator's 3G network from a non-secure network, extend wireless service coverage indoors, especially where access would otherwise be limited or unavailable, reduce the load on the macro wireless network, and make use of existing backhaul infrastructure to reduce the cost of carrying wireless calls. See **HNB**.

Foreign network. The network to which the mobile node is attached when it is not attached to its home network, and on which the care-of-address is reachable from the rest of the Internet.

Forward Tunnel. The direction of encapsulated data traveling from the Home Agent to the Foreign Agent.

FQDN. Fully Qualified Domain Name is a complete domain name as defined by the Domain Name System (DNS). A node can be known locally by a relative domain name that is a sub-string of its FQDN, but such a relative name cannot be resolved correctly by Internet nodes outside of the part of the domain name hierarchy indicated by the relative name. The FQDN can be resolved from anywhere in the Internet, subject to access control and ability to route of the resolution request.

Frame Relay. This is a standardized wide area network technology that specifies the physical and logical link layers of digital telecommunications channels for a variety of network interfaces. It can run on fractional T1/EI, or full T-carrier or E-carrier systems. Frame Relay complements and provides a mid-range service between basic rate ISDN and Asynchronous Transfer Mode (ATM).

Frequency layer. Set of cells with the same carrier frequency.

FSC. Fabric Storage Card (ASR 5500). The FSC sources fabric cross-bar connections to each MIO and DPC. It also includes two serial attached SCSI (SAS) solid state drives (SSDs) with a 6 Gbps SAS connection to each MIO.

FTN. FEC-To-NHLFE is an MPLS table structure. See **FEC**, **ILM** and **NHLFE**.

FTP. File Transfer Protocol is a standard network protocol used to transfer files from one host or to another host over a TCP-based network, such as the Internet.

G

G-CDR. GGSN charging data record.

Ga interface. The interface between GGSN or SGSN and the charging gateway (CG). It uses GTPP to communicate.

Gb interface. The interface between the SGSN and the 2G/2.5G RAN base station subsystem – usually the connection between the BSS and the PCU.

GBR bearer. Guaranteed Bit Rate Bearer uses dedicated network resources related to a GBR value, which are permanently allocated during EPS bearer establishment/modification.

Gc interface. The interface used by the GGSN to communicate with the HLR via a GTP-to-MAP (Mobile Application Part) protocol convertor.

Ge interface. The GPRS interface between the SGSN and the service control point (SCP). See **CAP protocol**.

GELC/GLC2. Gigabit Ethernet Line Card (Ethernet 100, ASR 5000). The GELC/GLC2 supports a variety of 1000 Mbps optical and copper interfaces based on the type of Small Form-factor Pluggable (SFP) modules installed on the card.

General Purpose PDP Context. A PDP context without associated TFT filters where all the traffic is allowed, including internet traffic. This may be a primary or a secondary PDP context. However, only one PDP context without associated TFT filters can exist.

GERAN. GSM/EDGE Radio Access Network.

Gf interface. The SS7 interface between the SGSN and an EIR.

GGSN. Gateway GPRS Support Node is a device in a GSM GPRS/UMTS data network that performs data session establishment, accounting, and traffic routing.

Gi interface. The interface used by the GGSN to communicate with Packet Data Networks (PDNs) external to the PLMN.

GMT. Greenwich Mean Time. See **UTC**.

Gn interface. The interface used between two GSN (GGSN and/or SGSN) in the same GPRS/UMTS Public Land Mobile Network (PLMN). This interface serves as both the signalling and data path for establishing and maintaining subscriber PDP contexts.

Go interface. The interface used by the GGSN to communicate with Policy Decision function (PDF) for provisioning of policy for a PDP context bearer used for IMS session media flow transport.

Gp interface. The IP-based interface used between a GGSN and a GPRS support nodes (GSNs, for example, GGSNs and/or SGSNs) in a different PLMNs.

GPRS. General Packet Radio Service is the GSM version of 2.5G wireless data communications.

GPS. Global Positioning System (is a space-based satellite navigation system that provides location and time information in all weather, anywhere on or near the Earth, where there is an unobstructed line of sight to four or more GPS satellites. See **Stratum clocking**.

GR. Generic Requirements provide the Telcordia view of proposed generic criteria for telecommunications equipment, systems, or services considering a wide variety of factors, including interoperability, network integrity, participating-client expressed needs, and other inputs.

Gr interface. The SS7 interface between the SGSN and an HLR.

GRE. Generic Routing Encapsulation is used to tunnel data between various networks (RFC 2784). This protocol is mandated for use in R-P and Mobile IP communications.

Gs interface. The SS7 interface between the SGSN and an MSC/VLR.

GSM. Global System for Mobile communications is one of three wireless technology classes that encompasses 2G, 2.5G, and 3G communications. The other two are CDMA and TDMA.

GSN. GPRS Support Node can be either an SGSN or a GGSN.

GSS. GTPP Storage Server. An external backup/storage server for one or more types of CDRs: eG-CDRs, G-CDRs, M-CDRs, S-CDRs, and/or SMS CDRs.

GT. Global Title is a unique SCCP address (such as a mobile phone number) used to identify a destination. A global title does not include routing information.

GTP. GPRS Tunneling Protocol is used between the GGSN and the SGSN.

GTP' or GTP-P. GPRS Tunneling Protocol-Prime uses the same message structure as GTP (GTP-C, GTP-U), but is largely a separate protocol. GTP' uses registered UDP/TCP port 3386. GTP' can be used for carrying charging data from the Charging Data Function (CDF) of the GSM or UMTS network to the Charging Gateway Function (CGF).

GTP-C. GPRS Tunneling Protocol for the control plane handles signalling between GSNs within the core network.

GTP-U. GTP User. Supports user data plane signalling that handles user data moving between the RAN and the Core Network (CN), and within the CN.

GTT. Global Title Translation is the SS7 mechanism that provides translation of the destination global titles to enable message routing to the appropriate end-point.

GUMMEI. Globally Unique MME Identifier consists of a PLMN Identity, an MME Group Identity and an MME Code. An MME logical node may be associated with one or more GUMMEI, but each GUMMEI uniquely identifies an MME logical node. (3GPP TS 23.003)

GUTI. Globally Unique Temporary Identifier is an identifier for the UE allocated by MME so that IMSI is not required to be transmitted over the radio interface (where it is susceptible to interception). The structure of GUTI is designed such that it is possible to identify the MME that allocated it: GUTI = GUMMEI + M-TMSI.

Gx interface. The interface used by the GGSN to communicate with Charging Rule Function (CRF). Gx interacts between GGSN, the TPF (Traffic Plane Function) and the CRF. It is based on the Diameter base protocol and DCCA standard. The GGSN acts as the client where as the CRF contains the Diameter server functionality.

H

HA. Home Agent is a node on the home network that effectively causes the mobile node to be reachable at its home address even when the mobile node is not attached to its home network.

Handoff. The process by which an air interface circuit between a mobile node and the network, including all signalling and transfer of user information.

Handover. Procedure that changes the serving cell of a UE in RRC_CONNECTED.

HAT. High Availability Task is a StarOS task that manages the operational state of the system.

HI. Handover Interface. See **LI**.

HLR. Home Location Register stores access service parameter information for users belonging to the particular GSM home network.

HNB. A Home Node B is the 3GPP term for a 3G femtocell. It is an element of a 3G macro RAN. A femtocell performs many of the function of a Node B, but is optimized for deployment in the home.

HNB-GW. Home Evolved NodeB Gateway is a gateway for Home NodeBs (HNBs) to access the core networks. The HNB-GW concentrates connections from a large number of HNBs through an IuH interface and terminates the connection to existing Core Networks (CS or PS) using an IuCS or IuPS interface.

Home address. The IP address assigned to the mobile node, making it logically appear attached to its home network.

Home network. The network at which the mobile node seems reachable, to the rest of the Internet, by virtue of its assigned IP address.

Home PLMN. This is a PLMN where the MCC and MNC of the PLMN identity match the MCC and MNC of the IMSI.

HRPD. High Rate Packet Data. See **EV-DO**.

HRPD Access. Combination of the eAN - PCF of the cdma2000 access.

HSDPA. High-Speed Downlink Packet Access is a 3G mobile telephony communications protocol in the High-Speed Packet Access (HSPA) family. It allows UMTS networks to have higher data transfer speeds and capacity.

HSGW. HRPD Serving Gateway terminates the HRPD access network interface from the Evolved Access Network/Evolved Packet Core Function (eAN/ePCF) and routes UE-originated or terminated packet data traffic. The HSGW supports interworking of the AT with the 3GPP EPS architecture and protocols specified in 3GPP 23.402.

HSPA. High Speed Packet Access is an amalgamation of two mobile telephony protocols, HSDPA and HSUPA, that extends and improves the performance of existing 3rd generation mobile telecommunication networks utilizing the WCDMA protocols.

HSS. Home Subscriber Service includes the master user database and software that support IMS network entities that handle calls and sessions. HSS contains user profiles, performs authentication and authorization of the user, and can provide information about the physical location of user. HSS is similar to the GSM Home Location Register (HLR).

HSUPA. High-Speed Uplink Packet Access is a 3G mobile telephony protocol in the HSPA family with up-link speeds up to 5.76 Mbps.

HTTP. Hypertext Transfer Protocol is an application protocol for distributed, collaborative, hypermedia information systems. It is the foundation of data communication for the World Wide Web.

I

IA. Intercepted Area. See **LI**.

ICMP. Internet Control Message Protocol (ICMP) is one of the core protocols of the Internet Protocol Suite. It is chiefly used by the operating systems of networked computers to send error messages indicating, for example, that a requested service is not available or that a host or router could not be reached. ICMP can also be used to relay query messages.

ICS. Internet Connection Sharing is the use of a device with Internet access such as 3G cellular service, broadband via Ethernet, or other Internet gateway as an access point for other devices. It makes use of DHCP and network address translation (NAT).

ICSR. Interchassis Session Recovery is a redundancy configuration that employs two identically configured ASR 5x00 chassis as a redundant pair. ICSR chassis share the same chassis key. If ICSR detects that the two chassis have incompatible chassis keys, an error message is logged but the ICSR system will continue to run. Without the matching chassis key, the standby ICSR chassis can recover services if the active chassis goes out of service; the standby chassis will still have access to the passwords in their decrypted form. See **SRP**.

IDL. Interface Definition Language refers to the application programming interface used to develop CORBA-based management interfaces as defined by the Object Management Group (OMG).

IE. Information Element.

IEEE. Institute of Electrical and Electronics Engineers is a professional association directed toward the advancement of the theory and practice of Electrical, Electronics, Communications and Computer Engineering, as well as Computer Science, the allied branches of engineering and the related arts and sciences.

IETF. Internet Engineering Task Force develops and promotes Internet standards, cooperating closely with the W3C and ISO/IEC standards bodies and dealing in particular with standards of the Internet protocol suite (TCP/IP). It is an open standards organization, with no formal membership or membership requirements.

IKE. Internet Key Exchange is an IPSec mechanism that is used to create SAs (Security Associations) between two entities in an IP-based Virtual Private Network (VPN).

ILM. Incoming Label Map is an MPLS table structure. See **FTN** and **NHLFE**.

IMEI. International Mobile Equipment Identity is a number that uniquely 3GPP mobile devices. The IMEI is used by a GSM network to identify valid devices and therefore can be used for stopping a stolen phone from accessing that network via an EIR. The IMEI is only used for identifying the device and has no permanent or semi-permanent relation to the subscriber. Instead, the subscriber is identified by transmission of an IMSI number, which is stored on a SIM card that can be transferred to any handset.

IMEI-SV. International Mobile Equipment Identity – Software Version.

IMS. IP Multimedia Subsystem provides wide application support for transport of voice, video, and data independent of the access support. IMS specifies a standard architecture for providing combined IP services (voice, data, multimedia) over the existing public switched domain.

IMSA. IP Multimedia Subsystem Authorization. For 3GPP networks this service requires specific support for a roaming IMS subscriber. Apart from other functionality sufficient, uninterrupted, consistent, and seamless user experience is

required to particular subscriber session for an application. It is also important that the subscriber gets charged only for the amount of resources consumed by the particular IMS application used.

IMSI. International Mobile Subscriber Identity uniquely identifies a subscriber to a mobile telephone service. It is a 50-bit field in GSM that identifies a mobile device's home country and carrier.

IN. Intelligent Network is the standard network architecture specified in the ITU-T Q.1200 series recommendations. It is intended for fixed as well as mobile telecom networks and allows operators to differentiate themselves by providing value-added services in addition to the standard telecom services such as PSTN, ISDN and GSM services on mobile phones. IN is supported by the SS7 protocol between telephone network switching centers and other network nodes owned by network operators.

Initial NAS message. A NAS message is considered as an initial NAS message, if this NAS message can trigger the establishment of a NAS signalling connection. For instance, the ATTACH REQUEST message is an initial NAS message.

Interface. As used in the context of system services, an interface is a virtual or logical assignment of a virtual router instance that provides higher-layer protocol transport. Interfaces are bound to physical ports within the system.

InTracer. InTracer is a high-performance subscriber troubleshooting and monitoring solution. It performs call tracing, control data acquisition, processing and analysis of both active and historical subscriber sessions. This provides a framework for operators to analyze and investigate call flows and call events for subscriber sessions in near real time.

IOT. Interoperability Testing.

IP. Internet Protocol is used for the transmission of packetized data. It is part of the TCP/IP suite of communications protocols.

IP in IP. Refers to the encapsulation of an inner IP header with an outer IP header for tunneling configuration.

IP-CAN. IP-Connectivity Access Network refers to 3GPP access networks such as GPRS or EDGE, but can be also used to describe wireless LAN (WLAN) or DSL networks. It was introduced in 3GPP IMS standards as a generic term referring to any kind of IP-based access network, as IMS put much emphasis on access and service network separation.

IPCF. Intelligent Policy Control Function runs under StarOS on ASR 5x00 chassis and is built around an intelligent rule configuration and execution system. Its policy rules engine is capable of acting on conditions such as the subscriber, the session state or network condition, or even time and day to decide on the corresponding treatment to be given to the subscriber. All information and rules fetched by querying the IPCF subscription plan stored in the SSC over an Sp interface. See **PCC**, **PPT** and **SSC**.

IPSec. IP Security is a multi-functional encryption technique used to transport packetized data in an un-readable fashion across multiple network devices.

IPSG. IP Services Gateway runs on the ASR 5x00 chassis and provides managed services to IP flows. The IPSG is situated on the network side of legacy, non-service capable GGSNs, PDSNs, HAs, and other subscriber management devices. It can provide per-subscriber services such as ECS, ADC, and others.

IPv4. Internet Protocol version 4 address consists of 32 bits divided into four octets. These four octets are written in decimal numbers, ranging from 0 to 255 that are concatenated as a character string with full stop delimiters (dots) between each number. For example, 192.165.34.55

IPv4v6 capability. Capability of the IP stack associated with a UE to support a dual stack configuration with both an IPv4 address and an IPv6 address allocated.

IPv6. Internet Protocol version 6 address has two logical parts: a 64-bit network prefix, and a 64-bit host address part. An IPv6 address is represented by eight groups of 16-bit hexadecimal values separated by colons (:). For example, 2001:0db8:85a3:0042:0000:8a2e:0370:7334.

IRI. Intercepted Related Information. See **LI**.

ISAKMP. Internet Security Association and Key Management Protocol. In IPSec negotiations, this protocol allows the receiver to obtain a public key and authenticate the sender using digital certificates.

ISDN. Integrated Services Digital Network is a set of communication standards for simultaneous digital transmission of voice, video, data, and other network services over the PSTN. ISDN provides access to packet switched networks by allowing digital transmission of voice and data over ordinary telephone copper wires. See **SS7**.

ISO. International Organization for Standardization is the world's largest developer of voluntary International Standards for products, services and good practice.

ISP. Internet Service Provider. A vendor, or telecommunications carrier, who provides Internet access services to customers.

ISR. Idle Mode Signalling Reduction is a feature that allows the UE to roam between LTE and 2G/3G. It aims at reducing the frequency of TAU and RAU procedures caused by UEs reselecting between EUTRAN and GERAN/UTRAN which are operated together.

ITU. International Telecommunication Union is the specialized agency of the United Nations which is responsible for information and communication technologies. ITU coordinates the shared global use of the radio spectrum, promotes international cooperation in assigning satellite orbits, works to improve telecommunication infrastructure in the developing world and establishes worldwide standards.

IuCS. The interface between the RNC and the Circuit Switched Core Network (CS-CN).

IuH. The interface between the femtocell (HNB) and the HNB-GW.

IuPS. The packet switched interface between the Radio Network Controller (RNC) in the UTRAN and a 3G SGSN. Supports both control plane and user data plane signalling, transmitting IP over ATM.

IWF. Inter-working Function is a device that is located between the MSC and the Internet, used to connect wireless subscribers to the Internet through 2G and 2.5G networks.

K

KPI. Key Performance Indicator is a type of performance measurement. KPIs are commonly used to measure the operational performance of a system based on such factors as availability, MTBF, MTTR and unplanned availability.

L

L2TP. Layer 2 Tunneling Protocol is a communications protocol used to establish tunnels between network devices to securely transport data.

LAC. L2TP Access Concentrator is a LAC connects an L2TP tunnel from a subscriber to a peer LNS.

LAC. Location Area Code identifies an area in a PLMN within which the MS/UE can move without the need of a location update to the VLR.

LAES. Lawfully Authorized Electronic Surveillance. See **LI**.

LAESP. LAES Protocol. See **LI**.

LAG. A Link Aggregation Group works by exchanging control packets via Link Aggregation Control Protocol (LACP) over configured physical ports with peers to reach agreement on an aggregation of links (IEEE 802.3ad). The LAG sends and receives the control packets directly on physical ports attached to different traffic bearing ports. Link aggregation (also called trunking or bonding) provides higher total bandwidth, auto-negotiation, and recovery by combining parallel network links between devices as a single link. A large file is guaranteed to be sent over one of the links, which removes the need to address out-of-order packets.

LAI. Location Area Identifier. Each location area of a PLMN has this internationally unique identifier used for location updating of mobile subscribers. It is composed of a three-decimal digit MCC, a two- to three-digit MNC that identifies the GSM PLMN in that country, and a LAC which is a 16-bit number thereby allowing 65536 location areas within one GSM PLMN.

LAN. Local Area Network denotes a group or groups of physically inter-connected network devices that are capable of sharing information with each other.

Last Visited Registered TAI. A TAI which is contained in the TAI list that the UE registered to the network and which identifies the tracking area last visited by the UE.

LAU. Location Area Update is a UMTS procedure that allows a mobile device to inform the cellular network whenever it moves from one location area to the next. Mobiles are responsible for detecting LACs. When a mobile finds that the LAC is different from its last update, it performs another update by sending a location update request to the network, together with its previous location and its TMSI. See **RAU** and **TAU**.

LBS. Location Based Services are a general class of computer program-level services used to include specific controls for location and time data as control features in computer programs.

LC. Line Card is a rear-installed card within an ASR 5000 chassis that provides physical network connectivity. Most LCs have physical external network interfaces.

LCP. Link Control Protocol forms part of the point-to-point protocol (PPP). In setting up PPP communications, both the sending and receiving devices send out LCP packets to determine the standards of the ensuing data transmission.

LDAP. Lightweight Directory Access Protocol is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol network.

LDP. Label Distribution Protocol is a protocol in which MPLS routers exchange label mapping information. Two routers with an established session are called LDP peers and the exchange of information is bidirectional. LDP is used to build and maintain LSP databases that are used to forward traffic through MPLS networks.

LEA. Law Enforcement Agency. See **LI**.

LER. Label Edge Router (also known as edge LSR) operates at the edge of an Multiprotocol Label Switching network. When forwarding IP datagrams into the MPLS domain, it uses routing information to determine appropriate labels to be affixed, labels the packet accordingly, and then forwards the labeled packets into the MPLS domain. Likewise, upon receiving a labeled packet which is destined to exit the MPLS domain, the LER strips off the label and forwards the resulting IP packet using normal IP forwarding rules. See **LSR** and **MPLS**.

LI. Lawful Intercept is an enhanced feature that enables Law Enforcement Agencies (LEAs) to intercept data sessions. The LI service gives the network operator the capability to intercept control and data messages of targeted mobile users. To invoke this support, the LEA will request the network operator to start the interception of a particular mobile user. This request will be supported by a court order or warrant. The ASR 5x00 supports a proprietary interface to either a third party Mediation Function (MF) or Delivery Function (DF).

Linked Bearer Identity. This identity indicates to which default bearer the additional bearer resource is linked.

LMA. Local Mobility Anchor (mobility server, HA-like, P-GW). See **PMIPv6**.

LNS. L2TP Network Server terminates an L2TP tunnel from a peer LAC and provides a network connection through the tunnel.

Logical Port. A subdivision of a physical port or interface within the system.

LR. Location Registration. A GPRS MS which is IMSI attached to GPRS services or to GPRS and non-GPRS services performs location registration by the Routing Area Update procedure only when in a network of network operation mode I. An MS which is IMSI attached to non-GPRS services only performs location registration by the Location Updating procedure. Both location updating and routing area update procedures are performed independently by the GPRS MS when it is IMSI attached to GPRS and non-GPRS services in a network of network operation mode II or III. See **TAU**.

LRSN. Local Record Sequence Number. The SGSN or GGSN includes this node-specific, unique sequential number in every partial or complete CDR.

LSA. Localized Service Area consists of a cell or a number of cells. The cells constituting a LSA may not necessarily provide contiguous coverage.

LSP. Label-Switched Path is a path through an MPLS network, set up by a signaling protocol such as LDP, RSVP-TE or BGP. The path is set up based on criteria in the forwarding equivalence class (FEC).

LSR. Label Switch Router performs routing based on MPLS labels. LSRs in an MPLS network regularly exchange label and reachability information with each other in order to build a complete picture of the network they can then use to forward packets. See **LER** and **MPLS**.

LTE. Long Term Evolution. Marketed as 4G LTE, is a standard for wireless communication of high-speed data for mobile phones and data terminals. It is based on the GSM/EDGE and UMTS/HSPA network technologies, increasing the capacity and speed using a different radio interface together with core network improvements. The standard is developed by the 3GPP (3rd Generation Partnership Project) and is specified in its Release 8 document series, with minor enhancements described in Release 9.

M

M-CDR. Mobility management CDR is generated by an SGSN.

M-TMSI. MME-TMSI identifies a UE served by an MME. See **TMSI**.

MAG. Mobile Access Gateway (mobility client, FA-like, HSGW, PMIP S-GW). The MAG function on the S-GW can maintain multiple PDN or APN connections for the same user session. The MAG runs a single node level Proxy Mobile IPv6 tunnel for all user sessions toward the LMA function of the P-GW.

MAP. Mobile Application Part is an SS7 protocol which provides an application layer for the various nodes in GSM and UMTS mobile core networks and GPRS core networks to communicate with each other in order to provide services to mobile phone users. MAP is the application-layer protocol used to access the HLR, VLR, MSC, EIR, Authentication Centre, SMS center and SGSN.

Mapped EPS security context. This is a mapped security context to be used in EPS.

MBMS. Multimedia Broadcast Multicast Services is a point-to-multipoint interface specification for existing and upcoming 3GPP cellular networks. It is designed to provide efficient delivery of broadcast and multicast services, both within a cell as well as within the core network.

MBMS-dedicated cell. A cell dedicated to MBMS transmission.

MBR. Maximum Bit Rate (QoS).

MCC. Mobile Country Code. See **PLMN**.

ME. Mobile Equipment.

MF. Mediation Function. See **LI**.

MGW. Media Gateway is a translation device or service that converts digital media streams between disparate telecommunications networks such as PSTN, SS7, Next Generation Networks (2G, 2.5G and 3G radio access networks) or PBX. Media gateways enable multimedia communications across Next Generation Networks over multiple transport protocols such as Asynchronous Transfer Mode (ATM) and Internet Protocol (IP).

MIB. Management Information Base is a virtual database used for managing the entities in a communications network. Most often associated with SNMP, the term is often used to refer to a particular subset, more correctly referred to as MIB-module. Objects in the MIB are defined using a subset of Abstract Syntax Notation One (ASN.1) [RFC 2578}. The software that performs the parsing is a MIB compiler. The database is hierarchical (tree-structured) and each entry is addressed through an object identifier (OID).

Minimal encapsulation. A variant encapsulation technique specified in RFC 2003 that temporarily alters the structure of the original IP header, but uses fewer bytes for tunneling packets to the care-of-address than the default method (IP-in-IP) uses.

MINE. Mobile and IP Network Enabler.

MIO. Management Input/Output card (ASR 5500). The MIO card performs chassis management, as well as local context (remote management) and non-local context (traffic bearing) external I/O operations.

MIP. Mobile IP is a protocol used to provide IP mobility to IPv4-based nodes (RFC-2002).

MME. Mobility Management Entity is an EPS element which manages mobility in EPC networks. The MME is the key control-node for the LTE access network. It works in conjunction with the eNodeB, S-GW within the EPC, or LTE/SAE core network to perform multiple functions.

MME area. An area containing tracking areas served by an MME. See **TAI list**.

MME Pool Area. An MME Pool Area is defined as an area within which a UE may be served without need to change the serving MME. An MME Pool Area is served by one or more MMEs ("pool of MMEs') in parallel. MME Pool Areas are a collection of complete Tracking Areas. MME Pool Areas may overlap each other.

MMS. Multimedia Messaging Service is a standard way to send messages that include multimedia content to and from mobile phones. It extends the core SMS (Short Message Service) capability.

MMS-C. MMS-Center is a carrier's store-and-forward server for MMS messages.

MN. Mobile Node is any device, handset, personal digital assistant, laptop, that connects to the Internet using wireless technology. A node that, as part of normal use, changes its point of attachment to the Internet. See **MS**.

MNC. Mobile Network Code. See **PLMN**.

MNO. Mobile Network Operator (also known as a wireless service provider, wireless carrier, cellular company, or mobile network carrier) is a provider of wireless communications services that owns or controls all the elements necessary to sell and deliver services to an end user, including radio spectrum allocation, wireless network infrastructure, back haul infrastructure, billing, customer care and provisioning computer systems and marketing, customer care, provisioning and repair organizations.

MNSRID. Mobile Node Session Reference ID. Denotes the calling number of the MN (the number from which the call is being made).

MO. A Mobile Originated call is initiated by a Mobile Station (MS).

Mobile hotspot. A mobile hotspot is a wireless device that connects to a cellular 3G or 4G network and then creates a personal Wi-Fi network, allowing access to the Internet. The hotspot can be a standalone device or a built-in smartphone function. See **Tethering**.

Mobility. The ability of a mobile node to change its point-of-attachment from one link to another while maintaining all existing communications and using only its IP home address.

Mobility Agent. A node (typically, a router) that offers support services to mobile nodes. A mobility agent can be either a Home Agent (HA) or a Foreign Agent (FA).

MP-eBGP. Multiprotocol-External BGP. See **BGP**.

MP-iBGP. Multiprotocol-Internal BGP. See **BGP**.

MPC. Mobile Packet Core. See **EPC**.

MPLS. Multiprotocol Label Switching is a mechanism in high-performance telecommunications networks that directs data from one network node to the next based on short path labels rather than long network addresses, avoiding complex lookups in a routing table. The labels identify virtual links (paths) between distant nodes rather than endpoints. MPLS can encapsulate packets of various network protocols. See **LER** and **LSR**.

MS. Mobile Station. See **MN**.

MSC. Mobile Switching Center. The MSC switches MS-originated or MS-terminated traffic. An MSC is usually connected to at least one base station. It may connect to other public networks PSTN, ISDN, etc., other MSCs in the same network. Another name used to identify the MSC is the Mobile Telephone Switching Office (MTSO). The MSC provides the interface for user traffic between the wireless network and other public switched networks, or other MSCs.

MSEG-3G. Mobile Service Edge Gateway - 3G.

MSID. Mobile Station Identification is a a number that is associated with the home service provider and the wireless phone number. This number is reprogrammed when the user changes home service providers.

MSISDN. Mobile Subscriber International ISDN Number is a number uniquely identifying a subscriber in a GSM or a UMTS mobile network. Essentially, it is the telephone number to the SIM card in a mobile/cellular phone.

MSS. Mobile Switching Center Server is a 3G core network element which controls the network switching subsystem elements. See **MSC**.

MT. A Mobile Terminated call is received by a Mobile Station (MS).

MTBF. Mean Time Between Failure. Synonymous with MTTF, this is the anticipated time between failures of the same component.

MTP-3. Message Transfer Part is the component of SS7 used for communication in PSTNs. MTP is responsible for reliable, unduplicated and in-sequence transport of SS7 messages between communication partners. MTP Level 3 provides the Signalling Network functional level for narrowband signalling links (ITU-T, Q.704) and broadband signalling links (ITU-T, Q.2210).

MTTF. Mean Time To Failure is the average interval of time that a component will operate before failing.

MTTR. Mean Time To Repair is the average amount of time needed to repair or replace a component, recover a system, or otherwise restore service after a failure.

MUR. Mobility Unified Reporting is a Web-based application providing a unified reporting interface for diverse data from in-line service and storage applications. It runs on a dedicated, high availability Solaris server.

MVG. Mobile Video Gateway is the central component of the Cisco Mobile Video Solution. It employs a number of video optimization techniques that enable mobile operators to enhance the video experience for their subscribers while optimizing the performance of video content transmission over the mobile network.

MVNE. Mobile Virtual Network Enabler supplies marketing/sales personnel, customer support and/or billing systems to an MVNO.

MVNO. Mobile Virtual Network Operator is a wireless communications services provider that does not own the radio spectrum or wireless network infrastructure over which the MVNO provides services to its customers. An MVNO enters into a business agreement with an MNO to obtain bulk access to network services at wholesale rates, then sets retail prices independently.

N

NAI. Network Address Identifier is used to create a new unique subscriber identifier, based on ESN or other identifiers, when a subscriber enters the network without a user name.

NAS. Non-access stratum s a functional layer in the UMTS wireless telecom protocol stack between a core network and user equipment. The layer supports signalling and traffic between those two elements.

NAS protocols. Non-access stratum protocols are the protocols between UE and MSC or SGSN that are not terminated in the UTRAN, and the protocols between UE and MME that are not terminated in the EUTRAN.

NAS signalling connection recovery. This is a mechanism initiated by the NAS to restore the NAS signalling connection on indication of "RRC connection failure" by the lower layers.

NAS signalling connection. This is a P2P S1 mode connection between UE and MME. A NAS signaling connection consists of the concatenation of an RRC connection via the "LTE-Uu" interface and an S1AP connection via the S1 interface. The UE considers the NAS signalling connection established when the RRC connection has been established successfully. The UE considers the NAS signalling connection released when the RRC connection has been released.

NAT. Network Address Translation is protocol (RFC 1631) that enables a LAN to use one set of IP addresses for internal traffic and another set of IP addresses for external traffic.

NE. Network Element is defined as a manageable logical entity uniting one or more physical devices. This allows distributed devices to be managed in a unified way using one management system. It includes a facility or equipment used in the provision of a telecommunications service, features, functions, and capabilities that are provided by means

of such facility or equipment, including subscriber numbers, databases, signaling systems, and information sufficient for billing and collection or used in the transmission, routing, or other provision of a telecommunications service.

NEBS. Network Equipment Building Standards is a rigid and extensive set of performance, quality, safety, electrical, and environmental recommendations that are applicable to devices installed in a carrier's Central Office (CO).

Network Type. The network type associated with HPLMN or a PLMN on the PLMN selector. The MS uses this information to determine what type of radio carrier to search for when attempting to select a specific PLMN. A PLMN may support more than one network type.

NHLFE. Next-Hop Label Forwarding Entry is an MPLS table structure. See **ILM** and **FTN**.

NMS. Network Management System consists of applications that provide overall management of all network elements. Defined by the third tier of the TMN model of telecommunications management networks.

Nomadicity. The full range of network technology being designed to come to the assistance of the mobile (or nomadic) computer user, not limited to network-layer protocols.

Non-GBR bearer. An EPS bearer that uses network resources that are not related to a Guaranteed Bit Rate (GBR) value.

NPU. Network Processor Unit is a high-speed, state-of-the-art processor optimized to perform packet forwarding functions.

NPU Manager. The StarOS NPU manager task provides NPU-related information to other software tasks and performs recovery services for the NPU. An NPU manager task is started for each processing card in the system.

NRI. Network Resource Identifier is a specific parameter used to identify the CN node assigned to serve a mobile station. The UE derives the NRI from TMSI, P-TMSI, IMSI or IMEI.

NTP. Network Time Protocol is a networking protocol for clock synchronization between computer systems over packet-switched, variable-latency data networks. NTP provides Coordinated Universal Time (UTC) including scheduled leap second adjustments. No information about time zones or daylight saving time is transmitted. See **Stratum clocking**.

Number Portability. This is a network feature that allows consumers to change service providers, locations, or service types without changing their telephone numbers. Number portability is a requirement for both wireline and wireless operators. An SS7 number portability database provides routing addresses for ported numbers. Network devices query the number portability database to resolve numbers during a call or service transaction. See **SRF**.

O

OAM. Operations, Administration and Maintenance is the processes, activities, tools, standards, etc. involved with operating, administering, managing, and maintaining any computer-based system.

OC-3. OC-3 is a network line with transmission data rate of up to 155.52 Mbps (payload: 148.608 Mbps; overhead: 6.912 Mbps, including path overhead) using fiber optics. Depending on the system, OC-3 is also known as STS-3 (electrical level) and STM-1 (SDH).

OCS. Online Charging System is a system that allows a communications service provider to charge their customers, in real time, based on service usage.

OFCS. Offline Charging System concurrently collects charging information for network resource usage. The charging information is then passed through a chain of logical charging functions. At the end of this process, CDR files are generated by the network, which are then transferred to the network operator's Billing Domain for subscriber billing and/or inter-operator accounting. See **CTF**, **CDF** and **CGF**.

OID. An OID, in the context of SNMP, consists of the object identifier for an object in a Management Information Base (MIB). Each OID identifies a variable that can be read or set via SNMP. MIBs use the notation defined by ASN.1.

OLC2. Optical Line Card (ATM, ASR 5000). Supports ATM over both SDH and SONET.

OMC. Operations and Maintenance Center is the central location to operate and maintain the network.

OMG. Object Management Group is an open membership, not-for-profit consortium that produces and maintains computer industry specifications for CORBA and other related protocols.

OOB. Out-of-band Management is a method wherein management information exchanged between the network element and its associated management application is carried on a separate communications path from the user data that is coming to/from the network element. Conversely, in-band management is management data that is carried across the same interface as user data.

OS. Operating System is a collection of software that manages computer hardware resources and provides common services for computer programs. See **StarOS**.

OSS. Operations Support System. Methods and procedures that support the daily operations of a carrier's network infrastructure. This includes order processing, equipment assignment, and other administrative functions related to the devices installed in the network.

P

P2P. Peer-to-Peer refers to a distributed computer network in which each computer in the network can act as a client or server for the other computers in the network, allowing shared access to various resources such as files, peripherals, and sensors without the need for a central server. All computers in the network to use the same or a compatible program to connect to each other and access files and other resources found on the other computer. P2P networks can be used for sharing content such as audio, video, data, or anything in digital format.

P-CSCF. Proxy-Call Session Control Function is the first point of contact for the UE in the IMS network. The UE needs to establish a bearer context using which the IMS signalling is carried by the UE with the P-CSCF.

P-CSCF Discovery. As part of the initial context establishment, the system may be required to select/discover a P-CSCF to be used by the UE and send the selected P-CSCF information to the UE in the create response for that PDP context. This procedure is called the P-CSCF discovery procedure.

P-GW. Packet Data Network Gateway is the node that terminates the SGi interface towards the PDN. If a UE is accessing multiple PDNs, there may be more than one P-GW for that UE. The P-GW provides connectivity to the UE to external packet data networks by being the point of exit and entry of traffic for the UE. A UE may have simultaneous connectivity with more than one P-GW for accessing multiple PDNs. The P-GW performs policy enforcement, packet filtering for each user, charging support, lawful interception and packet screening.

P-TMSI. Packet domain TMSI is an identifier for the UE allocated by the SGSN. P-TMSI contains an NRI (Network resource identifier) which is used to identify the node that allocated it. In 3G, LAC (present as part of RAI) is used to identify the SGSN pool and NRI is used to identify a node within the pool. See **TMSI**.

PBA. Proxy Binding Acknowledgement.

PBU. Proxy Binding Update (defined in RFC 5213 Proxy Mobile IPv6). See **PMIPv6**.

PC. Point Code (PC). A unique address for a node in an SS7 environment.

PCC. Policy and Charging Control. The Cisco PCC solution is comprised of an Intelligent Policy Control Function (IPCF) and Subscriber Service Controller (SSC). It also includes a Web-based GUI tool, Policy Provisioning Tool (PPT) to implement and control the policy based subscriber access in the existing wireless network, as well as service flow based charging implementation.

PCCM. PDN Connection Configuration Message.

PCEF. Policy and Charging Enforcement Function is implemented in the serving gateway to enforce gating and QoS for individual IP flows on the behalf of the PCRF. It also provides usage measurement to support charging.

PCF. Packet Control Function. A part of the 3G networking equipment that relays packet data and control signalling between the BSC and the PCF. In some cases, the PCF may be integrated into the BSC.

PCO. Protocol Configuration Options is a portion of an APN that contains external network protocol options that may need to be transferred between the GGSN and the MS (3GPP TS 24.008).

PCRF. Policy and Charging Rules Function is the node designated in real-time to determine policy rules in a multimedia network. It operates at the network core and efficiently accesses subscriber databases and other specialized functions, such as a charging systems, in a scalable, reliable, and centralized manner. The PCRF aggregates information to and from the network, operational support systems, and other sources (such as portals) in real time, supporting the creation of rules and then automatically making intelligent policy decisions for each subscriber active on the network. Such a network might offer multiple services, quality of service (QoS) levels, and charging rules. See **PCC**.

PCU. Packet Control Unit. Typically a component in the BSS that connects to the BSC to an SGSN in the core network of a GPRS/UMTS wireless network. Once the call is established, the PCU handles the packet data portion of a wireless call.

PDIF. Packet Data Interworking Function. A security gateway providing secure voice and data over a Wi-Fi network via an IPSec tunnel.

PDN. Packet Data Network is any packet-based data network, such as the Internet or an intranet, that a mobile subscriber would attempt to access.

PDN address. This is an IP address assigned to the UE by the Packet Data Network Gateway (P-GW).

PDN Connection. The association between a UE represented by one IPv4 address and/or one IPv6 prefix/address, and a PDN represented by an APN.

PDN-ID. PDN Identifier.

PDP. Packet Data Protocol. See **PDP Context**.

PDP Context. The Packet Data Protocol context is a data structure present on both the SGSN and GGSN that contains the subscriber's session information when the subscriber has an active session. When a mobile wants to use GPRS, it must first attach and then activate a PDP context.

PDP Session. The Packet Data Protocol session is a unique association of a subscriber with a network access service given by the combination of MSISDN, APN and IP address. A PDP session can consist of one or more PDP contexts (one primary and zero or more secondary).

PDSN. Packet Data Serving Node is the part of the 3G network that performs packet processing and re-direction to the mobile user's home network through communications with the Home Agent (HA).

PE. Provider Edge router is an MPLS VPN router that provides service provider peering service. See **PE**.

PEAP. Protected Extensible Authentication Protocol is a protocol that encapsulates the Extensible Authentication Protocol (EAP) within an encrypted and authenticated Transport Layer Security (TLS) tunnel.

PEP. Performance Enhancing Proxy is used to improve the performance of the Internet protocols (for example, TCP) on network paths where native performance suffers due to characteristics of a link or sub-network on the path.

PFU. Power Filter Unit (ASR 5000, ASR 5500). Supplies filtered -48VDC power to the ASR 5x00 chassis, cards and sub-components.

Pi Interface. The packet data interface from the Foreign Agent to the Internet or Home Agent.

Plain NAS message. A NAS message with a header that does not include a message authentication code or a sequence number.

PLMN. Public Land Mobile Network is a network that is established and operated by an administration or by a recognized operating agency (ROA) for the specific purpose of providing land mobile telecommunications services to the public. PLMN designates a GSM, GPRS or UMTS public mobile communications network that is identified by its MCC and MNC. PLMNs interconnect with other PLMNs and PSTNs for telephone communications or with internet service providers for data and internet access of which links are defined as interconnect links between providers.

PMIP. Proxy Mobile IP. WiMAX and CDMA2000 networks use PMIP to introduce a MAG (Mobility Access Gateway) into the MIP architecture which interacts with the HA. The MAG is termed a LMA (Local Mobility Anchor) in PMIP parlance. With PMIP, the use of MIP in the network is transparent to the mobile device.

PMIPv6. Proxy Mobile IP version 6 is a mobility management protocol that enables a single LTE-EPC core network to provide the call anchor point for user sessions as the subscriber roams between native EUTRAN and non-native e-HRPD access networks. S2a represents the trusted non-3GPP interface between the LTE-EPC core network and the eHRPD network anchored on the HSGW. In the eHRPD network, network-based mobility provides mobility for IPv6 nodes without host involvement. Proxy Mobile IPv6 extends Mobile IPv6 signaling messages and reuses the HA function (now known as LMA) on the P-GW.

Policy Decision. The set of policy information AGW receives from E-PDF in a Gx/Ty Diameter message. E-PDF constructs policy decision on the basis of Application Function (AF) events and events received over a Gx/Ty interface.

Policy Information. The set of policy related data stored in E-PDF associated to a user, including information determined via real-time analysis of an SDP offer/answer exchange derived information in the context of an IMS session, information derived from a pre configured charging rule and preconcerted rule set. These information includes at least charging rules, media component data, binding information and authorized QoS. Policy information such as charging rules and authorized QoS are sent in a policy decision by the E-PDF to the AGW for enforcement.

Pool-area. A pool area is an area within which an MS may roam without need to change the serving CN node. A pool area is served by one or more CN nodes in parallel. All the cells controlled by a RNC or BSC belong to the same one (or more) pool area(s).

Port. A defined physical or logical connection where data enters or leaves a network device.

POS. Packet over SONET.

PPC. Packet Processing Card (ASR 5000). The PPC is an application card providing memory and processing capabilities for handling subscriber sessions. The PPC is an economical alternative for smaller operators who do not require the throughput of PSC2s or PSC3s.

PPP. Point-to-Point Protocol. A protocol defined by RFC-1661 that allows for IP connectivity between network devices.

PPT. Policy Provisioning Tool is a Web-based client-server application that allows service providers or network operators to design policies for network usage and monitoring. These policies can then be used to monitor and control services rendered to subscribers, as well as their network usage. PPT runs on a Cisco UCS server and interfaces with other components of the PCC solution including IPCF and SSC to exchange data such as QoS profile or data plans.

Preconcerted Charging Rule. Charging rule created and configured in E-PDF by the operator.

Primary PDP Context. The first PDP context activated by a UE. At the primary PDP context activation an IP address (the PDP address) is assigned to a UE. When activated a primary PDP context is general purpose (i.e. with no associated TFT filters), during its lifetime may change to dedicated (i.e. with associated TFT filters).

PS. Packet Switched is a digital networking communications method that groups all transmitted data – regardless of content, type, or structure – into suitably sized blocks, called packets. It features delivery of variable-bit-rate data streams (sequences of packets) over a shared network. When traversing network adapters, switches, routers and other network nodes, packets are buffered and queued, resulting in variable delay and throughput depending on the traffic load in the network.

PSC. Packet Services Card (ASR 5000). The PSC is an application card providing memory and processing capabilities for handling subscriber sessions. There are multiple variants of the PSC designed to meet the processing requirements for individualized network deployments.

PSTN. Public Switched Telephone Network is the network of the world's public circuit-switched telephone networks. It consists of telephone lines, fiber optic cables, microwave transmission links, cellular networks, communications satellites, and undersea telephone cables, all inter-connected by switching centers, thus allowing any telephone in the world to communicate with any other. The technical operation of the PSTN utilizes standards created by the ITU-T.

PTI. Procedure Transaction Identity. An identity which is dynamically allocated by the UE for the UE requested ESM procedures. The procedure transaction identity is released when the procedure is completed.

Pull Model. A communication model where a policy decision is requested by the AGW.

Push Model. A communication model where a policy decision is sent unsolicited by the authorizing entity to the AGW.

Q

QCI. QoS Class Index is an operator provisioned value that controls bearer level packet forwarding treatments. Cisco EPC gateways also support the ability to map the QCI values to DiffServ codepoints in the outer GTP tunnel header of the S5/S8 connection. Additionally, the platform also provides configurable parameters to copy the DSCP marking from the encapsulated payload to the outer GTP tunnel header.

QGLC. Quad Gigabit Line Card (Ethernet 1000, ASR 5000). The QGLC is a four-port Gigabit Ethernet line card that supports the Star Channel (1 Gbps) for faster FPGA upgrades.

QoS. Quality of Service. A measure of the service quality provided to a subscriber. In the IP environment, this relates to acceptable levels of quality including bandwidth guarantees, latency, packet ordering, and other service-related levels of service.

R

R4 interface. This interface consists of a set of control and bearer plane protocols originating or terminating in various entities within the ASN that coordinate MS mobility between ASNs. See **ASN-GW**.

R6 interface. This interface consists of a set of control and bearer plane protocols for communication between the BS and the ASN-GW.

R-P interface. The Radio-Packet interface exists between the PCF and the PDSN in a CDMA2000 network.

R-P VPN. A routing domain for the ingress R-P protocol consisting of a group of physical or logical interfaces with an associated configuration. The system supports multiple R-P VPNs, and does not forward packets between multiple routing domains.

RA. Router Advertisement is an ICMP router discovery message that each router periodically multicasts from each of its multicast interfaces, announcing the IP address(es) of that interface. Hosts discover the addresses of their neighboring routers simply by listening for advertisements.

RAB. Radio Access Bearer is the entity responsible for transporting radio frames of an application over the access network in UMTS. The RAB carries the subscriber data between the handset and the core network. It is composed of one or more Radio Access Bearers between the handset and the Serving RNC, and one lu bearer between the Serving RNC and the core network. See **WCDMA**.

RADIUS. Remote Authentication Dial In User Service. A group of protocols used to provide AAA functionality for users through a defined server.

RAN or RN. Radio Access Network or Radio Network implements a radio access technology that resides between a mobile device and the core network (CN). Depending on the wireless standard being implemented, mobile phone and other wireless connected devices are varyingly known as user equipment (UE), terminal equipment or mobile station (MS). RAN functionality is typically provided by a silicon chip residing in both the core network, as well as the user equipment.

RAT. Radio Access Technologies/Radio Access Type is the air interface supported within LTE networks.

RAT-related TMSI. When the UE is camping on an EUTRAN cell, the RAT-related TMSI is the GUTI; when it is camping on a GERAN or UTRAN cell, the RAT-related TMSI is the P-TMSI.

RAU. Routing Area Update is the GPRS methodology for updating changes from one routing area to another by a mobile device. RAU is performed by the SGSN.

Rating Group. Information that identifies a user plane data traffic category and is used by the online and offline charging systems for rating purposes.

RCC. Redundancy Crossbar Card (ASR 5000). This is an interface card that provides redundant connectivity for line cards should a packet processing card fail.

RCT. Recovery Control Task is a StarOS task that controls the automatic failover and restart of other tasks within the system. Each recovery action is directed to the RCT from the HAT.

RD. Route Distinguisher is an address qualifier used only within a single internet service provider's MPLS network. It is used to distinguish the distinct VPN routes of separate customers who connect to the provider. See **VPNv6**.

Redirection. A message that is intended to cause a change in the routing behavior of the node receiving it.

REDR. Reporting Event Data Record.

Registration. The process by which the mobile node informs the Home Agent about its current care-of address.

Registration. This is the process of camping on a cell of the PLMN and doing any necessary location registrations (LRs).

Registration Area. A registration area is an area in which mobile stations may roam without a need to perform location registration. The registration area corresponds to location area (LA) for performing location updating procedure, to routing area for performing the GPRS attach or routing area update procedures, and to list of tracking areas (TAs) for performing the EPS attach or tracking area update procedure. The PLMN to which a cell belongs (PLMN identity) is given in the system information transmitted on the BCCH (MCC + MNC part of LAI). In a shared network a cell belongs to all PLMNs given in the system information transmitted on the BCCH.

Remote redirection. A redirect sent from a source not present on the local network. The source can be located anywhere in the global Internet and may have malicious intent and be untraceable.

Replay attacks. A security violation whereby a malicious entity attempts to imitate a transaction recorded during a previous and valid transaction between two protocol entities. Both protocol entities have to be aware that the subsequent identical traffic streams may no longer be valid. Since the previous transaction was valid, the algorithms for detecting replay attacks need to incorporate data that can never be reproduced in any correct subsequent transaction.

Reverse Tunnel. The direction of encapsulate data traveling from the Foreign Agent to the Home Agent.

Rf interface. This is the offline charging interface between the Charging Trigger Function (CTF) [for example, P-GW, S-GW, P-CSCF] and the Charging Collection Function (CCF). The Rf interface specification for LTE/GPRS/eHRPD offline charging is based on 3GPP TS 32.299 V8.6.0, 3GPP TS 32.251 V8.5.0 and other 3GPP specifications. The Rf interface specification for IMS) offline charging is based on 3GPP TS 32.260 V8.12.0 and 3GPP TS 32.299 V8.13.0.

RFC. Request for Comments is a document that contains Internet standards and protocols, along with other useful information that has relevance to the Internet community. RFCs provide developers the rules and directions on how to implement various Internet communications functions so that they adhere with, are interoperable to, other vendors' implementations of the same function. RFCs are controlled by the International Engineering Task Force (IETF).

RHEL. Red Hat Enterprise Linux is a Linux-based operating system developed by Red Hat and targeted toward the commercial market.

RM. Resource Management subsystem. This group of StarOS tasks assigns resources to other tasks within the system as they are initiated and monitors all resource allocations.

RMU. Rack Mounting Unit is a unit of measurement used in telecommunications to denote the amount of vertical space required to place a network device into an equipment cabinet or telecommunications rack. Each RMU is equivalent to 1.75 in. (4.45 cm.) in height.

RNC. Radio Network Controller is a governing element in UTRAN and is responsible for controlling the Node Bs that are connected to it. The RNC carries out radio resource management, some of the mobility management functions, and is the point where encryption is done before user data is sent to and from the mobile. The RNC connects to the Circuit Switched Core Network through a Media Gateway (MGW) and to the SGSN in the Packet Switched Core Network.

ROHC. RObust Header Compression is a standardized method for compressing the IP, UDP, RTP, and TCP headers of Internet packets. ROHC compresses the 40 bytes or 60 bytes of overhead into only 1 or 3 bytes by placing a compressor before the link that has limited capacity, and a decompressor after that link.

Route optimization. A process that enables the delivery of packets directly to the care-of address from a correspondent node without having to detour through the home network.

RPLMN. Registered PLMN. This is the PLMN on which certain LR outcomes have occurred. In a shared network the RPLMN is the PLMN defined by the PLMN identity of the CN operator that has accepted the LR.

RRC. Radio Resource Control belongs to the UMTS WCDMA protocol stack and handles the control plane signalling of Layer 3 between the UEs and the UTRAN.

RRM. Radio Resource Management is the system level control of co-channel interference and other radio transmission characteristics in wireless communication systems,

RS. Router Solicitation. The ICMP router discovery messages are called "Router Advertisements" and "Router Solicitations". Each router periodically multicasts a Router Advertisement from each of its multicast interfaces, announcing the IP address(es) of that interface. Hosts discover the addresses of their neighboring routers simply by listening for advertisements.

RSVP. Resource ReSerVation Protocol is a Transport Layer protocol designed to reserve resources across a network for an integrated services Internet. RSVP operates over an IPv4 or IPv6 Internet Layer and provides receiver-initiated setup of resource reservations for multicast or unicast data flows with scaling and robustness. RSVP is described in RFC 2205.

RSVP-TE. Resource Reservation Protocol - Traffic Engineering supports the reservation of resources across an IP network. Applications running on IP end systems can use RSVP to indicate to other nodes the nature (bandwidth, jitter, maximum burst, and so forth) of the packet streams they want to receive. RSVP-TE is described in RFC 5151.

RTCP. Real-Time Transport Control Protocol. See **RTP**.

RTP. Real-time Transport Protocol defines a standardized packet format for delivering audio and video over IP networks. RTP is used extensively in communication and entertainment systems that involve streaming media, such as telephony, video teleconference applications, television services and web-based push-to-talk features. It is used in conjunction with the RTCP. While RTP carries the media streams while RTCP monitors transmission statistics and quality of service (QoS) and aids synchronization of multiple streams.

RTSP. Real Time Streaming Protocol is a network control protocol designed for use in entertainment and communications systems to control streaming media servers. The protocol establishes and controls media sessions between end points. Clients of media servers issue VCR-like commands, such as play and pause, to facilitate real-time control of playback of media files from the server. Most RTSP servers use the RTP in conjunction with RTCP for media stream delivery.

Rule Base. A collection of static charging rules configured in a system.

Rule Base ID. The identifier of a rule base.

S

S1. An interface between an eNodeB and an EPC, providing an interconnection point between the EUTRAN and the EPC. It is also considered as a reference point.

S1-MME. An interface/reference point for the control plane protocol between EUTRAN and MME.

S101 mode. This mode applies to a system that operates with a functional division that is in accordance with the use of an S101 interface.

S12 Interface. A GTP-U direct tunnel interface/reference point between an S-GW and an RNC.

S13 Interface. A GTP-C/U interface/reference point between an MME and an EIR.

S2a interface. Provides the user plane with related control and mobility support between trusted non-3GPP IP access and the Gateway. S2a is based on Proxy Mobile IP. To enable access via trusted non-3GPP IP accesses that do not support PMIP, S2a also supports Client Mobile IPv4 FA mode.

S2b interface. Provides the user plane with related control and mobility support between evolved Packet Data Gateway (ePDG) and the PDN GW. It is based on Proxy Mobile IP.

S2c interface. Provides the user plane with related control and mobility support between UE and the P-GW. This reference point is implemented over trusted and/or untrusted non-3GPP Access and/or 3GPP access. This protocol is based on Client Mobile IP co-located mode.

S3 Interface. An interface/reference point between an MME and a release 8 SGSN.

S4 Interface. An interface/reference point between an S-GW and a release 8 SGSN.

S5/S8 Interface. A PMIPv6/GTP interface/reference point between a P-GW and an S-GW. S5 is the non-roaming (home network) interface between a home P-GW and a home S-GW. S8 is the roaming interface between a home P-GW and a visited S-GW.

S6a interface. Enables transfer of subscription and authentication data for authenticating/authorizing user access to the evolved system (AAA interface) between MME and HSS.

S-CDR. SGSN generated CDR.

S-GW. Serving Gateway routes and forwards data packets from the UE and acts as the mobility anchor during inter-eNodeB handovers. Signals controlling the data traffic are received on the S-GW from the MME which determines the S-GW that will best serve the UE for the session. Every UE accessing the EPC is associated with a single S-GW.

SAAU. Simultaneously Attached and Active Users.

SAE. System Architecture Evolution is the core network architecture of the 3GPP LTE wireless communication standard. It has a flat, all-IP architecture with separation of control plane and user plane traffic. The main component of the SAE architecture is the EPC with its MME, S-GW and P-GW subcomponents.

SAE-GW. System Architecture Evolution Gateway node is a collocation of S-GW and P-GW nodes on a single ASR 5x00 chassis. SAE-GW operates as a service on ASR5x00 and requires existing S-GW and P-GW services to be configured and mentioned inside SAE-GW service configuration. Existing standalone S-GW and P-GW services (when part of an SAE-GW umbrella service) work in tandem to present a single SAE-GW node as a black-box view.

SAFI. Subsequent Address Family Identifier provides additional information about the type of the Network Layer Reachability Information carried in the AFI attribute.

SAU. Simultaneously Attached Users.

SBLP. Service-based Local Policy. This term refers to the instantiation of a policy for use of bearer resources in the access network based on Authorization by a service. In the context of Go interface this is the combined QoS given to a set of IP flows for an IMS session.

SCCP. Signaling Connection Control Part is an SS7 transport layer, connection-oriented protocol that works with MTP-3 to provide routing.

SCCP Network. A proprietary concept designed to facilitate the creation and management of SCCP parameters specific to the SGSN routing.

SCM. Service Control Manager delivers and controls a multimedia services and provides an on-ramp for deploying Session Initiation Protocol (SIP)-based services, as well as a migration path to IP Multimedia Subsystem/Multimedia Domain (IMS/MMD) architectures.

SCP. Signaling Control Point is a computer database that receives information request messages from the SS7 network and returns information that is necessary for the completion of calls or services. The SCP usually receives requests for a service switching point (SSP) via signaling transfer points (STPs) that determine that additional information is necessary to complete the call.

SCT. Shared Configuration Task is a StarOS task for configuring system parameters, retrieving information, and notifying system components of configuration changes.

SCTP. Stream Control Transmission Protocol is a transport layer protocol that is message-oriented like UDP and ensures reliable, in-sequence transport of messages with congestion control like TCP. See **SIGTRAN**.

SDF. Service Data Flow Service Data Flow is an aggregate set of packet flows that matches a set of filters. Each flow is defined as a 5-tuple (source IP address, destination IP address, source port, destination port, protocol used above IP).

SDH. Synchronous Digital Hierarchy. See **SONET**.

SDP. Session Description Protocol is a format for describing streaming media initialization parameters. SDP is intended for describing multimedia communication sessions for the purposes of session announcement, session invitation, and parameter negotiation. SDP does not deliver media itself but is used for negotiation between end points of media type, format, and all associated properties. See **RTSP**.

SDT. Signalling De-Multiplexing Task. See **A11 Manager**.

Secondary PDP Context. A new activated PDP context reusing the PDP address and other PDP context information from an already active PDP context, but with a different QoS profile. A secondary PDP context may be dedicated (with associated TFT filters) or general purpose (with no associated TFT filters).

SectorID. Sector Address Identifier is used to identify an HRPD AN. The Network operator shall set the value of the SectorID according to the rules specified.

Selected PLMN. This is the PLMN that has been selected according to subclause 3.1, either manually or automatically.

Service Based Authorization. This term refers to the authorization for use of bearer resources in the access network based on a determination by the application, possibly due to negotiation involving the user. In general, bearer resources may be authorized if the resources requested at the bearer do not exceed the resources negotiated or requested at the service level.

Serving GW Service Area. A Serving GW Service Area is defined as an area within which a UE may be served without need to change the S-GW. A Serving GW Service Area is served by one or more S-GWs in parallel. Serving GW Service Areas are a collection of complete Tracking Areas. Serving GW Service Areas may overlap each other. See **TAI list**.

Session Manager. A group of StarOS tasks used by the system for subscriber processing services. Each CP can have multiple session managers. Each session manager is paired with an AAA manager, and can support multiple A11 managers.

SFTP. SSH File Transfer Protocol (also Secure File Transfer Protocol or Secure FTP) is a network protocol that provides file access, file transfer, and file management functionalities over any reliable data stream. It is an extension of the Secure Shell protocol (SSH) version 2.0 to provide secure file transfer capability.

SGs Interface. An interface between an MME and an MSC/VLR. Used for CSFB scenarios.

SGSN. Serving GPRS Support Node tracks the location of mobile devices in a GSM GPRS or UMTS network and routes packet traffic from the BSS to the GGSN.

Shared Network. An MS considers a cell to be part of a shared network, when multiple PLMN identities are received on the BCCH.

SI. International System of Units is the modern form of the metric system. It comprises a system of units of measurement devised around seven base units and the convenience of the number ten.

SID. System Identification. A number that uniquely identifies a network within a cellular or Personal Communication Service (PCS) system.

SIGTRAN. SIGTRAN (SIGnaling TRANsport) is the name of an IETF set of specifications for a family of protocols that provide reliable datagram service and user layer adaptations for Signaling System 7 (SS7) and ISDN communications protocols (RFC 2719). The SIGTRAN protocols are an extension of the SS7 protocol family. It supports the same application and call management paradigms as SS7 but uses an Internet Protocol (IP) transport called SCTP. SIGTRAN is used to carry PSTN signaling over IP.

SIM. Subscriber Identity Module stores the IMSI of an ME/MS/UE (SIM card).

Simple IP. The most commonly used routing protocol on the Internet. This is the IP portion of the TCP/IP suite of protocols used in wireless packet communications.

SIP. Session Initiation Protocol is an IETF-defined signaling protocol widely used for controlling communication sessions such as voice and video calls over Internet Protocol (IP). The protocol can be used for creating, modifying and terminating two-party (unicast) or multiparty (multicast) sessions. Sessions may consist of one or several media streams.

SIT. System Initiation Task is a critical StarOS task that is responsible for starting all tasks and system initialization on an ASR 5x00 chassis.

SLA. Service Level Agreement.

SMC. System Management Card (ASR 5000). This application card serves as the primary controller and is responsible for initializing the entire system and loading the StarOS configuration image into other cards in the chassis. It also provides out-of-band management interfaces and access to centralized chassis resources.

SMLC. Serving Mobile Location Center is a network element in GSM Networks that resides in the BSC (Base Station Controller) and calculates network-based location of mobile stations (MS).

SMS. Short Message Service is a text messaging service component of phone, web, or mobile communication systems, using standardized communications protocols that allow the exchange of short text messages between fixed line or mobile phone devices.

SMTP. Simple Mail Transfer Protocol is an Internet standard for e-mail transmission across IP networks [RFC 821, RFC 5321 (2008)].

SNMP. Simple Network Management Protocol is an Internet-standard protocol for managing devices on IP networks. It is used mostly in network management systems to monitor network-attached devices for conditions that warrant administrative attention. See **MIB**.

SOAP. Simple Object Access Protocol is a protocol specification for exchanging structured information in the implementation of Web Services in computer networks. It relies on XML for its message format.

SoLSA exclusive access. Cells on which normal camping is allowed only for MS with Localized Service Area (LSA) subscription.

SONET. Synchronous Optical Networking and SDH are standardized protocols that transfer multiple digital bit streams over optical fiber using lasers or highly coherent light from light-emitting diodes (LEDs).

Source Base Station. The BS that is in control of the call is designated the source BS and remains the source BS until it is removed from control of the call.

Source Context. The StarOS context that a mobile subscribers is placed into by the system when they connect to the system through a PCF.

Source routing. A routing technique that causes some or all intermediate routing points to be represented directly in the data packet to be forwarded. This is in contrast to the typical situation in which intermediate routers rely on acquired routing state information to forward incoming packets.

Sp interface This the interface that links the PCRF with the SPR.

SPIO. Switch Processor I/O card (ASR 5000). Interface card within the system that provides input/output and management interfaces for its corresponding management card. See **SMC**.

SPMS. Service Parts Management Service is the main component of a complete Strategic Service Management process that companies use to ensure that right spare part and resources are at the right place at the right time.

SPOF. Single Point Of Failure is a part of a system that, if it fails, will stop the entire system from working.

SPR. Subscriber Profile Repository. See **SSC**.

SRF. Signaling Relay Function delivers SS7 MAP queries to a number portability database. The SRF function queries the database and can route non-ported number queries to the appropriate HLR or return an acknowledgment to the originator (example, mobile switching center [MSC]) containing the location where the subscriber has been ported.

SRP. Service Redundancy Protocol. ICSR chassis use SRP to periodically check to see if the redundancy configuration matches with either decrypted passwords or DES-based two-way encryption strings. Since the configuration is generated internally to the software, users are not able to access the configuration used to check ICSR compatibility.

SRVCC. Single Radio Voice Call Continuity is an LTE functionality that allows a VoIP/IMS call in the LTE packet domain to be moved to a legacy voice domain (GSM/UMTS or CDMA 1x).

SS7. Signaling System Seven is a set of telephony signaling protocols which are used to set up most of the world's public switched telephone network telephone calls. There are many, country-specific variants of these protocols.

SS7 Routing Domain. A proprietary concept designed to facilitate the creation and management of SS7-based configuration parameters (such as, link IDs and application server processes) by organizing and grouping them.

SSC. Subscriber Service Controller provides extended centralized PCRF and Subscriber Profile Repository (SPR) functionality in the Cisco PCC solution and manages data related to service usage and subscriber profile for IP-CAN session. It is designed to be used in conjunction with IPCF running on ASR 5x00 chassis and the PPT running on Cisco© UCS platforms. SSC typically runs on a Cisco UCS or IBM Blade Center chassis capable of supporting multiple instances of database manager (RDBMS) along with multiple instances of SSC application.

SSC. System Status Card (ASR 5500). The SSC incorporates three alarm relays (Form C contacts), an audible alarm with front panel Alarm Cutoff (ACO) and system status LEDs

SSH. Secure Shell is a cryptographic network protocol for secure data communication, remote shell services or command execution and other secure network services between two networked computers that it connects via a secure channel over an insecure network: an SSH server and an SSH client. See **TLS**.

SSID. Service Set Identifier is chosen by the client device that starts the WLAN. The SSID is broadcast in a pseudo-random order by all devices that are members of the network. See **WLAN**.

SSL. Secure Sockets Layer. See **TLS**.

StarOS. This is the operating system that runs on ASR 5x00 chassis.

Static Charging Rule. The data within the charging rule (for example, service data flow filter information) is statically assigned by configuration.

STM-1. Synchronous Transport Module level-1 is the SDH ITU-T fiber optic network transmission standard. It has a bit rate of 155.52 Mbps and provides stratum-3 timing for both TDM and packet interfaces. See **Stratum clocking**.

STP. Signal Transfer Point (STP) is a router that relays SS7 messages between signaling end-points (SEPs) and other signaling transfer points (STPs). Typical SEPs include service switching points (SSPs) and service control points (SCPs).

Stratum clocking. Stratum levels define the distance from the NTP reference clock. A reference clock is a stratum-0 device that is assumed to be accurate and has little or no delay associated with it. The reference clock typically synchronizes to the correct time (UTC) using GPS transmissions, CDMA technology or other time signals such as Irig-B, WWV or DCF77. Stratum-0 servers cannot be used on the network, instead, they are directly connected to computers which then operate as stratum-1 servers. A server that is directly connected to a stratum-0 device is called a stratum-1 server; it acts as a primary network time standard. A stratum-2 server is connected to the stratum-1 server over a network path. Thus, a stratum-2 server gets its time via NTP packet requests from a stratum-1 server. A stratum-3 server gets its time via NTP packet requests from a stratum-2 server, and so on.

SWu interface. This is the reference point between the UE and the ePDG handling of IPSec tunnels. It represents the UE-initiated tunnel between the WLAN UE and the ePDG for user data packet transmission within the IPSec tunnel, and tear down of the tunnel and support for fast update of IPSec tunnels during handover between two untrusted non-3GPP IP accesses. See **ePDG**.

T

T1. A Ti circuit is a North American data circuit that runs at a 1.544 Mbps line rate. The T-1 circuit combines 24 x 64 Kbps channels into one single channel, using a technique known as Channel Bonding or Channel Aggregation. See **E1**.

TAC. The Cisco Technical Assistance Center provides around-the-clock, technical support services for Cisco products and technologies. Go to www.cisco.com/tac.

TAC. Type Allocation Code is part of the IMEI which is available after a call is established. The TAC of the bearer is looked up in a mobile smartphone database. If a match is found, a tethering detection decision for that subscriber session depends on a subsequent OS and/or User Agent (UA) match. If required, ECS performs tethering detection for all flows for that subscriber. See **Tethering**.

TAI. Tracking Area Identifier. A tracking area that consists of multiple eNodeBs.

TAI list. A list of TAIs that identify the tracking areas that the UE can enter without performing a tracking area updating procedure. The TAIs in a TAI list assigned by an MME to a UE pertain to the same MME area.

TAU. Tracking Area Update is a an LTE procedure used to ensure that the MME knows in which TAI or set of TAIs the UE is currently located. TAU is performed when the UE moves to a TAI outside the list of TAIs it was assigned in the previous TAI Update or Attach procedure. TAU is also performed periodically even if the UE remains in the assigned TAI.

TCAP. Transaction Capabilities Application Part (ITU-T, Q.771-Q.775 or ANSI T1.114) is a protocol for SS7 networks. Its primary purpose is to facilitate multiple concurrent dialogs between the same sub-systems on the same machines, using Transaction IDs to differentiate these, similar to the way TCP ports facilitate multiplexing connections between the same IP addresses on the Internet. See **SS7**.

TCP. Transmission Control Protocol provides reliable, ordered delivery of a stream of octets from a program on one computer to another program on another computer over an IP network.

TD-LTE. Time-Division Long-Term Evolution (TD-LTE), also referred to as Long-Term Evolution Time-Division Duplex (LTE TDD), is a 4G mobile-telecommunications technology and standard that offers asymmetric use of unpaired spectrum. It allocates separate channels for outgoing and incoming signals, emulating full-duplex transmission over a half-duplex communication link.

TDF. Traffic Detection Function.

TDM. Time Division Multiplex is a technique for simultaneously transmitting a number of separate data signals over a single communications medium by interleaving a part of each signal one after another.

TDMA. Time Division Multiple Access is one of the wireless technology classes that encompasses 2G, 2.5G, and 3G communications. The other is CDMA.

TEC. Telecommunication Engineering Center is a body under telecom commission and a nodal agency of the Department of Telecommunications, Ministry of Communications and Information Technology, Government of India which is responsible for drawing up of standards, generic requirements, interface requirements, service requirements and specifications for telecom products, services and networks.

Tethering. Refers to the use of a smartphone as a cellular data modem or a Wi-Fi hotspot to connect or "tether" tablets, notebooks or laptops to a wireless data network.

TFT. Traffic Flow Template. See **HSGW**.

TFTP. Trivial File Transfer Protocol is a file transfer protocol used for automated transfer of configuration or boot files between machines in a local environment. Compared to FTP, TFTP is extremely limited, providing no authentication, and is rarely used interactively by a user.

TIA. Tunnel Inner Address is an IP address assigned by a PDIF/FA and used to create the initial CHILD_SA. After authentication and the creation of a new IPSec_SA with the HoA, the initial CHILD_SA is torn down and the address returned to the pool.

TLLI. Temporary Logical Link Identifier is derived from the P-TMSI and the RA to uniquely identify an MS in a GPRS sub-network.

TLS. Transport Layer Security is a cryptographic protocol that provides communication security over the Internet. TLS encrypts the segments of network connections at the Application Layer for the Transport Layer, using asymmetric

cryptography for key exchange, symmetric encryption for confidentiality, and message authentication codes for message integrity.

TLV. Type Length Value. Within data communication protocols, optional information may be encoded as a type-length-value or TLV element inside of the protocol. The type and length are fixed in size (typically 1-4 bytes), and the value field is of variable size.

TMN. Telecommunications Management Network is a protocol model defined by ITU-T (series M.3000) for managing open systems in a communications network. TMN provides a framework for achieving interconnectivity and communication across heterogeneous operations system and telecommunication networks. To achieve this, TMN defines a set of interface points for elements which perform the actual communications processing (such as a call processing switch) to be accessed by elements, such as management workstations, to monitor and control them. The standard interface allows elements from different manufacturers to be incorporated into a network under a single management control.

TMSI. Temporary Mobile Subscriber Identity is the identity that is most commonly sent between the mobile and the network. TMSI is randomly assigned by the VLR to every mobile in the area, the moment it is switched on. The number is local to a location area, and so it has to be updated each time the mobile moves to a new geographical area.

TOS. Type of Service. Outdated, see **DSCP Marking**.

TPO. Traffic Performance Optimization.

Traffic Category. User plane data traffic subject to the same access cost and rating type. A traffic category is identified by a Rating-Group and gathers a set of services.

Traffic flow aggregate. A temporary aggregate of packet filters that are included in a UE requested bearer resource modification procedure and that is inserted into a traffic flow template (TFT) for an EPS bearer context by the network once the UE requested bearer resource modification procedure is completed.

Triangular routing. The path followed by a packet from a correspondent host to a mobile node that must first be routed to the mobile node's Home Agent (HA).

TTG. Tunnel Termination Gateway provides access and an attachment point to the mobile packet core for un-trusted non-3GPP access such as Wi-Fi. TTG allows service providers to deploy Wi-Fi networks in order o off-load macro cellular radio.

TTLS. Tunneled Transport LAyer Security (RFC 5281). See **EAP**.

Tunnel. A path followed by a first packet while it is encapsulated within the payload portion of a second packet.

Tunneling. The same as encapsulation, but with additional connotations about changing the effects of Internet routing on the original IP packet.

U

UCS. Unified Computing System is a Cisco data center platform offering unified, model-based management, end-to-end provisioning, and migration support to accelerate and simplify application deployment with greater reliability and security. See **RHEL**.

UDP. User Datagram Protocol (UDP) allows computer applications to send messages referred to as datagrams, to other hosts on an IP network without prior communications to set up special transmission channels or data paths.

UDR. Usage Data Records are generated by the ECS according to action statements in rule commands. Up to 32 different UDR schema types may be specified, each composed of up to 32 fields or analyzer parameter names. The records are written thresholds in comma-separated variable (CSV) format.

UE. User Equipment is a term commonly used in 3G/4G scenarios. Equivalent to MS or mobile station (commonly used in 2G/2.5G scenarios) and to MN or mobile node (commonly used in 2G/2.5G scenarios involving IP-level functions).

UE-associated logical S1-connection. The UE-associated logical S1-connection uses the identities MME UE S1AP ID and eNodeB UE S1AP ID. For a received UE associated S1-AP message the MME identifies the associated UE based

on the MME UE S1AP ID IE and the eNodeB identifies the associated UE based on the eNodeB UE S1AP ID IE. The UE-associated logical S1-connection may exist before the S1 UE context is setup in eNodeB.

UE-associated signalling. When S1-AP messages associated to one UE uses the UE-associated logical S1-connection for association of the message to the UE in eNodeB and EPC.

UID. Unique Identifier is any identifier which is guaranteed to be unique among all identifiers used for those objects and for a specific purpose.

UL. Uplink is the radio transmission path from the mobile station (cell phone) to a base station (cell site). See **DL**.

UMTS. Universal Mobile Telecommunications System is the GSM-based evolution for 3G wireless communications. This term is also referred to as WCDMA.

Unicast/MBMS-mixed cell. This is the cell supporting both unicast and MBMS transmissions

UPS. Uninterruptible Power Supply or Uninterruptible Power Source is an electrical apparatus that provides emergency power to a load when the input power source, typically mains power, fails.

USSD. Unstructured Supplementary Service Data is a protocol used by GSM cellular telephones to communicate with the service provider's computers. USSD can be used for WAP browsing, prepaid callback service, mobile-money services, location-based content services, menu-based information services, and as part of configuring the phone on the network.

UTC. Coordinated Universal Time is the primary time standard by which the world regulates clocks and time. It is one of several closely related successors to Greenwich Mean Time (GMT). Time zones around the world are expressed as positive or negative offsets (hours or minutes) from UTC.

UTRAN. UMTS Terrestrial Radio Access Network is a collective term for the NodeBs and Radio Network Controllers (RNCs) which make up the UMTS radio access network. This 3G communications network can carry many traffic types from real-time Circuit Switched to IP based Packet Switched. The UTRAN allows connectivity between the UE (user equipment) and the core network.

V

VLAN. Virtual Local Area Network is a concept of partitioning a physical network, so that distinct broadcast domains are created. This is usually achieved on switch or router devices that can mark packets through tagging, so that a single interconnect (trunk) may be used to transport data for various VLANs. VLAN membership can be configured through software instead of physically relocating devices or connections. Most enterprise-level networks today use VLANs.

VLR. Visited Location Register. The VLR caches access service parameter information (such as the MS/UE mobile number) that it obtains from a particular user's HLR upon call establishment. See **HLR**.

VoIP. Voice over IP. The protocol that describes the packetization of analog voice signals into digital data packets. See **SIP**.

VoLTE. Voice over LTE is an IMS-based specification devised by network operators seeking a standardised system for transferring voice traffic over LTE. Work on the definition of VoLTE is ongoing. See **CSFB**.

VPLMN. Visited PLMN is a PLMN different from the HPLMN (if the EHPLMN list is not present or is empty) or different from an EHPLMN (if the EHPLMN list is present).

VPN. Virtual Private Network is a virtual router or domain instance that enables secure communications between allowed network users and devices. Context is the work most commonly used to denote this type of connectivity.

VPNv6. This is an IPv6 VPN connected over an IPv6 interface or sub-interface to the Service Provider backbone via a PE router. The site can be both IPv4 and IPv6 capable. Each IPv6 VPN has its own address space which means a given address denotes different systems in different VPNs. This is achieved via a new address-family, VPN-IPv6 or VPNv6 address-family, which prepends a Route Distinguisher (RD) to the IP address.

VRF. Virtual Routing and Forwarding is a technology that allows multiple instances of a routing table to co-exist within the same router at the same time. Because the routing instances are independent, the same or overlapping IP addresses can be used without conflicting with each other.

VSA. Vendor-Specific Attribute is a RADIUS encapsulated AVP (IETF Attribute 26) that allows vendors to support attributes that are not suitable for general use.

VSNP. Vendor Specific Network Protocol. See **HSGW**.

VSNCP. Vendor Specific Network Control Protocol. See **HSGW**.

W

WAP. Wireless Application Protocol is a technical standard for accessing information over a mobile wireless network. A WAP browser is a web browser for mobile devices, such as mobile phones, that use the protocol.

WCDMA or W-CDMA. Wideband CDMA. The GSM-based evolution for 3G wireless communications. See **UMTS**.

WEM. Web Element Manager is a CORBA-based application that provides complete Fault, Configuration, Accounting, Performance, and Security (FCAPS) capability for the ASR 5x00 system under management. WEM application implements a client-server architecture that allows remote clients with Java-enabled Web browsers to manage one or more systems via the server component which implements the CORBA interfaces. The server component is fully compatible with the Solaris® and Cisco RHEL operating systems.

Wi-Fi. Wi-Fi is a popular technology that allows an electronic device to exchange data wirelessly (using radio waves) over a computer network, including high-speed Internet connections. The Wi-Fi Alliance defines Wi-Fi as any wireless local area network (WLAN) products that are based on IEEE 802.11.

WiMAX. Worldwide Interoperability for Microwave Access (IEEE 802.16). WiMAX is a wireless communications standard designed to provide 30 to 40 Mbps data rates (up to 1 Gbps for fixed stations).

WLAN. Wireless Local Area Network links two or more devices using some wireless distribution method (typically spread-spectrum or OFDM radio), and usually providing a connection through an access point to the wider internet. This gives users the mobility to move around within a local coverage area and still be connected to the network. Most modern WLANs are based on IEEE 802.11 marketed under the Wi-Fi brand name.

WLC. Wireless LAN Controller is used in combination with the Lightweight Access Point Protocol (LWAPP) to manage light-weight access points in large quantities by the network administrator or network operations center. The wireless LAN controller is part of the Data Plane within the Cisco Wireless Model. The WLAN controller automatically handles the configuration of anywhere from 6 to 500 wireless access-points, depending on the model.

X

X2 Interface. It is a logical interface between two eNodeBs. Whilst logically representing a point to point link between eNodeBs, the physical realization need not be a point to point link.

XGLC. 10 Gigabit Ethernet Line Card (Ethernet 10G, ASR 5000). The single-port XGLC supports the IEEE 802.3-2005 revision which defines full duplex operation of 10 Gigabit Ethernet. PSC2s or PSC3s are required to achieve maximum sustained rates with the XGLC.

XML. eXtensible Markup Language is a markup language created to structure, store, and transport data by defining a set of rules for encoding documents in a format that is both human-readable and machine-readable. XML emphasize simplicity, generality, and usability over the Internet. It is a textual data format that is widely used for the representation of arbitrary data structures, for example in web services.

Z

ZebOS. This is a scalable, robust and standards-based Layer 2 and Layer 3 carrier-class switching and routing software solution.