

**Oracle® Communications
Diameter Signaling Router**
Policy and Charging Application User's Guide
E63633 Revision 02

October 2015

Oracle® Communications Diameter Signaling Router Policy and Charging Application User's Guide
Copyright © 2011, 2015, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Table of Contents

Chapter 1: Introduction.....	16
Overview.....	17
Scope and Audience.....	17
Manual Organization.....	17
Documentation Admonishments.....	18
Related Publications.....	18
Locate Product Documentation on the Oracle Help Center Site.....	18
Customer Training.....	19
My Oracle Support (MOS).....	19
Emergency Response.....	20
Chapter 2: User Interface Introduction.....	21
User Interface Organization.....	22
User Interface Elements.....	22
Main Menu Options.....	24
Missing Main Menu options.....	28
Common Graphical User Interface Widgets.....	28
Supported Browsers.....	29
System Login Page.....	29
Main Menu Icons.....	30
Work Area Displays.....	31
Customizing the Splash Page Welcome Message.....	34
Column Headers (Sorting).....	35
Page Controls.....	35
Clear Field Control.....	36
Optional Layout Element Toolbar.....	36
Filters.....	37
Pause Updates.....	40
Max Records Per Page Controls.....	40
Chapter 3: Policy and Charging Introduction.....	41
Overview.....	42
The P-DRA Database.....	43

Bindings vs Sessions.....	43
The Binding Database.....	46
The Session Database.....	47
Binding Key Query Tool.....	48
The OC-DRA Database.....	48
Online Charging over Gy/Ro Reference Point.....	49
Binding-independent Interface.....	50
Deployment Topology.....	50
Policy DRA in Roaming Scenarios.....	52
PCA Configurable Components.....	53
IPFE.....	62
PCA Scalability.....	64
MP Growth.....	65
Mated Pair Growth.....	65
Small System Support.....	66
Redundancy.....	69
MP Server Redundancy.....	69
Site Redundancy.....	70
Data Redundancy.....	71
OAM Server Redundancy.....	72
IP Networking.....	74
PCA Routing.....	75
Ingress Routing.....	75
Egress Routing.....	77
PCA Data Auditing.....	80
PCA and Application Chaining.....	82
The Communication Agent.....	83
Diameter Routing and Communication with PCA.....	84
PCA and IDIH Metadata.....	87
PCA Capacity Constraints.....	94
PCA Assumptions and Limitations.....	95

Chapter 4: Policy DRA Overview.....98

The Policy DRA Function.....	99
PCRF Pools and Sub-Pools Concepts and Terminology.....	99
Policy DRA Functions.....	110
Diameter Request Message Processing.....	110
Query Subscriber's Binding Status.....	111
PCRF Selection and Routing.....	112
Topology Hiding Process.....	113

Diameter Answer Message Processing.....	113
Subscriber Session and Binding Database Management.....	113
Subscriber Identification and Binding.....	113
Binding-capable Sessions.....	115
Binding-dependent Sessions.....	120
In-session Message Processing.....	122
Topology Hiding.....	122
Session Integrity.....	124
Suspect Binding Removal.....	127
Chapter 5: Online Charging DRA Overview.....	129
Online Charging DRA Functions.....	130
OCS Selection and Routing.....	130
Single OCS Pool Mode.....	130
Multiple OCS Pools Mode.....	131
Regionalized Routing.....	132
Session State Maintenance.....	133
Gy/Ro Diameter Request Message Processing.....	136
Session Initiation Request Message Processing.....	137
In-Session Request Message Processing.....	138
Event Request Message Processing.....	139
Gy/Ro Diameter Answer Message Processing.....	139
Session Initiation Answer Message Processing.....	141
In-Session Answer Message Processing.....	141
Event Answer Message Processing.....	141
DRL-Initiated Answer Message Processing.....	142
Chapter 6: Policy and Charging Configuration.....	143
Policy and Charging Configuration Overview.....	144
Pre-Configuration Activities.....	148
Initial Installation for PCRF Pooling.....	148
Upgrading for PCRF Pooling.....	149
Diameter Common Configuration for PCA.....	150
Diameter Configuration for PCA.....	150
NOAM Configuration.....	155
General Options.....	155
SBR Databases.....	158
Access Point Names.....	160
SBR Database Resizing Plans.....	166
SBR Data Migration Plans.....	170

Policy DRA Configuration.....	172
Online Charging DRA.....	190
Alarm Settings.....	196
Congestion Options.....	201
SOAM Configuration.....	203
Policy DRA.....	204
Online Charging DRA.....	234
Error Codes.....	243
Alarm Settings.....	252
Congestion Options.....	257
Post-Configuration Activities.....	259
Enable the PCA Application.....	260
Enable Connections.....	260
Status Verification.....	261
Bulk Import and Export.....	261

Chapter 7: Policy and Charging Maintenance.....263

Introduction.....	264
Policy and Charging Maintenance Pages.....	264
SBR Database Status.....	264
SBR Status.....	266
Policy Database Query.....	267
SBR Database Reconfiguration Status.....	268
Alarms, KPIs, and Measurements.....	273
Policy and Charging and SBR Alarms and Events.....	273
PCA and SBR KPIs.....	274
Policy and Charging and SBR Measurements.....	274
Overload Management.....	274
Overload Controls.....	274
Shutdown.....	277
Diameter Maintenance and Status Data for Components, DSR Applications, and DA-MPs.....	278
Backup and Restore for Policy and Charging Configuration Data.....	279

Appendix A: PDRA PCRF Pooling Upgrade.....280

Upgrade Paths.....	281
Configuration After Upgrade.....	283
Concepts and Terminology.....	284
Configuring PCRF Pooling.....	285
Processing Phases.....	288

Binding Migration.....	290
------------------------	-----

Appendix B: PCA Error Resolution.....292

Introduction.....	293
Policy DRA Error Resolution Flowchart Summary.....	295
Diameter Message Validation Error Resolution Flowchart.....	297
Generic CCR Processing Error Resolution Flowchart.....	298
CCR-I Processing without PCRF Pool Error Resolution Flowchart.....	299
findSessionRef Processing Error Resolution Flowchart.....	301
findOrCreBindResShort Processing Error Resolution Flowchart.....	301
CCR-I Processing with PCRF Pool Error Resolution Flowchart.....	303
findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart.....	303
Early Bind Pool Error Resolution Flowchart.....	304
CCA-I Processing Error Resolution Flowchart.....	305
findSession Response Processing Error Resolution Flowchart.....	306
CCR-U Processing Error Resolution Flowchart.....	307
CCR-T Processing Error Resolution Flowchart.....	308
CCA-U/T Processing Error Resolution Flowchart.....	309
RAR Processing Error Resolution Flowchart.....	310
RAA Processing Error Resolution Flowchart.....	311
AAR Processing Error Resolution Flowchart.....	312
AAA Processing Error Resolution Flowchart.....	313
STR Processing Error Resolution Flowchart.....	314
STA Processing Error Resolution Flowchart.....	315
ASR/ASA Processing Error Resolution Flowchart.....	316
Online Charging DRA Error Resolution Flowchart Summary.....	317
Message Processing.....	319
Diameter Validation Processing.....	321
CCR Processing.....	321
CCR-I Processing.....	322
CCR-U Processing.....	324
CCR-T Processing.....	325
CCR-E Processing.....	326
RAR Processing.....	327
SBR SE Received Processing.....	328
CreateOcSessionResult SE Processing.....	329
FindAndRefreshOcSession SE Processing.....	329
FindAndRemoveOcSessionResult SE Processing.....	330
CCA Processing.....	331

CCA-I Processing.....	332
CCA-U Processing.....	334
CCA-T Processing.....	334
CCA-E Processing.....	335
RAA Processing.....	336
Glossary.....	338

List of Figures

Figure 1: Oracle System Login.....29

Figure 2: Paginated table32

Figure 3: Scrollable table.....32

Figure 4: Form page.....33

Figure 5: Tabbed pages.....33

Figure 6: Tabbed pages.....34

Figure 7: Report output.....34

Figure 8: Sorting a Table by Column Header.....35

Figure 9: Clear Field Control X.....36

Figure 10: Optional Layout Element Toolbar.....36

Figure 11: Automatic Error Notification.....37

Figure 12: Examples of Filter Styles.....38

Figure 13: Example of Gy/Ro Event-Based Charging.....49

Figure 14: Example of Gy/Ro Session-Based Charging.....50

Figure 15: Sites, Mated Pairs, and Region.....51

Figure 16: Policy DRA in Roaming Scenarios.....52

Figure 17: Example PCA Mated Pair - Hosting Binding SBRs.....54

Figure 18: Example PCA Mated Pair - Not Hosting Binding Policy SBRs.....55

Figure 19: Policy Client, PCRF, and Site Relationships.....56

Figure 20: Smallest Supported PCA Field Deployment.....67

Figure 21: Smallest Supported PCA Mated Pair.....68

Figure 22: Smallest Supported PCA Mated Triplets.....68

Figure 23: Comparing 2-Site and 3-Site Redundancy.....	71
Figure 24: Binding Table Partitioning Across Server Groups.....	72
Figure 25: Multi-Table Resources.....	72
Figure 26: Data Merging - Normal Case.....	73
Figure 27: Data merging - Redundant Site Failure.....	74
Figure 28: Communication between ComAgents, Policy DRA, and SBR.....	84
Figure 29: Request Processing at the Diameter Routing Function and PCA	85
Figure 30: Answer Processing at the Diameter Routing Function and PCA.....	85
Figure 31: PCA Generated Answer Routing.....	86
Figure 32: PCA Generated Request Routing.....	86
Figure 33: Event Diagram Trace - CCR Example.....	92
Figure 34: Update Request Policy DRA Example.....	93
Figure 35: Event Diagram - Hover (Mouse-over) Example.....	94
Figure 36: Relationship between APNs and PCRF Pools.....	100
Figure 37: Relationship between IMSIs and PCRF Pools.....	101
Figure 38: Multiple PCRF Pools.....	102
Figure 39: Multiple PCRF Versions in a PCRF Pool.....	103
Figure 40: PCRF Pools and Sub-Pools Routing Scenarios.....	104
Figure 41: Subscriber Key Usage.....	114
Figure 42: Find or Create a Binding.....	119
Figure 43: Binding Dependent Session Initiation Request Processing Overview.....	121
Figure 44: Local OCS Server Selection.....	131
Figure 45: Local OCS Server Pool Selection.....	132
Figure 46: Regionalized OCS Server Pool Selection.....	133
Figure 47: GUI Structure for 3-tiered DSR Topology with Policy and Charging for NOAM.....	145

Figure 48: GUI Structure for 3-tiered DSR Topology with Policy and Charging for SOAM.....	146
Figure 49: PCRF Pooling Data.....	175
Figure 50: PCRF Pooling Data.....	211
Figure 51: Error Resolution Flowchart Summary.....	297
Figure 52: Diameter Message Validation Error Resolution Flowchart.....	298
Figure 53: Generic CCR Processing Error Resolution Flowchart.....	299
Figure 54: CCR-I Processing without PCRF Pool Error Resolution Flowchart.....	300
Figure 55: findSessionRef Processing Error Resolution Flowchart.....	301
Figure 56: findOrCreBindResShort Processing Error Resolution Flowchart.....	302
Figure 57: CCR-I Processing with PCRF Pool Error Resolution Flowchart.....	303
Figure 58: findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart.....	304
Figure 59: Early Bind Pool Error Resolution Flowchart.....	305
Figure 60: CCA-I Processing Error Resolution Flowchart.....	306
Figure 61: findSession Response Processing Error Resolution Flowchart.....	307
Figure 62: CCR-U Processing Error Resolution Flowchart.....	308
Figure 63: CCR-T Processing Error Resolution Flowchart.....	309
Figure 64: CCA-U/T Processing Error Resolution Flowchart.....	310
Figure 65: RAR Processing Error Resolution Flowchart.....	311
Figure 66: RAA Processing Error Resolution Flowchart.....	312
Figure 67: AAR Processing Error Resolution Flowchart.....	313
Figure 68: AAA Processing Error Resolution Flowchart.....	314
Figure 69: STR Processing Error Resolution Flowchart.....	315
Figure 70: STA Processing Error Resolution Flowchart.....	316
Figure 71: ASR/ASA Processing Error Resolution Flowchart.....	317

Figure 72: Error Resolution Flowchart Summary.....	319
Figure 73: Message Processing.....	320
Figure 74: Diameter Validation Processing.....	321
Figure 75: CCR Processing.....	322
Figure 76: CCR-I Processing.....	323
Figure 77: CCR-U Processing.....	324
Figure 78: CCR-T Processing.....	325
Figure 79: CCR-E Processing.....	326
Figure 80: RAR Processing.....	327
Figure 81: SBR SE Received Processing.....	328
Figure 82: CreateOcSessionResult Processing.....	329
Figure 83: FindAndRefreshOcSession SE Processing.....	330
Figure 84: FindAndRemoveOcSessionResult Processing.....	331
Figure 85: CCA Processing.....	332
Figure 86: CCA-I Processing.....	333
Figure 87: CCA-U Processing.....	334
Figure 88: CCA-T Processing.....	335
Figure 89: CCA-E Processing.....	336
Figure 90: RAA Processing.....	337

List of Tables

Table 1: Admonishments.....	18
Table 2: User Interface elements.....	23
Table 3: Main Menu Options.....	24
Table 4: Main Menu icons.....	31
Table 5: Example Action buttons.....	35
Table 6: Submit buttons.....	36
Table 7: Filter Control elements.....	38
Table 8: Example of a Binding Key Priority Configuration.....	47
Table 9: Server Group Functions.....	57
Table 10: SBR Server Group Configuration and Data Redundance.....	57
Table 11: Client Connection Capability.....	59
Table 12: IP Traffic-to-Service Mapping.....	74
Table 13: P-DRA Application Routing Table Configuration.....	76
Table 14: Communication between the Diameter Routing Function and the DAI.....	85
Table 15: PCA Metadata-Generating Events.....	87
Table 16: Policy DRA Terminology.....	105
Table 17: Example Key Priority Configuration.....	114
Table 18: Topology Hiding Scope Configuration.....	123
Table 19: Policy DRA Error Scenarios for Session Integrity.....	125
Table 20: Session Integrity Conditions and Policy DRA Reaction.....	126
Table 21: Session State Configuration Settings.....	134
Table 22: Diameter AVPs used by OC-DRA for Request Message Processing.....	136

Table 23: Diameter AVPs used by OC-DRA for Answer Message Processing.....	140
Table 24: General Options elements.....	156
Table 25: SBR Databases elements.....	158
Table 26: Access Point Names elements.....	161
Table 27: SBR Database Resizing Plans elements.....	166
Table 28: SBR Data Migration Plans elements.....	170
Table 29: PCRF Pooling Concepts.....	174
Table 30: PCRF Pooling Configuration Summary.....	175
Table 31: PCRF Pools elements.....	177
Table 32: PCRF Sub-Pool Selection Rules elements.....	181
Table 33: Policy DRA Network-Wide Options elements.....	186
Table 34: Realms elements.....	191
Table 35: Online Charging DRA Network-Wide Options elements.....	193
Table 36: Alarm Settings elements.....	197
Table 37: Congestion Options elements.....	201
Table 38: PCRFs page elements.....	205
Table 39: Binding Key Priority elements.....	208
Table 40: PCRF Pooling Concepts.....	210
Table 41: PCRF Pooling Configuration Summary.....	211
Table 42: PCRF Pools elements.....	213
Table 43: PCRF Pools to PRT Mapping elements.....	216
Table 44: PCRF Sub-Pool Selection Rules elements.....	221
Table 45: Policy Clients elements.....	226
Table 46: Suspect Binding Removal Rules elements.....	229
Table 47: Site Options elements.....	233

Table 48: OCSs elements.....	235
Table 49: CTFs elements.....	238
Table 50: Realms elements.....	242
Table 51: PCA Error Conditions.....	244
Table 52: Interfaces Supported for Each Error Code.....	250
Table 53: Error Codes elements.....	250
Table 54: Alarm Settings elements.....	253
Table 55: Congestion Options elements.....	257
Table 56: SBR Database Status Pane elements.....	265
Table 57: SBR Status elements.....	267
Table 58: Policy Database Query elements.....	268
Table 59: Reconfiguration Status Pane elements.....	269
Table 60: PCA Default Overload Control Thresholds.....	275
Table 61: Diameter Routing Function Message Handling Based on PCA Operational Status.....	276
Table 62: Stack Event Load Shedding.....	277
Table 63: PCA Operational Status.....	278
Table 64: Processing During Transition Period.....	282
Table 65: Upgrading Policy DRA Terminology.....	285
Table 66: Error Resolution Attributes.....	293

Chapter 1

Introduction

Topics:

- *Overview.....17*
- *Scope and Audience.....17*
- *Manual Organization.....17*
- *Documentation Admonishments.....18*
- *Related Publications.....18*
- *Locate Product Documentation on the Oracle Help Center Site.....18*
- *Customer Training.....19*
- *My Oracle Support (MOS).....19*
- *Emergency Response.....20*

This chapter contains a brief description of the Policy and Charging Application, consisting of the Policy DRA and Online Charging DRA functions. The contents include sections about the document scope, audience, and organization; how to find related publications; and how to contact customer assistance.

Overview

The *Policy and Charging Application User's Guide* provides a conceptual overview of the Policy and Charging Application's purpose, architecture, and functionality.

It also describes the pages and elements on the PCA GUI (Graphical User Interface), as well as procedures for using the PCA interface.

Scope and Audience

This document is intended for anyone responsible for configuring and using the DSR Policy and Charging application and Session Binding Repository. Users of this manual must have a working knowledge of telecommunications and network installations.

Manual Organization

This manual is organized into the following chapters:

- [Introduction](#) contains general information about the DSR documentation, the organization of this manual, and how to get technical assistance.
- [User Interface Introduction](#) describes the organization and usage of the application user interface, including information about how the interface options are organized, how to use widgets and buttons, and how filtering and other page display options work.
- [Policy and Charging Introduction](#) describes the topology, architecture, components, and functions of the Policy and Charging application and the Session Binding Repository (SBR).
- [Policy DRA Overview](#) describes an overview of the Policy DRA feature and includes information about important fundamental concepts, as well as high-level functionality, including Pools and Sub-Pools.
- [Online Charging DRA Overview](#) describes an overview of the Online Charging DRA feature and includes information about important fundamental concepts, as well as high-level functionality, including OCSs and CTFs.
- [Policy and Charging Configuration](#) describes configuration of PCA application components.
- [Policy and Charging Maintenance](#) describes PCA Maintenance functions, and Diameter Maintenance functions that provide maintenance and status information for PCA and the SBR.
- [PDRA PCRF Pooling Upgrade](#) describes how to upgrade for PCRF Pooling from DSR Policy DRA release 5.0 to the 7.0 functionality.
- [PCA Error Resolution](#) describes information to help users diagnose and resolve PCA errors encountered while processing Diameter messages in the PCA application.

Documentation Admonishments

Admonishments are icons and text throughout this manual that alert the reader to assure personal safety, to minimize possible service interruptions, and to warn of the potential for equipment damage.

Table 1: Admonishments

Icon	Description
 DANGER	Danger: (This icon and text indicate the possibility of <i>personal injury</i> .)
 WARNING	Warning: (This icon and text indicate the possibility of <i>equipment damage</i> .)
 CAUTION	Caution: (This icon and text indicate the possibility of <i>service interruption</i> .)
 TOPPLE	Topple: (This icon and text indicate the possibility of <i>personal injury and equipment damage</i> .)

Related Publications

For information about additional publications that are related to this document, refer to the *Related Publications Reference* document, which is published as a separate document on the Oracle Technology Network (OTN) site. See [Locate Product Documentation on the Oracle Help Center Site](#) for more information.

Locate Product Documentation on the Oracle Help Center Site

Oracle Communications customer documentation is available on the web at the Oracle Help Center (OHC) site, <http://docs.oracle.com>. You do not have to register to access these documents. Viewing these files requires Adobe Acrobat Reader, which can be downloaded at <http://www.adobe.com>.

1. Access the Oracle Help Center site at <http://docs.oracle.com>.
2. Click **Industries**.

3. Under the Oracle Communications subheading, click the **Oracle Communications documentation** link.
The Communications Documentation page appears. Most products covered by these documentation sets will appear under the headings “Network Session Delivery and Control Infrastructure” or “Platforms.”
4. Click on your Product and then the Release Number.
A list of the entire documentation set for the selected product and release appears.
5. To download a file to your location, right-click the **PDF** link, select **Save target as** (or similar command based on your browser), and save to a local folder.

Customer Training

Oracle University offers training for service providers and enterprises. Visit our web site to view, and register for, Oracle Communications training:

<http://education.oracle.com/communication>

To obtain contact phone numbers for countries or regions, visit the Oracle University Education web site:

www.oracle.com/education/contacts

My Oracle Support (MOS)

MOS (<https://support.oracle.com>) is your initial point of contact for all product support and training needs. A representative at Customer Access Support (CAS) can assist you with MOS registration.

Call the CAS main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. When calling, make the selections in the sequence shown below on the Support telephone menu:

1. Select **2** for New Service Request
2. Select **3** for Hardware, Networking and Solaris Operating System Support
3. Select one of the following options:
 - For Technical issues such as creating a new Service Request (SR), Select **1**
 - For Non-technical issues such as registration or assistance with MOS, Select **2**

You will be connected to a live agent who can assist you with MOS registration and opening a support ticket.

MOS is available 24 hours a day, 7 days a week, 365 days a year.

Emergency Response

In the event of a critical service situation, emergency response is offered by the Customer Access Support (CAS) main number at 1-800-223-1711 (toll-free in the US), or by calling the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. The emergency response provides immediate coverage, automatic escalation, and other features to ensure that the critical situation is resolved as rapidly as possible.

A critical situation is defined as a problem with the installed equipment that severely affects service, traffic, or maintenance capabilities, and requires immediate corrective action. Critical situations affect service and/or system operation resulting in one or several of these situations:

- A total system failure that results in loss of all transaction processing capability
- Significant reduction in system capacity or traffic handling capability
- Loss of the system's ability to perform automatic system reconfiguration
- Inability to restart a processor or the system
- Corruption of system databases that requires service affecting corrective actions
- Loss of access for maintenance or recovery operations
- Loss of the system ability to provide any required critical or major trouble notification

Any other problem severely affecting service, capacity/traffic, billing, and maintenance capabilities may be defined as critical by prior discussion and agreement with Oracle.

Chapter 2

User Interface Introduction

Topics:

- [User Interface Organization.....22](#)
- [Missing Main Menu options.....28](#)
- [Common Graphical User Interface Widgets.....28](#)

This section describes the organization and usage of the application's user interface. In it you can find information about how the interface options are organized, how to use widgets and buttons, and how filtering and other page display options work.

User Interface Organization

The user interface is the central point of user interaction within an application. It is a Web-based graphical user interface (GUI) that enables remote user access over the network to an application and its functions.

The core framework presents a common set of Main Menu options that serve various applications. The common Main Menu options are:

- Administration
- Configuration
- Alarm & Events
- Security Log
- Status & Manage
- Measurements
- Help
- Legal
- Logout

Applications, such as DSR, build upon this framework to present features and functions. For example, the DSR Network OAM GUI may present the following Main Menu options in addition to the common options:

- Communication Agent
- Diameter Common
- Diameter
- Policy and Charging
- MAP-Diameter IWF

The DSR System OAM GUI may present even more Main Menu options as listed below. The end result is a flexible menu structure that changes according to the application needs and features activated.

- Transport Manager
- SS7/Sigtran
- RBAR
- FABR
- IPFE
- GLA

Note that the DSR System OAM Main Menu options differ from the Network OAM options. Some Main Menu options are configurable from the DSR Network OAM server and view-only from the System OAM server. This remains true for other applications.

User Interface Elements

Table 2: User Interface elements describes elements of the user interface.

Table 2: User Interface elements

Element	Location	Function
Identification Banner	Top bar across the web page	Displays the company name, product name and version, and the alarm panel.
Session Banner	Next bar across the top of the web page	<p>The left side of the banner just above the Main Menu provides the following session information:</p> <ul style="list-style-type: none"> • The name of the machine to which the user is connected, and whether the user is connected via the VIP or directly to the machine. • The HA state of the machine to which the user is connected. • The role of the machine to which the user is connected. <p>The right side of the banner:</p> <ul style="list-style-type: none"> • Shows the user name of the currently logged-in user. • Provides a link to log out of the GUI.
Main Menu	Left side of screen, under banners	<p>A tree-structured menu of all operations that can be performed through the user interface. The plus character (+) indicates a menu item contains subfolders.</p> <ul style="list-style-type: none"> • To display submenu items, click the plus character, the folder, or anywhere on the same line. • To select a menu item that does not have submenu items, click on the menu item text or its associated symbol.
Work Area	Right side of panel under status	<p>Consists of three sections: Page Title Area, Page Control Area (optional), and Page Area.</p> <ul style="list-style-type: none"> • Page Title Area: Occupies the top of the work area. It displays the title of the current page being displayed, date and time, and includes a link to context-sensitive help. • Page Control Area: Located below the Page Title Area, this area shows controls for the Page Area (this area is optional). When available as an option, filter controls display in this area. The Page Control Area contains the optional layout element toolbar, which displays different elements depending on which GUI page is selected. For more information, see Optional Layout Element Toolbar. • Page Area: Occupies the bottom of the work area. This area is used for all types of operations. It displays all options, status, data, file, and query screens. Information or error messages are displayed in a message box at the top of this section. A horizontal and/or vertical scroll bar is provided when the displayed information exceeds the page area of the screen. When a user first logs in, this area displays the application user interface page.

Element	Location	Function
		The page displays a user-defined welcome message. To customize the message, see Customizing the Splash Page Welcome Message .

Main Menu Options

Table 3: Main Menu Options describes all main menu user interface options.

Note: The menu options can differ according to the permissions assigned to a user's log-in account. For example, the Administration menu options do not appear on the screen of a user who does not have administrative privileges.

Note: Some menu items are configurable only on the Network OAM and view-only on the System OAM; and some menu options are configurable only on the System OAM.

Note: Some features do not appear in the main menu until the features are activated.

Table 3: Main Menu Options

Menu Item	Function
Administration	<p>The Administration menu allows the user to:</p> <ul style="list-style-type: none"> • General Options. Configure options such as password history and expiration, login message, welcome message, and the number of failed login attempts before an account is disabled • Set up and manage user accounts • Configure group permissions • View session information • Manage sign-on certificates • Authorize IP addresses to access the user interface • Configure SFTP user information • View the software versions report • Upgrade management including backup and reporting • Authenticate LDAP servers • Configure SNMP trapping services • Configure an export server • Configure DNS elements
Configuration	<p>On the NOAM, allows the user to configure:</p> <ul style="list-style-type: none"> • Network Elements • Network Devices • Network Routes • Services • Servers • Server Groups • Resource Domains • Places

Menu Item	Function
	<ul style="list-style-type: none"> • Place Associations • Interface and Port DSCP
Alarms and Events	<p>Allows the user to view:</p> <ul style="list-style-type: none"> • Active alarms and events • Alarm and event history • Trap log
Security Log	<p>Allows the user to view, export, and generate reports from security log history.</p>
Status & Manage	<p>Allows the user to monitor the individual and collective status of Network Elements, Servers, HA functions, Databases, KPIs, system Processes, and Tasks. The user can perform actions required for server maintenance, database management, data, and ISO file management.</p>
Measurements	<p>Allows the user to view and export measurement data.</p>
Transport Manager (optional)	<p>On the SOAM, allows the user to configure adjacent nodes, configuration sets, or transports. A maintenance option allows the user to perform enable, disable, and block actions on the transport entries.</p>
Communication Agent (optional)	<p>Allows the user to configure Remote Servers, Connection Groups, and Routed Services. The user can perform actions to enable, disable, and block connections. Also allows the user to monitor the status of Connections, Routed Services, and HA Services.</p>
SS7/Sigtran (optional)	<p>On the SOAM, allows the user to configure various users, groups, remote signaling points, links, and other items associated with SS7/Sigtran; perform maintenance and troubleshooting activities; and provides a command line interface for bulk loading SS7 configuration data.</p>
Diameter Common (optional)	<p>Allows the user to view or configure:</p> <ul style="list-style-type: none"> • Dashboard. Configure on the NOAM; view on both OAMs • Network Identifiers on the SOAM - MCC Ranges • Network Identifiers on the NOAM - MCCMNC and MCCMNC Mapping • MPs (on the SOAM) - editable Profile parameters and Profile assignments <p>The DSR Bulk Import and Export functions are available on both OAMs for the data that is configured on that OAM.</p>
Diameter (optional)	<p>Allows the user to configure, modify, and monitor Diameter routing:</p> <ul style="list-style-type: none"> • On the NOAMP, Diameter Topology Hiding configuration and Egress Throttle List • On the SOAM, Diameter Configuration, Maintenance, Reports, Troubleshooting with IDIH, AVP Dictionary, and Diameter Mediation configuration

Menu Item	Function
RBAR (Range-Based Address Resolution) (optional)	Allows the user to configure the following Range-Based Address Resolution (RBAR) settings: <ul style="list-style-type: none"> • Applications • Exceptions • Destinations • Address Tables • Addresses • Address Resolutions • System Options This is accessible from the SOAM only.
FABR (Full Address Based Resolution) (optional)	Allows the user to configure the following Full Address Based Resolution (FABR) settings: <ul style="list-style-type: none"> • Applications • Exceptions • Default Destinations • Address Resolutions • System Options This is accessible from the SOAM only.
Policy and Charging (optional)	On the NOAMP, allows the user to perform configuration tasks, edit options, and view elements for: <ul style="list-style-type: none"> • General Options • SBR Databases • Access Point Names • SBR Database Resizing Plans • SBR Data Migration Plans • Policy DRA <ul style="list-style-type: none"> • PCRF Pools • PCRF Sub-Pool Selection Rules • Network-Wide Options • Online Charging DRA <ul style="list-style-type: none"> • OCS Session State • Realms • Network-Wide Options • Alarm Settings • Congestion Options Additionally on the NOAMP, users are allowed to perform maintenance tasks, edit options, and view elements for: <ul style="list-style-type: none"> • Maintenance <ul style="list-style-type: none"> • SBR Database Status

Menu Item	Function
	<ul style="list-style-type: none"> • SBR Status • Policy Database Query • SBR Database Reconfiguration Status <p>On the SOAM, allows the user to perform configuration tasks, edit options, and view elements for:</p> <ul style="list-style-type: none"> • General Options • SBR Databases • Access Point Names • SBR Database Resizing Plans • SBR Data Migration Plans • Policy DRA <ul style="list-style-type: none"> • PCRFs • Binding Key Priority • PCRF Pools • PCRF Pool to PRT Mapping • PCRF Sub-Pool Selection Rules • Policy Clients • Suspect Binding Removal Rules • Site Options • Online Charging DRA <ul style="list-style-type: none"> • OCSs • CTFs • OCS Session State • Realms • Error Codes • Alarm Settings • Congestion Options
Gateway Location Application (optional)	<p>On the SOAM, allows the user to perform configuration tasks, edit options, and view elements for:</p> <ul style="list-style-type: none"> • Exceptions • Options <p>GLA can deploy with Policy DRA (in the same DA-MP or a separate DA-MP).</p>
IPFE (optional)	<p>Allows the user to configure IP Front End (IPFE) options and IP List TSAs. This is accessible from the SOAM server only.</p>
MAP-Diameter Interworking	<p>On the SOAM, allows the user to perform configuration tasks, edit options, and view elements for the DM-IWF DSR Application:</p> <ul style="list-style-type: none"> • DM-IWF Options • Diameter Exception

Menu Item	Function
	On the NOAMP, allows the user to perform configuration tasks, edit options, and view elements for the MD-IWF SS7 Application: <ul style="list-style-type: none"> • MD-IWF Options • Diameter Realm • Diameter Identity GTA • GTA Range to PC • MAP Exception • CCNDC Mapping
CPA (Charging Proxy Application) (optional)	Allows the user to perform configuration tasks, edit system options, and view elements for: <ul style="list-style-type: none"> • System Options • Message Copy • Session Binding Repository • SBR Subresource Mapping This is accessible from the SOAM only.
Help	Launches the Help system for the user interface
Legal Notices	Product Disclaimers and Notices
Logout	Allows the user to log out of the user interface

Missing Main Menu options

Permissions determine which Main Menu options are visible to users. Permissions are defined through the **Group Administration** page. The default group, **admin**, is permitted access to all GUI options and functionality. Additionally, members of the **admin** group set permissions for other users.

Main Menu options vary according to the group permissions assigned to a user's account. Depending on your user permissions, some menu options may be missing from the Main Menu. For example, Administration menu options do not appear on your screen if you do not have administrative permissions. For more information about user permissions, see *Group Administration* in the OAM section of the online help, or contact your system administrator.

Common Graphical User Interface Widgets

Common controls allow you to easily navigate through the system. The location of the controls remains static for all pages that use the controls. For example, after you become familiar with the location of the display filter, you no longer need to search for the control on subsequent pages because the location is static.

Supported Browsers

This application supports the use of Microsoft® Internet Explorer 8.0, 9.0, or 10.0.

System Login Page

Access to the user interface begins at the System Login page. The System Login page allows users to log in with a username and password and provides the option of changing a password upon login. The System Login page also features a current date and time stamp and a customizable login message.

The user interface is accessed via HTTPS, a secure form of the HTTP protocol. When accessing a server for the first time, HTTPS examines a web certificate to verify the identity of the server. The configuration of the user interface uses a self-signed web certificate to verify the identity of the server. When the server is first accessed, the supported browser warns the user that the server is using a self-signed certificate. The browser requests confirmation that the server can be trusted. The user is required to confirm the browser request.

Customizing the Login Message

Before logging in, the **System Login** page appears. You can create a login message that appears just below the **Log In** button on the **System Login** page.



ORACLE®

Oracle System Login Wed Jul 8 14:20:00 2015 EDT

Log In

Enter your username and password to log in

Username:

Password:

Change password

Welcome to the Oracle System Login.

Unauthorized access is prohibited. This Oracle system requires the use of Microsoft Internet Explorer 8.0, 9.0, or 10.0 with support for JavaScript and cookies.

Oracle and Java are registered trademarks of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Copyright © 2010, 2015, [Oracle](#) and/or its affiliates. All rights reserved.

Figure 1: Oracle System Login

1. From the **Main Menu**, select **Administration > General Options**.

The **General Options Administration** page appears.

2. Locate **LoginMessage** in the **Variable** column.
3. Enter the login message text in the **Value** column.
4. Click **OK** or **Apply** to submit the information.

A status message appears at the top of the Configuration Administration page to inform you if the operation was successful.

The next time you log in to the user interface, the login message text is displayed.

Accessing the DSR Graphical User Interface

In a DSR, some configuration is done at the NOAM server, while some is done at the SOAM server. Because of this, you will access the DSR graphical user interface (GUI) from two servers. Certificate Management (Single Sign-On) can be configured to simplify accessing the DSR GUI on the NOAM and the SOAM.

For information on configuring Single Sign-On certificates, see **OAM > Administration > Access Control > Certificate Management** in the DSR online help.

After the certificates have been configured, you can log into the DSR GUI on any NOAM or SOAM, and then access the DSR GUI on other servers (NOAM or other SOAMs) without having to re-enter your login credentials.

1. In the browser URL field, enter the fully qualified hostname of the NOAM server, for example `https://dsr-no.yourcompany.com`.

When using Single Sign-On, you cannot use the IP address of the server.

2. When prompted by the browser, confirm that the server can be trusted.

The System Login page appears.

3. Enter the Username and Password for your account.

The DSR GUI for the NOAM appears.

4. To access the DSR GUI for the SOAM, open another browser window and enter the fully qualified hostname of the SOAM.

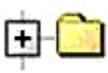
The DSR GUI for the SOAM appears

You can toggle between the DSR GUI on the NOAM and the DSR GUI on the SOAM as you perform configuration tasks.

Main Menu Icons

This table describes the icons used in the **Main Menu**.

Table 4: Main Menu icons

Icon	Name	Description
	Folder	Contains a group of operations. If the folder is expanded by clicking the plus (+) sign, all available operations and sub-folders are displayed. Clicking the minus (-) collapses the folder.
	Config File	Contains operations in an Options page.
	File with Magnifying Glass	Contains operations in a Status View page.
	File	Contains operations in a Data View page.
	Multiple Files	Contains operations in a File View page.
	File with Question Mark	Contains operations in a Query page.
	User	Contains operations related to users.
	Group	Contains operations related to groups.
	Help	Launches the Online Help.
	Logout	Logs the user out of the user interface.

Work Area Displays

In the user interface, tables, forms, tabbed pages, and reports are the most common formats.

Note: Screenshots are provided for reference only and may not exactly match a specific application's GUI.

Tables

Paginated tables describe the total number of records being displayed at the beginning and end of the table. They provide optional pagination with **First** | **Prev** | **Next** | **Last** links at both the beginning and

end of this table type. Paginated tables also contain action links on the beginning and end of each row. For more information on action links and other page controls, see [Page Controls](#).

Displaying Records 1-1 of 1 | [First](#) | [Prev](#) | [Next](#) | [Last](#)

Action	System ID	IP Address	Permission	Action
Edit Delete	lisa	10.25.62.4	READ_WRITE	Edit Delete

Displaying Records 1-1 of 1 | [First](#) | [Prev](#) | [Next](#) | [Last](#)

Figure 2: Paginated table

Scrollable tables display all of the records on a single page. The scroll bar, located on the right side of the table, allows you to view all records in the table. Scrollable tables also provide action buttons that operate on selected rows. For more information on buttons and other page controls, see [Page Controls](#).

Sequence #	Alarm ID	Timestamp	Severity	Product	Process	NE	Server	Type	Instance	Alarm Text
3498	31201	2009-Jun-11 18:07:41.214 UTC	MAJOR	MiddleWare	procmgr	OAMPNE	teks8011006	PROC	eclipseHelp	A managed process cannot be started or has unexpectedly terminated
5445	31201	2009-Jun-11 18:07:27.137 UTC	MAJOR	MiddleWare	procmgr	SOAMP	teks8011002	PROC	eclipseHelp	A managed process cannot be started or has unexpectedly terminated
5443	31107	2009-Jun-11 18:07:24.704 UTC	MINOR	MiddleWare	inetmerge	SOAMP	teks8011002	COLL	teks8011004	DB merging from a child Source Node has failed
5444	31107	2009-Jun-11 18:07:24.704 UTC	MINOR	MiddleWare	inetmerge	SOAMP	teks8011002	COLL	teks8011003	DB merging from a child Source Node has failed
5441	31209	2009-Jun-11 18:07:22.640 UTC	MINOR	MiddleWare	re.portmap	SOAMP	teks8011002	SW	teks8011003	Unable to resolve a hostname specified in the NodeInfo table.
										Unable to resolve a hostname specified in the NodeInfo table.

[Export](#)

Figure 3: Scrollable table

Note: Multiple rows can be selected in a scrollable table. Add rows one at a time using CTRL-click. Add a span of rows using SHIFT-click.

Forms

Forms are pages on which data can be entered. Forms are typically used for configuration. Forms contain fields and may also contain a combination of pulldown lists, buttons, and links.

Username: (5-16 characters)

Group: ▼

Time Zone: ▼

Maximum Concurrent Logins: Maximum concurrent logins for a user (0=no limit).
[Default = 1; Range = 0-50]

Session Inactivity Limit: Time (in minutes) after which login sessions expire (0 = never).
[Default = 120; Range = 0-120]

Comment: (max 64 characters)

Temporary Password: (8-16 characters)

Re-type Password: (8-16 characters)

Figure 4: Form page

Tabbed pages

Tabbed pages provide collections of data in selectable tabs. Click on a tab to see the relevant data on that tab. Tabbed pages also group Retrieve, Add, Update, and Delete options on one page. Click on the relevant tab for the task you want to perform and the appropriate fields populate on the page. Retrieve is always the default for tabbed pages.

Entire Network *		System.CPU_CoreUtilPct_Average		System.CPU_CoreUtilPct_Peak		
NOAMP	Timestamp	System CPU UtilPct Average	System CPU UtilPct Peak	System Disk UtilPct Average	System Disk UtilPct Peak	System RAM UtilPct Average
SOAM	10/22/2009 19:45	6.764068	44	0.520000	1	7.939407
	10/22/2009 20:00	7.143644	25	0.520000	1	8.523822

Figure 5: Tabbed pages

Retrieve Add Update Delete

Fields marked with a red asterisk (*) require a value.

Field	Value	Description
Network Entity	<input type="text"/>	* Numeric identifier for the Network Entity 1-15 DIGITS

Retrieve

Figure 6: Tabbed pages

Reports

Reports provide a formatted display of information. Reports are generated from data tables by clicking the **Report** button. Reports can be viewed directly on the user interface, or they can be printed. Reports can also be saved to a text file.

```

=====
User Account Usage Report
=====
Report Generated: Fri Jun 19 19:30:55 2009 UTC
From: Unknown Network OAM&P on host teks5001701
Report Version: 1.0
User: guiadmin
-----
Username      Date of Last Login  Days Since Last Login  Account Status
-----
guiadmin      2009-06-19 19:00:17  0                       enabled
-----
End of User Account Usage Report
=====

```

Figure 7: Report output

Customizing the Splash Page Welcome Message

When you first log in to the user interface, the **User Interface** splash page appears. You can display a customized welcome message on the **User Interface** splash page using this procedure.

1. From the **Main Menu**, select **Administration > General Options**.
The **General Options Administration** page appears.
2. Locate **WelcomeMessage** in the **Variable** column.
3. Enter the welcome message text in the **Value** column.

4. Click **Update**, **OK**, or **Apply** to submit the information.

A status message appears at the top of the Configuration Administration page to inform you if the operation was successful.

The next time you log in to the user interface, the welcome message text is displayed.

Column Headers (Sorting)

You can sort a table by a column by clicking the column header. However, sorting is not necessarily available on every column. Sorting does not affect filtering.

When you click the header of a column that the table can be sorted by, an indicator appears in the column header showing the direction of the sort. See [Figure 8: Sorting a Table by Column Header](#). Clicking the column header again reverses the direction of the sort.

Local Node Name	Realm	FQDN	SCTP Listen Port	TCP Listen Port	Connection Configuration Set	CEX Configuration Set	IP Addresses
-----------------	-------	------	------------------	-----------------	------------------------------	-----------------------	--------------

Figure 8: Sorting a Table by Column Header

Page Controls

User interface pages contain controls, such as buttons and links, that perform specified functions. The functions are described by the text of the links and buttons.

Note: Disabled buttons are grayed out. Buttons that are irrelevant to the selection or current system state, or which represent unauthorized actions as defined in **Group Administration**, are disabled. For example, **Delete** is disabled for users without Global Data Delete permission. Buttons are also disabled if, for example, multiple servers are selected for an action that can only be performed on a single server at a time.

[Table 5: Example Action buttons](#) contains examples of Action buttons.

Table 5: Example Action buttons

Action button	Function
Insert	Inserts data into a table.
Edit	Edits data within a table.
Delete	Deletes data from table.
Change	Changes the status of a managed object.

Some Action buttons take you to another page.

Submit buttons, described in [Table 6: Submit buttons](#), are used to submit information to the server. The buttons are located in the page area and accompanied by a table in which you can enter information. The Submit buttons, except for **Cancel**, are disabled until you enter some data or select a value for all mandatory fields.

Table 6: Submit buttons

Submit button	Function
OK	Submits the information to the server, and if successful, returns to the View page for that table.
Apply	Submits the information to the server, and if successful, remains on the current page so that you can enter additional data.
Cancel	Returns to the View page for the table without submitting any information to the server.

Clear Field Control

The clear field control allows you to clear the value from a pulldown list. The clear field control is available only on some pulldown fields.

Click the X next to a pulldown list to clear the field.



Figure 9: Clear Field Control X

Optional Layout Element Toolbar

The optional layout element toolbar appears in the Page Control Area of the GUI.

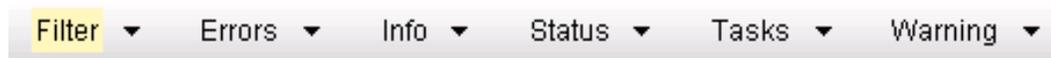


Figure 10: Optional Layout Element Toolbar

The toolbar displays different elements depending on which GUI page is selected. The elements of the toolbar that can appear include:

- Filter – Allows you to filter data in a table.
- Errors – Displays errors associated with the work area.
- Info – Displays information messages associated with the work area.
- Status – Displays short status updates associated with the main work area.
- Warning – Displays warnings associated with the work area.

Notifications

Some messages require immediate attention, such as errors and status items. When new errors occur, the Errors element opens automatically with information about the error. Similarly, when new status items are added, the Status element opens. If you close an automatically opened element, the element stays closed until a new, unacknowledged item is added.

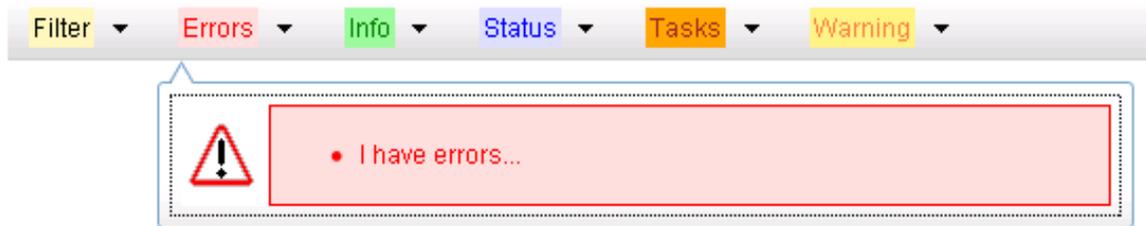


Figure 11: Automatic Error Notification

Note: Viewing and closing an error does not clear the Errors element. If you reopen the Errors element, previously viewed errors are still in the list.

When new messages are added to Warning or Info, the styling of the element changes to indicate new messages are available. The styling of the Task element changes when a task changes state (such as, a task begins or ends).

Opening an Element in the Toolbar

Use this procedure to open an element in the optional layout element toolbar.

1. Click the text of the element or the triangle icon to open an element.
The selected element opens and overlays the work area.
2. Click **X** to close the element display.

Filters

Filters are part of the optional layout element toolbar and appear throughout the GUI in the Page Control Area. For more information about optional layout element toolbar functionality, see [Optional Layout Element Toolbar](#).

Filters allow you to limit the data presented in a table and can specify multiple filter criteria. By default, table rows appear unfiltered. Three types of filters are supported, however, not all filtering options are available on every page. The types of filters supported include:

- Network Element – When enabled, the Network Element filter limits the data viewed to a single Network Element.
Note: Once enabled, the Network Element filter will affect all pages that list or display data relating to the Network Element.
- Collection Interval – When enabled, the collection interval filter limits the data to entries collected in a specified time range.
- Display Filter – The display filter limits the data viewed to data matching the specified criteria.

Once a field is selected, it cannot be selected again. All specified criteria must be met in order for a row to be displayed.

The style or format of filters may vary depending on which GUI pages the filters are displayed. Regardless of appearance, filters of the same type function the same.

Figure 12: Examples of Filter Styles

Filter Control Elements

This table describes filter control elements of the user interface.

Table 7: Filter Control elements

Operator	Description
=	Displays an exact match.
!=	Displays all records that do not match the specified filter parameter value.
>	Displays all records with a parameter value that is greater than the specified value.
>=	Displays all records with a parameter value that is greater than or equal to the specified value.
<	Displays all records with a parameter value that is less than the specified value.
<=	Displays all records with a parameter value that is less than or equal to the specified value.
Like	Enables you to use an asterisk (*) as a wildcard as part of the filter parameter value.
Is Null	Displays all records that have a value of Is Null in the specified field.

Note: Not all filterable fields support all operators. Only the supported operators will be available for you to select.

Filtering on the Network Element

The global Network Element filter is a special filter that is enabled on a per-user basis. The global Network Element filter allows a user to limit the data viewed to a single Network Element. Once enabled, the global Network Element filter affects all sub-screens that display data related to Network Elements. This filtering option may not be available on all pages.

1. Click **Filter** in the optional layout element toolbar.
The filter tool appears.
2. Select a Network Element from the **Network Element** pulldown menu.
3. Click **Go** to filter on the selection, or click **Reset** to clear the selection.

Records are displayed according to the specified criteria.

Filtering on Collection Interval

The Collection Interval filter allows a user to limit the data viewed to a specified time interval. This filtering option may not be available on all pages.

1. Click **Filter** in the optional layout element toolbar.
The filter tool appears.
2. Enter a duration for the **Collection Interval** filter.
The duration must be a numeric value.
3. Select a unit of time from the pulldown menu.
The unit of time can be seconds, minutes, hours, or days.
4. Select **Beginning** or **Ending** from the pulldown menu.
5. Click **Go** to filter on the selection, or click **Reset** to clear the selection.

Records are displayed according to the specified criteria.

Filtering using the Display Filter

Use this procedure to perform a filtering operation. This procedure assumes you have a data table displayed on your screen. This process is the same for all data tables. However, all filtering operations are not available for all tables.

1. Click **Filter** in the optional layout element toolbar.
The filter tool appears.
2. Select a field name from the **Display Filter** pulldown menu.
This selection specifies the field in the table that you want to filter on. The default is **None**, which indicates that you want all available data displayed.
The selected field name displays in the **Display Filter** field.
3. Select an operator from the operation selector pulldown menu.
The selected operator appears in the field.
4. Enter a value in the value field.
This value specifies the data that you want to filter on. For example, if you specify Filter=Severity with the equals (=) operator and a value of MINOR, the table would show only records where Severity=MINOR.
5. For data tables that support compound filtering, click **Add** to add another filter condition. Then repeat steps 2 through 4.
Multiple filter conditions are joined by an AND operator.
6. Click **Go** to filter on the selection, or click **Reset** to clear the selection.

Records are displayed according to the specified criteria.

Pause Updates

Some pages refresh automatically. Updates to these pages can be paused by selecting the **Pause updates** checkbox. Uncheck the **Pause updates** checkbox to resume automatic updates. The **Pause updates** checkbox is available only on some pages.

Max Records Per Page Controls

Max Records Per Page is used to control the maximum number of records displayed in the page area. If a page uses pagination, the value of Max Records Per Page is used. Use this procedure to change the Max Records Per Page.

1. From the **Main Menu**, select **Administration > General Options**.

The **General Options Administration** page appears.

2. Change the value of the **MaxRecordsPerPage** variable.

Note: **MaxRecordsPerPage** has a range of values from 10 to 100 records. The default value is 20.

3. Click **OK** or **Apply**.

OK saves the change and returns to the previous page.

Apply saves the change and remains on the same page.

The maximum number of records displayed is changed.

Chapter 3

Policy and Charging Introduction

Topics:

- *Overview.....42*
- *The P-DRA Database.....43*
- *The OC-DRA Database.....48*
- *Deployment Topology.....50*
- *PCA Scalability.....64*
- *Redundancy.....69*
- *IP Networking.....74*
- *PCA Routing.....75*
- *PCA Data Auditing.....80*
- *PCA and Application Chaining.....82*
- *The Communication Agent.....83*
- *Diameter Routing and Communication with PCA.....84*
- *PCA and IDIH Metadata.....87*
- *PCA Capacity Constraints.....94*
- *PCA Assumptions and Limitations.....95*

This section introduces the Policy and Charging application, key concepts, and basic functionality.

Policy and Charging is a feature of the Diameter Signaling Router (DSR) product, which is part of the Oracle product line of signaling products. Policy and Charging combines the existing functionality of Policy DRA (P-DRA) and with the enhanced functionality of Online Charging DRA (OC-DRA).

Overview

A PCA DSR consists of a number of PCA DA-MP servers, a number of SBR servers, OAM servers, and optional IPFE servers. The PCA DA-MPs are responsible for handling Diameter signaling and implementing the Policy DRA and Online Charging DRA functionalities, as well as the overall PCA application itself.

SBR servers host the policy session and policy binding databases for the P-DRA function and session database for the OC-DRA function. These are special purpose MP servers that provide an off-board database for use by the PCA feature hosted on the DA-MPs.

Each PCA DSR hosts connections to clients and to policy/charging servers such as OCSs and PCRFs. Clients are devices that request authorization for access to network resources on behalf of user equipment (such as mobile phones) from a PCRF, or request billing/charging instructions from an OCS. Policy Clients sit in the media stream and enforce Policy rules specified by the PCRF. Policy authorization requests and rules are carried in Diameter messages that are routed through Policy DRA. Policy DRA makes sure that all Policy authorization requests for a given subscriber in an APN are routed to the same PCRF. Charging clients (CTF) generates charging events based on the observation of network resource usage and collects the information pertaining to chargeable events within the network element, assembling this information into matching charging events, and sending these charging events towards the OCS. Online Charging DRA makes sure that these charging events are routed to the correct OCS.

PCA DSRs can be deployed in mated pairs such that policy and/or online charging session state is not lost, even if an entire PCA DSR fails or becomes inaccessible. When PCA mated pairs are deployed, the clients and PCRFs/OCSs are typically cross-connected such that both PCA DSRs have connections to all clients and all PCRFs/OCSs at both mated sites.

PCA DSRs can also be deployed in mated triplets such that session states are not lost, even if two PCA DSRs fail or become inaccessible. When a PCA mated triplet is deployed, clients and PCRFs/OCSs are cross-connected such that all three PCA DSRs have connections to all policy clients and all PCRFs/OCSs associated with the mated triplet.

“PCA network” is the term used to describe a set of PCA mated pairs and NOAM server pair/triplet. All clients and PCRFs/OCSs are reachable for Diameter signaling from any PCA DSR in the PCA network.

PCA is also designed to do the following:

- Reduce Diameter signaling latency where possible by doing the following:
 - Limiting the need to access off-board databases
 - Note:** "Off-board" in this context means on a server separate from the server handling the Diameter signaling
 - Limiting to a single WAN traversal to route a diameter message within the PCA network
 - Optimization of the most frequent "sunny day" scenarios, possibly at the expense of less common, or rainy day, scenarios
- Provide server redundancy by supporting clusters of active DA-MP servers
- Provides site redundancy by supporting mated pairs of P-DRA DSRs, as well as provide 3-site redundancy by supporting mated triplets of P-DRA DSRs

- Provide triple data redundancy for subscriber binding data by having geographically dispersed active, standby, and spare copies of each binding record for mated pair configuration
- Provide quadruple data redundancy for subscriber binding data by having geographically dispersed active, standby, spare, and spare copies of each binding record for mated triplet configuration
- Support scalability of each DSR by the addition of DA-MP servers, as well as support network scalability by the addition of PCA sites
- Limit network configuration complexity by making use of naming conventions for clients and PCRFs/OCSs
- Facilitate troubleshooting of network-wide database accesses and Diameter signaling by including correlation information in logs and traces

The P-DRA Database

The P-DRA function mainly uses databases. Subscribers are dynamically assigned to a PCRF. This assignment is called a binding. The binding exists as long as the subscriber has at least one policy Diameter session.

The following points describe a high-level view of SBR Binding and Session databases:

- There is one instance of the Binding database in the entire P-DRA network.
- There is one instance of the Session database per Policy DRA Mated Sites Place Association.
- Each binding record is associated with at least one Diameter session record. Binding records contain one Session Reference for each Diameter session that is associated with that binding.
- When a binding exists, there is at least one IMSI Anchor Key, Session, and Session Reference record.
- The IPv4, MISISDN, and IPv6 Alternate Keys are optional. They represent alternate ways, other than the IMSI, to identify a subscriber.

Bindings vs Sessions

While technically both are part of the P-DRA SBR database, the Binding database and the Session database are referred to separately because they serve different purposes and have different scopes within the P-DRA network.

Session Binding Repository (SBR) servers host the Session and Binding databases for use by the PCA application.

Bindings

In the most generic sense, a binding is a mapping between a subscriber and a PCRF assigned to handle policy decisions for that subscriber. In 3GPP networks, however, there is more than one way to identify a subscriber. So rather than having just one binding table mapping subscribers to PCRFs, there are four tables mapping subscriber identifiers to the PCRF that handles the subscriber's policy decisions.

P-DRA supports four subscriber identifiers: IMSI, MSISDN, IPv4 IP Address, and IPv6 IP Address. Of these, IMSI and MSISDN are relatively permanent in that they do not change from call to call. IP addresses, on the other hand, are assigned by PCEFs to a subscriber's device for temporary use in accessing the Internet or other IP services.

Regardless of the type of subscriber identifier, the relationship of a subscriber to a PCRF assigned by the P-DRA must be accessible from anywhere in the P-DRA network. This means that the tables in the binding database must be accessible from all P-DRA DSR sites. For example, a given IMSI, when bound, will appear in exactly one record in the binding database, but will be accessible from any P-DRA DSR in the P-DRA network

PCRF Pooling examines the APN along with the IMSI, in the mapping of the message to a Pool of PCRFs, but with the restriction that before a new binding is created, the logic must check for existence of another binding to the same PCRF Pool for the IMSI. If such a binding exists, the new APN is bound to the same PCRF as an existing APN mapped to the same PCRF Pool. After a binding exists, all sessions for that IMSI and APN are routed to the bound PCRF. Sessions for that IMSI and a different APN mapped to a different PCRF Pool can be routed to a different PCRF. With PCRF Pooling, an IMSI can have up to 10 binding-capable sessions, which can be bound to different PCRFs based on APN.

Binding-capable session initiation requests includes both IMSI and an APN. Policy DRA maps APNs, to a PCRF tool via **Policy and Charging > Configuration > Access Point Names**.

P-DRA then checks to determine whether a Sub-Pool exists by locating the PCRF Pool and the Origin-Host from the session initiation request via **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.

If the PCRF Pool and Origin-Host are mapped to a Sub-Pool, the Sub-Pool is used; otherwise, the PCRF Pool that was mapped to the APN is used.

The PCRF Pool or Sub-Pool is mapped to a PRT table via **Policy and Charging > Configuration > Policy DRA > PCRF Pool To PRT Mapping** on the SOAM GUI. The P-DRA application instructs the Diameter Routing Layer to use the PRT table associated with the PCRF Pool or Sub-Pool to route the request.

The Diameter Routing Layer selects the actual PCRF based on the Route Lists and Route Groups selected from the PRT Rules in the PRT table.

The following order is used to search for an existing binding:

- A binding for the IMSI and APN (from the ImsiApnAnchorKey table)
- A binding for the IMSI and suggested PCRF Pool or Sub-Pool (from the ImsiApnAnchorKey table)

If no binding exists, a new binding is created using the IMSI, APN, and PCRF Pool. For new bindings, the actual PCRF is not determined until a success answer is received from the PCRF that processed the session initiation request.

A split binding occurs when more than one PCRF has an active session for the same IMSI, APN combination. P-DRA avoids creation of split bindings by searching for and honoring applicable existing bindings before creating new bindings.

Sessions

In this context, a Session represents a Diameter session for a policy interface (Gx, Gxx, Gx-Prime, S9, or Rx). P-DRA maintains session state, for the following reasons:

- Subscriber identifiers used for bindings are created and destroyed as a result of Diameter Requests sent in the context of a Diameter session. In other words, subscriber identifiers are created by binding-capable session-initiating messages and removed by session-termination messages.
- If Topology Hiding is Enabled for a binding-dependent session, the bound PCRF is stored in the session state because binding keys are not guaranteed to exist in all Requests within a Diameter session.

Note: When topology hiding does not apply, the session state is not maintained for binding-dependent sessions.

There are two broad categories of Policy sessions:

Binding-capable sessions

A binding-capable session is a Policy session that is allowed to cause a new binding to be created for a subscriber.

Binding-capable sessions are created by Gx, Gxx, or the S9 versions of Gx and Gxx interfaces. If a CCR-I message arrives for a Binding Capable Interface, Policy DRA checks for an existing binding for the IMSI and APN in the message. If a binding exists, the CCR-I is routed to the bound PCRF.

Binding-capable sessions create and destroy alternate keys as the sessions are created and terminated.

Policy DRA APN-based PCRF Pool selection modifies the Policy DRA application logic to inspect the contents of binding-generating Gx CCR-I messages to select the type of PCRF to which the CCR-I messages are to be routed. This gives Policy DRA the ability to support service-specific PCRF sets. The APN used by the UE to connect to the network is used to determine the PCRF pool. The Origin-Host of the PCEF sending the CCR-I can then be used to select a PCRF sub-pool.

If additional subscriber identifiers, or Alternate Keys, are present in the CCR-I and configured in **Policy DRA -> Configuration -> Binding Key Priority**, binding records are created for each Alternate Key present in the CCR-I. For example, a binding-capable CCR-I may include a MSISDN and IPv4 and IPv6 addresses in addition to the IMSI. These Alternate Keys exist as long as the session exists.

Binding-dependent sessions

A binding-dependent session is a Policy session that cannot cause a binding to be created, and cannot be created unless a binding exists.

Binding-dependent sessions are created by Rx, Gx-Prime, or the S9 version of Rx binding-dependent session initiation request messages. If a binding dependent session initiation request message arrives for a Binding Dependent Interface, Policy DRA checks for an existing binding using a key in the binding dependent session initiation request message.

- If a binding is found, the AAR is routed to the bound PCRF.
- If no binding is found, Policy DRA answers the binding dependent session initiation request using an AAA with the error code configured for the “Binding Not Found” error condition.

Binding-dependent sessions can use Alternate Keys when locating a binding, but can neither create nor destroy Alternate Key Binding records.

The Policy DRA generally does not need to save session state for binding-dependent sessions. The exception is when the PCRF name is being topology hidden from the Policy Client. When Topology Hiding applies, the bound PCRF name is stored in the session. Storage of the PCRF name is necessary for the following reasons:

- If the Policy Client cannot learn the PCRF name from the AAA message because of the Topology Hiding.
- In-session messages (such as STR) are not guaranteed to include a subscriber identifier that could be used to look up the binding again.

The Binding Database

The Binding database consists of 4 tables: one Anchor Key table and three Alternate Key tables. Each binding table record maintains a list of one or more binding-capable sessions that contain a reference to the binding key. These sessions are referred to using a Session Reference (SessionRef) instance, which is just a shorter means of identifying a session (shorter than a Diameter Session Id string).

The more permanent keys (IMSI and MSISDN) can be referenced by more than one binding-capable session. These keys will not be removed until the last binding-capable session that included the key is terminated.

The transient keys (IP Addresses), on the other hand, can be referenced only by a single binding-capable session.

The metadata captured by IDIH for the PCA includes the results of each query that PCA makes to the binding database and the associated result. Whenever the result of a database query is captured in PCA metadata, it will include the identity of the specific server that generated the response.

Anchor Key

Because binding capable sessions can originate from different places in the network at nearly the same time, it is necessary to serialize the requests to prevent both from being assigned to different PCRFs. Serialization is accomplished by requiring that binding capable session origination messages (i.e. CCR-I) always contain an IMSI and that the IMSI is always used for creation of new bindings

See [Error Codes](#).

Alternate Keys

Alternate Keys provide different ways to identify a subscriber. Alternate Keys are created by binding-capable sessions and used by binding-dependent sessions.

For example, a UE attached to a binding-dependent interface like Rx might not have access to the subscriber's IMSI, but might have an IPv6 address that has been temporarily assigned to the subscriber. This IPv6 Alternate Key can be used to find the subscriber binding and the correct PCRF to route the Rx or Gx-Prime request to, only if that IPv6 Alternate Key record was previously created by a binding-capable session.

Alternate Keys are optional. If all interfaces have access to the IMSI, or Anchor Key, there is no need to create or use Alternate Keys. Alternate Keys are created when they are present in the binding-capable session creation message (CCR-I) and they are assigned a P-DRA Binding Key Priority.

If a binding-capable session initiation message includes multiple Alternate Keys that are also assigned with a Binding Key Priority, all of those Alternate Keys will be created when the binding-capable session is established. When a binding-dependent session creation message arrives, which Alternate Key will be used to find the binding depends to some degree on configuration.

P-DRA allows the handling of Alternate Keys to be configured. The configuration defines which Alternate Keys should be used, and the Priority order in which to use them. (Assignment of Priorities must be consecutive, without skipping a number between two other numbers.)

[Table 8: Example of a Binding Key Priority Configuration](#) illustrates an example configuration of Alternate Keys. Key types are assigned to the Priority values 1 through 4, where 1 is the highest Priority (IMSI, IPv4, IPv6, or MSISDN). If a particular type of key is not used, that key need not be assigned to a Priority. In the example, IPv4 is not being used as an Alternate Key, meaning that even if a

Framed-IP-Address is present in the binding-capable session initiation message, no IPv4 key will be created.

Table 8: Example of a Binding Key Priority Configuration

Priority	Key
1	IMSI
2	IPv6
3	MSISDN
4	<Not Configured>

The Priority order defines the order in which P-DRA looks for a given key type in a binding-dependent session initiating message. In the example in [Table 8: Example of a Binding Key Priority Configuration](#), P-A will look for keys in the following order and AVP:

1. IMSI: Subscription-Id AVP with Subscription-Id-Type of END_USER_IMSI
2. IPv6 Address: Framed-IPv6-Prefix AVP (only high order 64 bits used)
3. MSISDN: Subscription-Id AVP with Subscription-Id-Type of END_USER_E164

For each key found in the message and assigned a Binding Key Priority, P-DRA will attempt to find a binding record in the corresponding binding database table. If a key is not present, P-A will skip to the next highest Priority key type. Some keys can have more than one instance in a Diameter message, but only the first instance of a given key type will be used in the binding search.

- If no configured key is present in the Diameter message, an error response is returned to the originator.
- If keys are present in the Diameter message, but no corresponding binding is found, an error is returned to the originator. The configurable "Binding Not Found" error condition is used. See [Error Codes](#).

The Session Database

The Session database consists of 2 tables:

- A Session table
- A SessionRef table

Session

The Session table is keyed by a Diameter Session-Id, a long string that is defined by Diameter to be "globally and eternally unique". In addition, the Session table stores the values of any Alternate Keys defined by binding-capable sessions. The relationship between Diameter sessions and Alternate Keys must be maintained so that the Alternate Keys can be removed when sessions defining those Alternate Keys are terminated.

The PCRF identifier to which a session is bound is stored in the Session record. This may be used to route in-session messages if topology hiding is enabled. In-session messages are not guaranteed to contain the same keys as session initiating messages.

Each Session record has a corresponding SessionRef record. The SessionRef provides a more compact means of uniquely identifying a Diameter Session-Id. This allows for a more compact Binding database. Session and SessionRef records are created and destroyed in unison.

The metadata captured by IDIH for the PCA includes the results of each query that Policy DRA makes to the session database and the associated result. Whenever the result of a database query is captured in PXA metadata, it will include the identity of the specific server that generated the response.

Session Reference

SessionRef records are used to tie Binding records to Diameter sessions. This allows P-DRA to know when a Binding record should be removed. IMSI and MSISDN records are removed when the last binding-capable session that referenced them is removed. IP Address records are removed when the only binding-capable session that referenced them is removed.

Because each Binding record must be associated with at least one valid Session record, a Binding record can be removed if it is not associated with any existing SessionRef. Removal of orphaned Binding records is one of the jobs of the P-DRA database audit. See [PCA Data Auditing](#) for more information about the database audit.

Binding Key Query Tool

Due to the distributed nature of the binding and session databases, it can be difficult to determine if a given key is associated with a binding. P-DRA includes a GUI-based binding key query tool to help with troubleshooting policy problems.

To use the tool, the user inputs a binding key value. The tool queries the binding database to determine if the binding key exists. If the binding key exists, the tool generates a report that includes the PCRF that the key is bound to and information about which binding capable Diameter session or sessions are associated with that binding key. The session information, when returned, includes all other binding keys that were included in the session, the session creation time, and the session last touched time. If the binding key entered by the user does not exist, the report indicates that the binding key was not found.

Note: The binding key query tool only displays binding capable sessions (e.g. Gx, Gxx, S9, etc.) because only Gx sessions create and delete binding data. No tool exists for querying binding dependent sessions (e.g. Rx, Gx-Prime) associated with a given binding key.

The binding key query tool is intended for individual queries for binding keys specified by the user. It is not intended to dump all binding keys in the database, nor audit large numbers of binding keys.

The OC-DRA Database

The OC session database consists of two tables: an OcSession table and an OcClientHost table.

The OcSession table, keyed by a Diameter Session-Id, a long string that is defined by Diameter to be "globally and eternally unique" is used to store the session state info for a session. A subscriber Id is used to look up all sessions associated with the subscriber for session report. An OCS Id references another OC server table for the OCS servers that the sessions may be routed to. Each OcSession table also has a CTF Id referencing the OcClientHost table where the data of OC client host (i.e. FQDN) and OC client realm is stored.

Online Charging over Gy/Ro Reference Point

Early implementations of Online Charging were based on Diameter Credit Control Application (DCCA) messages.

The Ro reference point supports interaction between a Charging Trigger Function (CTF) and an Online Charging System (OCS). The Gy reference point is functionally equivalent to Ro, and hence is replaced by Ro within the common charging architecture. The Gy reference point is functionally equivalent to Ro, making Gy and Ro synonymous. Gy/Ro based DCCA is assigned Application-Id 4 (diameter credit-control) and used by DCCA compliant clients/servers.

The following basic scenarios are used for online charging:

- Event-Based Charging
- Session-Based Charging

Both online charging scenarios rely on Diameter Credit Control Credit-Control-Request/ Answer (CCR/CCA) messages and Re-Auth-Request/ Answer (RAR/RAA) messages.

Event-based charging is used for charging individual and independent events like SMS or MMS. For the event-based charging scenario, the CTF sends an OCS a Credit-Control-Request (CCR) with CC-Request-Type AVP set to EVENT_REQUEST (4).

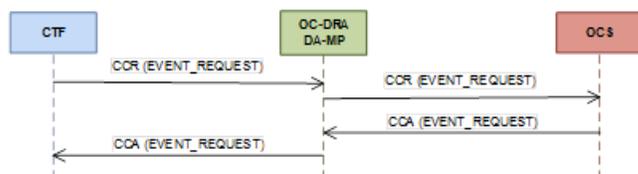


Figure 13: Example of Gy/Ro Event-Based Charging

Session-based charging is generally used for charging voice calls or data usage. For the session-based charging scenario, the session is initiated by the CTF sending an OCS a Credit-Control-Request (CCR) with CC-Request-Type AVP set to INITIAL_REQUEST (1), followed by zero, one or more CCRs with CC-Request-Type AVP set to UPDATE_REQUEST (2) until the session is terminated by the CTF sending the OCS a CCR with CC-Request-Type AVP set to TERMINATION_REQUEST (3). The OCS may also re-authorize multiple active resource quotas within a session by sending the CTF a Re-Auth-Request (RAR) message. All messages exchanged within a session use the same Session-Id value.

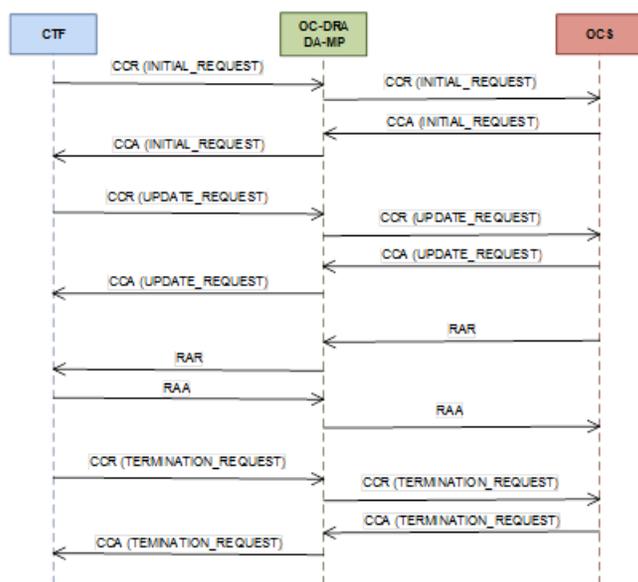


Figure 14: Example of Gy/Ro Session-Based Charging

Charging is typically based on MSISDN and all CCRs include the MSISDN in the Subscription-Id Grouped AVP or in the User-Name AVP. In the Subscription-Id Grouped AVP, the MSISDN is typically present in the Subscription-Id-Data AVP when the corresponding Subscription-Id-Type AVP is set to "END_USER_E164". However, it is also possible that the MSISDN could be packaged in the NAI format in the User-Name AVP or in the SIP URI format in the Subscription-Id-Data AVP of the Subscription-ID grouped AVP when the corresponding Subscription-Id-Type AVP is set to "END_USER_SIP_URI".

OC-DRA attempts to retrieve the subscriber's identity from the AVPs mentioned above (in the order listed) and store them as part of subscriber state if session state is maintained. If the subscriber's identity is in the form of a SIP URI, Tel URI or a NAI format, then OC-DRA will not extract the MSISDN or perform number conditioning from these formats. Instead, it saves the entire identity as it appears in the AVP.

Binding-independent Interface

A binding-independent session is an online charging session that is allowed to cause a new session to be created for a subscriber. Binding-independent sessions are created by Gy or Ro interfaces.

Deployment Topology

This section describes the makeup of a PCA network, regardless of its size. [Figure 15: Sites, Mated Pairs, and Region](#) illustrates an example PCA network.

- A PCA Network can have up to 16 mated pairs or 32 sites, or can be as small as a single site. A PCA Network can also have up to 5 mated triplets.
- The PCA Binding Region provides the scope of the Policy Binding database. Binding records are accessible from every PCA DSR where the PDRA function is enabled in the Region.

- The Binding database need not be confined to a single mated pair or mated triplet.
- A PCA Mated Sites Place Association provides the scope for an instance of the Session database. Session records are accessible from each PCA DSR in the Mated Sites.
- Clients and PCRFs/OCSs have primary connections to their local PCA and secondary/tertiary connections to the mate(s) of their local PCA.
- PCA DSRs are connected to each other on the External Signaling Network. Each PCA site must be reachable from every other PCA site in the Region for Diameter signaling.
- The external signaling network handles stack events, database replication, and Diameter signaling. All three are required for the Diameter signaling to function correctly and with the required level of redundancy. “Services” (configured using the **Configuration->Services** GUI page) can be used to enforce separation of different types of traffic.

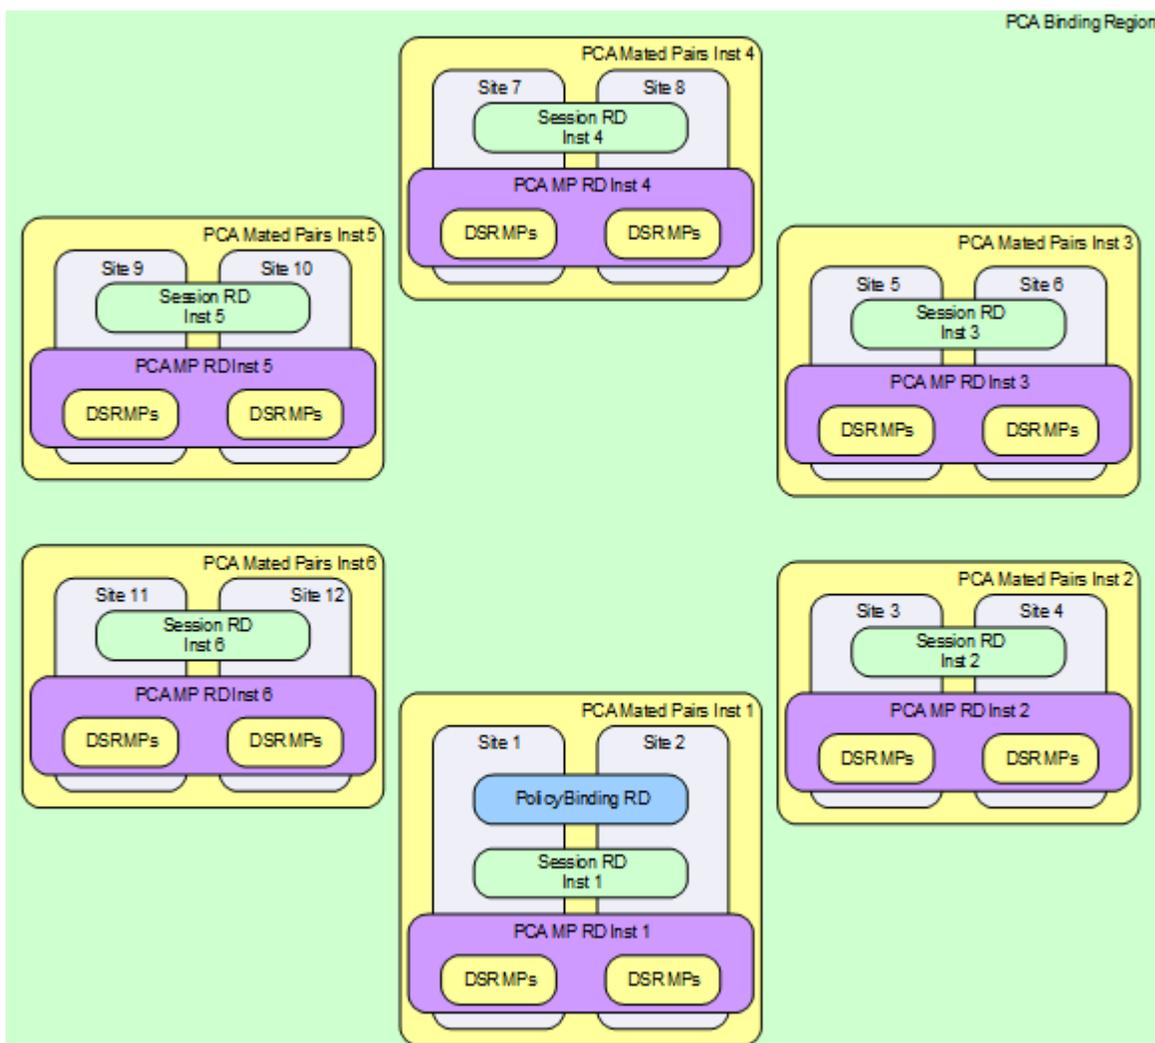


Figure 15: Sites, Mated Pairs, and Region

See [PCA Scalability](#) for details on how the Policy DRA feature can scale from very small lab and trial systems to large multi-site deployments.

Policy DRA in Roaming Scenarios

3GPP has defined two roaming scenarios with respect to Policy Control and Charging functions. The Policy DRA can be deployed for various network scenarios as a Policy routing agent, including the roaming scenarios.

In addition to communicating to the Policy Clients and Policy servers through Gx/Gxx, Gx-Prime, and Rx interfaces in their own networks, the Policy DRAs can communicate to each other across the Visited Access and Home Access (or Home Routed Access) networks through the S9 interface, for session binding purposes.

Figure 16: Policy DRA in Roaming Scenarios illustrates an example Diameter network where the Policy DRAs are located in Home Access and Visited Access networks.

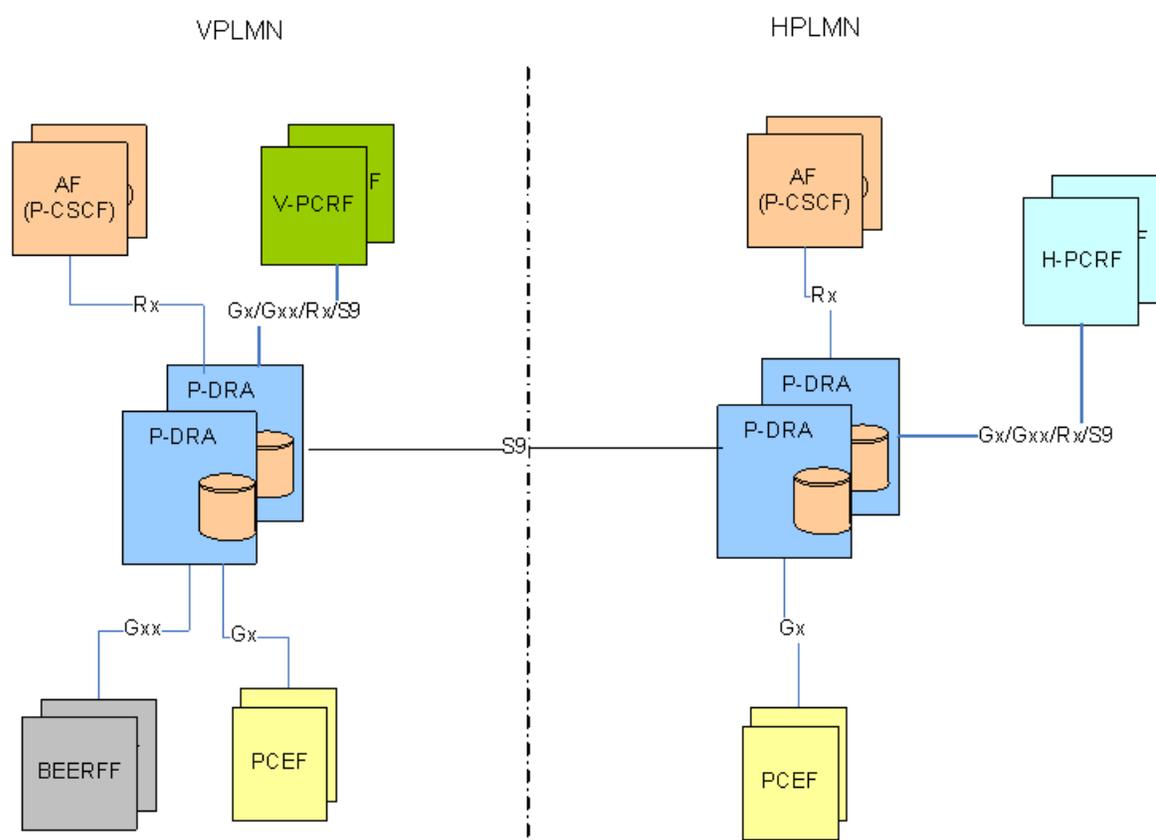


Figure 16: Policy DRA in Roaming Scenarios

The Visited Access (also known as Local Breakout) is one of the scenarios where UEs obtain access to the packet data network from the VPLMN where the PCEF is located.

The Home Routed Access is the roaming scenario in which the UEs obtain access to the Packet Data Network from the HPLMN where the PCEF is located.

The S9 reference point is defined in roaming scenarios between HPLMN and VPLMN over which two Diameter applications, S9 and Rx are used. The purpose of the S9 Diameter application is to install PCC or QoS rules from the HPLMN to the VPLMN and transport the events occurred in the VPLMN to the HPLMN.

The S9 protocol makes use of exactly the same commands and messages as the Gx/Gxx protocols, except that a V-PCRF in VPLMN will provide an emergency treatment for any incoming CC-Request (INITIAL_REQUEST) messages. This implies that the Policy DRA does not check the existence of the Called-Station-ID AVP if the IMSI is missing in a CC-Request (INITIAL_REQUEST) over the S9 interface.

PCA Configurable Components

Key PCA configurable components include:

- PCA Binding Region Place Association consisting of all PCA Sites
- PCA Mated Sites Place Associations, each consisting of redundant PCA Sites
- Policy Session Resource Domains - one per PCA Mated Sites Place Association consisting of all session SBR server groups at the mated sites.
- Policy Binding Resource Domain - one per PCA Binding Region Place Association consisting of all binding SBR server groups.
- Policy and Charging DRA Resource Domains - one per PCA Mated Sites Place Association consisting of all DSR (multi-active cluster) server groups at the mated sites.
- SBR Server Groups - enough to handle the load in stack events per second.
- Diameter Signaling Router (multi-active cluster) Server Groups - one per PCA DRA Site.
- SBR Databases - PCA supports two types of SBR Database: SBR Binding Database, used by the P-DRA function of PCA, and SBR Session Database, used by both P-DRA and OC-DRA functions of PCA. An SBR Database is hosted by Policy and Charging SBR Server Groups assigned to either a Policy Binding Resource Domain or a Policy Session Resource Domain.
- SBR Database Resizing Plans - An SBR Resizing Plan is the configuration that identifies the SBR Database to be resized and the Target Resource Domain to which the SBR data will be migrated.
- SBR Data Migration Plans - An SBR Data Migration Plan is the configuration that identifies the Initial SBR Database that is the source of SBR data to be migrated and the Target SBR Database to which the SBR data will be migrated.

For multiple mated pair/triplet deployments, there are two different configurations for mated pairs:

- One mated pair/triplet that hosts the PCA Binding database and an instance of the Session database
- N mated pairs/triplets that each host only an instance of the Session database

Figure 17: Example PCA Mated Pair - Hosting Binding SBRs illustrates two PCA DSR Sites configured as a Mated Pair:

- This Mated Pair hosts the PCA Binding database and an instance of the Session database.
- The SBR Database is configurable and contains a number of SBR Server Groups.
- Each SBR Server Group consists of 3 servers using the Active/Standby/Spare redundancy model, allowing for Site redundancy.
- The number of SBR Server Groups necessary to host the binding or session database will be determined by the application provider prior to feature activation based on expected policy signaling needs. The number of SBR Server Groups is configured in the SBR Database and can be increased or decreased as needed.
- Each Site has an SOAM Server Group consisting of 3 servers using the Active/Standby/Spare redundancy model, allowing for Site redundancy.
- The PCA network has an NOAM Server Group consisting of 2 servers using the Active/Standby redundancy model. If NOAM site redundancy is desired, another pair of Disaster Recovery NOAM servers can be deployed at a different Site.

- Each Site has a number of DA-MP servers sufficient to carry the desired Diameter signaling load.
- Each Site has two pairs of IPFE servers – one for use by Policy Clients and one for use by PCRFs. (IPFE is not required.)

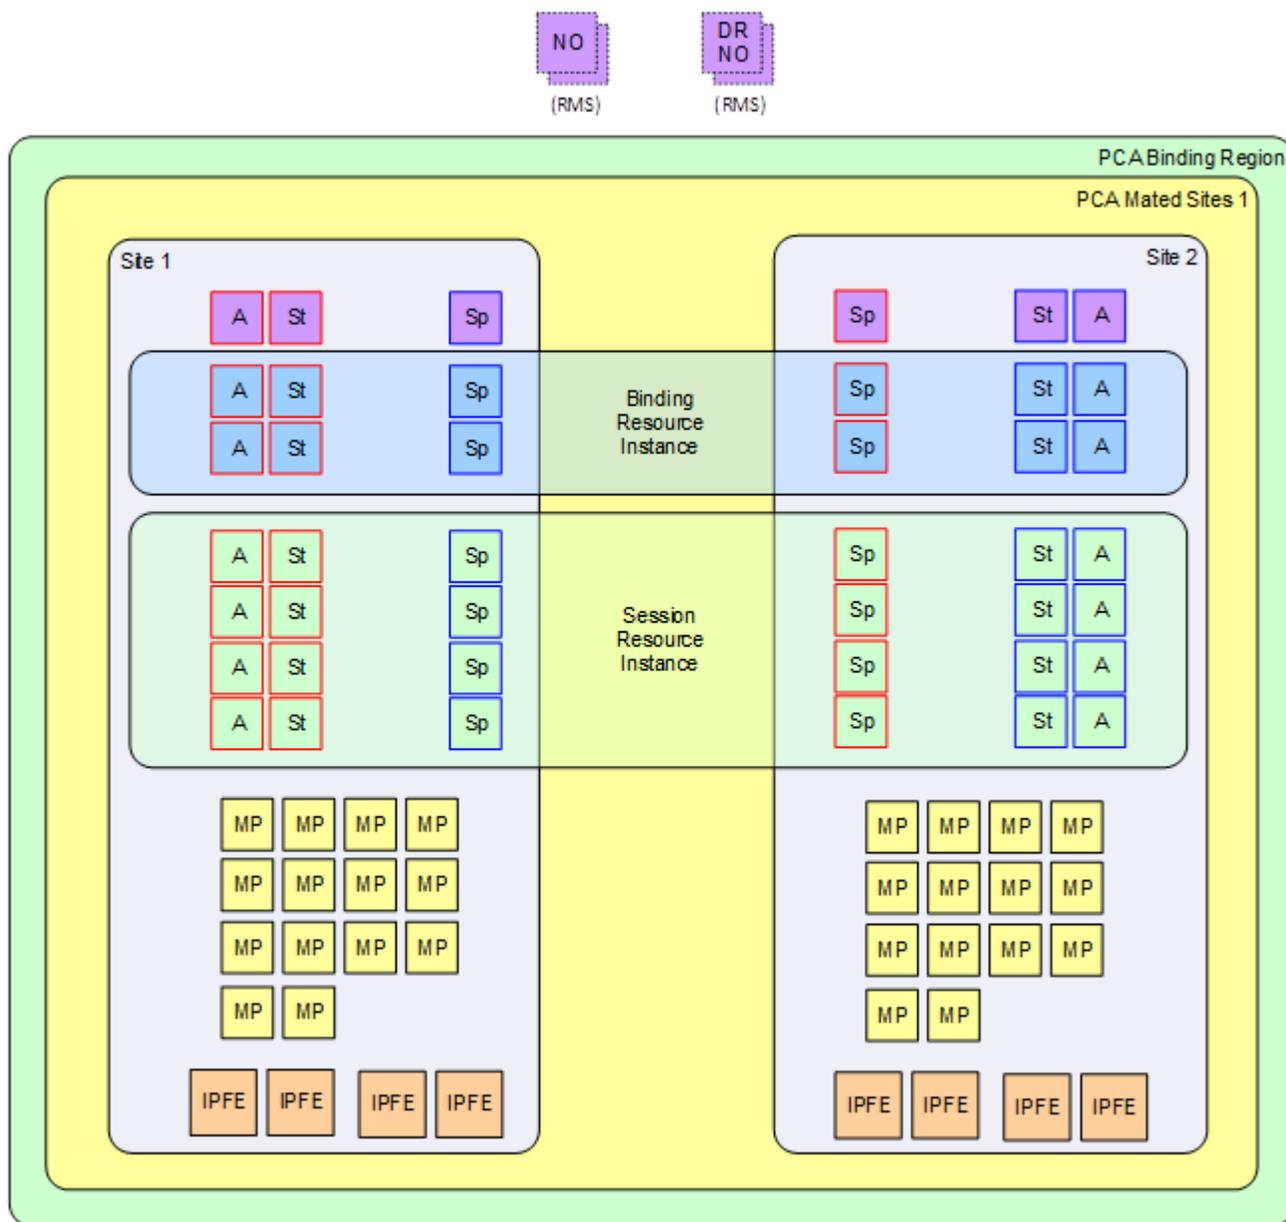


Figure 17: Example PCA Mated Pair - Hosting Binding SBRs

Figure 18: Example PCA Mated Pair - Not Hosting Binding Policy SBRs illustrates a possible configuration for additional mated pairs that do not host the Binding database:

- Each subsequent mated pair deployed after the set of mated pairs hosting the Binding database will host only an instance of the Session database (no Binding database).
- The number of DA-MPs can vary depending on the expected Diameter signaling load.

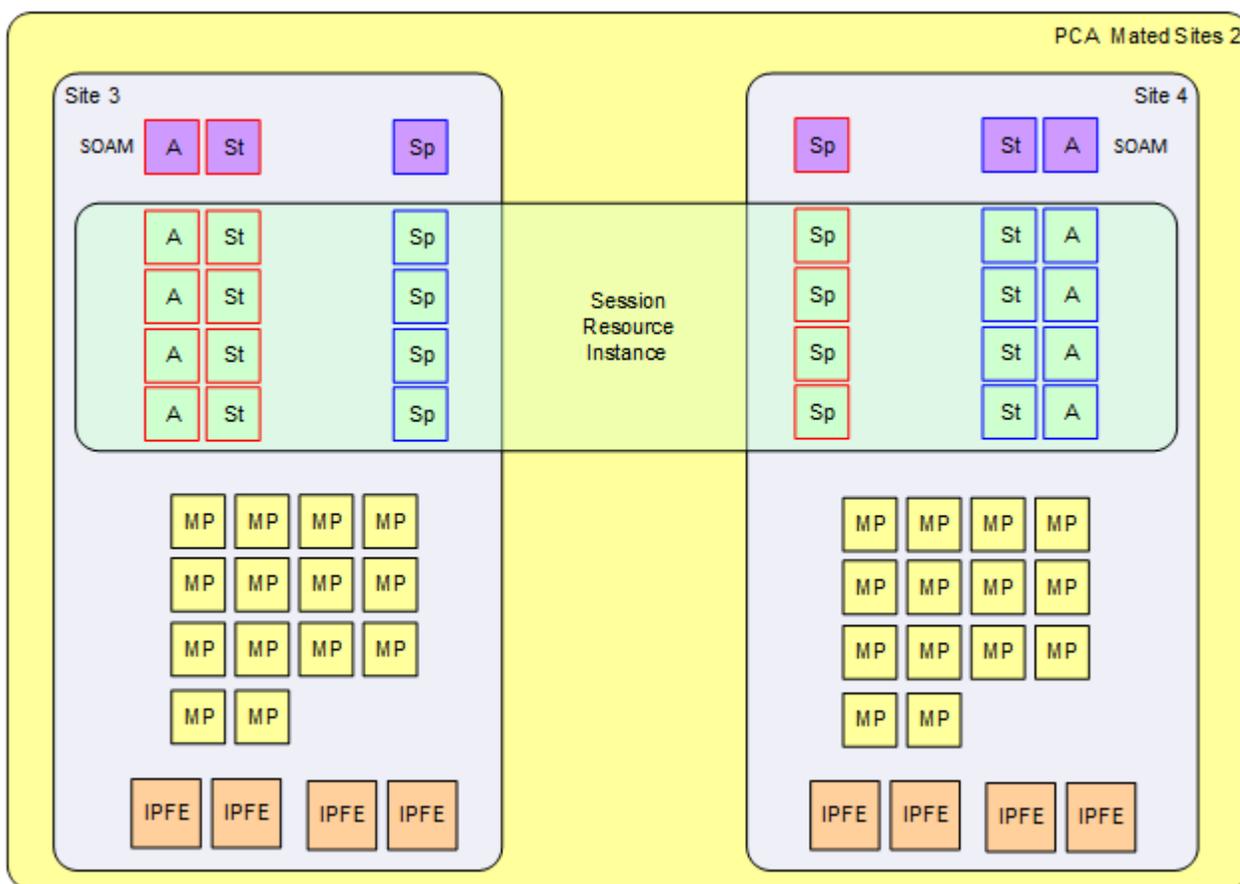


Figure 18: Example PCA Mated Pair - Not Hosting Binding Policy SBRs

Figure 19: Policy Client, PCRF, and Site Relationships illustrates example relationships between PCA DSR Sites and Policy Clients and PCRFs:

- Each PCA DSR Site has a set of Policy Clients whose primary connection is directed to that PCA.
- Each PCA DSR Site has a set of PCRFs to which it distributes new bindings. Each PCRF at this Site has a primary connection to the PCA DSR at that Site.
- Each policy client should have a secondary connection to the mate of the PCA DSR for which it has a primary connection. (Without this “cross-connect”, PCA site failure would leave the Policy Client with no access to any PCRF.)
- Each PCRF should have a secondary connection to the mate of the PCA DSR for which it has a primary connection. (Without this “cross-connect”, PCA site failure would leave the PCRF inaccessible.)
- Each Mated Pair of PCA DSRs shares an instance of the Policy Session database.
- All PCA DSRs share the Policy Binding database, conceptually in the middle of the network.
- If Diameter signaling must be sent to a PCRF for which the PCA DSR has no connection, the message must be routed to a PCA DSR that does have a connection. This routing is configured using the DSR routing tables.

See [Diameter Routing and Communication with PCA](#) for more details about Diameter routing for PCA

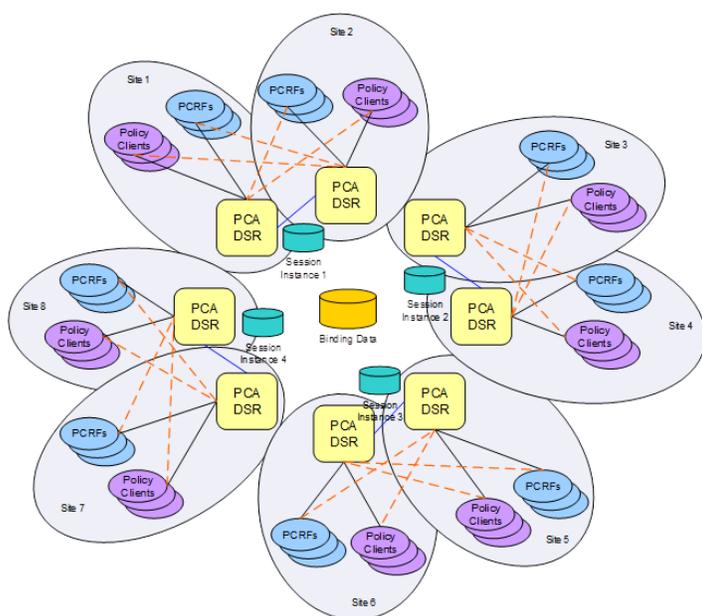


Figure 19: Policy Client, PCRF, and Site Relationships

Places

A “Place” allows servers or other Places to be associated with a physical location. The only Place type is “Site”. A Site Place allows servers to be associated with a physical site.

An OAM GUI is used to configure Sites that correspond to physical locations where equipment resides. For example, Sites may be configured for Atlanta, Charlotte, and Chicago. Exactly one Place can be associated with a server when the server is configured

Place Associations

A “Place Association” allows Places to be grouped in ways that make sense for DSR Applications. A Place Association is a collection of one or more Places that have a common “Type”. A Place can be a member of more than one Place Association.

The PCA application defines two Place Association Types:

- PCA Binding Region

As illustrated in [Figure 15: Sites, Mated Pairs, and Region](#), the PCA application defines a Region to include all Sites that are part of the PCA network. This provides a scope for the Binding database, which is accessible to all PCA Sites in the PCA network.

- PCA Mated Pair

As illustrated in [Figure 15: Sites, Mated Pairs, and Region](#), pairs of PCA Sites are grouped together as Mated Pairs. Each Place Association with Type of PCA Mated Pair includes exactly 2 sites. A PCA Mated Pair has the following attributes:

- Hosts an instance of the PCA Session database
- Hosts client Diameter connections for Policy Clients at both Sites in the Mated Pair
- Hosts PCRF/OCS Diameter connections for PCRFs/OCSs at both Sites in the Mated Pair

Server Groups

The PCA application makes use of several different types of Server Groups, as defined in [Table 9: Server Group Functions](#).

Table 9: Server Group Functions

Server Type	Server Group Function Name	Level
Diameter MP servers	DSR (multi-active cluster)	MP
SBR(S) and SBR(B) servers	SBR	MP
IPFE	IP Front End	MP
OAM server	DSR (active/standby pair)	NOAM,SOAM

- SBR Type

Server Groups with the “SBR” function type host either or both of the Policy Binding and Policy Session databases. The type of database hosted by a given Server Group depends on the Resource Domain or Domains with which the Server Group is associated.

Each SBR Server Group consists of one to four servers, depending on the type of deployment. [Table 10: SBR Server Group Configuration and Data Redundance](#) describes the supported configurations for SBR Server Groups. See [Redundancy](#) for details on policy data redundancy.

Table 10: SBR Server Group Configuration and Data Redundance

# of Servers	Redundancy	Typical Use
1	Active only. No Redundancy.	Labs and demos only.
2	Active/Standby. Server redundancy within a Site.	Single-site deployments or deployments not requiring Site redundancy.
3	Active/Standby/Spare	Mated Pair deployments to avoid a single-server failure from causing Session access requests to be routed to the mate Site. This is the target for large deployments. New sessions are equally distributed across all Session SBR Server Groups in the mated pair, meaning that ~50% of the Session accesses will be routed across the WAN. Note: SBR Server Groups must be configured with two WAN replication channels.
4	Active/Standby/Spare/Spare	Mated Triplet deployments where service is kept alive even two of the three sites with the service are down.

Because only the active server in a SBR Server Group is actually processing Stack Events, a SBR Server Group can be engineered to run at 80% of maximum capacity. This holds for Site failure as well since the Spare server at the mate site will take over.

- DSR (multi-active cluster) Type

For PCA, all of the DA-MPs at a Site (even if there is only one) must be included in one Server Group with the DSR (multi-active cluster) function type. This eliminates the need to have all clients and PCRFs/OCSs connected to every DA-MP.

The DA-MP servers in the Server Group will be treated as a cluster of active servers. There should be at least two DA- MP servers in the Server Group in order to support in-service maintenance or upgrade. The DA- MP servers in a Server Group should be engineered such that loss of a single server will not cause the remaining servers to go into overload.

If the PCA is being deployed in mated pairs, the DA- MP servers at one site need to be configured to handle the entire load of the other site (in case of a site failure) without causing the surviving DA-MPs to go into overload – typically 40% of engineered capacity.

Resource Domains

A Resource Domain allows Server Groups to be grouped together and associated with a type of application resource. Each Resource Domain has a “Profile” that indicates the application usage of the resource domain. The PCA application defines three Resource Domain Profiles: Policy Session, Policy Binding, and Policy and Charging DRA.

Once SBR Server Groups are configured to host the session and binding databases, those Server Groups can be added to Policy Binding and Policy Session Resource Domains. An SBR Server Group must be associated with either a Policy Session or Policy Binding Resource Domain.

DA-MP servers are configured in a single server group per PCA DSR with a server group function of "DSR (multi-active cluster)". For a mated pair deployment, the two DSR (multi-active cluster) server groups containing all of the DA-MPs at the two sites must be included in a PCA Resource Domain. For a non-mated deployment, the DSR (multi-active cluster) server group must be in its own P and Charging DRA Resource Domain. For a mated triplet deployment, the three DSR (multi-active cluster) server groups containing all of the DA-MPs at the three sites must be included in a Policy and Charging DRA Resource Domain.

Clients

Clients act on behalf of the user equipment (UEs) to request policy/charging authorization and enforce policy/charging rules received from the PCRFs/OCSs. The clients send requests to the PCA, which ensures that the request are sent to the PCRF/OCS in charge of policy/charging for the subscriber associated with the UE.

PCA supports four different types of Policy Clients, referred to by 3GPP as AF, PCEF, BBERF, and CTF:

- The AF uses the Rx Diameter interface.
- The PCEF uses the Gx Diameter interface.
- The BBERF uses the Gxx Diameter interface.
- The CTF uses the Gy/Ro Diameter interface

How many connections a Client might initiate towards the PCA and how those connections are used are in customer control. The capabilities of the client, however, affect the functionality of the solution; as shown in [Table 11: Client Connection Capability](#).

Table 11: Client Connection Capability

Number of Connections Supported by Policy and Charging Client (per Diameter host)	Effect on Solution Capability
1	<ul style="list-style-type: none"> • Site Redundancy cannot be supported. • Diameter signaling throughput is limited to the capacity of the connection. • Extra latency to reconnect in the event of a connection drop.
2	<ul style="list-style-type: none"> • Site Redundancy supported if secondary connection is configured to connect to PCA mate site. • If both connections go to a single site and the policy and charging client has the capability to use both connections simultaneously, Diameter signaling throughput may be doubled vs. only one connection. <p>This configuration requires multiple Diameter connections to a single Diameter host - something that is not supported by RFC 6733, but which many vendors support to allow capacity beyond what a single connection can support.</p> <ul style="list-style-type: none"> • Extra latency is avoided in the event of a single connection drop because the other connection can be used without waiting for reconnect and Capabilities Exchange. • PCA Mated Triplet is not supported, since the client would become isolated if the two DSRs with which it has connections fail.
3	<ul style="list-style-type: none"> • PCA Mated Triplet is supported if the client has one connection to each DSR in the triplet. • PCA Mated Pair is supported if and only if the client has at least one connection to each DSR in the pair. The extra connection to one of the two DSRs provides the opportunity for higher availability and throughput.
>3	<p>There are many scenarios possible, depending on the capabilities of the policy and charging client. For example, there might be two connections to the primary PCA (for capacity) and two to each mate PCA.</p>

Any Diameter Request can be sent to either PCA in the mated pair, but to avoid possible race conditions between signaling and replication, messages in a Diameter session should be sent to the same PCA Site when possible.

PCRFs

PCRFs are responsible for authorizing and making policy decisions based on knowledge of subscriber resource usage and the capabilities allowed by the subscriber's account. In order to perform this function, all policy requests for a given subscriber must be routed to the same PCRF.

Rather than provisioning a fixed relationship between a subscriber and a PCRF, the P-DRA function of PCA dynamically assigns subscribers to PCRFs using a load distribution algorithm and maintains state about which subscribers are assigned to which PCRF. The relationship between a subscriber and a PCRF can change any time the subscriber transitions from having no Diameter policy sessions to having one or more Diameter policy sessions. Once a policy session exists, however, all policy sessions for that subscriber are routed to the assigned PCRF.

PCA can interact with any 3GPP Release 9 compliant PCRF. Because these PCRFs come from different vendors, there are differences in how they are deployed in the network and how they “look” to the P-DRA function. The following PCRF configurations differ mainly in addressing and sharing of state across Diameter connections:

- A PCRF that shares state across different Diameter hostnames.
 - Each Diameter hostname can all support Gx, Gxx, S9, Gx-Prime and Rx Diameter interfaces. This type of PCRF is supported by PCA.
 - Each hostname has a different connection for each different interface type. This type of PCRF is supported by PCA.
 - There is a different Diameter hostname for each connection for a specific Diameter interface. All of the Diameter hostnames share state. This type of PCRF is supported by PCA.
 - There are different Diameter hostnames for different policy client vendors. Policy state is shared across the Diameter hostnames, but origin based routing is required to select a set of PCRFs for distribution of the initial binding depending on the policy client type. This type of PCRF is supported by PCA, but requires use of Diameter Routing Function PCRF selection as described in [PCRF Selection for New Bindings](#).
 - There is a different Diameter hostname for each connection. This type of PCRF is supported by PCA, but requires use of Diameter Routing Function PCRF selection based on the vendor type of the policy client as described in [PCRF Selection for New Bindings](#).
- A PCRF that has one Diameter hostname, but supports a number of connections to that hostname using different IP addresses.

Each connection can support Gx, Gxx, S9, Gx-Prime and Rx Diameter interfaces. This type of PCRF is supported by PCA.

OCSs

In the context of PCA deployment, OCS is referred to as Online Charging Server. OCSs are responsible for:

- Authorizing service, i.e. granting or denying the services to the subscribers who requested the services via Diameter online charging signaling.
- Charging in accordance with service provisioned in real time based on accounting/metering
- Making the decision to terminate the service if certain conditions are met.

An OCS is selected by the OC-DRA function of PCA for an incoming session initiation message, i.e. Gy/Ro CCR-I, using a load distribution algorithm. OC-DRA may store and maintain the session state for the subscriber to ensure all the in-session messages, i.e. CCR-U and CCR-T, will be routed to the same OCS for online charging processing for this session. The relationship between the subscriber and the OCS lasts during the lifetime of the session.

The SBR Database

An SBR Database consists of several fields:

- A user-configured Database Name
- A Database Type (either Session or Binding)
- The Resource Domain Name of the Resource Domain that contains the Policy and Charging SBR Server Groups
- The Number of Server Groups that define the capacity of the SBR Database
- The Place Association that uses the SBR Database

Note: An SBR Database must be configured after configuring Policy and Charging SBR Server Groups, Policy Binding or Policy Session Resource Domains, and Policy Binding Region or Policy and Charging Mated Sites Place Associations.

When using the Policy DRA function of PCA, a Binding type SBR Database is required, plus an instance of a Session type SBR Database for each Policy and Charging Mated Sites Place Association. The Binding SBR Database serves the entire network, or more specifically, the Site Places assigned to the Policy Binding Region Place Association. Each Session SBR Database serves the Site Places defined in the Policy and Charging Mated Sites Place Association.

When using only the Online Charging DRA function of PCA, no Binding type SBR Database is needed, but a Session type SBR Database must still be configured for each Policy and Charging Mated Sites Place Association.

When using both Policy DRA and Online Charging DRA functions of PCA, a single instance of Session type SBR Database supports both policy sessions and online charging sessions for a given Policy and Charging Mated Sites Place Association. Additionally, the SBR Binding Database is still necessary for the Policy DRA function to work.

SBR Databases cannot be edited. If a field must be changed, the SBR Database must be deleted and re-inserted with the corrected value.

For further information on configuring an SBR Database, refer to [SBR Databases](#).

For further information on maintaining an SBR Database, refer to [SBR Database Status](#).

Database Growth

SBR Databases can be created with different sized session databases for each Policy and Charging Mated Sites instance. Each instance of the Session SBR Database can be sized appropriately for the expected level of signaling traffic at that set of mated sites.

Session SBR Databases can be resized without loss of policy or charging data and without disruption to policy or charging signaling. Binding SBR Databases can be resized without loss of binding data and without disruption to policy signaling.

For further information about how to resize an SBR database, refer to [SBR Database Resizing Plans](#).

SBR Database Resizing

An SBR Resizing Plan is the configuration that identifies the SBR Database to be resized and the Target Resource Domain to which the SBR data will be migrated.

SBR Resizing involves changing the capacity (number of server groups) of an SBR Database by creating a new Resource Domain (called the Target Resource Domain) for the SBR Database and migrating the data from the Initial Resource Domain to the Target Resource Domain. The Target Resource Domain

may be larger or smaller than the Initial Resource Domain, allowing the SBR Database to grow or shrink in size.

For further information on configuring SBR Database Resizing Plans, refer to [SBR Database Resizing Plans](#).

SBR Data Migration

SBR Data Migration is used when the SBR Database topology for session databases is to be changed. SBR Data Migration is the process of moving data from the Initial to Target SBR Database by creating new records in the target and allowing records in the initial to be removed by normal PCA Diameter signaling.

There are multiple situations in which SBR Data Migration may occur:

- Mating (combining SBR Databases when transitioning from stand-alone PCA DSR sites to mated PCA DSR sites)
- Un-mating (Separating SBR Databases when transitioning from mated PCA DSR sites to stand-alone PCA DSR sites)
- Re-mating (Swapping mate relationships between two sets of mated sites)

For further information on configuring SBR Data Migration Plans, refer to [SBR Data Migration Plans](#).

IPFE

In order to simplify network connectivity, PCA will typically be deployed with one or two pairs of IPFEs per PCA DSR site. IPFE is not mandatory, however; it is up to the customer whether it should be included.

The following deployment scenarios involving IPFE are possible:

- A single site PCA in which the PCRFs and/or OCSs are not capable of initiating connections to the PCA. For example:
 - A PCA DSR Site with a pair of IPFE servers, 8 DA-MP server, and some SBR servers
 - Four Policy and Charging Clients connected to two IPFE TSAs, with primary connections and secondary connections
 - The DA-MP servers are split into two groups that host connections to TSA1 and TSA2 respectively. This is necessary to ensure that a Policy and Charging Client's primary and secondary connections do not end up being connected to the same DA-MP.
 - One IPFE server is primary for TSA1 and standby for TSA2; the other IPFE server is primary for TSA2 and standby for TSA1.
 - PCA MPs-to-PCRFs or MPs-to-OCSs connectivity need not be fully meshed.
- An IPFE configuration in which Policy and Charging Clients are connected to a PCA mated pair, but PCRFs and/or OCSs are not capable of initiating connections to the PCA. Each Policy Client has a primary connection to one PCA site and a secondary connection to the mate site. For example:
 - Two PCA DSR sites, each with a pair of IPFE servers and 4 DA-MP servers.
 - Three Policy and Charging Clients with a primary connection to PCA DSR Site 1 and secondary connections to PCA DSR Site 2.
 - Three Policy and Charging Clients with a primary connection to PCA DSR Site 2 and secondary connections to PCA DSR Site 1.
 - Two PCRFs or OCSs with primary connections to PCA DSR Site1 and secondary connections to PCA DSR Site 2.

- Two PCRFs or OCSs with primary connections to PCA DSR Site2 and secondary connections to PCA DSR Site 1.
- One IPFE at PCA DSR Site 1 is primary for TSA1. The other IPFE is standby for TSA1.
- One IPFE at PCA DSR Site 2 is primary for TSA2. The other IPFE is standby for TSA2.
- A single site PCA in which a single IPFE pair is used for both Policy and Charging Clients and PCRFs and/or OCSs. The use of IPFE for PCRFs is possible only if the PCRF can be configured to initiate connections towards the PCA. Some customers refer to an IPFE used by PCRFs as an IP Back-End, or IPBE, although there is no difference between an IPBE and an IPFE from a software or configuration perspective. For example:
 - One pair of IPFE servers, each server supporting two TSAs
 - Four Policy and Charging Clients connect to TSA1 with their secondary connection going to TSA3, or vice-versa.
 - The PCRFs or OCSs connect to TSA2 with their secondary connection going to TSA4, or vice-versa.
 - Six PCA MP servers, each capable of hosting connections from Policy and Charging Clients and PCRFs or OCSs
 - One IPFE server is primary for TSA1 and TSA2, and standby for TSA3 and TSA4.
 - The other IPFE server is primary for TSA3 and TSA4, and standby for TSA1 and TSA2.
- A single site PCA in which IPFE is used for both Policy Clients and PCRFs. In this case, two pairs of IPFE servers are deployed in order to support high Diameter signaling bandwidth. For example:
 - Two pairs of IPFEs, each supporting a two TSAs
 - The Policy and Charging Clients connect to either TSA1 or TSA2, with their secondary connection going to the other TSA.
 - The PCRFs or OCSs connect to either TSA3 or TSA4, with their secondary connection going to the other TSA.
 - Eight PCA DA-MPs, each capable of hosting connections from Policy and Charging Clients and PCRFs or OCSs
 - One IPFE server on the Policy and Charging Client side is primary for TSA1 and standby for TSA2. The other IPFE server is primary for TSA2 and standby for TSA1.
 - One IPFE server on the PCRF or OCS side is primary for TSA3 and standby for TSA4. The other IPFE server is primary for TSA4 and standby for TSA3.
- A PCA mated pair configured with an IPFE for Policy Clients and a separate IPFE for PCRFs. The Policy and Charging Clients and PCRFs have a primary connection to their local PCA DSR and a secondary connection to the mate PCA DSR. For example:
 - Two PCA DSR sites, each with a two pairs of IPFE servers and 6 DA-MP servers
 - Three Policy and Charging Clients with a primary connection to PCA DSR Site 1 and secondary connections to PCA DSR Site 2.
 - Three Policy and Charging Clients with a primary connection to PCA DSR Site 2 and secondary connections to PCA DSR Site 1.
 - Two PCRFs or OCSs with primary connections to PCA DSR Site1 and secondary connections to PCA DSR Site 2.
 - Two PCRFs or OCSs with primary connections to PCA DSR Site2 and secondary connections to PCA DSR Site 1.
 - One IPFE on the Policy and Charging Client side at PCA DSR Site 1 is primary for TSA1. The other IPFE is standby for TSA1.

- One IPFE on the Policy and Charging Client side at PCA DSR Site 2 is primary for TSA3. The other IPFE is standby for TSA3.
- One IPFE on the PCRF or OCS side at PCA DSR Site 1 is primary for TSA2. The other IPFE is standby for TSA2.
- One IPFE on the PCRF OCS side at PCA DSR Site 2 is primary for TSA4. The other IPFE is standby for TSA4.

PCA Scalability

The PCA application is highly scalable. In addition to scaling up to support large customer networks, PCA can scale down to support small customers, lab trials, and demos. This section describes supported configurations that illustrate how the PCA feature scales.

For large systems, PCA can scale up as follows:

- 16 mated pairs of PCA DSRs (32 sites) or up to 5 mated triplets (15 sites) per network
- Three enclosures per PCA DSR site using half-height servers

Each enclosure has 16 half-height slots.

- Two pairs of IPFE servers per PCA DSR
- Sixteen DA-MP servers per PCA DSR

Figure 15: Sites, Mated Pairs, and Region illustrates a sample PCA network consisting of 6 mated pairs, or 12 sites with components that must be configured as follows:

- An instance of a Site (Place with type Site) is created for each physical location of a PCA DSR.
- All MP servers (both SBRs and DA-MPs) are assigned to the Site where they are physically located.
- An instance of a PCA Mated Pair (Place Association with type PCA Mated Pair) is created for each pair of sites that are mates.
- A number of Policy Binding Server Groups are created on the PCA DSR nodes that are initially deployed.
 - Each Policy Binding Server Group, if configured for site redundancy, must have at least one Server at the home site and one Server at the mate site.
- A Policy Binding Resource Domain is created including all Policy Binding Server Groups.
- A pre-determined number of Session Server Groups are created at each mated pair.

Each Session Server Group, if configured for site redundancy, must have at least one server at the home site and one server at the mate site.

- A Policy Session Resource Domain is created for each mated pair including the Session Server Groups at the two mated sites.
- A DSR (multi-active cluster) Server Group is created for each Site, containing all of the DA-MP servers at the Site.
- A Policy and Charging DRA Resource Domain is created including the Diameter Signaling Router Server Group at each of the mated Sites.
- A PCA Binding Region (Place Association with type PCA Binding Region) is created containing all Sites.

MP Growth

The PCA application supports addition of DA-MPs as needed to support Diameter traffic. Each PDRA DA-MP server can support 12,000 MPS when engineered to run at 40% to support site redundancy. If site redundancy is not needed, PCA DA-MPs can be engineered at 80%, thereby supporting 24,000 MPS.

The DSR supports up to 16 DA-MPs per DSR site.

1. Insert the new Server on the NOAM GUI Main Menu **Configuration > Servers** page and assign it to a Site place.
2. Add the Server to the DSR (multi-active cluster) Server Group for the desired Signaling Network Element on the NOAM Main Menu **Configuration > Server Groups** page.
3. Restart the Server on the SOAM GUI Main Menu **Status & Manage > Server** page.
4. Enable the PCA application on the SOAM GUI Main Menu **Diameter > Maintenance > Applications** page.

Mated Pair Growth

A mated PCA DSR can be added to a PCA DSR.

A mated pair of PCA DSRs can be added to a PCA network.

A mated triplet of PCA DSRs can be added to a PCA network

Adding a Mated PCA DSR to an Existing PCA DSR

A PCA DSR deployed without a mate must host all of the SBR Server Groups that are planned for deployment across the mated pair when the mate is added. This action requires planning ahead for the eventual mate.

Note: SBR Server Groups with only one server represent a single point of failure for a portion of the SBR database.

A PCA DSR site could be configured as follows for eventually adding a mate:

- Site A has two SOAM Server Groups configured: the red one on the top left for use by Site A and the blue one on the top right for use by Site B.
 - The Site A SOAM Server Group is set up with two Servers in Active/Standby configuration.
 - The Site B SOAM Server Group is set up with one Server configured as Preferred Spare. Because there are no other Servers in this Server Group, the Server will become active.
- Site A has four SBR(B) Server Groups configured: the two red ones on the left for use by Site A and the two blue ones on the right for use by Site B.
 - The Site A SBR(B) Server Groups are set up with two Servers in Active/Standby configuration. These Server Groups have the Site A SOAM Server Group as parent.
 - The Site B SBR(B) Server Groups are set up with one Server configured as Preferred Spare. These Server Groups have the Site B SOAM Server Group as parent. Because there are no other Servers in these Server Groups, the single Server will become active.
- Site A has eight SBR(S) Server Groups configured: the four red ones on the left for use by Site A and the four blue ones on the right for use by Site B.

- The Site A SBR(S) Server Groups are set up with two servers in Active/Standby configuration. These Server Groups have the Site A SOAM Server Group as parent.
- The Site B SBR(S) Server Groups are set up with one Server configured as Preferred Spare. These Server Groups have the Site B SOAM Server Group as parent. Because there are no other Servers in these Server Groups, the single Server will become active.

Adding a Mated Pair of PCA DSRs

PCA network capacity can be expanded by adding mated pairs of PCA DSRs.

PCA MP servers can be added as needed (up to a maximum of 16 DA-MPs) to support the desired level of Diameter signaling traffic.

Adding a Mated Triplet of PCA DSRs

PCA network capacity can be expanded by adding up to five mated triplets of PCA.

PCA MP servers can be added as needed (up to a maximum of 16 DA-MPs) to support the desired level of Diameter signaling traffic.

Small System Support

In order to support small customers and lab and trial deployments, the PCA application can scale down to run on a small hardware footprint. This section describes the smallest supported PCA DSR deployments.

A lab or trial system may not be required to support in-service maintenance or have any hardware redundancy whatsoever. In the smallest supported lab/trial PCA DSR, IPFE is not included because it does not make sense to distribute ingress connections when there is only one DA-MP server.

The NOAM and SOAM servers are also running in simplex mode, meaning that no redundancy exists. In addition, the NOAM and SOAM are virtualized on a single physical server to save hardware. The binding and session SBR servers are also running in simplex mode, but must be configured to host either a Policy Binding and Policy Session database. A single DA-MP hosts all Diameter signaling. Signaling is not affected if one or both of the (virtual) OAM servers happens to fail.

The configuration of the smallest viable commercially deployable PCA DSR, illustrated in , has enough hardware redundancy to support in-service maintenance:

- Two DA-MPs are required to survive server failures and maintenance. These DA-MPs should be engineered at 40% load since in a failure or maintenance situation, one Server will have to handle the load for both.
- Both binding and session SBR Servers pairs use the Active/Standby redundancy model in order to support failures and maintenance.
- The NOAM/SOAM Server pair uses the Active/Standby redundancy model in order to support failures and maintenance.
- Both NOAM and SOAM are virtualized onto a single pair of physical servers. The NOAM instance is Active on one server and Standby on the other. The SOAM instance is Active on one server and Standby on the other.

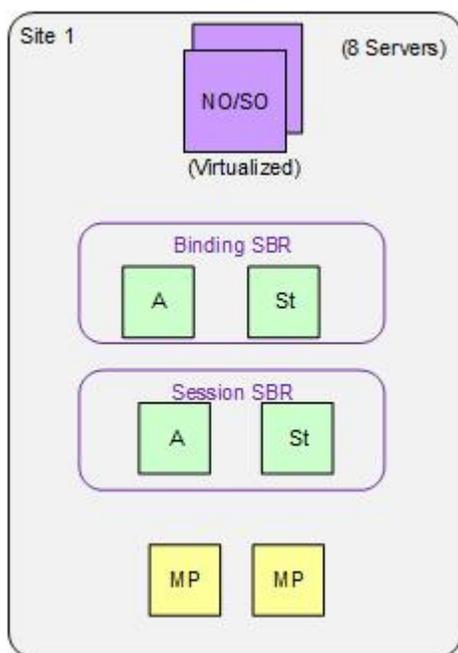


Figure 20: Smallest Supported PCA Field Deployment

The smallest supported Mated Pair of PCA DSRs, illustrated in [Figure 21: Smallest Supported PCA Mated Pair](#), has the following characteristics:

- The NOAM servers are deployed at Site 1 using Active/Standby redundancy.
- The Site 1 SOAM servers are deployed at Site 1, virtualized on the same servers with the NOAM servers. They, however, use the Active/Standby/Spare redundancy model, with the Spare server deployed at Site 2 and virtualized on the same server with one of the Site 2 SOAM servers.
- The Site 2 SOAM servers are deployed at Site 2 using the Active/Standby/Spare redundancy model. The Spare Site 2 SOAM server is virtualized at Site 1 on one of the servers already hosting an NOAM and a Site 1 SOAM server.
- A Binding SBR triplet is deployed with two servers at Site 1 and one server at Site 2.
- A Session SBR triplet is deployed with 1 server at Site 1 and two at Site 2
- Two DA-MPs are deployed at each site to support server redundancy at each site.

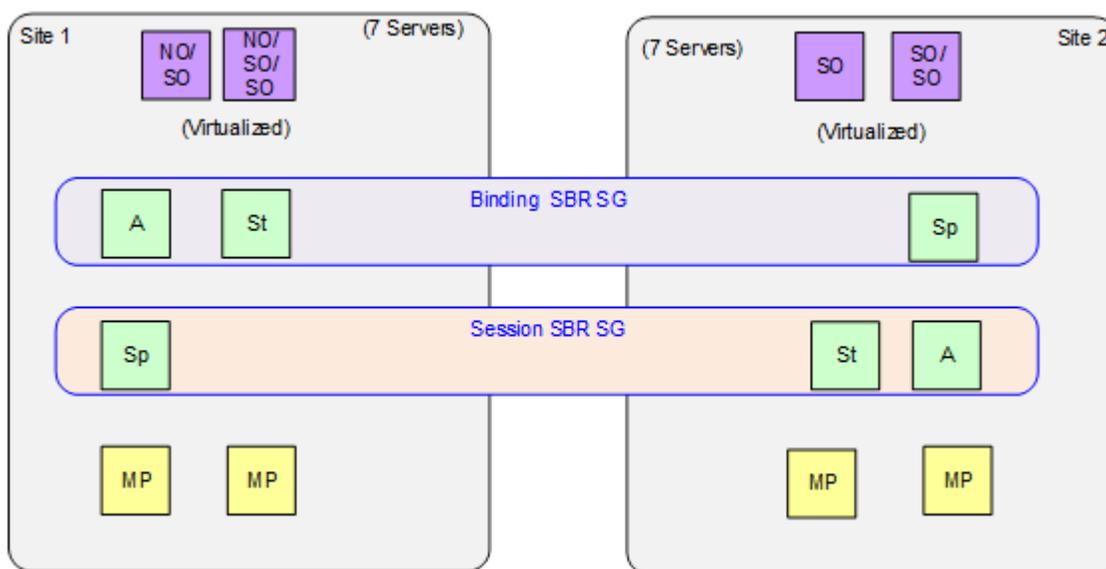


Figure 21: Smallest Supported PCA Mated Pair

The smallest supported mated triplets of PCA DSRs, illustrated in , has the following characteristics:

- The NOAM servers are deployed at Site 1 using Active/Standby redundancy.
- The Site 1 SOAM server are deployed at Site 1, virtualized on the same servers with the NOAM servers. However, the use the Active/Standby/Spare/Spare redundancy model, with the spare server deployed at Site 2 and Site 3 and virtualized on the same server with one of the Site 2 SOAM servers.
- The Site 2 (and Site 3) SOAM servers are deployed at Site 2 (and Site 3) using the Active/Standby/Spare/Spare redundancy model. The spare Site 2 SOAM server is virtualized at Site 1 on one of the servers already hosting an NOAM and a Site 1 SOAM server.
- A binding SBR triplet is deployed with two servers at Site 1 and one server each at Site 2 and Site 2 respectively.
- A session SBR triplet is deployed with two servers at Site 2 and one server each at Site 1 and Site 3 respectively.
- Two DA-MP servers are deployed at each site to support server redundancy at each site.

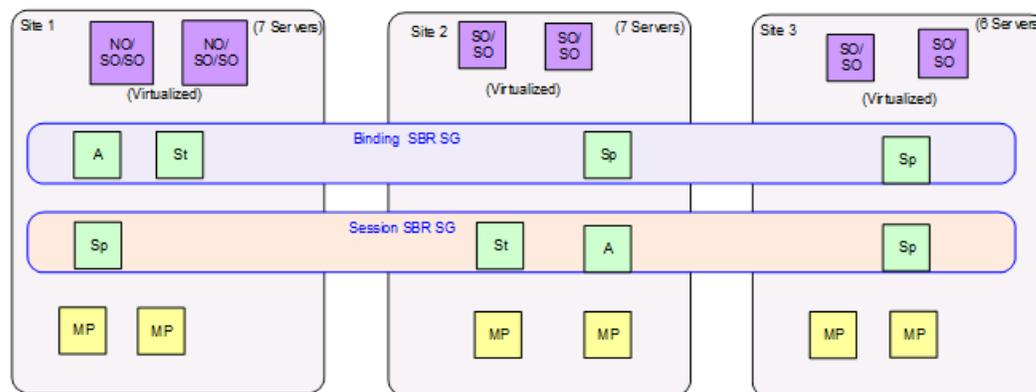


Figure 22: Smallest Supported PCA Mated Triplets

Redundancy

Making the PCA application highly available is accomplished by deploying enough hardware to eliminate single points of failure. Except for lab and trial deployments, OAM servers and MP servers must be deployed such that a single failure or maintenance activity will not prevent the feature from performing its function.

The PCA application also supports site redundancy, which is the ability for the feature to continue functioning even when an entire site is lost to disaster or network isolation.

MP Server Redundancy

The following redundancy models are supported for MP servers, whether deployed as DA-MPs or SBR MPs:

- DA-MP Multi-Active Cluster

PCA MP servers are deployed using an Active/Active redundancy model. This means that every DA-MP actively processes Diameter signaling. In order to avoid single points of failure, a minimum of two DA-MPs must be deployed (except for lab and trial deployments, where one DA-MP is acceptable). DA-MPs at a given site must be configured such that loss of a single DA-MP will not cause the remaining DA-MP servers to go into signaling overload.

- SBR Active Only

An SBR (either Session or Binding) can be deployed in simplex redundancy mode only for labs or trials. Otherwise this configuration represents a single point of failure for the SBR database being hosted by the Active-only Server Group. In this configuration, the SBR Server Groups consist of a single Server.

- SBR Active/Standby

The Active/Standby redundancy model should be used for single site PCA deployments, or for multi-site deployments when site redundancy is not important. In this configuration, the SBR Server Groups consist of two servers. On system initialization, one of the two servers in each SBR Server Group will be assigned the Active role and the other the Standby role. These roles will not change unless a failure or maintenance action causes a switch-over. For Active/Standby Server Groups, switch-overs are non-revertive, meaning that recovery of a formerly Active server will not cause a second switch-over to revert the Active role to that server.

- SBR Active/Spare

The Active/Spare redundancy model can be used for mated pair deployments in which it is acceptable for traffic to move from one site to the mate site on failure of a single server. In this configuration, the SBR Server Groups consist of two servers with one marked as "Preferred Spare". On system initialization, the server not marked as Preferred Spare will be assigned the Active role and the other the Spare role. These roles will not change unless a failure or maintenance action causes a switch-over. For Active/Spare Server Groups, switch-overs are revertive, meaning that recovery of a formerly Active server will cause a second switch-over to revert the Active role to that server.

- SBR Active/Standby/Spare

The Active/Standby/Spare redundancy model should be used for PCA mated pair deployments in which site redundancy is desired. In this configuration, each SBR Server Group is configured with two servers at one site and the third at the mate site. The server at the mate site is designated in the Server Group configuration as "Preferred Spare". On system initialization, one of the two servers that are located at the same site will be assigned the Active role and the other the Standby role. The server at the mate site will be assigned the Spare role (as was preferred). If the Active server can no longer perform its function due to failure or maintenance, the Standby Server will be promoted to Active. Only if both Active and Standby servers at a site are unable to perform their function will the Spare server at the mate site be promoted to Active. Active and Standby role changes within a site are non-revertive, but if the server at the mate site is Active and one of the other servers recovers, a switch-over will occur to revert the Active role back to the site with two servers.

- SBR Active/Standby/Spare/Spare

The Active/Standby/Spare/Spare redundancy model should be used for PCA mated triplet deployments. In this configuration, each SBR server group is configured with two Servers at one site and one server at each of two mate sites. The Server at each mate site is designated in the Server Group configuration as "Preferred Spare". on system initialization, one of the two Servers that are located at the same site will be assigned the spare role (as was preferred). If the active server can no longer perform its function due to failure or maintenance, the standby server will be promoted to active. Only if both active and standby servers at a site are unable to perform their function will a spare server at a mate site be promoted to active. Active and standby role changes within a site are non-revertive, but the server at a mate site is active and one of the servers at the site with two servers recovers, a switch-over will occur to revert the active role back to the site with two servers.

Site Redundancy

2-Site Redundancy

2-Site redundancy is the ability to lose an entire site, for example due to a natural disaster or major network failure, without losing signaling or application state data. For PCA, this means no loss of Policy Binding or Policy Session data. In order to achieve site redundancy, the following configuration applies:

- PCA is deployed on at least one mated pair of PCA DSRs.
- Clients and PCRFs/OCSs are able to connect to both sites in the mated pair.
- SBR server groups are set up to use the Active/Standby/Spare or Active/Spare redundancy model.
- System OAM (SOAM) server groups are set up to use the Active/Standby/Spare redundancy model.
- Diameter Agent MP servers are recommended to be engineered at 40% capacity across the mated pair.

3-Site Redundancy

3-Site redundancy is the ability to lose two entire PCA sites simultaneously, for example due to a natural disaster or major network failure, without losing signaling or application state data. For PCA, this means no loss of Policy Binding or Session data. In order to achieve 3-site redundancy, the following configuration applies:

- PCA is deployed on at least one mated triplet of PCA DSRs.

- Clients and PCRFs/OCSs connect to all sites in the mated triplet.
- SBR Server groups use the Active/Standby/Spare/Spare redundancy model.
- System OAM server groups use the Active/Standby/Spare/Spare redundancy model.
- Diameter Agent MP servers are recommended to be engineered at 26% capacity across the mated triplet, i.e. if two sites fail, then each remaining DA-MP operates at 80% capacity

Comparing 2-Site and 3-Site Redundancy

In order to achieve the same redundant capacity of a mated pair in a mated triplet, the following things need to be added:

- One 'Preferred Spare' SOAM to every existing SOAM server group.
- One SOAM quadruplet server group for the third DSR.
- One 'Preferred Spare' SBR to every existing SBR server group (total of two spares per group).
- [Optional] One new SBR quadruplet server group homed to the new site for every existing pair of SBR server groups. This increases total Session SBR capacity, but is primarily intended to balance flows.
- The same quantities of IPFEs and DA-MPs of the original DSR to the new DSR and Diameter connections to all of the Policy and Charging Clients and Policy and Charging Servers associated with the mated triplet.

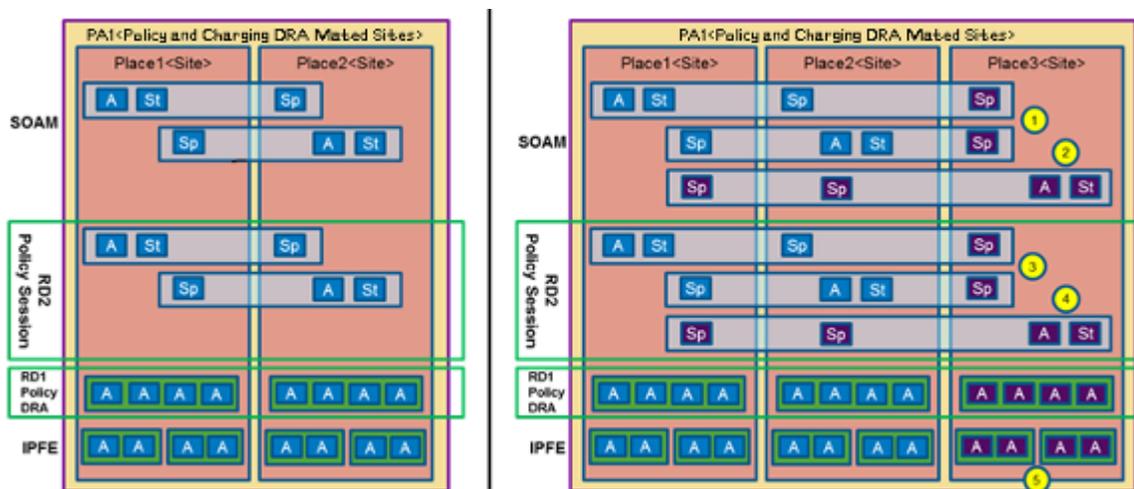


Figure 23: Comparing 2-Site and 3-Site Redundancy

Data Redundancy

The session and policy binding databases are partitioned such that each server group in a Session or Binding resource domain hosts a portion of the data. Because each server group consists of redundant servers (Active/Standby, Active/Spare, or Active/Standby/Spare), the data owned by each Server Group is redundant within the Server Group.

Active, Standby, and Spare servers within a SBR server group all maintain exact replicas of the data for the partition for which the server group is responsible. This data is kept in sync by using a form of signaling called replication. The synchronized tables on the Standby and Spare servers are continually audited and updated to match the master copy on the Active server.

Figure 24: Binding Table Partitioning Across Server Groups illustrates how a given Policy Binding table might be partitioned across four SBR Server Groups in a Policy Binding resource domain.

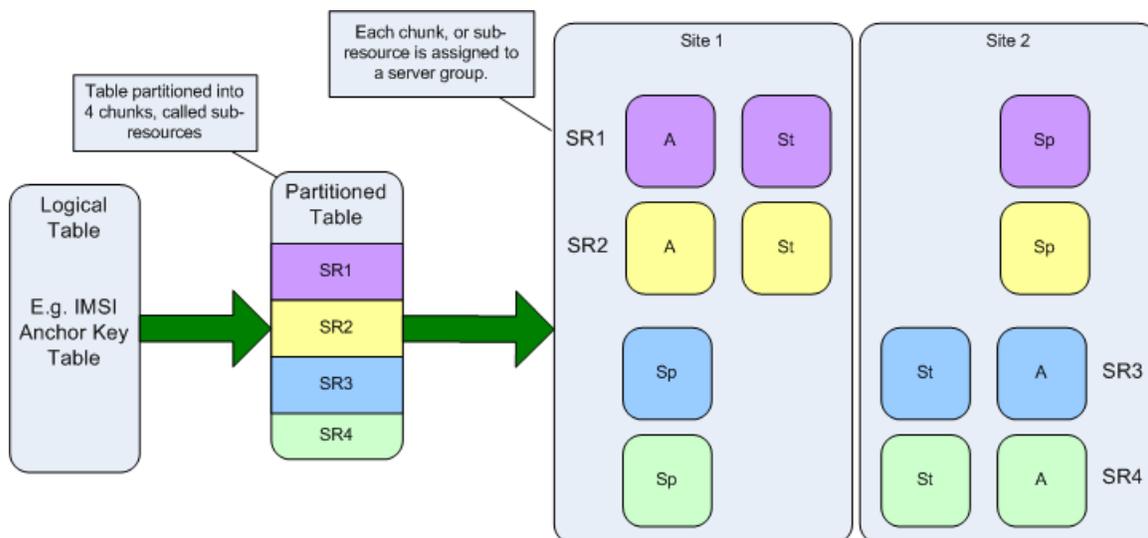


Figure 24: Binding Table Partitioning Across Server Groups

Figure 25: Multi-Table Resources illustrates how each SBR Server Group hosts a partition of several tables. Only the Active server within each server group can write to the database. The Standby and Spare servers replicate only actions (adds, changes, and deletes) performed by the Active server.

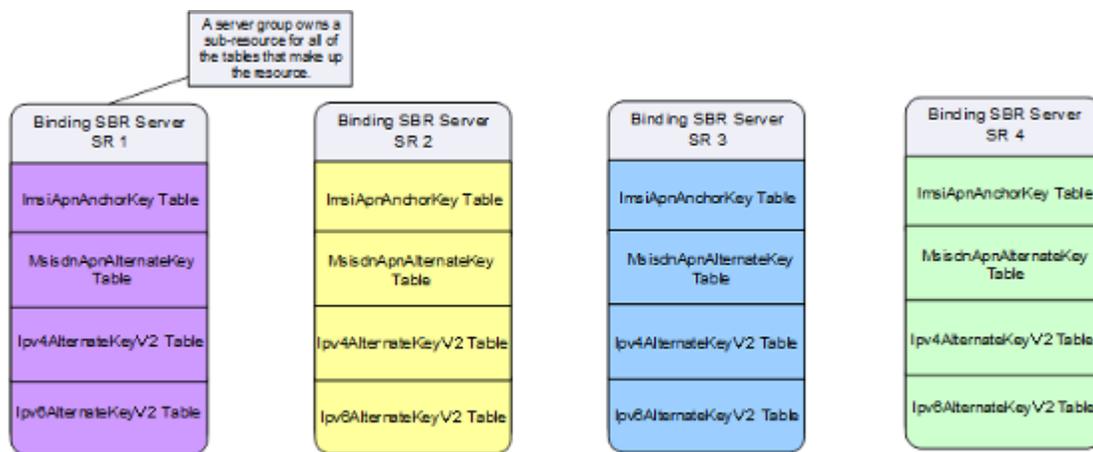


Figure 25: Multi-Table Resources

OAM Server Redundancy

The PCA application can be deployed with varying degrees of redundancy on the NOAM and SOAM servers. Like the SBR servers, the OAM servers can be configured to support site redundancy if desired.

Regardless of whether site redundancy is supported, the OAM servers must be deployed on redundant servers at a given site.

- Active/Standby NOAM and Active/Standby DR NOAM

The NOAM servers are deployed using the active/standby redundancy model at one of the sites in the PCA network. If site redundancy is desired, an optional pair of Disaster Recovery (DR) NOAM servers can be deployed at a different site. The DR NOAM servers are used only if manually brought into service following loss of the site where the original NOAM pair was located.

- Active/Standby/Spare SOAM

If site redundancy is desired for PCA mated pairs, the SOAM servers at each of the mate DSRs should be deployed using the Active/Standby/Spare redundancy model. In this configuration, two SOAM servers are deployed at one site and a third server is deployed at the mate site. The third server is configured as “Preferred Spare” in the SOAM Server Group. In the event of a site failure, the SBR Servers running at the surviving site of the mated pair will report measurements, events, and alarms to the SOAM server at that site. Without the Spare SOAM server, the Spare SBR servers would have no parent OAM server and would not be able to report measurements, events, and alarms.

SBR servers in a given SBR Server Group must be set up such that they belong to the Signaling Network Element of the site that has two of the three servers. This will allow all three servers in the Server Group to merge their measurements, events, and alarms to the same SOAM Server Group.

Figure 26: Data Merging - Normal Case illustrates how measurements, alarms, and events are merged. MP servers merge to the Active SOAM server for the signaling network element they belong to. The Active SOAM server then replicates the data to its Standby and Spare servers.

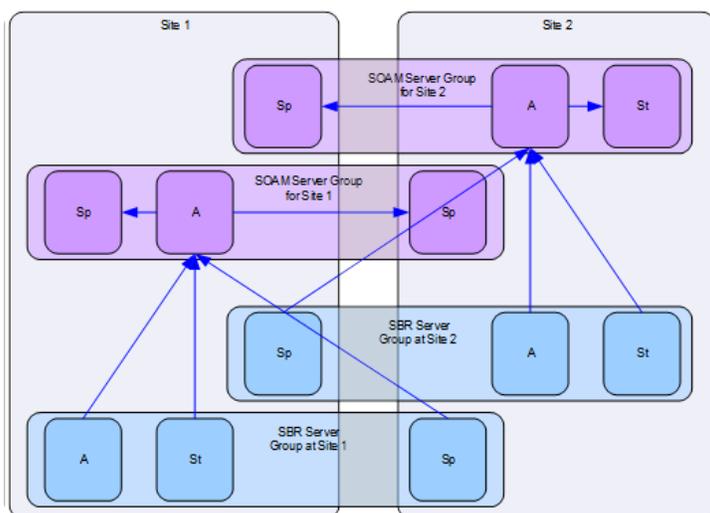


Figure 26: Data Merging - Normal Case

Figure 27: Data merging - Redundant Site Failure illustrates how a site failure affects merging of alarms, events, and measurements. When Site 2 fails, the servers at Site 1 that were marked as Preferred Spare are promoted to Active. The MP server that is now Active for the SBR Server Group for Site 2 will start merging its data to the SOAM server that is now Active for the SOAM Server Group for Site 2.

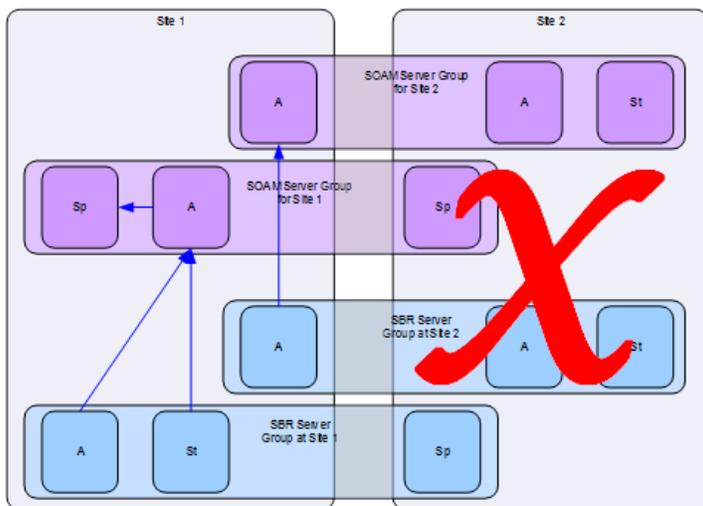


Figure 27: Data merging - Redundant Site Failure

IP Networking

The flexibility of the Diameter product results in many possible configurations for IP networking. This section focuses on IP network configurations that separate OAM functions from signaling functions such that signaling can continue to function normally if the OAM network is somehow disabled.

IP traffic is divided into categories called “Services”. For each Service, a network can be specified for both intra- and inter- Network Element IP traffic. [Table 12: IP Traffic-to-Service Mapping](#) illustrates a possible Services configuration for enabling signaling traffic from OAM traffic. In [Table 12: IP Traffic-to-Service Mapping](#), there are two physical networks, one for OAM traffic and one for signaling traffic. The signaling network is divided into two VLANs for separation of Diameter signaling from C-level replication and stack event signaling.

The OAM network is divided into intra-NE and inter-NE networks. Both signaling and OAM networks include a secondary path for HA heart-beating. (The secondary path for HA heart-beating was added to improve robustness for HA heart-beating going across WANs.) The primary path for HA heart-beating is always the same as the network used for replication.

Table 12: IP Traffic-to-Service Mapping

Traffic Type	Service Name	Intra-NE Network	Inter-NE Network
Signaling Traffic			
Diameter signaling	Signaling	Signaling VLAN 5	Signaling VLAN 5
Stack events sent between DA-MPs, between DA-MPs and SBRs, and between SBRs	ComAgent	Signaling VLAN 4	Signaling VLAN 4
Replication of data among DA-MPs	Replication_MP	Signaling VLAN 4	Signaling VLAN 4

Traffic Type	Service Name	Intra-NE Network	Inter-NE Network
Replication of data among SBRs	Replication_MP	Signaling VLAN 4	Signaling VLAN 4
HA Heartbeating among SBRs (Primary Path)	Replication_MP	Signaling VLAN 4	Signaling VLAN 4
HA Heartbeating among DA-MPs (Primary Path)	Replication_MP	Signaling VLAN 4	Signaling VLAN 4
HA Heartbeating among SBRs (Secondary Path)	HA_MP_Secondary	OAM VLAN 3	OAM VLAN 3
HA Heartbeating among DA-MPs (Secondary Path)	HA_MP_Secondary	OAM VLAN 3	OAM VLAN 3
OAM Traffic			
Replication of configuration data from NOAMs to SOAMs and from SOAMs to MPs	Replication	IMI	OAM VLAN 3
Merging of measurements, events, and alarms from MPs to SOAMs and from SOAMs to NOAMs	Replication	IMI	OAM VLAN 3
SNMP traps	Replication	IMI	OAM VLAN 3
SOAP Signaling	OAM	IMI	OAM VLAN 3
File Transfers to/from the File Management Area	OAM	IMI	
HA Heartbeating among OAM servers (Primary Path)	Replication	IMI	OAM VLAN 3
HA Heartbeating among OAM servers (Secondary Path)	HA_Secondary	Signaling VLAN 4	Signaling VLAN 4

PCA Routing

Routing of Diameter messages in a DSR with PCA activated is divided into ingress routing and egress routing. Ingress routing includes routing of Diameter requests and Diameter answers to the PCA application from a remote peer. Egress routing includes routing of Diameter requests and Diameter answers from PCA towards a remote peer. A remote peer can be a policy and charging client, a PCRF, an OCS or another DSR. Diameter requests can be initiated by both clients and PCRFs/OCSs.

Ingress Routing

This section describes how Diameter Request and Answer messages are routed to the PCA application.

Requests

Diameter Routing for Requests checks three conditions to determine whether to route a Request to a DSR Application:

1. Does the request include a DSR-Application-Invoked AVP, indicating that the request has already been processed and should not be processed again by a DSR application?

If this AVP is present, the request will not be routed to any DSR application. Otherwise the next condition is checked.

2. Does the request match a rule in the Application Routing Table (ART)?

If no rule is matched, the request is not routed to any DSR application. Otherwise the next condition is checked.

3. If the request matches an ART rule, is the application operational status for this DA-MP server set to Available?

If the DSR Application is not Available, then the "Unavailability action" is performed by Diameter. For PCA, the Unavailability action for Request is "Reject", which means PCA reject the Diameter Request messages by generating and sending an Answer message with a failure response with Result-Code AVP set to value configured for the error condition "PCA Unavailable or Degraded".

If the DSR Application is Available, then Diameter routes the Request to the DSR Application specified in the matching Application Routing Rule.

Ingress Requests are examined by Diameter to determine whether they should be routed to a DSR Application. The rules for deciding how to route ingress requests are defined in an Application Routing Table, or ART. [Table 13: P-DRA Application Routing Table Configuration](#) describes the expected configuration of Application Routing Rules for PCA. These rules will cause every Request that includes one of these values in the Application-Id in the Diameter Header to be routed to the PCA application. Some of these rules can be omitted, depending on which interfaces are used for PCA.

- The Rule Name can be any name that is chosen to identify the rule.
- The Priority is a value from 1 to 99 where 1 is the highest priority. Rules with higher priority will be evaluated for matching before rules with lower priority. Rules that overlap (i.e. one rule is more specific than another) should use the priority field remove ambiguity about which should be chosen. ART processing does not support "best match" semantics. Priority in an ART rule is an important factor to support DSR Applications interworking.
- Conditions can include Destination-Realm, Destination-Host, Application-Id, Command Code, Origin-Realm, and Origin-Host. If more than one condition is specified, the conditions are logically ANDed together to determine if a rule is matched.
- The Application Name will always be PCA for the Policy and Charging application. PCA will show up in the Application Name drop-down only if the PCA application has been activated

Table 13: P-DRA Application Routing Table Configuration

Rule Name	Priority	Conditions	Application Name
PcaGxRule	5	AppId Equal 16777238 - 3GPP Gx	PCA
PcaGxxRule	5	AppId Equal 16777266 - 3GPP Gxx	PCA
PcaRxRule	5	AppId Equal 16777236 - 3GPP Rx	PCA

Rule Name	Priority	Conditions	Application Name
PcaGxPRule	5	AppId Equal 16777238 - 3GPP Gx-Prime	PCA
PcaS9Rule	5	AppId Equal 16777267 - 3GPP S9	PCA
PcaRoGyRule	5	AppId Equal 4 - 3GPP	PCA

Answers

Diameter Answers, to which the related Diameter requests have been processed successfully by PCA, will be forwarded to PCA application for processing.

If PCA has requested to receive an answer, but the PCA has an operational status of "If PCA becomes Unavailable when an Answer" is received, the answer message will be relayed directly to the remote peer.

Note: Relaying an answer while the PCA application is unavailable may result in exposing a PCRF name that was supposed to be topology hidden for the P-DRA function

Note: All the Online Charging answer messages are always processed by the Online Charging DRA.

Egress Routing

Egress Answer messages are always routed according to the Connection-Id and Diameter Hop-By-Hop-Id of the Request they are answering.

PCRF Selection for New Bindings

For the P-DRA function, when a binding capable session initiation message (CCR-I) arrives for an IMSI that is not already bound to a PCRF, the P-DRA function selects a PCRF from the list of adjacent PCRFs configured on **Policy and Charging -> Configuration -> Policy DRA -> PCRFs** GUI page. This list of PCRFs generally contains only PCRFs that are local to the site with the P-DRA node. PCRFs that are local to the PCA node's mate should generally not be included. The reason to include only local PCRFs is to avoid the extra latency associated with selection of a PCRF separated across a WAN from the policy client that initiated the session.

If the PCRF has different hostnames for different 3GPP interfaces (e.g. Gx, Rx, Gxx, S9), only the binding capable hostnames should be configured on the **Policy and Charging -> Configuration -> Policy DRA -> PCRFs** GUI page.

Policy DRA uses a round-robin selection algorithm to distribute new bindings across the set of configured PCRFs. The round-robin algorithm runs independently on each DA-MP server, so predicting the next PCRF that will be used is difficult on a PCA node that has policy client connections to multiple DA-MP servers. In addition, the round-robin selection algorithm is executed for each CCR-I received, causing the next PCRF to be updated, even if the CCR-I is for a subscriber that already has a binding.

PCRF Selection for Existing Bindings

A binding becomes finalized when a successful CCA-I is received from a PCRF for a given subscriber. At this point, all Policy sessions for that subscriber must be routed to that PCRF Peer Node, or a Peer Node that shares state with the bound Peer Node. The subscriber remains bound to this PCRF until all of the subscriber's binding capable sessions (Gx, Gxx, S9) are terminated.

The architecture for many PCRFs is such that a single Diameter host is not a single point of failure for a subscriber's Policy sessions. This is generally accomplished by designating a set of Diameter hosts that all share a common database and can therefore all access the subscriber's Policy data and Resource usage.

If the PCRF supports multiple Diameter hosts that share state, routing can be set up as follows:

- A Peer Routing Rule that matches the Destination-Host equal to the bound PCRF name
- A Route List that has a Primary and a Secondary Route Group
 - The Primary Route Group routes only to the bound PCRF
 - The Secondary Route Group distributes across all PCRF Peers that share state with the bound PCRF.

Some PCRFs also have different Diameter hosts for different 3GPP interfaces. For example, they may have a hostname for Gx and a different hostname for Rx. This can be accommodated by splitting the PRT entry above into two entries as follows:

- A Peer Routing Rule that matches the Destination-Host equal to the bound PCRF name and Application-Id equal to Gx (16777238).
- A Peer Routing Rule that matches the Destination-Host equal to the bound PCRF name and Application-Id equal to Rx (16777236).

Routing In-Session Messages Without Topology Hiding

For the P-DRA function, when the PCRF name is not topology hidden, the policy client is expected to learn the PCRF name from the Origin-Host and Origin-Realm of the answer to the session initiation request (e.g. CCA-I or AAA). This PCRF name should be used as the Destination-Host and Destination-Realm of all subsequent in-session requests originated by the policy client.

Policy clients that are proxy-compatible (can learn the PCRF name) allow P-DRA to host-route in-session requests without the need for any binding or session database lookup. This behavior is desirable because it reduces the cost of the P-DRA by reducing the number of SBR servers needed to support a given Diameter traffic load.

There are, however, policy clients that are not proxy-compatible. Many of these always omit the Destination-Host AVP from requests, or worse, include the Destination-Host AVP with the PCA Diameter hostname. In order to support such policy clients, the P-DRA function must be configured to add or replace the Destination-Host and Destination-Realm of all requests with the PCRF that the subscriber is bound to. This can be accomplished by setting table PdraEngdValues entry CheckSessionForAllBindCapMessages value to the number one (1). This value defaults to zero (0), meaning that the P-DRA function does not replace the Destination-Host for in-session messages by default. Policy clients that are not proxy-compatible can also be accommodated by enabling topology hiding

Routing In-Session Message with Topology Hiding

When topology hiding is enabled, the PCRF name is hidden from the applicable policy client. If the PCRF name is hidden from the policy client, obviously the policy client cannot use the PCRF as the Destination-Host and Destination-Realm in its in-session requests. When topology hiding is in force for a policy client, PCA must route in-session requests to the bound PCRF by performing a session record lookup and using the PCRF information stored in the session record.

Use of topology hiding is expensive in terms of the increased stack event processing required and the increased latency required to lookup the bound PCRF in the session record. For these reasons, topology

hiding should be scoped as narrow as possible. For example, if topology should be hidden from only a few policy clients, choose the per policy client topology hiding scope instead of choosing to hide topology from all policy clients.

Topology hiding can also be used to "work around" a policy client that does not have the ability to learn the PCRF name (i.e. is not proxy-compatible). Turning on topology hiding for a subset of policy clients is more efficient than using the `CheckSessionForAllBindCapMessages` option

OCS Selection and Routing

When a Gy/Ro session-initiation request (i.e. CCR-I) is received at PCA, the OC-DRA function selects an OCS server among a collection of OCSs that are connected to the PCA DSR directly. The OC-DRA function selects an OCS based on one of the configured OCS selection mode:

- Single OCS Pool Mode
- Multiple OCS Pools Mode

If the "Single OCS Pool" mode is configured, OC-DRA removes the Destination-Host AVP, if present, from the session initiation request and forwards the message to DRL. PRT/RL will be used to route the request message to one of the OCS servers connected to the DSR based on some round-robin load balance algorithm.

If the "Multiple OCS Pools" mode is configured, OC-DRA forwards the session initiation request without any modification to DRL. DRL may use the Destination-Host info in the request message to match the PRT/RL to route the message to an OCS pool and then an OCS within the pool using priorities/weights configured in the Route List selected via PRT.

The decision of choosing one OCS pool mode over the other may be made by the assumption if the regionalized routing is used or not. The configuration of the multiple OCS pool mode may be based on the assumption that the DSR RBAR application is invoked prior to OC-DRA invocation. In this case, a correct Destination-Host AVP and/or Destination-Realm AVP may have been identified and populated in the session initiation requests by RBAR before forwarded to PCA. On the other hand, the configuration of the single OCS pool mode may be based on the assumption that RBAR is not invoked for processing the message beforehand. However, the regionalized routing and OCS mood selection are independently configured that it is quite possible the assumptions mentioned above may not be true. Therefore, the OC-DRA function should work properly on any combination of the following configurations:

- Single OCS Pool mode is configured, PCA is invoked without RBAR chaining,
- Single OCS Pool mode is configured, RBAR is invoked before OC-DRA receives the message,
- Multiple OCS Pool mode is configured, RBAR is invoked before OC-DRA receives the message,
- Multiple OCS Pool mode is configured, PCA is invoked without RBAR chaining

Naming Conventions and Hierarchical Routing

When PCA is deployed in large networks with multiple PCA mated pairs, the DRL routing tables can be greatly simplified by employing some simple naming conventions. For example, naming all clients and PCRFs/OCSs local to a particular PCA node such that they start with a common prefix allows PRT rules like "Destination-Host Starts-With xxx", where xxx is the site prefix for that PCA node. The "Starts-With" rule will point to a route list that routes to the PCA node where the equipment is located. Then if a new client or a PCRF/OCS is added at a given PCA node, routing changes are needed only at that node and that node's mate, which have peer node entries and Diameter connections (i.e. are adjacent) to the new client or PCRF/OCS. PCA nodes that are non-adjacent do not require any routing updates.

PCA Data Auditing

P-DRA Binding/Session Database

In most cases, Binding and Session database records are successfully removed as a result of signaling to terminate Diameter sessions. There are, however, instances in which signaling incorrectly removed a session and did not remove a database record that should have been removed. The following cases can result in stale Binding or Session records:

- No Diameter session termination message is received when the UE no longer wants the session.
- IP signaling network issues prevent communication between MPs that would have resulted in one or more records being deleted.
- SBR congestion could cause stack events to be discarded that would have resulted in removal of a Binding or Session record.

To limit the effects of stale Binding and Session records, all SBRs that own an active part of the database continually audit each table to detect and remove stale records. The audit is constrained by both minimum and maximum audit rates. The actual rate varies based on how busy the SBR server is. Audit has no impact on the engineered rate of signaling.

Generally, SBR servers are engineered to run at 80% of maximum capacity. The audit is pre-configured to run within the 20% of remaining capacity. Audit will yield to signaling. Audit can use the upper 20% only if signaling does not need it.

Binding table audits are confined to confirming with the Session SBR that the session still exists. If the session exists, the record is considered valid and the audit makes no changes. If the session does not exist, however, the record is considered to be an orphan and is removed by the audit. The Binding Audit Session Query Rate is the maximum rate at which a Binding SBR can send query messages to session servers to verify that sessions are still valid. This audit rate is configurable from the **Policy and Charging > Configuration > General Options** GUI page so that the audit maximum rate can be tuned according to network traffic levels.

Session table audits work entirely based on valid session lifetime. When a session is created, it is given a lifetime for which the session will be considered to be valid regardless of any signaling activity. Each time an RAA is processed, the lifetime is renewed for a session. The duration of the lifetime defaults to 7 days, but can be configured in one of two ways:

- A session lifetime can be configured per Access Point Name using the **Policy and Charging > Configuration > Access Point Names** GUI page.
- The Audit Operation Rate is configurable (with a default of 50,000) from the **Policy and Charging > Configuration > General Options** GUI page and depends on if Session or Binding SBRs are being audited:
 - For Session SBRs, the maximum rate at which Diameter sessions are checked for staleness
 - For Binding SBRs, the maximum rate at which binding session references are examined, if not already throttled by the Binding Audit Session Query Rate

If the SBR signaling load plus the audit load cause an SBR server to exceed 100% capacity, that SBR server will report congestion, which will cause an automatic suspension of auditing. Any SBR on which audit is suspended will have minor alarm 22715 to report the suspension. The alarm is cleared only when congestion abates.

- Local congestion refers to congestion at the SBR server that is walking through Binding or Session table records. Suspension of audit due to Local Congestion applies to both the binding audit and session audit.
- Remote congestion refers to congestion at one of the Session SBR servers that a Binding SBR server is querying for the existence of session data (using sessionRef). Suspension of audit due to Remote Congestion only applies to binding SBR servers because only binding SBRs send stack events to session SBR servers, while session SBR servers do not.
- Enhanced Suspect Binding Rate Control can also cause congestion. Suspension of audit due to Enhanced Suspect Binding Rate Control applies only to the binding audit.

When an SBR server starts up (i.e. SBR process starts), or when an SBR's audit resumes from being suspended, the audit rate ramps up using an exponential slow-start algorithm. The audit rate starts at 1500 records per second and is doubled every 10 seconds until the configured maximum audit rate is reached.

Note: If the binding audit resumes after a recovery from remote congestion, the slow-start algorithm is not applied.

In addition to the overall rate of record auditing described above, the frequency at which a given table audit can be started is also controlled. This is necessary to avoid needless frequent auditing of the same records when tables are small and can be audited quickly. A given table on an SBR server will be audited no more frequently than once every 10 minutes.

In order to have some visibility into what the audit is doing, the audit generates Event 22716 "SBR Audit Statistics Report" with audit statistics at the end of each pass of a table. The format of the report varies depending on which table the audit statistics are being reported for.

PCA Configuration Database

A number of Policy and Charging configuration database tables, i.e. PCRFs, Policy Clients, OCSs and CTFs are configured at the SOAM but contain data that are required network-wide. The site-wide portions of the data are stored at the SOAM servers. The network-wide portions of the data are stored globally at the NOAM. Due to the distributed nature of this data (the split between SOAM and NOAM), there is a PCA Configuration Database Audit which executes in the background to verify that all the related configuration tables for this data are in sync between SOAMs and the NOAM.

The PCA Configuration Database Audit executes on the SOAM periodically every 30 seconds in the background and will audit all the related configuration tables between SOAM and NOAM for PCRFs, Policy Clients, OCSs and CTFs. If the audit detects that there are any discrepancies among these tables, it will automatically attempt to resolve the discrepancies and validate that they are back in sync.

The configuration database can get out of sync due to a database transaction failure or due to operator actions. If an operator performs a database restore at the NOAM using a database backup that does not have all the network-wide data corresponding to the current SOAM configuration, then the database will not be in sync between SOAM and NOAM. Similarly, if an operator performs a database restore at an SOAM using a database backup that does not have the configuration records corresponding to network-wide data stored at the NOAM, then the database again will not be in sync. The audit is designed to execute without operator intervention and correct these scenarios where configuration data is not in sync between SOAM and NOAM.

If the audit fails to correct the database tables, the audit will assert Alarm 22737 (Configuration Database Not Synced). The audit continues to execute periodically every 30 seconds to attempt to correct the database tables. If the audit successfully corrects and validates the tables during an audit pass, it will clear Alarm 22737.

Note: All statements about database tables in this section only apply to configuration tables related to PCRFs, Policy Clients, OCSs and CTFs because the PCA Configuration Database Audit executes only on the database tables where it is necessary for the data to be split across SOAM and NOAM.

OC-DRA Session Database

The Session Database Audit is enhanced to detect and remove stale binding independent session (i.e., Gy/Ro session) data stored in the Session SBR. Session state maintained in the Session SBR for Gy/Ro session-based credit-control is considered stale when a CCR/CCA-U or RAR/RAA has not been exchanged for the session for a length of time greater than or equal to the Stale Session Timeout value (in hours) as configured by the Network OAM GUI. If the binding independent session is associated with an APN configured in the Network OAM GUI **Main Menu > Policy and Charging > Configuration > Access Point Names**, then the Stale Session Timeout value associated with the APN is used. Otherwise, the default Stale Session Timeout value configured in the Network OAM GUI **Main Menu > Policy and Charging > Configuration > General Options** is used.

Stale Gy/Ro sessions can occur for various reasons:

- OC-DRA did not receive the Diameter Credit-Control Session Termination Request (CCR-T) message from the OCS when the Gy/Ro session was to be terminated due to IP signaling network issues.
- Session SBR did not receive the findAndRemoveOcSession stack event from OC-DRA to find and remove the Gy/Ro session due to IP signaling network issues.
- Session SBR received the findAndRemoveOcSession stack event from OC-DRA, but discarded it due to congestion.
- Session SBR database access errors
- Internal software errors

PCA and Application Chaining

The PCA and RBAR chaining function provides the needed capability to enable operators to perform "regionalized routing" in such a way that a policy or charging server (e.g. a PCRF or an OCS) will only serve the subscribers whose subscriber identities, i.e. IMSIs or MSISDNs, are within the range of the IDs that has been assigned to this PCRF or OCS.

Some Diameter signaling networks may need to be segmented based on the ranges of the subscriber identities such as IMSI and MSISDN and associate the subscriber ID ranges to Diameter servers (HSS, OCS and PCRF etc.). With such a subscriber ID range <-> Diameter server mapping, a subscriber can be served by a pre-determined Diameter server or a group of servers such that all messages with this subscriber's ID (IMSI or MSISDN or both) will be routed to the pre-determined Diameter servers consequently. The routing based on the subscriber ID <-> Diameter server mapping is referred to as "regionalized routing"

It may also be necessary to be able to "bind" a subscriber to a policy server (e.g. PCRF) or correlate sessions such that all messages on behalf of the subscriber can be routed to the same PCRF regardless what the Diameter interfaces are used. DSR with Policy DRA functionality provides subscriber <-> PCRF binding capability. The same will occur for Online Charging DRA (OC-DRA) in the segmented network.

P-DRA may receive the incoming Diameter request messages over binding capable interfaces (i.e. Gx and Gxx) or over binding dependent interfaces (i.e. Rx and Gx-Prime). The RBAR and PCA application chaining function is applicable ONLY on binding capable interfaces for the P-DRA feature, and binding

independent interface (Ro/Gy) for the OC-DRA feature, but NOT on binding dependent interfaces. Specifically, the Diameter request messages over binding dependent interfaces (Rx or Gx-Prime) intending for being processed by P-DRA should never be routed to RBAR for address resolution.

With respect to DSR application chaining, an accessing region is a DSR network segment where the DSR has a direct connection to a policy and charging client who initiates and sends a Diameter request directly to the DSR and a serving region is a DSR network segment where the PCA (either P-DRA or OC-DRA functionality) actually receives and processes the Diameter request message as forwarded by DRL. Accessing region and serving region are all relative concepts, which make sense only relevant to a specific policy and charging client and a specific DSR application (i.e. PCA). A serving region can be an accessing region as well for a policy and online charging client and PCA application, i.e. the DSR that receives a Diameter requests hosts the PCA that processes the Diameter requests.

Request messages over binding capable interfaces (Gx/Gxx) and binding independent interfaces (Gy/Ro) are subject to RBAR and PCA application chaining while request messages over binding dependent interfaces (Rx or Gx-Prime) are not. Consequently, the topology hiding for Rx or Gx-Prime session of a subscriber may be performed by a P-DRA/DSR that is different from the P-DRA/DSR that has created the binding for the subscriber.

For the P-DRA function, the different treatment of binding capable and binding dependent sessions in the regionalized routing situation results in different requirements for topology hiding configuration. The configuration for enabling/disabling topology hiding and for the scope of topology hiding will be done on the NOAM GUI, which allows the management of the topology hiding to be handled on the NOAM level for all the P-DRA/DSRs within the same NOAM network. While the policy clients that are subject to topology hiding handling are still be configured on the SOAM, the configured data on the SOAM and NOAM will be communicated to each other such that a complete list of policy clients from all SOAMs can be consolidated on the NOAM. The consolidated list of policy clients for topology hiding can then be replicated to each of the SOAMs.

Note: For the OC-DRA function, topology hiding is not supported.

The Communication Agent

The Communication Agent (ComAgent) enables reliable communication between Policy and Charging DRA and SBRs and among SBRs in a scalable and high available PCA network. [Figure 28: Communication between ComAgents, Policy DRA, and SBR](#) depicts the communication paths between the Policy and Charging DRA, the SBR, and their ComAgents, and the communication paths between the ComAgents.

Note: The DTL uses ComAgent to transmit TTRs to DIH. The Diameter Troubleshooting Layer (DTL) is a component of the Diameter plug-In architecture that transmits TTRs to DIH.

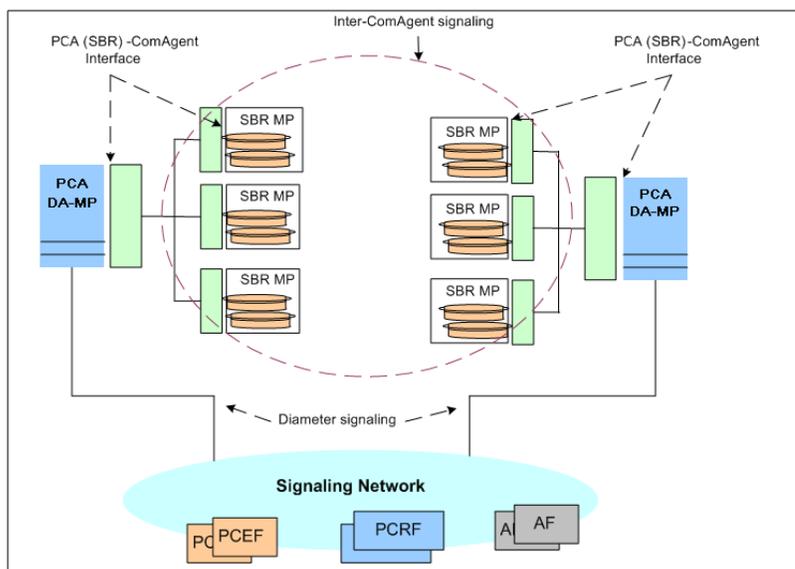


Figure 28: Communication between ComAgents, Policy DRA, and SBR

The ComAgent Direct Routing service, HA service, and the MP Overload Management Framework are used by the Policy and Charging DRA and SBR for communication and for SBR congestion control. (See [SBR Congestion](#) for information about the MP Overload Management Framework.)

Policy and Charging DRA automatically establishes TCP connections between all of the servers that need to communicate with the database. The following connections are established:

- All DA-MPs in the network connect to all binding SBRs in the network.
- All session SBRs in the network connect to all binding SBRs in the network
- All DA-MPs in a mated pair connect to all session SBRs in the mated pair

You can view and manage these connections using the ComAgent Connection Status GUI at the NOAM: **Communication Agent > Maintenance > Connection Status**. There is also a ComAgent HA Service Status, but the same information can be obtained from the **Policy and Charging > Maintenance > SBR Status** page.

Diameter Routing and Communication with PCA

The PCA Application uses the DSR Application Infrastructure (DAI), which provides a mechanism for Diameter messages routing and for status updates between the Diameter Routing Function and the DAI.

[Table 14: Communication between the Diameter Routing Function and the DAI](#) describes two functions for communication between the Diameter Routing Function and the DAI.

Table 14: Communication between the Diameter Routing Function and the DAI

Function	Communication Direction	Description
Application Data	PCA <-> Diameter Routing Function	Either a Request or an Answer with supporting information
Application Status	PCA <-> Diameter Routing Function	The PCA Operational Status of Available, Degraded, or Unavailable

Request Routing

As shown in the following figure, the Diameter Request messages are routed from the Diameter Routing Function to the PCA based on the configured Application Routing Rule, and routed from the PCA to the Diameter Routing Function, all using the Application-Data function. The PCA will return the Request to the Diameter Routing Function for Peer Routing Rule processing and routing.

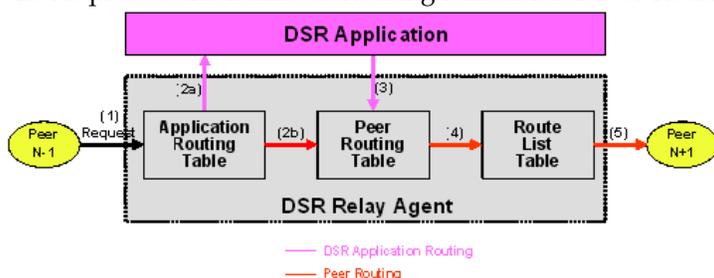


Figure 29: Request Processing at the Diameter Routing Function and PCA

Answer Routing

When the PCA forwards a Request message to the Diameter Routing Function for routing, it must inform the Diameter Routing Function how to process the corresponding Answer. It can inform the Diameter Routing Function either to route the Answer to the PCA or to route the Answer to the downstream Peer without involving the PCA. *Figure 30: Answer Processing at the Diameter Routing Function and PCA* shows the case where an Answer is transmitted back to the PCA. After the PCA completes processing of the Answer, it will send it to the Diameter Routing Function for transmission to the Diameter Transport Function so that it can be routed to the downstream Peer.

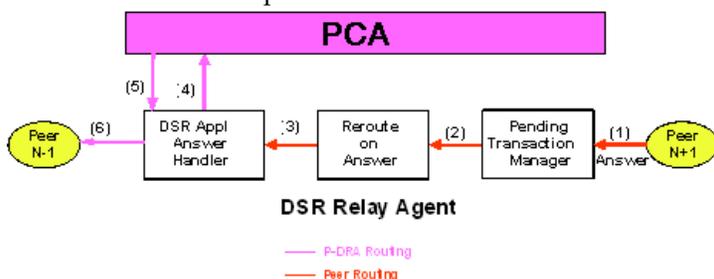


Figure 30: Answer Processing at the Diameter Routing Function and PCA

PCA Generated Answer

In some cases, the PCA needs to generate an Answer message in response to an incoming Request. For example, the Policy DRA function cannot find a PCRF to route the Request message to. *Figure 31: PCA Generated Answer Routing* shows the Diameter Routing Function routing for this scenario.

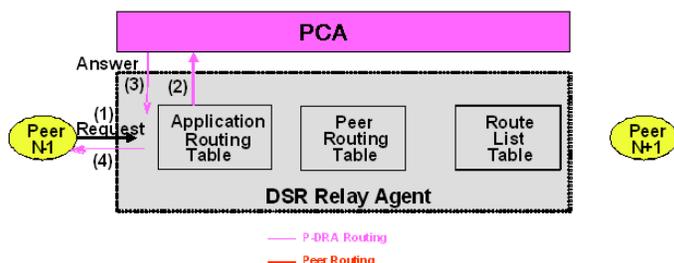


Figure 31: PCA Generated Answer Routing

PCA Generated Request

In some cases, the PCA needs to generate Diameter Requests. *Figure 32: PCA Generated Request Routing* shows the Diameter Routing Function routing for this scenario.

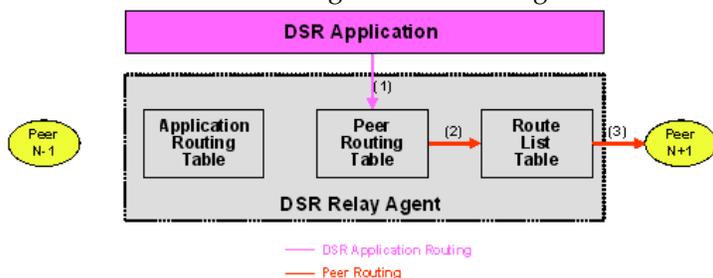


Figure 32: PCA Generated Request Routing

Policy DRA Function Use Cases

The following typical Policy DRA signaling use cases demonstrate the Policy DRA and SBR capabilities to establish subscriber binding to some PCRF, and update and terminate the sessions when requested:

- Binding and Session Creation and Session Termination over the Gx Interface - A Policy Client requests to bind a subscriber for policy provisioning over a Gx interface. The Policy DRA creates the binding to a selected PCRF, generates the binding and session records in the Policy SBR database, updates the session as requested, and eventually terminate session as requested.
- Subscriber Session Creation and Termination over the Rx Interface - A Diameter Request is sent to the Policy DRA over the Rx interface for the same subscriber that has established a binding with the PCRF over the Gx interface. The Policy DRA coordinates the sessions over the Gx and Rx interfaces and routes the Diameter messages to the same PCRF.
- Policy DRA in Roaming Scenarios - In addition to communicating to the Policy Clients and Policy servers through Gx/Gxx and Rx interfaces in their own networks, the Policy DRAs can communicate to each other across the Visited Access and Home Access networks through the S9 interface, for session binding purposes. See *Policy DRA in Roaming Scenarios*.

PCA and IDIH Metadata

The Diameter Routing Function and invoked DSR Applications record detailed information about each Diameter transaction - called transaction metadata. Each metadata record describes an important event in the lifetime of a Diameter transaction. Metadata appears in the Trace Transaction Record (TTR) in the order that the metadata-generating events actually occurred. Together, all of the metadata records combine to document the processing performed on the entire transaction, and can later be used to provide diagnostic information when performing troubleshooting. Metadata is recorded to a TTR for each transaction so that, even if the transaction is selected to be sent to IDIH at an Answer Troubleshooting Trigger Point (TTP-IA or TTP-EA), the metadata for all of the messages in the transaction will be present.

PCA will record the Application-specific metadata events described in [Table 15: PCA Metadata-Generating Events](#).

Table 15: PCA Metadata-Generating Events

Event	Instance Data	When Recorded
PCA Function Invoked	PCA Function Name	When an application with multiple functionality receives an ingress Diameter message (including both requests and answers) and routes it to one of its functions for processing Note: This metadata is only recorded if the application function is enabled and available processing messages.
PCRF Pool Selected	<ul style="list-style-type: none"> • PCRF Pool Name • PCRF Sub-Pool name • Sub-Pool selection rule name 	When P-DRA receives a binding-capable session initiation request
Binding Query Request Sent	<ul style="list-style-type: none"> • Anchor key • APN name • PCRF Pool name • Session reference 	When P-DRA sends a "Find or Create Binding" stack event to a SBR
Binding Query Result Received	<ul style="list-style-type: none"> • SBR IP Address (e.g., 10.240.55.25) • Result code • Binding state • Binding master • Master session reference • PCRF FQDN • Suspect duration 	When P-DRA receives a "Find or Create Binding Result" stack event from a SBR

Event	Instance Data	When Recorded
Topology Hiding Applied	N/A	When P-DRA receives a bind-capable or binding-dependent session initiation/in-session answer destined for a peer for which topology hiding is configured
Create Session Request Sent	<ul style="list-style-type: none"> • Session ID • Session reference • Anchor key • MSISDN • IPv4/IPv6 address key • PCRF FQDN 	When P-DRA sends a "Create Session" stack event to a SBR
Create Session Result Received	<ul style="list-style-type: none"> • SBR IP Address (e.g., 10.240.55.25) • Result code 	When P-DRA receives a "Create Session Result" stack event from a SBR
Update Binding Request Sent	<ul style="list-style-type: none"> • Operation • Anchor key • Final PCRF FQDN • Session reference 	When P-DRA sends an "Update Binding" stack event to a SBR
Update Binding Result Received	<ul style="list-style-type: none"> • SBR IP Address (e.g., 10.240.55.25) • Result code 	When P-DRA receives an "Update Binding Result" stack event from a SBR
Find Binding Request Sent	<ul style="list-style-type: none"> • Key type • Key value • APN name 	When P-DRA sends a "Find Binding" stack event to a SBR
Find Binding Result Received	<ul style="list-style-type: none"> • SBR IP Address (e.g., 10.240.55.25) • Result code • IMSI • PCRF FQDN 	When P-DRA receives a "Find Binding Result" stack event from a SBR
Refresh Session Request Sent	Session ID	When P-DRA sends a "Refresh Session" stack event to a SBR
Refresh Session Result Received	<ul style="list-style-type: none"> • SBR IP Address (e.g., 10.240.55.25) • Result code 	When P-DRA receives a "Refresh Session Result" stack event from a SBR
Delete Session Request Sent	Session ID	When P-DRA sends a "Remove Session" stack event to a SBR
Delete Session Result Received	<ul style="list-style-type: none"> • SBR IP Address (e.g., 10.240.55.25) • Result code 	When P-DRA receives a "Remove Session Result" stack event from a SBR

Event	Instance Data	When Recorded
	<ul style="list-style-type: none"> • Session reference • PCRF FQDN • Anchor key • MSISDN key • IPv4/IPv6 key 	
Find Session Request Sent	Session ID	When P-DRA sends a "Find Session" stack event to a SBR
Find Session Result Received	<ul style="list-style-type: none"> • SBR IP Address (e.g., 10.240.55.25) • Result code • Session reference • PCRF FQDN 	When P-DRA receives a "Find Session Result" stack event from a SBR
Remove Suspect Binding Request Sent	<ul style="list-style-type: none"> • Anchor key • PCRF FQDN 	When P-DRA sends a "remove Suspect Binding" stack event to a SBR
Remove Suspect Binding Result Received	<ul style="list-style-type: none"> • SBR IP Address (e.g., 10.240.55.25) • Result code 	When P-DRA receives a "Remove Suspect Binding Result" stack event from a SBR
Session Release Initiated	Application Name	When an "Update Binding" request, a "Create Session" request or, a "Create Alternate Key" request fails
Session Query Initiated	Application Name	When a stale Gx session is detected by a SBR
Routing Exception	<ul style="list-style-type: none"> • Routing Exception Type (e.g., "SBR Congestion") • Routing Exception Action (e.g., "Abandon Request") 	After any routing exception is encountered
SBR Request Failure	<ul style="list-style-type: none"> • After any routing exception is encountered • Resource name • Sub-resource ID • Failed Request Name 	When a PCA Function fails to send a request to the SBR
SBR Response Timeout	<ul style="list-style-type: none"> • Resource name • Sub-resource ID 	When a PCA Function times out waiting to receive a response from a SBR for a previous request
Routing Error Indication Received	Routing Error	When a PCA Function initiates a Diameter request (Session Release RAR) that is rejected by DRL due to a routing error.

Event	Instance Data	When Recorded
		Note: The Routing Error recorded is the Error-Message AVP value of the Answer message initiated by DRL.
Create OC Session Request Sent	<ul style="list-style-type: none"> • Session ID • CTF Realm • CTF FQDN • OCS Realm • OCS FQDN • Subscriber ID • APN Name 	When OC-DRA sends a "Create OC Session" stack event to the Session SBR
Create OC Session Result Received	<ul style="list-style-type: none"> • SBR IP Address • Result Code 	When OC-DRA receives a "Create OC Session Result" stack event from the Session SBR
Find and Refresh OC Session Request Sent	Session ID	When OC-DRA sends a "Find and Refresh OC Session" stack event to the Session SBR
Find and Refresh OC Session Result Received	<ul style="list-style-type: none"> • Session ID • Result Code • CTF Realm • CTF FQDN • OCS Realm • OCS FQDN • Subscriber ID • APN Name 	When OC-DRA receives the "Find and Refresh OC Session Result" stack event from the Session SBR
Find and Remove OC Session Request Sent	Session ID	When OC-DRA sends a "Find and Remove OC Session" stack event to the Session SBR
Find and Remove OC Session Request Received	<ul style="list-style-type: none"> • SBR IP Address • Result Code • CTF Realm • CTF FQDN • OCS Realm • OCS FQDN • Subscriber ID • APN Name 	When OC-DRA receives the "Find and Removed OC Session Result" stack event from the Session SBR

The metadata captured by IDIH for the PCA includes the results of each query that PCA makes to the session and binding database and the associated result. Whenever the result of a database query is captured in PCA metadata, it will include the identity of the specific server that generated the response.

The following are key concepts for PCA as it relates to IDIH:

- IDIH can display traces in multiple formats (for example, two- or three-way split screen or single screen). Because it is very difficult to display all of the information in a single screen, output columns slide out of view. Sliders allow the column views to be manipulated.
- There are two basic ways to view Event and metadata information:
 - A graphical display (for example, ladder and object)
 - An event list, which provides a listing of events

In graphical display mode, click on the bubble to view decode information for messages. To see metadata information, hover over it (for example, the PDRA bubble), and then use the slider to move up or down to see the information linked to that event. In event list mode, select the message or slide to the right to view the event/metadata information.

The following examples illustrate the type and format of information that is collected and used from Policy DRA IDIH traces.

Note: Although these examples are for Policy DRA, information collected and used from Online Charging DRA IDIH traces is similar in presentation.

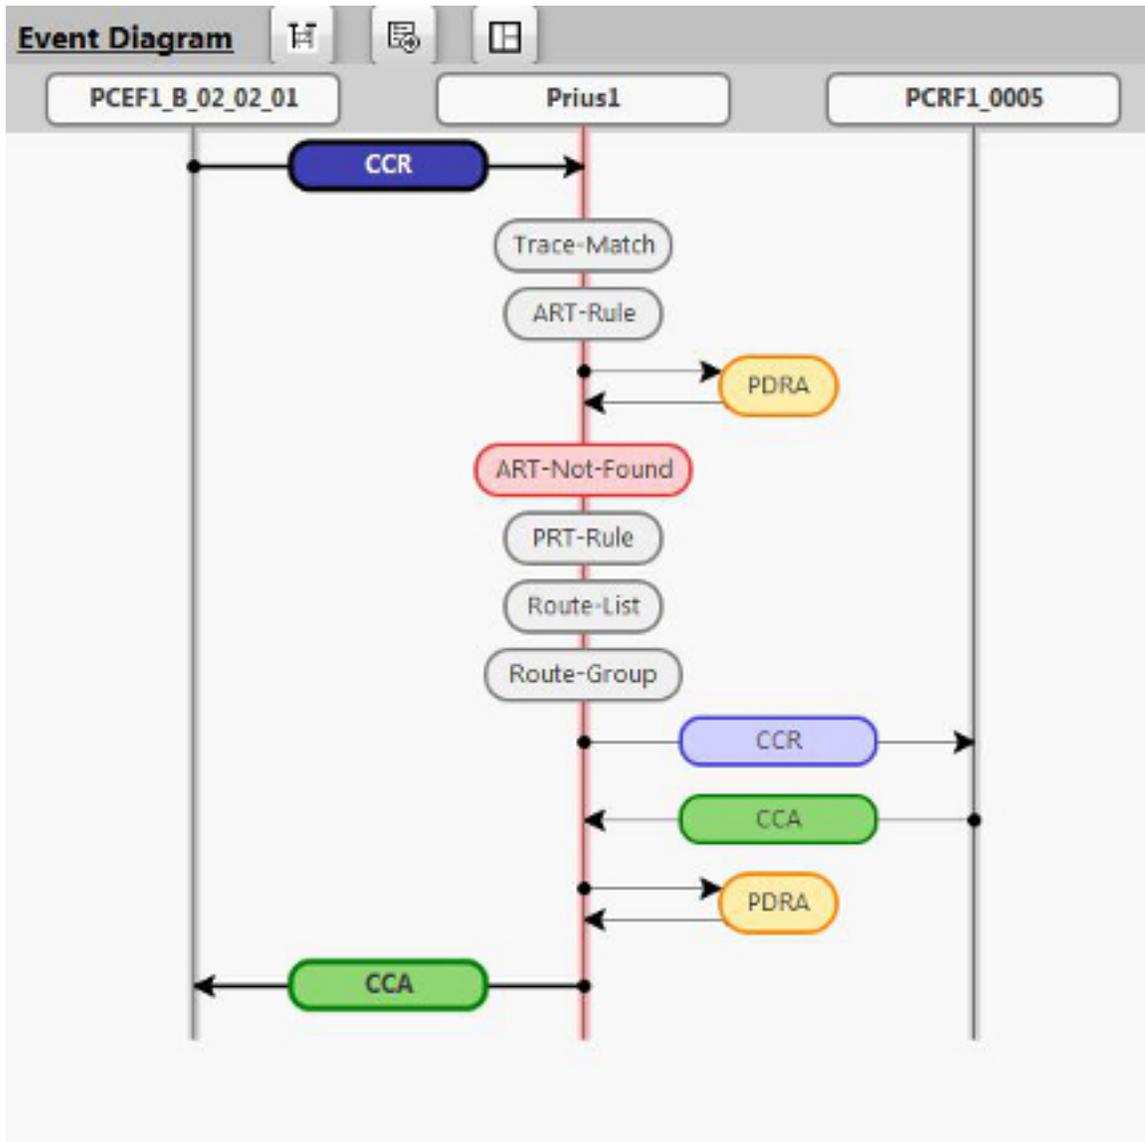


Figure 33: Event Diagram Trace - CCR Example

Rec #	Time	Event Type	Event Data
1	03/07/2014 09:47:30:000	Message Received	-
2	-	Trace Match	INIT, TTP-IR
3	-	ART Rule Match	rule=Gx, table=GXRule1
4	-	Routing To App	PDRA
5	-	PCRF Pool Selected	PCRF Pool Name=Default PCRF Sub-Pool Name=[Sub-Pool S
6	-	Binding Query Request Sent	Anchor Key=199754193097563 APN Name=APN1-B-02-02-01
7	-	Binding Query Result Received	pSBR IP Address=10.240.142.8 (IPv4) Binding State=Early Binc
8	-	App To Routing	requestRoutingAction=route via PRT rule overrideOriginRequ
9	-	ART Rule Not Found	GXRule1
10	-	PRT Rule Match	rule=Default_catch_all, table=Default
11	-	Route List Selected	Default_RL1
12	-	Route Group Selected	Default_RG1
13	03/07/2014 09:47:30:000	Message Sent	-
14	03/07/2014 09:47:30:000	Message Received	2001 - DIAMETER_SUCCESS
15	-	Routing To App	PDRA
16	-	Update Binding Request Sent	Operation=Update PCRF Anchor Key=199754193097563 Fina
17	-	Create Session Request Sent	Session Reference=f25eb553805e22c368032d00 Anchor Key-
18	-	App To Routing	requestRoutingAction=unknown routing action overrideOrigin
19	03/07/2014 09:47:30:000	Message Sent	2001 - DIAMETER_SUCCESS
20	-	Create Session Result Received	pSBR IP Address=172.18.24.14 (IPv4) Result Code=Success
21	-	Update Binding Result Received	pSBR IP Address=10.240.142.8 (IPv4) Result Code=Success

Figure 34: Update Request Policy DRA Example

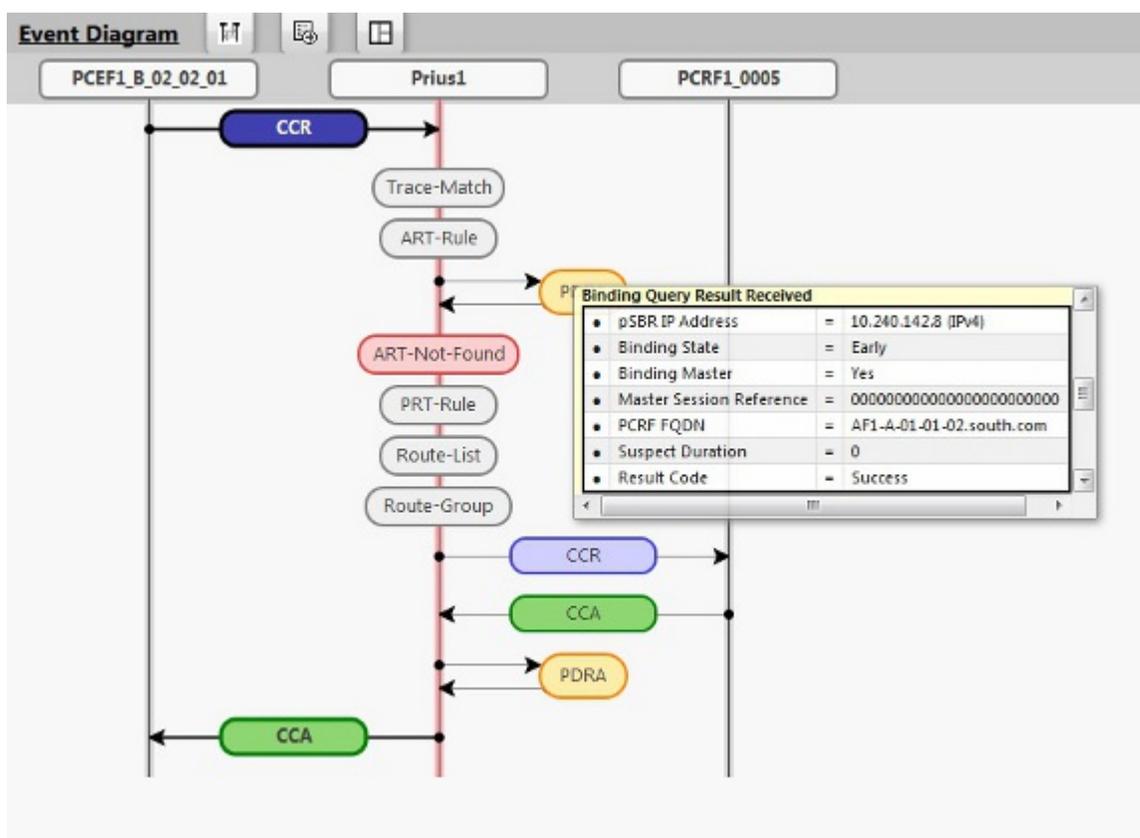


Figure 35: Event Diagram - Hover (Mouse-over) Example

PCA Capacity Constraints

PCA has the following engineered capacity constraints:

Constraint Name	Value
Maximum managed objects per Network	
Local PCRFs per Site	5000
APNs per Network	2500
PCRF Pools per Network	7
PCRF Sub-Pools per Network	14
Topo Hiding Policy Clients per Site	1000
Binding Server Groups per Network	0-8
Session Server Groups per Mated Pair	0-8
Binding Servers per Sg	8

Constraint Name	Value
Session Servers per Sg	8
Maximum Number of SBR Databases	65
Maximum Number of SBR Database Resizing Plans	33
Maximum Number of SBR Data Migration Plans	64
Maximum Number of Binding Region Place Associations	1
Maximum Number of Binding Resource Domains	2
Maximum Number of Mated Sites Place Associations	64
Maximum Number of PCA Sites	32
Maximum Number of Suspect Binding Rules Per Site	50
Cardinality relationships between managed objects	
Sub-Pool Rules per PCRF Pool	10
Maximum Number of Server Groups Per SBR Database	8
Maximum Number of Servers Per SBR Database	4

PCA Assumptions and Limitations

PCA has the following assumptions and limitations:

Assumptions

- For the P-DRA function, the anchor key that identifies all subscribers in the PCA network is the IMSI, while MSISDN is the key that identifies the subscribers for the OC-DRA Function
- All Gx and Gxx session initiating Diameter messages will always include the IMSI. The only exception is emergency calls from devices with no SIM card (UICC-less).
- Messages sharing a common Diameter Session-Id will never arrive out of sequence.
- PCRF/OCS names and client names start with characters that can be used to identify which PCA DSR hosts the primary connection to that equipment. This greatly simplifies routing configuration for the PCA network. The network can be configured to work without such a naming convention, but routing setup and maintenance will be unnecessarily complex.
- The PCA Gx-Prime interface support feature is backward compatible and functions whether or not PCRF pooling is available.
- A complete suspect binding and session removal from the network (not only from the binding and session SBRs) requires the involvement of the policy clients with corresponding actions in response

to PCA's RAR requests. PCA has no control on a policy client's behavior. When a PCEF receives a Gx RAR with Session-Release-Cause AVP, it is expected to send:

- A RAA (2001 Result Code AVP value) in response to the RAR
- A GX CCR-T to terminate the GX session

Limitations

- When a PCRF is selected for a new subscriber binding, a simple round-robin selection mechanism is employed. PCA PCRF selections can be overridden by DSR routing configuration. When PCRF selections are overridden by DSR, weighted load distribution can also be used.
- PCA does not support the 3GPP mechanism to redirect Policy Clients to a PCRF.
- Quota pooling is not supported. Quota pooling is a feature that would allow a number of subscribers to share a common pool of resources for policy decisions. For example, a family plan where all members of the family share access to resources such as bandwidth. PCA has no mechanism for identifying members of a quota pool such that their sessions could all be routed to the same PCRF.
- The P-DRA function supports only two of the Diameter Subscription-Id types: END_USER_IMSI (for IMSI) and END_USER_E164 (for MSISDN). Any other Subscription-Id type is ignored.
- PCA evenly distributes new sessions across the Session Policy SBR Server Groups at the mated pair, regardless of the physical location of the Active server. This results in ~50% of session accesses traversing the WAN between the mated pair sites. For mated triplet deployments having an even distribution of Session SBR server groups, ~66% of session accesses traverse the WAN between mated sites.
- In cases of Regionalized OCS deployments, where the Requests are routed to an OC-DRA on a remote DSR, RBAR may have to be invoked for subsequent Requests (CCR-U/Ts) as well. If MSISDN is not present in CCR-U/T messages, regionalized routing cannot be supported.
- Enabling/disabling the OC-DRA functionality or P-DRA functionality on a per site basis or on a per NE basis while enabling both at the NOAM is not supported.
- For customers upgrading from P-DRA to PCA, the customer team must ensure that there is enough spare capacity available in the session SBRs to support the additional online charging sessions.
- The P-DRA function will reject any binding capable session initiation request after the binding migration period if: (a) the message has no APN (i.e. Called-Station-Id AVP), or has an APN that is not configured in PCA, and (b) PCRF Pooling is enabled. The specifications point out a case in which a Gxx session initiation request could be sent with no APN, which will not work for PCRF Pooling

Note: This condition is not really considered a limitation, but it is important to understand how Policy DRA handles alternate keys. If more than one binding capable session initiation request is received having the same alternate key value, the alternate key is bound to the PCRF that the last received request having that key was bound to. For example, if CCR-I #1 arrives with IMSI X and IPv4 address a.b.c.d and is bound to PCRF A, then CCR-I #2 arrives with IMSI Y and the same IPv4 address and is bound to PCRF B, this will cause IPv4 address a.b.c.d to be bound to PCRF B.

- RBAR currently cannot extract the MSISDN from the User-Name AVP, but can extract it from other AVPs. If there is need to support regionalized routing for CCRs with MSISDN stored in the User-Name AVP, the Diameter Mediation feature will have to be used to extract the MSISDN from the User-Name AVP and include it in an AVP that is supported by RBAR.
- OC-DRA extracts the subscriber's identity from the session initiation request (CCR-I) for the purpose of including it with the session state information stored at the Session SBR when session state is required to be maintained. OC-DRA does not extract the MSISDN or perform number conditioning when the subscriber's identity is retrieved in a format other than E.164 (i.e., MSISDN) such as SIP

URI, TEL URI or NAI. The subscriber's identity is stored in the format in which it is retrieved from the Request.

- The following known error conditions could result in a split binding condition:
 1. A binding sessionRef is removed as a result of the Suspect Binding mechanism, but the actual Diameter session survived the PCRF inaccessibility. This condition is expected to be rare because for a Diameter session to survive the PCRF inaccessibility, there would have to be no signaling attempted for the session during the outage and the PCRF would have to maintain session state over the outage.
 2. A binding sessionRef was removed due to being discovered in an Early state for longer than the Maximum Early Binding Lifetime, but the actual Diameter session was successfully established. This condition is expected to be rare because the binding record is explicitly updated to Final when the master session succeeds or slave polling succeeds. This condition should only result from software errors or SBR congestion causing database update requests to be discarded.
 3. An attempt was made to create a binding-capable session record, but the attempt failed, which triggered a Session Integrity session teardown. However, this mechanism cannot succeed if no session record exists and topology hiding was in use for the policy client that tried to create the session (for example, because the resulting CCR-T cannot be routed to a topology hidden PCRF). This condition is unlikely to cause a split binding because PCA will request that the policy client tear down the session. If the policy client complies, the PCRF will have a hung session that must be audited out. If the policy client declines to tear down the session, a split binding could occur.
- Site redundancy that occurs while an SBR Re-Mating migration is in progress is not supported.
- A binding is associated with a subscriber key (IMSI) that is used to route the request Diameter message to a specific PCRF. If a request message contains multiple subscriber keys (MSISDN, IPv4, and IPv6), then all associated keys are part of the suspect binding removal procedure.
- In some specific situations, a Binding Capable Session can be established without an IMSI value. In such scenarios, additional Sessions may be established and routed using the alternate keys from the original Binding Capable session without an IMSI. If a Diameter message fails and matches the Suspect Binding Removal Rules, the original Session is removed immediately, regardless of the **Remove Immediately** attribute value of the matching rule.

Chapter 4

Policy DRA Overview

Topics:

- [The Policy DRA Function.....99](#)
- [PCRF Pools and Sub-Pools Concepts and Terminology.....99](#)
- [Policy DRA Functions.....110](#)
- [Subscriber Identification and Binding.....113](#)
- [Binding-capable Sessions.....115](#)
- [Binding-dependent Sessions.....120](#)
- [In-session Message Processing.....122](#)
- [Topology Hiding.....122](#)
- [Session Integrity.....124](#)
- [Suspect Binding Removal.....127](#)

This section gives an overview of the Policy DRA function, and includes important fundamental concepts, as well as high-level functionality. Information about PCRF Pools and Sub-Pools is included here as well.

Details about the user interface, feature components, and specific tasks is included in the configuration sections. See [Policy and Charging Configuration](#).

The Policy DRA Function

Policy DRA offers a scalable, geo-diverse Diameter function that creates a binding between a subscriber and a Policy and Charging Rules Function (PCRF) and routes all policy messages for a given subscriber and APN to the PCRF that currently hosts that subscriber's policy rules. Additionally, Policy DRA can perform Topology Hiding to hide the PCRF from specified Clients.

Policy DRA provides the following capabilities:

- Support for all DSR application IDIH requirements; Policy DRA captures metadata that can be used with IDIH to create traces (this assumes that the desired traces are configured in IDIH)
- Distribution of Gx, Gxx, and S9 Policy binding-capable sessions and distribution of Gx-Prime and Rx Policy binding-dependent sessions across available PCRFs

Note: Gx-Prime uses the same Application-Id and Vendor-Id as Gx.

- Binding of subscriber keys such as IMSI, MSISDN, and IP addresses to a selected PCRF when the initial Gx, Gxx, or S9 sessions are already established to that PCRF
- Network-wide correlation of subscriber sessions such that all Policy sessions for a given subscriber are routed to the same PCRF
- Creation of multiple pools of PCRFs, which are selected using the combination of IMSI and Access Point Name (APN). This capability allows you to route policy Diameter signaling initiating from a given APN to a designated subset of the PCRFs that can provide specialized policy treatment using knowledge of the APN.

Note: APNs must be configured before enabling the PCRF Pooling feature.

- Use of multiple binding keys that identify a subscriber, so that sessions with these binding keys can still be routed to the PCRF assigned to the subscriber
- Efficient routing of Diameter messages such that any Client in the network can signal to any PCRF in the network, and vice-versa, without requiring full-mesh Diameter connectivity
- Hiding of PCRF topology information from specified Clients
- The ability to divert a controlled amount of policy signaling to a small subset of the PCRFs in a PCRF Pool for purposes of testing new PCRF capabilities.

Use the Policy DRA GUI to perform configuration and maintenance tasks, edit System Options, and view elements for the Policy DRA Configuration and Maintenance components.

The Session Binding Repository (SBR) hosts the Session and Binding databases, which provide a distributed scalable and High Available (HA) database function to the Policy DRA function for storing and managing the Session data and the subscriber-PCRF Binding data.

PCRF Pools and Sub-Pools Concepts and Terminology

This section describes some basic Policy DRA PCRF Pools and Sub-Pools concepts, and includes useful acronyms and terminology.

Related Topics

- [Policy and Charging Configuration](#)

PCRF Pools

A PCRF Pool (one or more) is a set of PCRFs able to provide policy control for a specific set of services. Creating multiple pools requires that Policy DRA has the ability to select the pool to which a new-binding CCR-I belongs.

Note: Enabling the PCRF Pool function is a one-time operation used to begin a transition period from pre-PCRF Pool processing to PCRF Pool processing. After the function is enabled, it cannot be disabled.

Although the concept of a PCRF pool might appear to be a network-wide concept, PCRF pools configuration is done on a Policy DRA site-by-site basis. Policy DRAs in different sites must be able to have different PCRF Pool Selection configurations.

When deploying multiple PCRF pools, each pool supports either different policy-based services or different versions of the same policy based services. Each PCRF pool has a set of DSR Policy DRA peers that are a part of the pool.

As shown in [Figure 36: Relationship between APNs and PCRF Pools](#), there is a many to one relationship between APNs and PCRF pools. New sessions for the same IMSI can come from multiple APNs and map to the same PCRF Pool.

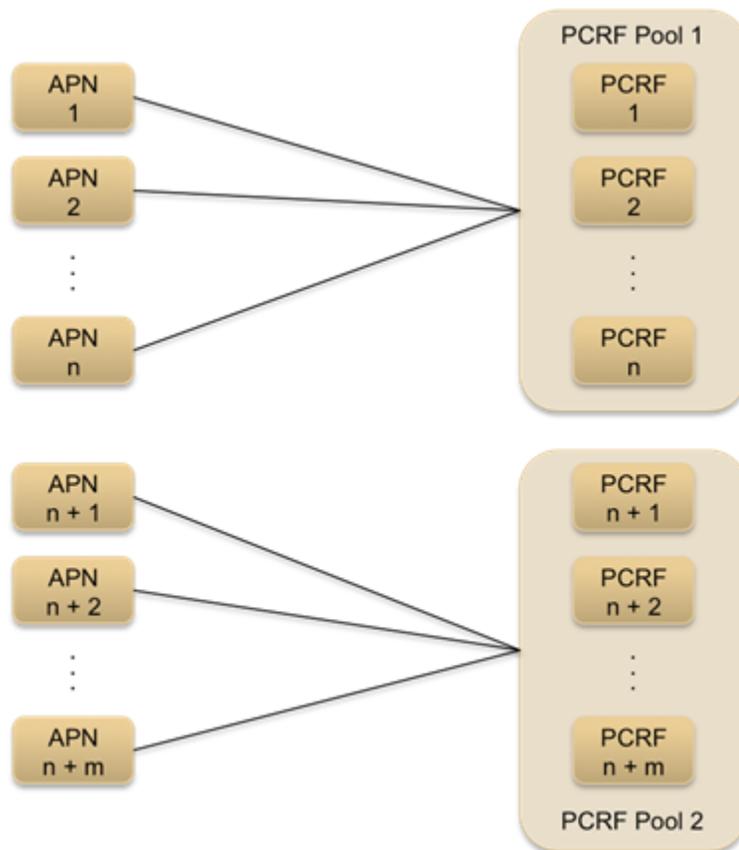


Figure 36: Relationship between APNs and PCRF Pools

[Figure 37: Relationship between IMSIs and PCRF Pools](#) illustrates the relationship between IMSIs and PCRF pool. The same IMSI must be able to have active bindings to multiple PCRF pools.

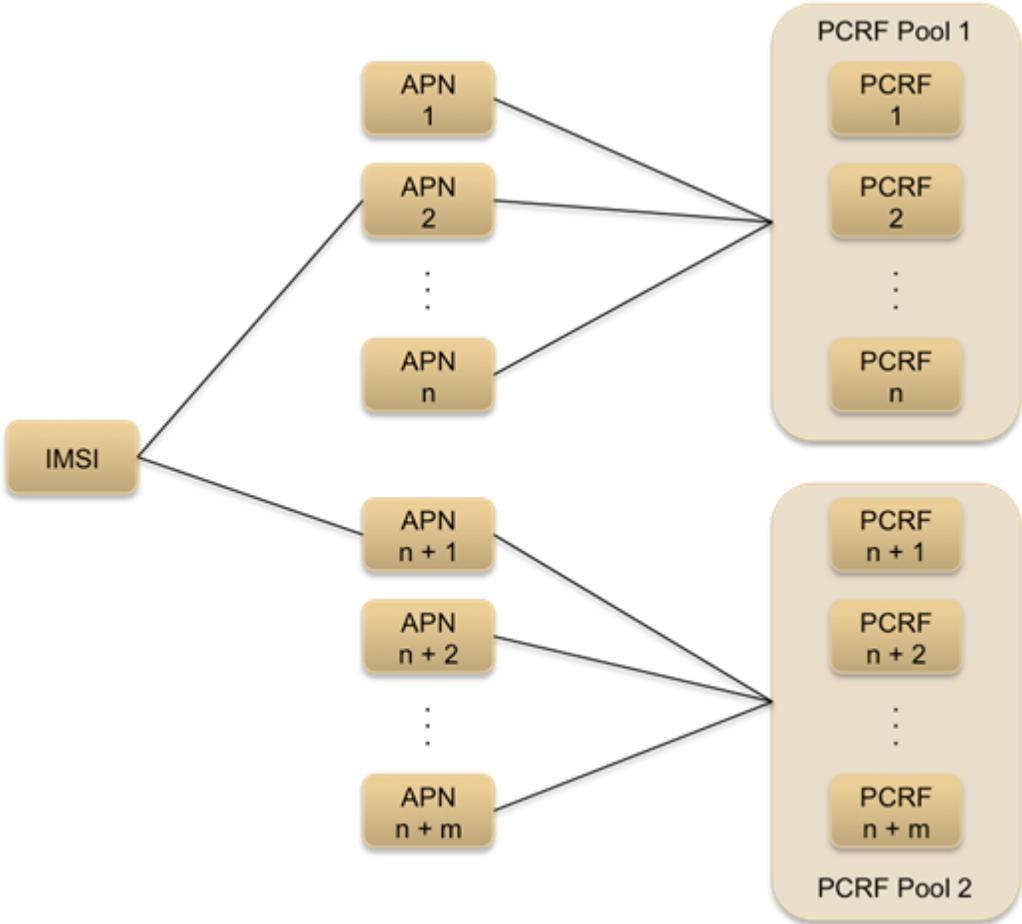


Figure 37: Relationship between IMSIs and PCRFS Pools

Figure 38: Multiple PCRFS Pools illustrates multiple PCRFS pools, each supporting a different service. In this example, PCRFS pool 1 might be dedicated to policy control over the usage of enterprise data services and PCRFS pool 2 might be dedicated to policy control over the usage of consumer data services. It is possible to deploy their policy control capabilities in this way to better enable capacity management of the two PCRFS pools.

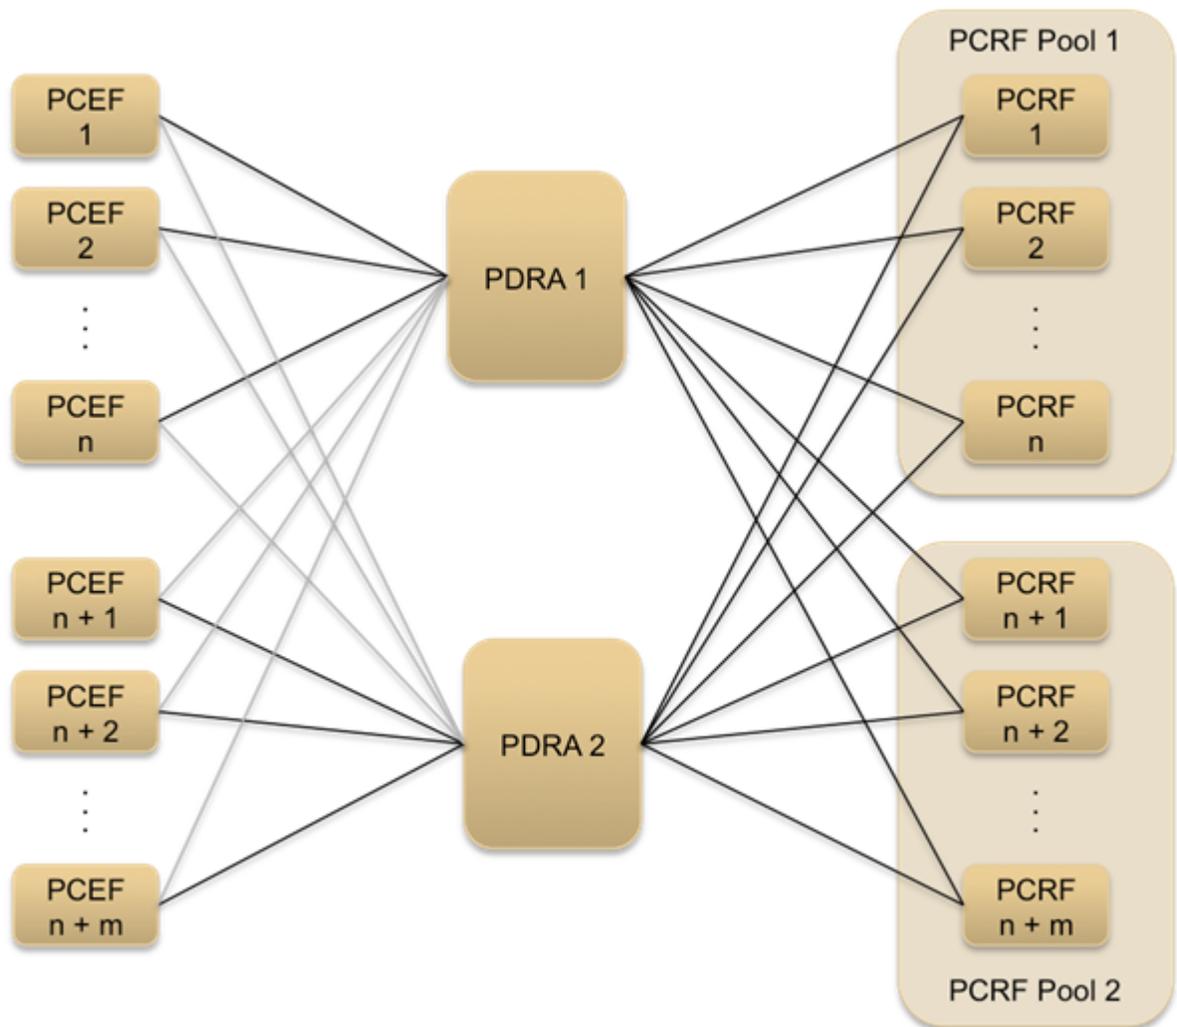


Figure 38: Multiple PCRF Pools

PCRF Sub-Pools

A PCRF Sub-Pool is a subset of a PCRF pool. This is required for scenarios that contain multiple versions of PCRF software within a PCRF pool. The PCRF sub-pool is selected based on the Origin-Host of the PCEF sending a CCR-I.

PCRF Sub-Pools configuration is an optional procedure. PCRF Sub-Pools are used to divert a controlled amount of traffic from a PCRF Pool to a subset of the PCRFs in that pool. This allows new PCRF capabilities or policies to be tested on a portion of the policy signaling prior to using them for the entire network.

Specification of what policy signaling should be routed to the PCRF Sub-Pool is accomplished by configuring PCRF Sub-Pool Selection Rules. Each rule specifies the PCRF Pool that is being subdivided and the Origin-Host of the PCEF, or PCEFs, whose traffic should be routed to the Sub-Pool. If no match is found in the PCRF Sub-Pool Selection Rules, then the original PCRF Pool, selected using the APN, is used for routing. Like PCRF Pool routing, Sub-Pool routing applies only to new bindings.

Figure 39: Multiple PCRF Versions in a PCRF Pool illustrates the concept of PCRF sub-pools. In this figure, there are multiple versions of PCRF Pool 1. This might be necessary when deploying a new version of a PCRF policy-based service and you need to target a subset of the overall sessions for that service to a PCRF running the new version of the PCRF Pools. All other sessions would be routed to the PCRF pool supporting the older version of the policy-based service.

A PCRF Sub-Pool is differentiated by the PCEF from which CCR-I messages originate. As such, PCRF sub-pools support requires adding origin-host to the selection criteria for identifying the PCRF pool.

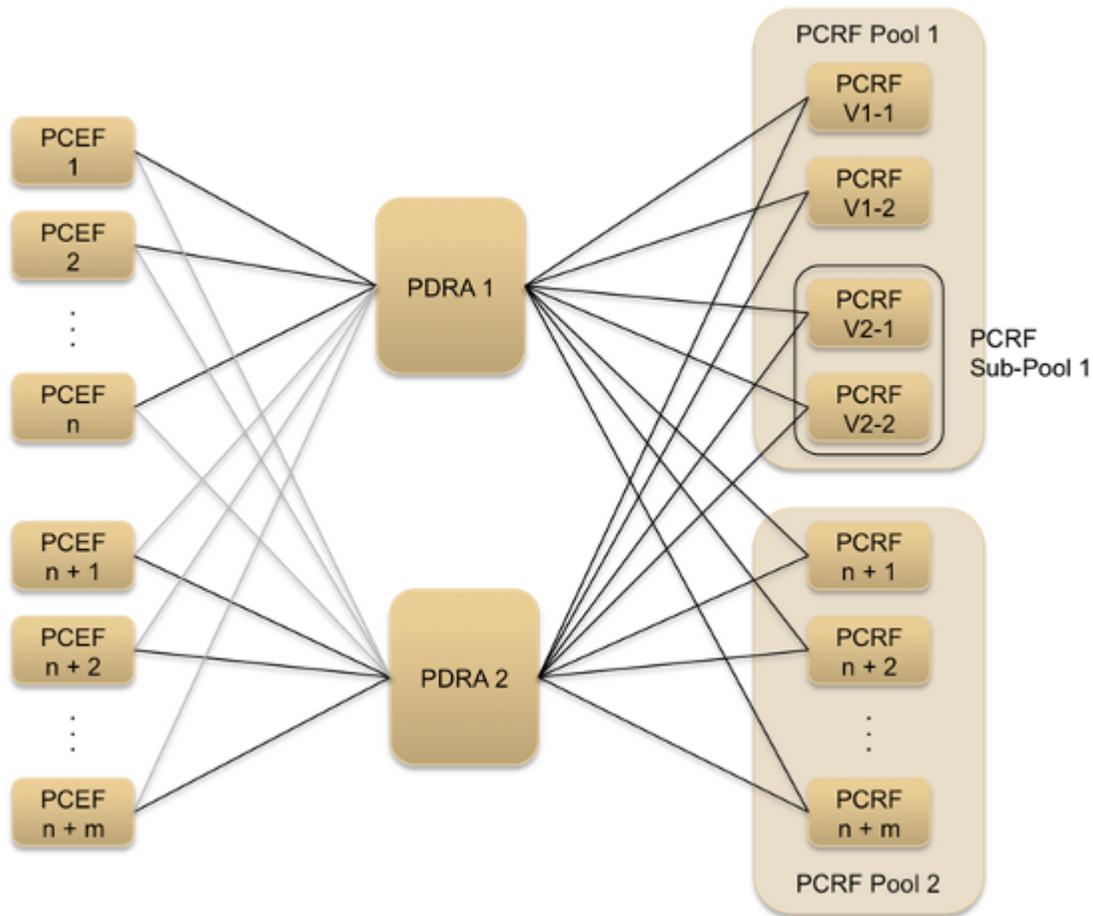


Figure 39: Multiple PCRF Versions in a PCRF Pool

To incrementally add service to a new version of the PCRF, PCRF pool configuration would progress as follows:

- PCRF Pool 1 is defined with the set of APNs that are to be routed to that PCRF pool.
- When a new version of the PCRF in Pool 1 is installed, the configuration is modified to have all new bindings from a specific subset of PCEFs route to the new PCRF in sub-pool 1. CCR-Is received from the remainder of the PCEFs are configured to continue to route to PCRF Pool 1.
- Over time, the configuration can be modified to so that bindings from other PCEFs will be routed to Sub-Pool 1. Alternatively, the sub-pool rule can be removed, resulting in all PCRF instances being part of the PCRF Pool.

- After the new version of the PCRF is proven confirmed, the configuration is modified so that all CCR-Is are routed to PCRF Pool 1.

Figure 40: PCRF Pools and Sub-Pools Routing Scenarios shows example routing scenarios using PCRF Pools and Sub-Pools.

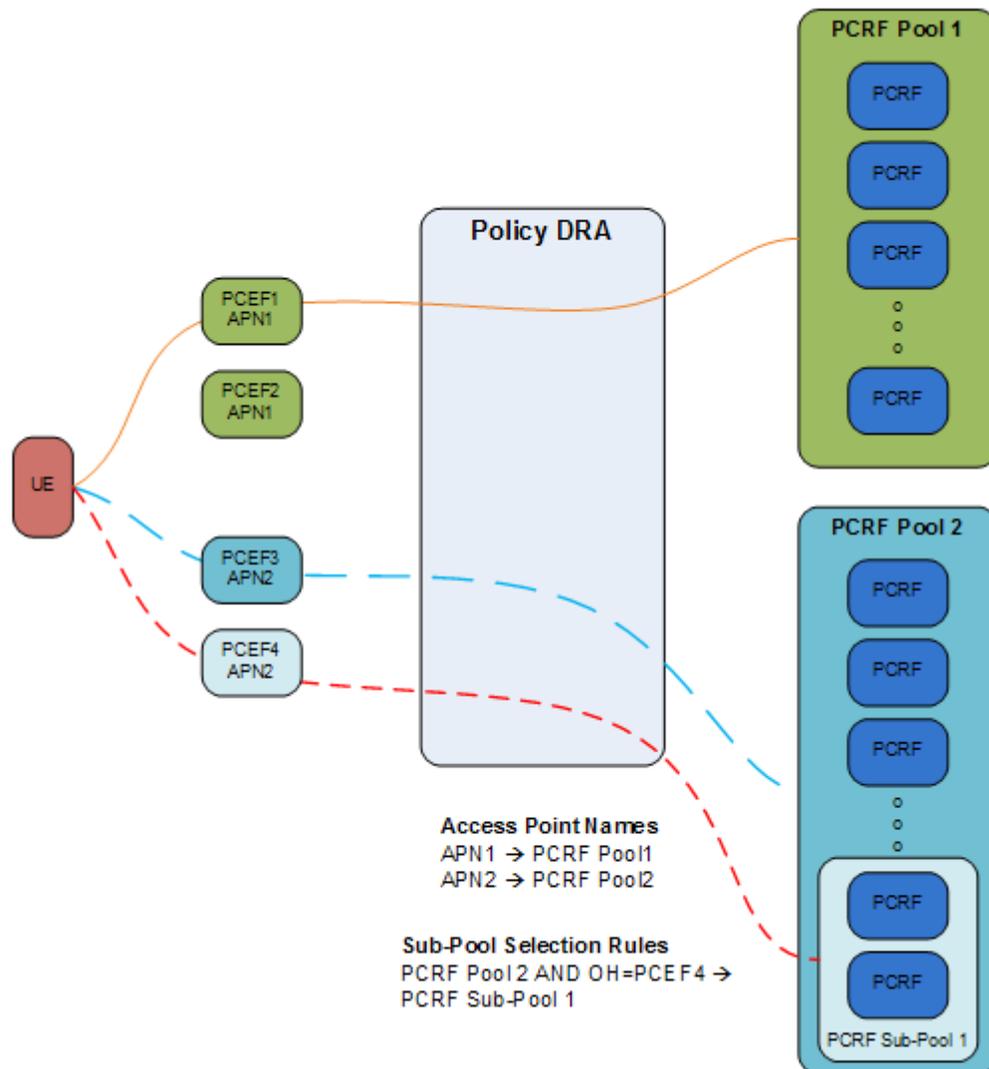


Figure 40: PCRF Pools and Sub-Pools Routing Scenarios

Planning for PCRF Sub-Pooling

To plan for PCRF Sub-Pooling, consider the following:

- Identify the PCRF (or PCRFs) on which the new functionality is to be proven. The PCRF could be an existing PCRF already in a pool, or a new PCRF not yet assigned to a PCRF Pool.
- Determine which PCRF Pool the PCRF belongs to.
 - This can be accomplished by examining routing data at the mated pair of DSRs that have connections to the PCRFs.
 - It is possible, though unlikely, that a PCRF could exist in more than one PCRF Pool.

- Determine which APNs map to the PCRF Pool.
 - This can be accomplished by examining the Access Point Names at the NOAM.
 - Filtering can be used to display only APNs that are mapped to the PCRF Pool of interest.
- Determine which PCEFs use the APNs.
- Determine which PCEFs host names you want to route signaling to the PCRF Sub-Pool that will contain the PCRFs from Step 1. Use caution not to overwhelm the PCRFs planned for the Sub-Pool by routing more signaling than they can reasonably support.

Session Binding

Without the PCRF Pool feature, bindings are accessed using the IMSI contained in the new-binding CCR-I request. In other words, before PCRF Pooling is enabled, IMSI was sufficient to find a binding. After PCRF Pooling, IMSI and APN both are required to find a binding.

Policy Sessions

There are two broad categories of Policy sessions: binding-dependent and binding-capable.

A binding-dependent session is a Policy session that cannot cause a binding to be created, and cannot be created unless a binding exists. Binding-dependent interfaces contain a specific PCRF peer to which sessions can be bound. A PCRF pool consists of multiple PCRF instances.

A binding-capable session is a Policy session that is allowed to cause a new binding to be created for a subscriber. binding-capable session initiation requests includes both IMSI and an APN. Policy DRA locates the APN, which is mapped to a PCRF Pool via **Policy and Charging > Configuration > Access Point Names**. The binding of a subscriber to a PCRF must remain intact as long as the subscriber has at least one active binding-capable Diameter session.

Binding-capable sessions are created by Gx, Gxx, or the S9 versions of Gx and Gxx interfaces. If a CCR-I message arrives for a Binding Capable Interface, Policy DRA checks for an existing binding for the IMSI and APN in the message.

Binding data is accessible from anywhere in the network. Session data is scoped to a mated pair, and is only accessible from that mated pair.

Policy DRA Terminology

Table 16: Policy DRA Terminology shows a list of some Policy DRA terms and their meanings as they apply to this document.

Table 16: Policy DRA Terminology

Term	Meaning
Ambiguous Rules	Two rules are ambiguous if they have equal priority, different conditions, different PCRF Pools, and a best-match cannot be determined for a single binding-capable request.
Binding	A mapping in the Policy DRA from an IMSI and APN to a PCRF for the purpose of routing policy Diameter signaling. Once a binding exists for an IMSI and APN, all policy Diameter sessions with that IMSI and APN are routed to the bound PCRF.

Term	Meaning
	A binding ceases to exist when the last Diameter session for that IMSI and APN is terminated. See also PCRF Pool Binding.
Binding-dependent Session	A specific PCRF peer to which sessions can be bound. A PCRF pool consists of multiple PCRF instances.
Condition Operator	A logical operator used to compare the Condition Parameter with the Condition Value. Only the Origin-Host parameter is supported in this release. Operators supported for Origin-Host are: Equals, Starts With, and Ends With.
Condition Parameter	The binding-capable session initiation request AVP to be used for PCRF Sub-Pool selection. The only supported Condition Parameters is Origin-Host.
Condition Value	The value of the Condition Parameter to be matched using the Condition Operator. For example, in the Condition "Origin-Host Starts With abc", "abc" is the Condition Value.
Conflicting Rules	Two rules conflict if everything in the rules is the same except for the PCRF Pool.
Duplicate Rules	Rules are duplicates if everything (Origin-Host operators and values, Priority, PCRF Pool, and PCRF Sub-Pool) in the two rules is the same.
Early Binding	An Early Binding is a binding for which a session initiation request has been received, but no session initiation answer has been received. The PCRF for an Early Binding is unknown. A given IMSI-APN combination can have only one early binding. The Early Binding serializes binding creation attempts for a given IMSI and APN. Subsequent session initiation requests for an IMSI-APN combination for which an Early Binding exists are held until the Early Binding becomes a Final Binding.
Early Binding Master	A binding-capable session initiation request that creates a new Early Binding is referred to as the Early Binding Master for that binding. A given Early Binding can have only one master. The term master is used to convey that no subsequent binding-capable session initiation requests for that binding can be routed until the master session is successfully answered by a PCRF.
Early Binding Slave	A binding-capable session initiation request that matches an Early Binding is referred to as an Early

Term	Meaning
	Binding Slave for that binding. There may be multiple slaves for a given Early Binding. The term slave is used to convey that the slave session request must wait for the master session request to be completed before it can be routed.
Enabling PCRF Pool Feature	<p>Enabling the PCRF Feature is a one-time operation used to begin a transition period from pre-PCRF Pool processing to PCRF Pool processing. This is a one-time operation and, after enabled, the PCRF Pool feature can no longer be disabled.</p> <p>Enabling the PCRF Pool feature only applies when upgrading from the Policy DRA 4.1.5 or 5.0 release. The PCRF Pool feature can only be enabled after all Policy DRA Network Elements are upgraded and those upgrades are committed. Only at this point is it possible to use the PCRF Pool feature logic, as the upgrade will result in changes to the handling of binding data.</p>
Existing-Binding CCR-I	A CCR-I request for a specific IMSI, APN combination that occurs when there is an Existing-Binding CCR-I binding SBR record for the IMSI+APN. In this case, the existing binding for the IMSI+APN is used to route the CCR-I request.
Final Binding	A Final Binding is a binding for which the PCRF is known because the PCRF sent a success answer in response to the session initiation request. When a binding-capable session initiation success answer is received, an Early Binding is explicitly marked as a Final Binding.
IPcan Session	A connection to the Enhanced Packet Core.
Migration Period	For customers upgrading from DSR 4.1 Policy DRA, a migration occurs from the IMSI-only binding table to a table that supports a binding per IMSI-APN combination. In order to avoid Split Bindings, bindings existing in the IMSI only table are honored until they naturally terminate. As existing IMSI-only bindings naturally terminate, they are replaced with IMSI-APN bindings. Once all IMSI-only bindings are gone, the migration period is complete. This data migration also applies to alternate key tables (MSISDN, IPv4 Address and IPv6 Address).
Non-Specific Binding Correlation Key	A binding correlation key value that may be specified in more than one binding-capable

Term	Meaning
	session initiation request is considered to be a non-specific binding correlation key. Non-Specific Binding Correlation Keys are generally associated with the subscriber vs. being associated with a particular session. IMSI and MSISDN are examples of non-specific binding correlation keys because multiple sessions may exist concurrently with the same IMSI or MSISDN value. IPv4 and IPv6 addresses are not "non-specific" because each binding-capable session is expected to have its own unique key value. (Note: There is a chance that Gx and Gxx sessions for the same IMSI could include the same IP addresses, but in this case the Gx and Gxx sessions are expected to have the same APN and should be routed to the same PCRF.)
PCRF Instance	A specific PCRF peer to which sessions can be bound. A PCRF pool consists of multiple PCRF instances.
PCRF Pool	A logical grouping of PCRFs intended to provide policy decisions for subscribers associated with a particular APN. Policy DRA supports 7 PCRF Pools per Policy DRA Network. A PCRF Pool is selected using the configured mapping between the APN and the PCRF Pool. More than one APN may point to the same PCRF Pool.
PCRF Pool Binding	For a given IMSI, if no binding exists for the APN present in the binding-capable session initiation request, the request must be routed to the same PCRF bound to another APN that maps to the same PCRF Pool, if one exists. For example, if APN X and APN Y both map to PCRF Pool "Maple" and there is already a final binding for APN X, a binding-capable session for APN Y must route to the same PCRF that APN X is bound to.
PCRF Sub-Pool	A logical sub-division of a PCRF Pool selected by Origin-Host. PCRF Sub-Pools can be used to selectively route policy traffic to a set of PCRFs for the purpose of proving in new PCRF capabilities. More than one PCRF Sub-Pool Selection Rule may point to the same PCRF Sub-Pool.
PCRF Sub-Pool Selection Rule	A rule that defines a mapping from PCRF Pool and Origin-Host to PCRF Sub-Pool. A set of values that must be matched against AVP values in a binding-capable session initiation request for the purpose of selecting a PCRF Sub-Pool. The

Term	Meaning
	number of PCRF Sub-Pool Selection Rules per PCRF Pool is limited to 10.
Primary PCRF Pool	A PCRF Pool that is mapped to an APN, as opposed to a PCRF Sub-Pool, which is mapped to a PCRF Pool and an Origin-Host.
Redundant Rules	Rules are redundant if the PCRF Sub-Pools are the same and a request matching the more specific rule always matches the less specific rule. Redundancy does not include the default rule. The PCRF Sub-Pool Selection Rules GUI does not prevent creation of redundant rules since the PCRF Sub-Pool is the same, leaving no ambiguity.
Rule Condition	Each PCRF Sub-Pool Selection Rule consists of a condition made up of a parameter (Origin-Host), an operator, and a value, for example Origin-Host Equals pcef015.tklc.com.
Rule Matching	Rule matching is the process of finding the best match among the configured PCRF Sub-Pool Selection Rules for a given binding-capable session initiation request. Rule matching occurs on the DA-MP that processes the binding-capable session initiation request.
Rule Priority	Each PCRF Sub-Pool Selection Rule has a priority value from 1 to 99, with 1 being the highest priority. The Rule Priority allows the user to give preference to one rule over another, regardless of which rule might be the "best match".
Split Binding	A Split Binding is defined as a situation in which a given subscriber has more than one binding for the same APN. Note: Split bindings would be created by addition of more specific PCRF Pool selection criteria. For example: Adding an explicit APN to PCRF Pool mapping when the "-Unrecognized-" APN mapping was previously being used. Adding a more specific PCRF Sub-Pool Selection Rule. Policy DRA prevents Split Bindings by always honoring existing bindings for an IMSI-APN combination. The presence of an existing binding for the IMSI-APN combination overrides the rule-based PCRF Pool selection. Prevention of Split Bindings is necessary to avoid having two PCRFs delivering possibly conflicting rules to one PCEF. Added benefit is avoidance of ambiguity in binding correlation for non-specific binding keys.

Term	Meaning
Suspect Binding	A Suspect Binding is a Binding for which Diameter messaging to a PCRF has detected a "Rule Match". Once a "Rule Match" has been detected, all Bindings for the subscriber(IMSI) to the PCRF are considered "Suspect". It is possible for a subscriber(IMSI) to have both "Suspect" and non-"Suspect" Bindings.

Policy DRA Functions

The Policy DRA functionality performs the following major functions:

- Processing Diameter Request messages
- Querying subscriber binding status
- Selecting an available PCRF and routing the Diameter Requests to a selected PCRF, including the ability to route new-binding CCR-I requests to one of a configured set of PCRF pools
- Topology Hiding
- Processing Diameter Answer messages
- Managing subscriber Session and Binding databases

Diameter Request Message Processing

Diameter Request messages from Policy clients (PCEF, BBERF, AF, and DPI/MOS) arrive at Policy DRA routed by the DSR Diameter Routing Function based on a prioritized list of Application Routing Rules. The Application Routing Rules are configured for the Policy DRA functionality based on the information in the Diameter Request message: Application ID, Command-Code, Destination-Realm and Host, and Origin-Realm and Host.

After receiving a Diameter Request, the Policy DRA retrieves and examines the relevant AVPs contained in the message. The Policy DRA-relevant AVPs vary depending on the Diameter interface on which a Diameter message is carried.

By retrieving and examining the contents of the relevant AVPs, the Policy DRA determines:

- The type of the Diameter Request: initiation, update, or termination
- The type of interface over which the Request message is carried and whether the session over this interface is binding-capable or binding-dependent.

A session over a binding-capable interface will be eligible to establish a binding to a PCRF, while a session over a binding-dependent interface will rely on an existing binding to a PCRF but cannot create a new binding by itself.

- The subscriber's IDs from the appropriate AVPs (Subscription-ID AVP, Framed-IP-Address AVP, and Framed-IPv6-Prefix AVP)
- The Origin-Host and Realm AVPs, and Destination-Host and Realm AVPs.
- The access point name (APN) from which the request was received.
- Session-Id AVPs

The Policy DRA will use the information to query the SBR database for binding and session status of the subscriber whose IDs are included in the Diameter Request message.

Query Subscriber's Binding Status

Binding-capable Session Initiation Requests

After processing an incoming Diameter Request message, the Policy DRA queries the SBR database for binding status based on the subscriber's IDs (keys) contained in the Request message. The query is done over the Policy DRA and SBR interface. A response to the request from the Active SBR to the Policy DRA provides a result on whether or not the queried binding or session record exists in the database.

When a session initiation Request message is received (Gx, Gxx or S9), the Policy DRA determines whether or not a binding exists for the Subscriber ID, an Anchor Key, included in the Request message. The Policy DRA queries the appropriate SBR for the binding status for this session. Depending on the output from the interactions with the SBRs, the Policy DRA might need to select an available PCRF to which the the Diameter Request message will be routed.

Special Cases

Occasionally, unique situations arise that require specialized attention. This section addresses, some of the more common ones.

Binding-capable Session Initiation Answers - Handling a Binding-Capable Session Initiation Request with No IMSI

The Policy DRA handles these calls by processing CCR-I messages that do not contain an IMSI and any Alternate Keys. When a CCR-I arrives with no IMSI, the Policy DRA selects a configured PCRF (see [Query Subscriber's Binding Status](#)) and routes the Request message to that PCRF. If a CCA-I is received from the selected PCRF, Policy DRA will invoke the SBR database to create a session and binding records based on any Alternate Keys included in the message.

Note: If the request contained more than one of a given type of key (for example, MSISDN, IPv4, or IPv6), only the first one of each type encountered in the request parsing is used. All other keys of that type are ignored.

If the session creation or any alternate key creation fails, the Session Integrity feature terminates the session.

Binding-capable Session Initiation Answers - Handling a Binding-Capable Session Initiation Request with an IMSI

When a binding-capable session initiation request is received, Policy DRA must check to see if the request matches an existing binding. If a matching binding exists, the request is relayed to the bound PCRF. If no existing binding is matched, a new binding is created.

Prior to checking for a matching binding; however, Policy DRA determines to which PCRF Pool or Sub-Pool the request belongs. This is determined as follows:

- The APN in the binding-capable session initiation request (for example, CCR-I) is mapped to a PCRF Pool. This mapping is configured in **Policy DRA -> Configuration -> Access Point Names**.

- Next, a check is performed to determine if an optional PCRF Sub-Pool applies to this request. If no Sub-Pool applies, the PCRF Pool mapped to the APN is used as the PCRF Pool for the request.

To determine if a Sub-Pool is configured for this request, the PCRF Pool mapped to the APN and the Origin-Host from the binding-capable session initiation request are compared against PCRF Sub-Pool Selection Rules. If a match is found, the specified PCRF Sub-Pool is used as the PCRF Pool for the request.

Now that a PCRF Pool has been selected for the request, the rules for determining if the new request matches an existing binding can be performed as follows:

- If a binding exists for the IMSI and APN, use that binding, else
- If a binding exists for the IMSI and suggested PCRF Pool or Sub-Pool, use that binding.

If no existing binding is found for the IMSI and APN or IMSI and PCRF Pool, a new binding is created, specifying the IMSI, APN, and PCRF Pool. This binding is referred to as an early binding because the actual PCRF will not be known until the binding-capable session initiation answer is received.

The binding-capable session initiation request message is then routed using the Peer Route Table (PRT) assigned to the PCRF Pool or Sub-Pool chosen above. The Diameter routing capabilities are used to load distribute the request across PCRFs in the specified pool.

Note: After PCRF Pooling capability is enabled, PCRF selection from within the pool is controlled entirely by the Diameter stack configuration. The Policy DRA functionality no longer performs a round-robin selection among all configured PCRFs. The Policy DRA functionality selects a PCRF Pool, which is mapped to a PRT. From that point onwards, routing logic proceeds as specified in the PRT rules, route lists, and route groups.

This binding becomes a final binding when a 2xxx response is received from the PCRF that answered the binding-capable session initiation request.

Early Binding

An Early Binding is a binding for which a session initiation request has been received, but no session Early Binding initiation answer has been received. The PCRF for an Early Binding is unknown. A given IMSI-APN combination can have only one early binding. The Early Binding serializes binding creation attempts for a given IMSI and APN. Subsequent session initiation requests for an IMSI-APN combination for which an Early Binding exists are held until the Early Binding becomes a Final Binding.

A binding-capable session initiation request that creates a new Early Binding is referred to as the Early Binding Master for that binding. A given Early Binding can have only one master. The term master means that no subsequent binding-capable session initiation requests for that binding can be routed until the master session is successfully answered by a PCRF.

A binding-capable session initiation request that matches an Early Binding is referred to as an Early Binding Slave for that binding. There may be multiple slaves for a given Early Binding. The term slave is used to convey that the slave session request must wait for the master session request to be completed before it can be routed.

PCRF Selection and Routing

PCRF selection involves distribution of subscriber bindings to PCRFs that are configured in advance. When a Diameter Request message arrives on a Gx, Gxx, or S9 interface aiming at generating a new session, the Policy DRA must determine if a binding already exists for the IMSI APN included in the Diameter message.

If a binding-capable session initiation request is received that would result in a new binding, and no PCRFs are configured at the site, Policy DRA generates an error response.

Note: This does not apply if a binding already exists for the IMSI and APN, or IMSI and PCRF Pool.

See [Query Subscriber's Binding Status](#) for a description of PCRF selection when PCRF Pooling is enabled.

Topology Hiding Process

See [Network-Wide Options](#).

Diameter Answer Message Processing

After the Policy DRA routes a Diameter Request message to a selected PCRF, and updates the SBR on binding status, the Policy DRA could find itself in one of the following situations:

1. An Answer is received from a PCRF and a response is received from a Policy SBR
2. An Answer is received from a PCRF, but no response is received from a Policy SBR after a configured time interval
3. A response is received from a Policy SBR, but no Answer is received after a configured time interval

For situations 1 and 2, the Policy DRA always forwards the Answer messages to the corresponding Requests initiators through the Diameter Routing Function, with or without Topology Hiding processing depending on the Topology Hiding status of the Policy Client.

For situation 3, the Policy DRA generates Diameter Answer messages with proper Error Codes and routes the Answers to the Request initiators through the Diameter Routing Function, with or without Topology Hiding processing depending on the Topology Hiding status of the Policy Client.

Subscriber Session and Binding Database Management

The Policy DRA will invoke the SBRs to perform relevant database operations after or in parallel with sending the Answer messages out. Which database operations to be performed depends on the Diameter interface type in the incoming Diameter Request, the Diameter Request message type (session initiation, session update, or session termination), and the results from the responses. The following operations can be performed:

- Finding, creating, or updating binding records
- Removing Suspect Binding records
- Creating or removing alternate key binding records
- Finding, creating, refreshing, or removing session records

Subscriber Identification and Binding

Policy sessions can be established using multiple Diameter interfaces such as Gx, Gxx, Gx-Prime, Rx and S9. A session can be characterized as binding-capable or binding-dependent, depending on whether or not a binding can be created over it.

- Gx, Gxx and S9 interfaces are binding-capable

- Rx, Rx over S9, and Gx-Prime interfaces are binding-dependent

A session over a binding-capable interface will be eligible to establish a binding to a PCRF, while a session over a binding-dependent interface will rely on an existing binding to a PCRF but cannot create a new binding by itself.

In order for the Policy DRA to route all messages from a subscriber (perhaps through multiple interfaces and devices) to the same PCRF, the Policy DRA should be able to identify the subscriber by the information in the incoming Diameter Request messages. One subscriber can be associated with multiple Subscriber Ids depending on the access networks and device types used. The Subscriber Ids are also called Subscriber Keys or keys. Messages that can cause creation of a subscriber-PCRF binding are required to contain the subscriber’s device IMSI, which can be used to uniquely identify the subscriber. IMSI is referred to as the subscriber Anchor Key in the SBR Binding database.

Session initiating messages may also contain additional information to identify the subscriber. This information, which may include an MSISDN, an IPv4 address, or an IPv6 address prefix, is referred to as subscriber Alternate Keys. Database records with Alternate Keys are always established by binding-capable sessions, and can be used to identify the subscriber in binding-dependent sessions. For example, a Gx CCR-I message must contain the IMSI Anchor Key under normal circumstance, and may also contain an MSISDN, an IPv4 address, and an IPv6 address. After a binding is established between the subscriber and a PCRF, binding-dependent sessions containing one or more of the subscriber keys can be routed to the PCRF using an Alternate Key.

In *Figure 41: Subscriber Key Usage*, a Gx CCR-I message created 3 subscriber keys: one Anchor Key and two Alternate Keys, all bound to a PCRF called PCRF5. When a binding-dependent Rx session (AAR message) is created containing only IP addresses with no Anchor Key, the Policy DRA functionality looks up the IPv4 address of the subscriber and is able to relate it to the same PCRF because the Gx session had defined those IP addresses.

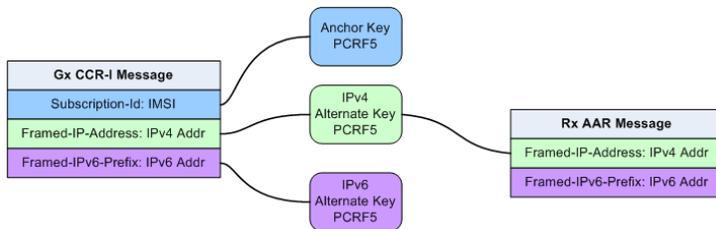


Figure 41: Subscriber Key Usage

Alternate Keys can be configured with a priority (values 1 through 5, where 1 is the highest Priority (IMSI, IPv4, IPv6, or MSISDN). This improves the chances of finding the data in the Diameter message and the chances of finding the Alternate Key in the Binding database. *Table 17: Example Key Priority Configuration* illustrates an example Binding Key configuration with priorities assigned to each key. IMSI, IPv4, IPv6, or MSISDN

Table 17: Example Key Priority Configuration

Priority	Key Type
1	IMSI
2	IPv4
3	MSISDN
4	IPv6

Priority	Key Type
5	<Not configured>

The example configuration in [Table 17: Example Key Priority Configuration](#) will affect how the keys are searched in the Diameter message for binding-dependent session initiating messages:

1. After the IMSI, the Framed-IP-Address AVP will be looked for first in the incoming Diameter Request message.
2. If the AVP is found, the Policy SBR database is searched for a binding with IPv4 address.
3. If the Framed-IP-Address AVP is not found, a Subscription-Id AVP containing an MSISDN will be looked for.
4. If the Subscription-Id AVP with an MSISDN is found, look for a binding with that MSISDN.
5. If a Subscription-Id AVP containing an MSISDN is not found, then no Alternate Keys are present in the message and no Alternate Key records will be created by the application.

Only the configured subscriber keys will be searched for. For example, an incoming Diameter message contains a MSISDN in the Subscription-ID AVP, but MSISDN is not configured in the priority configuration, the Policy DRA functionality will NOT look for MSISDN or use it in the Binding database.

Binding-capable Sessions

A binding is a relationship stored in the Binding SBR between various subscriber data session identities, such as MSIDN/IP Address(es)/IMSI and the assigned PCRF. A session is a relationship stored in the Session SBR that associate additional sessions with a binding.

Policy DRA allows distribution of Gx, Gxx, and S9 Policy binding-capable sessions and distribution of Gx-Prime and Rx Policy binding-dependent sessions across available PCRFs.

Binding-capable Session Initiation Request Processing Rules and Requirements

The following rules apply to the selection of a suggested PCRF Pool or Sub-Pool upon receipt of a binding-capable session initiation request. The request might be routed to an existing binding; only new bindings are guaranteed to route to the suggested PCRF Pool or Sub-Pool.

- Upon receipt of a binding-capable session initiation request containing no Called-Station-Id AVP (for example, no APN), Policy DRA generates and sends a binding-capable session initiation answer message using the Result Code configured for the Diameter interface for the Missing Or Unconfigured APN condition in the Error Codes GUI. The answer message shall include an Error-Message AVP with the 3-digit error code suffix of 500.
- Upon receipt of a binding-capable session initiation request containing no Called-Station-Id AVP (for example, no APN), Policy DRA asserts alarm-ID 22730 and increments measurement RxBindCapMissingApn by one.
- Upon receipt of a binding-capable session initiation request containing a Called-Station-Id AVP (for example, APN) that is not configured on the Access Point Names GUI page, Policy DRA generates and sends a binding-capable session initiation answer message using the Result Code configured for the Diameter interface for the Missing Or Unconfigured APN condition in the Error Codes GUI. The answer message includes an Error-Message AVP with the 3-digit error code suffix of 501.

- Upon receipt of a binding-capable session initiation request containing a Called-Station-Id AVP (for example, APN) that is not configured on the Access Point Names GUI page, Policy DRA asserts Alarm-ID 22730 and increments measurement RxBindCapUnknownApn by one.
- Upon receipt of a binding-capable session initiation request containing a Called-Station-Id AVP (for example, APN) that is configured on the Access Point Names GUI page, the Policy DRA application performs PCRF Pool selection. Measurement RxBindCapApn2PcrfPool is incremented by one for the APN.
- If no PCRF Sub-Pool Selection rule matches, the suggested PCRF Pool is the PCRF Pool configured for the APN on the Access Point Names GUI page.
- If no PCRF Sub-Pool Selection Rule exists for the PCRF Pool that was assigned to the APN from the binding-capable session initiation request, no match exists in the PCRF Sub-Pool Selection Rules.
- If no PCRF Sub-Pool Selection Rule exists where the PCRF Pool that was assigned to the APN from the binding-capable session initiation request matches and with an operator and value that match the Origin-Host of the binding-capable session initiation request, no match exists in the PCRF Sub-Pool Selection Rules.
- A PCRF Sub-Pool Selection Rule using the Equals operator is considered as a match if all of the following are true:
 - The PCRF Pool assigned to the APN from the binding-capable session initiation request matches.
 - All characters of the Origin-Host from the binding-capable session initiation request match the Value specified in the rule, ignoring case (for example, a.b.c is equivalent to A.B.C).
- A PCRF Sub-Pool Selection Rule using the Starts With operator is considered as a match if all of the following are true:
 - The PCRF Pool assigned to the APN from the binding-capable session initiation request matches.
 - All characters of the Value specified in the rule match the leading characters in the Origin-Host from the binding-capable session initiation request, ignoring case (for example, Fred is equivalent to FRED).
- A PCRF Sub-Pool Selection Rule using the Ends With operator is considered as a match if all of the following are true:
 - The PCRF Pool assigned to the APN from the binding-capable session initiation request matches.
 - All characters of the Value specified in the rule match the trailing characters in the Origin-Host from the binding-capable session initiation request, ignoring case (for example, Fred is equivalent to FRED).
- If more than one PCRF Sub-Pool Selection Rule matches and the matching rules have equal priority, the Policy DRA application prefers rules with the Equals operator to rules with the Starts With and Ends With operators.

Note: The GUI prevents ambiguous Starts With and Ends With rules.
- If more than one PCRF Sub-Pool Selection Rule matches according to requirements, the Policy DRA application selects the match having the highest priority (for example, the lowest numeric priority value).

Note: The GUI prevents creation of ambiguous, conflicting and duplicate rules.
- If a PCRF Sub-Pool Selection Rule matches according to requirements, Policy DRA application uses the PCRF Sub-Pool from the matching rule as the suggested PCRF Pool. Measurement RxBindCap2PcrfSubPool is incremented by one for the PCRF Sub-Pool Selection Rule that was matched.

- If a binding-capable session initiation request is received that would result in a new binding and no PCRFs are configured at the site, Policy DRA generates an error response with the 3002 Diameter Response-Code and Error-Message AVP including the string No PCRFs configured at this site.

Note: This requirement does not apply if a binding already exists for the IMSI and APN, or IMSI and PCRF Pool.

- If a binding-capable session initiation request is received and no PCRFs are configured at the site, Policy DRA generates timed alarm 22730, which indicates that no PCRFs are configured.

Note: The alarm is only generated if the binding-capable session initiation request results in a new binding being created.

The following requirements describe handling of binding-capable session initiation requests after a suggested PCRF Pool or Sub-Pool has been successfully selected.

- Upon receipt of a binding-capable session initiation request for an IMSI that has an existing Final binding, measurement SbrFinalBindingsFollowed (11351) is incremented by one and the Policy DRA application attempts to route the request to the PCRF from the selected binding.
- When checking for an existing binding, the Policy DRA searches in the following order, using the first binding that matches:
 - A binding for the IMSI and APN (from the ImsiApnAnchorKey table)
 - A binding for the IMSI and suggested PCRF Pool or Sub-Pool (from the ImsiApnAnchorKey table)
- Upon receipt of a binding-capable session initiation request for an IMSI for which no existing binding is found, the Policy DRA attempts to route the request using the suggested PCRF Pool or Sub-Pool.
- Upon receipt of a binding-capable session initiation request for an IMSI for which no existing binding is found, a new binding is created using the IMSI, APN, and suggested PCRF Pool or Sub-Pool.
- If, when creating the new binding, the record for the IMSI already contains 10 session references, the Policy DRA generates a Diameter error response using the response code configured for the SBR Error condition.

Note: The Error-Message AVP contains the reason for the failure.

- When a binding-capable session initiation request results in a new binding, the binding-capable session initiation request is routed to the Peer Routing Table mapped to the PCRF Pool or Sub-Pool at the site where the request was received. When the PCRF Pool or Sub-Pool is mapped to a configured PRT table, measurement RxBindCapPcrfPool2Prt is incremented by one for the PCRF Pool or Sub-Pool.
- If the PCRF Pool or Sub-Pool is not mapped to a Peer Routing Table (i.e. is mapped to the "-Select-" entry) at the site processing the request, the request shall be routed according to the routing layer PRT precedence. Measurement RxBindCapPcrfPoolNotMapped is incremented by one.

Note: When the PCA does not specify a PRT table to use, DRL looks for a PRT in the ingress Peer Node configuration; then, if still not specified, in the Diameter Application-Id configuration. This behavior is necessary for backwards compatibility for cases where the pre-PCRF Pooling release had the Site Options PRT table for new bindings set to "-Not Selected-".

- If a new binding is created after PCRF Pooling is Enabled and the GLA feature is activated in the Policy DRA Network, Policy DRA stores the Origin-Host of the Policy Client that originated the binding-capable session initiation request in the binding record for use by GLA.

PCRF Pool Selection

The following configuration data are needed to support the PCRF Pools feature:

- PCRF Pool Definition - Definition of the logical concept of a PCRF pool. This includes configuring the following information about PCRF pools:

PCRF Pool Name	A string naming the PCRF Pool.
PCRF Pool Description	A string describing the PCRF Pool.
Subpool Indicator	An indicator that a sub-pool is defined for this PCRF Pool.
PRT Table ID	The PRT Table to be used for this PCRF pool.

- PCRF Sub-Pool Selection Rules - Rules to determine the PCRF sub-pool, if any, to which a new-session CCR-I is routed. This further qualifies the PCRF Pool based on the Origin-Host of the PCEF that originates the CCR-I. Note that absence of sub-pool rules for a PCRF Pool means that there are no sub-pools for the PCRF Pool and all new-session CCR-Is are routed to the PCRF Pool selected using the PCRF Pool Selection rules.

PCRF Pool	One of the PCRF Pools in the PCRF Pool Selection Rules. This is used as a key to determine a new PCRF Pool to be used for the subpool.
Priority	Rule priority
FQDN (PCEF Origin-Host FQDN)	An FQDN value or partial match.
PCRF Sub-Pool	One of the configured PCRF Pools.

A default PCRF Pool will be configured into the system upon installation of the PCRF Pool Feature. All configured APNs will be configured to map to the default PCRF Pool.

If there is an existing binding for the IMSI that matches the APN, the existing binding will always be used. This occurs even if there is a more specific rule that was configured after the binding was created. This avoids a split-binding scenario. A split binding exists when more than one PCRF is managing Gx sessions for the same PCEF.

If there are no existing bindings that match the Gx session, Policy DRA uses the PCRF Pool Selection Rules to determine the PCRF Pool to which the CCR-I message is to be routed.

After selecting the PCRF Pool, Policy DRA determines whether there are PCRF sub-pool rules for the selected PCRF Pool. The PCRF Sub-Pool rules consist of the FQDN of the Diameter peer that originated the new-binding CCR-I and a priority. If multiple rules match, the highest priority rule is used. If all of the matching rules have the same priority, the more specific rule takes precedence.

Note: The PCA GUI ensures that no two rules with the same specificity having the same priority.

The following list indicates the order of precedence, from most specific to least:

1. Origin-Host full FQDN value
2. Origin-Host partial match

If there is a matching PCRF sub-pool rule then the PCRF pool id indicated in the PCRF sub-pool rules is used for routing the CCR-I. If there are no matching PCRF sub-pool rules then the CCR-I is handled based on the PCRF Pool selection rules.

Finding or Creating a Binding

Figure 42: Find or Create a Binding shows the logic used for this task.

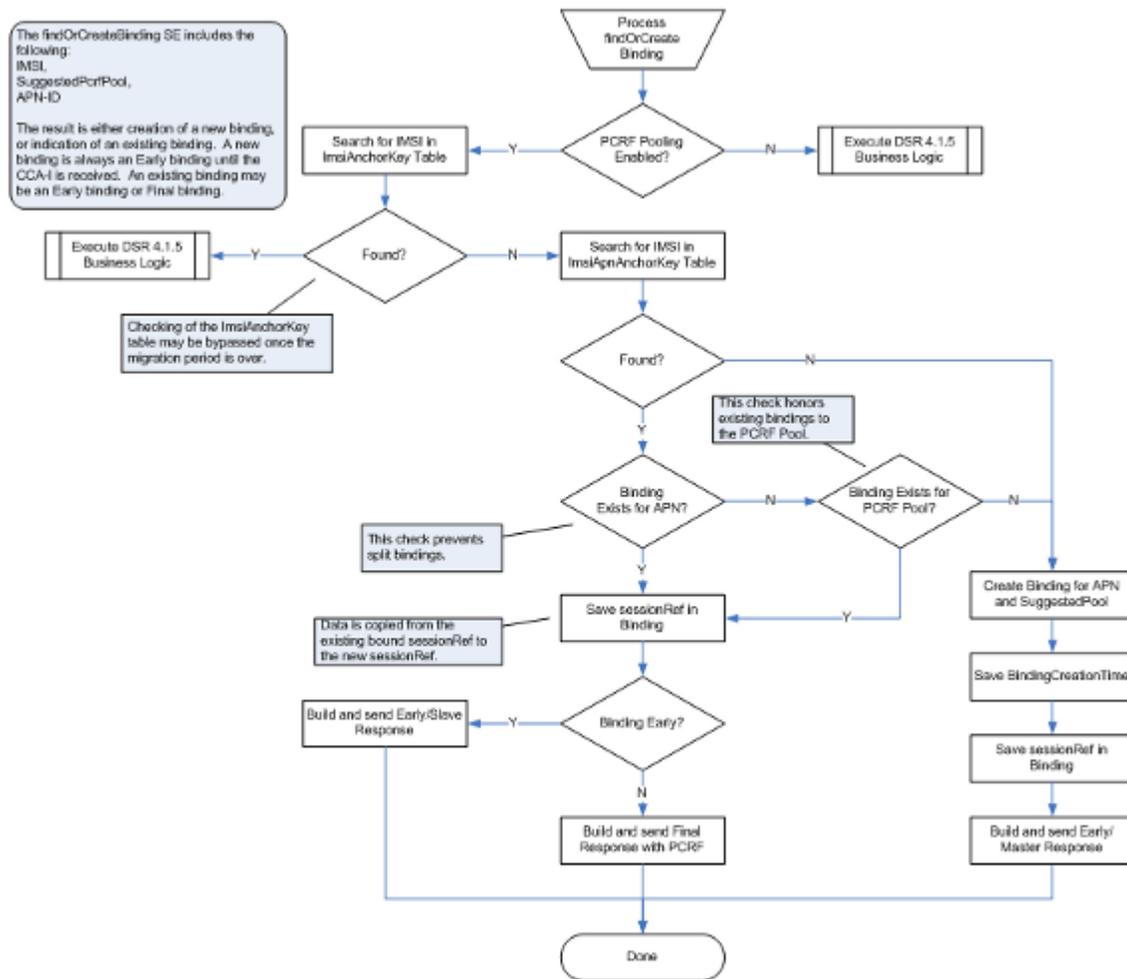


Figure 42: Find or Create a Binding

Routing to the selected PCRF Pool

If an existing binding is to be used to route a CCR-I, then the PCRF in that binding is used. If a new binding is to be created, after Policy DRA has selected PCRF Pool through a combination of the PCRF Pool Selection Rules and the PCRF Sub-Pool selection rules, then Policy DRA must select the PCRF peer that will own the binding.

The PRT Table ID mapped to the PCRF Pool points to the PRT table to be used for routing the CCR-I message.

All existing PRT functionality, including all valid PRT rules and load balancing capabilities, can be used for routing of the CCR-I to an instance within the PCRF pool.

Binding-capable Session Initiation Answer Processing

If a success response (for example, 2xxx) is received in a binding-capable session initiation answer (for example, CCA-I) the following actions occur:

- The answer message is relayed to the Policy Client that sent the request.
- A Session record is created with information related to the Diameter session.
- Alternate key binding records are created for the intersection of alternate keys configured in **Policy and Charging > Configuration > Policy DRA > Binding Key Priority** and alternate keys present in the binding-capable session initiation request.
- If the binding-capable session initiation request created a new binding, the early binding record is updated with the PCRF identified in the Origin-Host of the answer message and marked as a final binding.

If a failure response (for example, non 2xxx) is received in a binding-capable session initiation answer (if example, CCA-I) the following actions occur:

- The answer message is relayed to the Policy Client that sent the request.
- No session or alternate key records are created.
- If the binding-capable session initiation request created a new binding, the early binding record is removed.

Related Topic

[The P-DRA Database](#)

Binding-dependent Sessions

A binding is a relationship stored in the SBR-B between various subscriber data session identities, such as MSIDN/IP Address(es)/IMSI and the assigned PCRF. A session is a relationship stored in the Session SBR that associate additional sessions with a binding.

Binding-dependent sessions are created by Rx, Gx-Prime, or S9 version of Rx AAR messages.

Figure 43: Binding Dependent Session Initiation Request Processing Overview shows an overview of binding-dependent session initiation requests using IPv4 or IPv6 as correlation keys .

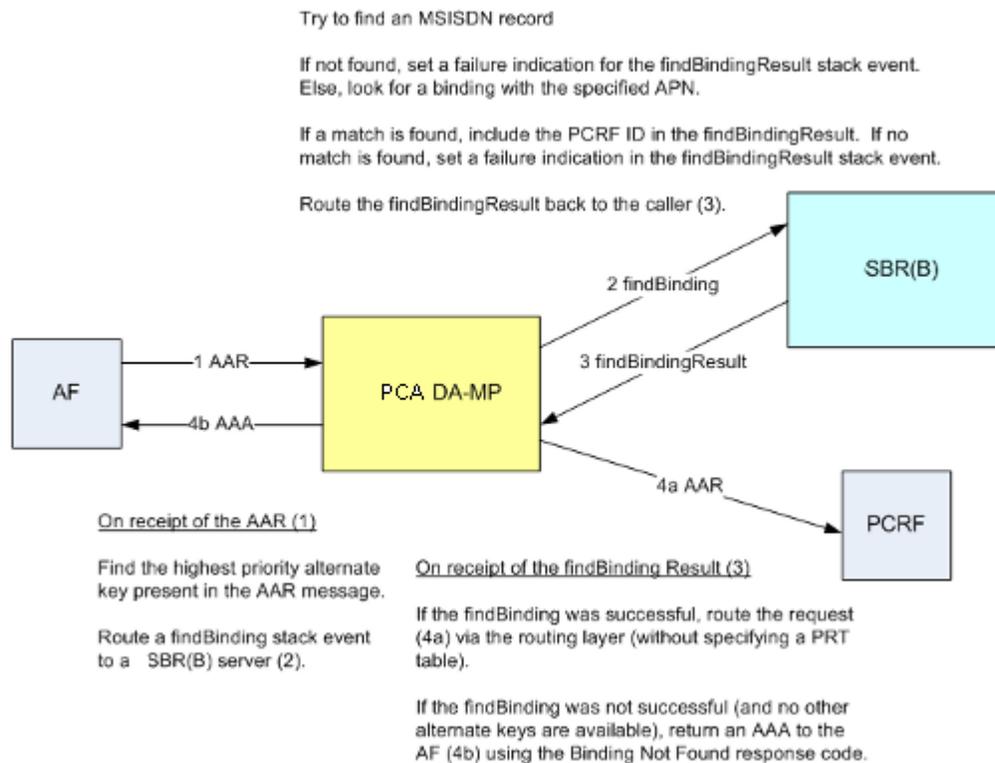


Figure 43: Binding Dependent Session Initiation Request Processing Overview

The following logic is used to locate an IP address Binding (used by binding-dependent interfaces):

- If PCRF Pooling is enabled, search the IpXAlternateKey table for a match, and if found, establish the alternate key. If the IP address is not found in the IpXAlternateKey table, search the IpXAlternateKeyV2 table for a match. If a match is found, the result is a binding found to PCRF X, which completes the process. If a match is not found, the result is binding not found, which completes the process.
- Binding-dependent session initiation requests using MSISDN as correlation key.
- Both MSISDN-Only and MSISDN+APN binding tables are audited.
- Both old and new IPv4 and IPv6 binding tables are audited.

Note: It is possible to determine the progress of data migration from the IMSI Only table by looking at the "Records Visited" statistic in the audit reports contained in event 22716. The records visited number shows how many IMSI Only records still remain. If no event 22716 occurs for the ImsiAnchorKey table, the migration is complete.

Binding-dependent Session Initiation Request Processing Rules and Requirements

Binding-capable request processing uses the binding key priority table to determine which keys present in the message should have alternate keys created in the binding database. Binding-dependent processing uses the binding key priority configuration to determine which keys to use and in what priority order when attempting binding correlation.

Related Topics

- [The P-DRA Database](#)

- [In-session Message Processing](#)
- [Topology Hiding Process](#)

In-session Message Processing

An in-session message is any message other than a session initiation request or session initiation answer for both binding-capable and binding-dependent interfaces.

The SBR Session Database holds session information that is used for routing in-session messages. A given session record is accessible on every SBR server a P-DRA Mateset. The Policy DRA application only adds a session record to the database when necessary. The P-DRA application always maintains session records for binding capable sessions (Gx, Gxx, and the S9 versions of Gx and Gxx), Gx-Lite sessions, and binding dependent sessions for which topology hiding is in effect.

Policy DRA has a mechanism similar to that of the PCRF (see [Session Integrity](#)), but the P-DRA does not need to process every in-session message. For example, the CCR-U message only has to be routed from the policy client to the PCRF. As a result, the Policy DRA does not contact the session record on CCR-U messages. Policy DRA only contacts the session record on RAR/RAA exchanges. Because the PCRF scheme for contacting sessions might differ from the Policy DRA mechanism for contacting sessions, it is possible that the Policy DRA could determine that a session is stale when the PCRF does not consider it to be stale.

If the Policy DRA simply removed a binding capable session that it considered to be stale, any keys associated with that session are also removed. In turn, this causes binding-capable (for example, Rx) sessions that rely on those keys to fail. The policy client and PCRF have no idea that there is a problem with the binding capable session and therefore does not re-create it, which causes the session and keys to be added back to the Policy DRA database.

Instead of removing a session considered to be stale, Policy DRA queries the policy client by sending an RAR message. If the policy client still thinks the session is valid, it responds with a success RAA (for example, 2xxx result code). This causes Policy DRA to contact the session and give it another interval of time before it can be considered to be stale again. If the policy client responds to the Policy DRA with an error indicating that the session is unknown (for example, 5002), Policy DRA removes its session and frees all resources associated with the session, including any keys that the session created.

Topology Hiding

For security reasons, network operators require the Diameter Routing Agents to be able to hide the PCRF topology from selected Policy Clients. When a Policy Client is configured to have the PCRF topology hidden from it, all Diameter messages (Request or Answer) that are sent to it need to be processed by the Policy DRA for Topology Hiding. The Policy DRA will place some configured Origin-Host and Origin-Realm values into the messages instead of the PCRF's real Origin-Host and Origin-Realm values.

Topology Hiding configuration is done on each Policy DRA DSR using the Policy and Charging section of the NOAM GUI. The configuration enables users to set the Topology Hiding function to be Enabled or Disabled for the Policy DRA node. After being enabled, the Topology Hiding function can be further

configured to apply for a specific Topology Hiding Scope, as summarized in [Table 18: Topology Hiding Scope Configuration](#):

- The Policy Clients with specific FQDNs
- All of the Policy Clients with Foreign Realm
- All the Policy Clients with Foreign Realm and the local Policy Clients with specific FQDNs
- All Policy Clients

The Host Name used for hiding PCRF topology is also configured. If a Policy Client is configured to use Topology Hiding, the Origin Host and Realm of all messages sent to the Policy Client will be changed to the configured Host Name.

The Diameter messages to be topology hidden from certain Policy Clients can be initiated from either Policy Clients (by a CCR from a PCEF) or Policy servers (by an RAR from a PCRF), or initiated by the Policy DRA (by an RAR generated by the Policy DRA). The handling of the Diameter messages for Topology Hiding will be different depending on the specific scenarios. To determine whether or not Topology Hiding is applicable for a Policy Client:

- For messages initiated from Policy Clients, the Policy DRA will compare the Origin-Host and Origin-Realm values in the incoming messages to the configured values.
- For messages initiated from Policy servers or by the Policy DRA, the Policy DRA compares the Destination-Host and Destination-Realm values to the configured values.
- For messages initiated by the Policy DRA, the Policy DRA will compare the Destination-Host and Destination-Realm of the Policy Client with the configured values to determine whether or not the Topology Hiding is applicable to the Policy Client.

Table 18: Topology Hiding Scope Configuration

Topology Hiding System Setting	Topology Hiding Scope Setting	Result
Disabled	N/A	No Topology Hiding is performed
Enabled	Specific Clients	Topology Hiding is performed for messages destined to the policy clients that are configured from the SOAM GUI Main Menu Policy and Charging > Configuration > Policy DRA > Policy CLients screen
	All Foreign Realms	Topology Hiding is performed for messages destined to the policy clients whose realms are different from the realm of the PCRF to be bound
	All Foreign Realms + Specific Clients	Topology Hiding is performed if either "All Foreign Realms" or "Specific Clients" condition is met
	All Messages	Topology Hiding is performed for all messages destined to all policy clients

Session Integrity

The Policy DRA application provides a capability called "Session Integrity" that addresses two potential problems:

Session Audit Premature Removal of Sessions

Policy DRA uses the mechanism of the Session Audit (see [PCA Data Auditing](#)), by which session-related resources can be freed in the event that the session is not torn down properly by Diameter signaling.

Session state synchronization between Policy DRA and Policy Client for binding capable sessions prevents the Session Audit (see [PCA Data Auditing](#)) from removing valid sessions that could be considered as "stale" .

If the Policy DRA simply removed a binding capable session that it considered to be stale, any keys associated with that session would also be removed. This in turn would cause binding capable dependent Rx or Gx-Prime sessions that rely on those keys to fail. The Policy Client and PCRF have no idea that there is a problem with the binding capable session and therefore will not re-create it, causing the session and keys to be added back to the Policy DRA database.

Instead of just removing a session that could be considered to be stale, Policy DRA queries the Policy Client. If the Policy Client responds indicating that the session is valid, Policy DRA waits for an interval of time before the session can be considered to be stale again. If the Policy Client responds indicating that the session is unknown, the Policy DRA will remove its session and free all resources associated with the session, including any keys that the session created.

Incomplete Session Data

In order to reduce Diameter signaling latency for policy signaling, Policy DRA attempts to relay Diameter messages before updating its various database tables. Provided that all database updates are created successfully and in a timely manner, this works very well. There are scenarios in which records cannot be successfully updated and the Policy Client and the PCRF are not aware of any problem. [Table 19: Policy DRA Error Scenarios for Session Integrity](#) describes specific scenarios where Policy DRA record creation failure can occur and the consequences of the failures for policy signaling.

In the case in which Policy DRA fails to create a binding record when a binding capable session is created, Policy DRA has already relayed the CCA-I message back to the PCEF (to reduce latency). The PCEF is unaware that one of the binding keys that it requested to be correlated with the subscriber's session does not exist in the Policy DRA. When a binding dependent Rx session attempts to use the failed binding key, the Rx or Gx-Prime session will fail because Policy DRA does not know which PCRF it should be routed to.

Incomplete or incorrect binding capable session data could persist for days because binding capable sessions can last as long as the UE (the subscriber's phone) is powered up and attached to the network. The PCEF that set up the binding capable session does not know that there is any problem with the correlation keys.

The solution for incomplete or incorrect data in the P-DRA is to compel the PCEF to tear down and reestablish the binding capable session in hopes that all P-DRA data updates will be created successfully on the next attempt. This is accomplished by P-DRA sending an RAR message containing a Session-Release-Cause AVP indicating that the session should be torn down.

Table 19: Policy DRA Error Scenarios for Session Integrity describes the specific scenarios in which the Policy DRA Session Integrity mechanism is required to remove a broken session. The first scenario is included to describe why Session Integrity does not apply to creation of an IMSI Anchor Key for a new binding.

Table 19: Policy DRA Error Scenarios for Session Integrity

Error Scenario	Policy DRA Behavior
Failed to create IMSI Anchor Key for new binding	<p>Because the CCR-I has not yet been forwarded to the PCRF, this scenario can be handled by sending a failure Answer to the Policy Client in the CCA-I response. In this case, no session is ever established.</p> <p>The Policy Client will attempt to re-establish the binding capable session.</p>
Failed to create binding capable session	<p>By the time Policy DRA creates a session record, the CCA-I has already been relayed to the Policy Client. If the session record cannot be created, no Alternate Keys are created. Policy DRA must cause the Policy Client to terminate the binding capable session (and re-create it).</p> <p>If the session record is not created, and no Alternate Keys are created, a binding dependent session that needs to use those keys will fail.</p>
Failed to create an alternate key	<p>By the time Policy DRA creates an alternate key record, the CCA-I has already been relayed to the Policy Client. If the Alternate Key record cannot be created, Policy DRA must cause the Policy Client to terminate the binding capable session (and re-create it).</p> <p>If Alternate Keys are not created, a binding dependent session that needs to use those keys will fail.</p>
Failed to update a new binding with the answering PCRF	<p>By the time Policy DRA updates the binding with the new PCRF (the PCRF that actually originated the CCA-I), the CCA-I has already been relayed to the Policy Client. If the IMSI Anchor Key record cannot be updated, Policy DRA must cause the Policy Client to terminate the binding capable session (and re-create it).</p> <p>If the IMSI Anchor Key cannot be updated with the PCRF that sent the CCA-I, the binding will still point to the Suggested PCRF, while the original Policy Client will have a session with the answering PCRF. This could lead to a subscriber (IMSI) having sessions with 2 different PCRFs.</p>

Note: Although Policy DRA maintains session state for binding dependent sessions when Topology Hiding applies to the Policy Client that created the session, the Policy DRA Session Integrity solution does not apply to binding dependent Rx sessions. The Rx or Gx-Prime RAR message differs from the Gx RAR message in that the Rx or Gx-Prime RAR message processing does not provide either a means to query a session or a means to cause a session to be released. If an Rx or Gx-Prime session is considered by Policy DRA to be stale, Policy DRA simply removes the session. If an Rx or Gx-Prime session is removed by Policy DRA audit or never successfully created, the next message in the Rx session will fail, causing the Policy Client to recreate the session.

Session Integrity Common Solution

The common solution for these two problems is based on the ability of Policy DRA to initiate binding capable Gx RAR Requests toward the Policy Client involved in the binding capable session. (Policy DRA does not relay an RAA received from a Policy Client to the PCRF associated with the session; the RAA is locally consumed by Policy DRA.)

Table 20: Session Integrity Conditions and Policy DRA Reaction describes the conditions that trigger the Policy DRA to send an RAR to the Policy Client. For each condition, the type of RAR is listed (Query or Release), and whether sending of the RAR is subject to throttling.

Table 20: Session Integrity Conditions and Policy DRA Reaction

Condition	RAR Type	Throttled	Comments
Session determined to be stale	Query	Y	See throttling description below.
Failed to create alternate key	Release	Y	Throttling is not needed in this case, but the error is detected on the Policy SBR server which already has the throttling mechanism for auditing and is therefore free for use.
Failed to create session record	Release	N	Quick teardown is desirable.
Failed to update binding when the answering PCRF differed from the Suggested PCRF	Release	N	Quick teardown is desirable.

When an RAR is subject to throttling, certain rules apply:

- Each session SBR performs RAR throttling independently
- The rate of RARs initiated by any given session SBR varies depending on the number of sessions in the list of sessions pending RAR treatment
- The rate of RARs initiated by a given session SBR is bounded to be no faster than on RAR every 20 ms and no slower than on RAR every 500 ms

- The fastest rate of RARs for a given session SBR is 1 per 20ms (50 per sec) when the list of sessions pending RAR treatment is 200 or more deep
- The slowest rate of RARs for a given session SBR is 1 per 500 ms (2 per sec) when the list of sessions pending RAR treatment is very small
- When the list of sessions pending RAR treatment is between 1 and 200 deep, the rate for sending RARs varies proportionally to the list depth between 2 per second and 50 per second
- If no response is received for an RAR for a given session, that RAR is not re-attempted more frequently than every 2.5 seconds
- If no response is received from the RARs for a given session, the session is removed from the P-DRA database after 12 attempts

When an RAR is not subject to throttling, the RAR is subject to transaction processing rules configured in the Diameter Routing Function.

When a query-type RAR is sent to ask the Policy Client if the session is valid, Policy DRA is looking for two result codes:

- An RAA response with a success result code indicates that the Policy Client still has the session. This causes Policy DRA to refresh the time the session can be idle before being considered as stale again.
- An RAA response with a result code of Unknown Session-Id indicates that the Policy Client no longer has the session. This causes the Policy DRA to remove the session and all of the session's keys.

An RAA response with any other result code is ignored.

Suspect Binding Removal

The Policy DRA function provides the capability to bind and correlate between policy service subscribers and policy servers (PCRFs) specified by the subscriber identifiers such as IMSIs, MSISDNs or IP addresses. All service requests on behalf of a subscriber from all relevant Diameter applications (i.e. Diameter interfaces such as Gx or Rx) will be routed to the bound PCRF so that policy decisions can be made with the knowledge of all the policies being applicable for that particular subscriber. Routing failures occur if a bound PCRF fails to operate or loses connection to the PCA after the binding is established.

P-DRA contains a set of configurable rules that invoke a series of suspect binding cleanup steps if the suspect binding removal rules are matched by Diameter Answer messages with error response.

These rules monitor the error responses received from PCRF or from the DSR routing layer on behalf of the PCRF, comparing the error responses (types and the number of error responses etc.) with the configured rules. Bindings are labelled as suspect if some rules are matched, cleaning up the suspect binding records in the SBR and invoking sending RAR procedure to the policy clients.

The rules also include the configuration of the timing when an action needs to be taken to start a suspect binding cleanup if it is detected, be it immediately or after a provisioned threshold number of detections is reached.

P-DRA initiates the suspect binding removal procedure if a matched suspect binding removal rule for a given subscriber is configured as **Yes to Remove Suspect Binding Immediately** in the PCA configuration, or the **Suspect Binding Removal Event Threshold** value is exceeded by the number of matched rules. The suspect binding removal procedure involves the following components/actions:

- coordination among the P-DRA MP, Binding and Session SBRs via back-end communications (stack event messages) between them
- clean up the binding SBR databases for the suspect binding records
- sending Diameter RAR messages to the policy clients (PCEF) to trigger the policy clients to initiate the removal of the sessions related to those suspect bindings

Specifically, the initiation of the suspect binding removal procedure can be triggered by either P-DRA MP, if the matched suspect binding removal rule is configured as **Yes to Remove Suspect Binding Immediately** in the PCA configuration, or by the binding SBR, if the **Suspect Binding Removal Event Threshold** value is exceeded by the number of Suspect Binding Removal Events. In either case, the session SBR that receives the request to clean up the sessions related to the suspect bindings invokes the Session Integrity Service (SIS) functionality.

The SIS is used by session SBRs to verify session integrity and to trigger policy clients to remove the problematic sessions. This functionality allows external servers such as a binding SBR or DA-MP to invoke the SIS service on the session SBR. SIS on the session SBRs can be used to provide two types of session integrity handling:

1. requesting DA-MP to send Query type RAR messages to policy clients for verifying session integrity of some session records stored on the Session SBR. Based on the RAA responses, the sessions in question will either be re-authorized or torn down in a later time by the policy clients
2. requesting DA-MP to send Release type RAR messages to policy clients for tearing down the sessions that the Session SBR has records. A Release RAR will request the policy clients to tear down the session in question by sending CCR-T messages to the bound PCRF.

After receiving a session release RAR, PCEF may initiate a CCR-T request to remove the binding capable (Gx or Gxx) session and the corresponding IP-CAN session that in turn triggers the involved UE to re-initiate the IP-CAN session. This may result in a new binding capable session that are bound to a working PCRF this time by P-DRA. If the subscriber has multiple binding capable sessions (Gx or Gxx) associated with the same PCRF, P-DRA initiates session release RAR requests for each of these sessions.

If a subscriber has multiple Bound Sessions associated with a failed PCRF, it is possible that one Bound Session is torn down and is being re-established before cleaning up other remaining binding capable sessions. The suspect binding removal procedure marks the Bound Sessions to this failed PCRF in the binding SBR as pending delete such that no newly created bindings could use the failed PCRF anymore until it returns functional. The execution of the suspect binding removal procedure leads to removing the Bound Sessions associated with the suspect bindings for a given subscriber and a PCRF. Therefore, the Bound Sessions of the same subscriber associated with different PCRFs, if they exist, are not impacted.

For information on how to configure Suspect Binding Removal Rules, refer to [Suspect Binding Removal Rules](#).

Chapter 5

Online Charging DRA Overview

Topics:

- *Online Charging DRA Functions.....130*
- *Session State Maintenance.....133*
- *Gy/Ro Diameter Request Message Processing.136*
- *Gy/Ro Diameter Answer Message Processing.139*

This section gives an overview of the Online Charging DRA function.

Online Charging DRA Functions

The OC-DRA functionality of PCA provides the following functions for processing Diameter messages over Gy/Ro reference points for Online Charging:

- OCS Selection and Routing
- Session State Maintenance

If regionalized routing is required, DSR Range Based Address Resolution (RBAR) application can also be optionally invoked prior to OC-DRA invocation.

OCS Selection and Routing

Gy/Ro session initiation request (i.e. CCR-I) messages received from online charging clients to initiate credit-control sessions are load balanced across a collection of OCS servers connected to the PCA DSR that can serve the Diameter Request. Subsequent Gy/Ro CCR-U/T messages within the session are routed to the same OCS that served the CCR-I either by means of destination-host routing or by the stateful mechanism.

OC-DRA supports the following OCS Pool Selection modes for selecting the specific collection of OCS servers connected to the PCA DSR in which session initiation requests are to be load balanced across:

- Single PCS Pool
- Multiple OCS Pools

OCS Pool Selection modes are configurable via the NOAM GUI Main Menu: **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options > OCS Pool Selection Mode.**

One-time Gy/Ro credit-control events received from online charging clients are handled in the same manner as session initiation request (CCR-I) messages and are load balanced across a collection of OCS servers connected to the DSR that can serve the Diameter request.

Single OCS Pool Mode

When OC-DRA is operating in "Single OCS Pool" mode, session initiation requests are load balanced across all available OCS servers connected to the PCA DSR.

Note: OC-DRA removing the Destination-Host AVP from the Online Charging Diameter Request when operating in the Single OCS Pool mode ensures that the request will not be accidentally rejected by the PRT or by the OCS Server if the Destination-Host AVP contained the PCA DSR's hostname.

In this mode, OC-DRA removes the Destination-Host AVP (if present) from the session initiation request and forwards it to DRL where PRT/RL is used to route the session initiation request to one of the available OCS servers connected to the PCA DSR.

Note: OC-DRA does not specify the PRT/RL that is used by DRL to route Diameter request messages to an available OCS. The PRT selected for routing is based on DRL's PRT precedence rules

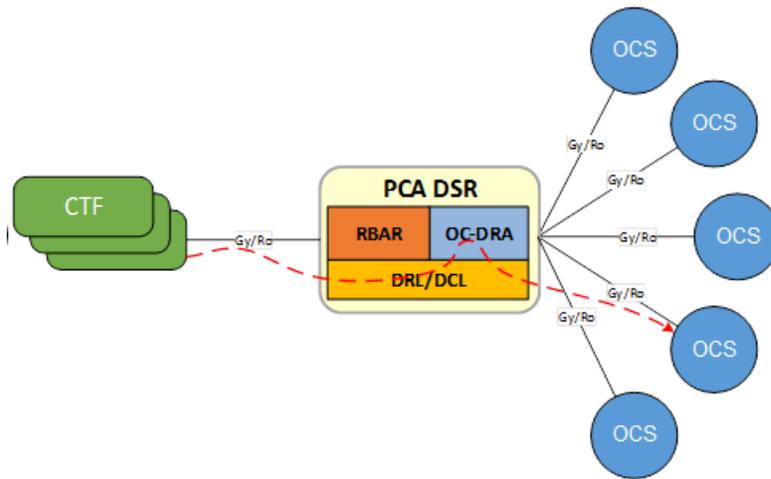


Figure 44: Local OCS Server Selection

Multiple OCS Pools Mode

When OC-DRA is operating in "Multiple OCS Pools" mode, session initiation requests are loaded balanced across a pool of available OCS servers connected to the PCA DSR that can serve the request. OC-DRA relies on the RBAR application to be invoked prior to the invocation of PCA to populate the Destination-Host AVP and/or Destination-Realm AVP in session initiation requests. The hostname that RBAR uses to populate session initiation request's Destination-Host AVP can be a real hostname or a virtual hostname that is used to represent a pool of OCS servers that can serve the request. OC-DRA forwards the session initiation request without any modification to DSRL where PRT/RL is used to route the session initiation request to one of the available OCSs within the selected pool of OCS servers.

Note: OC-DRA removing the Destination-Host AVP from the Online Charging Diameter Request when operating in the Single OCS Pool mode ensures that the request will not be accidentally rejected by the PRT or by the OCS Server if the Destination-Host AVP contained the PCA DSR's hostname

Note: OC-DRA does not specify the PRT/RL that is used by DRL to route Diameter request messages to an available OCS. The PRT selected for routing is based on DRL's PRT precedence rules.

Note: When operating in "Multiple OCS Pools" mode, OC-DRA assumes (i.e., does not verify) that RBAR was previously invoked to populate Destination-Host and/or Destination-Realm AVP of CCR-I/E messages received. It is entirely up to the operator to ensure that RBAR is invoked prior to PCA invocation as PCA will forward CCR-I/E messages received without any modification to DRL for routing using the PRT. A failure of RBAR invocation prior to PCA when OC-DRA is operating in "Multiple OCS Pools" mode may lead to unexpected routing results (e.g., unable to route), each depending on PRT/RL configuration.

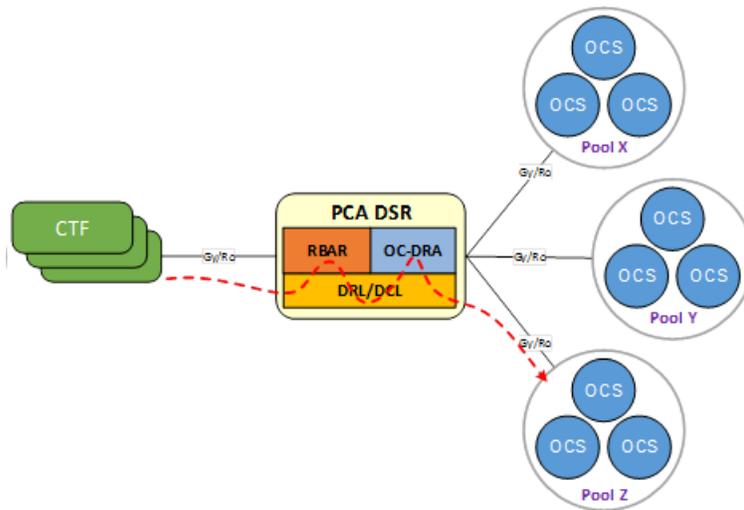


Figure 45: Local OCS Server Pool Selection

Regionalized Routing

Operators/service providers may have OCS deployments which are segmented based on ranges of subscriber identities (i.e. MSISDNs) such that a given group of OCSs can only serve the subscriber range it has been assigned to serve. To support this architecture, the DSR RBAR application can be provisioned with higher priority ART rules to be invoked prior to the invocation of PCA to perform regionalized routing based on subscriber's identity. If RBAR invocation fails, the DAL configuration should be provisioned such that RBAR is invoked on the mate DSR or an Answer response with a non-successful Result-Code/Experimental-Result AVP is generated and sent to the originator of the Diameter transaction.

In regionalized OCS deployments, it is likely that RBAR is invoked at one DSR NE (DSR that has direct peer connectivity with the online charging client) while PCA OC-DRA is invoked at another NE (DSR that has direct peer connectivity with OCSs in the serving region). ART rules corresponding to PCA for OC-DRA invocation should be configured such that PCA is invoked only if the Destination-Host and/or Destination-Realm is served by the same DSR (where RBAR was invoked). In cases where the Destination-Host and/or Destination-Realm are not served by the same DSR, the Request is routed to the DSR serving the Destination-Host and/or Destination-Realm (called the target DSR) and PCA is invoked for OC-DRA on the target DSR. OC-DRA invoked on the target DSR can be configured to operate in any of the OCS Pool Selection modes for routing within the target region.

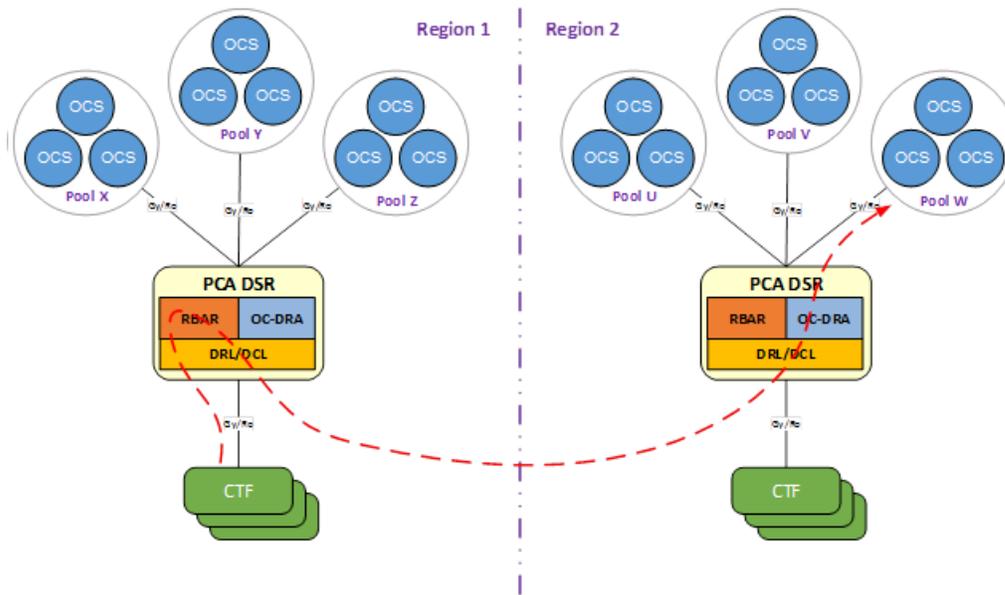


Figure 46: Regionalized OCS Server Pool Selection

Note: In cases where a session initiation request (CCR-I) is routed to OC-DRA on a target DSR, RBAR may have to be invoked for subsequent in-session request (CCR-U/T) messages for that session as well (e.g., for non-proxy-compatible online charging clients). If MSISDN is not present in CCR-U/T messages, regionalized routing cannot be supported

Session State Maintenance

Online charging clients (CTFs) are expected to be proxy-compatible, thus capable of learning the OCS name from the Origin-Realm and Origin-Host of the answer to the session initiation request (i.e. CCA-I). This OCS name should be used as the Destination-Realm and Destination-Host of all subsequent in-session requests originated by the online charging client. Online charging clients that are proxy-compatible allow OC-DRA to host-route in-session requests to an OCS. However, there are online charging clients that are not proxy-compatible. These online charging clients may omit the Destination-Host AVP from requests or include the Destination-Host AVP with the OC-DRA Diameter hostname.

In addition to non-proxy-compatible online charging clients, there may also be online charging servers (OCSs) that are not capable of learning the name of the online charging client that originated the session initiation request, but need to be for the purpose of sending re-authorization requests (RARs).

In order to support such online charging clients and servers and to ensure that in-session requests (i.e., CCR-Us and CCR-Ts) are sent to the same online charging server that answered the session initiation request and re-authorization request (i.e., RARs) messages are sent to the online charging client that originated the session initiation request, OC-DRA provides the capability to maintain session state based on configuration and message content.

Session state is only applicable when Session-based charging is used and does not apply to Event-based online charging. As such, session state only applies to Diameter messages used for session based

charging, which include CCR-I/U/T, CCA-I/U/T, RAR and RAA messages. Event-based charging Diameter messages, CCR-E and CCA-E do not create sessions. Thus, session state is not maintained for these messages.

Given the varying capabilities of online charging clients and servers from various vendors, OC-DRA provides the ability to configure and maintain session state for selective clients and servers while not maintaining session state for clients and servers that are capable of learning server and client names from previous Diameter transactions. OC-DRA supports the Session State configuration settings.

Table 21: Session State Configuration Settings

Session State Setting	Scope	Description
No Messages	-	Session State is not maintained for any Gy/Ro sessions.
All Messages	-	Session State is maintained for all Gy/Ro sessions.
Specific Messages	Specific Realms	Session State is maintained for all sessions that are originated from an online charging client in a specific realm or being sent to an online charging server in the specific realm as configured in Policy and Charging > Configuration > Online Charging DRA > Realms.
	Specific Clients	Session State is maintained for all sessions that are originated from specific online charging clients as configured in Policy and Charging > Configuration > Online Charging DRA > CTFs.
	Specific Realms + Specific Client	Session State is maintained for all sessions that are originated from specific realms or specific online charging clients as scoped above.
	Specific Servers	Session State is maintained for all sessions that are destined to specific online charging servers as configured in Policy and Charging > Configuration > Online Charging DRA > OCSs.
	Specific Realms + Specific Servers	Session State is maintained for all sessions that are originated from specific realms or destined to specific online charging servers.

Session State Setting	Scope	Description
	Specific Realms + Specific Clients + Specific Servers	Session State is maintained for all sessions that are originated from specific realms or specific online charging clients or destined to specific online charging servers.

OC-DRA relies on the SBR s that may be local to the site or remotely located to store the session state information for the life of the session when session state is maintained for a Gy/Ro session. The session state table is keyed by the Diameter Session-Id, a long string that is defined by Diameter to be "globally and eternally unique".

Session State Information	Description
Session ID	Session Identifier from Session initiation request (CCR-I) Session-Id AVP
CTF Realm	Online Charging Client Realm
CTF Hostname	Online Charging Client FQDN
OCS Realm	Online Charging Server Realm
OCS Hostname	Online Charging Server FQDN
Subscriber Identifier	Identification of the user that is going to access the service e.g. MSISDN
Access Point Name	Access Point Name (APN) the user is connected to.
Creation Time	The timestamp the session state was created in the Session SBR.
Last Touch Time	The timestamp of when the session state was last accessed for CCR-U or RAR message processing.

On receipt of an in-session request for a Gy/Ro session whose state is required to be maintained based on session state configuration and message content, OC-DRA replaces the in-session request Destination-Realm and Destination-Host with the OCS or CTF Realm and Host (depending on the direction of the Diameter request message) obtained from the session state associated with the received Session-Id maintained in the Session SBR prior to forwarding it to DRL for routing.

Session state maintained in the Session SBR is considered active as long as CCR-U's and RARs continue to be received with the same Session-Id and session state continues to be configured to be maintained on behalf of either the online charging client or server. Session state is considered stale if the time between requests for a particular session exceeds the Stale Session Time-out value (in hours) as configured by the following GUI screens:

- **Main Menu: Policy and Charging > Configuration > Access Point Names** or
- **Main Menu: Policy and Charging > Configuration > General Options** (if session is not associated with a configured Access Point Name)

All stale session states maintained in the Session SBR database are automatically removed by the Session Audit.

Gy/Ro Diameter Request Message Processing

On receipt of a Gy/Ro Diameter Credit Control Application Request message, OC-DRA performs validation checks on the contents of the message before it attempts to route the message. Validation is limited to header information and routable Attribute Value Pairs (AVPs) that are used by OC-DRA for making processing decisions for routing

OC-DRA validates the Application ID and Command Code in the Diameter Request message for consistency. OC-DRA supports the Gy/Ro DCCA messages. If OC-DRA receives a Diameter Request message with a Command Code that is not supported, PCA will abandon message processing and send an Answer message with Result-Code set to DIAMETER_COMMAND_UNSUPPORTED (3001) and Error-Message AVP to the downstream peer that initiated the Diameter transaction.

OC-DRA also makes Diameter Request message processing decision based on the small subset of AVPs for online charging. Those Diameter AVPs that are used specifically by OC-DRA for making routing decisions and maintaining session state. These AVPs are shown in [Table 22: Diameter AVPs used by OC-DRA for Request Message Processing](#) and marked "M", "O", "O_M", or "-" to indicate which ones are mandatory, optional, optional-mandatory or not used for each of the supported Gy/Ro Diameter Credit Control Application Request messages.

Note: AVPs marked as "O_M" are optional, but are mandatory if the optional Grouped AVP in which they are a member is present in the Diameter message.

Table 22: Diameter AVPs used by OC-DRA for Request Message Processing

AVP Name	AVP Code	Used In		Value Type	Description
		CCR	RAR		
Session-Id	263	M	M	UTF8String	Contains the session identifier.
Origin-Host	264	M	M	DiamIdent	Contains the end point that originated the Diameter message.
Origin-Realm	296	M	M	DiamIdent	Contains the realm of the originator of the Diameter message.
Destination-Host	293	O	O	DiamIdent	Contains the end point to which the Diameter message is to be routed.
Destination-Realm	283	M	M	DiamIdent	Contains the realm to which the Diameter message is to be routed.
Auth-Application-Id	258	M	M	Unsigned32	Contains the application ID of the Diameter Credit Control Application which is 4.
CC-Request-Type	416	M	-	Enumerated	Contains the transfer type: event for event based charging and

AVP Name	AVP Code	Used In		Value Type	Description
		CCR	RAR		
					initial, update, terminate for session based charging.
User-Name	1	O	-	UTF8String	Contains the user name in the format of a NAI according to RFC 6733.
Subscription-Id	443	O	-	Grouped	Contains the identification of the user that is going to access the service in order to be identified by the OCS.
Subscription-Id-Type	450	O _M	-	Enumerated	Contains the type of the identifier, e.g. value 0 is used for the international E.164 format according to ITU-T E.164 numbering plan. This AVP is a member of Subscription-Id Grouped AVP.
Subscription-Id-Data	444	O _M	-	UTF8String	Contains the user data content e.g. the MSISDN. This AVP is a member of Subscription-Id Grouped AVP.
Called-Station-Id	30	O	-	UTF8String	Contains the Access Point Name (APN) the user is connected to.

OC-DRA validates all the AVP listed except for those that have already been validated by DCL and DRL prior to the invocation of OC-DRA which include Origin-Host AVP, Origin-Realm AVP and Destination-Realm AVP.

Once validation of the Diameter Request content is complete, OC-DRA performs Diameter Request message processing and routing.

Session Initiation Request Message Processing

Gy/Ro Credit-Control-Requests (CCRs) with CC-Request-Type AVP set to INITIAL_REQUEST (1) received from online charging clients to initiate credit-control sessions are load balanced across a collection of OCS servers connected to the DSR that can serve the Diameter request.

OC-DRA supports the following operating modes for selecting the specific collection of OCS servers connected to the DSR in which session initiation requests are to be load balanced across:

- Single OCS Pool
- Multiple OCS Pools

If OC-DRA is configured to operate in "Single OCS Pool" mode, OC-DRA removes the Destination-Host AVP from the received session initiation request (if present) and forwards it to DRL where PRT/RL will be used to route the session initiation request to one of the available OCS servers connected to the DSR.

If OC-DRA is configured to operate in "Multiple OCS Pools" mode, OC-DRA forwards the session initiation request without modification to DRL where PRT/RL will be used to load balance the session initiation request across a subset (i.e., one of several pools) of available OCS servers connected to the DSR that can serve the request. In this mode, OC-DRA relies on RBAR to be invoked prior to OC-DRA invocation to populate the Destination-Host AVP and/or Destination-Realm AVP in session initiation requests. The hostname that RBAR uses to populate session initiation request's Destination Host AVP can be a real hostname or a virtual hostname that is used to represent a pool of OCS servers that can serve the request.

Subsequent Gy/Ro CCR messages with CC-Request-Type AVP set to UPDATE_REQUEST | TERMINATION_REQUEST within the session are routed to the same OCS that served the CCR-I either by means of destination-host routing or by the stateful mechanism

In-Session Request Message Processing

CCR

Credit-Control-Requests (CCRs) with CC-Request-Type set to UPDATE_REQUEST (2) received from online charging clients to update existing credit-control sessions are routed to the same online charging server that served the session initiation (i.e., CCR-I) request.

OC-DRA determines whether session state is maintained based on session state configuration and message content. If session state is not maintained, OC-DRA routes the CCR-U without modification, expecting the online charging client to have set its Destination-Host AVP value to the hostname of the same online charging server that served the session initiation request. If session state is maintained, OC-DRA queries the Session SBR to retrieve and refresh the session state associated with the received Session-Id by sending a findAndRefreshOcSession stack event to the Session SBR. If session state is found, OC-DRA replaces the Destination-Realm and Destination-Host in the CCR-U with the realm and hostname of the online charging server obtained from the session state and forwards it to DRL for routing. If session state is not found or an SBR error is encountered, OC-DRA will determine how to handle the message based upon the user-configurable "Session State Unavailable Action".

CCR-T

Credit-Control-Requests (CCRs) with CC-Request-Type set to TERMINATION_REQUEST (3) received from online charging clients to terminate existing credit-control sessions are routed to the same online charging server that served the session initiation request (i.e., CCR-I).

OC-DRA determines whether session state is maintained based on session state configuration and message content. If session state is not maintained, OC-DRA routes the CCR-T without modification, expecting the online charging client to have set its Destination-Host AVP value to the hostname of the same OCS that served the session initiation request. If session state is maintained, OC-DRA queries the Session SBR to retrieve and remove the session state associated with the received Session-Id by sending a findAndRemoveOcSession stack event to the Session SBR. If session state is found, OC-DRA replaces the Destination-Realm and Destination-Host in the CCR-T with the realm and hostname of the online charging server obtained from the session state and forwards it to DRL for routing. If session state is not found or an SBR error is encountered, OC-DRA will determine how to handle the message based upon the user-configurable "Session State Unavailable Action".

RAR

Re-Auth-Request (RARs) received from online charging servers to re-authorized existing credit-control sessions are routed to the online charging client that originated the session initiation request (i.e., CCR-I).

OC-DRA determines whether session state is maintained on the behalf of the online charging server based on session state configuration and message content. If session state is not maintained, OC-DRA routes the RAR without modification, expecting the online charging server to have set its Destination-Host AVP value to the hostname of the online charging client that originated the session initiation request. If session state is maintained, OC-DRA queries the Session SBR to retrieve and refresh the session state associated with the received Session-Id by sending a findAndRefreshOcSession stack event to the Session SBR. If session state is found, OC-DRA replaces the Destination-Realm and Destination-Host in the RAR with the realm and hostname of the online charging client obtained from the session state and forwards it to DRL for routing. If session state is not found or an SBR error is encountered, OC-DRA will determine how to handle the message based upon the "Session State Unavailable Action".

Routing In-Session Request Messages when Unable to Retrieve Session State

When OC-DRA cannot successfully process an in-session request (i.e., CCR-U/T and RAR) due to its inability to retrieve session state associated with the received Session-Id from the SBR (either session state is not found or an SBR error is encountered), the operator is provided the flexibility to determine how to handle the message based upon **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options** "Session State Unavailable Action" user-configurable setting.

The following user-configurable Session State Unavailable Actions are supported by OC-DRA:

- Route To Peer (via PRT)
- Send an Answer response with a user-defined Result-Code/Experimental-Result AVP value (default)

When configured to forward route the message, OC-DRA forwards the in-session request message to DRL for routing using PRT. When configured to reject the message, OC-DRA abandons request message processing, generates and sends an Answer response using the Result-Code configured for error condition to the peer that initiated the Diameter transaction.

Event Request Message Processing

Credit-Control-Requests (CCRs) with CC-Request-Type AVP set to EVENT_REQUEST (4) received from online charging clients (CTFs) are load balanced across a collection of OCS servers connected to the DSR that can serve the Diameter request in the same manner as Credit-Control-Requests (CCRs) with CC-Request-Type AVP set to INITIAL_REQUEST (1).

Gy/Ro Diameter Answer Message Processing

When OC-DRA forwards a Diameter Request message to DRL for routing, it requests that the corresponding Answer response message is forwarded back to PCA for Answer message processing.

Note: PCA requests that DRL forward all Gy/Ro Answers back to PCA for Answer message processing in order to maintain measurements. PCA processing all Gy/Ro Answers for accounting purposes is

not expensive since the OC-DRA and the routing layer are guaranteed to be on the same physical server (different threads in the same process) and avoids lots of explaining about why some measurements are not pegged.

On receipt of a Gy/Ro Diameter Credit Control Application Answer message, OC-DRA performs validation checks on the contents of the message before it attempts to relay the message. Validation is limited to header information and routable Attribute Value Pairs (AVPs) that are used by OC-DRA for making processing decisions for Answer message routing.

OC-DRA validates the Application ID and Command Code in the Diameter Answer message for consistency. OC-DRA supports the Gy/Ro DCCA Answer messages. If OC-DRA receives a Diameter Answer message with a Command Code that is not supported, PCA will send the Answer message without modification to the downstream peer that initiated the Diameter transaction.

Command Name	Abbreviation	Code	Source	Destination
Credit-Control-Answer	CCA	272	OCS	CTF
Re-Auth-Answer	RAA	258	CTF	OCS

OC-DRA makes Diameter Answer message processing decisions based on a small subset of AVPs defined in the Diameter protocol for online charging. Those Diameter AVPs that are used specifically by OC-DRA for making routing decisions. These AVPs are shown in [Table 23: Diameter AVPs used by OC-DRA for Answer Message Processing](#) and marked "M" or "-" to indicate which ones are mandatory or not used for each of the supported Gy/Ro DCCA Answer messages.

Table 23: Diameter AVPs used by OC-DRA for Answer Message Processing

AVP Name	AVP Code	Used In		Value Type	Description
		CCR	RAR		
Session-Id	263	M	M	UTF8String	Contains the session identifier.
Result-Code	268	M	M	Unsigned32	Contains whether a particular request was completed successfully (i.e., 2xxx) or an error occurred (non-2xxx).
Origin-Host	264	M	M	DiamIdent	Contains the end point that originated the Diameter message.
Origin-Realm	296	M	M	DiamIdent	Contains the realm of the originator of the Diameter message.
Auth-Application-Id	258	M	M	Unsigned32	Contains the application ID of the Diameter Credit Control Application which is 4.
CC-Request-Type	416	M	-	Enumerated	Contains the transfer type: event for event based charging and initial, update, terminate for session based charging.

OC-DRA validates all the AVPs listed except for those that have already been validated by DCL and DRL prior to the invocation of OC-DRA which includes the Origin-Host AVP and the Origin-Realm AVP. OC-DRA Diameter message header field and AVP validation requirements.

Once validation of the Diameter Answer content is complete, OC-DRA performs Diameter Answer message processing and routing.

Session Initiation Answer Message Processing

Credit-Control-Answer (CCA) messages with CC-Request-Type AVP set to INITIAL_REQUEST (1) received from online charging servers (OCSs) are routed without any modifications to the online charging client (CTF) that initiated the Diameter transaction

If a CCA-I is received with a successful Result-Code AVP (i.e., 2xxx), OC-DRA verifies that the Origin-Host of the answering OCS is configured as an OCS at the local site (**Policy and Charging > Online Charging DRA > Configuration > OCSs**). If the OCS is not configured at the local site, OC-DRA asserts timed Alarm 22730 Policy and Charging Configuration Error (Refer to the *DSR Alarms and KPIs Reference* for further details). If the answering OCS is configured at the local site and session state needs to be maintained as based on session state configuration and message content, OC-DRA stores the session information in the Session SBR by sending a createOcSessions tack event to the Session SBR.

In-Session Answer Message Processing

CCA-U

Credit-Control-Answer (CCA) messages with CC-Request-Type AVP set to UPDATE_REQUEST (2) received from online charging servers (OCSs) are routed to the online charging client (CTF) that initiated the Diameter transaction.

If a CCA-U message with Result-Code AVP set to DIAMETER_UNKNOWN_SESSION_ID (5002) is received and session state is maintained based on session state configuration and message content, OC-DRA will remove the session state associated with the received Session-Id by sending a findAndRemoveOcSession stack event to the Session SBR.

CCA-T

Credit-Control-Answer (CCA) messages with CC-Request-Type AVP set to TERMINATION_REQUEST (3) received from online charging servers (OCSs) are routed to the online charging client (CTF) that initiated the Diameter transaction.

RAA

Re-Auth-Answer (RAA) messages received from online charging clients (CTFs) are routed to the online charging server (OCS) that initiated the Diameter transaction.

If a RAA message with Result-Code AVP set to DIAMETER_UNKNOWN_SESSION_ID (5002) is received and session state is maintained based on session state configuration and message content, OC-DRA will remove the session state associated with the received Session-Id by sending a findAndRemoveOcSession stack event to the Session SBR.

Event Answer Message Processing

Credit-Control-Answer (CCA) messages with CC-Request-Type AVP set to EVENT_REQUEST (4) received from online charging servers (OCSs) are routed to the online charging client (CTF) that initiated the Diameter transaction.

DRL-Initiated Answer Message Processing

Answer messages can also be initiated by DRL for a variety of reasons. For example, when DRL is processing a Diameter Request message, it may encounter a routing failure or an operator instruction (e.g., PRT rule) which requires abandoning transaction routing and sending an Answer response.

On receipt of a DRL-initiated Gy/Ro Diameter Answer, OC-DRA updates the Diameter Answer's Result-Code AVP using the "Unable To Route" Result-Code and Error-Message AVP and sends it to the downstream peer that initiated the Diameter transaction.

Chapter 6

Policy and Charging Configuration

Topics:

- [Policy and Charging Configuration Overview.144](#)
- [Pre-Configuration Activities.....148](#)
- [NOAM Configuration.....155](#)
- [SOAM Configuration.....203](#)
- [Post-Configuration Activities.....259](#)

The **Policy and Charging > Configuration** GUI pages for Policy and Charging components provide fields for entering the information needed to manage Policy and Charging configuration in the DSR.

Policy and Charging Configuration Overview

The **Policy and Charging > Configuration** GUI pages for Policy and Charging components provide fields for entering the information needed to manage Policy and Charging configuration in the DSR.

The Policy and Charging application must be activated in the system before Policy and Charging configuration can be performed.

The DSR 3-tiered Operations, Administration, and Maintenance (OAM) topology is required for the Policy and Charging application. 3-tiered OAM topology consists of the following tiers:

- A pair of NOAM servers running in active/standby redundancy

OAM configuration is performed on the NOAM.

As shown in *Figure 47: GUI Structure for 3-tiered DSR Topology with Policy and Charging for NOAM*, network-wide Policy and Charging configuration is performed on the NOAM.

- A pair, triplet, or quadruplet of SOAM servers at each site running in active/standby, active/standby/spare redundancy, or or active/standby/spare/spare redundancy

Diameter protocol configuration is done on the SOAM.

Most of the OAM configuration components are viewable on the SOAM.

Most DSR Application configuration is done on the SOAM.

As shown in *Figure 48: GUI Structure for 3-tiered DSR Topology with Policy and Charging for SOAM*, site-specific configuration for Policy and Charging is performed on the SOAM; some network-wide Policy and Charging configuration components are viewable on the SOAM.

- A set of MP servers, which can host signaling protocol stacks (for example, DA-MPs), or in-memory database servers (for example, Session Binding Repository [SBR])

An optional pair of Disaster Recovery NOAMs can be configured to manually take over in the event of loss of both the active and standby NOAMs

The three tiers allow configured data to be replicated down to the MP servers, and measurements, events, and alarms to be merged up to the OAM servers.

3-tiered topology allows administrators to access all DSR GUI pages from a single sign-on. An administrator can access the DSR SOAM when logged into the DSR NOAM, without needing to re-enter login credentials.

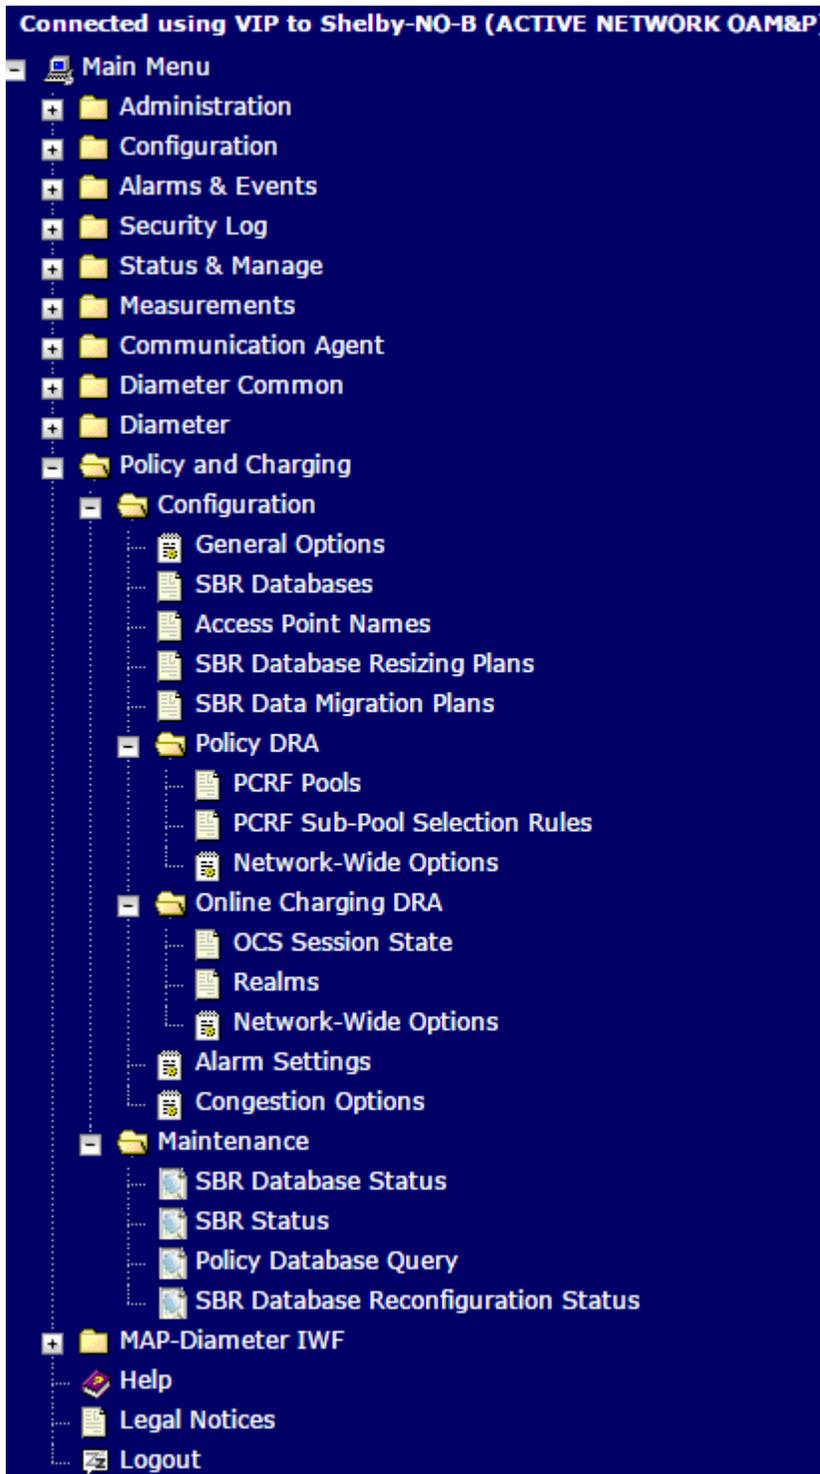


Figure 47: GUI Structure for 3-tiered DSR Topology with Policy and Charging for NOAM

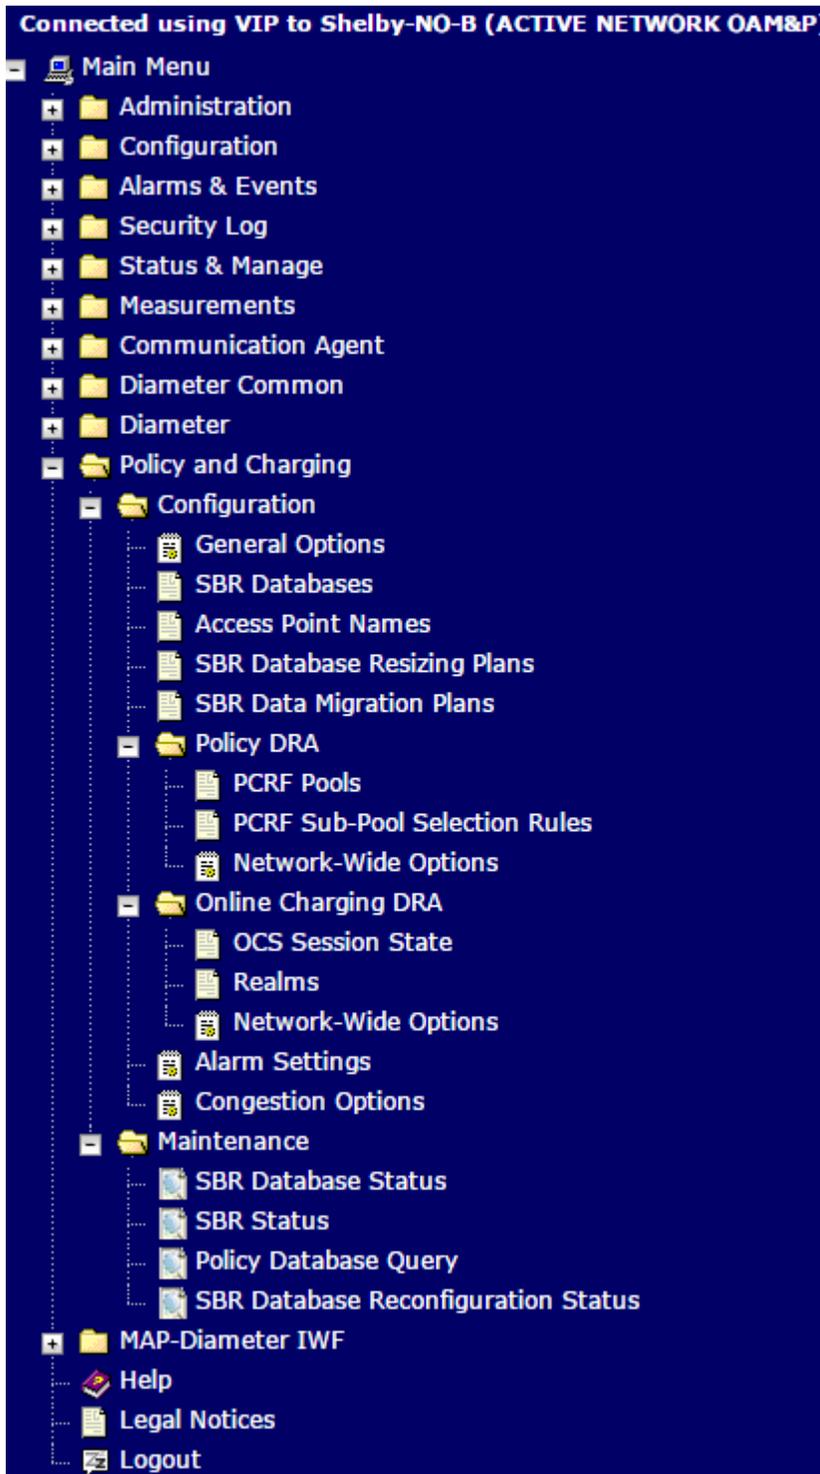


Figure 48: GUI Structure for 3-tiered DSR Topology with Policy and Charging for SOAM

NOAM and SOAM Configuration

Configuration data is divided into two categories depending on the scope of the data:

- Network-wide data is configured at the NOAM and is called A-scope data.
- Per-site data is configured at the SOAM for a given site and is called B-scope data.

In general, topology data like creation of sites, assignment of servers to sites, creation of server groups, and so on is A-scope data. DSR data configuration is generally site-scoped, or B-scope data.

Some Policy and Charging Application data must be configured at the A-scope level and some data must be configured at the B-scope level.

Policy related data configured at the NOAM include:

- Assignment of Servers to Site Places
- Assignment of Servers to SBR Server Groups
- Assignment of SBR Server Groups to Session and/or Binding Resource Domains
- Assignment of DSR Multi-active Cluster Server Groups to PCA Resource Domains
- Assignment of Site Places to PCA Mated Pair Place Associations
- Assignment of Site Places to PCA Binding Region Place Associations

PCA-specific data configured at the NOAM include:

- Alarm Thresholds for:
 - PCA Application Ingress Message Rate
 - Session Database Capacity
 - Policy Binding Database Capacity
- Policy DRA and Online Charging DRA function disabling/enabling
- Default Audit options
- OCS/Realm Session State
- Access Point Names (APN)
- Maximum Session Inactivity Time per APN
- PCRF Pools and PCRF Sub-Pool Selection Rules
- SBR Databases
- SBR Database Resizing Plans
- SBR Data Migration Plans

PCA-specific data configured at the SOAM include:

- PCRFs adjacent to the site
- Binding Key Priority for the site
- Topology Hiding configuration for the site
- OCSs and CTFs for OC-DRA
- Error response configuration for the site
- Congestion handling options
- Suspect Binding Removal Rules

For more information, see [PCA Capacity Constraints](#).

Pre-Configuration Activities

Before PCA configuration can be performed, the following activities need to be performed in the system:

- Verify that the PCA application is activated in the system. (This is usually performed as part of the installation or upgrade activities.)

Policy and Charging appears in the left-hand GUI menu on the NOAM and the SOAM after the application is activated.

- Verify that the following NOAM configuration is complete for PCA:

- Places

Select **Configuration > Places**.

Click **Report** to generate a report about the configured Places.

Click **Print** to print the report, or **Save** to save the report as a text file.

- Place Associations

Select **Configuration > Place Associations**.

Click **Report** to generate a report about the configured Place Associations

Click **Print** to print the report, or **Save** to save the report as a text file.

- Resource Domains

Select **Configuration > Resource Domains**.

Click **Report** to generate a report about the configured Resource Domains

Click **Print** to print the report, or **Save** to save the report as a text file.

Note: A Resource Domain cannot be deleted that is part of a Binding or Session profile, unless the PCA is deactivated. Resource Domains that are part of PCA profiles can be deleted when the PCA application is activated.

- Gather component information that is required for Diameter, Diameter Common, and PCA configuration, including component item naming conventions and names, IP addresses, hostnames, and numbers of items to be configured.

Naming Conventions and Hierarchical Routing illustrates the use of a naming convention.

- Configure Diameter Common components that are required for PCA configuration. See *Diameter Common Configuration for PCA* for PCA configuration information.
- Configure Diameter Configuration components that are required for PCA configuration. See *Diameter Configuration for PCA*.

Initial Installation for PCRF Pooling

Note: PCRF Pools and PCRF Sub-Pool Selection Rules are only configured at the NOAM.

When a DSR release, including PCRF Pooling is initially installed (not upgraded from a previous release that did not include PCRF Pooling) and Policy DRA is activated, PCRF Pooling is enabled by default.

Note: Use the explanations and procedures in the Diameter Configuration online help and the *Diameter User Guide* to complete the configuration of the Diameter Configuration components for the system.

The following must be performed prior to using the software for policy signaling:

1. Diameter must be configured according to the appropriate release documentation.
2. Policy and Charging Application feature must be activated.
3. Policy DRA must be enabled.
4. PCRF Pooling must be configured; consider the following:
 - The PCRF Pooling capability is enabled by default and cannot be disabled.
 - A Default PCRF Pool is pre-configured and cannot be deleted. This PCRF Pool can be used or not used, similar to the Default PRT table.
 - The Default PCRF Pool is not mapped to a PRT table by default. The PCRF Pool to PRT Mapping table uses the Not Selected choice for PRT by default.
 - When Access Point Names are configured, they must be mapped to a configured PCRF Pool.

Consider the following for PCRF Pooling function:

- The PCRF Pooling capability is enabled by default for initial installations, and it cannot be disabled.
- A Default PCRF Pool is pre-configured, and it cannot be deleted. This PCRF Pool can be used or not used, similar to the Default PRT table.
- The Default PCRF Pool is not mapped to a PRT table by default. The PCRF Pool to PRT Mapping table uses the Not Selected choice for PRT (this is the default).

If PCA is activated on a DSR that was upgraded to a release that supports PCRF Pooling and the PCA activation occurs after the upgrade is completed and accepted, the considerations listed above apply to the initial install. Activation of PCA on a network where the upgrade is not completed and accepted on all servers is prohibited by the activation script.

Upgrading for PCRF Pooling

This section discusses upgrade of a Policy DRA network from a release that supports Policy DRA, but does not support PCRF Pooling (for example, DSR 4.1.5 or DSR 5.0), to a release that does support PCRF Pooling. This information is limited to cases in which the Policy DRA feature has been activated prior to upgrading to a release that supports PCRF Pooling.

A graceful upgrade is supported when moving to a DSR release that includes PCRF Pooling.

A graceful update ensures the following:

- Existing bindings are not adversely affected by the in-service upgrade.
- Policy DRA business logic continues to execute the previous release logic until PCRF Pooling is explicitly enabled.
- Split bindings are not created.
- PCRF Pooling configuration can be performed before or after enabling PCRF Pooling with no unexpected behavior.

- If PCRF Pooling is enabled with no configuration changes, the Policy DRA behavior will be the same as prior to PCRF Pooling being enabled (assuming that all APNs were already configured).
- If PCRF Pooling is configured prior to enabling PCRF Pooling, existing bindings are honored until they are released normally. Only new bindings are routed according to the PCRF Pooling behavior.

Configuration Changes Caused by Upgrading to PCRF Pooling

After upgrading to a DSR release that supports PCRF Pooling, but prior to enabling the PCRF Pooling functionality, the following changes to Policy DRA configuration must be complete:

- All configured APNs are mapped to the Default PCRF Pool. This can be seen on the NOAMP GUI at Policy DRA à Configuration à Access Point Names. START HERE section 2.5.1 in 7930

Diameter Common Configuration for PCA

The following Diameter Common configuration must be done before PCA configuration can be performed.

Use the explanations and procedures in the Diameter Common configuration help and the *Diameter Common User's Guide* to complete the Diameter Common configuration, including the Diameter Common components needed for use with PCA.

SOAM Diameter Common Configuration

Diameter Common configuration for MP Profile assignment for PCA is done from the SOAM GUI
Main Menu: Diameter Common > MPs > Profile Assignments

Select **Diameter Common > DA-MPs > Profile Assignments**, and verify that the correct Session MP Profiles have been assigned for PCA DA-MPs. If assignments need to be made or changed:

- Use the **Diameter > Configuration > DA-MPs > Profile Assignments** page to assign an **MP Profile** for each configured PCA DA-MP shown in the **DA-MP** list.
- From the pulldown list, select the MP Profile that is for the correct blade type and for a Session application (such as **G6 Session** or **G8 Session**).

Diameter Configuration for PCA

The PCA application requires configuration of several Diameter Configuration components before the PCA configuration can be performed.

All Diameter Configuration components are configured using the SOAM GUI.

Use the explanations and procedures in the Diameter Configuration online help and the *Diameter User Guide* to complete the configuration of the Diameter Configuration components for the system, including the following Diameter Configuration components for use with the PCA application.

1. Application Ids

Use the **Diameter > Configuration > Application Ids [Insert]** page to define an Application Id for each Diameter interface that will be used by PCA in the system.

PCA supports the following values that can be selected in the **Application Id Value** pulldown list:

- 4 - Diameter Credit Control
- 16777236 – 3GPP Rx
- 16777238 – 3GPP Gx
- 16777238 – 3GPP Gx-Prime
- 16777266 – 3GPP Gxx
- 16777267 – 3GPP S9
- 4294967295 – Relay

Note: Gx-Prime shares the same Application Id as Gx. To distinguish between them, the content of the Diameter message is checked against a configured Application Routing Table to determine if the message originates from a Gx or Gx-Prime interface.

PCA always attempts to route using Peer Route Tables. The Peer Route Table can be configured here for each Application Id, or can be configured for Peer Nodes. If neither is configured, the Default Peer Route Table will be used. See [PCA Routing of Diameter Messages](#).

2. CEX Parameters

Use the **Diameter > Configuration > CEX Parameters [Insert]** page to define the Capability Exchange parameters for each Application Id that was configured for use by PCA:

For each Application Id, select or enter:

- **Application Id Type** – Authentication
- **Vendor Specific Application Id**, if the Application Id and Vendor Id will be grouped in a Vendor-specific Application Id AVP
- **Vendor Id** – if **Vendor Specific Application Id** is selected

The Vendor ID 10415 is defined in 3GPP as follows:

- Gx: 16777238 with Vendor-Id of 10415 (Defined in 3GPP 29.212)
- Gx-Prime: 16777238 with Vendor-Id of 10415 (Defined in 3GPP 29.212)
- Gxx: 16777266 with Vendor-Id of 10415 (Defined in 3GPP 29.212)
- Rx: 16777236 with Vendor-Id of 10415 (Defined in 3GPP 29.214)
- S9: 16777267 with Vendor-Id of 10415 (Defined in 3GPP 29.215)

3. CEX Configuration Sets

Use the **Diameter > Configuration > Configuration Sets > CEX Configuration Sets [Insert]** page to select the configured CEX parameters to use in:

- A CEX Configuration Set to be used for connections with the PCEF nodes (Gx)
- A CEX Configuration Set to be used for connections with the AF nodes (Rx)
- A CEX Configuration Set to be used connections with the PCRF/OCS nodes (Gx and Rx)
- CEX Configuration Sets to be used with any other types of nodes, such as BBERF (Gxx)
- A CEX Configuration Set named Default is provided for the Relay Application Id; it can be edited if needed.

4. Transaction Configuration Sets

Use the **Diameter > Configuration > Configuration Sets > Transaction Configuration Sets [Insert]** page to configure Transaction Configuration parameters.

The transaction configuration, once configured, needs to be set in the **Diameter > Configuration > Peer Nodes** page.

5. Local Nodes (PCA DA-MPs)

Use the **Diameter > Configuration > Local Nodes [Insert]** page to configure the PCA DA-MPs as Local Nodes in the system.

The pulldown list of **IP Addresses** contains the XSI addresses configured on DSR MP Servers.

6. Peer Nodes

Use the **Diameter > Configuration > Peer Nodes [Insert]** page to configure PCEFs, AFs, BBERFs, and any other types of nodes as Peer Nodes to the PCA DA-MPs in the system. (PCA DA-MPs can also be Peer Nodes to each other at different sites). Additionally, select a **Transaction Configuration Set** from the pulldown list.

See [PCA Routing of Diameter Messages](#) for details on routing of messages for PCA.

7. Connections

Use the **Diameter > Configuration > Peer Nodes [Insert]** page to configure connections between the PCA DA-MPs and the Peer Nodes.

Any IPFE Target Set Address (TSA) that is used to configure a connection must use the same **Transport Protocol** (SCTP or TCP) that is selected to configure the connection.

8. Route Groups

Use the **Diameter > Configuration > Route Groups [Insert]** page to configure Route Groups for use with PCA Peers.

For priority-based initial CCR-I routing, configure N+1 Route Groups where N is the number of PCRF/OCSs in the system. The first N Route Groups contain one corresponding PCRF/OCS Peer Node in each one, and the last Route Group contains all PCRF/OCSs.

The goal is to setup a routing configuration such that if there is no route available to the suggested PCRF/OCS in an initial (binding capable) session Request, Diameter automatically sends the Request messages to any other available PCRF/OCS.

Define a Route Group for each PCRF/OCS; enter the **Route Group Name**, select the **Peer Node** name (PCRF/OCS name) and enter the **Provisioned Capacity** as 1.

Define a last Route Group for all PCRF/OCSs; enter the **Route Group Name**, then add a **Peer Node, Connection and Capacity** entry for every PCRF/OCS. Select the **Peer Node** (PCRF/OCS) and enter the **Provisioned Capacity** as 1 for each PCRF/OCS entry.

9. Route Lists

Use the **Diameter > Configuration > Route Lists [Insert]** page to configure Route Lists for use with the configured Route Groups.

For priority-based initial session binding, configure N Route Lists where N is the number of PCRF/OCSs in the system.

All Route Lists must contain at least two Route Groups, one for a single PCRF/OCS and one for all PCRF/OCSs.

Assign **Priority** value 1 to each Route Group for a single PCRF/OCS; assign **Priority** value 2 to the Route Group containing all the PCRF/OCSs.

Enter 1 for the **Minimum Route Group Availability Weight** in all of the Route Lists.

10. Peer Route Table and Peer Routing Rules

Use the **Diameter > Configuration > Peer Route Tables [Insert]** page to configure new Peer Route Tables if needed, and the **Viewing Rules for Peer Route Table** page to configure Peer Routing Rules, such that DSR forwards messages based on the PCRF/OCS preference.

Peer Routing Rules can be added to the Default Peer Route Table (PRT) or to new Peer Route Tables.

See [PCA Routing of Diameter Messages](#) for details on PRT routing of PCA messages.

The routing configuration will ensure that whenever PCA requests Diameter to route to a particular PCRF/OCS based on the PRT:

- If the PCRF/OCS is available, Diameter will route to it.
- If the PCRF/OCS is not available, Diameter will route the message to any other available PCRF/OCS.

11. Application Route Tables and Application Routing Rules

Use the **Diameter > Configuration > Application Route Tables [Insert]** page to configure new Application Routing Rules, if needed for each Diameter interface (such as GxGx-Prime, or Rx) that is configured in an Application Name, to be used for Diameter routing of messages to the PCA application. PCA must receive all Diameter Requests.

Use the **Viewing Rules for Application Route Table** page to view existing Rule Names, configure new rules, or edit and delete existing Application Routing Rules.

Application Routing Rules can be added to the Default Application Route Table or to new Application Route Tables.

For each rule, enter or select:

- **Rule Name** for a configured Application Id (Diameter interface)
- **Priority**
- In **Conditions**, select a hyperlink to view the associated **Diameter > Configuration > Application Ids (Filtered)** page for configured for PCA.
- **Application Name - PCA**
- **Gx-Prime**
- **Application Route table**

PCA Routing of Diameter Messages

PCA routes Diameter messages depending on the following criteria:

- Answer message or Request message
- New session Request or in-session Request
- New binding or existing binding new session Request

Peer Routing

PCA always attempts to route using Peer Route Tables. The Diameter Routing Function attempts to use Peer Route Tables in the following predefined precedence:

1. Peer Route Table configured for the originating Peer Node (Diameter->Configuration->Peer Nodes)

If a match is found, the specified Peer Route Table is used.

- Peer Route Table configured for the Diameter Application-ID of the policy session initiation request being routed (Diameter->Configuration->Application Ids)

If the ingress Peer Node is configured as "Not Selected", that entry is skipped and the Application Ids configuration is checked.

- Default Peer Route Table

If no match is found in the Application-Ids configuration, the Default Peer Route Table is used.

- Destination-Host Routing

If no Peer Routing Rule matches in the Default Peer Route Table, PCA will attempt to route the Request using Destination-Host routing (for example, to a connection or Alternate Implicit Route List associated with the destination Peer Node).

Routing of Session Initiation Requests for New Bindings

PCA allows a Peer Route Table to be configured for use when a new binding is created. This Peer Route Table can specify Peer Routing Rules to:

- Allow new bindings to be routed, for example, based on the Origin-Host or Origin-Realm of the PCEF
- Cause new bindings to be load-shared across all local PCRFs.

The Peer Route Table to use for new bindings is specified in the **Policy and Charging > Configuration > Policy DRA > Site Options** GUI page on the SOAM at each site.

If the Peer Route Table for new bindings is set to "Not Selected", the Diameter Routing Function uses the precedence described in [Peer Routing](#).

Routing of Session Initiation Requests for Existing Bindings

Sessions for subscribers that are already bound to a PCRF must be routed to the bound PCRF, or to a PCRF that shares state with the bound PCRF if the PCRF supports sharing of policy state. For existing bindings, no Peer Route Table is configured in the PCA application Site Options. Instead, the Diameter Routing Function uses the precedence described in [Peer Routing](#).

Routing of Requests from PCRF to a Policy Client

In order to route Requests initiated by the PCRF, routing must be configured such that Requests from any PCRF can be routed to any Policy Client in the network. This type of routing is used to route RAR and ASR requests. For Requests from PCRFs to Policy Clients, no Peer Route Table is configured in the PCA application Site Options. Instead, the Diameter Routing Function uses the precedence described in [Peer Routing](#).

Routing of In-Session Requests

In-session Requests are Requests within a Diameter session other than the Request that established the Diameter session. CCR-U, CCR-T, and STR are all examples of in-session Requests. In-session Requests are routed using the predefined precedence of Peer Route Tables described in [Peer Routing](#).

Routing of Answer Messages

All Diameter Answer messages are routed over the same path on which the Request was routed, using hop-by-hop routing. No routing configuration is necessary to route Answer messages.

NOAM Configuration

This section describes the **Policy and Charging > Configuration** GUI pages on the NOAM.

General Options

On the **Policy and Charging > Configuration > General Options** page on an Active NOAM, the following **General Options** can be configured:

Note: **General Options** is also available to be viewed on the SOAM. However, these options are only able to be sorted and filtered on the SOAM. Modifying these options is only permissible on the NOAM.

The fields are described in [General Options elements](#).

General Options

- Indicate whether or not the Policy DRA function of PCA is enabled.
- Indicate whether or not the Online Charging DRA Function of PCA is enabled.

Audit Options

- Change the **Default Stale Session Timeout** value to a value other than the default value in the field.

This setting is a time value (in hours), after which a session is considered to be stale. For PDRA, a session is considered stale only if no RAR/RAA messages are received in longer than this configured time. For OCDRA, a session is considered stale if no any in session messages are received in longer than this configured time. If a session's age exceeds this value, that session is eligible to be audited out of the database.

This value is only used if a session is not associated with a configured Access Point Name in the Access Point Names configuration table. For sessions that are associated with a configured Access Point Name, the appropriate Stale Session value in the Access Point Name configuration table is used.

- Change the **Binding Audit Session Query Rate**, which is the maximum rate at which a binding SBR can send query messages to session servers to verify that sessions are still valid.
- Change the **Audit Operation Rate**

For session SBRs - the maximum rate at which Diameter sessions are checked for staleness.

For binding SBRs - the maximum rate at which binding session references are examined, if not already throttled by the Binding Audit Session Query Rate.

General Options elements

[Table 24: General Options elements](#) describes the elements on the **Policy and Charging > Configuration > General Options** page on the NOAM.

Table 24: General Options elements

Fields (* indicates a required field)	Description	Data Input Notes
General Options		
Policy DRA Enabled	Indicates whether the Policy DRA Function of PCA is enabled	Format: Check box Range: Enabled (Checked) or Disabled (Unchecked) Default: Disabled (Unchecked)
Online Charging DRA Enabled	Indicates whether the Online Charging DRA Function of PCA is enabled	Format: Check box Range: Enabled (Checked) or Disabled (Unchecked) Default: Disabled (Unchecked)
Audit Options		
* Default Stale Session Timeout	<p>This setting is a time value (in hours), after which a session is considered to be stale. For PDRA, a session is considered stale only if no RAR/RAA messages are received in longer than this configured time. For OCDRA, a session is considered stale if no any in session messages are received in longer than this configured time. If a session's age exceeds this value, that session is eligible to be audited out of the database.</p> <p>This value is only used if a session is not associated with a configured Access Point Name in the Access Point Names configuration table. For sessions that are associated with a configured Access Point Name, the appropriate Stale Session value in the Access Point Name configuration table is used.</p>	Format: Text box Range: 1-2400 hours (1 hour to 100 days) Default: 168 hours (7 days)
* Binding Audit Session Query Rate	The maximum rate at which a binding SBR can send query messages to session servers to verify that sessions are still valid.	Format: Text box Range: 5000-25000 records per second Default: 12000 per second
* Audit Operation Rate	For session SBRs: the maximum rate at which Diameter sessions are checked for staleness.	Format: Text box Range: 25000-50000 per second

Fields (* indicates a required field)	Description	Data Input Notes
	For binding SBRs: the maximum rate at which binding session references are examined, if not already throttled by the Binding Audit Session Query Rate.	Default: 50000 per second

Viewing General Options

Use this task to view configured General Options on the NOAM.

Select **Policy and Charging > Configuration > General Options**.

The **Policy and Charging > Configuration > General Options** page appears with a list of configured General Options.

The fields are described in [General Options elements](#).

Setting General Options

Use this task to set General Options on the NOAM.

The fields are described in [General Options elements](#).

The following general options can apply to the configuration of Policy and Charging:

- Policy DRA Enabled
- Online Charging DRA Enabled

1. Select **Policy and Charging > Configuration > General Options**.

The **General Options** page appears.

2. The **Policy DRA Enabled** check box allows the user to enable or disable Policy DRA.

3. The **Online Charging DRA Enabled** check box allows the user to enable or disable Online Charging DRA.

4. Enter a number in the **Default Stale Session Timeout** field.

5. Enter a number in the **Binding Audit Session Query Rate** field.

6. Enter a number in the **Audit Operation Rate** field

7. Click:

- **Apply** to save the changes and remain on this page.
- **Cancel** to discard changes and remain on the **Policy and Charging > Configuration > General Options** page.

If **Apply** is clicked and the following condition exists, an error message appears:

- The entered Default Stale Session Timeout value contains invalid characters, is out of the allowed range, or the field is empty.

SBR Databases

On the **Policy and Charging > Configuration > SBR Databases** page on an Active NO, an SBR Database can be configured. For each database the Resource Domain, number of Server Groups, and Place Association are defined.

Note: SBR Database is also available to be viewed on the SO GUI. However, these options are only able to be sorted and filtered on the SO. Modifying these options is only permissible on the NO.

The fields are described in [SBR Databases elements](#).

SBR Databases elements

[Table 25: SBR Databases elements](#) describes the elements on the **Policy and Charging > Configuration > SBR Databases** page on the NO GUI.

Table 25: SBR Databases elements

Fields (* indicates a required field)	Description	Data Input Notes
Database Name*	A name that uniquely identifies the SBR Database.	Format: Text box Default: N/A Range: A 32-character string. Valid characters are alphanumeric and underscore. Must contain at least one alpha and must not start with a digit.
Database Type*	The type of SBR Database.	Format: Pull-down list Default: N/A Range: Binding or Session
Resource Domain*	The Policy Session or Policy Binding Resource Domain that contains the SBR Server Groups configured for use by this database. Select the Resource Domain that will host this database.	Format: Pull-down list Default: N/A Range: Configured Resource Domains matching the selected Database Type that have not already been assigned to a Database.
Number of Server Groups*	The number of SBR Server Groups required to host this database. Enter or change the number of Server Groups necessary to support the desired capacity of the database. If the selected Resource Domain already contains Server Groups, the number of Server Groups in the	Format: Text box Default: N/A Range: 1 to 8

Fields (* indicates a required field)	Description	Data Input Notes
	Resource Domain is displayed in the field, but can be overridden as desired.	
Place Association*	The Policy Binding Region or Policy and Charging Mated Sites Place Association that contains the Sites that will use this database. Select the Place Association that is to use this SBR Database.	Format: Pull-down list Range: Configured Place Associations matching the selected Database Type that have not already been assigned to a Database

Viewing SBR Databases

Use this task to view all configured SBR Databases on the NO GUI.

Select **Policy and Charging > Configuration > SBR Databases**.

The **Policy and Charging > Configuration > SBR Databases** page appears.

The fields are described in [SBR Databases elements](#).

Inserting SBR Databases

Use this task to insert an SBR Databases on the NO GUI.

The fields are described in [SBR Databases elements](#).

1. Select **Policy and Charging > Configuration > SBR Databases**.
The **Policy and Charging > Configuration > SBR Databases** page appears.
2. Click **Insert**.
3. Enter a unique Database Name in the **Database Name** field.
4. Select a **Database Type** from the pull-down list.
5. Select a **Resource Domain** from the pull-down list.
6. Enter a number of Server Groups in the **Number of Server Groups** field.
7. Select a **Place Association** from the pull-down list.
8. Click:
 - **Ok** to save the new SBR Database and return to the **Policy and Charging > Configuration > SBR Databases** page.
 - **Apply** to save the new SBR Database and remain on the page.
 - **Cancel** to return to the **Policy and Charging > Configuration > SBR Databases** page without saving any changes.

Note: If **Ok** or **Apply** is clicked and the maximum number of SBR Databases has already been created, an error message will appear.

Deleting an SBR Database

Use this task to delete an SBR Database.

1. Select **Policy and Charging > Configuration > SBR Databases**.
The **Policy and Charging > Configuration > SBR Databases** page appears.
2. Select the **SBR Database** to be deleted.
3. Click the **Delete** button.
A popup window appears to confirm the delete.
4. Click:
 - **Ok** to delete the SBR Database.
 - **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > SBR Databases** page.

If **Ok** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entry no longer exist.
- A specific SBR Database must be Disabled prior to deletion.
- A specific SBR Database is referenced by at least one SBR Database Resizing Plan.
- A specific SBR Database is reference by at least one SBR Data Migration Plan.

Access Point Names

An Access Point Name (APN) is a unique Packet Data network identifier. The PCA uses configured Access Point Names to validate APN entries received in Diameter signaling, and to apply appropriate Stale Session Timeout values during database audits.

PCRF pool selection allows the APN used by the UE to connect to the network is used to determine the PCRF pool. This allows multiple bindings to exist for a single IMSI, one for each PCRF pool. The Origin-Host of the PCEF sending the CCR-I can then be used to select a PCRF sub-pool. Each APN is mapped to a PCRF Pool designated to manage policy bindings originated from that APN. In addition, a stale session timeout is assigned to the APN to control how long a session from the APN can remain idle before being subject to audit.

When an APN entry is added, new bindings from that APN are routed to a PCRF in the specified PCRF Pool (or a Sub-Pool if a matching PCRF Sub-Pool Selection Rule also exists). When an APN is mapped to a PCRF Pool using the Access Point Names GUI, a check is performed to determine if the selected PCRF Pool is configured with a PRT mapping at each site. If at least one site does not have a mapping for the selected PCRF Pool, a confirmation dialog displays a warning as follows:

- If a PCRF Pool is not mapped to a PRT table for a site, a confirmation dialog is displayed on the APN GUI warning that Site X does not have a mapping defined for this PCRF Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN that is mapped to that PCRF Pool.
- If a site cannot be reached due to network errors, a confirmation dialog is displayed on the APN GUI warning that it cannot be determined whether Site X has a mapping defined for this PCRF Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN that is mapped to that PCRF Pool.

Single PCRF pool support is achieved by using the default pool, with all APNs mapped to that pool. This results in all bindings pointing to a single PCRF Pool.

If an APN is successfully deleted from the NOAMP GUI, the entry is internally marked as retired. Retired entries are not displayed on the GUI, but cannot be removed from the internal tables because that APN could still be referenced by any number of bindings. If you add a new APN with the same

name as one that has been retired, the record comes out of retirement, but with the PCRF Pool and Stale Session Lifetime configured when the record was re-added.

The fields are described in [Access Point Names elements](#).

On the **Policy and Charging > Configuration > Access Point Names** page on the Active NOAM, you can perform the following actions:

- Filter the list of Access Point Names, to display only the desired Access Point Names.
- Sort the list entries in ascending or descending order by Access Point Names or by Stale Session Timeout, by clicking the column heading. By default, the list is sorted by Access Point Names in ascending numerical order.
- Work with PCRF Pool Names and Sub-Pools for PDRA APNs.
- Click the **Insert** button.

The **Policy and Charging > Configuration > Access Point Names [Insert]** page opens. You can add a new Access Point Name. See [Inserting Access Point Names](#). If the maximum number of Access Point Names (2500) already exists in the system, the **Policy and Charging > Configuration > Access Point Names [Insert]** page will not open, and an error message is displayed.

- Select an Access Point Name in the list, and click the **Edit** button.

The **Policy and Charging > Configuration > Access Point Names [Edit]** page opens. You can edit the selected Access Point Name. See [Editing Access Point Names](#).

- Select an Access Point Name in the list, and click the **Delete** button to remove the selected Access Point Name. See [Deleting an Access Point Name](#).

On the **Policy and Charging > Configuration > Access Point Names** page on the SOAM, you can view the configured Access Point Names, and perform the following actions:

- Filter the list of Access Point Names, to display only the desired Access Point Names.
- Sort the list entries in ascending or descending order by Access Point Names or by Stale Session Timeout, by clicking the column heading. By default, the list is sorted by Access Point Names in ascending numerical order.

Access Point Names elements

[Table 26: Access Point Names elements](#) describes the elements on the **Policy and Charging > Configuration > Access Point Names** page.

Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 26: Access Point Names elements

Elements (* indicates required field)	Description	Data Input Notes
* Access Point Name	The unique network identifier of a Packet Data Access Point.	Format: Text box; valid characters are alphabetic characters (A-Z and a-z), digits (0-9), hyphen (-), and period (.). Must begin and end with an

Policy and Charging Configuration

Elements (* indicates required field)	Description	Data Input Notes
		alphabetic character or a digit. Default: N/A Range: 1 to 100
Function	The PCA function what uses this Access Point. PCRF Pool is required to be configured for PDRA only.	Format: Dropdown menu Range: PDRA Only, OCDRA Only, and PDRA and OCDRA Default: PDRA Only
PCRF Pool Name	The PCRF Pool associated with the Access Point Name. PCRF Pool Names in the row are hyperlinks to the Policy DRA -> Configuration -> PCRF Pools (Filtered) view screen filtered by the PCRF Pool Name.	Format: List Range: Configured PCRF Pools Default: Default PCRF Pool
Number of Sub-Pools	This read-only field displays the number of Sub-Pools within the corresponding PCRF Pool Name. The mapping between PCRF Pool and PCRF Sub-Pool is configured from the Policy DRA -> Configuration -> PCRF Sub-Pool Selection Rules page. If the value is not zero, each Sub-Pool in the row is a hyperlink to the Policy DRA -> Configuration -> PCRF Sub-Pools Selection Rules (Filtered) view screen filtered by the PCRF Sub-Pool Selection Rule. If the number of Sub-Pools is zero, this is not a hyperlink field.	Format: List Range: N/A
Stale Session Timeout (Hrs)	This setting is a time value (in hours), after which a session is considered to be stale. For PDRA, a session is considered stale only if no RAR/RAA messages are received in longer than this configured time. For OCDRA, a session is considered stale if no any in session messages are received in longer than this configured time. If a session's age exceeds this value, that session is eligible to be audited out of the database. This value is used for sessions associated with this Access Point Name. For sessions which are not associated with any configured Access Point Names, the Default Stale Session Timeout value	Format: Text box. Value must be numeric. Range: 1-2400 (1 hour to 100 days) Default: 168 hours (7 days)

Elements (* indicates required field)	Description	Data Input Notes
	in the Policy and Charging Configuration General Options table is used.	
Last Updated	<p>This read-only field displays a timestamp of the time the Access Point Name was created or last updated, whichever occurred most recently.</p> <p>For APNs that existed prior to the upgrade to PCRF Pooling, the Last Updated timestamp reflects the time of the upgrade of the NOAMP, or the last time the APN's PCRF Pool was updated via Edit.</p> <p>For APNs added after the upgrade to PCRF Pooling, the Last Updated timestamp reflects the time when the APN was inserted, or the last time the APN's PCRF Pool was updated via Edit.</p>	<p>Format: List</p> <p>Range: N/A</p>

Viewing Access Point Names

Use this task to view all configured Access Point Names on the NOAM or SOAM.

Select **Policy DRA > Configuration > Access Point Names**.

The **Policy DRA > Configuration > Access Point Names** page appears with a list of configured Access Point Names.

The fields are described in [Access Point Names elements](#).

Inserting Access Point Names

Use this task to insert Access Point Names.

Note: Access Point Names are configurable only on Active NOAM servers, and are viewable on NOAM and SOAM servers.

The fields are described in [Access Point Names elements](#).

1. Select **Policy and Charging > Configuration > Access Point Names**.
The **Policy and Charging > Configuration > Access Point Names** page appears.
2. Click **Insert**.
The **Policy and Charging > Configuration > Access Point Names [Insert]** page appears.
3. Enter a unique Access Point Name in the Access Point Name **Value** field.
4. Select the **Function**.
5. Select a PCRF Pool Name from the **PCRF Pool Name** dropdown menu. This field contains all the qualified PCRF Pools configured from **Policy and Charging > Configuration > Policy DRA > PCRF Pools**. A qualified PCRF Pool is non-retired and has not been marked as Sub-Pool.

Note: This step is only valid for **PDRA Only** or **PDRA and OCDRA**.

This identifies the PCRF Pool to which new bindings initiated from the Access Point Network are to be routed.

Note: A retired PCRF Pool entry can be created by first adding a new PCRF Pool and then deleting it.

The Number of Sub-Pools field is a read-only field that displays the number of PCRF Sub-Pools associated with the selected PCRF Pool. The mapping between PCRF Pool and PCRF Sub-Pool is configured from the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page.

6. If a value other than the default Stale Session Timeout value is desired, enter the desired length of time in hours in the Stale Session Timeout (Hrs) **Value** field.

For sessions that are not associated with any configured Access Point Names, the default Stale Session Timeout value in the **Policy and Charging > Configuration > Policy DRA > Network-Wide Options** table is used. The default is 168 hours (7 days), and the range is 1-2400 hours (1 hour to 100 days).

The Last Updated field is a read-only field that displays the date and time that this APN was created, or the last time the PCRF Pool Name was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding-capable session initiation requests. You can compare this date and time to the binding creation times when troubleshooting using the Binding Key Query Tool.

7. Click:

- **OK** to save the new Access Point Name and return to the **Policy and Charging > Configuration > Access Point Names** page.
- **Apply** to save the new Access Point Name and remain on this page.
- **Cancel** to return to the **Policy and Charging > Configuration > Access Point Names** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entered Access Point Name is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range
- Any required field is empty (not entered)
- Adding the new Access Point Name would cause the maximum number of Access Point Names (2500) to be exceeded

Editing Access Point Names

Use this task to edit Access Point Stale Session Timeout values.

Note: The Access Point Name **Value** cannot be edited.

Note: Access Point Names are configurable only on Active NOAM servers, and are viewable on NOAM and SOAM servers.

The fields are described in [Access Point Names elements](#).

1. Select **Policy and Charging > Configuration > Access Point Names**.

The **Policy and Charging Configuration Access Point Names** page appears.

2. Click **Edit**.

The **Policy and Charging > Configuration > Access Point Names [Edit]** page appears.

3. Select the **Function**.
4. Select a PCRF Pool Name from the **PCRF Pool Name** dropdown menu. This is the PCRF Pool to which new bindings initiated from the Access Point Network are to be routed. The default is Default PCRF Pool, and the range is Configured PCRF Pools.
The Number of Sub-Pools field is a read-only field that displays the number of PCRF Sub-Pools associated with the selected PCRF Pool. The mapping between PCRF Pool and PCRF Sub-Pool is configured from the **Policy DRA -> Configuration -> PCRF Sub-Pool Selection Rules** page.
Note: This step is only valid for **PDRA Only** or **PDRA and OCDRA**.
5. Enter the desired length of time in hours in the Stale Session Timeout (Hrs) **Value** field.
For sessions that are not associated with any configured Access Point Names, the default Stale Session Timeout value in the **Policy and Charging > Configuration > Network-Wide Options** table is used. The default is 168 hours (7 days), and the range is 1-2400 hours (1 hour to 100 days).
The Last Updated field is a read-only field that displays the date and time that this APN was created, or the last time the PCRF Pool Name was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding-capable session initiation requests. You can compare this date and time to the binding creation times when troubleshooting using the Policy Database Query Tool.
6. Click:
 - **OK** to save the changes and return to the **Policy and Charging > Configuration > Access Point Names** page.
 - **Apply** to save the edited Access Point Name and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Access Point Names** page without saving any changes.If **OK** or **Apply** is clicked and the following condition exists, an error message appears:
 - The edited Access Point Name no longer exists (for example, it has been deleted by another user), and no changes are made to the database.

Deleting an Access Point Name

Use this task to delete an Access Point Name.

Note: Access Point Names are configurable only on Active NOAM servers, and are viewable on NOAM and SOAM servers.

1. Select **Policy and Charging > Configuration > Access Point Names**.
The **Policy and Charging > Configuration > Access Point Names** page appears.
2. Select the **Access Point Name** to be deleted.
3. Click the **Delete** button.
A popup window appears to confirm the delete.
4. Click:
 - **OK** to delete the Access Point Name.
 - Click **Cancel** to cancel the delete function and return to the **Policy DRA > Configuration > Access Point Names** page.

If **OK** is clicked and the selected Access Point Name no longer exists (it was deleted by another user), an error message is displayed. The Access Point Names view is refreshed and the deleted Access Point Name no longer appears on the page.

SBR Database Resizing Plans

On the **Policy and Charging > Configuration > SBR Database Resizing Plans** page on an Active NO, SBR Database Resizing Plans can be configured. An SBR Database Resizing plan allows for changing the size of a running SBR Database.

Each database is to be resized by migrating records from the Initial Resource Domain to the Target Resource Domain. Creation of an SBR Database Resizing Plan has no effect on the running database. After creating a Resizing Plan, go to the **SBR Database Reconfiguration Status** screen to prepare, start and monitor the migration. Once migration has completed, or has been cancelled, the SBR Database Resizing Plan should be deleted.

Note: SBR Data Migration Plans are configurable only on an Active NO, but are viewable on NO and SO servers.

Resizing Plans can be configured prior to the start of the plan's execution.

The fields are described in *SBR Database Resizing Plans elements*.

SBR Database Resizing Plans elements

Table 27: SBR Database Resizing Plans elements describes the elements on the **Policy and Charging > Configuration > SBR Database Resizing Plans** page on the NO GUI.

Table 27: SBR Database Resizing Plans elements

Fields (* indicates a required field)	Description	Data Input Notes
Plan Name*	A name that uniquely identifies the SBR Database Resizing Plan	Format: Text box Default: N/A Range: A 32-character string. Valid characters are alphanumeric and underscore. Must contain at least one alpha and must not start with a digit.
Database Name*	The SBR Database to be resized	Format: Pull-down list Default: N/A Range: List of configured Databases eligible to be resized
Initial Resource Domain	The Resource Domain configured for the selected database.	This read only field displays the Resource Domain associated with the selected database.

Fields (* indicates a required field)	Description	Data Input Notes
Database Type	The type of SBR Database	This read-only field displays the Database Type associated with the selected database.
Initial Number of Server Groups	The number of Server Groups hosting this database.	This read only field displays the number of Server Groups contained in the Initial Resource Domain.
Target Resource Domain*	The Resource Domain into which the database is to be migrated. Select the Resource Domain that will host the resized database	Format: Pull-down list Default: N/A Range: Configured Resource Domains eligible for use as a target resource domain for this resizing plan
Target Number of Server Groups*	The total number of Policy SBR Server Groups required to support the migrated or resized database. Enter the total number of Server Groups necessary to support the capacity of the resized database, including any server groups that will be reused from the original resource domain.	Format: Text box Default: N/A Range:1 to 8

Viewing SBR Database Resizing Plans

Use this task to view all configured SBR Database Resizing Plans on the NO or SO GUIs.

Select **Policy and Charging > Configuration > SBR Database Resizing Plans**.

The **Policy and Charging > Configuration > SBR Database Resizing Plans** page appears.

The fields are described in [SBR Database Resizing Plans elements](#).

Inserting SBR Database Resizing Plans

Use this task to insert an SBR Database Resizing Plans on the NO GUI.

The fields are described in [SBR Database Resizing Plans elements](#).

1. Select **Policy and Charging > Configuration > SBR Database Resizing Plans**.
The **Policy and Charging > Configuration > SBR Database Resizing Plans** page appears.
2. Click **Insert**.
3. Enter a unique Plan Name in the **Plan Name** field.
4. Select a **Database Name** from the pull-down list.
5. The **Initial Resource Domain** field is read only.
6. The **Database Type** field is read only.

7. The **Initial Number of Server Groups** field is read only.
8. Select a **Target Resource Domain** from the pull-down list.
9. The **Target Number of Server Groups** field initially displays the number of Server Groups configured for the selected **Target Resource Domain**. However, this value may be modified to any value from 1 to 8. Refer to [Editing SBR Database Resizing Plans](#) for information on modifying the **Target Number of Server Groups** value.
10. Click:
 - **Ok** to save the new SBR Database Resizing Plan and return to the **Policy and Charging > Configuration > SBR Database Resizing Plans** page.
 - **Apply** to save the new SBR Database Resizing Plan and remain on the page.
 - **Cancel** to return to the **Policy and Charging > Configuration > SBR Database Resizing Plans** page without saving any changes.

If **Ok** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Field value is missing
- Syntax is invalid
- A Field value isn't unique
- The operation failed because the selection for a field is no longer valid
- An Insert failed because the maximum number of SBR Database Resizing Plans have already been created
- An Insert failed because the SBR Database selected for resizing is not an active database
- An Insert failed because the SBR Database selected for resizing is already configured in another SBR Database resizing Plan
- An Insert failed because the SBR Database selected for resizing is already configured in another SBR Data Migration Plan
- An Operation failed because the Target Resource Domain selected is already configured as the Resource Domain for another SBR Database
- An Operation failed because the Target Resource Domain selected is already configured as the Target Resource Domain for another SBR Resizing Plan
- An Operation failed because the Target Resource Domain selected must have the same Resource Domain Profile as the Resource Domain configured for the selected SBR Database

Editing SBR Database Resizing Plans

Use this task to edit an SBR Database Resizing Plans on the NO GUI.

The fields are described in [SBR Database Resizing Plans elements](#).

1. Select **Policy and Charging > Configuration > SBR Database Resizing Plans**.
The **Policy and Charging > Configuration > SBR Database Resizing Plans** page appears.
2. Click **Edit**.
3. Edit the Plan Name in the **Plan Name** field.
4. Select a **Database Name** from the pull-down list.
5. The **Initial Resource Domain** field is read only.
6. The **Database Type** field is read only.
7. The **Initial Number of Server Groups** field is read only..
8. Select a **Target Resource Domain** from the pull-down list.

9. If needed, edit the total number of Server Groups needed to support the capacity of the resized database in the **Target Number of Server Groups** field.

10. Click:

- **Ok** to save the edited SBR Database Resizing Plan and return to the **Policy and Charging > Configuration > SBR Database Resizing Plans** page.
- **Apply** to save the edited SBR Database Resizing Plan and remain on the page.
- **Cancel** to return to the **Policy and Charging > Configuration > SBR Database Resizing Plans** page without saving any changes.

If **Ok** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Field value is missing
- Syntax is invalid
- A Field value isn't unique
- The operation failed because the selection for a field is no longer valid
- An Edit failed because the maximum number of SBR Database Resizing Plans have already been created
- An Edit failed because the SBR Database selected for resizing is not an active database
- An Edit failed because the SBR Database selected for resizing is already configured in SBR Database resizing Plan <Plan Name>
- An Edit failed because the SBR Database selected for resizing is already configured in SBR Data Migration Plan <Plan Name>
- An Operation failed because the Target Resource Domain selected is already configured as the Resource Domain for SBR Database <SBR Database Name>
- An Operation failed because the Target Resource Domain selected is already configured as the Target Resource Domain for SBR Resizing Plan <Plan Name>
- An Operation failed because the Target Resource Domain selected must have the same Resource Domain Profile as the Resource Domain configured for the selected SBR Database

Deleting an SBR Database Resizing Plan

Use this task to delete an SBR Database Resizing Plan.

1. Select **Policy and Charging > Configuration > SBR Database Resizing Plans**.
The **Policy and Charging > Configuration > SBR Database Resizing Plans** page appears.
2. Select the **SBR Database Resizing Plan** to be deleted.
3. Click the **Delete** button.
A popup window appears to confirm the delete.
4. Click:
 - **Ok** to delete the SBR Database Resizing Plan.
 - **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > SBR Database Resizing Plans** page.

If **Ok** is clicked, but the SBR Database Resizing Plan is not in the Planned administrative state or in the Complete administrative state with an operational status of Migration Plan Completed, an error is displayed.

SBR Data Migration Plans

On the **Policy and Charging > Configuration > SBR Data Migration Plans** page on an Active NO, an SBR Data Migration Plan can be configured.

SBR Data Migration is the process of moving data from an Initial SBR Database to a Target SBR Database by creating new records in the target and allowing records in the initial to be removed by normal PCA Diameter signaling.

Creation of an SBR Data Migration Plan has no effect on either database. After creating a Data Migration Plan, go to the SBR Database Reconfiguration Status screen to prepare, start and monitor the migration. Once migration has completed, or has been cancelled, the SBR Data Migration Plan should be deleted.

Note: SBR Data Migration Plans are configurable only on an Active NO, but are viewable on NO and SO servers.

Data Migration Plans can be configured prior to the start of the plan's execution.

The fields are described in [SBR Data Migration Plans elements](#).

SBR Data Migration Plans elements

Table 28: SBR Data Migration Plans elements describes the elements on the **Policy and Charging > Configuration > SBR Data Migration Plans** page on the NO GUI.

Table 28: SBR Data Migration Plans elements

Fields (* indicates a required field)	Description	Data Input Notes
Plan Name*	A name that uniquely identifies the SBR Data Migration Plan.	Format: Text box Default: N/A Range: A 32-character string. Valid characters are alphanumeric and underscore. Must contain at least one alpha and must not start with a digit.
Migration Type*	The type of data migration plan that this plan is to be a part of.	Format: Pull-down list Default: N/A Range: Un-Mating, Mating, Re-Mating
Initial Database*	A name that uniquely identifies the database from which records are to be migrated.	Format: Pull-down list Default: N/A Range: List of configured Databases eligible to be the source of a migration

Fields (* indicates a required field)	Description	Data Input Notes
Initial Place Association	The Place Association configured to use the Initial Database.	This read only field displays the Place Association configured for the Target Database.
Target Database*	A name that uniquely identifies the database to which records are to be migrated.	Format: Pull-down list Default: N/A Range: List of configured Databases eligible to be the target of a migration
Target Place Association	The Place Association configured to use the Target database.	This read only field displays the Place Association configured for the Target Database.

Viewing SBR Data Migration Plans

Use this task to view all configured SBR Data Migration Plans on the NO or SO GUIs.

Select **Policy and Charging > Configuration > SBR Data Migration Plans**.

The **Policy and Charging > Configuration > SBR Data Migration Plans** page appears.

The fields are described in [SBR Data Migration Plans elements](#).

Inserting SBR Data Migration Plans

Use this task to insert an SBR Data Migration Plans on the NO GUI.

The fields are described in [SBR Data Migration Plans elements](#).

1. Select **Policy and Charging > Configuration > SBR Data Migration Plans**.
The **Policy and Charging > Configuration > SBR Data Migration Plans** page appears.
2. Click **Insert**.
3. Enter a unique Plan Name in the **Plan Name** field.
4. Select a **Migration Type** from the pull-down list.
5. Select an **Initial Database** from the pull-down list.
6. The **Initial Place Association** field is read only.
7. Select a **Target Database** from the pull-down list.
8. The **Target Place Association** field is read only.
9. Click:
 - **Ok** to save the new SBR Data Migration Plan and return to the **Policy and Charging > Configuration > SBR Data Migration Plans** page.
 - **Apply** to save the new SBR Data Migration Plan and remain on the page.
 - **Cancel** to return to the **Policy and Charging > Configuration > SBR Data Migration Plans** page without saving any changes.

If **Ok** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- A Field value is missing
- Syntax is invalid
- A Field value must be unique
- An Operation failed because the selection for the field is invalid
- An Insert failed because the maximum number of SBR Data Migration Plans have already been created
- An Insert failed because the SBR Database selected as the Initial Database is not an active database
- An Insert failed because the SBR Database selected as the Initial Database is already configured in a specific SBR Database Resizing Plan
- An Insert failed because the SBR Database selected and Target Database must have the same Database Type
- An Insert failed because a specific SBR Data Migration Plan already exists with the same Initial and Target Databases
- An Insert failed because SBR Data Migration is not supported for Binding Databases
- An Insert failed because a new SBR Data Migration Plan conflicts with an existing SBR Data Migration Plan because the Initial Database of the new Plan is configured as the Target Database of the existing plan
- An Insert failed because a new SBR Data Migration Plan conflicts with the SBR Data Migration Plan because the Target Database of the new Plan is configured as the Initial Database of the existing plan

Deleting an SBR Data Migration Plan

Use this task to delete an SBR Data Migration Plan.

1. Select **Policy and Charging > Configuration > SBR Data Migration Plans**.
The **Policy and Charging > Configuration > SBR Data Migration Plans** page appears.
2. Select the **SBR Data Migration Plan** to be deleted.
3. Click the **Delete** button.
A popup window appears to confirm the delete.
4. Click:
 - **Ok** to delete the SBR Data Migration Plan.
 - **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > SBR Data Migration Plans** page.

If **Ok** is clicked, but the SBR Data Migration Plan is not in the Planned administrative state or in the Complete administrative state with an operation status of Migration Plan Complete, an error message is displayed. The Data Migration Plan can be deleted only when the Data Migration Plan is in the Planned administrative state or is in the Completed administrative state with operational status of Migration Plan Completed.

Policy DRA Configuration

This section describes the **Policy and Charging > Configuration > Policy DRA** GUI pages on the NOAM and the SOAM.

PCRF Pools

Policy DRA continues to support a single pool of PCRFs at each PCA site over which policy Diameter signaling is distributed using the subscriber's IMSI. This allows the incorporation of new services or new PCRF infrastructure without disturbing existing services. For example, one set of PCRF servers handle policy control for all consumer data accesses to their network and a second set of PCRF servers handle all enterprise data accesses for their network. The policy rules and/or PCRF implementations might be different enough to necessitate that these two services are segregated at the PCRF level.

This means that a given IMSI might concurrently have a binding to one PCRF for APN *A* and a binding to a different PCRF for APN *B*. Each APN is mapped to a set of PCRFs; this is called a PCRF Pool. In addition, if a binding to a PCRF Pool already and a new session is created that maps to that same PCRF Pool, the request must be routed to the same PCRF. When new bindings are created for different IMSIs and a given APN, the binding-capable session initiation requests are distributed across the PCRFs in the PCRF Pool assigned to that APN.

PCRF Pooling expands this capability for the creation of multiple pools of PCRFs, which are selected using the combination of IMSI and Access Point Name (APN). This allows you to route policy Diameter signaling initiating from a given APN to a designated subset of the PCRFs that can provide specialized policy treatment using knowledge of the APN.

PCRF Pooling modifies the logic in the Policy DRA to inspect the contents of binding generating Gx CCR-I messages to select the type of PCRF to which the CCR-I messages are to be routed. In the initial P-DRA, it was assumed that all PCRFs could handle all Gx session bindings. PCRF Pooling provides service-specific sets of PCRFs. In this release, the APN used by the UE to connect to the network is used to determine the PCRF pool. The Origin-Host of the PCEF sending the CCR-I can then be used to select a PCRF sub-pool.

Multiple PCRF pools requires differentiation among the binding records in the binding SBR. It is possible for the same UE, as indicated by the IMSI, to have multiple active IPcan sessions spread across the different pools.

Note: Although the concept of a PCRF pool is a network-wide concept for a service provider, PCRF pools configuration is done on a PCA site-by-site basis. PCAs in different sites can support different PCRF Pool Selection configurations.

When deploying multiple PCRF pools, each pool supports either different policy-based services or different versions of the same policy based services. Each PCRF pool has a set of DSR PCA peers that are a part of the pool.

On the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page on the NOAM or SOAM, you can perform the following actions:

- Create new PCRF Pools
- Edit existing PCRF Pools
- Delete PCRF Pools
- Identify PCRF Sub-Pools
- Add optional comments for Pools

When a binding-capable session initiation request is received, the Policy DRA uses the following high-level logic to route the request:

- If a binding exists for the IMSI and APN or PCRF Pool, route the request to the bound PCRF.
- Otherwise, distribute the request to a PCRF in the configured PCRF Pool.

When determining if a binding exists, the following logic is used:

- If the IMSI and APN are bound to a PCRF, use that binding.
- Else, if the IMSI and PCRF Pool are bound to a PCRF, create a binding for the APN to the same PCRF as already bound to the PCRF Pool.
- Else, no binding exists for the IMSI and APN or PCRF Pool, so a new binding can be created.

The following table illustrates the major differences between PCRF Pooling and non-pooling functionality.

Table 29: PCRF Pooling Concepts

Concept	Before PCRF Pooling	After PCRF Pooling
PCRF Pools	One PCRF Pool for all APNs.	Up to 7 PCRF Pools selected for new bindings using APN. More than one APN can be mapped to a given PCRF Pool, but a given APN can only be mapped to one PCRF Pool.
Subscriber Bindings	A binding is a simple mapping between an IMSI and a PCRF. Once a binding exists, all sessions for that IMSI are routed to the bound PCRF.	A binding is a mapping from an IMSI and APN to a PCRF, but with the caveat that before a new binding is created, the logic must check for existence of another binding to the same PCRF Pool for the IMSI. If such a binding exists, the new APN is bound to the same PCRF as an existing APN mapped to the same PCRF Pool. Once a binding exists, all sessions for that IMSI and APN are routed to the bound PCRF. Sessions for that IMSI and a different APN mapped to a different PCRF Pool can be routed to a different PCRF.
Number of Sessions per Binding	An IMSI may have up to 10 binding capable sessions.	An IMSI may have up to 10 binding capable sessions, which may be bound to different PCRFs based on APN.
Origin Based Routing	PRT table for new bindings specified in Site Options allows for selection of route list based on origin-host/realm.	After PCRF Pool selection, Sub-Pool rule matching is performed to select a PCRF Sub-Pool given the PCRF Pool and the origin-host of the PCEF.
PRT Table for New Bindings	Each site defines one PRT table to be used for all new bindings.	Each site can define a PRT table to be used for new bindings for each PCRF Pool.

Additionally, Pooling provides the ability to route to subsets of PCRFs in a PCRF Pool on the basis of the Diameter hostname of the PCEF that originated the binding capable session initiation request. These subsets are called PCRF Sub-Pools. This capability allows a controlled amount of policy Diameter signaling to be routed to one or more PCRFs within the PCRF Pool.

The following figure illustrates a sample PCA network configured for PCRF Pooling. The upper third of the figure shows data that is configured with the Policy and Charging GUI at the NOAM server. This data, including PCRF Pools, APN to PCRF Pool mapping, and PCRF Sub-Pool Selection Rules applies to all sites in the Policy DRA network.

The middle third of the figure shows data configured at the SOAM Policy and Charging GUI at each of two PCA sites. This data includes the PCRF Pool to PRT mappings, PCRFs, PRT tables, Route Lists, Route Groups, Peer Nodes, and Connections. This data can differ at each PCA site.

The bottom third of the figure shows the PCRFs logically grouped into PCRF Pools as defined by the network operator.

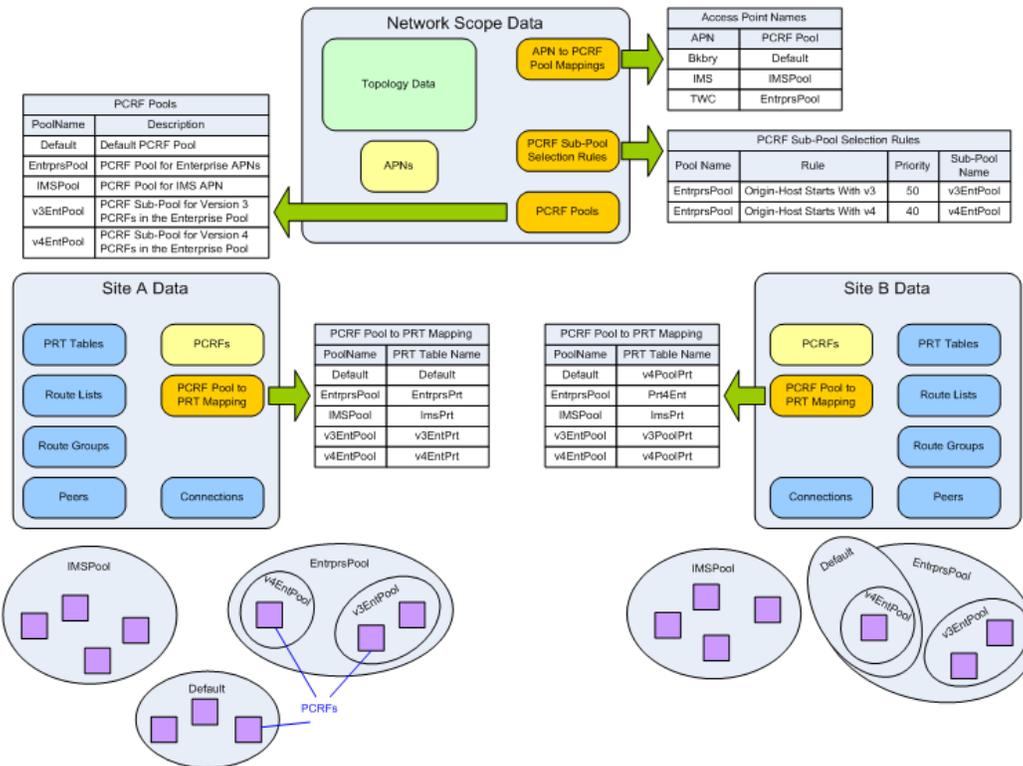


Figure 49: PCRF Pooling Data

Table 30: PCRF Pooling Configuration Summary describes each of the new PCRF Pooling configuration tables, including the order in which they should be configured.

Table 30: PCRF Pooling Configuration Summary

Configuration Order	GUI Page	Purpose
1	PCRF Pools	Define the names of the PCRF Pools and Sub-Pools that are

Policy and Charging Configuration

Configuration Order	GUI Page	Purpose
		needed for grouping PCRFs to handle policy signaling for the various APNs.
2	PCRF Pool to PRT Mapping	<p>At each site, select a PRT table that is used to route binding-capable session initiation requests for new bindings destined for each PCRF Pool. Each PCRF Pool should be configured with a PRT table, unless it is known that the PCRF Pool will never be selected at the site being configured.</p> <p>Note: Before this step can be performed, PRT tables must be defined in the Diameter folder.</p>
3	PCRF Sub-Pool Selection Rules	<p>An optional table. If it is necessary to subdivide a PCRF Pool so that policy requests from a limited number of policy clients (based on Origin-Host) are routed differently, configure appropriate rules in the PCRF Sub-Pool Selection Rules table. During routing, this table is examined after the APN is mapped to a PCRF Pool.</p> <p>If a matching PCRF Sub-Pool Selection Rule exists, the request is routed to the PCRF Sub-Pool. Otherwise, the PCRF Pool selected by the APN mapping is used.</p>
4	Access Point Names	<p>After all Diameter configuration is completed (including PRT Rules, Route Lists, Route Groups, Peer Nodes, and Connections), each APN can be mapped to a PCRF Pool. After an APN is mapped to a PCRF Pool, binding-capable session initiation requests that result in creation of a new binding are routed using the PCRF Pool.</p>

PCRF Pools elements

Table 31: PCRF Pools elements describes the elements on the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.

Note: Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

The PCRF Pools table contains the list of configured PCRF Pools and Sub-Pools settings that you can use when selecting a set of PCRFs to host a new subscriber binding. The PCRF Pool to be used for a given subscriber binding attempt is determined based on the APN-to-PCRF Pool mappings configured in **Policy and Charging > Configuration > Access Point Names** and the PCRF Sub-Pool Selection Rules configured in **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.

Table 31: PCRF Pools elements

Fields (* indicates required field)	Description	Data Input Notes
* PCRF Pool Name	A unique name for the PCRF Pool assigned by the network operator. A PCRF Pool identifies a set of PCRFs that should be used for policy requests from a specified APN. The mapping from APN-to-PCRF Pool is configured from the Policy and Charging -> Configuration -> Access Point Names page.	Format: List Range: 1 to 32 characters, must start with an upper or lower case letter, and can contain digits and underscores; a maximum of 7 PCRF Pool Names can be defined
Sub-Pool	A setting that indicates that the PCRF Pool is to be used as a PCRF Sub-Pool (for example, the target of a PCRF Sub-Pool Selection Rule). Note: If the check box on the PCRF Pools > [Insert] page is not checked, this PCRF Pool is a pool, not a sub-pool.	Format: Check box Range: Yes (Checked for Sub-Pool), No (Unchecked for Sub-Pool) Default: No (Unchecked for Sub-Pool)
Comments	An optional comment to provide more information about the purpose of this PCRF Pool or Sub-Pool.	Format: Text box Range:0-64 characters

Inserting PCRF Pools

Use this task to insert (create new) PCRF Pools.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Pools**. The **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page appears.
2. Click **Insert**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Pools [Insert]** page opens.
3. Enter a unique PCRF Pool Name in the **PCRF Pool Name** field.

4. Check the **Sub-Pool** check box if the PCRF Pool is to be used as a Sub-Pool.

A Sub-Pool is used if policy requests from specified origin-hosts should be routed to a different set of the PCRFs from those in the PCRF Pool selected by the APN. Sub-Pool Selection Rules are configured in **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.

The choices are Default = No (Unchecked for Sub-Pool); the range is Yes (Checked for Sub-Pool) and No (Unchecked for Pool).

5. You can type an optional comment in the **Comments** field to describe the Pool or Sub-Pool. The entry must be characters in the range of 0 to 64, and the default is N/A.
6. Click:
 - **OK** to save the new PCRF Pool name and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.
 - **Apply** to save the new PCRF Pool name and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entered PCRF Pool name is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- Adding the new PCRF Pool would cause the maximum number of PCRF Pools (7) to be exceeded.

Editing PCRF Pools

Use this task to edit PCRF Pools comments. After a PCRF Pool is created, only the comment can be edited, and the Sub-Pool Indicator can only be changed by deleting the PCRF Pool and creating a new one.

Note: The PCRF Pool Name cannot be edited.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Pools**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page appears. The page displays a list of the configured PCRF Pools that are used when a new subscriber binding was created.
2. Select a PCRF Pool Name to edit.
3. Click **Edit**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Pools [Edit]** page appears.
4. Click in the **Comments** field.
5. Edit the **Comments** field for the selected PCRF Pool. The comment must be characters in the range of 0 to 64, and the default is N/A.
6. Click:
 - **OK** to save the change and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.
 - **Apply** to save the change and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page without saving any changes.

If **Apply** or **OK** is clicked and the selected **PCRF Pool Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting PCRF Pools

Use this task to delete a PCRF Pool.

A PCRF Pool can be deleted only if no APN is mapped to that PCRF Pool. A PCRF Sub-Pool can be deleted only if no PCRF Sub-Pool Selection Rule refers to that PCRF Sub-Pool.

If a PCRF Pool or Sub-Pool is successfully deleted from the NOAMP GUI, the entry is internally marked as retired. Retired entries are not displayed on the GUI, but they cannot be removed from the internal tables because that PCRF Pool or Sub-Pool might still be referenced by any of number of bindings. If you add a new PCRF Pool or Sub-Pool with the same name as one that has been retired, the record is reactivated.

When a PCRF Pool or Sub-Pool is deleted (retired), the entry no longer appears on the **PCRF Pool to PRT Mapping** pages at any of the sites.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Pools**. The **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page appears.
2. Select the **PCRF Pool Name** or **PCRF Sub-Pool Name** to be deleted.
3. Click **Delete**.

A window appears to confirm the delete.

4. Click:
 - **OK** to delete the PCRF Pool or PCRF Sub-Pool.
 - **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.

If **OK** is clicked and the selected PCRF Pool or Sub-Pool no longer exists (it was deleted by another user), an error message is displayed, and the PCRF Pools page is refreshed. The row that was selected is no longer displayed in the list.

PCRF Sub-Pool Selection Rules

The PCRF Sub-Pool Selection table contains rules for selection of a PCRF Sub-Pool for a given PCRF Pool and Origin-Host value.

It is sometimes necessary to subdivide a PCRF Pool into sub-pools; for example, to support controlled routing of traffic to a new PCRF. In such a case, you can configure PCRF Sub-Pool Selection Rules to a selected a sub-pool on the basis of the Origin-Host of the binding capable session initiation request.

A PCRF Sub-Pool Selection Rule has the following attributes:

- The Default PCRF Pool can have sub-pools.
- The **PCRF Pool Name** column contains hyperlinks to the **PCRF Pools** page filtered by the PCRF Pool Name.
- Origin-Host is the only supported PCRF Sub-Pool Selection parameter.
- Supported Origin-Host operators are: Equals, Starts With, and Ends With.
- Priority values can range from 1 to 99, with 1 being the highest priority.

An APN-to-PCRF Pool mapping specifies that all binding-capable session initiation requests that result in creation of a new binding should be routed to a PCRF in PCRF Pool 'X'.

A PCRF Sub-Pool Selection Rule can override the APN-to-PCRF Pool mapping by specifying binding-capable session initiation requests that result in new bindings that were destined for PCRF Pool 'X', but come from PCEF 'Y', should be routed to a PCRF in PCRF Sub-Pool 'Z'.

A PCRF Sub-Pool Selection Rule will never be considered if no APN is mapped to its PCRF Pool. As a result, it is safe to add PCRF Sub-Pool Selection Rules prior to mapping APNs to the PCRF Pool that is being subdivided. It is also acceptable to add PCRF Sub-Pool Selection Rules for a PCRF Pool that is already mapped to an APN. However, if this is done, bindings that were created prior to the existence of the PCRF Sub-Pool Selection Rule take precedence over the PCRF Sub-Pool chosen for new binding-capable session initiation requests that arrive after the new rule is in place. This behavior is necessary to prevent split bindings.

PCRF Sub-Pool Selection Rules are configured using the NOAMP GUI as a network-wide managed object.

The creation of a new PCRF Sub-Pool Selection Rule does not affect P-DRA signaling in any way until both of the following conditions exist:

- An APN is mapped to the PCRF Pool using the Access Point Names GUI
- A binding-capable session initiation request arrives with an APN mapped to that PCRF Pool and an Origin-Host that matches the Condition specified in the PCRF Sub-Pool Selection Rule.

When a PCRF Sub-Pool Selection Rule entry is added, new bindings from that APN and Origin-Host will be routed to a PCRF in the specified PCRF Sub-Pool. When a PCRF Sub-Pool Selection Rule is mapped to a PCRF Sub-Pool, a check is performed to determine if the selected PCRF Sub-Pool is configured with a PRT mapping at each site. If at least one site does not have a mapping for the selected PCRF Sub-Pool, a confirmation dialog is displayed that including a warning as follows:

- If a site does not have the PCRF Sub-Pool mapped to a PRT table, a confirmation dialog is displayed on the APN GUI warning that Site 'X' does not have a mapping defined for this PCRF Sub-Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN and Origin-Host that is mapped to that PCRF Sub-Pool.
- If a site cannot be reached due to network errors, a confirmation dialog is displayed on to warn you that it cannot be determined whether Site 'X' has a mapping defined for this PCRF Sub-Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN and Origin-Host that is mapped to that PCRF Pool.

The PCRF Sub-Pool Selection Rule GUI prevents creation of rules that are:

- Ambiguous
- Conflicting
- Duplicate

Two rules are considered as ambiguous if the following criteria are met:

- The rules have the same PCRF Pool values and
- The rules have the same Priority values and
- The rules have different PCRF Sub-Pool values and one of the following is true:
 - One rule has an Origin-Host with a "Starts With" operator and the other rule has an Origin-Host with an "Ends With" operator -- OR --

- For example, starts With ab and Ends With xyz
- Value length is not considered as a factor in the best match decision at this time.
- Both rules have an Origin-Host with a "Starts With" operator and all of the value characters of the shorter value match the first characters of the longer value -- OR --
 - For example, starts With abc and Starts With ab
- Both rules have an Origin-Host with a "Ends With" operator and all of the value characters of the shorter value match the last characters of the longer value.
 - For examples, ends With xyz and Ends With yz

Two rules are considered to be conflicting if all of the following criteria are met:

- The rules have the same PCRF Pool values.
- The rules have the same Priority values.
- The rules have the same Origin-Host operators and values.
- The rules have different PCRF Sub-Pool values.

Two rules are considered to be duplicate if all of the following criteria are met:

- The rules have the same PCRF Pool values.
- The rules have the same Origin-Host operators and values.
- The rules have the same PCRF Sub-Pool values.

PCRF Sub-Pool Selection Rules elements

Table 32: PCRF Sub-Pool Selection Rules elements describes the elements on the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page.

Table 32: PCRF Sub-Pool Selection Rules elements

Fields (* indicates required field)	Description	Data Input Notes
PCRF Sub-Pool Selection Rule Name	A unique name for the PCRF Sub-Pool Selection Rule assigned by the network operator.	Format: Text box; string 1-32 characters, must start with an upper or lower case letter, and can contain digits and underscores; maximum number of Sub-Pool Selection Rules is of 70 Range: Valid name
Priority	A priority value. The priority value is used to break ties when more than one PCRF Sub-Pool Selection Rule matches a given binding-capable session initiation request. Multiple rules can match a request when more than one rule using a "Starts	Format: Text box Range: 1-99, inclusive, where a lower value equates to a higher priority Default: 50

Fields (* indicates required field)	Description	Data Input Notes
	With" or "Ends With" condition exists.	
PCRF Pool Name	The PCRF Pool that is being subdivided by this PCRF Sub-Pool Selection Rule. A pulldown menu contains the names of all available PCRF Pools. The PCRF Pool does not need to have any APN mapped to it when this PCRF Sub-Pool Selection Rule is created. This field is a hyper-link to the PCRF Pools view screen, filtered by the PCRF Pool name.	Format: Dropdown menu Range: Configured PCRF Pools that have not been specified as PCRF Sub-Pool Names
Conditions	A condition allows for configuration of a value to be compared to a given Diameter AVP using the specified operator. The only condition currently supported for PCRF Sub-Pool Selection Rules is for the Origin-Host AVP. The value field allows you to enter a string to be compared to the Origin-Host using the operator.	Format: Textbox Range: Equals, Starts With, and Ends With.
PCRF Sub-Pool Name	The PCRF Sub-Pool name that is used for routing new bindings created from binding-capable session initiation requests that matched this PCRF Sub-Pool Selection Rule. A match occurs when the APN in the request mapped to the PCRF Pool in the rule AND the Origin-Host condition matched. This field is a hyper-link to the PCRF Pools view screen, filtered by the PCRF Sub-Pool name.	Format: Hyperlink Range: Assigned PCRF Sub-Pool Name

Fields (* indicates required field)	Description	Data Input Notes
Last Updated	The PCRF Sub-Pool Selection Rules view page also includes a timestamp of the time the rule was created or last updated, whichever occurred most recently. This field can help you troubleshoot by allowing comparison of existing binding session creation time stamps (displayed using the binding key query tool) with rule creation time stamps. Use this capability to determine whether a binding was created before or after a rule was created	Format: Read-only field Range: N/A

Inserting PCRF Sub-Pool Selection Rules

Use this task to insert (create new) PCRF Sub-Pool Selection Rules.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page appears.
2. Click **Insert**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules [Insert]** page opens.
3. Enter a unique PCRF Sub-Pool Selection Rules Name in the **PCRF Pool Selection Rule Name** field.
Enter a unique name that identifies the PCRF Sub-Pool Selection Rule. The default is N/A, and the range is a 32-character string. Valid characters are alphanumeric and underscore, and must contain at least one alpha character and must not start with a digit..
4. Enter a priority value for this rule in **Priority**.
The lower the value means the higher the priority. The default is 50, and the range is 1 to 99.
5. Select a PCRF Pool Name from the **PCRF Pool Name** pulldown menu.
This is the name of the PCRF Pool for which a Sub-Pool is being defined The default is N/A, and the range is Configured PCRF Pools that have not been specified as PCRF Sub-Pool Names.
6. Select a condition from the **Operator** pulldown menu to associate the selected condition with this rule.
The range is Equals, Starts With, or Ends With.
FQDN is a case-insensitive string consisting of a list of labels separated by dots, where a label can contain alphanumeric characters, dashes, underscores. A label must start with a letter, digit or underscore and must end with a letter or digit. Underscores can be used as the first character only.

A label range is 1 to 64, and an FQDN range is 1 to 255 characters in length. The default is N/A, and the range is Substring or complete string of a valid FQDN.

7. Enter a value in the **Value** field.
8. Select a PCRF Sub-Pool Name in the **PCRF Sub-Pool Name** pulldown menu. Choices include all the qualified PCRF Sub-Pool configured from the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page. A qualified PCRF Sub-Pool is a PCRF Pool that is non-retired and has been marked as Sub-Pool. A retired PCRF Sub-Pool entry can be created by first adding a new PCRF Sub-Pool and then deleting it.
This is the PCRF Sub-Pool that is to be used for Gx and Gxx session initiation request messages that match this Rule. The default is N/A and the range is the choice of configured PCRF Pools.
9. The **Last Updated** field is a read-only field that displays the date and time that this rule was created, or the last time the rule was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding-capable session initiation requests. This date and time can be compared against binding creation times when troubleshooting using the Binding Key Query Tool.
10. Click:
 - **OK** to save the new PCRF Pool name and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.
 - **Apply** to save the new PCRF Sub-Pool Selection Rule and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rule** page without saving any changes.

Editing PCRF Sub-Pool Selection Rules

Use this task to edit PCRF Sub-Pool Selection Rules.

The PCRF Sub-Pool Selection Rule edit page allows a network operator to change all fields except the PCRF Sub-Pool Selection Rule Name. Changes take effect on the next binding-capable session initiation request received after the rule is successfully committed.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.
The **Policy and Charging > Configuration > Policy DRA > PCRF PSub-Pool Selection Rules** page appears. The PCRF Sub-Pool Selection table contains rules for selection of a PCRF Sub-Pool for a given PCRF Pool and Origin-Host value.
2. Select a PCRF Sub-Pool Selection Rule to edit.
DO NOT click the blue PCRF Pool Name or the PCRF Sub-Pool Name (except to see the configuration of the PCRF Pool Name or PCRF Sub-Pool Name). The blue color indicates a hyper-link that opens the **Diameter > Configuration > Peer Nodes [Filtered]** page to display the configuration information for the Peer Node.
3. Click **Edit**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pools Selection Rules [Edit]** page appears. You cannot edit the **PCRF Sub-Pool Selection Rule** value. This is a name that uniquely identifies the PCRF Sub-Pool Selection Rule. The default is N/A, and the range is a 32-character string. Valid characters are alphanumeric and underscore, and must contain at least one alpha character and must not start with a digit.
4. Enter a priority value for this rule in **Priority**.
The lower the value means the higher the priority. The default is 50, and the range is 1 to 99.

5. Enter a PCRF Pool Name.
The name of the PCRF Sub-Pool Selection Rules for which a Sub-Pool is being defined. The default is N/A, and the range is Configured PCRF Sub-Pool Selection Rules that have not been specified as PCRF Sub-Pool Names.
6. Specify the condition associated with this rule.
Select a Host-Origin Operator value from the pulldown menu. FQDN is a case-insensitive string consisting of a list of labels separated by dots, where a label can contain letters, digits, dashes ('-') and underscores ('_'). A label must start with a letter, digit, or underscore, and it must end with a letter or digit. Underscores can be used as the first character only. A label cannot exceed 63 characters in length and an FQDN cannot exceed 255 characters in length. The default is N/A, and the range is a substring or complete string of a valid FQDN.
7. Select a PCRF Sub-Pool Name value from the pulldown menu.
This PCRF Sub-Pool that will be used for Gx and Gxx session initiation request messages matching this Rule. The default is N/A, and the range is the choice of configured PCRF Sub-Pool Selection Rules.
8. **Last Updated** is a read-only field that displays the date and time that this rule was created, or the last time the rule was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding capable session initiation requests. This date and time can be compared against binding creation times when troubleshooting using the Binding Key Query Tool.
9. Click:
 - **Ok** to save the change and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page.
 - **Apply** to save the change and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF PCRF Sub-Pool Selection Rules** page without saving any changes.

If **Apply** or **OK** is clicked and the selected **PCRF Peer Node Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting PCRF Sub-Pool Selection Rules

Use the following procedure to delete a PCRF.

A PCRF Sub-Pool Selection Rule can be deleted at any time.

1. Select **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page appears.
2. Select the **PCRF Sub-Pool Selection Rule Name** to be deleted.
3. Click **Delete**.
A popup window appears to confirm the delete.
4. Click:
 - **OK** to delete the PCRF Sub-Pool Selection Rule Name.
 - **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page.

If **OK** is clicked and the selected PCRF no longer exists (it was deleted by another user), an error message is displayed and the PCRF Sub-Pool Selection Rules page is refreshed. The row that was selected is no longer displayed in the list.

Network-Wide Options

On the **Policy and Charging > Configuration > Policy DRA > Network-Wide Options** page on an Active NOAM, the following **Network-Wide Options** can be configured:

The fields are described in [Network-Wide Options elements](#).

General Options

- Indicate whether to use the Local Host Origin-Host and Origin-Realm or the PCRF Origin-Host and Origin-Realm as the Origin-Host and Origin-Realm in RAR messages that are constructed and sent by Policy DRA to the Policy Clients.
- Enable PCRF Pooling.

Early Binding Options

- Set the **Early Binding Polling Interval** value (number of milliseconds between sending queries to the early binding master).
- Set the **Maximum Early Binding Lifetime** value (the maximum time that a binding is allowed to remain as an early binding).

Topology Hiding Options

- Enable Topology Hiding
- Set the Topology Hiding Scope
- Set the Default Topology Hiding Virtual Name

Network-Wide Options elements

[Table 33: Policy DRA Network-Wide Options elements](#) describes the elements on the **Policy and Charging > Configuration > Policy DRA > Network-Wide Options** page on the NOAM.

Table 33: Policy DRA Network-Wide Options elements

Fields (* indicates a required field)	Description	Data Input Notes
General Options		
Origin-Host and Origin-Realm for Policy DRA generated RAR messages	The selected option's Origin-Host and Origin-Realm will be used as the Origin-Host and Origin-Realm in the RAR messages constructed and sent by Policy DRA to the Policy Clients.	Format: Radio buttons Range: Local Host or PCRF Default: Local Host
Enable PCRF Pooling	Indicates whether the PCRF Pooling feature is enabled. Check the box to allow a subscriber's policy sessions to be routed to different PCRFs	Format: Checkbox Range: Yes (Checked) or No (Unchecked)

Fields (* indicates a required field)	Description	Data Input Notes
	<p>depending on the Access Point Network the session originated from. this box must be checked following acceptance of upgrade or future upgrades will be disallowed.</p>	<p>Default: PCRF Pooling Enabled (checked) for initial installs; PCRF Pooling Disabled (Unchecked) for upgrades from activated pre-5.1 releases.</p>
<p>Early Binding Options</p>		
<p>Early Binding Polling Interval</p>	<p>The number of milliseconds between sending queries to the early binding master to determine which PCRF the master session was routed to so that the slave session can be routed to the same PCRF.</p> <p>The goal is to set this value such that the master session has time to receive an answer a high percentage of the time. Choosing a low value increases database queries, but may reduce latency. A high value does the opposite</p> <p>Note: This values is used only when PCRF Pooling is enabled.</p>	<p>Format: Text box</p> <p>Range: 50 to 10000 milliseconds</p> <p>Default: 200 milliseconds</p>
<p>Maximum Early Binding Lifetime</p>	<p>The maximum time that a binding is allowed to remain as an early binding.</p> <p>The ideal setting for this value is 100 - 200 msec longer than the Diameter transaction timeout. This value prevents bindings from becoming stuck for long periods in the early binding state due to congestion or other error conditions. If a new Diameter request or polling attempt discovers a binding session that has been in the early state for longer than this time, the binding session is removed.</p> <p>Note: This value is used only when PCRF Pooling is Enabled.</p>	<p>Format: Text box</p> <p>Range: 500 to 15000 milliseconds</p> <p>Default: 2500 milliseconds</p>
<p>Topology Hiding Options</p>		
<p>Enable Topology Hiding</p>	<p>Enable or disable topology hiding using the check box. Once enabled or disabled here, the Topology Hiding is enabled or disabled at all SOAMs under this NOAM.</p>	<p>Format: Check box</p> <p>Range: Enabled (checked), Disabled (unchecked)</p> <p>Default: Disabled (unchecked)</p>

Fields (* indicates a required field)	Description	Data Input Notes
Topology Hiding Scope	This sets the scope of messages where topology hiding will be applied. Select 'All Messages' to perform topology hiding for all messages destined to policy clients. Select 'All Foreign Realms' to perform topology hiding for messages destined to the policy clients whose realms are different from the realm of the PCRF to be bound. Select 'Specific Clients' to perform topology hiding for the policy clients that are configured in one of SOAM GUI Main Menu: Policy and Charging->Configuration->Policy DRA->Policy Clients screen. Select 'All Foreign Realms + Specific Clients' to perform topology hiding if either condition ('All Foreign Realms' or 'Specific Clients') is met.	Format: Pulldown list Range: All Messages, All Foreign Realms, Specific Clients, All Foreign Realms + Specific Clients Default: N/A
Default Topology Hiding Virtual name	<ul style="list-style-type: none"> FQDN - This FQDN is used as a default value in the Origin-Host AVP for answer messages routed from a PCRF to a policy client, or in the Destination-Host AVP for request messages routed from a PCRF to a policy client, only if Topology Hiding Virtual Name FQDN is not configured at a SOAM relevant to the policy client and PCRF. Realm - This Realm is used as a default value in the Origin-Realm AVP for answer messages routed from a PCRF to a policy client, or in the Destination-Realm AVP for request messages routed from a PCRF to a policy client, only if Topology Hiding Virtual Name Realm is not configured at a SOAM relevant to the policy client and PCRF. 	Format: Text box Range: FQDN and Realm - a case-insensitive string consisting of a list of labels separated by dots, where a label may contain letters, digits, dashes ('-') and underscore ('_'). A label must start with a letter, digit or underscore and must end with a letter or digit. Underscores may be used only as the first character. A label must be at most 63 characters long and a FQDN must be at most 255 characters long. Default: N/A

Viewing Network-Wide Options

Use this task to view configured Network-Wide Options on the NOAM.

Select **Policy and Charging > Configuration > Policy DRA > Network-Wide Options**.

The **Policy and Charging > Configuration > Policy DRA > Network-Wide Options** page appears with a list of configured Network-Wide Options.

The fields are described in [Network-Wide Options elements](#).

Setting Network-Wide Options

Use this task to set Network-Wide Options on the NOAM.

The fields are described in [Network-Wide Options elements](#).

The following Policy DRA configuration options apply to the entire Policy DRA Network:

- Origin-Host and Origin-Realm for Policy DRA generated RAR messages
 - Enable PCRF Pooling
 - Early Binding
 - Topology Hiding
1. Select **Policy and Charging > Configuration > Policy DRA > Network-Wide Options**.
The **Policy DRA Network-Wide Options** page appears.
 2. Select the **Local Host** or **PCRF** radio button.
This sets the Origin-Host and Origin-Realm that will be used in the RAR messages constructed and sent by Policy DRA to policy clients.
 3. Select the **Enable PCRF Pooling** check box.
Checking the box allows a subscriber's policy sessions to be routed to different PCRFs depending on the Access Point Network the session originated from. This box must be checked following acceptance of upgrade or future upgrades will be disallowed.
 4. Set the **Early Binding Polling Interval**.
The number of milliseconds between sending queries to the early binding master to determine which PCRF the master session was routed to so that the slave session can be routed to the same PCRF.

The goal is to set this value such that the master session has time to receive an answer a high percentage of the time. Choosing a low value increases database queries, but may reduce latency. A high value does the opposite.

Note: This value is used only when PCRF Pooling is Enabled.
 5. Set the **Maximum Early Binding Lifetime**.
The maximum time that a binding is allowed to remain as an early binding.

The ideal setting for this value is 100 - 200 msec longer than the Diameter transaction timeout. This value prevents bindings from becoming stuck for long periods in the early binding state due to congestion or other error conditions. If a new Diameter request or polling attempt discovers a binding session that has been in the early state for longer than this time, the binding session is removed.

Note: This value is used only when PCRF Pooling is Enabled.
 6. Select the **Enable Topology Hiding** check box.
Enable or disable topology hiding using the check box. Once enabled or disabled here, the Topology Hiding is enabled or disabled at all SOAMs under this NOAM.
 7. Select a **Topology Hiding Scope**.
This sets the scope of messages where topology hiding will be applied. Select 'All Messages' to perform topology hiding for all messages destined to policy clients. Select 'All Foreign Realms' to perform topology hiding for messages destined to the policy clients whose realms are different from the realm of the PCRF to be bound. Select 'Specific Clients' to perform topology hiding for the policy clients that are configured in one of SOAM GUI Main Menu: **Policy and Charging >**

Configuration > Policy DRA > Policy Clients screen. Select 'All Foreign Realms + Specific Clients' to perform topology hiding if either condition ('All Foreign Realms' or 'Specific Clients') is met.

8. Enter a **Default Topology Hiding Virtual Name**.

The entered FQDN is used as a default value in the Origin-Host AVP for answer messages routed from a PCRF to a policy client, or in the Destination-Host AVP for request messages routed from a PCRF to a policy client, only if Topology Hiding Virtual Name FQDN is not configured at a SOAM relevant to the policy client and PCRF.

The entered Realm is used as a default value in the Origin-Realm AVP for answer messages routed from a PCRF to a policy client, or in the Destination-Realm AVP for request messages routed from a PCRF to a policy client, only if Topology Hiding Virtual Name Realm is not configured at a SOAM relevant to the policy client and PCRF.

9. Click:

- **Apply** to save the changes and remain on this page.
- **Cancel** to discard changes and remain on the **Policy and Charging > Configuration > Policy DRA > Network-Wide Options** page.

If **Apply** is clicked and the following condition exists, an error message appears:

Online Charging DRA

This section describes the **Policy and Charging > Configuration > Online Charging DRA** GUI pages on the NOAM and the SOAM.

OCS Session State

On an Active NOAM, the **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page lists the network-wide list of Online Charging Servers (OCSs), listed by their Realm and FQDN. It is used to configure the Session State setting for OCSs.

The list of OCSs is updated by inserting or deleting an OCS from the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page at each site's SOAM. Additionally, the Realm and FQDN are configured from each site's **Diameter > Configuration > Peer Nodes** page on the SOAM.

Once the list of OCSs is populated, the following options become available:

- Editing whether or not OCS Session State is enabled
- Pausing the updating of the OCS list

Editing OCS Session State

Use this task to edit a Realm.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > OCS Session State**.

The **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page appears.

2. Click **Edit**.

The **Policy and Charging > Configuration > Online Charging DRA > OCS Session State [Edit]** page opens.

3. The **Realm** and **FQDN** fields are disabled and cannot be edited from this screen.
4. Check or uncheck the box to Enable or Disable **OCS Session State**.
5. Click:
 - **OK** to save the edited OCS Session State and return to the **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page.
 - **Apply** to save the edited OCS Session State and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any required field no longer exists

Realms

The **Policy and Charging > Configuration > Online Charging DRA > Realms** page on an NOAM contains the list of Online Charging network realms for which the Session state is stored.

Note: This page is only use if **Session State Scope** is set to "Specific Messages" on the **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options** page.

Realms elements

[Table 34: Realms elements](#) describes the elements on the **Policy and Charging > Configuration > Online Charging DRA > Realms** page.

Note: Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

The Realms table lists the Online Charging network realms for which the Session state is to be stored. This table is only used if Session State Scope is set to "Specific Messages" in the Network-Wide Options Configuration.

Table 34: Realms elements

Fields (* indicates required field)	Description	Data Input Notes
*Realm Name	Realm name is a case-insensitive string consisting of a list of lables separated by dots, where a label may contain letter, digits, dashes('-') and underscore('_'). A label must start with a letter, digit or underscore and must end with a letter or digit. Underscores may be used only as the first character. A label must be at most 63 characters long and a Realm must be at most 255 characters long.	Format: text box Range: 1-1000 entries Default: N/A
Comments	An optional comment to provide more information about the purpose of this PCRF Pool or Sub-Pool.	Format: Text box Range:0-64 characters

Inserting Realms

Use this task to insert (create new) Realms.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > Realms**.
The **Policy and Charging > Configuration > Online Charging DRA > Realms** page appears.
2. Click **Insert**.
The **Policy and Charging > Configuration > Online Charging DRA > Realms [Insert]** page opens.
3. Enter a unique Realm Name in the **Realm Name** field.
4. If desired, enter an optional comment in the **Comments** field to describe the Realm. The entry must be characters in the range of 0 to 64, and the default is N/A.
5. Click:
 - **OK** to save the new Realm name and return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page.
 - **Apply** to save the new Realm name and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any required field is empty (not entered).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Adding a new Realm would cause the maximum number of Realms (1000) to be exceeded.
- The entered Realm name is not unique (already exists).

Editing Realms

Use this task to edit a Realm.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > Realms**.
The **Policy and Charging > Configuration > Online Charging DRA > Realms** page appears.
2. Click **Edit**.
The **Policy and Charging > Configuration > Online Charging DRA > Realms [Edit]** page opens.
3. Edit the unique Realm Name in the **Realm Name** field.
4. If desired, edit an optional comment in the **Comments** field.
5. Click:
 - **OK** to save the edited Realm name and return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page.
 - **Apply** to save the edited Realm name and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any required field no longer exists
- Any fields contain a value that contains invalid characters or is out of the allowed range.

Deleting Realms

Use this task to delete a PCRF Pool.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > Realms**.

The **Policy and Charging > Configuration > Online Charging DRA > Realms** page appears.

2. Select the **Realm** to be deleted.
3. Click **Delete**.

A window appears to confirm the delete.

4. Click:

- **OK** to delete the Realm.
- **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page.

If **OK** is clicked and the selected Realm no longer exists (it was deleted by another user), an error message is displayed, and the Realms page is refreshed. The row that was selected is no longer displayed in the list.

Network-Wide Options

On the **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options** page on an Active NOAM, the following Network-Wide Options can be configured:

The fields are described in [Network-Wide Options elements](#).

Session Options

- Set the scope of messages for which Session State will be stored.
- Set the action to be performed if an in-session Request message cannot be successfully processed due to the inability to retrieve session state associated with the received Session-Id from the Session SBR (i.e., session state is not found or an SBR error is encountered).

OCS Selection Options

- Set the operating mode for selecting the OCS Server for routing the Session Initiation Request messages.

Network-Wide Options elements

[Table 35: Online Charging DRA Network-Wide Options elements](#) describes the elements on the **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options** page on the NOAM.

Table 35: Online Charging DRA Network-Wide Options elements

Fields (* indicates a required field)	Description	Data Input Notes
Session Options		
Session State Scope	This sets the scope of messages for which Session State will be stored.	Format: Pulldown menu

Fields (* indicates a required field)	Description	Data Input Notes
	Select 'All Messages' to store Session State for all messages. Select 'None' to disable Session State for all messages. Select 'Specific Messages' to store Session State only if the CTF client is configured in the CTFs configuration or OCS is configured with Session State as enabled in OCSs configuration or realm is configured in Realms configuration.	Range: None, All Messages, Specific Messages Default: None
Session State Unavailable Action	Sets the action to be performed if an in-session Request message cannot be successfully processed due to the inability to retrieve session state associated with the received Session-Id from the Session SBR (i.e., session state is not found or an SBR error is encountered). 'Route to Peer' will route the message to a peer using the Peer Routing Table. 'Send Answer' will abandon message processing and send an Answer response containing Answer Result-Code value configured for 'Session Not Found' or 'SBR Error'.	Format: Pulldown menu Range: Send Answer, Route To Peer Default: Send Answer
OCS Selection Options		
OCS Pool Selection Mode	This sets the operating mode for selecting the OCS Server for routing the Session Initiation Request messages. When 'Single Pool' mode is selected, the Session Initiation Requests are distributed in a weighted round-robin scheme among all available OCS servers connected to this Node. When 'Multiple Pools' mode is selected, the Session Initiation Requests are routed to an OCS server identified by RBAR in a specific pool of OCS servers.	Format: Pulldown menu Range: Single Pool, Multiple Pools Default: Single Pool

Note: Keep these consideration in mind when working with network-wide options:

- If **Apply** is clicked and the 'Session State Scope' transitioned from 'None' to 'All Messages', a confirmation dialog with a checkbox shall be displayed containing the text: "IMPORTANT! The Session State for all the messages will be enabled. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected. Check the checkbox and click OK to continue, otherwise click **Cancel**."
- If **Apply** is clicked and the 'Session State Scope' transitioned from 'None' to 'Specific Messages', a confirmation dialog with a checkbox shall be displayed containing the text: "IMPORTANT! The Session State Scope 'Specific Messages' requires the OCSs, CTFs or Realms to be configured for maintaining Session State. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected."
- If **Apply** is clicked and the 'Session State Scope' transitioned from 'All Messages' to 'Specific Messages', a confirmation dialog with a checkbox shall be displayed containing the text:

"IMPORTANT! The Session State Scope 'Specific Messages' requires the OCSs, CTFs or Realms to be configured for maintaining Session State. Some or all subsequent in-session messages may be rejected if Destination-Host AVP is not present in them. Check the checkbox and click OK to continue, otherwise click **Cancel**."

- If **Apply** is clicked and the 'Session State Scope' transitioned from 'Specific Messages to All Messages', a confirmation dialog with a checkbox shall be displayed containing the text: "IMPORTANT! The Session State for all the messages will be enabled. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected. Check the checkbox and click **OK** to continue, otherwise click **Cancel**."
- If the confirmation dialog for 'Session State Scope' is cancelled by clicking **Cancel**, control is returned to the Network-Wide Options screen with no data committed.
- If the confirmation dialog for 'Session State Scope' is confirmed by checking the checkbox and clicking OK and no OCS or CTF or Realm is configured for session state maintenance, a Warning Box is displayed on the Network-Wide Options screen containing the text: "Session State Scope is configured as 'Specific Messages' but no OCS, CTF or Realm is configured for Session State maintenance." The configured data is saved in the configuration database.

Viewing Network-Wide Options

Use this task to view configured Network-Wide Options associated with Online Charging DRA on the NOAM.

Select **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options**.

The **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options** page appears with a list of configured Network-Wide Options.

The fields are described in .

Setting Network-Wide Options

Use this task to set Online Charging DRA Network-Wide Options on the NOAM.

The fields are described in [Network-Wide Options elements](#).

The following Network-Wide Options associated with Online Charging DRA can be set:

- Setting the Session State Scope
 - Setting the action to be taken if the Session State is Unavailable
 - Setting the OCS Pool Selection Mode
1. Select **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options**. The **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options** page appears.
 2. Select a **Session State Scope** from the pulldown list.
 3. Select a **Session State Unavailable Action** from the pulldown list.
 4. Select an **OCS Pool Selection Mode** from the pulldown list.
 5. Click:
 - **Apply** to save the changes and remain on this page.
 - **Cancel** to discard changes and remain on the **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options** page.

If **Apply** is clicked and the following condition exists, a warning message appears:

- If **Apply** is clicked and the 'Session State Scope' transitioned from 'None' to 'All Messages', a confirmation dialog with a checkbox shall be displayed containing the text: "IMPORTANT! The Session State for all the messages will be enabled. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected. Check the checkbox and click OK to continue, otherwise click **Cancel**."
- If **Apply** is clicked and the 'Session State Scope' transitioned from 'None' to 'Specific Messages', a confirmation dialog with a checkbox shall be displayed containing the text: "IMPORTANT! The Session State Scope 'Specific Messages' requires the OCSs, CTFs or Realms to be configured for maintaining Session State. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected."
- If **Apply** is clicked and the 'Session State Scope' transitioned from 'All Messages' to 'Specific Messages', a confirmation dialog with a checkbox shall be displayed containing the text: "IMPORTANT! The Session State Scope 'Specific Messages' requires the OCSs, CTFs or Realms to be configured for maintaining Session State. Some or all subsequent in-session messages may be rejected if Destination-Host AVP is not present in them. Check the checkbox and click OK to continue, otherwise click **Cancel**."
- If **Apply** is clicked and the 'Session State Scope' transitioned from 'Specific Messages to All Messages', a confirmation dialog with a checkbox shall be displayed containing the text: "IMPORTANT! The Session State for all the messages will be enabled. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected. Check the checkbox and click **OK** to continue, otherwise click **Cancel**."
- If the confirmation dialog for 'Session State Scope' is cancelled by clicking **Cancel**, control is returned to the Network-Wide Options screen with no data committed.
- If the confirmation dialog for 'Session State Scope' is confirmed by checking the checkbox and clicking OK and no OCS or CTF or Realm is configured for session state maintenance, a Warning Box is displayed on the Network-Wide Options screen containing the text: "Session State Scope is configured as 'Specific Messages' but no OCS, CTF or Realm is configured for Session State maintenance." The configured data is saved in the configuration database.

Alarm Settings

Note: Alarm Settings are configurable only on Active NOAM servers, and are viewable on NOAM and SOAM servers.

On the **Policy and Charging > Configuration > Alarm Settings** page on an SOAM, the user can view the configured Alarm Thresholds and Suppress indications.

Each alarm can be configured with Minor, Major, and Critical threshold percentages.

The fields are described in [Alarm Settings elements](#).

On the **Policy and Charging > Configuration > Alarm Settings** page on the NOAM, you can change the Alarm Thresholds and the Suppress indications for the following alarms:

- DSR Application Ingress Message Rate

The DSR Application Ingress Message Rate alarm is raised when the average Policy and Charging ingress messages rate exceeds the configured Alarm Threshold. The thresholds are based on the engineered system value for Ingress Message Capacity.

- SBR Sessions Threshold Exceeded

The SBR Sessions Threshold Exceeded alarm percent full is based on the number of Session records compared to an engineered maximum that varies according to the number of session SBR Server Groups per mated pair.

The SBR Sessions Threshold Exceeded alarm is raised when number of concurrent policy and Charging sessions exceeds the configured threshold.

- SBR Bindings Threshold Exceeded

The SBR Bindings Threshold Exceeded alarm measures the number of IMSI Anchor Key records against an engineered maximum value that varies according to the number of binding SBR Server Groups.

The Policy SBR Bindings Threshold Exceeded alarm works similarly to the session capacity alarm except that the scope of the binding capacity alarm is network-wide.

Alarm Settings elements

Table 36: Alarm Settings elements describes the elements on the **Policy and Charging > Configuration > Alarm Settings** page. The elements can be configured and viewed on the NOAM, and only viewed on the SOAM. Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

The page contains three sets of input fields for the following alarms:

- DSR Application Ingress Message Rate
- SBR Sessions Threshold Exceeded
- SBR Bindings Threshold Exceeded

The element labels are the same for each input field set, but some serve different purposes and have different values. These distinctions are noted in the table.

Table 36: Alarm Settings elements

Elements (* indicates required field)	Description	Data Input Notes
DSR Application Ingress Message Rate		
* Alarm Name	This alarm is raised when average Policy and Charging ingress messages rate exceeds the configured threshold. The thresholds are based on the engineered system value for Ingress Message Capacity.	Format: Non-editable text box Range: DSR Application Ingress Message Rate
* Critical Alarm Threshold (Percent)	The Policy and Charging ingress message rate threshold for this alarm to be raised as Critical. The threshold is a percentage of the Ingress Capacity Capability.	Format: Text box Range: 100-200 Default: 160
Suppress Critical	Controls whether this alarm is raised as Critical.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)

Policy and Charging Configuration

Elements (* indicates required field)	Description	Data Input Notes
* Major Alarm Threshold (Percent)	The Policy and Charging ingress message rate threshold for this alarm to be raised as Major. The threshold is a percentage of the Ingress Capacity Capability.	Format: Text box Range: 100-200 Default: 140
Suppress Major	Controls whether this alarm is raised as Major.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Minor Alarm Threshold (Percent)	The Policy and Charging ingress message rate threshold for this alarm to be raised as Minor. The threshold is a percentage of the Ingress Capacity Capability.	Format: Text box Range: 100-200 Default: 110
Suppress Minor	Controls whether this alarm is raised as Minor.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
SBR Sessions Threshold Exceeded		
* Alarm Name	This alarm is raised when the number of concurrent Policy and Online Charging SBR sessions exceeds the configured threshold.	Format: Non-editable text box Range: Policy SBR Sessions Threshold Exceeded
* Critical Alarm Threshold (Percent)	The concurrent sessions threshold for this alarm to be raised as Critical. The threshold is a percentage of the Maximum SBR Sessions.	Format: Text box Range: 1-99 Default: 95
Suppress Critical	Controls whether this alarm is raised as Critical.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Major Alarm Threshold (Percent)	The concurrent sessions threshold for this alarm to be raised as Major. The threshold is a percentage of the Maximum SBR Sessions.	Format: Text box Range: 1-99 Default: 90

Policy and Charging Configuration

Elements (* indicates required field)	Description	Data Input Notes
Suppress Major	Controls whether this alarm is raised as Major.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Minor Alarm Threshold (Percent)	The concurrent sessions threshold for this alarm to be raised as Minor. The threshold is a percentage of the Maximum SBR Sessions.	Format: Text box Range: 1-99 Default: 80
Suppress Minor	Controls whether this alarm is raised as Minor.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
SBR Bindings Threshold Exceeded		
* Alarm Name	This alarm is raised when the number of concurrent Policy SBR bindings exceeds the configured threshold.	Format: Non-editable text box Range: Policy SBR Bindings Threshold Exceeded
* Critical Alarm Threshold (Percent)	The concurrent bindings threshold for this alarm to be raised as Critical. The threshold is a percentage of the Maximum Policy SBR Bindings.	Format: Text box Range: 1-99 Default: 95
Suppress Critical	Controls whether this alarm is raised as Critical.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Major Alarm Threshold (Percent)	The concurrent bindings threshold for this alarm to be raised as Major. The threshold is a percentage of the Maximum Policy SBR Bindings.	Format: Text box Range: 1-99 Default: 90
Suppress Major	Controls whether this alarm is raised as Major.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)

Elements (* indicates required field)	Description	Data Input Notes
* Minor Alarm Threshold (Percent)	Te concurrent bindings threshold for this alarm to be raised as Minor. The threshold is a percentage of the Maximum Policy SBR Bindings.	Format: Text box Range: 1-99 Default: 80
Suppress Minor	Controls whether this alarm is raised as Minor.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)

Viewing Alarm Settings

Use this task to view configured Alarm-Settings on either the NOAM or SOAM.

Select **Policy and Charging > Configuration > Alarm Settings**.

The **Policy and Charging > Configuration > Alarm Settings** page appears with a list of configured Alarm Settings.

The fields are described in [Alarm Settings elements](#).

Defining Alarm Settings

Use this task to define Alarm Settings on an Active NOAM.

Note: Alarm Settings are configurable only on Active NOAM servers, and are viewable on NOAM and SOAM servers.

The fields are described in [Alarm Settings elements](#).

1. Select **Policy and Charging > Configuration > Alarm Settings**.

The **Policy and Charging > Configuration > Alarm Settings** page appears.

2. Enter values in the editable fields to define the alarm settings.

3. Click:

- **Apply** to save the changes and remain on this page.
- **Cancel** to discard the changes and remain on the **Policy and Charging > Configuration > Alarm Settings** page.

If **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entered values contain the wrong data type or is out of the allowed range.
- The value entered for **Critical Alarm Threshold (Percent)** is less than or equal to the value entered for **Major Alarm Threshold (Percent)**.
- The value entered for **Major Alarm Threshold (Percent)** is less than or equal to the value entered for **Minor Alarm Threshold (Percent)**.

Congestion Options

Congestion Options are configurable on Active NOAM servers.

The following Congestion Options can be configured:

- Alarm Thresholds, which are used to:
 - Set the percentage of the Policy and Charging ingress message rate capacity at which an alarm is raised with Critical, Major, or Minor severity.
 - Set the percentage of the Policy and Charging ingress message rate capacity at which a Critical, Major, or Minor severity alarm is cleared.

The percentages control the onset and abatement of the corresponding Congestion Levels.

Default thresholds are based on the engineered system value for Ingress Policy and Charging Request Message Capacity.

- Message Throttling Rules, which determine the percentage of Session Creation, Update, and Terminate Request messages that are discarded when Congestion Levels 1, 2, and 3 exist.

The fields are described in [Congestion Options elements](#).

Congestion Options elements

[Table 37: Congestion Options elements](#) describes the elements on the **Policy and Charging > Configuration > Congestion Options** page. The elements can be configured and viewed on the NOAM.

The page contains two sets of input fields:

- Alarm Thresholds
- Message Throttling Rules

Table 37: Congestion Options elements

Fields (* indicates required field)	Description	Data Input Notes
Alarm Thresholds		
Alarm Name	Alarm is raised when average Policy and Charging ingress request messages rate exceeds the configured threshold. The thresholds are based on the engineered system value for Ingress Policy and Charging Request Message Capacity.	Format: Non-editable text box Range: Policy and Charging Server in Congestion
* Critical Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 3.	Format: Text box Range: 100-200 Default: 160
* Critical Alarm Abatement Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 3.	Format: Text box Range: 100-200 Default: 150

Policy and Charging Configuration

Fields (* indicates required field)	Description	Data Input Notes
* Major Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 2.	Format: Text box Range: 100-200 Default: 140
* Major Alarm Abatement Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 2.	Format: Text box Range: 100-200 Default: 130
* Minor Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 1.	Format: Text box Range: 100-200 Default: 110
* Minor Alarm Abatement Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 1.	Format: Text box Range: 100-200 Default: 100
<p>Message Throttling Rules</p> <p>Tabs for Congestion Level 1, Congestion Level 2, and Congestion Level 3</p>		
* Discard Session Creation Requests	Percentage of Request messages that result in new session creation, to be discarded when this congestion level exists.	Format: Text box Range: 0-100 Default: Level 1 - 25 Level 2 - 50 Level 3 - 100
* Discard Session Update Requests	Percentage of Request messages that result in updating existing sessions, to be discarded when this congestion level exists.	Format: Text box Range: 0-100 Default: Level 1 - 0 Level 2 - 25 Level 3 - 50
* Discard Session Terminate Requests	Percentage of Request messages that result in terminating existing sessions, to be discarded when this congestion level exists.	Format: Text box Range: 0-100 Default: Level 1 - 0 Level 2 - 0

Fields (* indicates required field)	Description	Data Input Notes
		Level 3 - 0

Viewing Congestion Options

Use this task to view configured Congestion Options on the NOAM.

Select **Policy and Charging > Configuration > Congestion Options**.

The **Policy and Charging > Configuration > Congestion Options** page appears with a list of configured Congestion Options.

The fields are described in [Congestion Options elements](#).

Setting Congestion Options

Use this task to set the following Congestion Options on the Active NOAM:

- **Alarm Thresholds** for the **Policy and Charging Server in Congestion** onset and abatement alarm for Critical, Major, and Minor severities
 - **Message Throttling Rules** for discarding Session Creation, Update, and Terminate Requests for Congestion Levels 1, 2, and 3
1. Select **Policy and Charging > Configuration > Congestion Options**.
The **Policy and Charging > Configuration > Congestion Options** page appears.
 2. Enter changes for the **Alarm Thresholds**.
 3. Enter changes for the **Message Throttling Rules**.
 4. Click:
 - **Apply** to save the Congestion Options changes and refresh the page to show the changes.
 - **Cancel** to discard the changes and refresh the page.

If **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- A **Major Alarm Onset Threshold** value is greater than the corresponding **Critical Alarm Onset Threshold**.
- A **Minor Alarm Onset Threshold** value is greater than the corresponding **Major Alarm Onset Threshold**.
- An **Alarm Abatement Threshold** value is greater than the corresponding **Alarm Onset Threshold** of a particular severity.

SOAM Configuration

This section describes the **Policy and Charging > Configuration** GUI pages on the SOAM.

Policy DRA

This section describes the **Policy and Charging > Configuration > Policy DRA** GUI pages on the SOAM.

PCRFs

The **Policy and Charging > Configuration > Policy DRA > PCRFs** page contains the list of PCRF Peer Nodes that are to be used when a subscriber binding is created at this site. New bindings created at this Policy and Charging DSR are distributed evenly among the configured PCRFs.

PCRFs are responsible for authorizing and making policy decisions based on knowledge of subscriber resource usage and the capabilities allowed by the subscriber's account. All policy requests for a given subscriber must be routed to the same PCRF. Policy and Charging dynamically assigns subscribers to PCRFs using a load distribution algorithm, and maintains state about which subscribers are assigned to which PCRF. The relationship between a subscriber and a PCRF can change any time the subscriber transitions from having no Diameter policy sessions to having one or more Diameter policy sessions. After a policy session exists, all policy sessions for that subscriber are routed to the assigned PCRF.

The fields are described in [PCRFs elements](#).

Note: For details about configuring Peer Nodes, refer to the *Diameter User Guide* and Diameter online help.

On the **Policy and Charging > Configuration > Policy DRA > PCRFs** page on the SOAM, you can perform the following actions:

- Filter the list of PCRFs, to display only the desired PCRFs.
- Sort the list entries by column in ascending or descending order by clicking the column heading. By default, the list is sorted by PCRFs in ascending numerical order.
- Click the **Insert** button.

The **Policy and Charging > Configuration > Policy DRA > PCRFs [Insert]** page opens and allows the user to add a PCRF. See [Inserting PCRFs](#). If the maximum number of PCRFs (5000) already exists in the system, the **Policy and Charging > Configuration > Policy DRA > PCRFs [Insert]** page will not open, and an error message is displayed.

- Select a PCRF in the list, and click the **Edit** button.

The **Policy and Charging > Configuration > Policy DRA > PCRFs [Edit]** page opens and allows the user to edit the selected PCRF. See [Editing PCRFs](#).

- Select a PCRF in the list, and click the **Delete** button to remove the selected PCRF. See [Deleting a PCRF](#).

PCRFs elements

[Table 38: PCRFs page elements](#) describes the elements on the **Policy and Charging > Configuration > Policy DRA > PCRFs** page on the Active SOAM.

Note: Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 38: PCRFs page elements

Fields (* indicates required field)	Description	Data Input Notes
* PCRF Peer Node Name	The name of a configured Diameter Peer Node that identifies the PCRF Peer Node to be included in the distribution of new bindings to PCRFs. Selecting a PCRF Peer Node name (blue hyperlink) displays the Diameter > Configuration > Peer Nodes (Filtered) page where Diameter Peer Nodes are filtered by the PCRF Peer Node Name.	Format: List Range: Configured Diameter Peer Nodes Note: The PCRF Peer Node Name cannot be changed on the [Edit] page.
Comments	An optional comment to describe the PCRF Peer Node.	Format: Text box Range:0-64 characters

Viewing PCRFs

Use this task to view all configured PCRFs on the SOAM.

Select **Policy and Charging > Configuration > Policy DRA > PCRFs**.

The **Policy and Charging > Configuration > Policy DRA > PCRFs** page appears with a list of configured PCRF Peer Nodes.

The fields are described in [PCRFs elements](#).

Inserting PCRFs

Use this task to insert (create new) PCRFs.

The fields are described in [PCRFs elements](#).

1. On the Active SOAM, select **Policy and Charging > Configuration > Policy DRA > PCRFs**.
The **Policy and Charging > Configuration > Policy DRA > PCRFs** page appears.
 2. Click **Insert**.
The **Policy and Charging > Configuration > Policy DRA > PCRFs [Insert]** page opens.
 3. Enter a unique PCRF Peer Node Name in the **PCRF Peer Node Name** field.
This name uniquely identifies the PCRF Peer Node to be included in the load distribution of new bindings to PCRFs.
 4. Enter an optional comment in the **Comments** field.
 5. Click:
 - **OK** to save the new PCRF and return to the **Policy and Charging > Configuration > Policy DRA > PCRFs** page.
 - **Apply** to save the new PCRF and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRFs** page without saving any changes.
- If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:
- The entered PCRF is not unique (already exists).

- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- Adding the new PCRF would cause the maximum number of PCRFs (5000) to be exceeded.

Editing PCRFs

Use this task to edit PCRF Comments.

Note: The PCRF Pool Name cannot be edited.

1. On the Active SOAM, select **Policy and Charging > Configuration > Policy DRA > PCRFs**.
The **Policy and Charging > Configuration > Policy DRA > PCRFs** page appears. The page displays a list of the configured PCRF Peer Nodes that are used when a new subscriber binding is created.
2. Click in the **Comments** field of the row to select the PCRF to edit.
DO NOT click the blue PCRF Peer Node Name (except to see the configuration of the Peer Node). The blue color indicates a hyper-link that opens the **Diameter > Configuration > Peer Nodes [Filtered]** page to display the configuration information for the Peer Node.
3. Edit the **Comments** field for the selected PCRF.
The PCRF Peer Node name cannot be changed.
4. Click:
 - **OK** to save the change and return to the **Policy and Charging > Configuration > Policy DRA > PCRFs** page.
 - **Apply** to save the change and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRFs** page without saving any changes.

If **Apply** or **OK** is clicked and the selected **PCRF Peer Node Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting a PCRF

Use the following procedure to delete a PCRF.

This procedure describes the recommended steps for deleting a PCRF from a Policy and Charging configuration. In this procedure, PCRF refers to a Diameter peer of the PCA, which is sometimes referred to as a PCRF Front-end.

The PCRF procedure minimizes disruption to policy signaling by:

- Preventing sessions from creating new bindings to a PCRF that has been removed
- Allowing sessions with existing bindings to continue to use a PCRF that has been removed until those sessions terminate normally

The following procedure describes the recommended steps for deletion of a PCRF from a Policy and Charging configuration. In this procedure, PCRF refers to a Diameter peer of the PCA, sometimes referred to as a PCRF Front-End.

Note: The PCRF removal procedure is restricted to SOAM servers.

1. Use **Main Menu > Diameter > Configuration > Peer Nodes** from the SOAM GUI page to determine the Peer Node name of the PCRF(s) being removed.

2. Use **Main Menu > Diameter > Route Groups** from the SOAM GUI page, use the GUI filter by Peer Node with the corresponding Peer Node name of the PCRF. This will display only the Route Groups that are associated with the PCRF.
3. From the same GUI page, determine if there are any Route Groups that contain other Peer Nodes in addition to the PCRF to be removed.
There are generally at least two Route Groups for each PCRF. One Route Group with only the specified PCRF peer, and one or more Route Groups with the specified PCRF peer plus other PCRF peers. The goal is to leave the route group with only the specified PCRF peer, but delete the PCRF peer from the other route groups. This allows routing for existing bindings to the PCRF peer, but prevents alternate routing to the PCRF peer.
4. From the same GUI page, edit each of the determined Route Groups and remove the PCRF/PCRF Front-End Peer Nodes from the Route Group.
This prevents alternate routing selection of the PCRF peer being removed.
5. Use **Main Menu > Policy and Charging > Policy DRA > Configuration > PCRFs** from the SOAM GUI page to delete the PCRF.
This prevents new Bindings from using the PCRF peer being removed.
6. After enough time has elapsed such that all Diameter sessions that could be bound to the PCRF peer should have terminated normally, use **Main Menu > Policy and Charging > Policy DRA > Configuration > PCRFs** on the SOAM GUI page to delete the route group containing only the PCRF peer being removed.
7. Use **Main Menu > Diameter > Maintenance > Connections** from the SOAM GUI page to find the connection for the PCRF Peer Node and disable it
8. Use **Main Menu > Diameter > Maintenance > Connections** from the SOAM GUI page to delete the connection to the PCRF Peer Node.
9. Use **Main Menu > Diameter > Configuration > Peer Nodes** from the SOAM GUI page to delete the Diameter Peer Node for the PCRF being removed.

Binding Key Priority

The Binding Key Priority defines search priorities for Alternative Keys that can be used to locate a subscriber binding.

The Binding Key Priority controls:

- Which keys are stored for binding correlation
- The order in which keys are searched for purposes of binding correlation

The priority determines the order used to find a binding for subsequent sessions. Alternative Keys with an assigned priority will be created with the binding if they are present in the session initiation message that created the binding. The Alternative Keys must be assigned a priority in order to be used to locate subscriber bindings. If any Alternative Keys are not assigned a priority, they will not be used to locate subscriber bindings even if the Alternative Key is present in the session initiation message.

The fields are described in [Binding Key Priority elements](#).

On the **Policy and Charging > Configuration > Policy DRA > Binding Key Priority** page on the Active SOAM, you can change the Binding Key Type for Binding Key Priority 2, 3, and 4.

Note: Priority 1 for Binding Key Type IMSI is the highest priority and cannot be modified.

Enabling and disabling the binding key field depends on the value that you select for the Binding Key type.

Binding Key Priority elements

Table 39: Binding Key Priority elements describes the elements on the **Policy and Charging > Configuration > Policy DRA > Binding Key Priority** page.

Table 39: Binding Key Priority elements

Field (* indicates a required field)	Description	Data Input Notes
* Binding Key Type	The Binding Key Type which is assigned to a Binding Key Priority. Note: The first row is Priority 1 and the corresponding Binding Key Type is IMSI. This row is read-only.	Format: Pulldown list Range: MSISDN, IPv4, or IPv6 for Priority 2, 3, and 4 Default: -Select- (No Binding Key Type selected)

Viewing Binding Key Priority

Use this task to view configured Binding Key Priority settings on the SOAM.

Select **Policy and Charging > Configuration > Policy DRA > Binding Key Priority**.

The **Policy and Charging > Configuration > Policy DRA > Binding Key Priority** page appears with a list of configured Binding Key Priority settings.

The fields are described in *Binding Key Priority elements*.

Setting Binding Key Priority

Use this task to set Binding Key Priority values.

The fields are described in *Binding Key Priority elements*.

1. On the Active SOAM, select **Policy and Charging > Configuration > Policy DRA > Binding Key Priority**.
The **Policy and Charging > Configuration > Policy DRA > Binding Key Priority** page appears.
2. Make Binding Key Type selections for Priority 2 - 4 as needed. Priority 1 is non-editable (it is the Anchor Key and is always IMSI).
3. Click:
 - **Apply** to save the selected Binding Key Type values and remain on this page.
 - **Cancel** to remain on the **Policy and Charging > Configuration > Policy DRA > Binding Key Priority** page without saving any changes.

If **Apply** is clicked and any of the following conditions exist, an error message appears:

- A Binding Key Priority Type is selected for more than one Priority
- Binding Key Types are not selected for consecutive Priority values

PCRF Pools

Policy DRA continues to support a single pool of PCRFs at each PCA site over which policy Diameter signaling is distributed using the subscriber's IMSI. This allows the incorporation of new services or new PCRF infrastructure without disturbing existing services. For example, one set of PCRF servers handle policy control for all consumer data accesses to their network and a second set of PCRF servers handle all enterprise data accesses for their network. The policy rules and/or PCRF implementations might be different enough to necessitate that these two services are segregated at the PCRF level.

This means that a given IMSI might concurrently have a binding to one PCRF for APN A and a binding to a different PCRF for APN B. Each APN is mapped to a set of PCRFs; this is called a PCRF Pool. In addition, if a binding to a PCRF Pool already and a new session is created that maps to that same PCRF Pool, the request must be routed to the same PCRF. When new bindings are created for different IMSIs and a given APN, the binding-capable session initiation requests are distributed across the PCRFs in the PCRF Pool assigned to that APN.

PCRF Pooling expands this capability for the creation of multiple pools of PCRFs, which are selected using the combination of IMSI and Access Point Name (APN). This allows you to route policy Diameter signaling initiating from a given APN to a designated subset of the PCRFs that can provide specialized policy treatment using knowledge of the APN.

PCRF Pooling modifies the logic in the Policy DRA to inspect the contents of binding generating Gx CCR-I messages to select the type of PCRF to which the CCR-I messages are to be routed. In the initial P-DRA, it was assumed that all PCRFs could handle all Gx session bindings. PCRF Pooling provides service-specific sets of PCRFs. In this release, the APN used by the UE to connect to the network is used to determine the PCRF pool. The Origin-Host of the PCEF sending the CCR-I can then be used to select a PCRF sub-pool.

Multiple PCRF pools requires differentiation among the binding records in the binding SBR. It is possible for the same UE, as indicated by the IMSI, to have multiple active IPcan sessions spread across the different pools.

Note: Although the concept of a PCRF pool is a network-wide concept for a service provider, PCRF pools configuration is done on a PCA site-by-site basis. PCAs in different sites can support different PCRF Pool Selection configurations.

When deploying multiple PCRF pools, each pool supports either different policy-based services or different versions of the same policy based services. Each PCRF pool has a set of DSR PCA peers that are a part of the pool.

On the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page on the NOAM or SOAM, you can perform the following actions:

- Create new PCRF Pools
- Edit existing PCRF Pools
- Delete PCRF Pools
- Identify PCRF Sub-Pools
- Add optional comments for Pools

When a binding-capable session initiation request is received, the Policy DRA uses the following high-level logic to route the request:

- If a binding exists for the IMSI and APN or PCRF Pool, route the request to the bound PCRF.
- Otherwise, distribute the request to a PCRF in the configured PCRF Pool.

When determining if a binding exists, the following logic is used:

- If the IMSI and APN are bound to a PCRF, use that binding.
- Else, if the IMSI and PCRF Pool are bound to a PCRF, create a binding for the APN to the same PCRF as already bound to the PCRF Pool.
- Else, no binding exists for the IMSI and APN or PCRF Pool, so a new binding can be created.

The following table illustrates the major differences between PCRF Pooling and non-pooling functionality.

Table 40: PCRF Pooling Concepts

Concept	Before PCRF Pooling	After PCRF Pooling
PCRF Pools	One PCRF Pool for all APNs.	Up to 7 PCRF Pools selected for new bindings using APN. More than one APN can be mapped to a given PCRF Pool, but a given APN can only be mapped to one PCRF Pool.
Subscriber Bindings	A binding is a simple mapping between an IMSI and a PCRF. Once a binding exists, all sessions for that IMSI are routed to the bound PCRF.	A binding is a mapping from an IMSI and APN to a PCRF, but with the caveat that before a new binding is created, the logic must check for existence of another binding to the same PCRF Pool for the IMSI. If such a binding exists, the new APN is bound to the same PCRF as an existing APN mapped to the same PCRF Pool. Once a binding exists, all sessions for that IMSI and APN are routed to the bound PCRF. Sessions for that IMSI and a different APN mapped to a different PCRF Pool can be routed to a different PCRF.
Number of Sessions per Binding	An IMSI may have up to 10 binding capable sessions.	An IMSI may have up to 10 binding capable sessions, which may be bound to different PCRFs based on APN.
Origin Based Routing	PRT table for new bindings specified in Site Options allows for selection of route list based on origin-host/realm.	After PCRF Pool selection, Sub-Pool rule matching is performed to select a PCRF Sub-Pool given the PCRF Pool and the origin-host of the PCEF.
PRT Table for New Bindings	Each site defines one PRT table to be used for all new bindings.	Each site can define a PRT table to be used for new bindings for each PCRF Pool.

Additionally, Pooling provides the ability to route to subsets of PCRFs in a PCRF Pool on the basis of the Diameter hostname of the PCEF that originated the binding capable session initiation request. These subsets are called PCRF Sub-Pools. This capability allows a controlled amount of policy Diameter signaling to be routed to one or more PCRFs within the PCRF Pool.

The following figure illustrates a sample PCA network configured for PCRF Pooling. The upper third of the figure shows data that is configured with the Policy and Charging GUI at the NOAM server. This data, including PCRF Pools, APN to PCRF Pool mapping, and PCRF Sub-Pool Selection Rules applies to all sites in the Policy DRA network.

The middle third of the figure shows data configured at the SOAM Policy and Charging GUI at each of two PCA sites. This data includes the PCRF Pool to PRT mappings, PCRFs, PRT tables, Route Lists, Route Groups, Peer Nodes, and Connections. This data can differ at each PCA site.

The bottom third of the figure shows the PCRFs logically grouped into PCRF Pools as defined by the network operator.

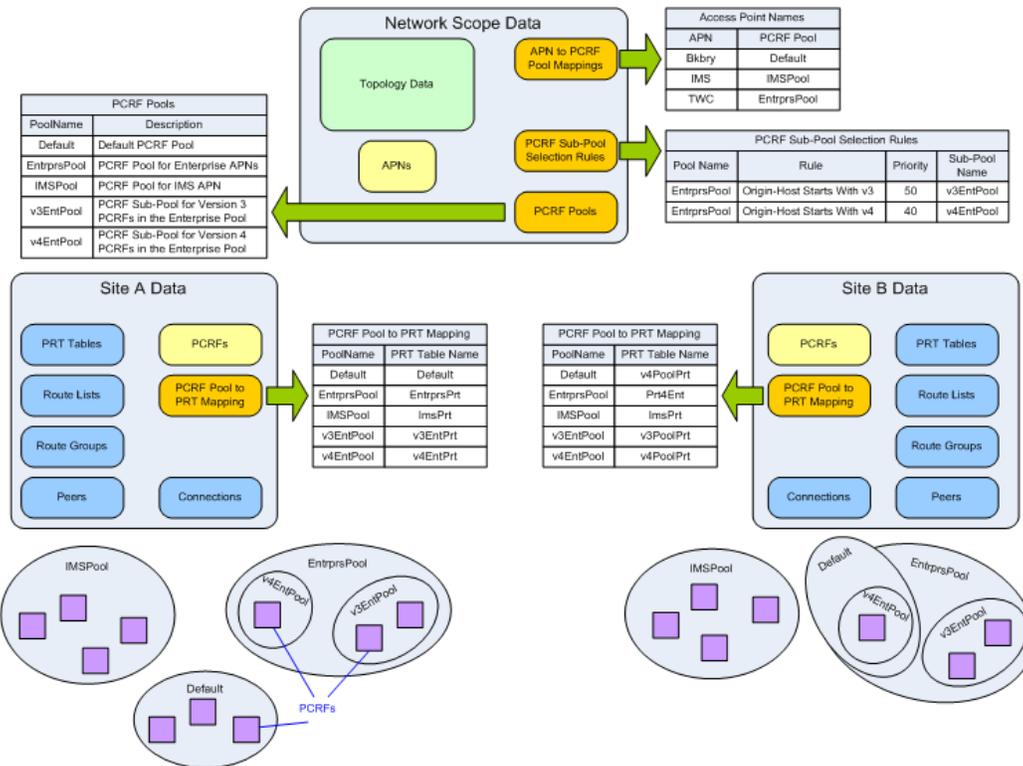


Figure 50: PCRF Pooling Data

Table 41: PCRF Pooling Configuration Summary describes each of the new PCRF Pooling configuration tables, including the order in which they should be configured.

Table 41: PCRF Pooling Configuration Summary

Configuration Order	GUI Page	Purpose
1	PCRF Pools	Define the names of the PCRF Pools and Sub-Pools that are

Policy and Charging Configuration

Configuration Order	GUI Page	Purpose
		needed for grouping PCRFs to handle policy signaling for the various APNs.
2	PCRF Pool to PRT Mapping	<p>At each site, select a PRT table that is used to route binding-capable session initiation requests for new bindings destined for each PCRF Pool. Each PCRF Pool should be configured with a PRT table, unless it is known that the PCRF Pool will never be selected at the site being configured.</p> <p>Note: Before this step can be performed, PRT tables must be defined in the Diameter folder.</p>
3	PCRF Sub-Pool Selection Rules	<p>An optional table. If it is necessary to subdivide a PCRF Pool so that policy requests from a limited number of policy clients (based on Origin-Host) are routed differently, configure appropriate rules in the PCRF Sub-Pool Selection Rules table. During routing, this table is examined after the APN is mapped to a PCRF Pool.</p> <p>If a matching PCRF Sub-Pool Selection Rule exists, the request is routed to the PCRF Sub-Pool. Otherwise, the PCRF Pool selected by the APN mapping is used.</p>
4	Access Point Names	<p>After all Diameter configuration is completed (including PRT Rules, Route Lists, Route Groups, Peer Nodes, and Connections), each APN can be mapped to a PCRF Pool. After an APN is mapped to a PCRF Pool, binding-capable session initiation requests that result in creation of a new binding are routed using the PCRF Pool.</p>

PCRF Pools elements

Table 42: PCRF Pools elements describes the elements on the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.

Note: Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

The PCRF Pools table contains the list of configured PCRF Pools and Sub-Pools settings that you can use when selecting a set of PCRFs to host a new subscriber binding. The PCRF Pool to be used for a given subscriber binding attempt is determined based on the APN-to-PCRF Pool mappings configured in **Policy and Charging > Configuration > Access Point Names** and the PCRF Sub-Pool Selection Rules configured in **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.

Table 42: PCRF Pools elements

Fields (* indicates required field)	Description	Data Input Notes
* PCRF Pool Name	A unique name for the PCRF Pool assigned by the network operator. A PCRF Pool identifies a set of PCRFs that should be used for policy requests from a specified APN. The mapping from APN-to-PCRF Pool is configured from the Policy and Charging -> Configuration -> Access Point Names page.	Format: List Range: 1 to 32 characters, must start with an upper or lower case letter, and can contain digits and underscores; a maximum of 7 PCRF Pool Names can be defined
Sub-Pool	A setting that indicates that the PCRF Pool is to be used as a PCRF Sub-Pool (for example, the target of a PCRF Sub-Pool Selection Rule). Note: If the check box on the PCRF Pools > [Insert] page is not checked, this PCRF Pool is a pool, not a sub-pool.	Format: Check box Range: Yes (Checked for Sub-Pool), No (Unchecked for Sub-Pool) Default: No (Unchecked for Sub-Pool)
Comments	An optional comment to provide more information about the purpose of this PCRF Pool or Sub-Pool.	Format: Text box Range:0-64 characters

Inserting PCRF Pools

Use this task to insert (create new) PCRF Pools.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Pools**. The **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page appears.
2. Click **Insert**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Pools [Insert]** page opens.
3. Enter a unique PCRF Pool Name in the **PCRF Pool Name** field.

4. Check the **Sub-Pool** check box if the PCRF Pool is to be used as a Sub-Pool.

A Sub-Pool is used if policy requests from specified origin-hosts should be routed to a different set of the PCRFs from those in the PCRF Pool selected by the APN. Sub-Pool Selection Rules are configured in **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.

The choices are Default = No (Unchecked for Sub-Pool); the range is Yes (Checked for Sub-Pool) and No (Unchecked for Pool).

5. You can type an optional comment in the **Comments** field to describe the Pool or Sub-Pool. The entry must be characters in the range of 0 to 64, and the default is N/A.
6. Click:
 - **OK** to save the new PCRF Pool name and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.
 - **Apply** to save the new PCRF Pool name and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entered PCRF Pool name is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- Adding the new PCRF Pool would cause the maximum number of PCRF Pools (7) to be exceeded.

Editing PCRF Pools

Use this task to edit PCRF Pools comments. After a PCRF Pool is created, only the comment can be edited, and the Sub-Pool Indicator can only be changed by deleting the PCRF Pool and creating a new one.

Note: The PCRF Pool Name cannot be edited.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Pools**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page appears. The page displays a list of the configured PCRF Pools that are used when a new subscriber binding was created.
2. Select a PCRF Pool Name to edit.
3. Click **Edit**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Pools [Edit]** page appears.
4. Click in the **Comments** field.
5. Edit the **Comments** field for the selected PCRF Pool. The comment must be characters in the range of 0 to 64, and the default is N/A.
6. Click:
 - **OK** to save the change and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.
 - **Apply** to save the change and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page without saving any changes.

If **Apply** or **OK** is clicked and the selected **PCRF Pool Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting PCRF Pools

Use this task to delete a PCRF Pool.

A PCRF Pool can be deleted only if no APN is mapped to that PCRF Pool. A PCRF Sub-Pool can be deleted only if no PCRF Sub-Pool Selection Rule refers to that PCRF Sub-Pool.

If a PCRF Pool or Sub-Pool is successfully deleted from the NOAMP GUI, the entry is internally marked as retired. Retired entries are not displayed on the GUI, but they cannot be removed from the internal tables because that PCRF Pool or Sub-Pool might still be referenced by any of number of bindings. If you add a new PCRF Pool or Sub-Pool with the same name as one that has been retired, the record is reactivated.

When a PCRF Pool or Sub-Pool is deleted (retired), the entry no longer appears on the **PCRF Pool to PRT Mapping** pages at any of the sites.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Pools**. The **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page appears.
2. Select the **PCRF Pool Name** or **PCRF Sub-Pool Name** to be deleted.
3. Click **Delete**.

A window appears to confirm the delete.

4. Click:
 - **OK** to delete the PCRF Pool or PCRF Sub-Pool.
 - **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.

If **OK** is clicked and the selected PCRF Pool or Sub-Pool no longer exists (it was deleted by another user), an error message is displayed, and the PCRF Pools page is refreshed. The row that was selected is no longer displayed in the list.

PCRF Pool to PRT Mapping

In initial DSR release installations, PCRF Pools and PRT tables must be configured as part of configuring the PCA application. For initial installs and upgrades from releases in which PCA was not activated, the Default PCRF Pool is created and mapped to the Not Selected PRT.

When a PCRF Pool or Sub-Pool is added at the NOAMP, the data is replicated on the SOAM servers at each site. When a user opens the **PCRF Pool to PRT Mapping** page, a row is displayed for each configured PCRF Pool or Sub-Pool. If the PCRF Pool or Sub-Pool has already been mapped to a PRT, the mapping is shown. If the PCRF Pool or Sub-Pool has not yet been mapped, the PRT field shows Not Selected in red text.

Note: The screen does not automatically refresh if a new PCRF Pool or Sub-Pool is added at the NOAMP after the PCRF Pool to PRT Mappings screen is displayed at a given site.

In general, every PCRF Pool and Sub-Pool should be mapped to a PRT table, but there is an exception. If the network operator knows that binding-capable session initiation requests will never originate at that site from an APN (and optionally Origin-Host) that is mapped to that PCRF Pool or Sub-Pool.

A PCRF Pool or Sub-Pool that is deleted from the NOAMP GUI is not actually deleted, but rather retired. When a PCRF Pool or Sub-Pool is deleted from the NOAMP GUI, the entry disappears from the **PCRF Pool to PRT Mapping** GUI page at each site (the next time the screen is manually refreshed). If the PCRF Pool or Sub-Pool entry is restored (added again) at the NOAMP, the entry reappears on the **PCRF Pool to PRT Mapping** page, and it will have the same PRT choice as was previously configured, provided the PRT table still exists.

A Peer Route Table cannot be deleted from a site if that Peer Route Table is referenced by a current **PCRF Pool to PRT Mapping** entry. Entries for retired PCRF Pools or Sub-Pools are not included in this restriction. As a result, if a PCRF Pool A had a mapping to PRT table X, then PCRF Pool A was deleted at the NOAMP, it is possible to delete PRT X (provided no other active PCRF Pool to PRT Mappings referenced PRT X). If PCRF Pool A was added back at the NOAM after the deletion of PRT X, PCRF Pool A would appear on the PCRF Pool to PRT Mapping GUI with its PRT entry set to the default of Not Selected.

If a PCRF Pool or Sub-Pool is changed from being mapped to a PRT table to the -Select- value in the PRT pulldown menu, you might see a confirmation window that includes a warning of one of the following conditions applies:

- If an APN is mapped to the PCRF Pool being changed, a confirmation window is displayed on the **PCRF Pool to PRT Mapping** page that warns that this PCRF Pool is being used by one or more APNs. You can choose to continue, but know that a call might fail at that site if a binding-capable session initiation request arrives with an APN that is mapped to that PCRF Pool.
- If the PCRF Pool is included as a Sub-Pool in a PCRF Sub-Pool Rule, a confirmation window is displayed on the **PCRF Pool to PRT Mapping** page that warns that this PCRF Pool is being used by one or more PCRF Sub-Pool Rules. You can choose to continue, but know that a call might fail at that site if a binding-capable session initiation request arrives with an APN and Origin-Host that is mapped to that PCRF Sub-Pool.

PCRF Pools to PRT Mapping elements

Table 43: PCRF Pools to PRT Mapping elements describes the elements on the **Policy and Charging > Configuration > Policy DRA > PCRF Pools to PRT Mapping** page.

Note: Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

The PCRF Pool To PRT Mapping table displays the list of PCRF Pools or Sub-Pools configured at the NOAMP and allows each to be mapped to a Peer Routing Table that is used when a new binding is created for the PCRF Pool. The PCRF Pool or Sub-Pool to be used for a given subscriber binding attempt is determined based on Access point Name to PCRF Pool mappings, or by rules configured at the NOAMP in **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.

Use this table to configure (at each site) the mapping between the selected PCRF Pool or PCRF Sub-Pool and a PRT table that defines the routing for the pool at that site.

Table 43: PCRF Pools to PRT Mapping elements

Field	Description	Data Input Notes
PCRF Pool Name	The name of the PCRF Pool or Sub-Pool that is defined for the network in the PCRF Pools GUI. When a PCRF Pool or PCRF Sub-Pool is configured at the	Format: Text box; string of 1-32 alphanumeric characters, must contain at least one alpha character, must not start with a

Field	Description	Data Input Notes
	NOAMP, it automatically appears on the PCRF Pool to PRT Mappings page so that a PRT can be defined for it if needed. This field is a hyper-link to the PCRF Pools (Filtered) view page, filtered by the PCRF Pool or Sub-Pool name.	digit, and can contain underscores Range: Valid name
Peer Route Table Name	The name of a configured Peer Route Table that should be used to route new binding requests destined to the PCRF Pool or PCRF Sub-Pool. This field is a hyper-link to the Diameter > Configuration > Peer Route Tables view page, filtered by the PRT name.	Format: String Range: All Peer Route Tables configured at this site Default: Not Selected

Editing PCRF Pool to PRT Mapping

Use this task to edit PCRF Pool to PRT Mapping settings.

1. On the Active SOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Pool to PRT Mapping**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Pool to PRT Mapping** page appears. The page displays a list of PCRF Pools or Sub-Pools configured at the NOAMP .
2. Select a row to edit (click in the row, but do not click on a specific element within the row).
DO NOT click the blue PCRF Pool Name or the Peer Route Table Name (except to view the PCRF Pools (Filtered) page or the Peer Routes Table (Filtered) page. The blue color indicates a hyper-link. The PCRF Pool Name hyper-link opens the **Policy and Charging > Configuration > Policy DRA > PCRF Pools (Filtered)** page and the Peer Route Table hyper-link opens the **Diameter > Configuration > Peer Routes Table (Filtered)** page.
If the PCRF Pool has NOT been assigned a Peer Route Table record, Not Selected is displayed in red in the **Peer Route Table Name** column. This helps to inform the SOAM user that the PCRF Pool should be mapped to a Peer Route Table.
3. (optional) Click **Pause updates** to suppress the automatic page refresh function. The default is Unchecked.
Pause updating applies to all rows on the screen. If you add a new PCRF Pool at the NOAMP, a new row automatically appears on the SOAM **PCRF Pool to PRT Mapping** page the next time an update occurs.
4. Click **Edit**. The **PCRF Pool To PRT Mapping [Edit]** page is displayed.
The **Peer Route Table Name** pulldown menu initially displays the Peer Route Table from the row being edited and contains all configured Peer Route Tables and Not Selected. Not Selected provides backwards compatibility for users who had the **Site Options Peer Route Table Name** set to Not Selected. When Not Selected is chosen, PCA does not instruct DRL to use an application specified PRT, but enables DRL use its normal PRT precedence for PRT selection instead. If **Edit** is clicked

and the PCRF Pool Name of the selected row has been deleted, an error is displayed and this row is no longer displayed. If **Edit** is clicked and the PCRF Pool Name of the selected row still exists (has not been retired), the **PCRF Pool To PRT Mapping [Edit]** page is displayed with data populated from the selected row.

5. Select an item from the **Peer Route Table Name** pulldown menu. The default is Not Selected, and the range is All Peer Route Tables configured at this site.
6. Click:
 - **OK** to save the selection and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pool to PRT Mapping** page.
 - **Apply** to save the selection and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pool to PRT Mapping** page without saving any changes.

Additionally, the following can occur as a result of clicking **Ok** or **Apply**:

- If the selected PCRF Pool Name or the Peer Route Table Name entry no longer exists (it was deleted by another user from the NOAMP), an error message is displayed on the **PCRF Pool To PRT Mapping [Edit]** page and no changes are made to the database.
- If all the data syntax validation as per each field's description does not meet requirements, an error message is displayed.
- If the PRT selection has changed from a PRT name to Not Selected and the corresponding PCRF Pool is mapped to an APN, a confirmation message is displayed with the text: "PCRF Pool <PCRF Pool Name> is currently used for bindings originating from at least one APN. Changing the PRT entry to 'Not Selected' may cause these bindings to fail if originated at this site. Click Ok to continue or Cancel to return to the PCRF Pool To PRT Mapping screen."
- If the PRT selection has changed from a PRT name to Not Selected and the corresponding PCRF Pool is specified as the PCRF Sub-Pool in a PCRF Sub-Pool Selection Rule, a confirmation dialog is displayed with the text: "PCRF Pool <PCRF Sub-Pool Name> is currently used for bindings that match PCRF Sub-Pool Selection Rule <PCRF Sub-Pool Selection Rule Name>. Changing the PRT entry to 'Not Selected' may cause these bindings to fail if originated at this site. Click Ok to continue or Cancel to return to the PCRF Pool To PRT Mapping screen."

Pausing Updates to PCRF Pool to PRT Mapping

Use this task to pause updates to PCRF Pool to PRT Mapping.

The **PCRF Pool To PRT Mapping** page is automatically refreshed every *N* seconds to show the latest PCRF Pools configured at the NOAMP **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.

Pausing update applies to all rows in the table on the **Policy and Charging > Configuration > PCRF Pool to PRT Mapping** page. Selecting this check box pause the automatic update function for all items in the table.

1. On the Active SOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Pool to PRT Mapping**.

The **Policy and Charging > Configuration > PCA DRA > PCRF Pool to PRT Mapping** page appears. The page displays a list of the configured PCRF Pool Names and corresponding Peer Route Table Names.

2. (optional) Click **Pause updates** to suppress the automatic page refresh function. The default is Unchecked. This function remains in effect until the **Pause updates** check box is unchecked.

Pause updating applies to all rows on the screen. If you add a new PCRF Pool at the NOAMP, a new row automatically appears on the SOAM **PCRF Pool to PRT Mapping** page the next time an update occurs.

PCRF Sub-Pool Selection Rules

The PCRF Sub-Pool Selection table contains rules for selection of a PCRF Sub-Pool for a given PCRF Pool and Origin-Host value.

It is sometimes necessary to subdivide a PCRF Pool into sub-pools; for example, to support controlled routing of traffic to a new PCRF. In such a case, you can configure PCRF Sub-Pool Selection Rules to a selected a sub-pool on the basis of the Origin-Host of the binding capable session initiation request.

A PCRF Sub-Pool Selection Rule has the following attributes:

- The Default PCRF Pool can have sub-pools.
- The **PCRF Pool Name** column contains hyperlinks to the **PCRF Pools** page filtered by the PCRF Pool Name.
- Origin-Host is the only supported PCRF Sub-Pool Selection parameter.
- Supported Origin-Host operators are: Equals, Starts With, and Ends With.
- Priority values can range from 1 to 99, with 1 being the highest priority.

An APN-to-PCRF Pool mapping specifies that all binding-capable session initiation requests that result in creation of a new binding should be routed to a PCRF in PCRF Pool 'X'.

A PCRF Sub-Pool Selection Rule can override the APN-to-PCRF Pool mapping by specifying binding-capable session initiation requests that result in new bindings that were destined for PCRF Pool 'X', but come from PCEF 'Y', should be routed to a PCRF in PCRF Sub-Pool 'Z'.

A PCRF Sub-Pool Selection Rule will never be considered if no APN is mapped to its PCRF Pool. As a result, it is safe to add PCRF Sub-Pool Selection Rules prior to mapping APNs to the PCRF Pool that is being subdivided. It is also acceptable to add PCRF Sub-Pool Selection Rules for a PCRF Pool that is already mapped to an APN. However, if this is done, bindings that were created prior to the existence of the PCRF Sub-Pool Selection Rule take precedence over the PCRF Sub-Pool chosen for new binding-capable session initiation requests that arrive after the new rule is in place. This behavior is necessary to prevent split bindings.

PCRF Sub-Pool Selection Rules are configured using the NOAMP GUI as a network-wide managed object.

The creation of a new PCRF Sub-Pool Selection Rule does not affect P-DRA signaling in any way until both of the following conditions exist:

- An APN is mapped to the PCRF Pool using the Access Point Names GUI
- A binding-capable session initiation request arrives with an APN mapped to that PCRF Pool and an Origin-Host that matches the Condition specified in the PCRF Sub-Pool Selection Rule.

When a PCRF Sub-Pool Selection Rule entry is added, new bindings from that APN and Origin-Host will be routed to a PCRF in the specified PCRF Sub-Pool. When a PCRF Sub-Pool Selection Rule is mapped to a PCRF Sub-Pool, a check is performed to determine if the selected PCRF Sub-Pool is configured with a PRT mapping at each site. If at least one site does not have a mapping for the selected PCRF Sub-Pool, a confirmation dialog is displayed that including a warning as follows:

- If a site does not have the PCRF Sub-Pool mapped to a PRT table, a confirmation dialog is displayed on the APN GUI warning that Site 'X' does not have a mapping defined for this PCRF Sub-Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a

binding-capable session initiation request arrives with an APN and Origin-Host that is mapped to that PCRF Sub-Pool.

- If a site cannot be reached due to network errors, a confirmation dialog is displayed on to warn you that it cannot be determined whether Site 'X' has a mapping defined for this PCRF Sub-Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN and Origin-Host that is mapped to that PCRF Pool.

The PCRF Sub-Pool Selection Rule GUI prevents creation of rules that are:

- Ambiguous
- Conflicting
- Duplicate

Two rules are considered as ambiguous if the following criteria are met:

- The rules have the same PCRF Pool values and
- The rules have the same Priority values and
- The rules have different PCRF Sub-Pool values and one of the following is true:
 - One rule has an Origin-Host with a "Starts With" operator and the other rule has an Origin-Host with an "Ends With" operator -- OR –
 - For example, starts With ab and Ends With xyz
 - Value length is not considered as a factor in the best match decision at this time.
 - Both rules have an Origin-Host with a "Starts With" operator and all of the value characters of the shorter value match the first characters of the longer value -- OR –
 - For example, starts With abc and Starts With ab
 - Both rules have an Origin-Host with a "Ends With" operator and all of the value characters of the shorter value match the last characters of the longer value.
 - For examples, ends With xyz and Ends With yz

Two rules are considered to be conflicting if all of the following criteria are met:

- The rules have the same PCRF Pool values.
- The rules have the same Priority values.
- The rules have the same Origin-Host operators and values.
- The rules have different PCRF Sub-Pool values.

Two rules are considered to be duplicate if all of the following criteria are met:

- The rules have the same PCRF Pool values.
- The rules have the same Origin-Host operators and values.
- The rules have the same PCRF Sub-Pool values.

PCRF Sub-Pool Selection Rules elements

[Table 44: PCRF Sub-Pool Selection Rules elements](#) describes the elements on the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page.

Table 44: PCRF Sub-Pool Selection Rules elements

Fields (* indicates required field)	Description	Data Input Notes
PCRF Sub-Pool Selection Rule Name	A unique name for the PCRF Sub-Pool Selection Rule assigned by the network operator.	Format: Text box; string 1-32 characters, must start with an upper or lower case letter, and can contain digits and underscores; maximum number of Sub-Pool Selection Rules is of 70 Range: Valid name
Priority	A priority value. The priority value is used to break ties when more than one PCRF Sub-Pool Selection Rule matches a given binding-capable session initiation request. Multiple rules can match a request when more than one rule using a "Starts With" or "Ends With" condition exists.	Format: Text box Range: 1-99, inclusive, where a lower value equates to a higher priority Default: 50
PCRF Pool Name	The PCRF Pool that is being subdivided by this PCRF Sub-Pool Selection Rule. A pulldown menu contains the names of all available PCRF Pools. The PCRF Pool does not need to have any APN mapped to it when this PCRF Sub-Pool Selection Rule is created. This field is a hyper-link to the PCRF Pools view screen, filtered by the PCRF Pool name.	Format: Dropdown menu Range: Configured PCRF Pools that have not been specified as PCRF Sub-Pool Names
Conditions	A condition allows for configuration of a value to be compared to a given Diameter AVP using the specified operator. The only condition currently supported for PCRF Sub-Pool Selection Rules is for the Origin-Host AVP. The value field allows you to enter a string to be compared to the	Format: Textbox Range: Equals, Starts With, and Ends With.

Fields (* indicates required field)	Description	Data Input Notes
	Origin-Host using the operator.	
PCRF Sub-Pool Name	The PCRF Sub-Pool name that is used for routing new bindings created from binding-capable session initiation requests that matched this PCRF Sub-Pool Selection Rule. A match occurs when the APN in the request mapped to the PCRF Pool in the rule AND the Origin-Host condition matched. This field is a hyper-link to the PCRF Pools view screen, filtered by the PCRF Sub-Pool name.	Format: Hyperlink Range: Assigned PCRF Sub-Pool Name
Last Updated	The PCRF Sub-Pool Selection Rules view page also includes a timestamp of the time the rule was created or last updated, whichever occurred most recently. This field can help you troubleshoot by allowing comparison of existing binding session creation time stamps (displayed using the binding key query tool) with rule creation time stamps. Use this capability to determine whether a binding was created before or after a rule was created	Format: Read-only field Range: N/A

Inserting PCRF Sub-Pool Selection Rules

Use this task to insert (create new) PCRF Sub-Pool Selection Rules.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page appears.
2. Click **Insert**.

The **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules [Insert]** page opens.

3. Enter a unique PCRF Sub-Pool Selection Rules Name in the **PCRF Pool Selection Rule Name** field. Enter a unique name that identifies the PCRF Sub-Pool Selection Rule. The default is N/A, and the range is a 32-character string. Valid characters are alphanumeric and underscore, and must contain at least one alpha character and must not start with a digit..
4. Enter a priority value for this rule in **Priority**. The lower the value means the higher the priority. The default is 50, and the range is 1 to 99.
5. Select a PCRF Pool Name from the **PCRF Pool Name** pulldown menu. This is the name of the PCRF Pool for which a Sub-Pool is being defined The default is N/A, and the range is Configured PCRF Pools that have not been specified as PCRF Sub-Pool Names.
6. Select a condition from the **Operator** pulldown menu to associate the selected condition with this rule.
The range is Equals, Starts With, or Ends With.
FQDN is a case-insensitive string consisting of a list of labels separated by dots, where a label can contain alphanumeric characters, dashes, underscores. A label must start with a letter, digit or underscore and must end with a letter or digit. Underscores can be used as the first character only. A label range is 1 to 64, and an FQDN range is 1 to 255 characters in length. The default is N/A, and the range is Substring or complete string of a valid FQDN.
7. Enter a value in the **Value** field.
8. Select a PCRF Sub-Pool Name in the **PCRF Sub-Pool Name** pulldown menu. Choices include all the qualified PCRF Sub-Pool configured from the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page. A qualified PCRF Sub-Pool is a PCRF Pool that is non-retired and has been marked as Sub-Pool. A retired PCRF Sub-Pool entry can be created by first adding a new PCRF Sub-Pool and then deleting it.
This is the PCRF Sub-Pool that is to be used for Gx and Gxx session initiation request messages that match this Rule. The default is N/A and the range is the choice of configured PCRF Pools.
9. The **Last Updated** field is a read-only field that displays the date and time that this rule was created, or the last time the rule was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding-capable session initiation requests. This date and time can be compared against binding creation times when troubleshooting using the Binding Key Query Tool.
10. Click:
 - **OK** to save the new PCRF Pool name and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Pools** page.
 - **Apply** to save the new PCRF Sub-Pool Selection Rule and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rule** page without saving any changes.

Editing PCRF Sub-Pool Selection Rules

Use this task to edit PCRF Sub-Pool Selection Rules.

The PCRF Sub-Pool Selection Rule edit page allows a network operator to change all fields except the PCRF Sub-Pool Selection Rule Name. Changes take effect on the next binding-capable session initiation request received after the rule is successfully committed.

1. On the Active NOAM, select **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.

The **Policy and Charging > Configuration > Policy DRA > PCRF PSub-Pool Selection Rules** page appears. The PCRF Sub-Pool Selection table contains rules for selection of a PCRF Sub-Pool for a given PCRF Pool and Origin-Host value.

2. Select a PCRF Sub-Pool Selection Rule to edit.

DO NOT click the blue PCRF Pool Name or the PCRF Sub-Pool Name (except to see the configuration of the PCRF Pool Name or PCRF Sub-Pool Name). The blue color indicates a hyper-link that opens the **Diameter > Configuration > Peer Nodes [Filtered]** page to display the configuration information for the Peer Node.

3. Click **Edit**.

The **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pools Selection Rules [Edit]** page appears. You cannot edit the **PCRF Sub-Pool Selection Rule** value. This is a name that uniquely identifies the PCRF Sub-Pool Selection Rule. The default is N/A, and the range is a 32-character string. Valid characters are alphanumeric and underscore, and must contain at least one alpha character and must not start with a digit.

4. Enter a priority value for this rule in **Priority**.

The lower the value means the higher the priority. The default is 50, and the range is 1 to 99.

5. Enter a PCRF Pool Name.

The name of the PCRF Sub-Pool Selection Rules for which a Sub-Pool is being defined The default is N/A, and the range is Configured PCRF Sub-Pool Selection Rules that have not been specified as PCRF Sub-Pool Names.

6. Specify the condition associated with this rule.

Select a Host-Origin Operator value from the pulldown menu. FQDN is a case-insensitive string consisting of a list of labels separated by dots, where a label can contain letters, digits, dashes ('-') and underscores ('_'). A label must start with a letter, digit, or underscore, and it must end with a letter or digit. Underscores can be used as the first character only. A label cannot exceed 63 characters in length and an FQDN cannot exceed 255 characters in length. The default is N/A, and the range is a substring or complete string of a valid FQDN.

7. Select a PCRF Sub-Pool Name value from the pulldown menu.

This PCRF Sub-Pool that will be used for Gx and Gxx session initiation request messages matching this Rule The default is N/A, and the range is the choice of configured PCRF Sub-Pool Selection Rules.

8. **Last Updated** is a read-only field that displays the date and time that this rule was created, or the last time the rule was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding capable session initiation requests. This date and time can be compared against binding creation times when troubleshooting using the Binding Key Query Tool.

9. Click:

- **Ok** to save the change and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page.
- **Apply** to save the change and remain on this page.
- **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > PCRF PCRF Sub-Pool Selection Rules** page without saving any changes.

If **Apply** or **OK** is clicked and the selected **PCRF Peer Node Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting PCRF Sub-Pool Selection Rules

Use the following procedure to delete a PCRF.

A PCRF Sub-Pool Selection Rule can be deleted at any time.

1. Select **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules**.
The **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page appears.
2. Select the **PCRF Sub-Pool Selection Rule Name** to be deleted.
3. Click **Delete**.
A popup window appears to confirm the delete.
4. Click:
 - **OK** to delete the PCRF Sub-Pool Selection Rule Name.
 - **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > Policy DRA > PCRF Sub-Pool Selection Rules** page.

If **OK** is clicked and the selected PCRF no longer exists (it was deleted by another user), an error message is displayed and the PCRF Sub-Pool Selection Rules page is refreshed. The row that was selected is no longer displayed in the list.

Policy Clients

Topology hiding configuration is performed at both the network level using the NOAM GUI and at the site level using the SOAM GUI.

The fields are described in [Policy Clients elements](#).

SOAM Options

On the SOAM GUI, use the **Policy and Charging > Configuration > Policy DRA > Policy Clients** to define the list of Policy Client Peer Nodes from which the PCRF name is to be hidden. This page can be used only if Topology Hiding is **Enabled** and the **Topology Hiding Scope** option is either **Specific Clients** or **All Foreign Realms + Specific Clients** on the **Policy and Charging > Policy DRA > Configuration > Network-Wide Options** page on the NOAM GUI. See [Site Options](#) for additional information.

- Filter the list of Policy Client Peer Node Names, to display only the desired Policy Client Peer Node Names.
- Sort the list entries in ascending or descending order by Policy Client Peer Node Names or by Comments, by clicking the column heading. By default, the list is sorted by Policy Client Peer Node Names in ascending numerical order.
- Click the **Insert** button.

The **Policy and Charging > Configuration > Policy DRA > Policy Clients [Insert]** page opens. You can add a Policy Client Peer Node Name and Comment. See [Adding a new Policy Client for Topology Hiding](#). If the maximum number of Policy Client Peer Nodes (1000) already exists in the system, the **Policy and Charging > Configuration > Policy DRA > Policy Clients [Insert]** page will not open, and an error message is displayed.

- Select the **Comment** cell in the row for a Policy Client Peer Node Name in the list, and click the **Edit** button. (Clicking the blue **Policy Client Peer Node Name** will open the filtered **Diameter > Configuration > Peer Nodes** page for the Peer Node.)

The **Policy and Charging > Configuration > Policy DRA > Policy Clients [Edit]** page opens. You can edit the **Comment** for the selected **Policy Client Peer Node Name**. (The Policy Client Peer Node Name cannot be changed). See [Editing Policy Clients for Topology Hiding](#).

- Select the **Comment** in the row for a Policy Client Peer Node Name in the list, and click the **Delete** button to remove the selected **Policy Client Peer Node Name**. See [Deleting a Topology Hiding Policy Client Peer Node](#).

Policy Clients elements

[Table 45: Policy Clients elements](#) describes the elements on the **Policy and Charging > Configuration > Policy DRA > Policy Clients** page. Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 45: Policy Clients elements

Elements	Description	Data Input Notes
Policy Client Peer Node Name	The name of a configured Diameter Peer Node that identifies a Policy Client Peer Node. Selecting a Policy Client Peer Node name (blue hyperlink) displays the Diameter > Configuration > Peer Nodes (Filtered) page where Diameter Peer Nodes are filtered by the Policy Client Peer Node Name.	Format: Pulldown list Note: The Policy Client Peer Node Name cannot be changed on the [Edit] page. Range: Configured Diameter Peer Nodes
Topology Hiding Enabled	A read-only check box with default 'checked' to indicate the Topology Hiding for the policy client peer node being enabled. It is the only option currently supported.	Format: Check box Range: N/A (Read Only)
Comments	An optional comment that describes the Policy Client Peer Node.	Format: Text box Range 0-64 characters

Viewing Policy Clients

Use this task to view all configured Policy Client Peer Nodes on the SOAM.

Select **Policy and Charging > Configuration > Policy DRA > Policy Clients**.

The **Policy and Charging > Configuration > Policy DRA > Policy Clients** page appears.

The fields are described in [Policy Clients elements](#).

Adding a new Policy Client for Topology Hiding

Use this task to add a new Policy Client for Topology Hiding.

Note: Topology Hiding is performed only if it is Enabled and the Topology Hiding **Scope** option is defined as **Specific Clients** or **All Foreign Realms + Specific Clients** in the **Policy and Charging > Configuration Policy DRA > Network-Wide Options** page on the NOAM.

The fields are described in [Policy Clients elements](#).

1. On the Active SOAM, select **Policy and Charging > Configuration > Policy DRA > Policy Clients**.
The **Policy and Charging > Configuration > Policy DRA > Policy Clients** page appears.
2. Click **Insert**.
The **Policy and Charging > Configuration > Policy DRA > Policy Clients [Insert]** page appears.
3. Select a Policy Client Peer Node Name from the **Value** pulldown list.
4. Check **Topology Hiding Enabled** if Topology hiding is needed for the Policy Client.
5. Enter an optional comment in the **Comments** field.
6. Click:
 - **OK** to save the changes and return to the **Policy and Charging > Configuration > Policy DRA > Policy Clients** page.
 - **Apply** to save the changes and remain on this page.
 - **Cancel** to return to the Policy DRA **Policy and Charging > Configuration > Policy DRA > Policy Clients** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entered comment exceeds 64 characters in length or contains something other than 7-bit ASCII characters.
- The **Policy Client Peer Node Name** is missing.
- The selected **Policy Peer Node Name** is already configured in the system.
- Any fields contain invalid input (for example, the wrong type of data was entered or a value is out of range).
- The maximum number (1000) of **Topology Hiding** records has already been configured.

Editing Policy Clients for Topology Hiding

Use this task to edit a Policy Client for Topology Hiding.

Note: Topology Hiding is performed only if it is Enabled and the Topology Hiding **Scope** option is defined as **Specific Clients** or **All Foreign Realms + Specific Clients** in the **Policy and Charging > Configuration Policy DRA > Network-Wide Options** page on the NOAM.

The fields are described in [Policy Clients elements](#).

1. On the Active SOAM, select **Policy and Charging > Configuration > Policy DRA > Policy Clients**.
The **Policy and Charging > Configuration > Policy DRA > Policy Clients** page appears.
2. Click **Edit**.
The **Policy and Charging > Configuration > Policy DRA > Policy Clients [Edit]** page appears.
A read-only value is displayed in the Policy Client Peer Node Name **Value** field.
3. Check **Topology Hiding Enabled** if Topology hiding is needed for the Policy Client.
4. Edit or enter an optional comment in the **Comments** field.
5. Click:
 - **OK** to save the edited Comment and return to the **Policy and Charging > Configuration > Policy DRA > Policy Clients** page.
 - **Apply** to save the edited Comment and remain on this page.

- **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > Policy Clients** page without saving any changes.

If **OK** or **Apply** is clicked and the following condition exists, an error message appears:

- The selected Policy Client Code Name no longer exists (for example, it has been deleted by another user), and no changes are made to the database.

Deleting a Topology Hiding Policy Client Peer Node

Use the following procedure to delete a Topology Hiding Policy Client Peer Node.

1. On the Active SOAM, select **Policy and Charging > Configuration > Policy DRA > Policy Clients**. The **Policy and Charging > Configuration > Policy DRA > Policy Clients** page appears.
2. Select the Comment in the line for a Policy Client Peer Node Name to be deleted. (Clicking the blue Policy Client Peer Node Name will open the filtered **Diameter > Configuration > Peer Nodes** page for the Peer Node.)
3. Click the **Delete** button.

A popup window appears to confirm the delete.

4. Click:
 - **OK** to delete the Policy Client Peer Node Name.
 - Click **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > Policy DRA > Topology Hiding** page.

If **OK** is clicked and the selected Policy Client Peer Node no longer exists (it was deleted by another user), an error message is displayed and the **Policy and Charging > Configuration > Policy DRA > Policy Clients** page is refreshed. The row that was selected is no longer displayed in the list.

Suspect Binding Removal Rules

On the **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules** page on the SO GUI, Suspect Binding Removal Rules are listed. During Diameter messaging, when a Suspect Binding Removal Rule is matched, the rule configuration determines if the Subscriber's bindings should be considered suspect or be removed immediately.

When configuring Suspect Binding Removal Rules:

- The (Experimental) Result Code attribute is applicable only if the Error Scenario Category is set to External Result.
- Vendor ID is applicable only if the Result Code is an experimental result code.
- If Remove Suspect Binding Immediately is set to Yes, a single match on this rule will cause the removal of all sessions for the IMSI and PCRF.
- If Remove Suspect Binding Immediately is set to No, a rule match causes all sessions for the IMSI and PCRF to be considered suspect and the number of rule matches is incremented.

Note: If the number of rule matches reaches or exceeds the configured in the Suspect Binding Removal Events Threshold field on the Policy and Charging > Configuration > Policy, all sessions for the IMSI and PCRF will be removed.

The fields are described in [Suspect Binding Removal Rules elements](#).

Suspect Binding Removal Rules elements

Table 46: *Suspect Binding Removal Rules elements* describes the elements on the **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules** page on the SO GUI.

Table 46: Suspect Binding Removal Rules elements

Field (* indicates a required field)	Description	Data Input Notes
Rule Name*	A name that uniquely identifies the Suspect Binding Removal Rule	Format: Text box Default: N/A Range: A 32-character string. Valid characters are alphanumeric and underscore. Must contain at least one alpha and must not start with a digit.
Application Name*	The Diameter Application name and Id to which this Suspect Binding Removal Rule applies. Session initiation answer messages including the Application-Id are candidates to match this rule.	Format: Pull-down list Default: N/A Range: Supported P-DRA Application-Ids
Command Code*	The Diameter Command Code or Extended Command Code name and value to which this Suspect Binding Removal Rule applies. Session initiation answers messages including this Command Code are candidates to match this rule.	Format: Pull-down list Default: N/A Range: Supported P-DRA session initiation answer messages
Error Scenario Category*	The error category to which the Suspect Binding Removal Rule applies. Category "Unable to Route" is for when no session initiation answer is received from the PCRF (possibly because the request could not be routed). If "Unable to Route" is chosen, the (Experimental) Result Code sent to the policy clients is the one configured on the Policy and Charging > Configuration > Error Codes screen for the specific interface.	Format: Pull-down list Default: N/A Range: External Result, Unable to Route

Field (* indicates a required field)	Description	Data Input Notes
	Category "External Result" is for when a specified session initiation error answer is received from the PCRF. If "External Result" is chosen, a Result Code must be specified, otherwise no Result Code is necessary.	
Result Code	The session initiation error answer (Experimental) Result Code to which this Suspect Binding Removal rule applies if the Error Scenario Category is "External Result". This field is not applicable when Error Scenario Category is set to "Unable to Route".	Format: Text box Default: N/A Range: 1-9999
Vendor ID	If a Result Code is entered in the Result Code field and that Result Code is an experimental result code, enter a Vendor-Id in this field. Otherwise, leave the field blank.	Format: Text box Default: N/A Range: 1-4294967295
Remove Suspect Binding Immediately	Check this box if a single occurrence of this rule match means that the binding should be removed. Uncheck this box if multiple occurrences of this rule match are required before the binding should be removed. If this box is unchecked , the Suspect Binding Removal Events Threshold field in Policy and Charging -> Configuration -> Policy DRA -> Network-Wide Options on the NOAM controls how many Suspect Binding Removal Events must occur before a Session-Release RAR will be sent to the policy client to request removal of the binding.	Format: Check box Default: No (Unchecked) Range: Yes (Checked), No (Unchecked)
Comments	An optional comment to describe this suspect binding removal rule	Format: Text box Default: N/A

Field (* indicates a required field)	Description	Data Input Notes
		Range: 0-64 characters

Viewing Suspect Binding Removal Rules

Use this task to view all configured Suspect Binding Removal Rules on the SO GUI.

The fields are described in [Suspect Binding Removal Rules elements](#).

Select **Policy and Charging > Configuration > Suspect Binding Removal Rules**.

The **Policy and Charging > Configuration > Suspect Binding Removal Rules** page appears.

The fields are described in LINK.

Inserting Suspect Binding Removal Rules

Use this task to insert a Suspect Binding Removal Rule.

The fields are described in [Suspect Binding Removal Rules elements](#).

1. Select the **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules**. The **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules** page appears.
2. Click **Insert**.
3. Enter a unique Rule Name in the **Rule Name** field.
4. Select an **Application Name** from the pull-down list.
5. Select a **Command Code** from the pull-down list.
6. Select an **Error Scenario Category** from the pull-down list.
7. Enter a **Result Code** unless **Error Scenario Category** is "Unable to Route".
8. Enter an optional Vendor ID in the **Vendor ID** field only if a Result Code is found in the **Result Code** field.
9. Check or uncheck the **Remove Suspect Binding Immediately** box as needed.
10. Add an optional comment in the **Comments** box.
11. Click:
 - **Ok** to save the new Suspect Binding Removal Rule to the **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules** page.
 - **Apply** to save the new Suspect Binding Removal Rule and remain on the page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules** page without any changes.

If **Ok** or **Apply** is clicked and any of the following conditions exist, an error message appear:

- A Field value is missing
- The syntax is invalid
- A Field value must be unique
- The maximum number of suspect binding removal rules allowed to be configured
- An operation failed because the suspect binding removal rule already exists
- An operation failed because the new suspect binding removal rule conflicts with another existing suspect binding removal rule

- An operation failed because inconsistent field values are configured
- An operation failed because the Rule Name can't be "Total" because "Total" is a reserved Rule Name

Editing Suspect Binding Removal Rules

Use this task to edit a Suspect Binding Removal Rule.

The fields are described in [Suspect Binding Removal Rules elements](#).

1. Select the **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules**. The **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules** page appears.
2. Click **Edit**.
3. Select an **Application Name** from the pull-down list.
4. Select a **Command Code** from the pull-down list.
5. Select an **Error Scenario Category** from the pull-down list.
6. Enter a **Result Code** unless **Error Scenario Category** is "Unable to Route".
7. Enter an optional Vendor ID in the **Vendor ID** field only if a Result Code is found in the **Result Code** field.
8. Check or uncheck the **Remove Suspect Binding Immediately** box as needed.
9. Add an optional comment in the **Comments** box.
10. Click:
 - **Ok** to save the edited Suspect Binding Removal Rule to the **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules** page.
 - **Apply** to save the edited Suspect Binding Removal Rule and remain on the page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Policy DRA > Suspect Binding Removal Rules** page without any changes.

If **Ok** or **Apply** is clicked and any of the following conditions exist, an error message appear:

- A Field value is missing
- The syntax is invalid
- A Field value must be unique
- The maximum number of suspect binding removal rules allowed to be configured
- An operation failed because the suspect binding removal rule already exists
- An operation failed because the edited suspect binding removal rule conflicts with another existing suspect binding removal rule
- An operation failed because inconsistent field values are configured
- An operation failed because the Rule Name can't be "Total" because "Total" is a reserved Rule Name

Deleting a Suspect Binding Removal Rule

Use this task to delete a Suspect Binding Removal Rule.

1. Select **Policy and Charging > Configuration > Suspect Binding Removal Rules**. The **Policy and Charging > Configuration > Suspect Binding Removal Rules** page appears.
2. Select the **Suspect Binding Removal Rule** to be deleted.
3. Click the **Delete** button.

A popup window appears to confirm the delete.

4. Click:

- **Ok** to delete the Suspect Binding Removal Rule.
- **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > Suspect Binding Removal Rules** page.

Site Options

The Policy DRA Site Options apply independently to each Policy DRA site. The following Site Options can be configured on **Policy and Charging > Configuration > Policy DRA > Site Options** page on the active SOAM server:

- Topology Hiding Virtual Name - FQDN and Realm. See [Site Options elements](#)
- Peer Route Table Name - The name of the Peer Route Table to be used for routing new binding requests. This entry is no longer used once PCRF Pooling is Enabled.

The fields are described in [Site Options elements](#).

Site Options elements

[Table 47: Site Options elements](#) describes the elements on the SOAM **Policy and Charging > Configuration > Policy DRA > Site Options** page. Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 47: Site Options elements

Field (* indicates field is required)	Descriptions	Data Input Notes
Topology Hiding Virtual Name	<p>FQDN</p> <p>Value used to populate the Diameter Origin-Host AVP for Answer messages routed from a PCRF to a Policy Client, or the Diameter Destination-Host AVP for Request messages routed from a PCRF to a Policy Client.</p>	<p>Format: Text box</p> <p>Range: 1 - 255 characters. Valid characters are letters, digits, dots (.), and hyphens (-). At least one alpha character is required.</p>
	<p>Realm</p> <p>Value used to populate the Origin-Realm AVP for Answer messages routed from a PCRF to a policy client, or the Diameter Destination-Realm AVP for Request messages routed from a PCRF to a Policy Client.</p>	<p>Format: Text box</p> <p>Range: 1 - 255 characters. Valid characters are letters, digits, dots (.), and hyphens (-). At least one alpha character is required.</p>

Field (* indicates field is required)	Descriptions	Data Input Notes
Peer Route Table Name	<p>The name of the Diameter Peer Route Table to be used for routing new binding requests.</p> <p>The Default PRT is always available, but must be selected from the list to be used.</p>	<p>Format: Pulldown list</p> <p>Range: Not Selected, Default, configured Diameter Peer Route Tables</p> <p>Default: Not Selected</p>

Viewing Site Options

Use this task to view all configured Site Options on an SOAM.

Select **Policy and Charging > Configuration > Policy DRA > Site Options**.

The **Policy and Charging > Configuration > Policy DRA > Site Options** page appears with a list of configured Site Options.

The fields are described in [Site Options elements](#).

Setting Site Options

Use this task to set Site Options on the Active SOAM server.

The fields are described in [Site Options elements](#).

1. Select **Policy and Charging > Configuration > Policy DRA > Site Options**.

The **Policy and Charging > Configuration > Policy DRA > Site Options** page appears.

2. Enter an **FDQN** and **Realm**.

Note: If no values are configured here when Topology Hiding is enabled, the FQDN and Realm values of the Default Topology Hiding Virtual Name configured in NOAM GUI Main Menu: **Policy and Charging -> Configuration -> Policy DRA -> Network-Wide Options** will be used.

3. Select a Peer Route Table from the **Peer Route Table Name** pulldown list.

This entry is no longer used once PCRF Pooling is Enabled.

4. Click:

- **Apply** to save the changes and refresh this page.
- **Cancel** to discard the changes and remain on the **Policy and Charging > Configuration > Policy DRA > Site Options** page.

If **Apply** is clicked and any entered value contains the wrong data type or is out of the allowed range, an error message appears.

Online Charging DRA

This section describes the **Policy and Charging > Configuration > Online Charging DRA** GUI pages on the SOAM.

OCSs

On an Active SOAM, the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page lists the Online Charging Servers (OCS) Peer Nodes configured on a site.

The list of OCS Peer Nodes is updated by inserting, editing, or deleting an OCS Peer Node from the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page at each site's SOAM.

OCSs elements

Table 48: OCSs elements describes the elements on the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page on the Active SOAM.

Note: Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 48: OCSs elements

Fields (* indicates required field)	Description	Data Input Notes
* OCS Peer Node Name	The name of a configured Diameter Peer Node that identifies the OCS Peer Node to be included in the distribution of new bindings to OCSs. Selecting a OCS Peer Node name (blue hyperlink) displays the Diameter > Configuration > Peer Nodes (Filtered) page where Diameter Peer Nodes are filtered by the OCS Peer Node Name.	Format: List Range: Configured Diameter Peer Nodes Note: The OCS Peer Node Name cannot be changed on the [Edit] page.
Comments	An optional comment to describe the OCS Peer Node.	Format: Text box Range:0-64 characters

Viewing OCSs

Use this task to view all configured OCSs on the SOAM.

Select **Policy and Charging > Configuration > OCSs**.

The **Policy and Charging > Configuration > Online Charging DRA > OCSs** page appears with a list of configured OCS Peer Nodes.

The fields are described in *PCRFs elements*.

Inserting OCSs

Use this task to insert (create new) OCSs.

The fields are described in *OCSs elements*.

1. On the Active SOAM, select **Policy and Charging > Configuration > Online Charging DRA > OCSs**.

The **Policy and Charging > Configuration > Online Charging DRA > OCSs** page appears.

2. Click **Insert**.

The **Policy and Charging > Configuration > Online Charging DRA > OCSs [Insert]** page opens.

3. Select an OCS Peer Node Name from the **OCS Peer Node Name** dropdown menu.
4. Enter an optional comment in the **Comments** field.
5. Click:
 - **OK** to save the new OCS and return to the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page.
 - **Apply** to save the new OCS and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entered OCS is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- Adding the new OCS would cause the maximum number of OCSs (2500) to be exceeded.

Editing OCSs

Use this task to edit OCSs.

1. On the Active SOAM, select **Policy and Charging > Configuration > Online Charging DRA > OCSs**.

The **Policy and Charging > Configuration > Online Charging DRA > OCSs** page appears. The page displays a list of the configured OCS Peer Nodes that are used when a new subscriber binding is created.

2. Click in the **Comments** field of the row to select the OCS to edit.

DO NOT click the blue OCS Peer Node Name (except to see the configuration of the Peer Node). The blue color indicates a hyper-link that opens the **Diameter > Configuration > Peer Nodes [Filtered]** page to display the configuration information for the Peer Node.

3. Edit the **Comments** field for the selected OCS.

The OCS Peer Node name cannot be changed.

4. Click:

- **OK** to save the change and return to the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page.
- **Apply** to save the change and remain on this page.
- **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page without saving any changes.

If **Apply** or **OK** is clicked and the selected **OCS Peer Node Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting an OCS

Use the following procedure to delete an OCS.

This procedure describes the recommended steps for deleting an OCS from a Policy and Charging configuration. In this procedure, OCS refers to a Diameter peer of the PCA, which is sometimes referred to as an OCS Front-end.

The OCS procedure minimizes disruption to policy signaling by:

- Preventing sessions from creating new bindings to an OCS that has been removed
- Allowing sessions with existing bindings to continue to use an OCS that has been removed until those sessions terminate normally

The following procedure describes the recommended steps for deletion of an OCS from a Policy and Charging configuration. In this procedure, OCS refers to a Diameter peer of the PCA, sometimes referred to as an OCS Front-End.

Note: The PCRF removal procedure is restricted to SOAM servers.

1. Use **Main Menu > Diameter > Configuration > Peer Nodes** from the SOAM GUI page to determine the Peer Node name of the OCS(s) being removed.
2. Use **Main Menu > Diameter > Configuration > Route Groups** from the SOAM GUI page, use the GUI filter by Peer Node with the corresponding Peer Node name of the OCS. This will display only the Route Groups that are associated with the OCS.
3. From the same GUI page, determine if there are any Route Groups that contain other Peer Nodes in addition to the OCS to be removed.

There are generally at least two Route Groups for each OCS. One Route Group with only the specified OCS peer, and one or more Route Groups with the specified OCS peer plus other OCS peers. The goal is to leave the route group with only the specified OCS peer, but delete the OCS peer from the other route groups. This allows routing for existing bindings to the OCS peer, but prevents alternate routing to the OCS peer.

4. From the same GUI page, edit each of the determined Route Groups and remove the OCS/OCS Front-End Peer Nodes from the Route Group.
This prevents alternate routing selection of the OCS peer being removed.
5. Use **Main Menu > Policy and Charging > Configuration > Online Charging DRA > OCSs** from the SOAM GUI page to delete the OCS.
This prevents new Bindings from using the OCS peer being removed.
6. After enough time has elapsed such that all Diameter sessions that could be bound to the OCS peer should have terminated normally, use **Main Menu > Policy and Charging > Configuration > Online Charging DRA > OCSs** on the SOAM GUI page to delete the route group containing only the OCS peer being removed.
7. Use **Main Menu > Diameter > Maintenance > Connections** from the SOAM GUI page to find the connection for the OCS Peer Node and disable it
8. Use **Main Menu > Diameter > Maintenance > Connections** from the SOAM GUI page to delete the connection to the OCS Peer Node.
9. Use **Main Menu > Diameter > Configuration > Peer Nodes** from the SOAM GUI page to delete the Diameter Peer Node for the OCS being removed.

CTFs

On an Active SOAM, the **Policy and Charging > Configuration > Online Charging DRA > CTFs** page lists the CTF Peer Nodes for which the Session state is to be stored. This page is only used if Session State Scope is set to "Specific Messages" in the Network-Wide Options Configuration on the NOAM

The list of CTF Peer Nodes is updated by inserting, editing, deleting a CTF Peer Node from the **Policy and Charging > Configuration > Online Charging DRA > CTFs** page at each site's SOAM.

CTFs elements

Table 49: CTFs elements describes the elements on the **Policy and Charging > Configuration > Online Charging DRA > CTFs** page on the Active SOAM.

Note: Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 49: CTFs elements

Fields (* indicates required field)	Description	Data Input Notes
* CTF Peer Node Name	The name of a configured Diameter Peer Node that identifies the CTF Peer Node to be included in the distribution of new bindings to CTFs. Selecting a CTF Peer Node name (blue hyperlink) displays the Diameter > Configuration > Peer Nodes (Filtered) page where Diameter Peer Nodes are filtered by the CTF Peer Node Name.	Format: List Range: Configured Diameter Peer Nodes Note: The CTF Peer Node Name cannot be changed on the [Edit] page.
Comments	An optional comment to describe the CTF Peer Node.	Format: Text box Range:0-64 characters

Viewing CTFs

Use this task to view all configured CTFs on the SOAM.

Select **Policy and Charging > Configuration > CTFs**.

The **Policy and Charging > Configuration > Online Charging DRA > CTFs** page appears with a list of configured CTF Peer Nodes for which the Session state is to be stored.

The fields are described in *PCRFs elements*.

Inserting CTFs

Use this task to insert (create new) CTFs.

The fields are described in *CTFs elements*.

1. On the Active SOAM, select **Policy and Charging > Configuration > Online Charging DRA > CTFs**.

The **Policy and Charging > Configuration > Online Charging DRA > CTFs** page appears.

2. Click **Insert**.

The **Policy and Charging > Configuration > Online Charging DRA > CTFs [Insert]** page opens.

3. Enter a unique CTF Peer Node Name in the **CTF Peer Node Name** field.

This name uniquely identifies the CTF Peer Node.

4. Enter an optional comment in the **Comments** field.

5. Click:

- **OK** to save the new CTF and return to the **Policy and Charging > Configuration > Online Charging DRA > CTFs** page.

- **Apply** to save the new CTF and remain on this page.
- **Cancel** to return to the **Policy and Charging > Configuration > Onling Charging DRA > CTFs** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entered CTF is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- Adding the new CTF would cause the maximum number of CTFs (2500) to be exceeded.

Editing CTFs

Use this task to edit CTFs.

1. On the Active SOAM, select **Policy and Charging > Configuration > Online Charging DRA > CTFs**.

The **Policy and Charging > Configuration > Online Chargin DRA > CTFs** page appears. The page displays a list of the configured CTF Peer Nodes that are used when a new subscriber binding is created.

2. Click in the **Comments** field of the row to select the CTF to edit.
3. Edit the **Comments** field for the selected CTF.

The CTF Peer Node name cannot be changed.

4. Click:

- **OK** to save the change and return to the **Policy and Charging > Configuration > Online Charging DRA > CTFs** page.
- **Apply** to save the change and remain on this page.
- **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > CTFs** page without saving any changes.

If **Apply** or **OK** is clicked and the selected **CTF Peer Node Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting a CTF

Use the following procedure to delete an CTF.

This procedure describes the recommended steps for deleting an CTF from a Policy and Charging configuration. In this procedure, CTF refers to a Diameter peer of the PCA, which is sometimes referred to as an CTF Front-end.

The CTF procedure minimizes disruption to policy signaling by:

- Preventing sessions from creating new bindings to an CTF that has been removed
- Allowing sessions with existing bindings to continue to use an CTF that has been removed until those sessions terminate normally

The following procedure describes the recommended steps for deletion of an CTF from a Policy and Charging configuration. In this procedure, CTF refers to a Diameter peer of the PCA, sometimes referred to as an CTF Front-End.

Note: The PCRF removal procedure is restricted to SOAM servers.

1. Use **Main Menu > Diameter > Configuration > Peer Nodes** from the SOAM GUI page to determine the Peer Node name of the CTF(s) being removed.
2. Use **Main Menu > Diameter > Configuration > Route Groups** from the SOAM GUI page, use the GUI filter by Peer Node with the corresponding Peer Node name of the CTF. This will display only the Route Groups that are associated with the CTF.
3. From the same GUI page, determine if there are any Route Groups that contain other Peer Nodes in addition to the CTF to be removed.
There are generally at least two Route Groups for each CTF. One Route Group with only the specified CTF peer, and one or more Route Groups with the specified CTF peer plus other CTF peers. The goal is to leave the route group with only the specified CTF peer, but delete the CTF peer from the other route groups. This allows routing for existing bindings to the OCS peer, but prevents alternate routing to the OCS peer.
4. From the same GUI page, edit each of the determined Route Groups and remove the OCS/OCS Front-End Peer Nodes from the Route Group.
This prevents alternate routing selection of the OCS peer being removed.
5. Use **Main Menu > Policy and Charging > Configuration > Online Charging DRA > OCSs** from the SOAM GUI page to delete the OCS.
This prevents new Bindings from using the OCS peer being removed.
6. After enough time has elapsed such that all Diameter sessions that could be bound to the OCS peer should have terminated normally, use **Main Menu > Policy and Charging > Configuration > Online Charging DRA > OCSs** on the SOAM GUI page to delete the route group containing only the OCS peer being removed.
7. Use **Main Menu > Diameter > Maintenance > Connections** from the SOAM GUI page to find the connection for the OCS Peer Node and disable it
8. Use **Main Menu > Diameter > Maintenance > Connections** from the SOAM GUI page to delete the connection to the OCS Peer Node.
9. Use **Main Menu > Diameter > Configuration > Peer Nodes** from the SOAM GUI page to delete the Diameter Peer Node for the OCS being removed.

OCS Session State

On an Active NOAM, the **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page lists the network-wide list of Online Charging Servers (OCSs), listed by their Realm and FQDN. It is used to configure the Session State setting for OCSs.

The list of OCSs is updated by inserting or deleting an OCS from the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page at each site's SOAM. Additionally, the Realm and FQDN are configured from each site's **Diameter > Configuration > Peer Nodes** page on the SOAM.

Once the list of OCSs is populated, the following options become available:

- Editing whether or not OCS Session State is enabled
- Pausing the updating of the OCS list

Viewing OCS Session State

Use this task to view an OCS Session State.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > OCS Session State**.

The **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page appears.

2. Check or uncheck the **Pause updates** box to pause or unpause updates to the list of OCSs. The OCSs listed were inserted or deleted on the **Policy and Charging > Configuration > Online Charging DRA > OCSs** page on the SOAM, as well as the more specific Realm and FQDN that are configured from each site's **Diameter > Configuration > Peer Nodes** page on the SOAM.

Editing OCS Session State

Use this task to edit a Realm.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > OCS Session State**.

The **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page appears.

2. Click **Edit**.

The **Policy and Charging > Configuration > Online Charging DRA > OCS Session State [Edit]** page opens.

3. The **Realm** and **FQDN** fields are disabled and cannot be edited from this screen.
4. Check or uncheck the box to **Enable or Disable OCS Session State**.
5. Click:
 - **OK** to save the edited OCS Session State and return to the **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page.
 - **Apply** to save the edited OCS Session State and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > OCS Session State** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any required field no longer exists

Realms

The **Policy and Charging > Configuration > Online Charging DRA > Realms** page on an NOAM contains the list of Online Charging network realms for which the Session state is stored.

Note: This page is only use if **Session State Scope** is set to "Specific Messages" on the **Policy and Charging > Configuration > Online Charging DRA > Network-Wide Options** page.

Realms elements

[Table 50: Realms elements](#) describes the elements on the **Policy and Charging > Configuration > Online Charging DRA > Realms** page.

Note: Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

The Realms table lists the Online Charging network realms for which the Session state is to be stored. This table is only used if Session State Scope is set to "Specific Messages" in the Network-Wide Options Configuration.

Table 50: Realms elements

Fields (* indicates required field)	Description	Data Input Notes
*Realm Name	Realm name is a case-insensitive string consisting of a list of labels separated by dots, where a label may contain letter, digits, dashes('-') and underscore('_'). A label must start with a letter, digit or underscore and must end with a letter or digit. Underscores may be used only as the first character. A label must be at most 63 characters long and a Realm must be at most 255 characters long.	Format: text box Range: 1-1000 entries Default: N/A
Comments	An optional comment to provide more information about the purpose of this PCRF Pool or Sub-Pool.	Format: Text box Range:0-64 characters

Inserting Realms

Use this task to insert (create new) Realms.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > Realms**.
The **Policy and Charging > Configuration > Online Charging DRA > Realms** page appears.
2. Click **Insert**.
The **Policy and Charging > Configuration > Online Charging DRA > Realms [Insert]** page opens.
3. Enter a unique Realm Name in the **Realm Name** field.
4. If desired, enter an optional comment in the **Comments** field to describe the Realm. The entry must be characters in the range of 0 to 64, and the default is N/A.
5. Click:
 - **OK** to save the new Realm name and return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page.
 - **Apply** to save the new Realm name and remain on this page.
 - **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any required field is empty (not entered).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Adding a new Realm would cause the maximum number of Realms (1000) to be exceeded.
- The entered Realm name is not unique (already exists).

Editing Realms

Use this task to edit a Realm.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > Realms**.

The **Policy and Charging > Configuration > Online Charging DRA > Realms** page appears.

2. Click **Edit**.

The **Policy and Charging > Configuration > Online Charging DRA > Realms [Edit]** page opens.

3. Edit the unique Realm Name in the **Realm Name** field.

4. If desired, edit an optional comment in the **Comments** field.

5. Click:

- **OK** to save the edited Realm name and return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page.
- **Apply** to save the edited Realm name and remain on this page.
- **Cancel** to return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page without saving any changes.

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any required field no longer exists
- Any fields contain a value that contains invalid characters or is out of the allowed range.

Deleting Realms

Use this task to delete a PCRF Pool.

1. On the Active NOAM, select **Policy and Charging > Configuration > Online Charging DRA > Realms**.

The **Policy and Charging > Configuration > Online Charging DRA > Realms** page appears.

2. Select the **Realm** to be deleted.

3. Click **Delete**.

A window appears to confirm the delete.

4. Click:

- **OK** to delete the Realm.
- **Cancel** to cancel the delete function and return to the **Policy and Charging > Configuration > Online Charging DRA > Realms** page.

If **OK** is clicked and the selected Realm no longer exists (it was deleted by another user), an error message is displayed, and the Realms page is refreshed. The row that was selected is no longer displayed in the list.

Error Codes

For each Policy and Charging Site, the Diameter Error Result Code value to send to the Request initiator for policy related errors can be configured according to which condition applies. Each condition can be mapped to a different Result Code for each supported interface. Result Codes can be Diameter IANA defined or experimental.

When PCRF Pooling is enabled, new binding cannot be created unless the binding-capable session initiation request contains a configured APN. If the binding-capable session initiation request arrives with either no APN, or an APN that is not configured in the Access Point Names table, the request is answered by Policy DRA using a configurable error response code. To configure the Diameter response

code for this scenario, a new Missing Or Unconfigured APN error condition has been added to the existing SOAM error. This error response applies to all binding capable interfaces (for example, Gx, Gxx, and S9) and can be configured with either an IANA Diameter result code, or an experimental result code and vendor-id.

Three-digit error codes in Diameter Error-Message AVPs indicate exactly why a slave session could not be routed. This provides more robust troubleshooting using Diameter capture tools.

A 3-digit error code is an identifier to uniquely identify a specific error scenario (not error category) encountered in a Diameter Answer message generated by PCA. 3-digit codes are unique across all DSR layers (DSR connection layer, routing layer and application layer) and all DSR applications (PCA, RBAR, FABR, and IDIH, etc.) for errors they represent. The ranges of 500-549 and 850-899 are for the PCA application, while the DSR connection layer, routing layer and other DSR applications uses other non-overlapping ranges. Multiple errors may belong to a same error category and are associated with a same Result-Code. It is the 3-digit code that can distinguish an error from others. Users should search for the 3-digit code when identifying an error if possible and available.

Note: The error conditions in this table are GUI-configurable.

Table 51: PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error-Message Suffix	Error Text
Missing or Unconfigured APN	PDRA	Policy-related binding capable session initiation request messages	3002	500	Missing or Unconfigured APN. Binding capable session initiation request is received with no APN
Missing or Unconfigured APN	PDRA	Policy-related binding capable session initiation request messages	3002	501	Missing or Unconfigured APN. Binding capable session initiation request is received with APN, but the APN is not configured in the APN configuration.
Unable To Route	PDRA	Policy-related binding capable and dependent session initiation request messages	3002	502	Unable To Route. Request message is received and a binding with a PCRF was found. P-DRA can't route the request to PCRF due to DSR queue full error.
Unable To Route	PDRA	Policy-related binding capable and dependent session initiation request messages	3002	3-digit error code from DRL	Unable To Route. Request message is received and a binding with a PCRF was found. P-DRA can't route the request to PCRF due to PCRF being unreachable. DRL Text string.
No Usable Keys In	PDRA	Policy-related binding	3002	503	No Usable Keys In Binding Dependent Message. No

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error-Message Suffix	Error Text
Binding Dependent Message		dependent session initiation request messages			binding key in Binding Key Priority GUI can be matched, or no key is included in the binding dependent message.
Binding Not Found	PDRA	Policy-related binding dependent session initiation request messages	3002	505	Binding Not Found. Binding record is not found after examining all configured binding keys in Diameter message.
SBR Error	PDRA	Policy-related binding capable and dependent session initiation, update or terminate answer messages,	3002	507	SBR Error. ComAgent timeout.
SBR Error	PDRA	Policy-related binding capable and dependent session initiation, update or terminate answer messages	3002	508	SBR Error. SBR database error prevents SBR from reading, writing or deleting a record,
Session Not Found	PDRA	Policy-related binding capable and dependent session update or terminate request messages	3002	509	Session Not Found. Session record doesn't exist for given session ID
PCA Unavailable Or Degraded	PDRA/ OCDRA	Any Diameter Requests forwarded to PCA	3002	305	PCA Unavailable or Degraded
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	520	SBR Error. Binding capable session initiation request received, but no PCRFs are configured at the site, or PCRF ID is not found in PCRF table.
SBR Error	PDRA	Policy-related binding capable	3002	521	SBR Error. Maximum number of sessions per binding is exceeded that

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error-Message Suffix	Error Text
		session initiation request messages			fails the binding creation for given subscriber's key.
Session ID is missing from Request	PDRA	Any Policy-related Diameter Requests forwarded to P-DRA	5005	522	Session ID is missing from Request
CC - Request - Type AVP is missing from CCR message	PDRA	Policy-related binding capable session initiation, update or terminate request messages	5005	523	CC-Request-Type AVP is missing from CCR message
Not In Use					
Invalid AVP value in request message	PDRA	Any Policy-related Diameter Requests forwarded to P-DRA	5004	525	Invalid AVP value in request message
Destination - Host AVP is missing in in-session request	PDRA	Policy-related binding capable update and terminate request and dependent session initiation update or terminate request messages	5012	506	Destination-Host AVP is missing in in-session request
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	510	Unable To Route. A slave session could not be routed because on polling the slave sessionRef was no longer in the binding database.
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	511	Unable To Route. A slave session could not be routed because on polling the master sessionRef was no longer in the binding database.

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error-Message Suffix	Error Text
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	512	Unable To Route. A slave session could not be routed because on polling the master sessionRef was early too long.
SBR Error	PDRA	Policy-related Requests and Answers	3002	504	SBR Error. ComAgent unavailable when sending stack event to SBR
Unsupported Application ID	PDRA/ OCDRA	Diameter Requests	3007	530	Application ID unsupported by PCA
Command Code and App ID no match	PDRA	Policy-related Requests and Answers	5019	531	Command Code does not match App ID or not exist
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	513	Unable To Route. A slave session could not routed because on polling the master session and internal error occurred.
PCA Functionality Unavailable or Disabled	PDRA	Policy related binding capable and dependent session update or terminate request messages	3002	532	PCA Functionality Unavailable or Disabled. Policy DRA Function Disabled.
PCA Functionality Unavailable or Disabled	PDRA	Policy related binding capable and dependent session update or terminate request messages	3002	533	PCA Functionality Unavailable or Disabled. Policy DRA Function Unavailable.
PCA Functionality Unavailable or Disabled	OCDRA	Online Charging related binding independent session request messages	3002	534	PCA Functionality Unavailable or Disabled. Online Charging DRA Function Disabled.
PCA Functionality Unavailable or Disabled	OCDRA	Online Charging related binding independent session request messages	3002	535	PCA Functionality Unavailable or Disabled. Online Charging DRA Function Unavailable

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error-Message Suffix	Error Text
Session ID is missing from Request	OCDRA	Any Online Charging -related Diameter Requests forwarded to OC-DRA	5005	536	Session ID is missing from Request
CC - Request - Type AVP is missing from CCR message	OCDRA	Any Online Charging-related Diameter Requests forwarded to OC-DRA	5005	537	CC-Request-Type AVP is missing from CCR message
Invalid AVP value in request message	OCDRA	Any Online Charging-related Diameter Requests forwarded to OC-DRA	5004	538	Invalid AVP value in request message
Not In Use					
Unable To Route	OCDRA	Online Charging-related binding independent session request messages	3002	540	Unable To Route. Request message is received, OC-DRA can't route the request to OCS due to DSR queue full error.
Unable To Route	OCDRA	Online Charging-related binding independent session initiation request messages	3002	539	Unable To Route. Request message can not be routed to peer node. DIAM-ERR-XXXX-XXX:DRL Text string.
SBR Error	OCDRA	Online Charging-related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	541	SBR Error. ComAgent timeout.

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error-Message Suffix	Error Text
SBR Error	OCDRA	Online Charging-related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	542	SBR Error. SBR database error prevents SBR from reading, writing or deleting a record,
SBR Error	OCDRA	Online Charging-related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	543	SBR Error . ComAgent unavailable when sending stack event to SBR,
Session Not Found	OCDRA	Online Charging-related session update or terminate request messages, if session state or topology hiding applies	5002	544	Session Not Found. Session record doesn't exist for given session ID
Command Code and App ID no match	OCDRA	Online Charging-related Requests.	3001	545	Command Code does not match App ID or not exist

On the **Policy and Charging > Configuration > Error Codes** page on the SOAM, you can perform the following action:

- Select an **Error Condition** in the list, and click the **Edit** button.

The **Policy and Charging > Configuration > Error Codes [Edit]** page opens. You can edit the selected Error Code. See [Editing Error Codes](#).

The fields are described in [Error Codes elements](#).

Error Codes elements

Table 53: Error Codes elements describes the elements on the **Policy add Charging > Configuration > Error Codes** pages. Data Input Notes apply to the [Edit] page; the View page is read-only.

The Error Codes define the Result Codes to be returned for various Policy and Charging Error Conditions. Each Error Condition will return the Result Code configured for each applicable Diameter interface.

Table 52: Interfaces Supported for Each Error Code indicates the Diameter interfaces that are supported for each Error Code.

The default Result Code is 3002-DIAMETER_UNABLE_TO_DELIVER.

Table 52: Interfaces Supported for Each Error Code

Error Code	Result Code	Vendor ID
PCA Unavailable Or Degraded	Gx/Gxx, Gx-Prime, Rx, S9, Gy/Ro	Gx/Gxx, Gx-Prime, Rx, S9, Gy/Ro
PCA Functionality Unavailable or Disabled	Gx/Gxx, Rx, S9, Gx-Prime, Gy/Ro	Gx/Gxx, Rx, S9, Gx-Prime, Gy/Ro
Binding Not Found	Rx, Gx-Prime	Rx, Gx-Prime
Unable To Route	Gx/Gxx,Gx-Prime, Rx, S9, Gy/Ro	Gx/Gxx,Gx-Prime, Rx, S9, Gy/Ro
SBR Error	Gx/Gxx,Gx-Prime, Rx, S9, Gy/Ro	Gx/Gxx,Gx-Prime, Rx, S9, Gy/Ro
No Usable Keys In Binding Dependent Message	Rx,Gx-Prime	Rx,Gx-Prime
Session Not Found	Gx/Gxx,Gx-Prime, Rx, S9, Gy/Ro	Gx/Gxx,Gx-Prime, Rx, S9, Gy/Ro
Missing or Unconfigured APN	Gx/Gxx, S9	Gx/Gxx, S9

Table 53: Error Codes elements

Fields (* indicates required field)	Description	Data Input Notes
Error Condition	The name of the selected Policy and Charging Error Condition.	View only; cannot be edited
* Gx/Gxx, Result Code	The Result Code to be returned on the Gx and Gxx interfaces.	Format: Text box Range: 1-9999 Default: 3002

Fields (* indicates required field)	Description	Data Input Notes
Gx/Gxx Vendor ID	The Vendor ID that corresponds with the Gx and Gxx interfaces. The Vendor ID '---' means the RFC standard error code will be sent.	Format: Text box Range: 1-4294967295
* Rx Result Code	The Result Code to be returned to the Rx interface.	Format: Text box Range: 1-9999 Default: 3002
Rx Vendor ID	The Vendor ID that corresponds with the Rx interface. The Vendor ID '---' means the RFC standard error code will be sent.	Format: Text box Range: 1-4294967295
* S9 Result Code	The Result Code to be returned to the S9 interface.	Format: Text box Range: 1-9999 Default: 3002
S9 Vendor ID	The Vendor ID that corresponds the S9 interface. The Vendor ID '---' means the RFC standard error code will be sent.	Format: Text box Range: 1-4294967295
Gx-Prime Result Code	The Result Code to be returned on the Gx-Prime interface	Format: Text Box Range: 1-9999 Default: 3002
Gx-Prime Vendor ID	The Vendor ID that corresponds with the Gx-Prime interface. The Vendor ID '---' means the RFC standard error code will be sent.	Format: Text Box Range: 1-4294967295
Gy/Ro Result Code	The Result code to be returned on the Gy/Ro interface.	Format: Text Box Range: 1-9999
Gy/Ro Vendor ID	The Vendor ID which corresponds with the experimental code for the Gy/Ro interface.	Format: Text Box Range: 1-4294967295

Viewing Error Codes

Use this task to view configured Error Codes on the SOAM.

Select **Policy and Charging > Configuration > Error Codes**.

The **Policy and Charging > Configuration > Error Codes** page appears with a list of configured Error Codes. The fields are described in [Error Codes elements](#).

Editing Error Codes

Use this task to edit Error Codes on the Active SOAM.

The fields are described in [Error Codes elements](#).

1. Select **Policy and Charging > Configuration > Error Codes**.

The **Policy and Charging > Configuration > Error Codes** page appears

2. Select the **Error Condition** that you want to edit.
3. Click **Edit**.

The **Policy and Charging > Configuration > Error Codes [Edit]** page opens

The fields that appear on the **Policy and Charging > Configuration > Error Codes [Edit]** page are dependent on the Error Condition that was selected.

4. Edit the fields to define the selected Error Condition.
5. Click:

- **Ok** to save the changes and return to the **Policy and Charging > Configuration > Error Codes** page
- **Apply** to save the changes and remain on this page
- **Cancel** to discard the changes and return to the **Policy Charging > Configuration > Error Codes** page

If **OK** or **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any required field value is missing (not entered or selected)
- Any fields contain invalid input (for example, the wrong type of data was entered or a value is out of range).

Alarm Settings

Note: Alarm Settings are configurable only on Active NOAM servers, and are viewable on NOAM and SOAM servers.

On the **Policy and Charging > Configuration > Alarm Settings** page on an SOAM, the user can view the configured Alarm Thresholds and Suppress indications.

Each alarm can be configured with Minor, Major, and Critical threshold percentages.

The fields are described in [Alarm Settings elements](#).

On the **Policy and Charging > Configuration > Alarm Settings** page on the NOAM, you can change the Alarm Thresholds and the Suppress indications for the following alarms:

- DSR Application Ingress Message Rate

The DSR Application Ingress Message Rate alarm is raised when the average Policy and Charging ingress messages rate exceeds the configured Alarm Threshold. The thresholds are based on the engineered system value for Ingress Message Capacity.

- SBR Sessions Threshold Exceeded

The SBR Sessions Threshold Exceeded alarm percent full is based on the number of Session records compared to an engineered maximum that varies according to the number of session SBR Server Groups per mated pair.

The SBR Sessions Threshold Exceeded alarm is raised when number of concurrent policy and Charging sessions exceeds the configured threshold.

- SBR Bindings Threshold Exceeded

The SBR Bindings Threshold Exceeded alarm measures the number of IMSI Anchor Key records against an engineered maximum value that varies according to the number of binding SBR Server Groups.

The Policy SBR Bindings Threshold Exceeded alarm works similarly to the session capacity alarm except that the scope of the binding capacity alarm is network-wide.

Alarm Settings elements

Table 54: Alarm Settings elements describes the elements on the **Policy and Charging > Configuration > Alarm Settings** page. The elements can be configured and viewed on the NOAM, and only viewed on the SOAM. Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

The page contains three sets of input fields for the following alarms:

- DSR Application Ingress Message Rate
- SBR Sessions Threshold Exceeded
- SBR Bindings Threshold Exceeded

The element labels are the same for each input field set, but some serve different purposes and have different values. These distinctions are noted in the table.

Table 54: Alarm Settings elements

Elements (* indicates required field)	Description	Data Input Notes
DSR Application Ingress Message Rate		
* Alarm Name	This alarm is raised when average Policy and Charging ingress messages rate exceeds the configured threshold. The thresholds are based on the engineered system value for Ingress Message Capacity.	Format: Non-editable text box Range: DSR Application Ingress Message Rate
* Critical Alarm Threshold (Percent)	The Policy and Charging ingress message rate threshold for this alarm to be raised as Critical. The threshold is a percentage of the Ingress Capacity Capability.	Format: Text box Range: 100-200 Default: 160
Suppress Critical	Controls whether this alarm is raised as Critical.	Format: Check box Range: Unchecked (No) or Checked (Yes)

Policy and Charging Configuration

Elements (* indicates required field)	Description	Data Input Notes
		Default: Unchecked (No)
* Major Alarm Threshold (Percent)	The Policy and Charging ingress message rate threshold for this alarm to be raised as Major. The threshold is a percentage of the Ingress Capacity Capability.	Format: Text box Range: 100-200 Default: 140
Suppress Major	Controls whether this alarm is raised as Major.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Minor Alarm Threshold (Percent)	The Policy and Charging ingress message rate threshold for this alarm to be raised as Minor. The threshold is a percentage of the Ingress Capacity Capability.	Format: Text box Range: 100-200 Default: 110
Suppress Minor	Controls whether this alarm is raised as Minor.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
SBR Sessions Threshold Exceeded		
* Alarm Name	This alarm is raised when the number of concurrent Policy and Online Charging SBR sessions exceeds the configured threshold.	Format: Non-editable text box Range: Policy SBR Sessions Threshold Exceeded
* Critical Alarm Threshold (Percent)	The concurrent sessions threshold for this alarm to be raised as Critical. The threshold is a percentage of the Maximum SBR Sessions.	Format: Text box Range: 1-99 Default: 95
Suppress Critical	Controls whether this alarm is raised as Critical.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Major Alarm Threshold (Percent)	The concurrent sessions threshold for this alarm to be raised as Major. The threshold is a percentage of the Maximum SBR Sessions.	Format: Text box Range: 1-99

Policy and Charging Configuration

Elements (* indicates required field)	Description	Data Input Notes
		Default: 90
Suppress Major	Controls whether this alarm is raised as Major.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Minor Alarm Threshold (Percent)	The concurrent sessions threshold for this alarm to be raised as Minor. The threshold is a percentage of the Maximum SBR Sessions.	Format: Text box Range: 1-99 Default: 80
Suppress Minor	Controls whether this alarm is raised as Minor.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
SBR Bindings Threshold Exceeded		
* Alarm Name	This alarm is raised when the number of concurrent Policy SBR bindings exceeds the configured threshold.	Format: Non-editable text box Range: Policy SBR Bindings Threshold Exceeded
* Critical Alarm Threshold (Percent)	The concurrent bindings threshold for this alarm to be raised as Critical. The threshold is a percentage of the Maximum Policy SBR Bindings.	Format: Text box Range: 1-99 Default: 95
Suppress Critical	Controls whether this alarm is raised as Critical.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Major Alarm Threshold (Percent)	The concurrent bindings threshold for this alarm to be raised as Major. The threshold is a percentage of the Maximum Policy SBR Bindings.	Format: Text box Range: 1-99 Default: 90
Suppress Major	Controls whether this alarm is raised as Major.	Format: Check box Range: Unchecked (No) or Checked (Yes)

Elements (* indicates required field)	Description	Data Input Notes
		Default: Unchecked (No)
* Minor Alarm Threshold (Percent)	Te concurrent bindings threshold for this alarm to be raised as Minor. The threshold is a percentage of the Maximum Policy SBR Bindings.	Format: Text box Range: 1-99 Default: 80
Suppress Minor	Controls whether this alarm is raised as Minor.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)

Viewing Alarm Settings

Use this task to view configured Alarm-Settings on either the NOAM or SOAM.

Select **Policy and Charging > Configuration > Alarm Settings**.

The **Policy and Charging > Configuration > Alarm Settings** page appears with a list of configured Alarm Settings.

The fields are described in [Alarm Settings elements](#).

Defining Alarm Settings

Use this task to define Alarm Settings on an Active NOAM.

Note: Alarm Settings are configurable only on Active NOAM servers, and are viewable on NOAM and SOAM servers.

The fields are described in [Alarm Settings elements](#).

1. Select **Policy and Charging > Configuration > Alarm Settings**.

The **Policy and Charging > Configuration > Alarm Settings** page appears.

2. Enter values in the editable fields to define the alarm settings.

3. Click:

- **Apply** to save the changes and remain on this page.
- **Cancel** to discard the changes and remain on the **Policy and Charging > Configuration > Alarm Settings** page.

If **Apply** is clicked and any of the following conditions exist, an error message appears:

- The entered values contain the wrong data type or is out of the allowed range.
- The value entered for **Critical Alarm Threshold (Percent)** is less than or equal to the value entered for **Major Alarm Threshold (Percent)**.
- The value entered for **Major Alarm Threshold (Percent)** is less than or equal to the value entered for **Minor Alarm Threshold (Percent)**.

Congestion Options

Congestion Options are configurable on Active NOAM servers.

The following Congestion Options can be configured:

- Alarm Thresholds, which are used to:
 - Set the percentage of the Policy and Charging ingress message rate capacity at which an alarm is raised with Critical, Major, or Minor severity.
 - Set the percentage of the Policy and Charging ingress message rate capacity at which a Critical, Major, or Minor severity alarm is cleared.

The percentages control the onset and abatement of the corresponding Congestion Levels.

Default thresholds are based on the engineered system value for Ingress Policy and Charging Request Message Capacity.

- Message Throttling Rules, which determine the percentage of Session Creation, Update, and Terminate Request messages that are discarded when Congestion Levels 1, 2, and 3 exist.

The fields are described in [Congestion Options elements](#).

Congestion Options elements

[Table 55: Congestion Options elements](#) describes the elements on the **Policy and Charging > Configuration > Congestion Options** page. The elements can be configured and viewed on the NOAM.

The page contains two sets of input fields:

- Alarm Thresholds
- Message Throttling Rules

Table 55: Congestion Options elements

Fields (* indicates required field)	Description	Data Input Notes
Alarm Thresholds		
Alarm Name	Alarm is raised when average Policy and Charging ingress request messages rate exceeds the configured threshold. The thresholds are based on the engineered system value for Ingress Policy and Charging Request Message Capacity.	Format: Non-editable text box Range: Policy and Charging Server in Congestion
* Critical Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 3.	Format: Text box Range: 100-200 Default: 160
* Critical Alarm Abatement Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 3.	Format: Text box Range: 100-200 Default: 150

Policy and Charging Configuration

Fields (* indicates required field)	Description	Data Input Notes
* Major Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 2.	Format: Text box Range: 100-200 Default: 140
* Major Alarm Abatement Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 2.	Format: Text box Range: 100-200 Default: 130
* Minor Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 1.	Format: Text box Range: 100-200 Default: 110
* Minor Alarm Abatement Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 1.	Format: Text box Range: 100-200 Default: 100
<p>Message Throttling Rules</p> <p>Tabs for Congestion Level 1, Congestion Level 2, and Congestion Level 3</p>		
* Discard Session Creation Requests	Percentage of Request messages that result in new session creation, to be discarded when this congestion level exists.	Format: Text box Range: 0-100 Default: Level 1 - 25 Level 2 - 50 Level 3 - 100
* Discard Session Update Requests	Percentage of Request messages that result in updating existing sessions, to be discarded when this congestion level exists.	Format: Text box Range: 0-100 Default: Level 1 - 0 Level 2 - 25 Level 3 - 50
* Discard Session Terminate Requests	Percentage of Request messages that result in terminating existing sessions, to be discarded when this congestion level exists.	Format: Text box Range: 0-100 Default: Level 1 - 0 Level 2 - 0

Fields (* indicates required field)	Description	Data Input Notes
		Level 3 - 0

Viewing Congestion Options

Use this task to view configured Congestion Options on the NOAM.

Select **Policy and Charging > Configuration > Congestion Options**.

The **Policy and Charging > Configuration > Congestion Options** page appears with a list of configured Congestion Options.

The fields are described in [Congestion Options elements](#).

Setting Congestion Options

Use this task to set the following Congestion Options on the Active NOAM:

- **Alarm Thresholds** for the **Policy and Charging Server in Congestion** onset and abatement alarm for Critical, Major, and Minor severities
 - **Message Throttling Rules** for discarding Session Creation, Update, and Terminate Requests for Congestion Levels 1, 2, and 3
1. Select **Policy and Charging > Configuration > Congestion Options**.
The **Policy and Charging > Configuration > Congestion Options** page appears.
 2. Enter changes for the **Alarm Thresholds**.
 3. Enter changes for the **Message Throttling Rules**.
 4. Click:
 - **Apply** to save the Congestion Options changes and refresh the page to show the changes.
 - **Cancel** to discard the changes and refresh the page.

If **Apply** is clicked and any of the following conditions exist, an error message appears:

- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- A **Major Alarm Onset Threshold** value is greater than the corresponding **Critical Alarm Onset Threshold**.
- A **Minor Alarm Onset Threshold** value is greater than the corresponding **Major Alarm Onset Threshold**.
- An **Alarm Abatement Threshold** value is greater than the corresponding **Alarm Onset Threshold** of a particular severity.

Post-Configuration Activities

After PCA configuration is complete, the following activities need to be performed to make the Policy DRA application fully operational in the system:

- Enable the PCA application
- Enable Diameter Connections with Peer Nodes
- Status Verification

Enable the PCA Application

Use this task to enable the PCA application. For each Active SOAM,

1. Select **Diameter > Maintenance > Applications**.
The **Diameter > Maintenance > Applications** page appears.
2. Under **DSR Application Name**, select each **PCA** row.
To select more than one row, press and hold **Ctrl** while you click each row.
3. Click **Enable**.
4. Verify the application status on the page.
The **Admin State**, **Operational Status**, **Operational Reason**, and **Congestion Level** in each of the selected rows should change respectively to **Enabled**, **Available**, **Normal**, **Normal**.

Enable Connections

Use the following task to enable one or more connections to Peer Nodes.

1. At the Active SOAM, select **Diameter > Maintenance > Connections**.
The **Diameter > Maintenance > Connections** page appears.
2. Select 1 - 20 connections to enable.
To select multiple connections, press and hold the **Ctrl** key while you select each connection.
To select multiple contiguous connections, click the first connection you want, press and hold the **Shift** key, and select the last connection you want. All the connections between are also selected.
3. Click **Enable**.
A confirmation box appears.
4. Click **OK**.
The selected connections are enabled.
If any of the selected connections no longer exist (they have been deleted by another user), an error message is displayed, but any selected connections that do exist are enabled.
5. Verify Connection status on the page.
Verify that the **Admin State** of all connections changes to **Enabled** and the **Operational Reason** shows **Connecting** for connections to PCRF nodes and **Listening** for connections to other nodes (such as policy clients - PCEF, AF, and others). nodes.
For connections of type **Responder Only** (Policy Client nodes), the **Operational Status** and **Operational Reason** will be "Unk" if IPFE TSA connections are used.

Status Verification

Use the following task to verify PCA and SBR status after configuration is complete.

1. Verify Communication Agent (ComAgent) HA Services Status.
 - a) At the Active NOAM, select **Communication Agent > Maintenance > Connection Status**.
 - b) Verify that **Resource Routing Status** is **Available** for all listed **User/Provider** entries.
2. Verify the ComAgent Automatic Connection Status.
 - a) At the Active NOAM, select **Communication Agent > Maintenance > HA Services Status**
 - b) Verify that **Automatic Connection Count** is **X of Y In Service**, where $Y \geq X$ and $X = Y$ indicate successful Automatic Connection setup.
3. Verify SBR Status.
 - a) At the Active NOAM, select **Policy and Charging > Maintenance > SBR Status**.
 - b) Verify that the server **Resource HA Role** is **Active/Standby/Spare** and **Congestion Level** is **Normal** for all Servers in each Server Group in the Binding Region and Mated Site tab entries.

Bulk Import and Export

The *Diameter Common User's Guide* describes the use and operation of Bulk Import and Export functions:

- **Help > Diameter Common > Bulk Import**
- **Help > Diameter Common > Bulk Export**

The Bulk Import and Export functions can be used to export Diameter, IPFE, and Application configuration data in CSV files to a location outside the system, and to import the files (usually edited) into the system where the Import function is executed.

Bulk Import

The Bulk Import operations use configuration data in ASCII Comma-Separated Values (CSV) files (.csv), to insert new data into, update existing data in, or delete existing data from the configuration data in the system.

Note: Some configuration data can be imported only with the Update operation, and other data can be imported with Insert and Delete operations but not Update. Refer to the *Diameter Common User's Guide* or the **Diameter Common > Import Help** for valid Import operations.

Import CSV files can be created by using a Bulk Export operation, or can be manually created using a text editor.

Note: The format of each Import CSV file record must be compatible with the configuration data in the release that is used to import the file.

Files that are created using the Bulk Export operation can be exported either to the local Status & Manage File Management Directory (**Status & Manage > Files** page), or to the local Export server Directory.

CSV files that are in the local File Management area can be used for Bulk Import operations on the local system.

Files can be created manually using a text editor; the files must be uploaded to the File Management area of the local system before they can be used for Import operations on the local system.

The following Import operations can be performed:

- Insert new configuration data records that do not currently exist in the system
- Update existing configuration data in the system
- Delete existing configuration data from the system

Each Import operation creates a log file. If errors occur, a Failures CSV file is created that appears in the File Management area. Failures files can be downloaded, edited to correct the errors, and imported to successfully process the records that failed. >Failures files that are unchanged for more than 14 days and log files that are older than 14 days are automatically deleted from the File Management area.

Bulk Export

The Bulk Export operation creates ASCII Comma-Separated Values (CSV) files (.csv) containing Diameter, IPFE, and Application configuration data. Exported configuration data can be edited and used with the Bulk Import operations to change the configuration data in the local system without the use of GUI pages. The exported files can be transferred to and used to configure another system.

Each exported CSV file contains one or more records for the configuration data that was selected for the Export operation. The selected configuration data can be exported once immediately, or exports can be scheduled to periodically occur automatically at configured times.

The following configuration data can be exported in one Export operation:

- All exportable configuration data in the system
- All exportable configuration data from the selected Application, IPFE, or Diameter (each component's data is in a separate file)
- Exportable configuration data from a selected configuration component for the selected Application, IPFE, or Diameter

Exported files can be written to the File Management Directory in the local File Management area (**Status & Manage > Files** page), or to the Export server Directory for transfer to a configured remote Export server.

CSV files that are in the local File Management area can be used for Bulk Import operations on the local system.

If the export has any failures or is unsuccessful, the results of the export operation are logged to a log file with the same name as the exported file but with a ".log" extension. Successful export operations will not be logged.

Policy and Charging Maintenance

Topics:

- *Introduction.....264*
- *Policy and Charging Maintenance Pages.....264*
- *Alarms, KPIs, and Measurements.....273*
- *Overload Management.....274*
- *Shutdown.....277*
- *Diameter Maintenance and Status Data for Components, DSR Applications, and DA-MPs.....278*
- *Backup and Restore for Policy and Charging Configuration Data.....279*

This chapter describes or indicates where to find the following information that can be used for the Policy and Charging application and SBR:

- Maintenance and status information that is maintained by the Policy and Charging Configuration and Maintenance components and displayed on the **Policy and Charging Maintenance** pages.
- Maintenance and status data that is maintained by Diameter for Diameter Configuration components, DSR Applications, and DA-MPs and displayed on the **Diameter Maintenance** GUI pages.
- Descriptions of Policy and Charging and SBR alarms, KPIs, and measurements
- Auditing of the Session and Binding databases
- Policy and Charging and SBR overload management
- Database Backup and Restore of Policy and Charging configuration data

Introduction

This chapter describes:

- *Policy and Charging Maintenance Pages* describes maintenance and status data that is maintained by the Policy and Charging application and by Policy and Charging DA-MPs.

On the **Policy and Charging > Maintenance** pages, the user can:

- Monitor and maintain the SBR Database Status
- Monitor and maintain the SBR Status
- Define and execute a Policy Database Query
- Monitor and maintain the SBR Database Reconfiguration Status
- *Diameter Maintenance and Status Data for Components, DSR Applications, and DA-MPs* describes maintenance and status information that is maintained by the Diameter Routing Function and the Diameter Transport Function for the Diameter Configuration components that are used to make egress Request message routing decisions.

The **Diameter > Maintenance** pages include status information for:

- Peer Nodes
- Connections
- DSR Applications (including Policy and Charging)
- DA-MPs
- *Alarms, KPIs, and Measurements* describes Policy and Charging-specific database alarms, and indicates the location of descriptions of PCA and SBR alarms, KPIs, and measurements.
- *PCA Data Auditing* describes the auditing of the Session and Binding databases.
- *Overload Management* describes overload controls and load shedding and for PCA and SBR.
- *Backup and Restore for Policy and Charging Configuration Data* describes the OAM database backup and restore of Policy and Charging configuration data.

Policy and Charging Maintenance Pages

The Policy and Charging > Maintenance GUI pages on the NOAM display SBR status information and provide access to the Policy Database Query tool.

SBR Database Status

The **Policy and Charging > Maintenance > SBR Database Status** page is used to prepare, enable, disable, and monitor the status of SBR Databases that have been configured on the **Policy and Charging > Configuration > SBR Databases**.

Fields are described in *SBR Database Status elements*.

SBR Database Status elements

Table 56: SBR Database Status Pane elements describes the elements on the **SBR Database Status** page.

The main SBR Database Status pane shows the details for a configured SBR Database.

Table 56: SBR Database Status Pane elements

Elements	Description
Database Name	The name of the configured SBR Database
Administrative State	The state that the operator commands the SBR Database to be in (Disable, Prepare, or Enable)
Operational Status	The state that the SBR Database is actually in, as reported by the resource users and resource providers of the database
Resource User Operational Reason	The status of the SBR Database as reported by all resource users of the database
Resource Provider Operational Reason	The status of the SBR Database as reported by all resource providers (such as the active server of each SBR server group hosting the database) of the database
Reconfiguration in Progress	Set to No if no SBR Database Reconfiguration is in progress that contains this database. Set to Resizing if a Resizing Plan for that database is in progress. Set to Data Migration if an SBR Data Migration Plan specifying the SBR Database as either the Initial Database or the Target Database is in progress
Database Type	Set to Session if the SBR Database Type is Session. Set to Binding if the SBR Database Type is Binding. Suffixed with (Target) if the SBR Database is the target of an SBR Data Migration Plan

The Resource User pane and the Resource Provider panes provide additional details when the Resource User Operational Status or Resource Provider Operational Status for an SBR Database are colored yellow or red, which indicates a problem. If no problems exist, the Resource User and Resource Provider panes are empty.

Managing an SBR Database

In order to use the SBR Database, perform the following steps using the **SBR Database Status** screen.

1. Select the new SBR Database row and click the **Prepare** button. Preparing the database allows the user to verify connectivity between the SBR Database and all users of that SBR Database before the SBR Database is used by the PCA application.

If the operational status changes to Prepared, then

- All ComAgent connections applicable to the database are connected
- An active SBR server is available for each Policy and Charging SBR Server Group in the SBR Database's Resource Domain
- All resource users have received notification that all portions of the SBR Database (one portion per SBR server group) are available

Once the SBR Database is in the Prepare administrative state, then the operational status can take on one of the following values:

- Preparing - At least one resource user of the SBR Database cannot successfully access at least one portion of the SBR Database
 - Prepared - All resource users of the SBR Database can successfully access all portions of the SBR Database
2. Select the Prepared SBR Database row and click the **Enable** button. Enabling the database allows the PCA application to begin using the database to store and retrieve binding or session records

If the operational status was Preparing at the time the **Enable** button was clicked, then some resource users cannot successfully access the SBR Database. A confirmation dialog appears to warn that the database is not entirely ready to be Enabled. The user can choose to cancel the Enable or proceed despite the warning. If the user proceeds, the operational status in the Enable state will likely be either Unavailable or Degraded.

Once the SBR Database is in the Enable administrative state, the **SBR Database Status** screen can be used to monitor the status of the database. An Enabled database will have on of the following operational statuses:

- Normal - Everything related to the SBR Database is functioning normally
 - Degraded - At least one resource user of the SBR Database cannot successfully access at least one portion of the database
 - Unavailable - None of the resource users of the SBR Database can successfully access any portion of the database
 - Congested - At least one of the resource users of the SBR Database is experiencing congestion for at least one portion of the SBR Database. This problem could be due to congestion at an active SBR server of the database, or due to congestion of a ComAgent connection between the resource user and an active SBR server. Both Degraded and Unavailable take precedence over Congested. For example, if an SBR Database is both Degraded and Congested, then the Operational Status will be Degraded.
3. If it becomes necessary to Disable an SBR Database, select the row for the desired SBR Database on the **SBR Database Status** screen and click the **Disable** button.

If the SBR Database was in the Enable administrative state prior to clicking **Disable**, a confirmation dialog appears to warn about impacts on the PCA application. The user can choose to cancel the **Disable** or proceed despite the warning.

If the SBR Database was in the Prepare administrative state prior to clicking **Disable**, no confirmation dialog appears since the database is not being used by PCA.

Once the **Disable** button is clicked, the operational status takes on one of the following values:

- Users Detaching - waiting for all resource users to detach from the SBR Database
- Providers Detaching - waiting for all resource providers to detach from the SBR Database
- Disabled - all resource users and resource providers are detached from the SBR Database and the SBR Database may be safely deleted

SBR Status

The **Policy and Charging > Maintenance > SBR Status** page displays a collapsed or expanded detailed report for SBR. The data is displayed within Server Groups by configured Place Associations.

Fields are described in [SBR Status elements](#).

SBR Status elements

[Table 57: SBR Status elements](#) describes the elements on the **SBR Status** page, which displays SBR Server Status data within Server Groups that are assigned to each type of Place Association.

Each tab name was configured on the **Configuration > Place Associations** GUI page.

Table 57: SBR Status elements

Elements	Description	Data Input Notes
PCA Binding Region tab	<p>A list of all configured Server Groups that are assigned to the Binding Region Place Association.</p> <p>The Resource Domain Name and the Resource Domain Profile of each Server Group is shown.</p> <p>The Server Name, Resource HA Role of the Server, the server's Congestion Level, and a list of Sub Resources Hosted by the server are shown for each Server in the expanded list.</p>	<p>The page is view-only.</p> <p>The Server Group in each row under the tab can be expanded or collapsed by clicking on the + symbol, to list the Servers that are assigned to that Server Group.</p>
PCA Mated Site tab	<p>A list of all configured Server Groups that are assigned to that Mated Pair Place Association.</p> <p>The Resource Domain Name and the Resource Domain Profile of each Server Group are shown.</p> <p>The Server Name, Resource HA Role of the Server, the server's Congestion Level, and a list of Sub Resources Hosted by the server are shown for each Server in the expanded list.</p> <p>Note: There can be multiple Mated Site tabs</p>	

Policy Database Query

Use the **Policy and Charging > Maintenance > Policy Database Query** page to enter a value for an individual query for a specified binding key. The tool queries the Binding database to determine if the binding key exists.

- If the binding key exists, a report is generated that includes the PCRF that the key is bound to and information about which Diameter session or sessions are associated with that binding key.

The returned session information includes all other binding keys that were included in the session, the session creation time, and the session last touched time.

- If the queried binding key does not exist, an error message is displayed..

Note: The Policy Database Query tool can be used only with Gx sessions. It is not applicable to Rx sessions.

The fields are described in [Policy Database Query elements](#).

To use the Policy Database Query tool,

1. On the Active NOAM, select **Policy and Charging > Maintenance > Policy Database Query**.
2. Select the **Binding Key Type** in the pulldown list.
3. Enter the **Binding Key** value to search for.
4. Click **Search**.

The report appears on the page.

To enter another query, click **Clear**, and select and enter the values for the new search.

Policy Database Query elements

Table 58: Policy Database Query elements describes the elements on the **Policy and Charging > Maintenance > Policy Database Query** page.

Table 58: Policy Database Query elements

Elements (* indicates a required field)	Description	Data Input Notes
* Binding Key Type	Select the type of binding key data entered in the Binding Key field.	Format: Pulldown list Range: IMSI, MSISDN, IPv4 Address, IPv6 Address Default: N/A
* Binding Key	Enter the binding key string to search for. Note: If Binding Key Type field is set to '--Select--', the Binding Key field is disabled.	Format: Text box. Valid characters are letters (a-z, A-Z), digits (0-9), dots (.), colons (:), and hyphens (-). Range: 1-256 characters. <ul style="list-style-type: none"> • IMSI (1-15 digits) • MSISDN (1-15 digits) • Valid IPv4 Address • IPv6 Address (Address representation type 2 as described in RFC 4291 Section 2.2.) Note: If the complete IPv6 Address is not known, enter only the first 4 sets of 16-bit words, followed by a double-colon; for example, .db3:1234:1a:23c::

SBR Database Reconfiguration Status

The **Policy and Charging > Maintenance > SBR Database Reconfiguration Status** page is used to prepare, cancel, start, complete, and monitor SBR Database Resizing Plans and SBR Data Migration

Plans (collectively referred to as Reconfiguration Plans) that were configured on the **Policy and Charging > Configuration > SBR Database Resizing Plans** and **Policy and Charging > Configuration > SBR Data Migration Plans** screens.

Fields are described in *SBR Database Reconfiguration Status elements*.

SBR Database Reconfiguration Status elements

Table 59: Reconfiguration Status Pane elements describes the elements on the **SBR Database Reconfiguration Status** page.

The main Reconfiguration Status pane shows the details for a configured SBR Reconfiguration Plan.

Table 59: Reconfiguration Status Pane elements

Elements	Description
Plan Name	A name that uniquely identifies the Reconfiguration Plan
Plan Type	Indicates whether a Reconfiguration Plan is Resizing or Data Migration
Administrative State	The state that the reconfiguration Plan is desired to be in (Planned, Prepare, Start, Cancel)
Progress	Indicates the number of records remaining in the Initial Resource Domain or Initial SBR Database as the migration progresses out of the number of records that were in the Initial Resource Domain or Initial SBR Database when the Reconfiguration Plan was started
Operational Status	The state that the Reconfiguration Plan is actually in, as reported by the resource users and resource providers of the database
Resource User Operational Reason	The status of the Reconfiguration Plan as reported by all resource users of the database
Resource Provider Operational Reason	The status of the Reconfiguration Plan as reported by all resource providers (i.e. the active server of each SBR server group hosting the database) of the database.

The Resource User pane and the Resource Provider panes are empty when no row is selected in the Reconfiguration Status pane. If a row is selected in the Reconfiguration Status pane and the Resource User Operational reason cell for that row has a red or yellow background, the Resource User pane displays rows showing the operation status for all resource users that are not in the desired state according to the Administrative State of the row. Similarly, if the Resource Provider Operational Reason cell for that row has a red or yellow background, the Resource Provider pane displays rows for all resource providers that are not in the desired state.

Managing an SBR Reconfiguration Plan

The **SBR Database Reconfiguration Status** screen allows Reconfiguration Plans to be managed and monitored. As soon as a Reconfiguration Plan is successfully configured, it appears on the **SBR Database Reconfiguration Status** screen in the Planned administrative state. From this point onwards

until the reconfiguration plan is completed, the **SBR Database Reconfiguration Status** screen is used to control and monitor the migration.

The basic steps for any Reconfiguration Plan are as follows:

1. Prepare the Reconfiguration Plan
2. Start the Reconfiguration Plan
3. Monitor the progress of Reconfiguration Plan
4. Complete the Reconfiguration Plan

All rows of the **SBR Database Reconfiguration Status** screen are automatically refreshed every 10 seconds. The automatic refresh can be paused by clicking the **Pause Updates** checkbox. The screen can also be refreshed manually by navigating away from the **SBR Database Reconfiguration Status** screen, then returning.

The five buttons on the screen are disabled unless a row is selected in the Reconfiguration Status pane. If a button is clicked when the row is not in the appropriate state for that button, an error message is displayed indicating why the button is not valid for the row. For example, a Reconfiguration Plan cannot be Started unless it has been Prepared first.

The NOAM knows from configuration data and the Reconfiguration Plan Administrative State which servers are resource users and resource providers for a given Reconfiguration Plan. If the NOAM detects missing merge data from any of the expected resource users or resource providers, the NOAM server reports these servers as Unknown because the status of the servers is not known. Server merge data can be missing at the NOAM if there is a problem with the server itself, or if there is a problem with the merge data communication between the MP server and the NOAM server.

Each Reconfiguration Plan starts out in the Planned Administrative State with Progress indicating Not Started. The Operational status is Ready if all resource users have received notification from all relevant resource providers that the resource provider is available. If any resource user has not received a notification that its resource providers are available (for example, due to problems with the resource provider, or to ComAgent connection problems), the Operational Status is Not Ready.

Reconfiguration Plans in the Planned Administrative State have Resource User and Resource Provider Operational Reason values that correspond to the information seen on the **SBR Database Status** screen for the SBR Database of a Resizing Plan, or the Initial SBR Database for a Data Migration Plan. These values because only the status of the initial resource is examined when in the Planned State.

As long as a Reconfiguration Plan is in the Planned Administrative State, the Reconfiguration Plan may be deleted or changed. Deletion of a Reconfiguration Plan causes the row for that plan to disappear from the **SBR Database Reconfiguration Status** screen. Changing a Reconfiguration Plan has no effect until the plan is prepared for execution again.

Predicting the Duration of an SBR Reconfiguration Plan

During execution of an SBR Database Resizing Plan or SBR Data Migration Plan, data is migrated by creating new records on the target Resource Domain or SBR Database. The old record created prior to starting the reconfiguration plan in the initial Resource Domain or SBR Database are removed over time as the sessions or bindings are terminated normally by policy or charging Diameter signaling.

Because the migration is driven by normal Diameter signaling, predicting the duration of an SBR Reconfiguration can be difficult unless accurate information is available for the hold-times of sessions and bindings.

Some sessions may last longer and prevent the migration from completing in the predicted time frame. In such cases, the migration can be forced to end or accelerated by requesting that remaining sessions be removed by sending RAR messages including a Session-Release-Cause AVP.

Histogram measurements provide reliable information on session and binding hold-times. Each measurement provides a count for a specific time duration. Refer to the *DSR Measurements Reference* for details on the specific measurements.

Preparing a Reconfiguration Plan

The first step in executing a Reconfiguration Plan is to prepare the plan. This is accomplished by selecting the desired Reconfiguration Plan and clicking the **Prepare** button. If **Prepare** is clicked when the Reconfiguration Plan is in the Not Ready Operational Status, a confirmation dialog appears to explain that it is recommended to resolve resource user and resource provider problems prior to preparing the Reconfiguration Plan. The confirmation dialog allows the user to cancel or proceed anyway.

When a Reconfiguration Plan is prepared, some validations are performed to ensure that the Reconfiguration Plan is configured correctly and that the initial SBR Database is enabled. If these checks pass, the resource users are instructed to begin communicating with the resource providers of the Target Resource Domain or Target SBR Database. This action allows for resolution of any configuration or communication problems prior to actually starting the Reconfiguration Plan.

Cancelling a Reconfiguration Plan

Prior to clicking the **Start** button to start execution of a Reconfiguration Plan, the **Cancel** button can be used to cancel preparation for the Reconfiguration Plan if necessary.

Selecting a Reconfiguration Plan that is in the Prepare Administrative State and clicking **Cancel** causes first the resource users, then the resource providers, to detach from the Target Resource Domain or Target SBR Database and return the plan to the Planned Administrative State. This has no effect on the running SBR Database that was to be reconfigured.

Note: Detaching normally occurs very quickly. It is unlikely that Operational Status of Users Detaching or Providers Detaching will be captured by refreshing the **Reconfiguration Status** screen.

If some of the Resource Users do not detach from the Target resources, then there are likely replication problems preventing the detach command from reaching those resource users, or merging problems preventing those resource users from reporting that they have detached.

If a Resource Provider does not detach from the Target resources, then there are likely replication problems preventing the detach command from reaching that resource provider, or merging problems preventing that resource provider from reporting that it has detached.

After a Reconfiguration Plan is cancelled, the plan can be altered or deleted.

Starting a Reconfiguration Plan

After the Reconfiguration Plan is in the Prepare Administrative State, the SBR Reconfiguration Plan can be started by selecting the row and clicking the **Start** button.

Note: Once migration is started, it cannot be reversed.

If the **Start** button is clicked for a Reconfiguration Plan with Operational Status of Preparing, a confirmation dialog appears advising to resolve resource user and resource provider problems prior to starting the migration. The user can choose to cancel the starting of the migration and resolve problems, or choose to continue and start the migration anyway. If the **Start** button is clicked for a Reconfiguration Plan with Operational Status of Prepared, the Administrative State is set to Start, the Operational Status is set to Migration Started, and the migration of SBR data begins.

After starting a Reconfiguration Plan the Operational Status changes to Migration Started for a short period before transitioning to Migration Coordinated. The reason for the Migration Started status is

to solve race conditions that could occur as all resource users begin executing data migration logic. When the **Start** button is clicked, the NOAM replicates a command down to all resource users to begin data migration. Since this command cannot reach every resource user at the exact same instant, some resource users may begin migration before others are aware that migration is starting. This difference can result in split bindings in which one session for a subscriber and APN is routed to one PCRF and another session for the same subscriber and APN is routed to a different PCRF. To avoid this race condition, the resource users are put into a special mode where they create records in both the initial and target resources. This mode lasts until all resource users report that they have started migration, or until a timer expires in less than a minute and the NOAM commands them to go to the Migration Coordinated mode.

Once in the Migration Coordinated mode, the resource users create all new records in the Target Resource Domain or Target SBR Database and look for all records there first. If a record is not found in the Target, another query is launched to the Initial Resource Domain or Initial SBR Database. The Migration Coordinated mode continues until the data migration is completed.

Monitoring the Progress of an SBR Reconfiguration Plan

As soon as a Reconfiguration Plan is in the Start Administrative State, the Progress cell of the row is updated to indicate the number of records remaining in the Initial Resource Domain or Initial SBR Database as the migration progresses out of the number of records that were in the Initial Resource Domain or Initial SBR Database when the Reconfiguration Plan was started.

The progress numbers are refreshed by the SBR audit at a statically configured interval (defaulting to every 10 seconds). These progress updates can be seen each time the **SBR Database Reconfiguration Status** screen is refreshed, allowing the progress of the Reconfiguration Plan to be monitored.

The number of records is decremented in the initial SBR database when a session or binding records is deleted through Policy and Charging Diameter signaling.

The Reconfiguration Plan is complete when the Progress indicates there are zero records remaining and the Operational Status is changed Data Migration Complete - Click Complete to Finish.

Note: Because the SBR audit is designed to audit one table at a time, there will be periods when no progress updates are reported if there are multiple tables to be audited.

Completing an SBR Reconfiguration Plan

Once a Reconfiguration Plan is started, records are migrated from the Initial Resource Domain or Initial SBR Database to the Target Resource Domain or Target SBR Database as session or binding records are created and removed as a result of PCA Diameter signaling. Once all the records are removed from the initial SBR resource domain and the operation status changes to Data Migration Complete, click the **Complete** button to finish the migration. Clicking **Complete** once the migration has finished causes the following to occur:

- Administrative State changes to Complete
- Resource Users detach from the initial resource
- After Resource Users have detached, Resource Providers detach
- For an SBR Database Resizing Plan, the Resource Domain of the SBR Database is replaced with the Target Resource Domain of the Resizing Plan
- For an SBR Data Migration Plan, the Initial SBR Database is Disabled and the Target SBR Database is Enabled
- For an SBR Data Migration Plan, the Initial SBR Database is set to the Target SBR Database in the Data Migration Plan

Some long lived sessions or bindings can cause the data migration to take longer to complete. To force the migration to complete, the **Complete** button can be clicked while the Reconfiguration Plan is still in the Migration Coordinated Operational Status. If the **Complete** button is clicked, a confirmation dialog appears with options:

- Cancel the attempt to Complete the migration and let the migration proceed
- Force an end to the migration, abandoning any un-migrated records in the Initial Resource Domain or the Initial SBR Database

If there are a lot of records remaining to be migrated, cancelling is recommended.

Forcing an end to the migration is not recommended because any records that have not yet been migrated will be deleted even though the sessions or bindings they represent may still be valid. If these records are deleted, signaling failures or split bindings may result.

Accelerating an SBR Reconfiguration Plan

Accelerating the migration of an SBR Reconfiguration Plan (by clicking the **Accelerate** button) is generally a better option than cancelling an SBR Reconfiguration Plan.

Accelerating the migration for a binding database reconfiguration causes the binding audit to look for non-migrated bindings. If such a binding is found, an RAR is sent to the policy client that created the binding, including a Session-Release-Cause AVP asking the policy client to terminate the session. When the policy client terminates the session and creates a new session, that record is created in the target and is successfully migrated.

Accelerating the migration for a session database reconfiguration causes the session to look for non-migrated binding capable sessions. If such a session is found, an RAR is sent to the policy client that created the binding capable session, including a Session-Release-Cause AVP asking the policy client to terminate the session. When the policy client terminates the session and creates a new session, that record is created in the target and is successfully migrated. Acceleration of a session database reconfiguration has no effect on binding dependent or binding independent sessions. However, such sessions are short-lived and are less likely to hinder migration.

Accelerating migration causes signaling disruptions since binding capable sessions are being torn down and recreated. However, split bindings do not occur.

Alarms, KPIs, and Measurements

This section describes the type of alarm, KPI, and measurements information that is available for the Policy and Charging Application's combination of Policy DRA, Online Charging DRA and SBR, as well as how to access the information in the DSR GUI.

Policy and Charging and SBR Alarms and Events

The Policy and Charging Application SBR alarms and events are described in the *Alarms and KPIs Reference* and the DSR online help for alarms and events.

Active alarms and events and alarm and event history can be displayed on the **Alarms & Events > View Active** and **Alarms & Events > View History** GUI pages.

PCA and SBR KPIs

Key Performance Indicators, or KPIs, provide a means to convey performance information to the user in near real-time. All the KPIs for the Policy and Charging Application and SBR are displayed on the **Status & Manage > KPIs** GUI page. Selecting the tab for a server and a label under the tab displays the KPI information for the selected server.

The PCA and SBR KPIs are described in the *DSR Alarms and KPIs Reference* and the DSR Alarms and KPIs online help.

Policy and Charging and SBR Measurements

Measurements for the Policy and Charging Application and SBR are collected and reported in various measurement groups.

A measurement report and a measurement group can be associated with a one-to-one relationship. A measurements report can be generated with report criteria selected on the **Measurements -> Reports** GUI page.

The *DSR Measurements Reference* and online help explain the report selection criteria, and describe each measurement in each measurement group.

Overload Management

The Policy and Charging Application (PCA) provides mechanisms to manage the overload and congestion that can occur on the Policy and Charging Application and SBR. The PCA might receive ingress messages at a rate higher than the engineered capacity. The internal queues on the PCA might experience higher utilization level than configured. The same might happen on the SBR servers, directly or indirectly resulting from the overloaded traffic from the network or from the PCA.

Overload Controls

The SBRs that implement the Session and Binding databases must protect themselves from becoming so overloaded that they cease to perform their function. There are two parts to achieving this goal:

- Detecting the overload condition and severity
- Shedding work to reduce load.

Policy and Charging DA-MP Overload Control

The number of ingress messages (both Requests and Answers) per second received by PCA is counted as input to PCA ingress message processing capacity. The capacity is an engineering number of ingress messages per second processed by PCA. The number of Request messages received at PCA per second is also measured separately.

PCA defines alarms on the queue utilization levels based on configured threshold values. Thresholds (in percentage) are configured in association with the PCA ingress message capacity. If the ingress message rate received at PCA exceeds the configured percentage of the maximum capacity, alarms

will be raised. PCA ingress Request capacity can be engineering configured to provide the value based on which thresholds (in percentage) are configured. See [Alarm Settings](#).

The PCA congestion is then defined by the ingress Request messages capacity and the configured threshold values. PCA will be considered in congestion if the ingress Request rate at PCA exceeds the configured percentages (thresholds) of PCA ingress Request capacity.

Three PCA congestion levels (CL_1, CL_2 and CL_3) are defined, each of them is associated with onset and abatement threshold values. The onset and abatement values are configurable (see [Congestion Options](#)). When PCA is in congestion, a PCA congestion alarm will be raised at the severity (Minor, Major or Critical) corresponding to the congestion level (CL_1, CL_2 or CL_3).

When congestion is detected, PCA will perform overload control by throttling a portion of incoming messages to keep PCA from being severely impacted. The type and percentage of the messages to be throttled will be configurable through the PCA GUI as displayed in [Table 60: PCA Default Overload Control Thresholds](#):

Table 60: PCA Default Overload Control Thresholds

PCA Operational Status	Alarm-ID 22721			PCA Congestion Level	PCA Message Throttling Rules
	Severity	Onset Threshold	Abatement Threshold		
Available	N/A	N/A	N/A	CL0	No messages are discarded (Accept and process 100% Request and Answer messages)
Available	Minor	110%	100%	CL1	<ul style="list-style-type: none"> Discard 25% of requests for creating new sessions Discard 0% of requests for updating existing sessions Discard 0% of requests for terminating existing sessions Discard 0% of answer messages
Available	Major	140%	130%	CL2	<ul style="list-style-type: none"> Discard 50% of requests for creating new sessions Discard 25% of requests for updating existing sessions Discard 0% of requests for terminating existing sessions Discard 0% of answer messages

The PCA's internal congestion state contributes to PCA's Operational Status directly, along with its Admin state and Shutdown state. Consequently, the congestion state of the PCA impacts the Diameter Routing Function message transferring decision. Depending on the PCA's Operational Status (Unavailable, Degraded, Available), the Diameter Routing Function will forward all the ingress messages to the PCA when the PCA's Operational Status is Available, or discard some or all of the ingress messages when the Operational Status is Degraded or Unavailable. [Table 61: Diameter Routing Function Message Handling Based on PCA Operational Status](#) describes the Diameter Routing Function handling of the messages to the PCA.

Table 61: Diameter Routing Function Message Handling Based on PCA Operational Status

PCA Operational Status	Diameter Routing Function Message Handling
Available	Forward all Request and Answer messages to PCA
Degraded	Forward all Answer messages only to PCA
Unavailable	Discard all messages intended for PCA

SBR Congestion

SBR relies on ComAgent for resource monitoring and overload control. The ComAgent Resource Monitoring and Overload Framework monitors local MP's resource utilizations, defines MP congestion based on one or multiple resource utilizations, communicates the MP congestion levels to Peers, and reports local MP congestion level to the local application (SBR).

Messages called "stack events" are used for communication to and from ComAgent.

ComAgent defines MP congestion levels based on a CPU utilization metric and ingress stack event rate (number of stack events received per second at local ComAgent), whichever is higher than the pre-defined congestion threshold, and broadcasts the MP congestion state to all its Peers. ComAgent provides APIs that the local SBR can call for receiving congestion level notifications.

SBR congestion is measured based on the SBR CPU utilization level. There are four SBR congestion levels: CL0 (normal), CL1 (Minor), CL2 (Major) and CL3 (Critical). There are related Onset and Abatement threshold values, and Abatement time delays.

The SBR congestion state (CPU utilization) is managed and controlled by the ComAgents on both PCA and SBR MPs based on the ComAgent MP Overload Management Framework. Messages to a SBR from a PCA are handled based on the congestion state of the SBR. A SBR congestion alarm will be raised when MP congestion notification is received from ComAgent. The appropriate alarm severity information will be included in the notification. The alarm will be cleared if the congestion level is changed to Normal, also indicated in the notification from ComAgent.

In order to manage the overload situation on a SBR, all stack event messages are associated with pre-defined priorities. Before a stack event message is sent, its priority will be compared with the congestion level of the SBR to which the stack event is sent. If the priority is higher than or equal to the SBR current congestion level, the message will be forwarded. Otherwise, it will be discarded.

The stack events may also be routed from a SBR to another SBR in some scenarios. The congestion control in this case should be conducted based on the congestion state of the receiving SBR, i.e. the ComAgent on the sending SBR is responsible to compare the stack event priority with the congestion level of the receiving SBR and make the routing decision accordingly.

Load Shedding

After the SBR has determined that it is in overload (CL1 – CL3), it informs ComAgent that its resources and sub-resources are in congestion. ComAgent then broadcasts this information to all of the resource users for the specified resources and sub-resources. The resource users now begin to shed load by sending only certain requests for database updates. The resource users determine which database requests to discard based on the current congestion level of the resource provider.

Database requests are delivered to SBRs using ComAgent stack events. Each stack event has a priority. The resource user software (on either DA-MPs or SBRs) sets the stack event priority for every Stack

Event it sends, depending on the type of stack event and the circumstances under which the Stack Event is being used. For example, the same stack event may be used for signaling and for audit, but may have a different priority in each circumstance. The Stack Event priority is compared with the congestion level of the server that is the target of the stack event to determine whether stack event should be sent, as shown in [Table 62: Stack Event Load Shedding](#).

Table 62: Stack Event Load Shedding

Congestion Level	Description
CL0	The resource provider is not congested. No load shedding occurs. Send all Stack Events.
CL1	Minor congestion. Auditing is suspended. Send all Stack Events not related to auditing.
CL2	Major congestion. No new bindings or sessions are created. Existing bindings and sessions are unaffected. Send only Stack Events related to existing sessions.
CL3	Critical congestion. Send only Stack Events already started and Stack Events that remove sessions or bindings.

Shutdown

DA-MP- The Policy and Charging Application running on DA-MPs supports the DSR Application Infrastructure graceful shutdown with 5 seconds grace period. This means that when PCA is Disabled (using the **Diameter->Maintenance->Applications** GUI page), the application will transition to the Degraded Operational Status for 5 seconds to allow in-flight messages to be processed without accepting any new Requests before transitioning to the Unavailable Operational Status. In the Unavailable status, neither Requests nor Answers are processed by the PCA.

SBR - Because SBR servers use the Active/Standby/Spare redundancy model, and ComAgent supports reliable transactions, there is no need for a graceful shutdown mode. Shutdown of a SBR server will cause a failover to another server in the same Server Group. (The exception is if the Server Group only has one server, as might be the case in a small demo system.)

The PCA Operational Status (Unavailable, Degraded and Available) is determined by its Admin State, Congestion Level, and the Shutdown State. The PCA calculates and maintains its own operational status and reports it to the Diameter Routing Function.

When the PCA is not processing requests (in Operational Status of Degraded or Unavailable), the Diameter Routing Function will attempt to route new Requests using the rules in the Peer Routing Tables. If the Request has no Destination-Host AVP, as would be the case for session-initiating Requests, the routing will fail and the Diameter Routing Function will respond with a 3002 DIAMETER_UNABLE_TO_DELIVER Answer.

When a Server is "Stopped" using the Stop function on the **Status & Manage -> Server** GUI page, Diameter will terminate all Diameter connections by sending a DPR and waiting for the DPA. If all DPAs have not been received within 15 seconds, Diameter begins termination of its layers and queues. If Diameter is still not shut down after another 15 seconds, the process is abruptly terminated.

To properly shut down a PCA DA-MP server,

1. Go to the Diameter -> Maintenance -> Applications GUI page and Disable the PCA application.
The Operational Status of the application will transition to Unavailable
2. Go to the **Status & Manage -> Server** page and Stop the Server's application processes.
After 30 seconds maintenance can proceed as necessary.

Table 63: PCA Operational Status shows an example of the PCA Operational Status determination where the Shutdown mode is Graceful Shutdown. The Shut down and Shutting down in the Operational Reason column indicate the states where the (Graceful) shutdown process has been completed (Shut down) and is in progress (Shutting down) respectively. While the Graceful Shutdown is in progress, the PCA continues to process the messages in its queue for a time interval that is engineering configurable.

Table 63: PCA Operational Status

Admin State	Congestion Level	Shutdown State	Operational Status	Operational Reason
N/A	N/A	N/A	Unavailable	Not initialized
Disabled	0 ,1, 2, 3	False	Unavailable	Shut down
Disabled	0 ,1, 2, 3	True	Degraded	Shutting down
Enabled	0 1 2	N/A	Available	Normal Available with CL_1 Available with CL_2
Enabled	3	N/A	Degraded	Congested with CL_3

SBR - Because SBR servers use the Active/Standby/Spare redundancy model, and ComAgent supports reliable transactions, there is no need for a graceful shutdown mode. Shutdown of a SBR server will cause a failover to another server in the same Server Group. (The exception is if the Server Group only has one server, as might be the case in a small demo system.)

Diameter Maintenance and Status Data for Components, DSR Applications, and DA-MPs

Maintenance and status data is maintained and displayed on the following Diameter > Maintenance GUI pages for Diameter Configuration components, DSR Applications including Policy and Charging, and DA-MPs including those that run the Policy and Charging application:

- **Route Lists Maintenance** - The **Diameter > Maintenance > Route Lists** page displays information about the Route Groups assigned to Route Lists. Route List maintenance and status data is maintained and merged to the OAMs. The data is derived from the current Operational Status of Route Groups assigned to a given Route List. The Operational **Status** of each Route List determines whether the Route List can be used for egress routing of Request messages.
- **Route Groups Maintenance** - The **Diameter > Maintenance > Route Groups** page displays the configured and available capacity for Route Groups and displays information about Peer Nodes or Connections assigned to a Route Group.

This information can be used to determine if changes need to be made to the Peer Node or Connection assignments in a Route Group in order to better facilitate Diameter message routing. Additionally, this information is useful for troubleshooting alarms.

Note:

Policy and Charging will create and add one metadata record to the TTR for each event that occurs while any Diameter message in the transaction is being processed. This function is achieved through Policy and Charging support of IDIH.

- **Peer Nodes Maintenance** - The **Diameter > Maintenance > Peer Nodes** page provides the Operational Status of Peer Node connections, including a Reason for the status.
- **Connections Maintenance** - The **Diameter > Maintenance > Connections** page displays information about existing connections, including the Operational Status of each connection.

The **Diameter > Maintenance > Connections > SCTP Statistics** page displays statistics about paths within an SCTP connection. Each line on the page represents a path within an SCTP connection.

- **Applications Maintenance** - The **Diameter > Maintenance > Applications** page displays status, state, and congestion information about activated DSR Applications. The data is refreshed every 10 seconds.

On the **Diameter > Maintenance > Applications** page, you can change the Admin State of the selected DSR Application to Enabled or Disabled.

- **DA-MPs Maintenance** - The **Diameter > Maintenance > DA-MPs** page provides state and congestion information about Diameter Agent Message Processors.

On the **Diameter > Maintenance > DA-MPs** page,

- The Peer DA-MP Status tab displays Peer status information for the DA-MPs.
- The DA-MP Connectivity tab displays information about connections on the DA-MPs.
- The tab for each individual DA-MP displays DA-MP and connection status from the point-of-view of that DA-MP.

The **Diameter > Reports > MP Statistics (SCTP) Reports** GUI page displays the Message Processor (MP) SCTP statistics per MP, for all MPs or for a selected set of MPs. Each row shows the statistics for one MP.

Diameter Maintenance is described in more detail in the *Diameter User Guide* and in the Diameter Help.

Backup and Restore for Policy and Charging Configuration Data

Because the Policy and Charging Application is required to run on a 3-tier OAM topology where some data is mastered at the NOAM and some data is mastered at SOAMs at each site, backup and restore must be performed on the NOAM and on the SOAMs at each site.

Only configured data is backed up and restored. Dynamic data such as policy and policy charging sessions and policy bindings that is mastered on SBR MP servers is not backed up or restored.

The PCA feature uses the capabilities of the Backup and Restore functions provided by the OAM **Status & Manage > Database** GUI page, as described in the "Database Backups and Restores" chapter of the *DSR Administration Guide*.

PDRA PCRF Pooling Upgrade

Topics:

- [Upgrade Paths.....281](#)
- [Configuration After Upgrade.....283](#)
- [Concepts and Terminology.....284](#)
- [Configuring PCRF Pooling.....285](#)
- [Processing Phases.....288](#)
- [Binding Migration.....290](#)

This section provides information about, and includes procedures for upgrading from an installed and activated, non-PCRF Pooling Policy DRA release.

PCRF Pool enablement applies when upgrading from the DSR 5.0 Policy DRA to the 7.0 functionality. PCRF Pooling can be enabled only after all Policy DRA network elements are upgraded and those upgrades are accepted. Then, it is possible to use PCRF Pooling logic, as the upgrade changes the way that binding data is handled.

Note: In releases prior to 7.0, the application is called "PDRA." In the 7.0 release, the application is called "PCA" and PDRA is a functionality of the PCA application.

Upgrade Paths

Upgrading on Previously Activated Policy DRA Releases

If PCA is already activated on a DSR that has been upgraded to the release that supports PCRF Pooling and the PCA activation occurs after the upgrade is completed and accepted, the following conditions apply:

- Diameter must be configured according to the appropriate release documentation.
- PCA feature must be activated.
- Policy DRA function must be enabled
- Policy DRA must be configured.
- PCRF Pooling must be configured.
 - The PCRF Pooling capability is enabled by default and cannot be disabled.
 - A Default PCRF Pool is pre-configured and cannot be deleted. This PCRF Pool can be used or not used, similarly to the Default PRT table.
 - The Default PCRF Pool is not mapped to a PRT table by default. The PCRF Pool to PRT Mapping table uses the **Not Selected** for PRT by default.
 - When Access Point Names are configured, they must be mapped to a configured PCRF Pool.

Activation of PCA on a network where the upgrade is not completed and accepted on all servers is prohibited.

Assuming an upgrade from a previously activated Policy DRA release, the following conditions apply:

- After upgrade to the release that supports PCRF Pooling from a release that did not support PCRF Pooling, all APNs configured prior to the upgrade will be mapped to the PCRF Pool called Default. This can be seen on the NOAMP GUI at **Policy DRA > Configuration > Access Point Names**.
- After upgrade to the release that supports PCRF Pooling from a release that has Policy DRA activated, but did not support PCRF Pooling, the PCRF Pooling functionality is not enabled. This can be seen on the NOAMP GUI at **Policy DRA > Configuration > Network-Wide Options**.
- After upgrade to the release that supports PCRF Pooling from a release that did not support PCRF Pooling, there shall be no PCRF Sub-Pool Selection Rules configured on **Policy DRA > Configuration > PCRF Sub-Pool Selection Rules**.

Upgrading on Previously Non-Activated Policy DRA Releases

This section discusses supported upgrade paths and information related to previously-activated and non-activated Policy DRA releases.

The PCRF Pooling release supports upgrades from 4.1.5 to 5.1 and 5.0 to 5.1.

Note: Incremental upgrades that skip builds might be supported, but only as justified by business needs.

After upgrade to the release that supports PCRF Pooling from a release that did not have Policy DRA activated, the PCRF Pooling functionality is Enabled when Policy DRA is activated. This can be seen on the NOAMP GUI at **Policy DRA > Configuration > Network-Wide Options**.

If Policy DRA was not activated on the release being upgraded to the release that supports PCRF Pooling, the activation is treated like an initial install of PCRF Pooling.

The following steps are required to initiate support of the PCRF Pool feature:

- The Policy DRA application on all DSRs in the network must be upgraded to the point where the upgrade will not be backed out to a version that supports the PCRF Pool feature. The Policy DRA application must be upgraded and the upgrade accepted on all Policy DRA DSR Network Elements.
- After upgrading and prior to enabling, the Policy DRA continues to use 4.1.5 logic. PCRF pooling can be configured at this point, but it is not required.
- As a result of upgrading to a version of the Policy DRA that supports the PCRF Pool feature, a default pool will be in place and all existing APNs will be configured to map to the default pool. The default PCRF pool will point to the existing PRT used for handling new-binding CCR-Is.
- In the case of a new install, the PCRF Pool and PRT must be configured as part of configuring the Policy DRA application.

Note: In the case of an upgrade, existing bindings may have been created before the upgrade.

After all Policy DRA NEs have been upgraded, requests will proceed as shown in the following table.

Table 64: Processing During Transition Period

Request Type	Processing during transition period
CCR-I	<p>If an existing binding matches for the IMSI+APN combination, then route to the PCRF indicated in the binding.</p> <p>Note: Any binding for the IMSI that existed prior to enabling PCRF Pooling will match any IMSI+APN combination for that IMSI.</p> <p>The following logic applies:</p> <ul style="list-style-type: none"> • If binding exists for IMSI from prior to enabling PCRF Pooling, use that binding. • Else if binding exists for IMSI and APN, use that binding. • Else if binding exists for IMSI and PCRF Pool, use that binding. • Else create a new binding using both APN and PCRF Pool.
CCR-U	No change - uses Destination-Host routing
CCR-T	No change - uses binding created by CCR-I
AAR (IPv6)	No change - query IPv6 correlation binding
AAR (MSISDN)	<p>If an existing secondary key matches for the MSISDN+APN combination, then route to the PCRF indicated by the secondary key.</p> <p>Note: Any binding for the MSISDN that existed prior to enabling PCRF Pooling will match any MSISDN+APN combination for that MSISDN.</p> <p>Else, existing behavior for invalid request (binding not found).</p>

Request Type	Processing during transition period
RAR and all other requests	No change

After upgrading to the release that supports PCRF Pooling, but prior to enabling the PCRF Pooling functionality, the following changes to Policy DRA configuration are in place:

- A single PCRF Pool called Default has been created. This is done on the NOAM GUI at **Policy DRA > Configuration > PCRF Pools**.
- All configured APNs are mapped to the Default PCRF Pool. This is done on the NOAM GUI at **Policy DRA > Configuration > Access Point Names**.
- The PCRF Pooling functionality is not Enabled. This is done on the NOAM GUI at **Policy DRA > Configuration > Network-Wide Options**.
- There are no PCRF Sub-Pool Selection Rules configured on the NOAM GUI at **Policy DRA > Configuration > PCRF Sub-Pool Selection Rules**.
- The Default PCRF Pool is mapped at each site to the same PRT table that was configured for new bindings on the SOAM GUI at **Policy DRA > Configuration > Site Options** in the field called **Peer Route Table Name**. The new mapping can be seen on the SOAM GUI at **Policy DRA > Configuration > PCRF Pool to PRT Mapping**.
- The new Error Condition to be used when a binding capable session initiation request arrives with an unconfigured APN or no APN defaults to IANA Diameter response code 3002. This can be seen on the SOAM GUI at **Policy DRA > Configuration > Error Codes** for Error Condition **Missing Or Unconfigured APN**.

Configuration After Upgrade

When a network that is already running Policy DRA is upgraded to DSR 5.1, several changes are performed automatically to prepare for PCRF Pooling, but backwards compatibility is maintained for all aspects of Policy DRA. The following changes occur automatically:

1. New entries appear in the **Policy DRA > Configuration** folder at the NOAM.
 - PCRF Pools
 - PCRF Sub-Pools Selection Rules
2. New entries appear in the **Policy DRA -> Configuration** folder at each SOAM.
 - PCRF Pools (read only at the SOAM)
 - PCRF Pool To PRT Mapping
 - PCRF Sub-Pool Selection Rules (read only at the SOAM)
3. A PCRF Pool called Default is created in **Policy DRA -> Configuration -> PCRF Pools**.
4. All configured APNs are mapped to the Default PCRF Pool. This can be seen on the NOAM GUI at **Policy DRA > Configuration > Access Point Names**.
5. The PCRF Pooling functionality is not Enabled. This can be seen on the NOAM GUI at **Policy DRA > Configuration > Network-Wide Options**.
6. There are no PCRF Sub-Pool Selection Rules configured on the NOAM GUI at **Policy DRA > Configuration > PCRF Sub-Pool Selection Rules**.
7. The Default PCRF Pool is mapped at each site to the same PRT table that was configured for new bindings on the SOAM GUI at **Policy DRA > Configuration > Site Options** in **Peer Route Table**

Name. The new mapping can be seen on the SOAM GUI at **Policy DRA > Configuration > PCRF Pool to PRT Mapping**.

Note: The GUI field for the PRT table for new bindings previously configured at Policy DRA > Configuration > Site Options is deprecated by Pooling.

8. The new Error Condition to be used when a binding capable session initiation request arrives with an unconfigured APN or no APN, defaults to IANA Diameter response code 3002. This can be seen on the SOAM GUI at **Policy DRA > Configuration > Error Codes** for Error Condition "Missing Or Unconfigured APN".

Note:

After upgrading to the release that supports PCRF Pooling, but prior to enabling the PCRF Pooling functionality, the following changes to Policy DRA configuration are in place:

Prior to enabling PCRF Pooling, no request will be rejected due to a missing or unconfigured APN.

Refer to [NOAM Configuration](#) for additional information.

Concepts and Terminology

Policy DRA upgrading incorporates new services or new PCRF infrastructure without disturbing existing services. In releases prior to 5.1, a binding was a mapping between an IMSI and a single PCRF. After a binding existed, all sessions for that IMSI are routed to the bound PCRF. Upgrading to PCRF Pooling allows for multiple bindings to exist for a single IMSI, one for each PCRF pool.

When the release that supports PCRF Pooling is installed, a PCRF Pool called "Default" is automatically created. This PCRF Pool cannot be deleted. It allows backwards compatibility with prior releases in which there was only one logical group of PCRFs, which could be thought of as a single PCRF Pool. When a network is upgraded that already has Policy DRA activated, all configured APNs are mapped to the Default PCRF Pool. The Default PCRF Pool is in turn mapped to whatever PRT table was defined to handle new bindings in the prior release.

Upgrading to a release of Policy DRA that supports PCRF Pools or PCRF Sub-Pools, requires that the Default pool must point to the PRT used for routing of new binding requests prior to the upgrade. The Default PCRF Pool is mapped to the PRT defined to manage new bindings in the prior release.

A graceful upgrade ensures the following:

- Existing bindings are not adversely affected by the in-service upgrade.
- Policy DRA business logic continues to execute the previous release logic until PCRF Pooling is explicitly enabled.
- Split bindings are not created.
- PCRF Pooling configuration can be performed before or after enabling PCRF Pooling with no unexpected behavior.
 - If PCRF Pooling is enabled with no configuration changes, the Policy DRA behavior will be the same as prior to PCRF Pooling being enabled (assuming that all APNs were already configured).
 - If PCRF Pooling is configured prior to enabling PCRF Pooling, existing bindings are honored until they are released normally. Only new bindings are routed according to the PCRF Pooling behavior.

Note: Migration from the pre-PCRF Pools pSBR database to the PCRF Pools pSBR database must be finished prior to upgrading beyond DSR 5.1.

Terminology for Upgrading Policy DRA

The following table shows a list of some common Policy DRA acronyms related to upgrading.

Table 65: Upgrading Policy DRA Terminology

Acronym/Term	Description
Enabling PCRF Pool Feature	Enabling the PCRF Pool Feature is a one-time operation used to begin a transition period from pre PCRF Pool processing to PCRF Pool processing. This is a one-time operation and, once enabled, the PCRF Pool feature can no longer be disabled.
Graceful upgrade	The ability to accomodate upgrades for PCRF Pooling functionality without disruption of existing configurations.
Migration Period	For customers upgrading from a release prior to DSR 5.1 Policy DRA, a migration occurs from the IMSI-only binding table to a table that supports a binding per IMSI-APN combination. In order to avoid Split Bindings, bindings existing in the IMSI only table are honored until they naturally terminate. As existing IMSI-only bindings naturally terminate, they are replaced with IMSI-APN bindings. Once all IMSI-only bindings are gone, the migration period is complete. This data migration also applies to alternate key tables (MSISDN, IPv4 Address and IPv6 Address).
New-Binding CCR-I	A CCR-I request for a specific IMSI, APN combination that occurs when there is not an existing binding SBR record for the IMSI+APN. In this case, a new binding is created for the IMSI+APN.

Configuring PCRF Pooling

Use this task to configure PCRF Pooling.

In order to configure PCRF Pooling, the following steps should be carried out in the order specified. Read through all of the steps prior to beginning configuration.

Note: This procedure assumes that the network is already configured as used in DSR 4.1.5.

1. Configure all Access Point Names.

After upgrade to DSR 5.1 and prior to enabling the PCRF Pooling capability, all APNs used in the network must be configured.

To ensure that all APNs are configured, perform the following from the NOAM:

- a) Use the alarm history at **Alarms & Events > View History** with filter setting as follows:
 1. Scope set to the NOAM Network Element
 2. Display Filter set to Event ID=22730
 3. Collection Interval set to N Days, where N is long enough to cover the period after the upgrade, or something shorter if required.
- b) If the Event History shows any instances of alarm 22730, check the Additional Info portion of the alarm to determine if the configuration problem is related to missing or unconfigured APNs. If this is the case, either configure the unconfigured APN, or determine why the policy client is not sending an APN in the request.
- c) Repeat item 2 in this step for each instance of alarm 22730.

2. Configure DSR routing for each PCRF Pool.

This step pre-configures all of the DSR routing necessary to route new binding requests to PCRF Pools. This configuration will not be used until further configuration is completed later in this procedure.

The routing configuration controls which PCRFs are in a given pool. Policy DRA application software selects a PCRF Pool name, but the DSR routing configuration at each site determines which PCRFs are part of the PCRF Pool.

For each Policy DRA node in the network, perform the following steps using that site's SOAM.

- a) Determine a set of PCRF Peer Nodes that will be grouped into a PCRF Pool.
- b) At each site's SOAM, for the PCRF Peer Nodes from item 1, create a Route Group containing those PCRFs.
- c) For the Route Group in item 2 in this step, create a Route List that uses that Route Group.
- d) For the Route List in item 3 in this step, create a Peer Route Table.
- e) For the Peer Route Table created in item 4 in this step, create a peer routing table rule that will choose the Route List created in item 3.
- f) Repeat items 1 through 5 in this step for each required PCRF Pool.

At this point, DSR routing is configured to route to the PCRF Pools, but the Policy DRA application is not yet aware of the PRT tables, and will not use them.

Note: It is possible that not every pool will be used at every site. If this is the case, you do not need to create routing for that pool.

3. Configure PCRF Pools.

This step creates PCRF Pools. These pools will not be used yet because no APNs are mapped to these pools.

Note: Perform this step at the NOAM.

- a) At **Policy DRA -> Configuration -> PCRF Pools**, insert a PCRF Pool using a descriptive name.
- b) Repeat item 1 in this step for each PCRF Pool.

4. Configure PCRF Pool to PRT Mappings.

This step maps the PCRF Pools created in step 3 above to the PRT tables created in step 2, items 4 and 5. Even though PCRF Pools are now mapped to the DSR routing configuration, none of this configuration will be used until one APN (at a minimum) is changed to use one of the new PCRF Pools.

For each Policy DRA node in the network, perform the following using that site's SOAM.

- a) Use **Policy DRA -> Configuration -> PCRF Pool to PRT Mapping** to map a PCRF Pool to the PRT table created for that pool in Step 2, items 4 and 5.
- b) Repeat item 1 in this step for each PCRF Pool that indicates a Peer Route Table of "Not Selected" until all PCRF Pools are mapped to the appropriate PRT table.

Note: It is possible that not every pool will be used at every site. If this is the case, the PCRF Pool can be left-mapped to the "Not Selected" PRT. For these entries, a warning is issued at the NOAM when an APN is mapped to the PCRF Pool that is not mapped to a PRT. If you are sure that no signaling will be received for any of the APNs mapped to that PCRF Pool at that site, then you can confirm the operation; thus, overriding the warning.

5. Configure the Error Codes for a Missing or Unconfigured APN

This step allows configuration of the Diameter Response Code to use if a request is rejected after PCRF Pooling is enabled because the request contains no APN, or contains an APN that is not configured in the Policy DRA. This step is not required if the default result code of 3002 is appropriate for this error condition.

Perform this step at the SOAM for each site in the network.

- a) Use **Policy DRA -> Configuration -> Error Codes**.
- b) Select the row for Missing or Unconfigured APN and click **Edit**.
- c) Configure the Diameter Result Code to be used for each interface. Leave **Vendor ID** blank if the Result Code is IANA defined.

6. Enable PCRF Pooling

This step enables the PCRF Pooling capability. No routing changes should occur yet because all APNs are still mapped to the Default PCRF Pool.

Perform this step at the NOAM.

- a) At **Policy DRA -> Configuration -> Network-Wide Options**, check **Enable PCRF Pooling**.

Note: PCRF Pooling can be enabled only after every server at every site in the network has been upgraded to DSR 5.1 and the upgrade has been accepted for all servers. If these conditions are not met, PCRF Pooling cannot be enabled.

7. Edit APNs to Begin Using the New PCRF Pools

This step will begin to use the PCRF Pooling functionality by mapping APNs to the newly created PCRF Pools.

Perform this step at the NOAM.

- Use **Policy DRA -> Configuration -> Access Point Names** to edit an APN and change its PCRF Pool from the Default PCRF Pool to the desired PCRF Pool.
- Commit the change. This causes the system to verify that the PCRF Pool is mapped to a PRT table at every site. If the PCRF Pool is not mapped to a PRT at any site, or if the NOAM cannot communicate with one or more SOAMs, a warning is displayed in a confirmation dialog indicating which case applies.
 1. If the NOAM cannot communicate with all SOAMs, investigate and resolve the communications issue before proceeding.
 2. If the PCRF Pool is not mapped to a PRT table at one or more sites, verify that the mapping was intentionally omitted. The mapping should be omitted only if no signaling will occur at the site or sites that do not have the PCRF Pool mapped to a PRT using any of the APNs that are mapped to the PCRF Pool.

- After the mapping from APN to PCRF Pool (other than Default) is committed, verify that new bindings are routed correctly to the PCRFs in the new PCRF Pool, according to the APN. Note that any existing bindings that match the IMSI, or IMSI and APN will be honored until those bindings are terminated by a CCR-T for the last session for the binding.
- Repeat items 1 through 3 for each APN until all are mapped to the required PCRF Pool.

After this step is complete, Policy DRA is fully functioning using PCRF Pooling to route new binding requests.

Processing Phases

On a system that is upgraded to PCRF Pooling, consider the following phases (note the dependencies):

1. System upgraded, but PCRF Pooling not yet enabled
2. PCRF Pooling enabled and database migration in progress (this is applicable if you have upgraded from a prior release where Policy DRA was activated)
3. PCRF Pooling enabled and database migration completed (this phase is equivalent to a new install with PCRF Pooling and is applicable if you have upgraded from a prior release where Policy DRA was activated)

System Ungraded, but PCRF Pooling Not Yet Enabled

After and during the upgrade, but prior to enabling PCRF Pooling, no behavior changes from the prior release.

- All signaling business logic from the prior release is still used.
- All PCRF Pooling data can be safely configured without affecting ongoing signaling.
- All bindings and sessions are maintained over the upgrade.
- All new bindings are created in the old binding tables.
- APN to PCRF Pool mappings are not yet used.
- The APN present in session initiation requests is ignored except for the purpose of establishing the proper Stale Session Lifetime as was done in the prior release.
- All sessions with the same binding key are routed to the same PCRF.

PCRF Pooling Enabled and Database Migration in Progress

PCRF Pooling functionality is enabled from the NOAMP GUI at **Policy DRA > Configuration > Network-Wide Options** by checking **Enable PCRF Pooling**. PCRF Pooling can only be enabled after all servers in the network have been successfully upgraded to the release supporting PCRF Pooling and the upgrade has been accepted on all servers. The GUI will not allow PCRF Pooling to be enabled until this state has been achieved.

After PCRF Pooling is enabled, the following occurs:

- Binding capable session initiation requests arriving with no APN, or an APN that is not configured in **Policy DRA > Configuration > Access Point Names**, are responded to using the Diameter error response configured for the Missing Or Unconfigured APN condition at **Policy DRA > Configuration > Error Codes**.

Note: This does not apply if a binding exists for the IMSI prior to enabling PCRF Pooling. In that case, the signaling will succeed even with no APN.

- PCRF Pool selection occurs.
- The binding database is consulted to determine if a suitable existing binding should be used as follows:
 - If there is a binding in the IMSI-Only table for the IMSI, that binding is used, else
 - If there is a binding in the IMSI+APN table for the IMSI and APN, that binding is used, else
 - If there is a binding in the IMSI+APN table for the IMSI and PCRF Pool, that binding is used, else
 - Create a new binding.
- If a new binding is necessary, it is created using the new IMSI+APN table, and includes the IMSI, APN, and PCRF Pool.
- If a new binding was created, the Policy DRA application asks the routing layer to route using the PRT table mapped to the selected PCRF Pool or Sub-Pool.
- If an existing binding was selected, the Policy DRA application asks the routing layer to route using the PRT precedence as follows:
 - PRT associated with the ingress Peer Node, OR
 - PRT associated with the Diameter application-id, OR
 - The Default PRT, OR finally
 - Connections associated with the egress Peer Node
- When a binding-capable session is successfully established (for example, by success response from PCRF):
 - The PCRF that answered is written to the binding such that all subsequent requests that match the binding are routed to the same PCRF.
 - See [The Binding Database](#).
- Old and new tables for IMSI, IPv4, IPv6, and MSISDN are all audited during the migration period.

PCRF Pooling Enabled and Database Migration Completed

After there are no more records in the old binding tables (`ImsiAnchorKey`, `MsisdnAlternateKey`, `Ipv4AlternateKey`, and `Ipv6AlternateKey`), the migration period is considered to be complete. Note that because these tables are partitioned across a number of binding pSBR server groups, each server group makes the determination independently as to whether migration has completed. There is no global indicator that shows that migration has completed across the entire binding database.

After migration has completed for a binding pSBR server group:

- All new bindings are created in the IMSI+APN, MSISDN+APN, and the new IP Address tables (`ImsiApnAnchorKey`, `MsisdnApnAlternateKey`, `Ipv4AlternateKeyV2`, and `Ipv6AlternateKeyV2`).
- Early Binding Master sessions are explicitly updated when they become Final; there is no more implicit transition to Final.
- All Early Binding polling occurs at the binding database, eliminating the need to route an Early Binding Slave Diameter request to the mated pair of the Early Binding Master session with the PDRA-Early-Binding AVP included.
- Binding dependent session initiation requests using MSISDN as correlation key must include a configure APN, or binding correlation for the MSISDN key will fail.
- Auditing of the IMSI-Only, MSISDN-Only, and old IP Address tables ceases.
- Memory for the portion of the database owned by that server group for the IMSI-Only, MSISDN-Only, and old IP Address tables (actually, the old DB Part fragments) is freed.

Binding Migration

A binding migration period is required in order to successfully create new bindings without interfering with existing bindings.

Handling of Binding-capable Session Initiation Requests

This section describes the Policy DRA handling of binding capable session initiation requests during the binding migration period.

For bindings created after PCRF Pooling is enabled, Policy DRA enforces the requirements for handling missing and unconfigured APN values in binding capable session initiation requests.

Policy DRA allows binding-capable session initiation requests for an IMSI that have no APN, or have an unconfigured APN to be routed according to existing bindings for that same IMSI created before PCRF Pooling was enabled.

Note: The software attempts to find a binding created prior to enabling PCRF Pooling. If such a binding is found, it can be used for routing the new request. If no such binding exists, the binding capable session initiation request is rejected.

Upon receipt of a binding-capable session initiation request for an IMSI that has an existing Final binding, the Policy DRA application attempts to route the request to the PCRF from the selected binding. This process is described below.

When checking for an existing binding, the Policy DRA application searches in the following order, using the first binding that matches:

1. A binding for the IMSI created prior to enabling PCRF Pooling (from the ImsiAnchorKey table)
2. A binding for the IMSI and APN (from the ImsiApnAnchorKey table)
3. A binding for the IMSI and suggested PCRF Pool or Sub-Pool (from the ImsiApnAnchorKey table)

Upon receipt of a binding capable session initiation request at a site that has no PCRFs configured, the following requirements apply:

- If a binding capable session initiation request is received that would result in a new binding and no PCRFs are configured at the site, Policy DRA shall generate an error response with the 3002 Diameter Response-Code and Error-Message AVP including the string "No PCRFs configured at this site."

Note: This requirement does not apply if a binding already exists for the IMSI and APN, or IMSI and PCRF Pool.

- If a binding-capable session initiation request is received and no PCRFs are configured at the site, Policy DRA generates timed alarm 22730, which indicating that no PCRFs are configured.

Note: The alarm is only generated if the binding-capable session initiation request results in a new binding being created.

When routing a binding-capable session initiation request, Policy DRA behaves according to the following requirements:

- Upon receipt of a binding-capable session initiation request for an IMSI for which no existing binding is found, a new binding is created using the IMSI, APN, and suggested PCRF Pool or Sub-Pool.

- If, when creating the new binding, the record for the IMSI already contains 10 session references, the Policy DRA application generates a Diameter error response using the response code configured for the Policy SBR Error condition.

Note: The Error-Message AVP contains the reason for the failure.

- After a binding is successfully created, the Policy DRA application attempts to route the request using the suggested PCRF Pool or Sub-Pool.
- When a binding-capable session initiation request results in a new binding, the binding-capable session initiation request is routed to via the Peer Routing Table mapped to the PCRF Pool or Sub-Pool at the site where the request was received.
- If the PCRF Pool or Sub-Pool is not mapped to a Peer Routing Table (for example, is mapped to the "-Select-" entry) at the site processing the request, the request is routed according to the routing layer PRT precedence.

Note: When the P-DRA application does not specify a PRT table to use, DRL looks for a PRT in the ingress Peer Node configuration, then, if still not specified, in the Diameter Application-Id configuration. This behavior is necessary for backwards compatibility for cases in which the pre-PCRF Pooling release had the Site Options PRT table for new bindings set to "-Not Selected-".

Binding-capable session initiation requests containing no IMSI are handled accordingly:

- If a binding-capable session initiation request is received and the request contains no IMSI, but does contain a configured APN, Policy DRA executes the PCRF Pool selection logic and routes the request using the selected PCRF Pool or Sub-Pool.

Note: A malformed Subscription-Id is treated as if it did not exist. No binding database lookup is attempted here because IMSI is required to do a binding lookup.

- If a binding-capable session initiation request is received and the request contains no IMSI, and does not contain an APN, the request is treated as described below:
 - Upon receipt of a binding-capable session initiation request containing no Called-Station-Id AVP (for example, no APN), Policy DRA generates and sends a binding capable session initiation answer message using the Result Code configured for the Diameter interface for the "Missing Or Unconfigured APN" condition in the Error Codes GUI. The answer message shall include an Error-Message AVP with the 3-digit error code suffix of 500.
 - Upon receipt of a binding-capable session initiation request containing no Called-Station-Id AVP (for example, no APN), Policy DRA asserts Alarm-ID 22730.

Appendix B

PCA Error Resolution

Topics:

- [Introduction.....293](#)
- [Policy DRA Error Resolution Flowchart Summary.....295](#)
- [Online Charging DRA Error Resolution Flowchart Summary.....317](#)

This section provides information to support the Policy DRA error resolution process, including a business logic flowchart summary, a list of flowchart procedures attributes, and individual flowchart examples.

This information focuses on errors that are directly related to the signaling processing Diameter messages for the Policy DRA application. More specifically, these are errors that cause Policy DRA to generate Diameter Answers with Error-Message AVPs where the errors are included.

The flowcharts in this section are intended to support the step-by-step actions that you can take to resolve errors. The corresponding error resolution process is documented in the *Alarms, KPIs, and Measurements Reference*.

The information presented here does not represent a comprehensive error resolution solution, but should be used with the *Alarms, KPIs, and Measurements Reference* and [Error Codes](#) for a more complete understanding of Policy DRA error resolution.

Introduction

Error resolution flowcharts illustrate PCA application business logic that includes information such as when in the processing that the error occurred, where in the logic flow the error occurred, and what type of error has occurred. A relationship tree on each of the flowcharts helps you understand the entire business logic of PCA application and the error location.

You should use the information in the flowcharts with the additional error resolution steps included in the *Measurements Reference*. That document contains detailed error recovery procedures and corresponding alarm, events, and measurements, as well as actions that you can take to resolve errors. In general use the flowcharts as a guide to investigate and understand the circumstances about where the error occurred and potential paths to resolution.

The flowcharts can be used as a navigation tool to guide you through the PCA GUI during error resolution efforts.

Note: See [Error Codes](#) for more information about PCA error conditions and error code numbers.

Understanding Error Resolution Procedure Attributes

[Table 66: Error Resolution Attributes](#) lists the attributes in the resolution procedures that you should be familiar with before using this appendix. See *Measurements Reference* for that specific information about resolution steps.

Note: Although these terms are used in the flowcharts in this section, they are actually more accurately associate with the error resolution process itself.

These attributes provide information to define, categorize, and associate additional data with the errors. When an error is found from tracing functionality or by raised alarms, you can use these attributes (for example, 3-digit error code, Alarm-ID, and so forth) to navigate the error resolution procedure to locate all relevant information about this particular error. Use the information provided by these error-related attributes related to take specific actions for further error resolution.

Note: The error resolution flowcharts do not contain all of these attributes; they are included here for additional clarification.

Table 66: Error Resolution Attributes

Attribute	Procedure
Error Categories / Names	<p>This attribute defines an error category where multiple specific errors belong, or a single error scenario. If the value in the corresponding GUI Configurable attribute is Yes, it is an configured error category defined in the Policy and Charging SOAM GUI (Policy and Charging > Configuration > Error Codes in Error Condition . You can refer to a Result Code related to the error, in addition to the relevant data listed in the recovery procedure.</p> <p>In a single error scenario, if the value in the corresponding GUI Configurable attribute is No,</p>

Attribute	Procedure
	it is an error that causes Policy DRA to generate an error Answer, but it does not correspond to any of the configured error categories. The default Result Code related to this error is listed in Default Result Code .
3-digit Error Code	<p>A 3-digit error code is an identifier that uniquely identifies a specific error scenario (not error category) encountered in a Diameter Answer message generated by Policy DRA.</p> <p>3-digit codes are unique across all DSR layers (DSR connection layer, routing layer, and application layer) and all DSR applications (PCA, RBAR, FABR, IDIH, and so on) for errors the codes represent. The ranges of 500-549 and 850-899 are for PCA application, while the DSR connection layer, routing layer, and other DSR applications use other non-overlapping ranges.</p> <p>Multiple errors can belong to a same error category and are associated with a same Result Code. It is the 3-digit code that distinguishes an error from other errors. Users should search for the 3-digit code when identifying an error if possible and available.</p>
Result Code in PCA Generated Answer	If the error is not GUI configurable, this attribute is the value of the Result Code AVP in the Answer generated by PCA. If the error is GUI configurable, the Result-Code value can be found in PCA SOAM GUI (Policy and Charging > Configuration > Error Codes).
Error Scenario Description	This field contains a detailed description of the error scenario that can be identified by a 3-digit error code.
Applicable Diameter Interface/Message Type	<p>This attribute provides information about the applicable Diameter interfaces or Diameter message types on which the corresponding error could occur.</p> <p>The Diameter interfaces are categorized as binding-capable (Gx/Gxx) or binding-dependent (Rx or Gx-Prime). Together with the PCA business logic flowcharts in this appendix, the Diameter interface and message information lets you narrow the scope where the errors could happen, which helps you locate the root cause of an error.</p>

Attribute	Procedure
Direct PCA Alarm/Event	The PCA Alarm/Event listed in this field is the most direct alarm or event that is launched due to the occurrence of the error. Multiple alarms or events can be generated due to the error, but only the most direct triggering alarm or event is used in this column.
Direct Exception Measurement and Measurement Group	The measurement in this field is the most direct measurement pegged for this error. Multiple measurements can be pegged also, due to the occurrence of the error.
Error Resolution Flowcharts	This attribute lists the chart numbers of the PCA business logic flowcharts indicating specifically where and when in the business logic that the corresponding error might occur. PCA business logic flowcharts are included in this appendix.
Troubleshooting Steps/Customer Actions	See <i>Measurements Reference</i> .

Policy DRA Error Resolution Flowchart Summary

Figure 51: Error Resolution Flowchart Summary shows a summary of the error resolution flowcharts in this section and their relationships to each another. This relationship tree should help you understand the P-DRA business logic as it relates to error resolution task.

Each reference in the flowchart summary points to a corresponding error resolution flowchart in this appendix. Use the following table to reconcile those references:

Chart number	Flowchart Name	Link
Chart 1	Diameter Message Validation	Figure 52: Diameter Message Validation Error Resolution Flowchart
Chart 2	Generic CCR Processing	Figure 53: Generic CCR Processing Error Resolution Flowchart
Chart 3	CCR-I Processing without PCRF	Figure 54: CCR-I Processing without PCRF Pool Error Resolution Flowchart
Chart 4	findSessionRef Processing	Figure 55: findSessionRef Processing Error Resolution Flowchart
Chart 5	findOrCreBindResShort Processing	Figure 56: findOrCreBindResShort Processing Error Resolution Flowchart

Chart number	Flowchart Name	Link
Chart 6	CCR-I Processing with PCRF Pool	Figure 57: CCR-I Processing with PCRF Pool Error Resolution Flowchart
Chart 7	findOrCreateBinding Response Processing with PCRF Pool	Figure 58: findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart
Chart 8	Early Bind Pool	Figure 59: Early Bind Pool Error Resolution Flowchart
Chart 9	CCA-I Processing	Figure 60: CCA-I Processing Error Resolution Flowchart
Chart 10	findSession Response Processing	Figure 61: findSession Response Processing Error Resolution Flowchart
Chart 11	CCR-U Processing	Figure 62: CCR-U Processing Error Resolution Flowchart
Chart 12	CCR-T Processing	Figure 63: CCR-T Processing Error Resolution Flowchart
Chart 13	CCA-U/T Processing	Figure 64: CCA-U/T Processing Error Resolution Flowchart
Chart 14	RAR Processing	Figure 65: RAR Processing Error Resolution Flowchart
Chart 15	RAA Processing	Figure 66: RAA Processing Error Resolution Flowchart
Chart 16	AAR Processing	Figure 67: AAR Processing Error Resolution Flowchart
Chart 17	AAA Processing	Figure 68: AAA Processing Error Resolution Flowchart
Chart 18	STR Processing	Figure 69: STR Processing Error Resolution Flowchart
Chart 19	STA Processing	Figure 70: STA Processing Error Resolution Flowchart
Chart 20	ASR/ASA Processing	Figure 71: ASR/ASA Processing Error Resolution Flowchart

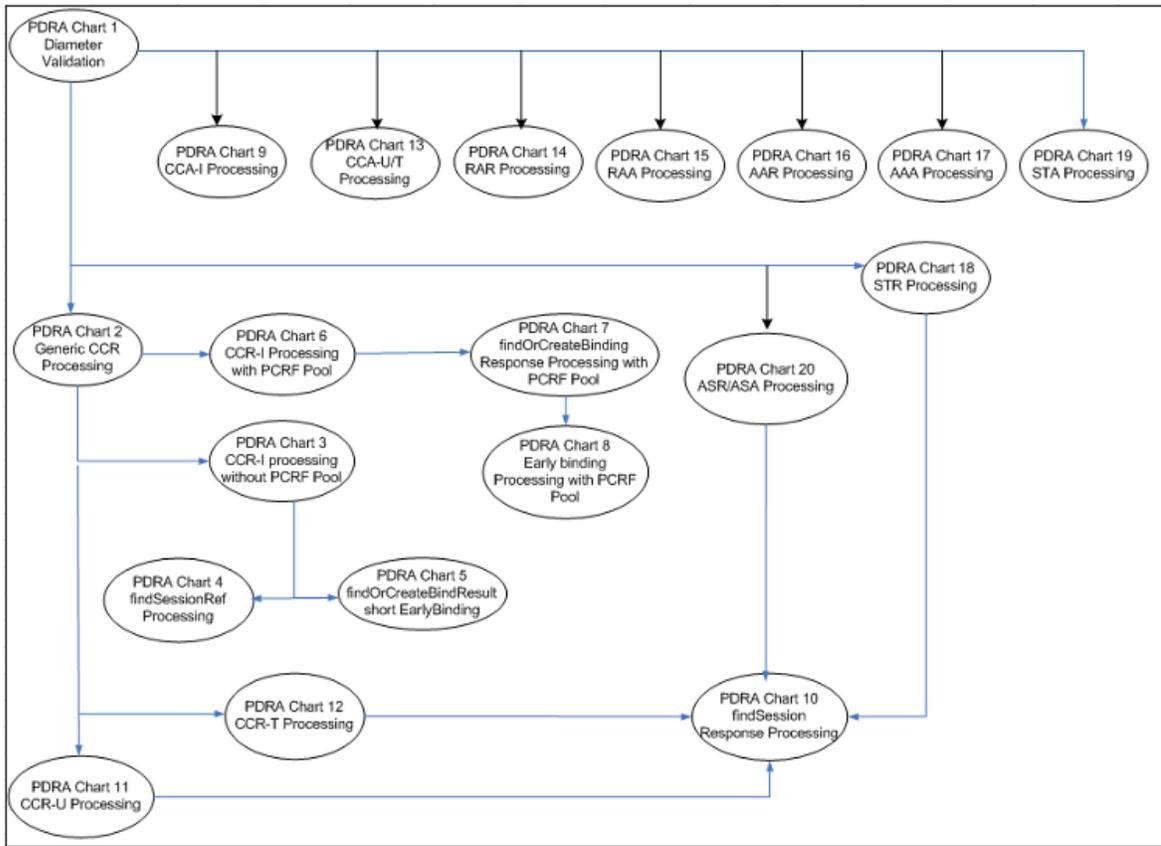


Figure 51: Error Resolution Flowchart Summary

Diameter Message Validation Error Resolution Flowchart

Figure 52: Diameter Message Validation Error Resolution Flowchart shows an error resolution flowchart that illustrates where diameter message validation errors can occur.

Note: See Policy DRA Error Resolution Flowchart Summary for the entire business logic flowchart summary.

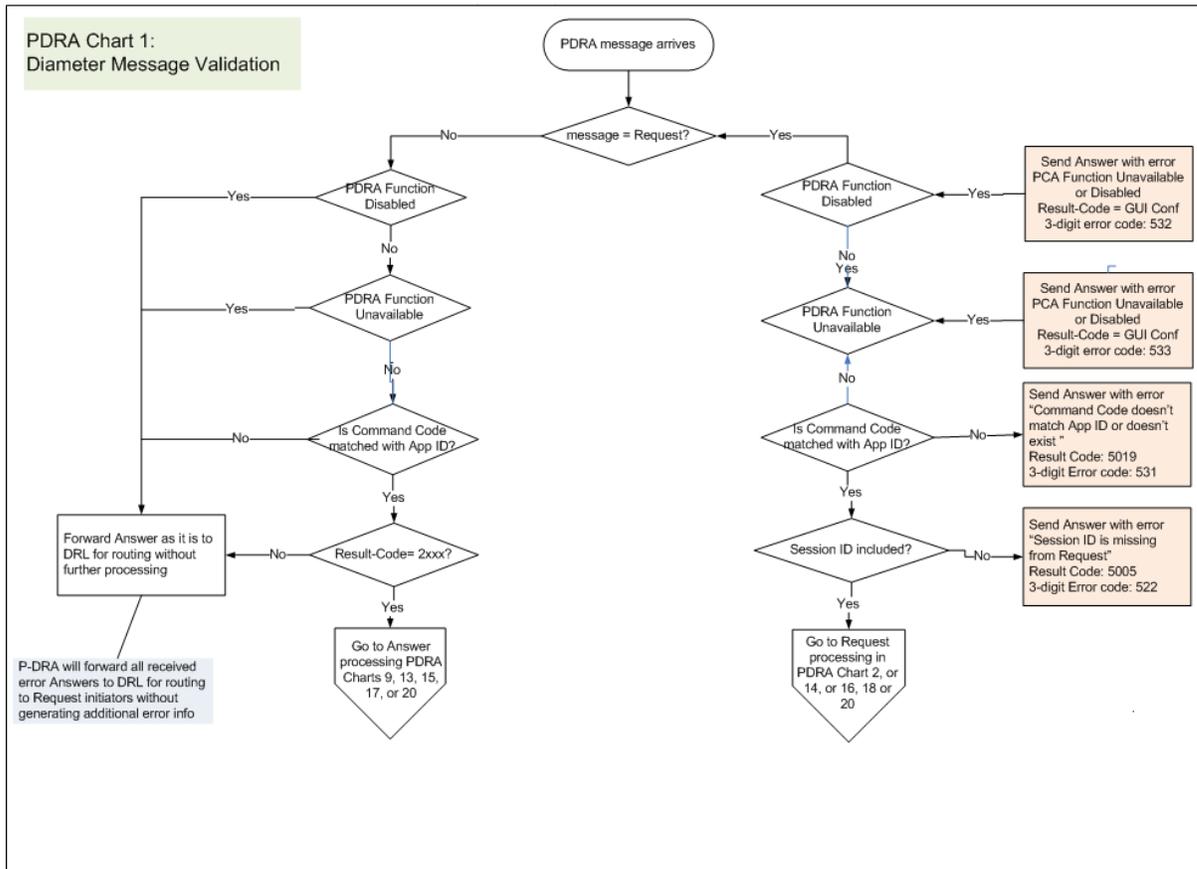


Figure 52: Diameter Message Validation Error Resolution Flowchart

Generic CCR Processing Error Resolution Flowchart

Figure 53: *Generic CCR Processing Error Resolution Flowchart* shows an error resolution flowchart that illustrates where generic CCR processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

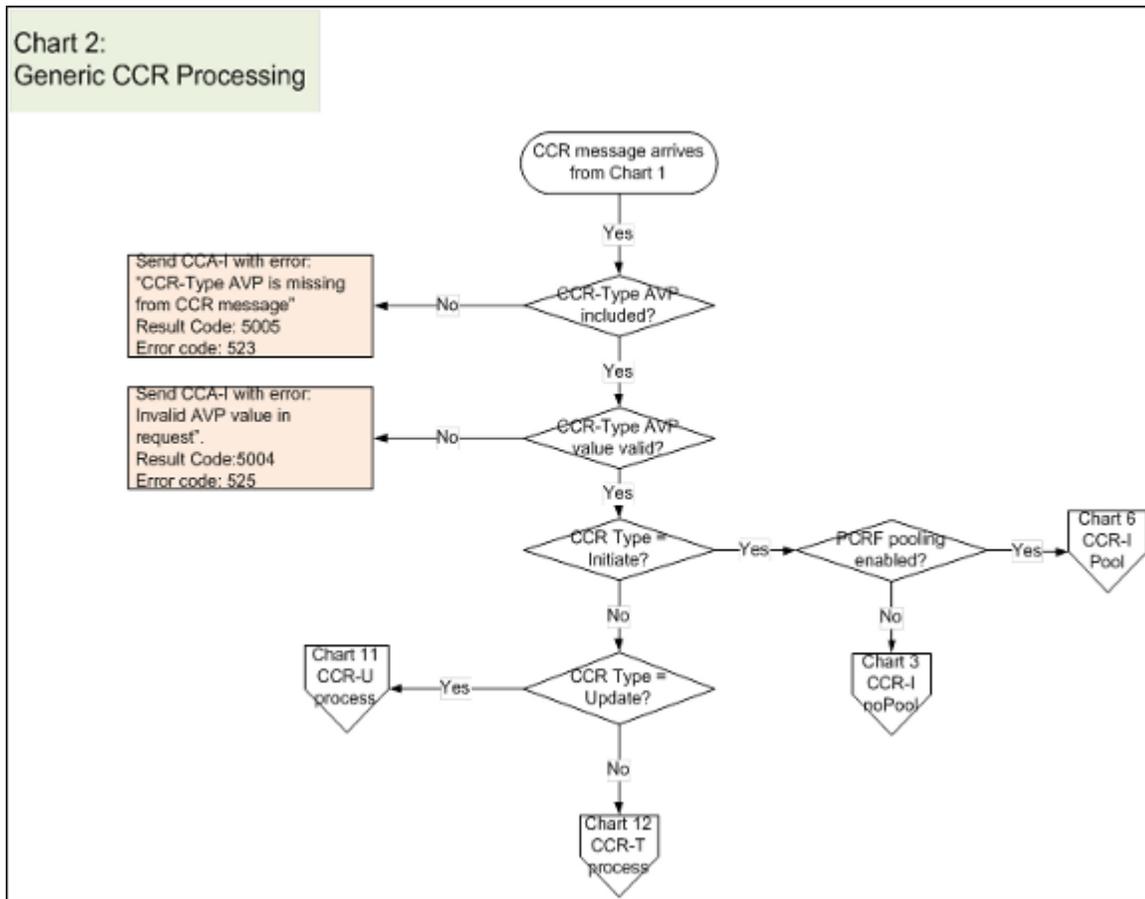


Figure 53: Generic CCR Processing Error Resolution Flowchart

CCR-I Processing without PCRF Pool Error Resolution Flowchart

Figure 54: CCR-I Processing without PCRF Pool Error Resolution Flowchart shows an error resolution flowchart that illustrates where CCR-I processing without PCRF pool errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

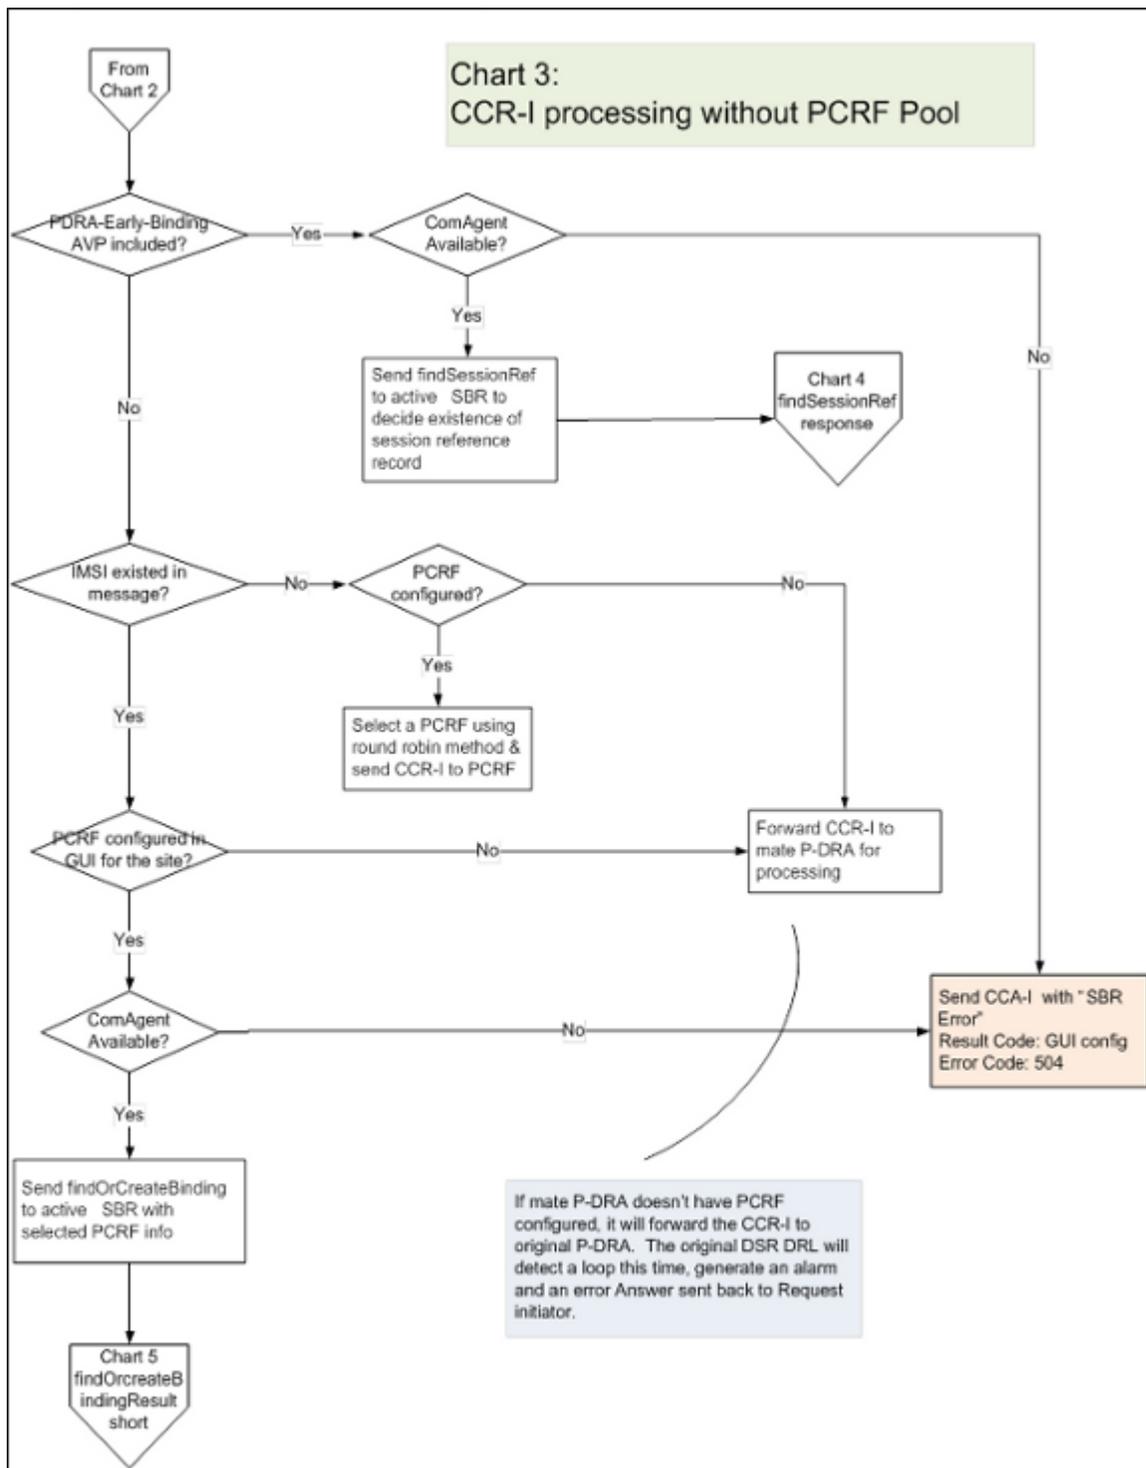


Figure 54: CCR-I Processing without PCRF Pool Error Resolution Flowchart

findSessionRef Processing Error Resolution Flowchart

Figure 55: *findSessionRef Processing Error Resolution Flowchart* shows an error resolution flowchart that illustrates where findSessionRef processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

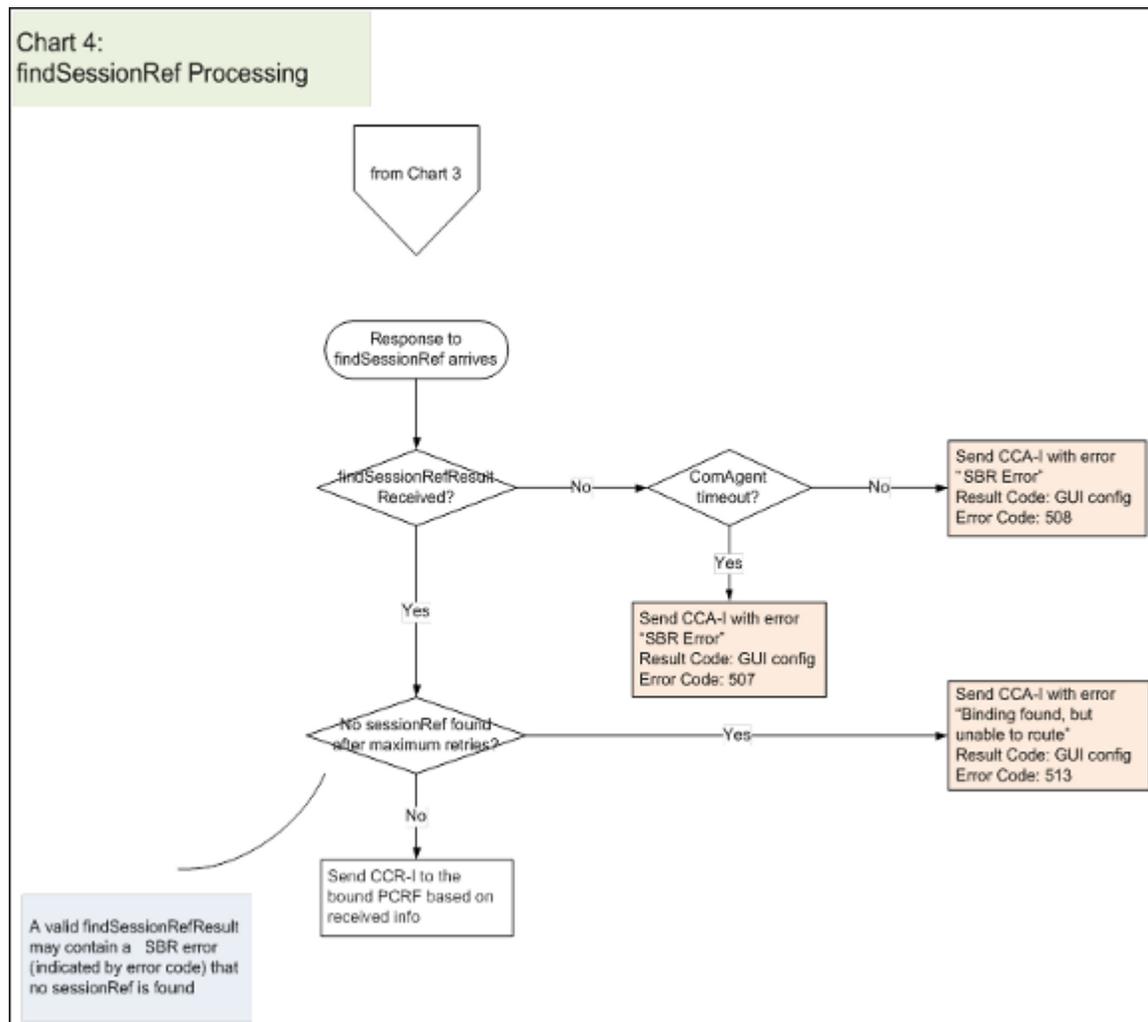


Figure 55: findSessionRef Processing Error Resolution Flowchart

findOrCreBindResShort Processing Error Resolution Flowchart

Figure 56: *findOrCreBindResShort Processing Error Resolution Flowchart* shows an error resolution flowchart that illustrates where findOrCreBindResShort processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

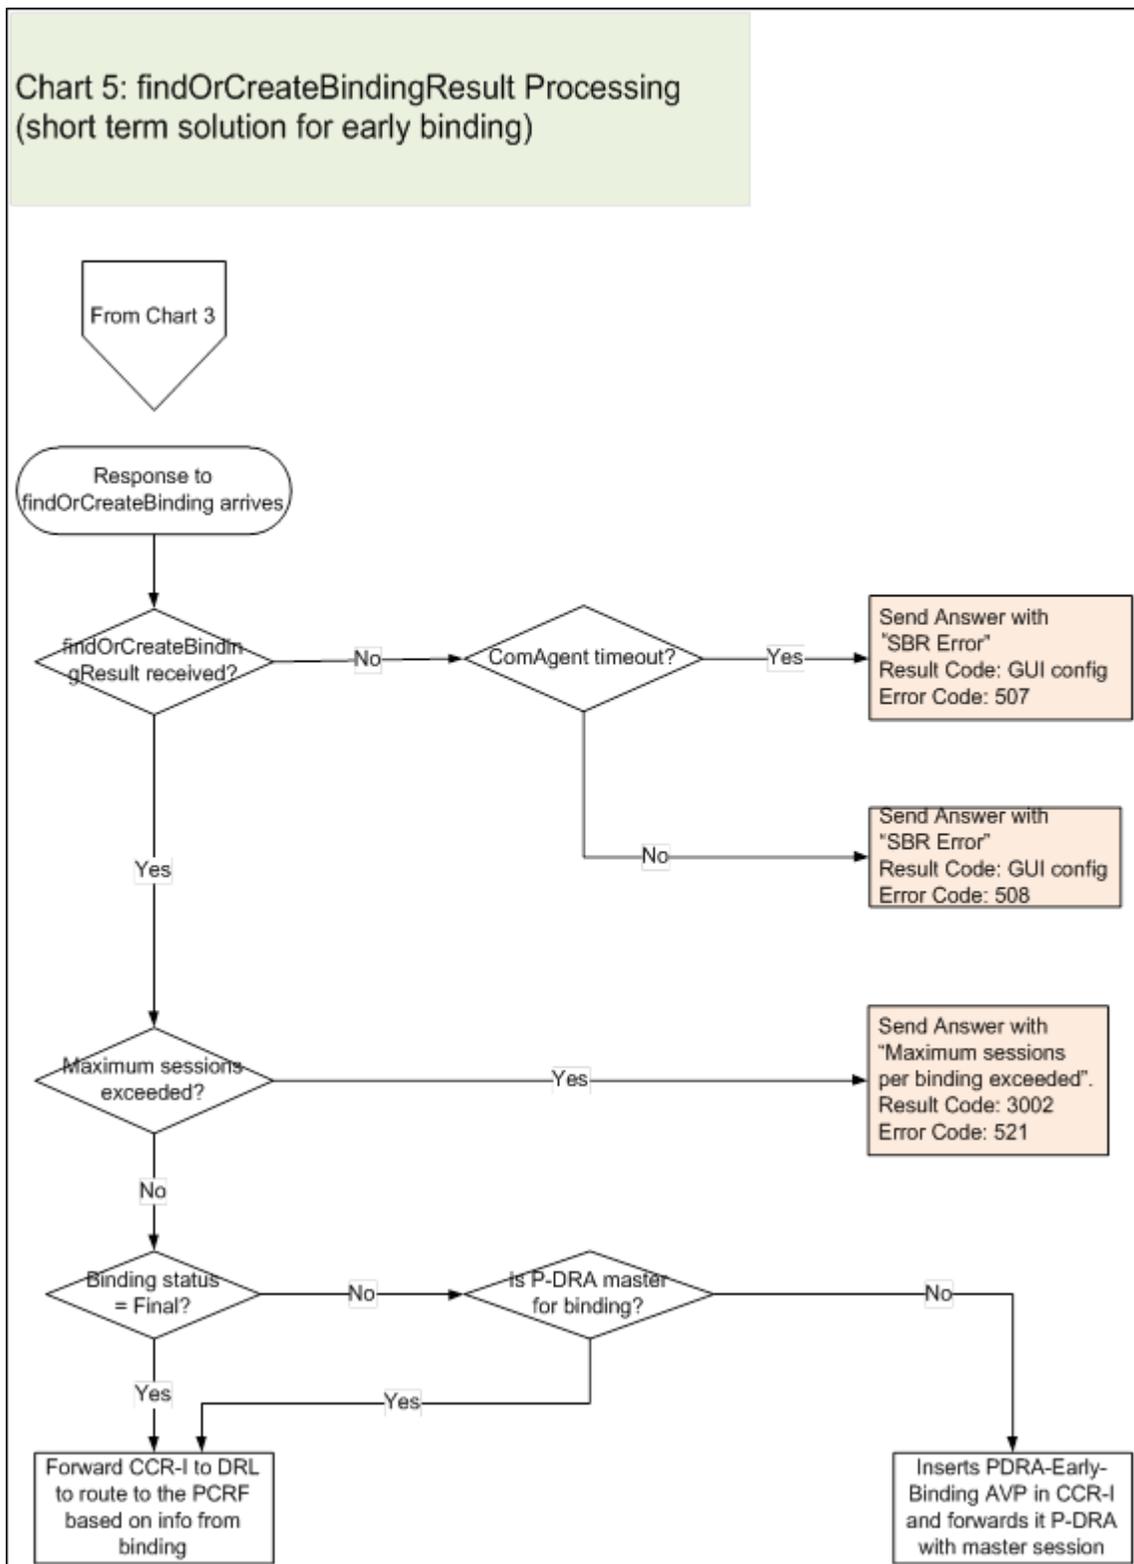


Figure 56: findOrCreBindResShort Processing Error Resolution Flowchart

CCR-I Processing with PCRF Pool Error Resolution Flowchart

Figure 57: CCR-I Processing with PCRF Pool Error Resolution Flowchart shows an error resolution flowchart that illustrates where CCR-I Processing with PCRF Pool errors can occur.

Note: See Policy DRA Error Resolution Flowchart Summary for the entire business logic flowchart summary.

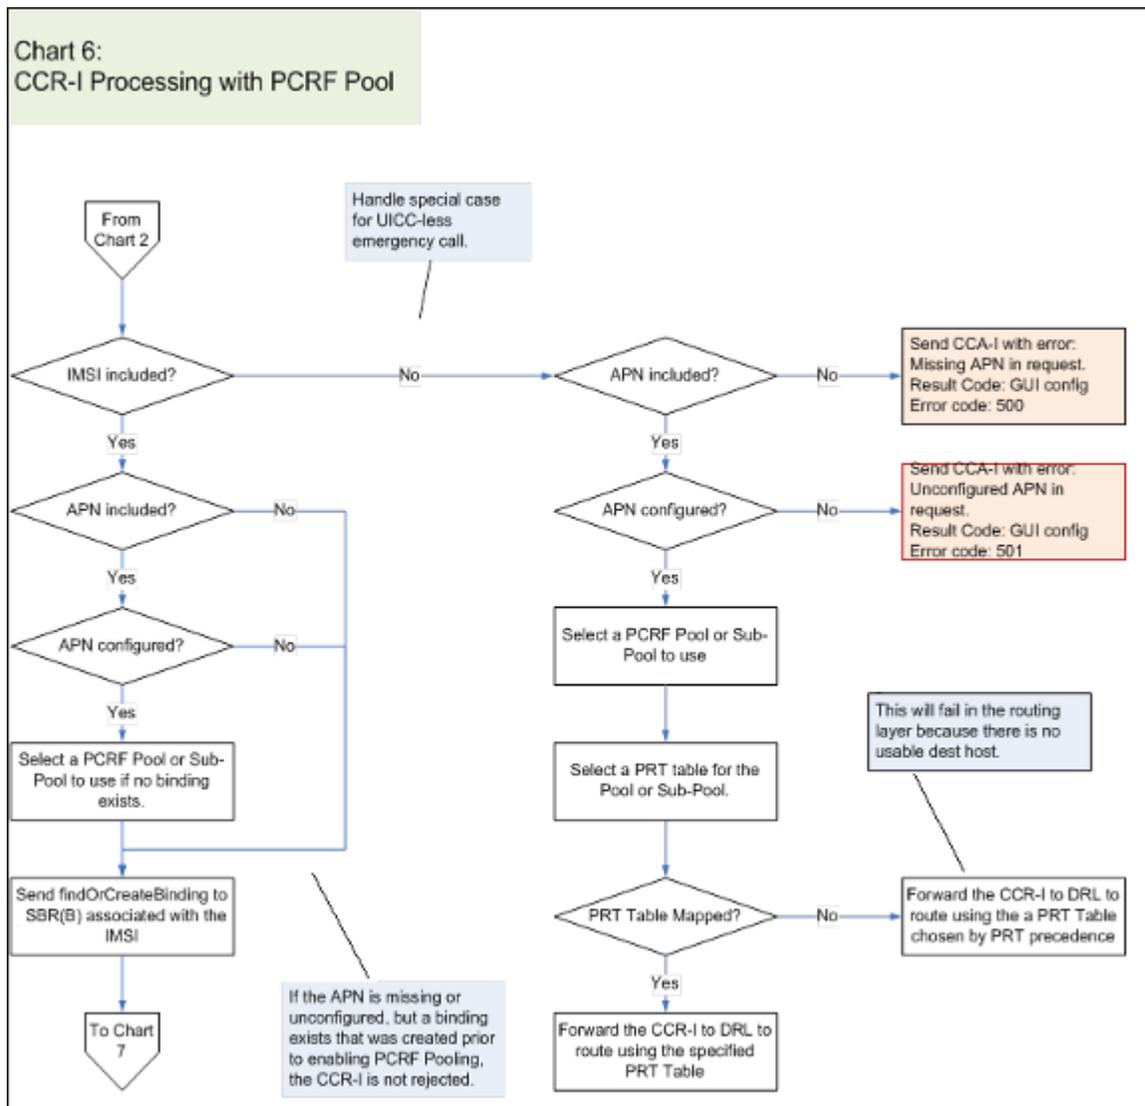


Figure 57: CCR-I Processing with PCRF Pool Error Resolution Flowchart

findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart

Figure 58: findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart shows an error resolution flowchart that illustrates where findOrCreateBinding Response Processing with PCRF Pool errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

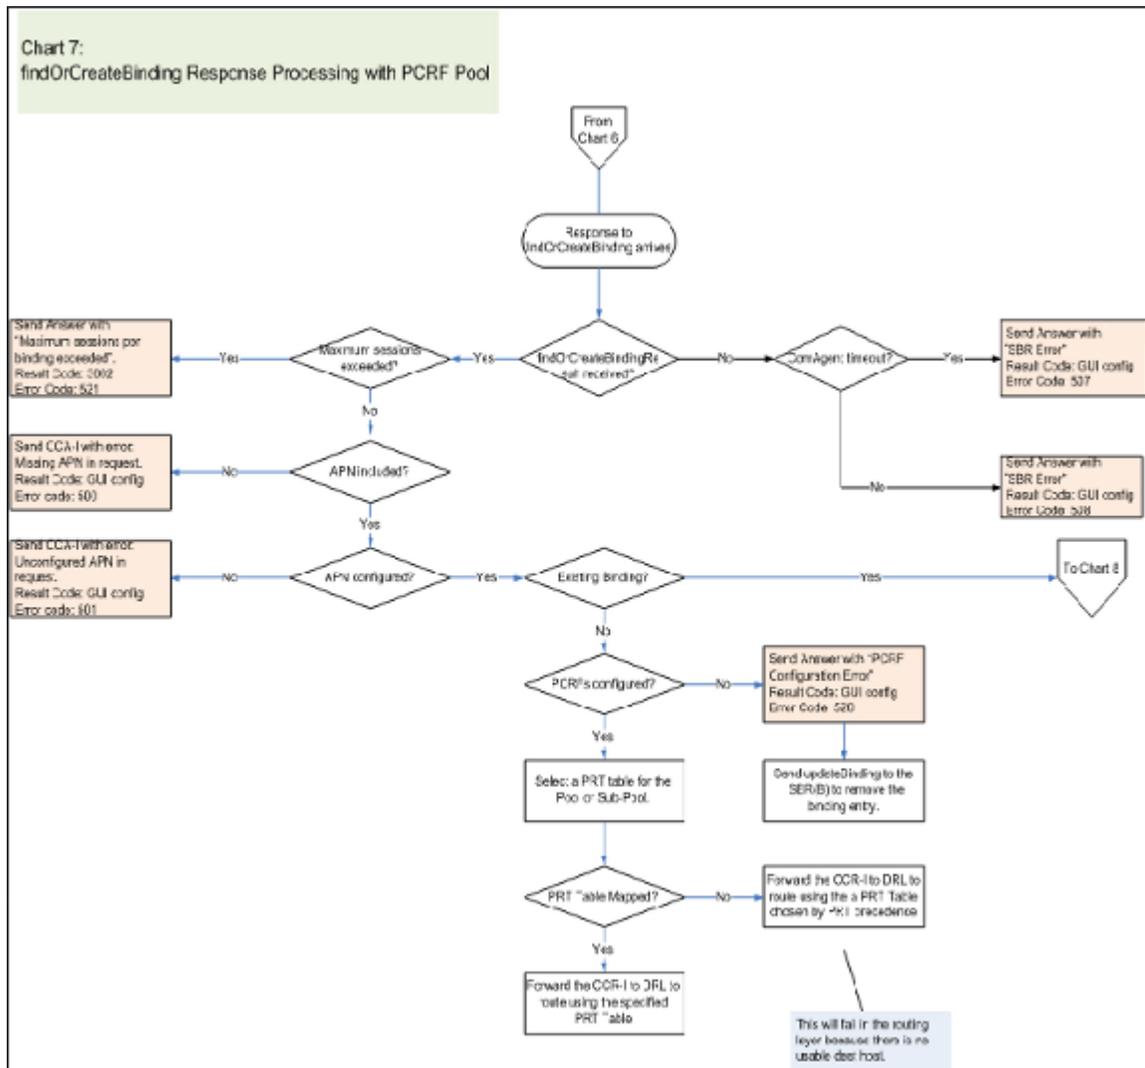


Figure 58: findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart

Early Bind Pool Error Resolution Flowchart

Figure 59: Early Bind Pool Error Resolution Flowchart shows an error resolution flowchart that illustrates where Early Bind Pool errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

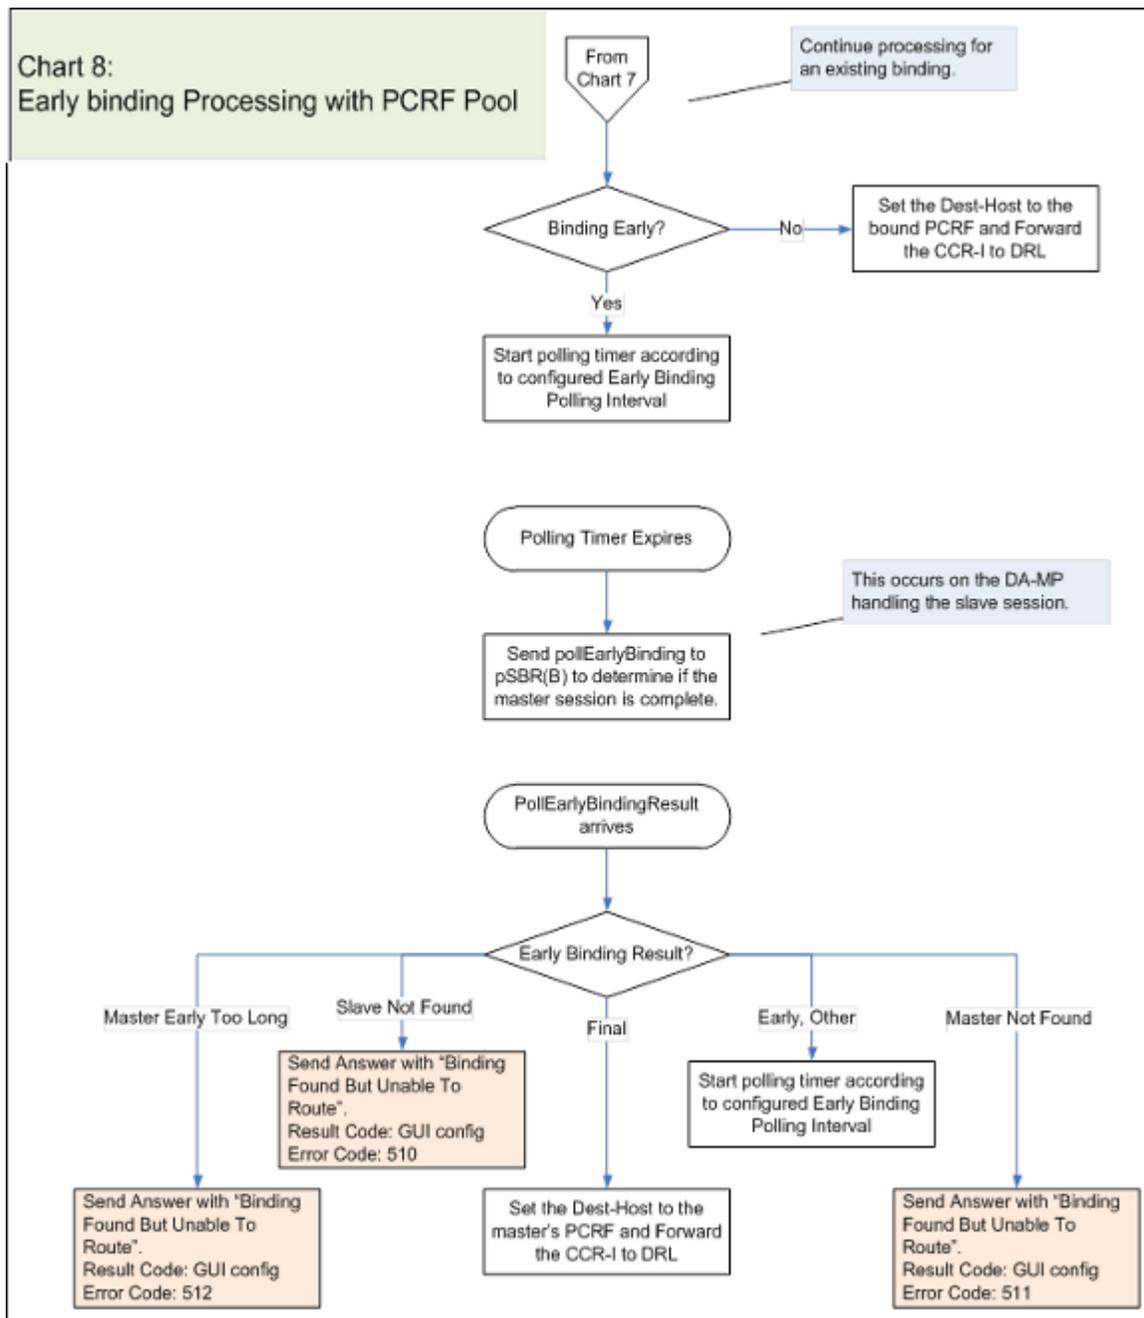


Figure 59: Early Bind Pool Error Resolution Flowchart

CCA-I Processing Error Resolution Flowchart

Figure 60: CCA-I Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where Early Bind Pool errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

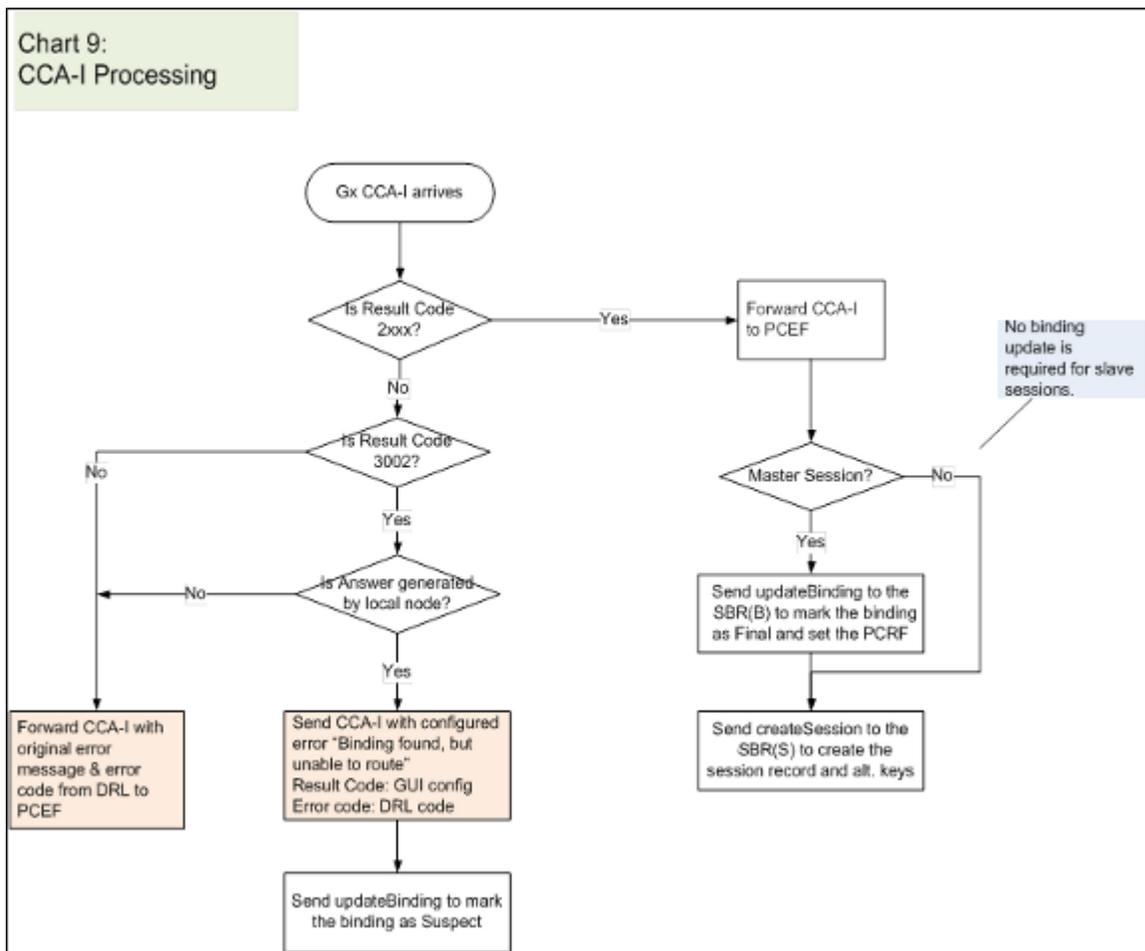


Figure 60: CCA-I Processing Error Resolution Flowchart

findSession Response Processing Error Resolution Flowchart

Figure 61: *findSession Response Processing Error Resolution Flowchart* shows an error resolution flowchart that illustrates where findSession Response Processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

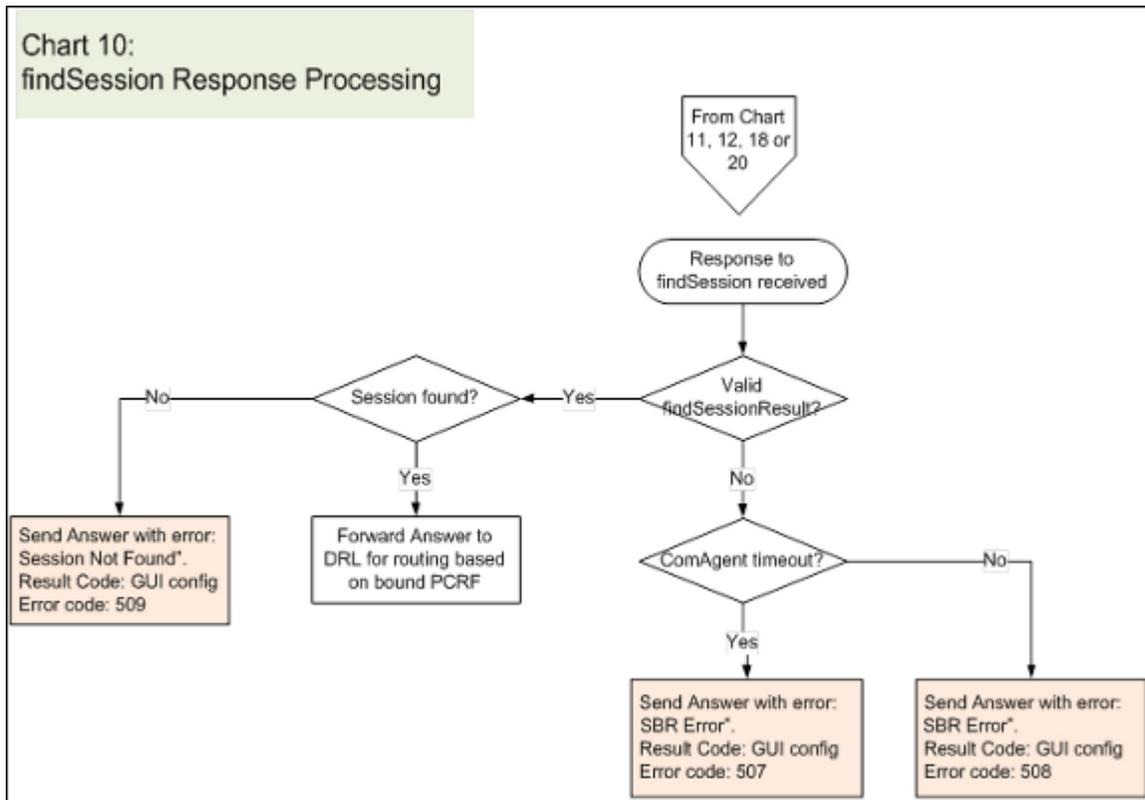


Figure 61: findSession Response Processing Error Resolution Flowchart

CCR-U Processing Error Resolution Flowchart

Figure 62: *CCR-U Processing Error Resolution Flowchart* shows an error resolution flowchart that illustrates where CCR-U Processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

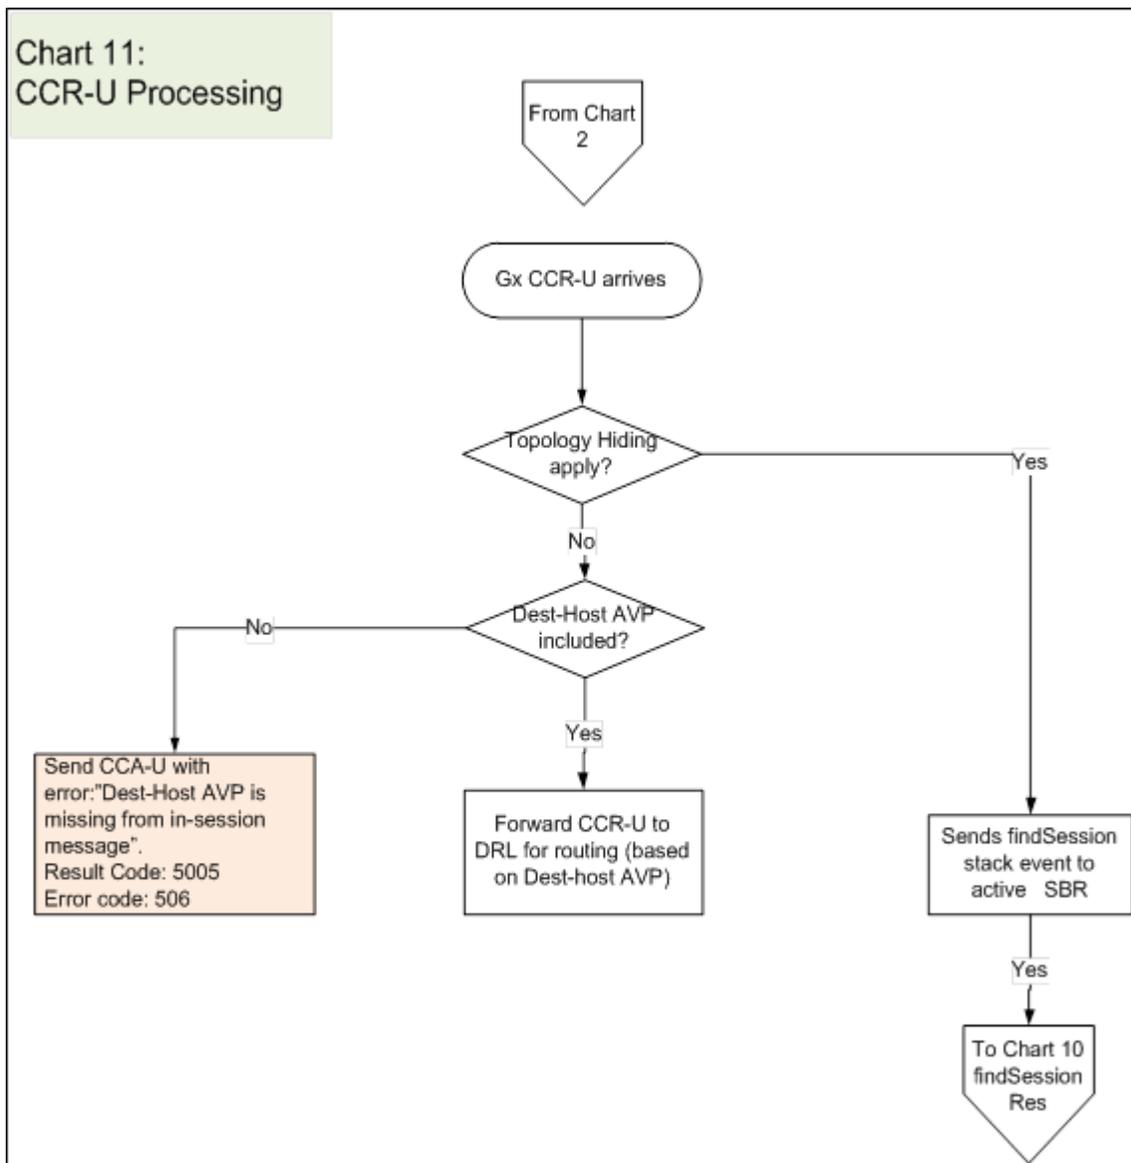


Figure 62: CCR-U Processing Error Resolution Flowchart

CCR-T Processing Error Resolution Flowchart

Figure 63: CCR-T Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where CCR-T Processing errors can occur.

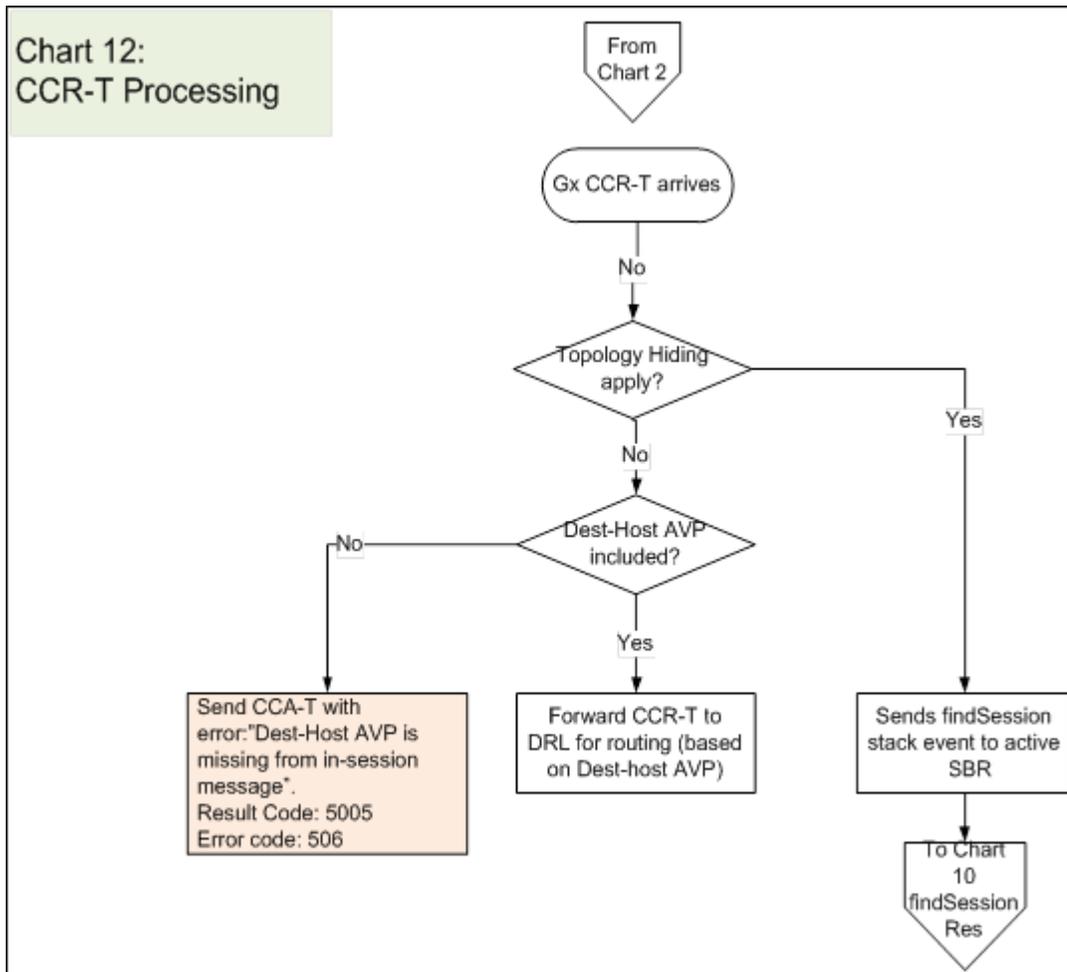


Figure 63: CCR-T Processing Error Resolution Flowchart

CCA-U/T Processing Error Resolution Flowchart

Figure 64: CCA-U/T Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where CCA-U/T Processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

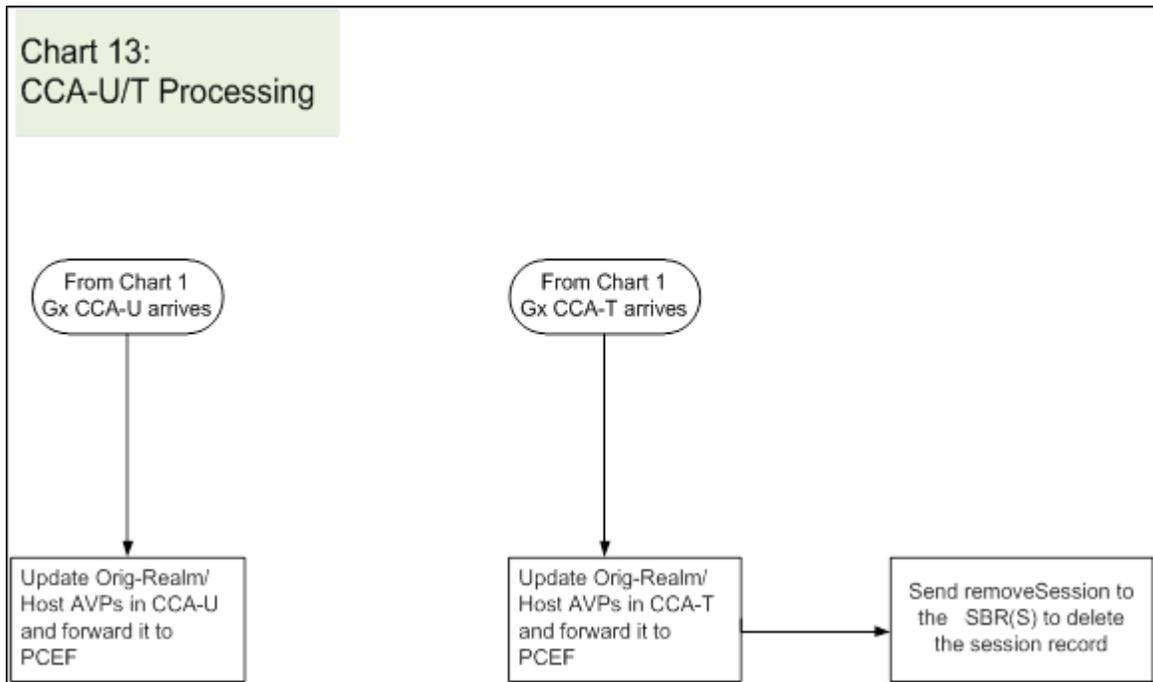


Figure 64: CCA-U/T Processing Error Resolution Flowchart

RAR Processing Error Resolution Flowchart

Figure 65: RAR Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where RAR Processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

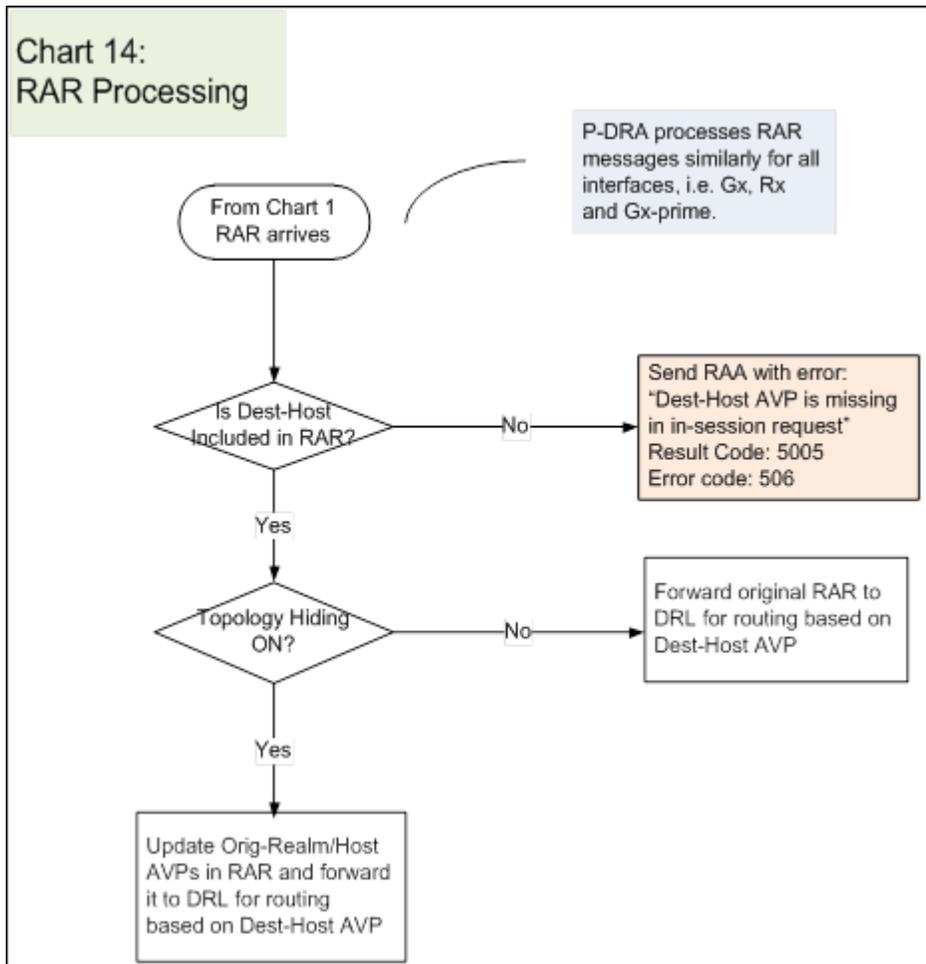


Figure 65: RAR Processing Error Resolution Flowchart

RAA Processing Error Resolution Flowchart

Figure 66: RAA Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where RAA Processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

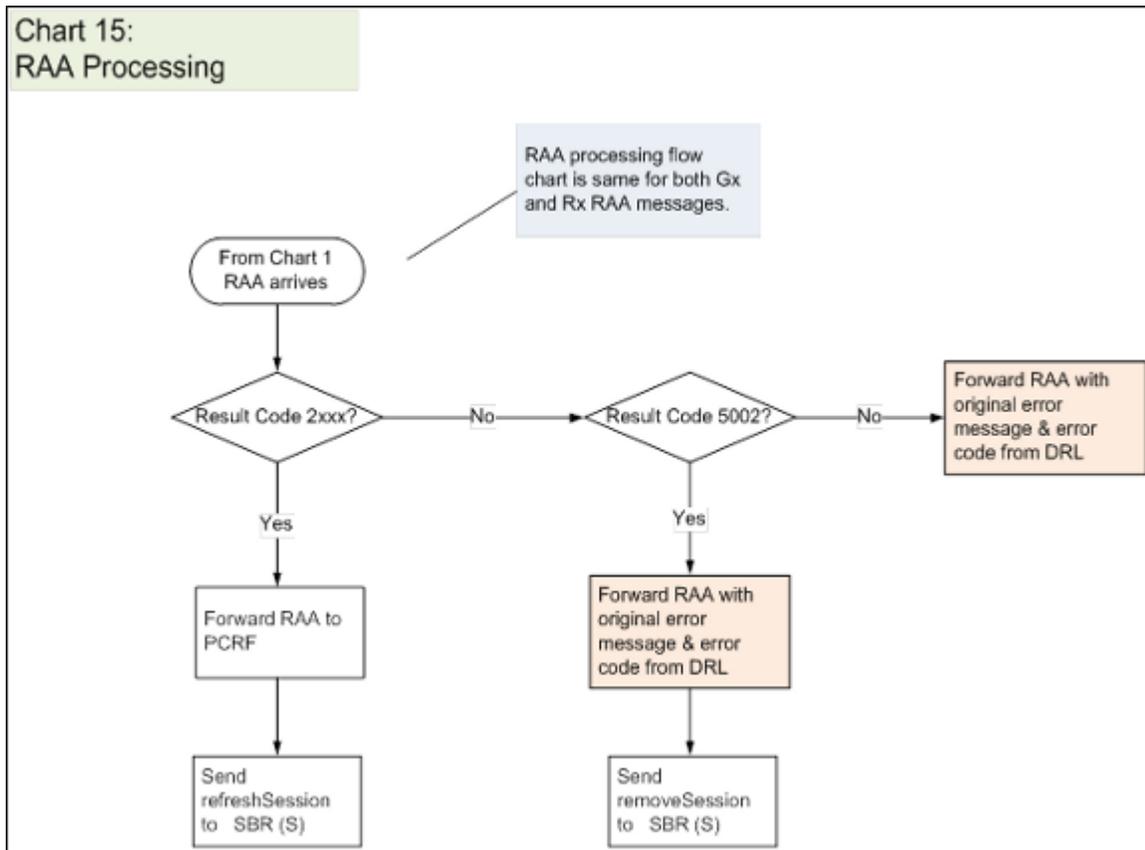


Figure 66: RAA Processing Error Resolution Flowchart

AAR Processing Error Resolution Flowchart

Figure 67: AAR Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where AAR Processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

Chart 16: AAR Processing

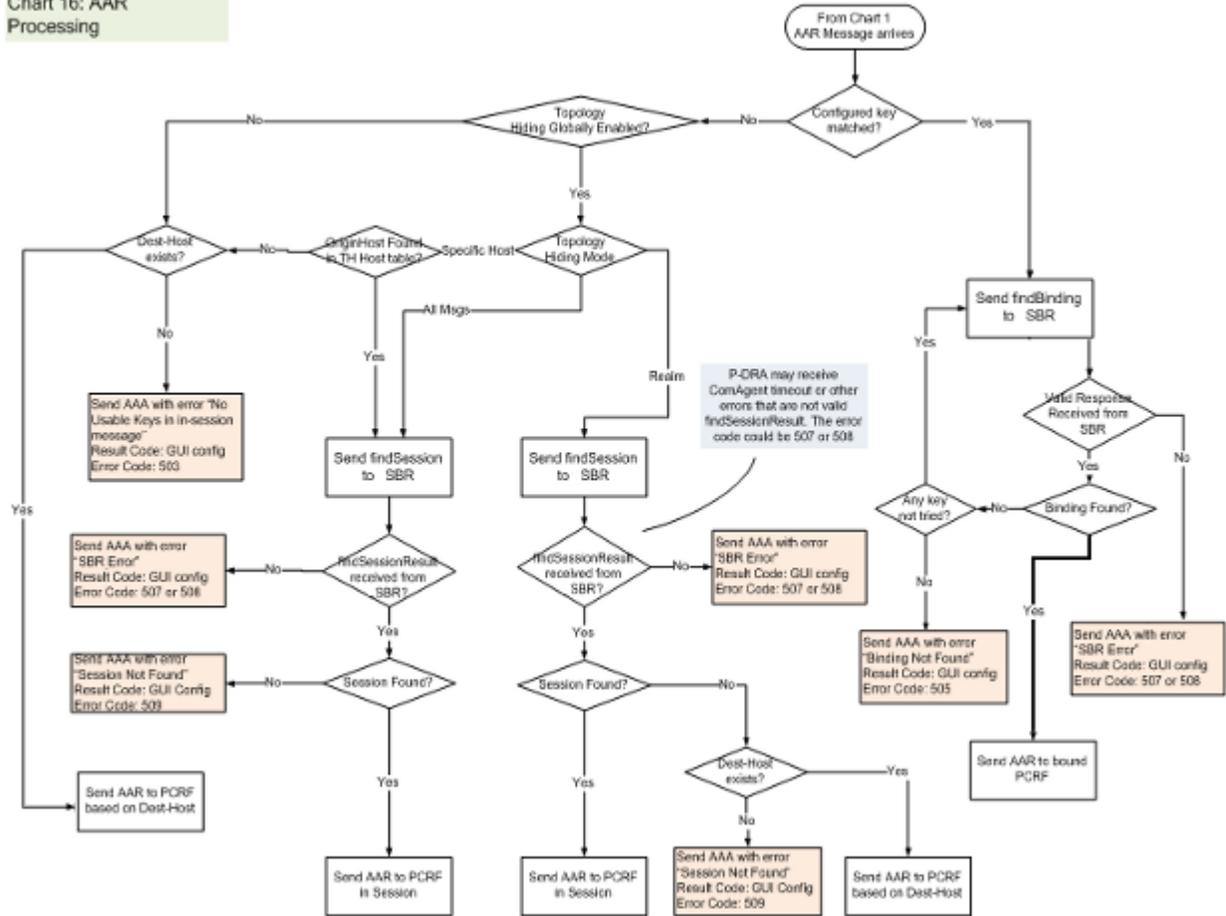


Figure 67: AAR Processing Error Resolution Flowchart

AAA Processing Error Resolution Flowchart

Figure 68: AAA Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where AAA Processing errors can occur.

Note: See Policy DRA Error Resolution Flowchart Summary for the entire business logic flowchart summary.

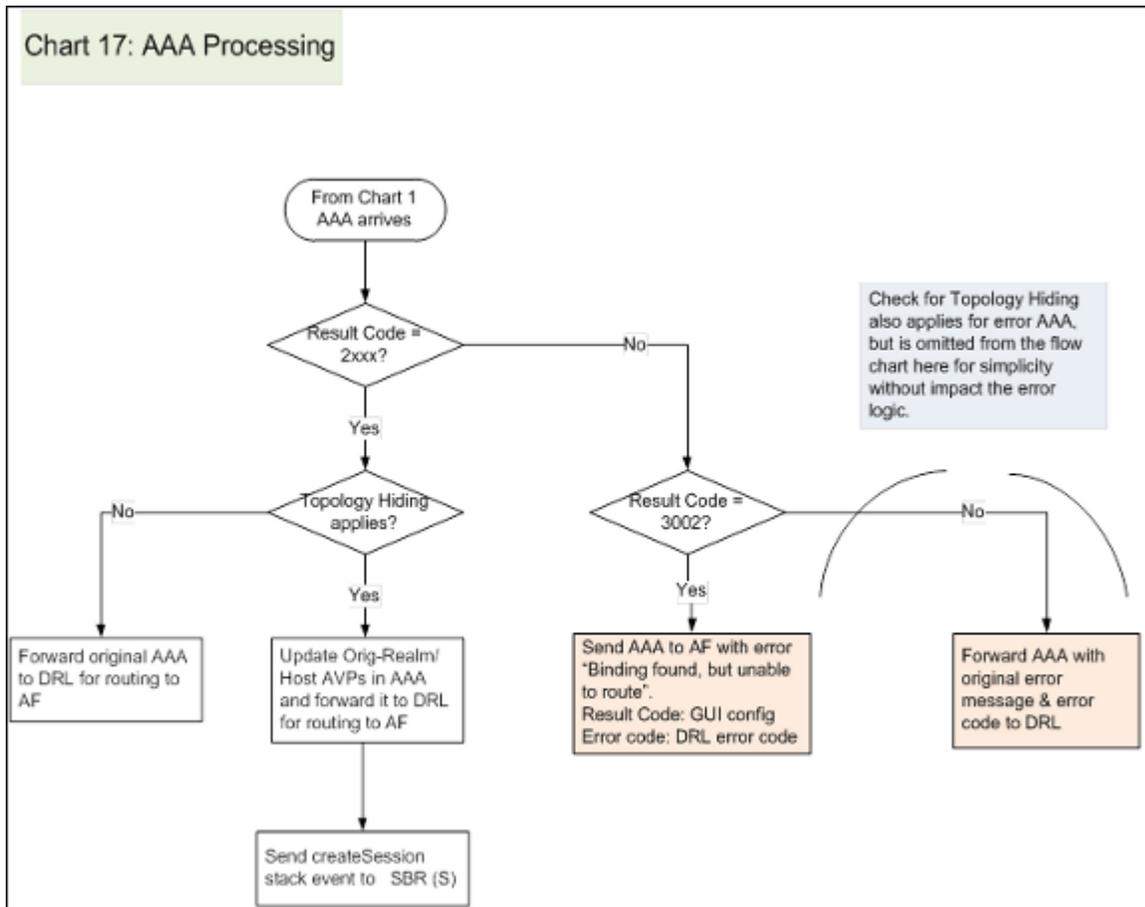


Figure 68: AAA Processing Error Resolution Flowchart

STR Processing Error Resolution Flowchart

Figure 69: STR Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where STR Processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

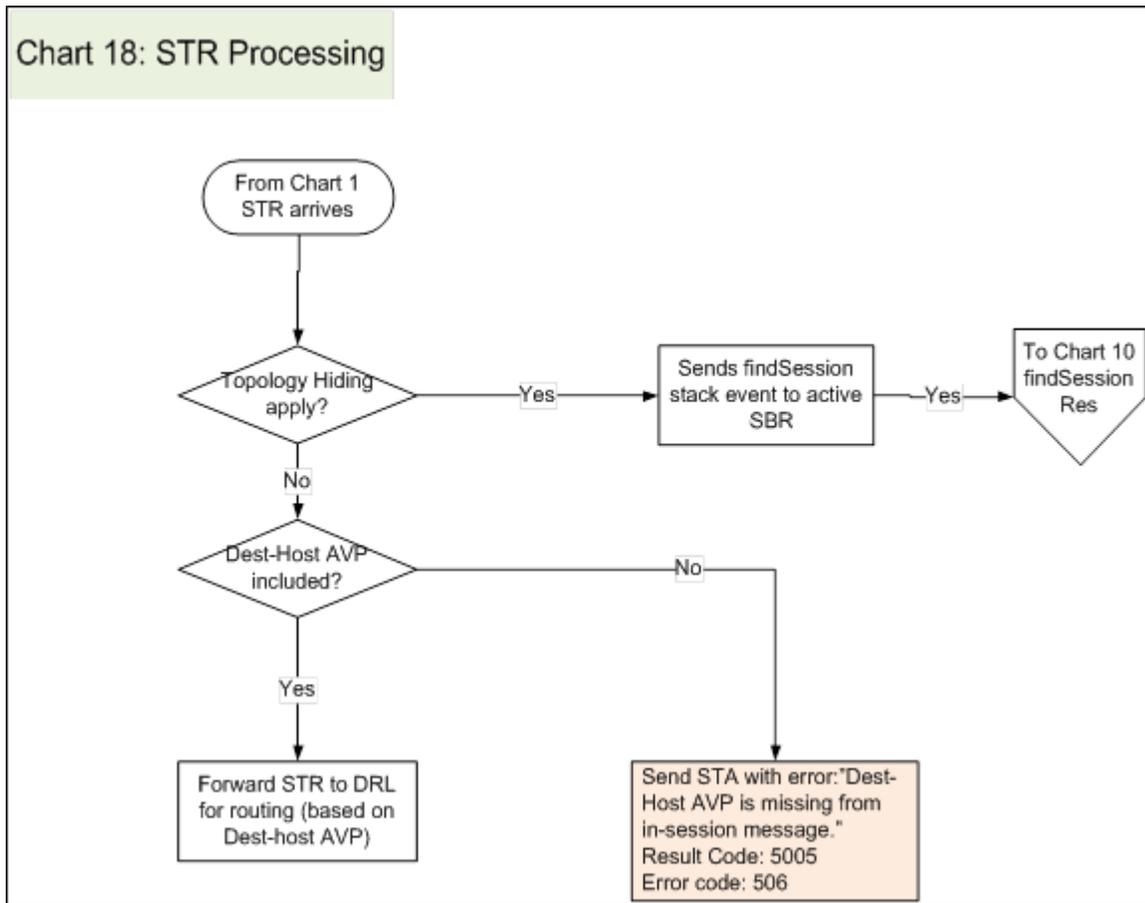


Figure 69: STR Processing Error Resolution Flowchart

STA Processing Error Resolution Flowchart

Figure 70: *STA Processing Error Resolution Flowchart* shows an error resolution flowchart that illustrates where STA Processing errors can occur.

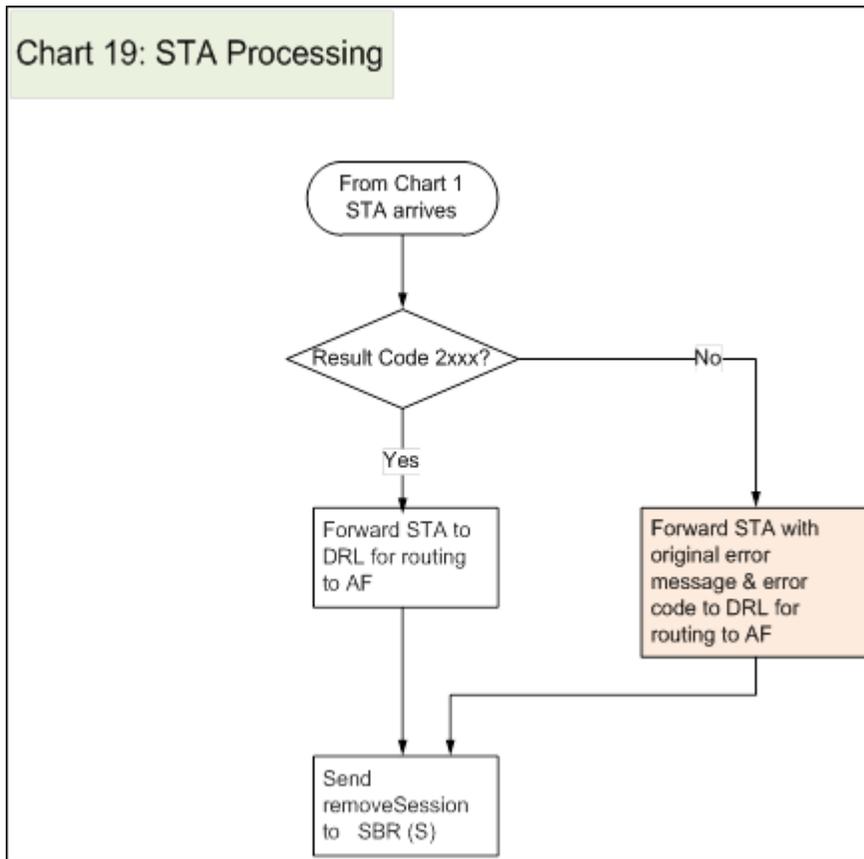


Figure 70: STA Processing Error Resolution Flowchart

ASR/ASA Processing Error Resolution Flowchart

Figure 71: ASR/ASA Processing Error Resolution Flowchart shows an resolution flowchart map that illustrates where ASR/ASA Processing errors can occur.

Note: See *Policy DRA Error Resolution Flowchart Summary* for the entire business logic flowchart summary.

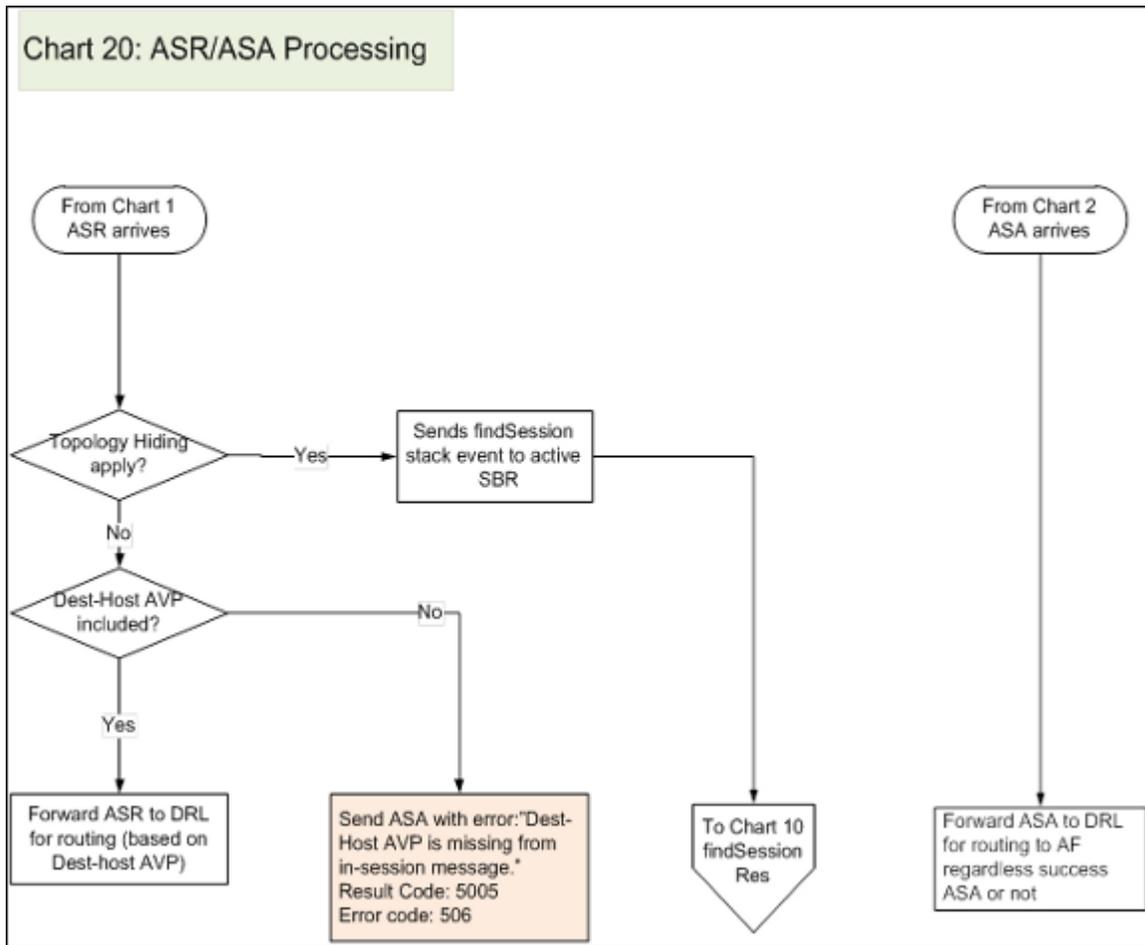


Figure 71: ASR/ASA Processing Error Resolution Flowchart

Online Charging DRA Error Resolution Flowchart Summary

Figure 72: Error Resolution Flowchart Summary shows a summary of the Online Charging DRA error resolution flowcharts in this section and their relationships to each other. This relationship tree should help you understand the PCA business logic as it relates to OC-DRA error resolution tasks.

Each reference in the flowchart summary points to a corresponding error resolution flowchart in this appendix. Use the following table to reconcile those references:

Chart number	Flowchart Name	Link
Chart 1	Message Processing	Message Processing
Chart 2	Diameter Validation Processing	Diameter Validation Processing
Chart 3	CCR Processing	CCR Processing
Chart 4	CCR-I Processing	CCR-I Processing

Chart number	Flowchart Name	Link
Chart 5	CCR-U Processing	<i>CCR-U Processing</i>
Chart 6	CCR-T Processing	<i>CCR-T Processing</i>
Chart 7	CCR-E Processing	<i>CCR-E Processing</i>
Chart 8	RAR Processing	<i>RAR Processing</i>
Chart 9	SBR SE Received Processing	<i>SBR SE Received Processing</i>
Chart 10	CreateOcSessionResult SE Processing	<i>CreateOcSessionResult SE Processing</i>
Chart 11	FindAndRefreshOcSession SE Processing	<i>FindAndRefreshOcSession SE Processing</i>
Chart 12	FindAndRemoveOcSessionResult SE Processing	<i>FindAndRemoveOcSessionResult SE Processing</i>
Chart 13	CCA Processing	<i>CCA Processing</i>
Chart 14	CCA-I Processing	<i>CCA-I Processing</i>
Chart 15	CCA-U Processing	<i>CCA-U Processing</i>
Chart 16	CCA-T Processing	<i>CCA-T Processing</i>
Chart 17	CCA-U Processing	<i>CCA-U Processing</i>
Chart 18	RAA Processing	<i>RAA Processing</i>



Figure 72: Error Resolution Flowchart Summary

Message Processing

Figure 73: *Message Processing* shows an error resolution flowchart that illustrates where Message Processing error.

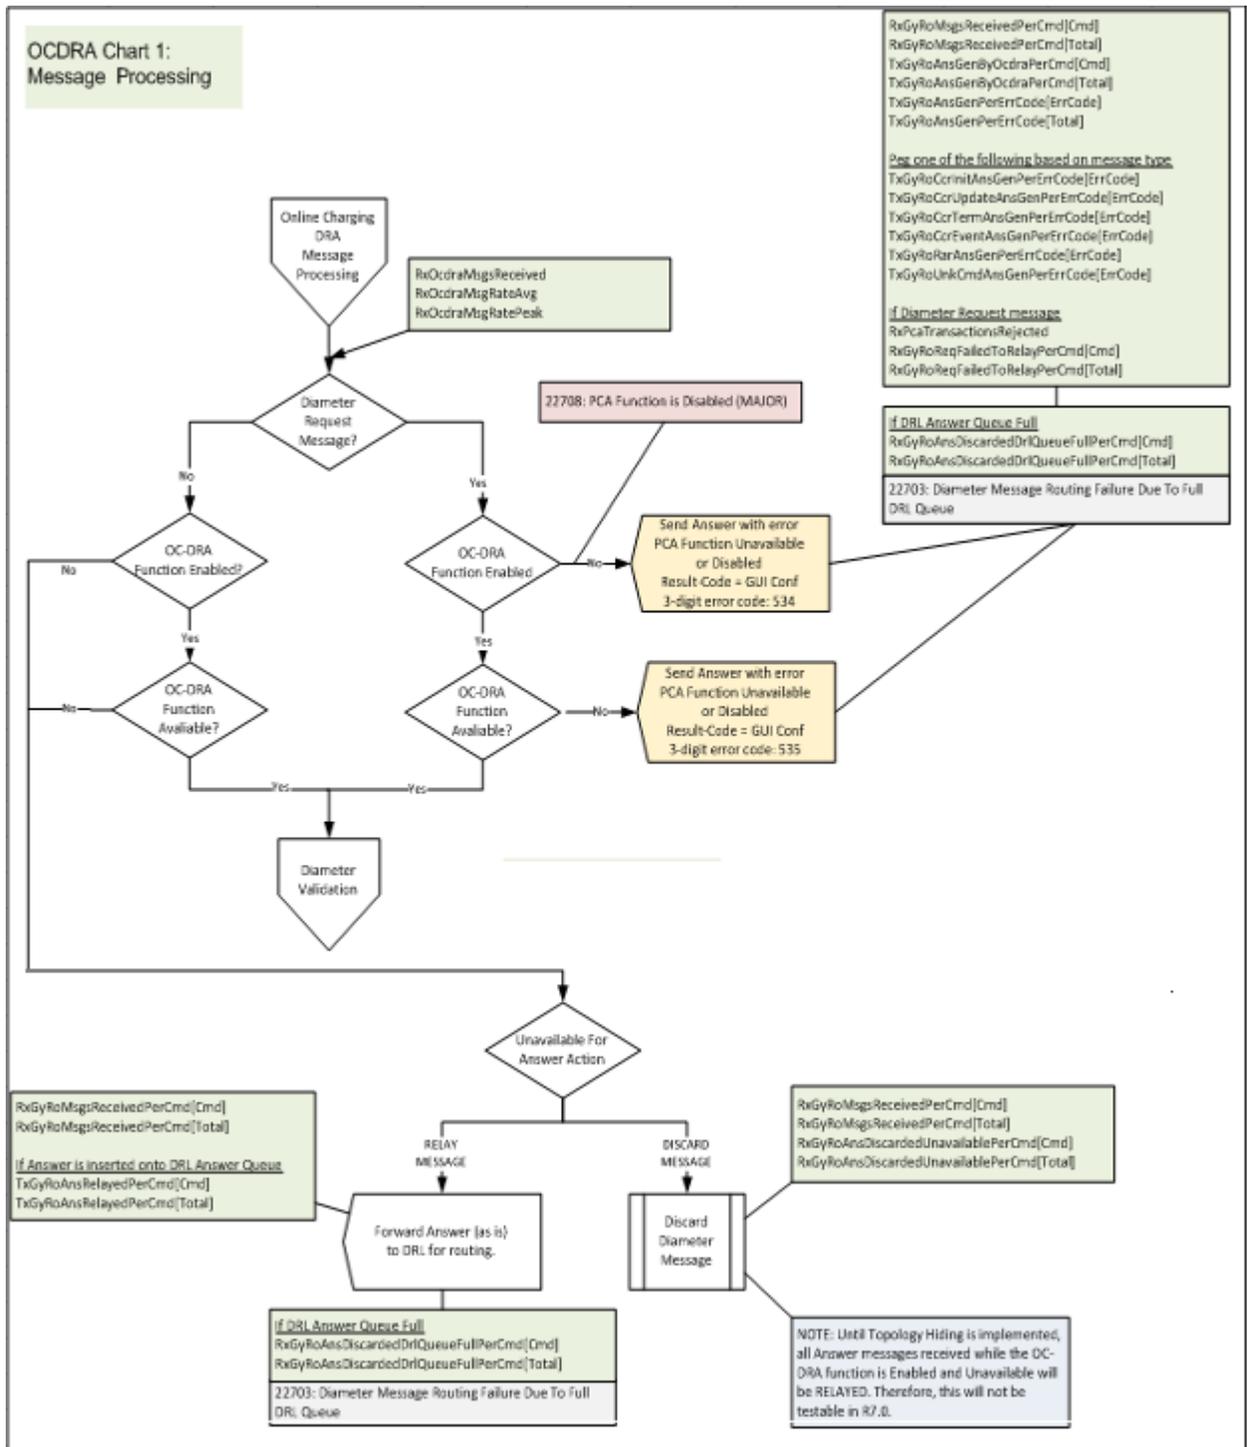


Figure 73: Message Processing

Diameter Validation Processing

Figure 74: Diameter Validation Processing shows an error resolution flowchart that illustrates where Diameter Validation Processing error.

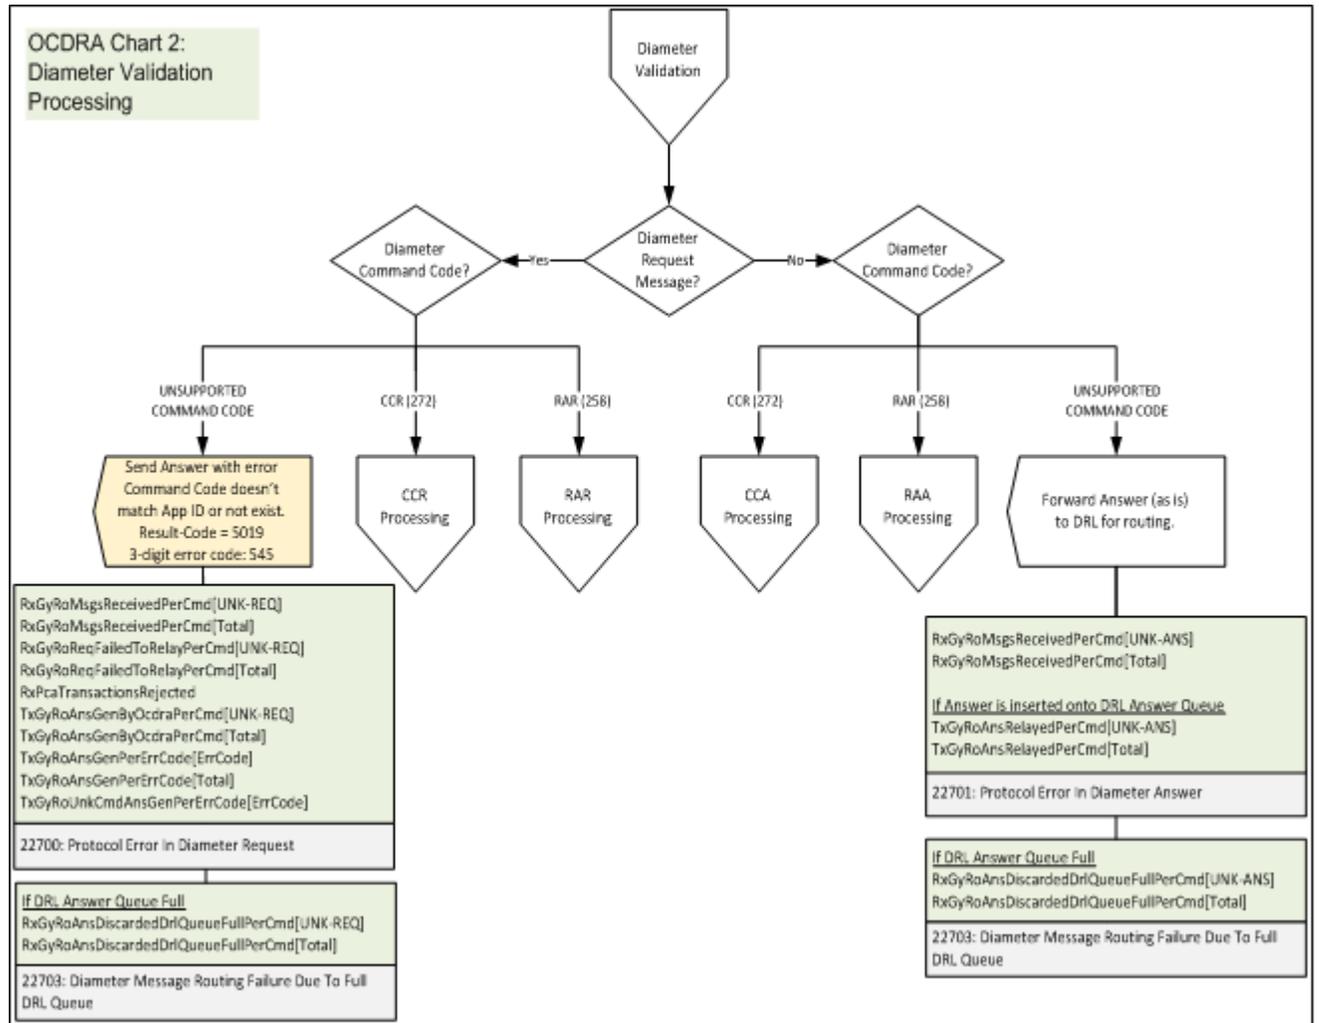


Figure 74: Diameter Validation Processing

CCR Processing

Figure 75: CCR Processing shows an error resolution flowchart that illustrates where CCR Processing error.

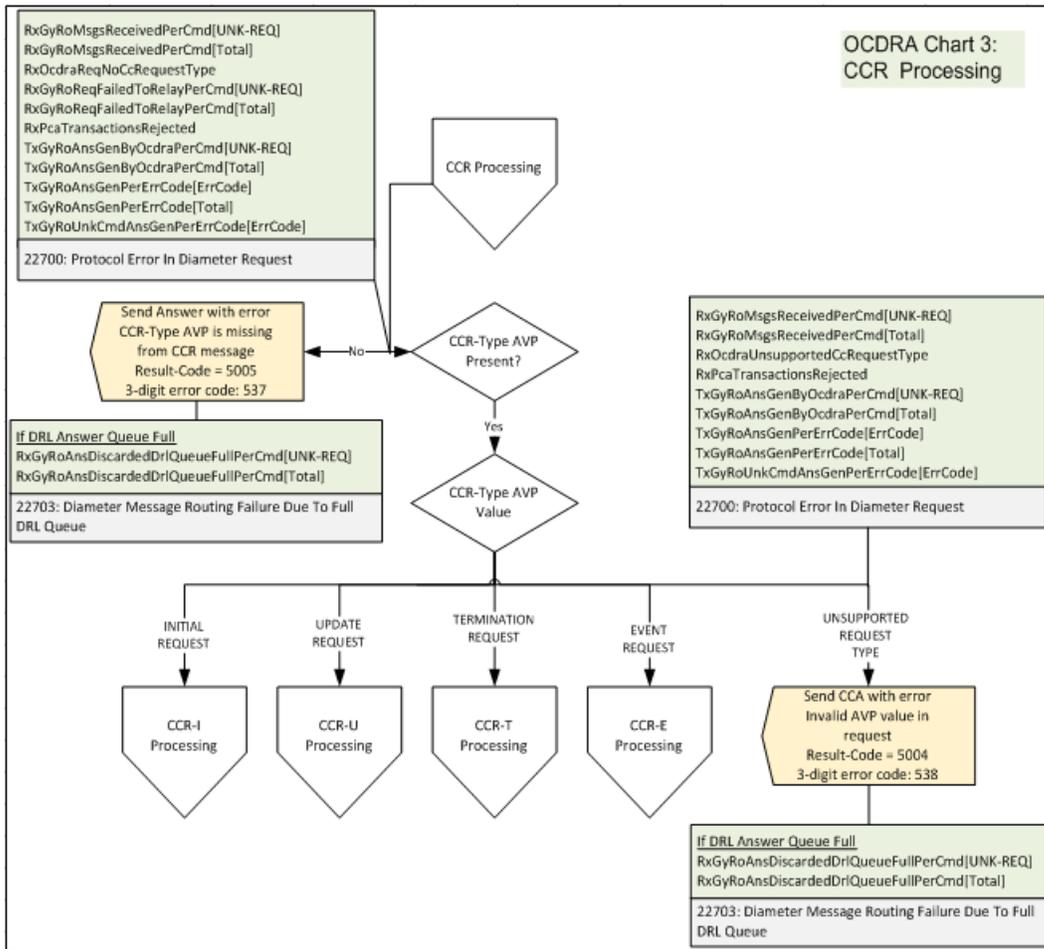


Figure 75: CCR Processing

CCR-I Processing

Figure 76: CCR-I Processing shows an error resolution flowchart that illustrates where CCR-I Processing error.

CCR-U Processing

Figure 77: CCR-U Processing shows an error resolution flowchart that illustrates where CCR-U Processing error.

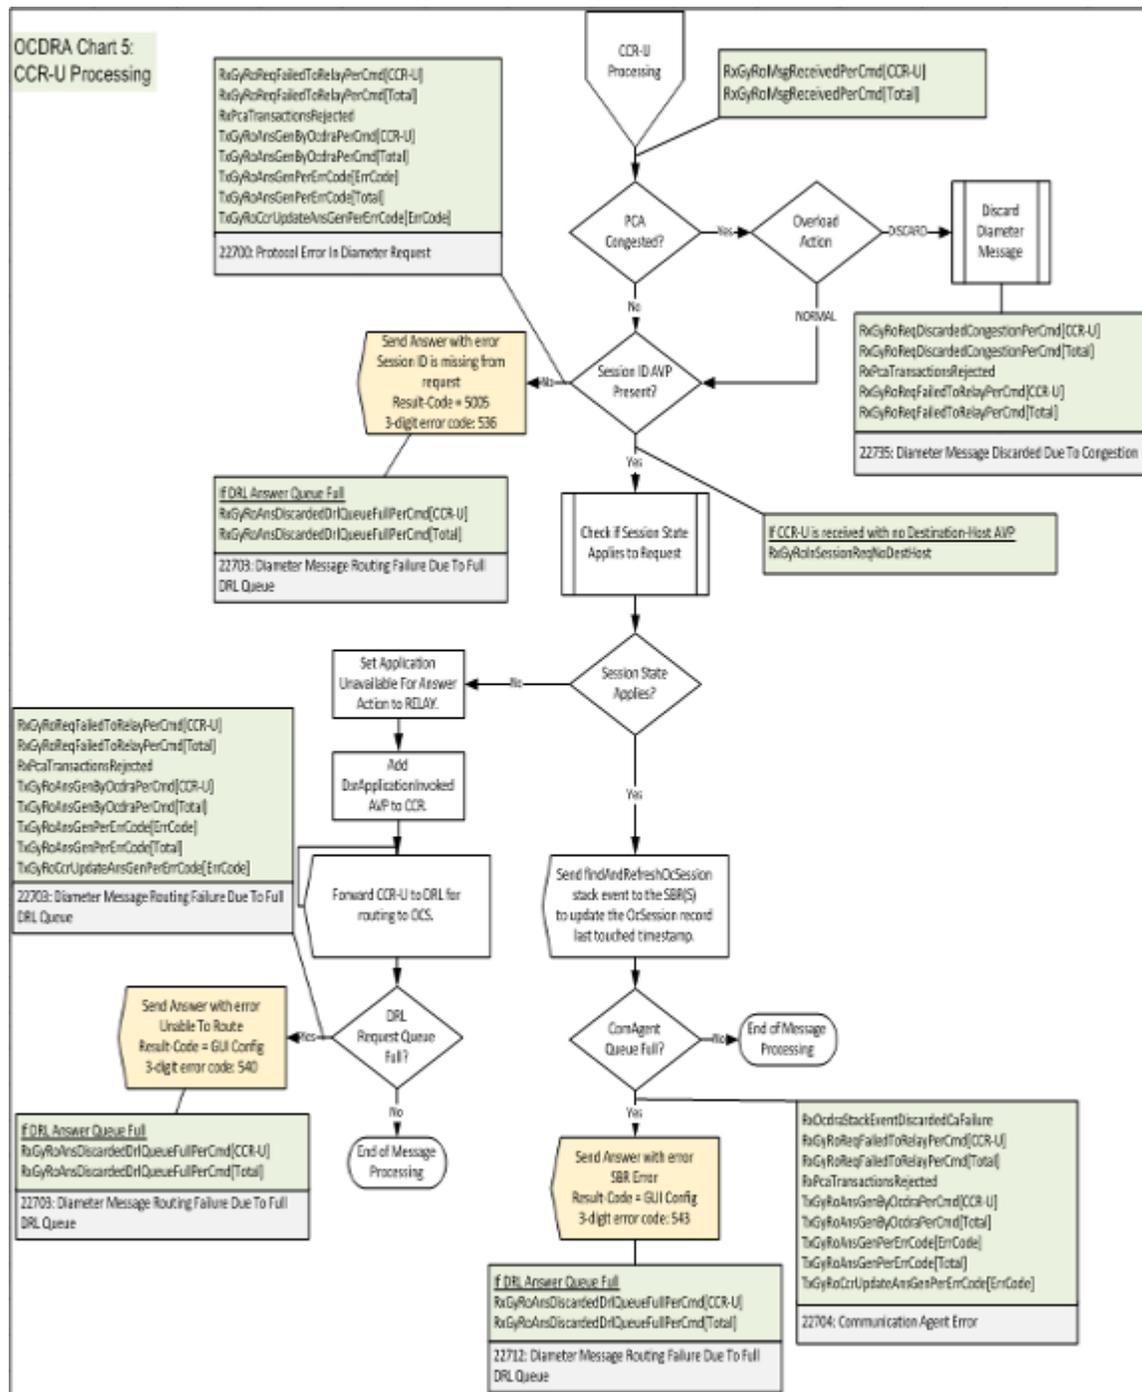


Figure 77: CCR-U Processing

CCR-T Processing

Figure 78: CCR-T Processing shows an error resolution flowchart that illustrates where CCR-T Processing error.

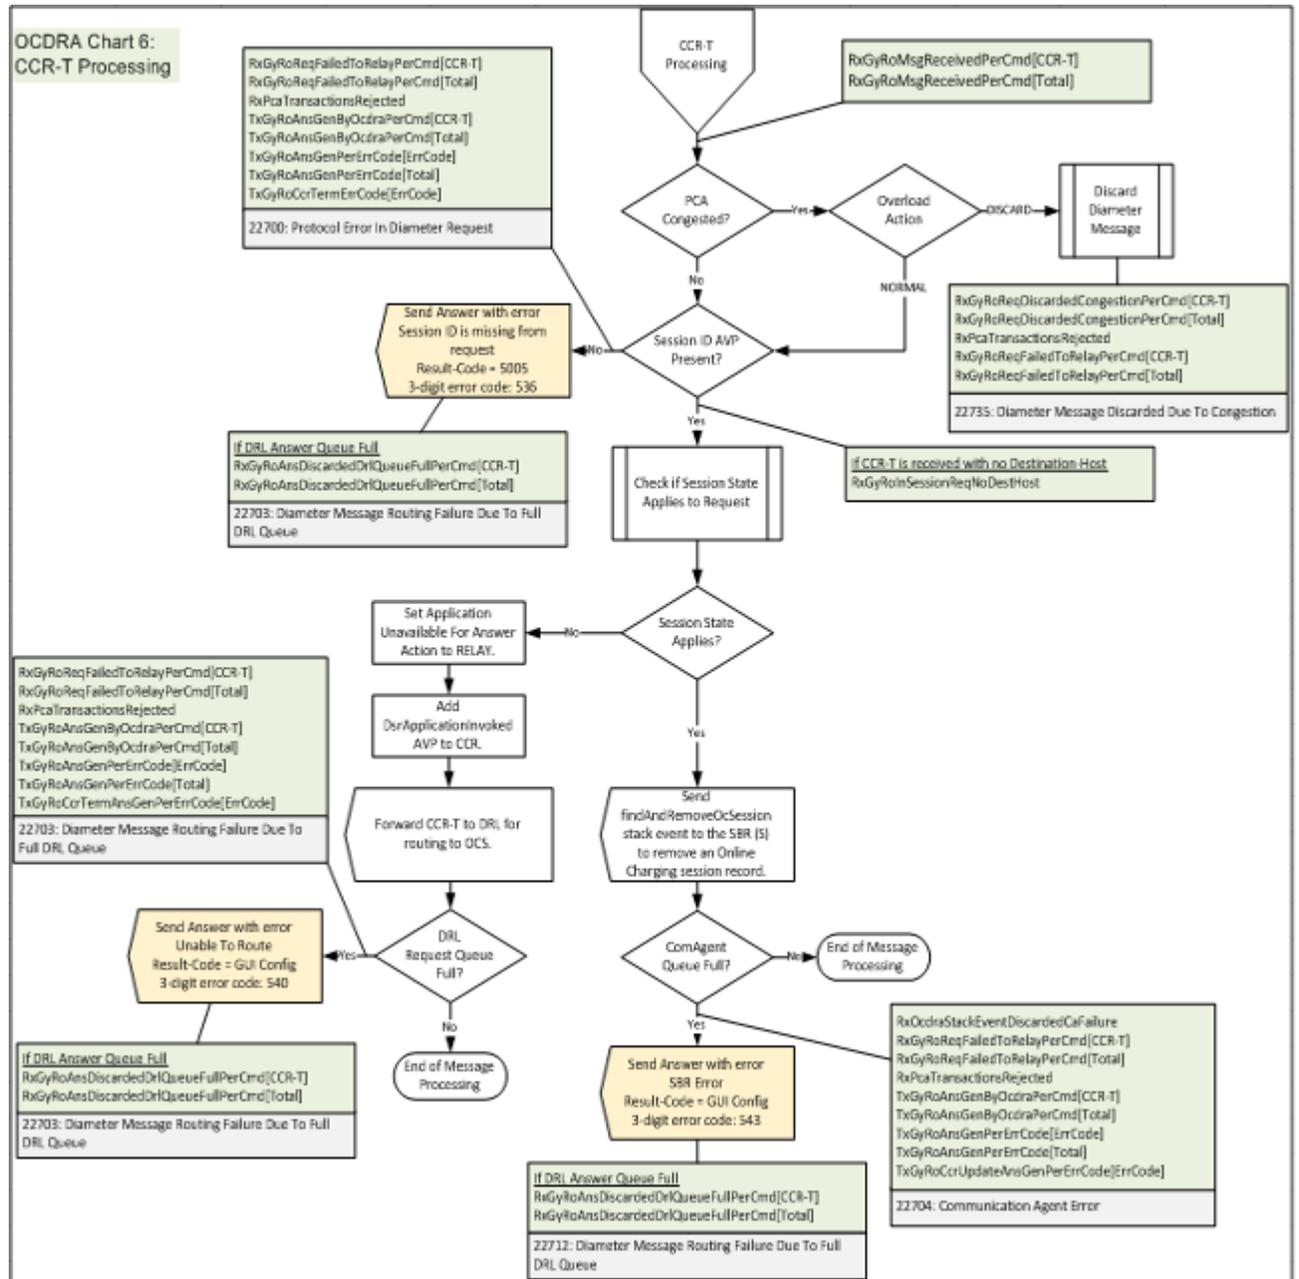


Figure 78: CCR-T Processing

CCR-E Processing

Figure 79: CCR-E Processing shows an error resolution flowchart that illustrates where CCR-E Processing error.

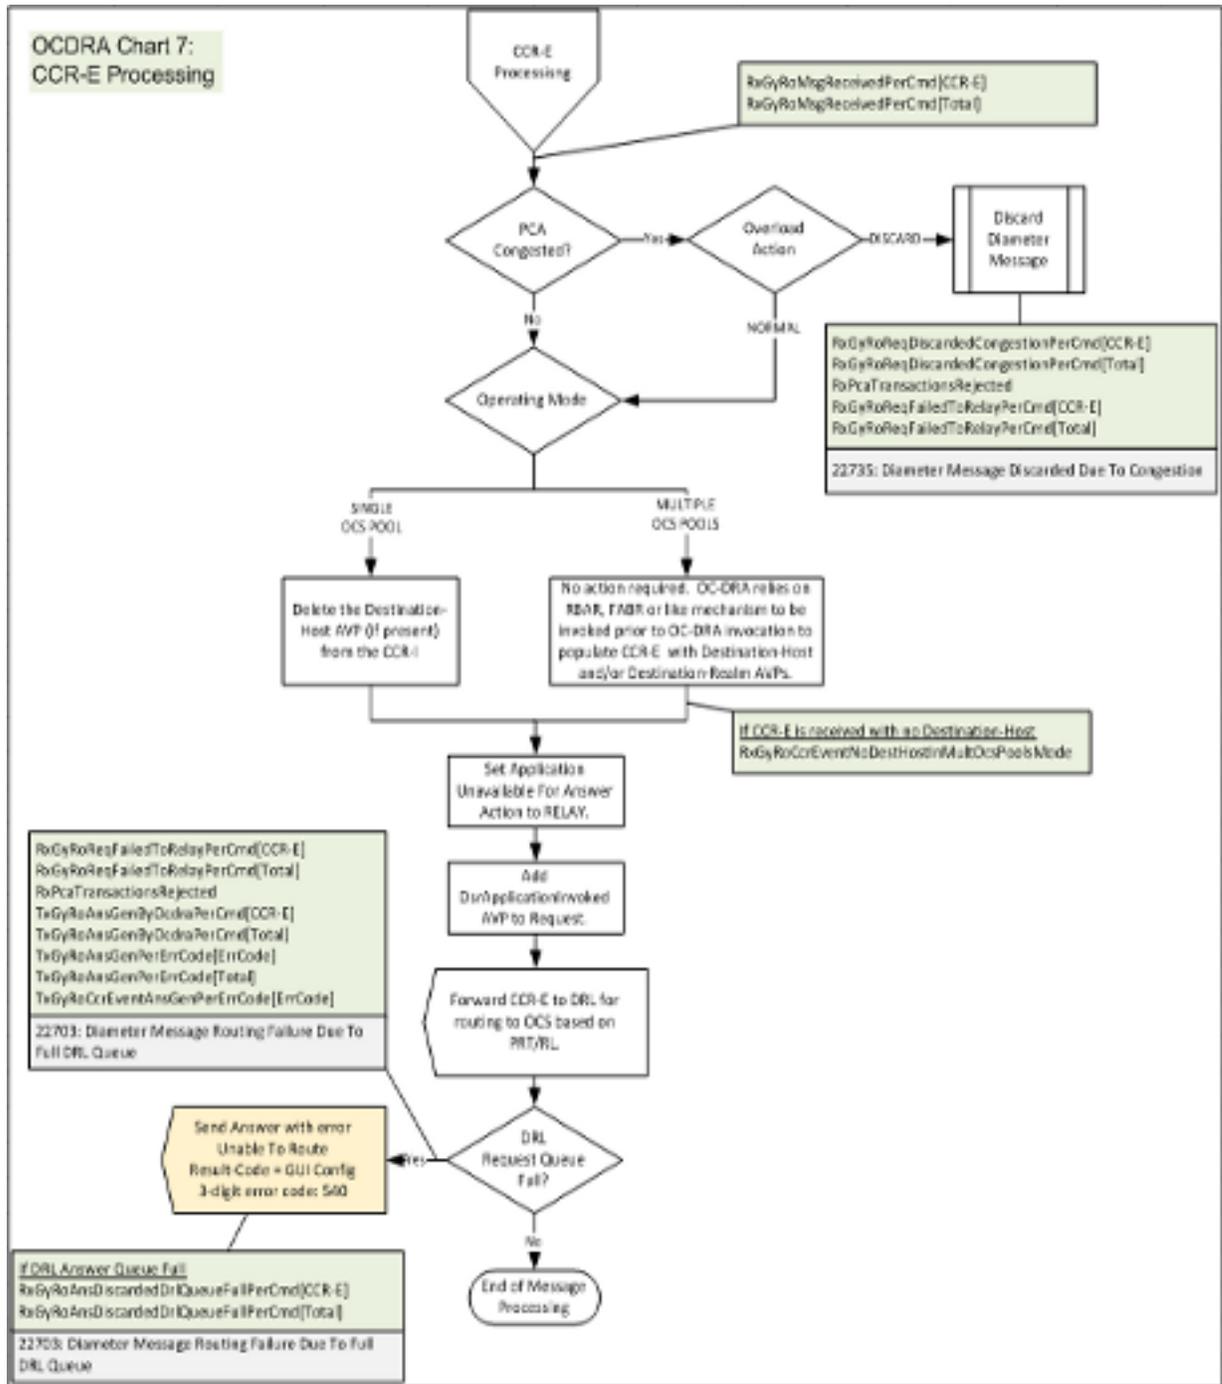


Figure 79: CCR-E Processing

RAR Processing

Figure 80: RAR Processing shows an error resolution flowchart that illustrates where RAR Processing error.

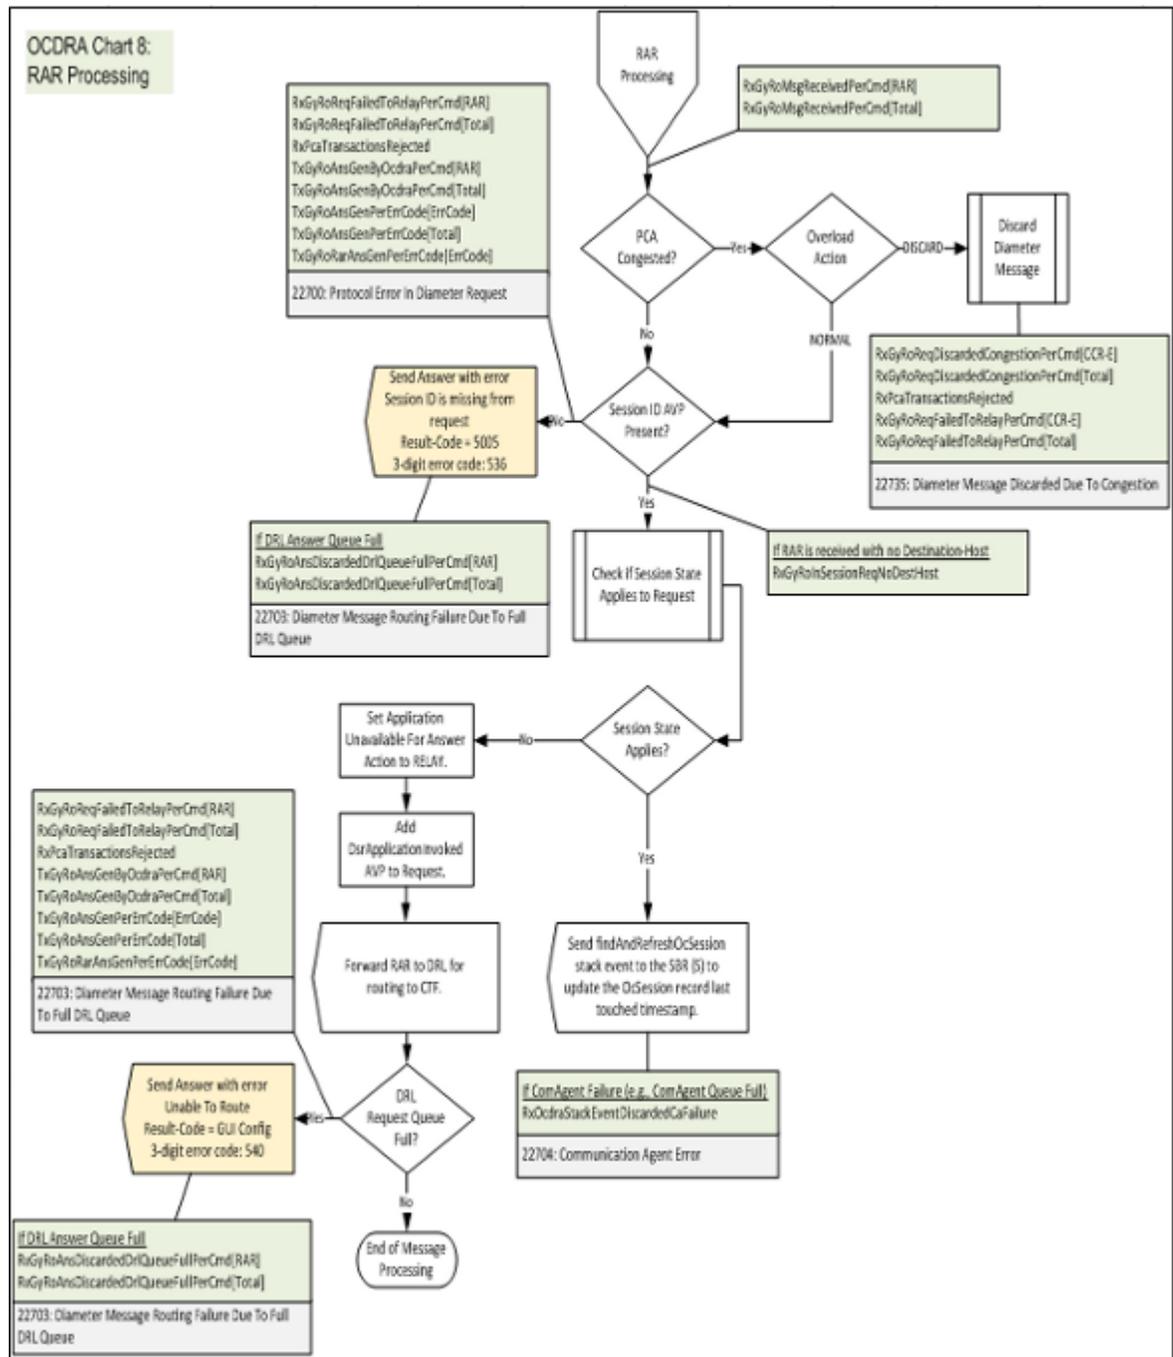


Figure 80: RAR Processing

SBR SE Received Processing

Figure 81: SBR SE Received Processing shows an error resolution flowchart that illustrates where SBR SE Received Processing error.



Figure 81: SBR SE Received Processing

CreateOcSessionResult SE Processing

Figure 82: *CreateOcSessionResult Processing* shows an error resolution flowchart that illustrates where CreateOcSessionResult SE Processing error.

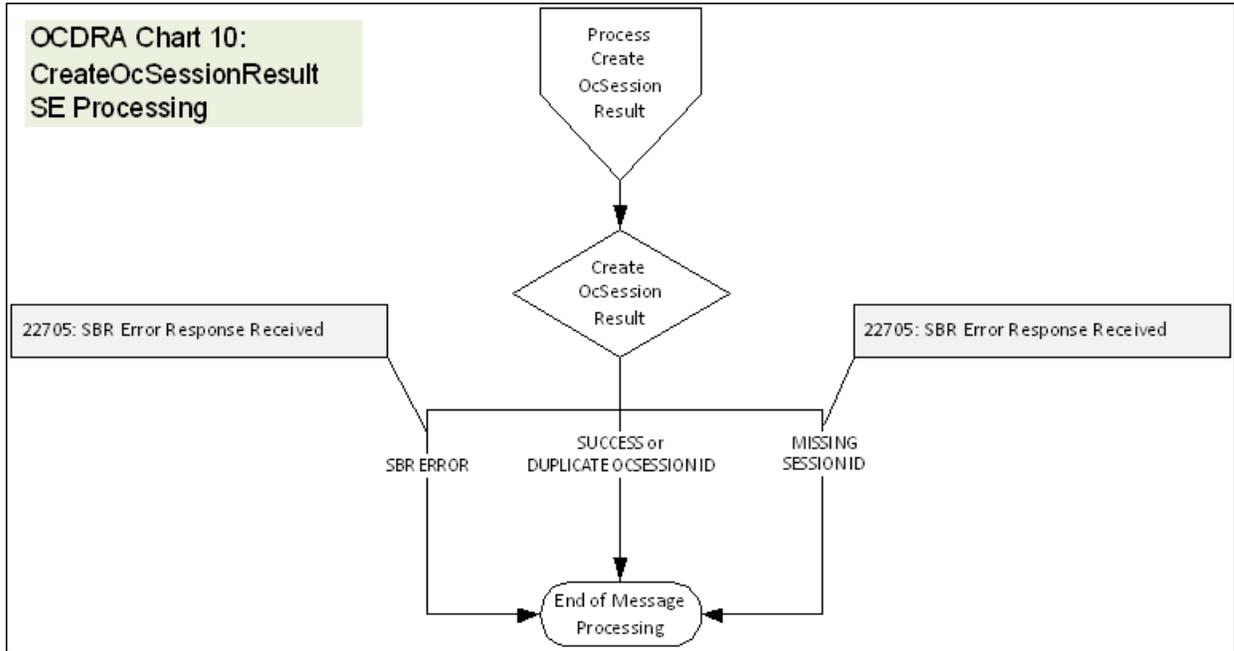


Figure 82: CreateOcSessionResult Processing

FindAndRefreshOcSession SE Processing

FindAndRefreshOcSession SE Processing shows an error resolution flowchart that illustrates where FindAndRefreshOcSession SE Processing error.

OCDRA Chart 11:
FindAndRefreshOcSession
SE Processing

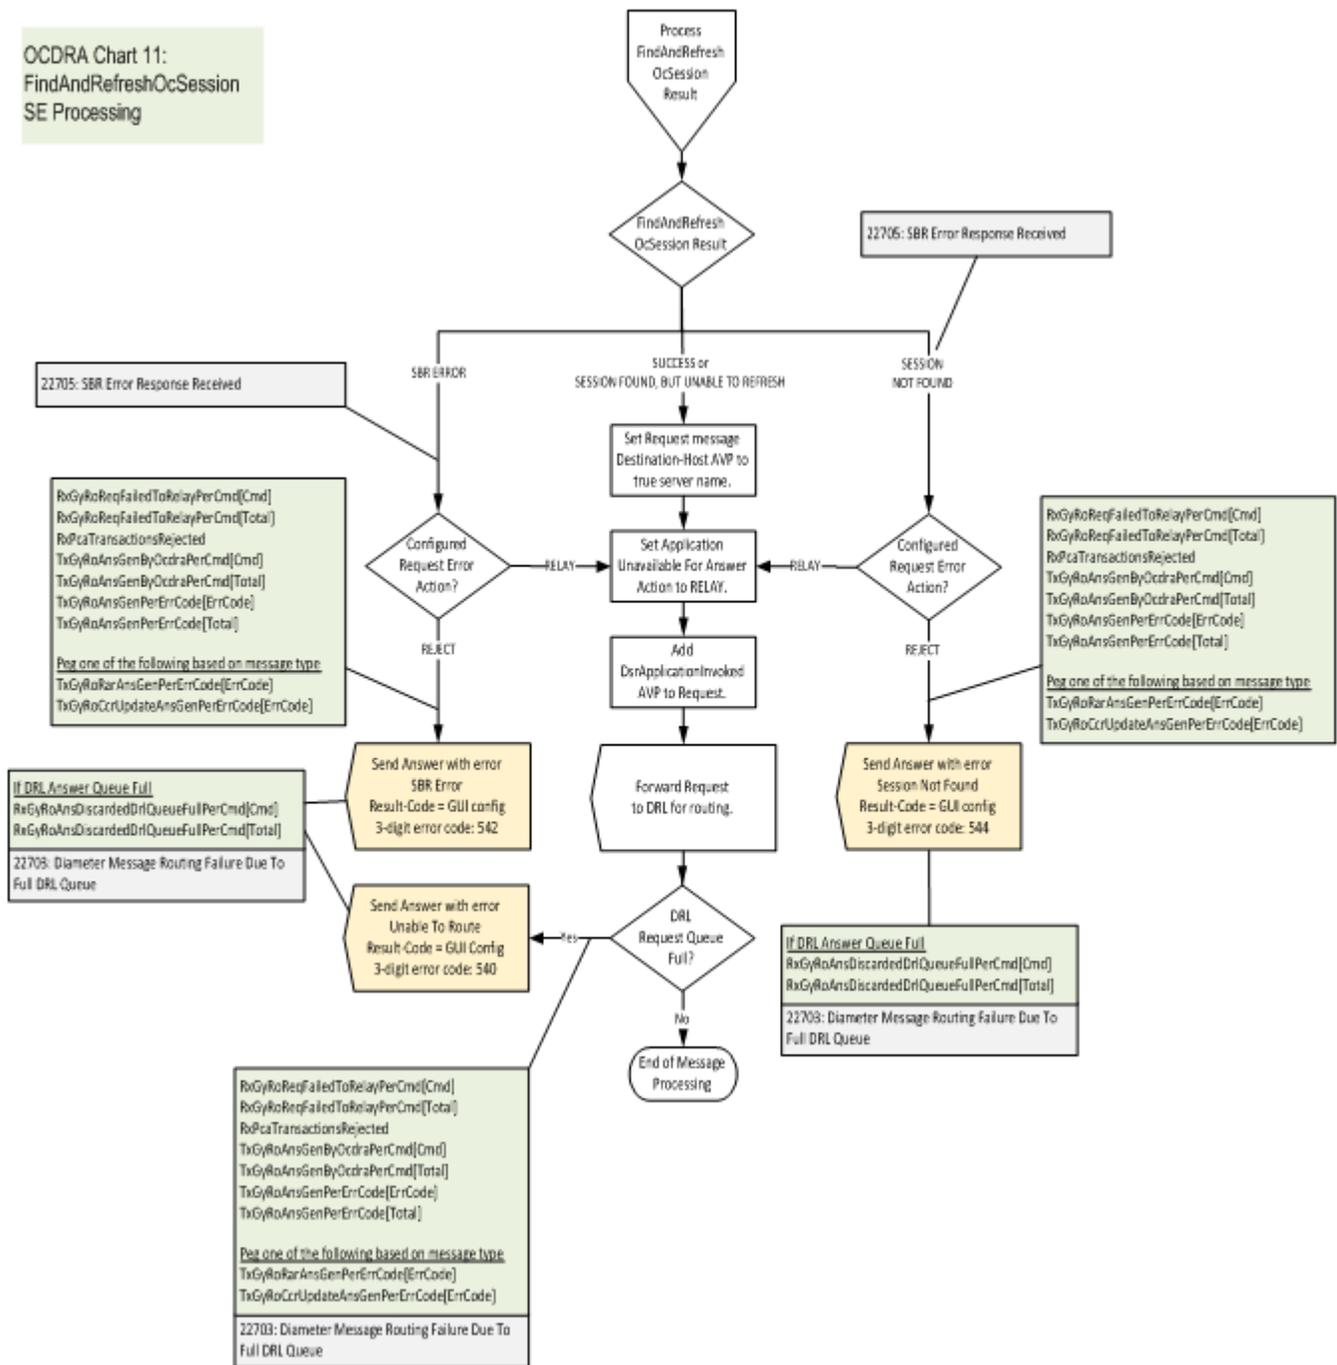


Figure 83: FindAndRefreshOcSession SE Processing

FindAndRemoveOcSessionResult SE Processing

Figure 84: FindAndRemoveOcSessionResult Processing shows an error resolution flowchart that illustrates where FindAndRemoveOcSessionResult SE Processing error.

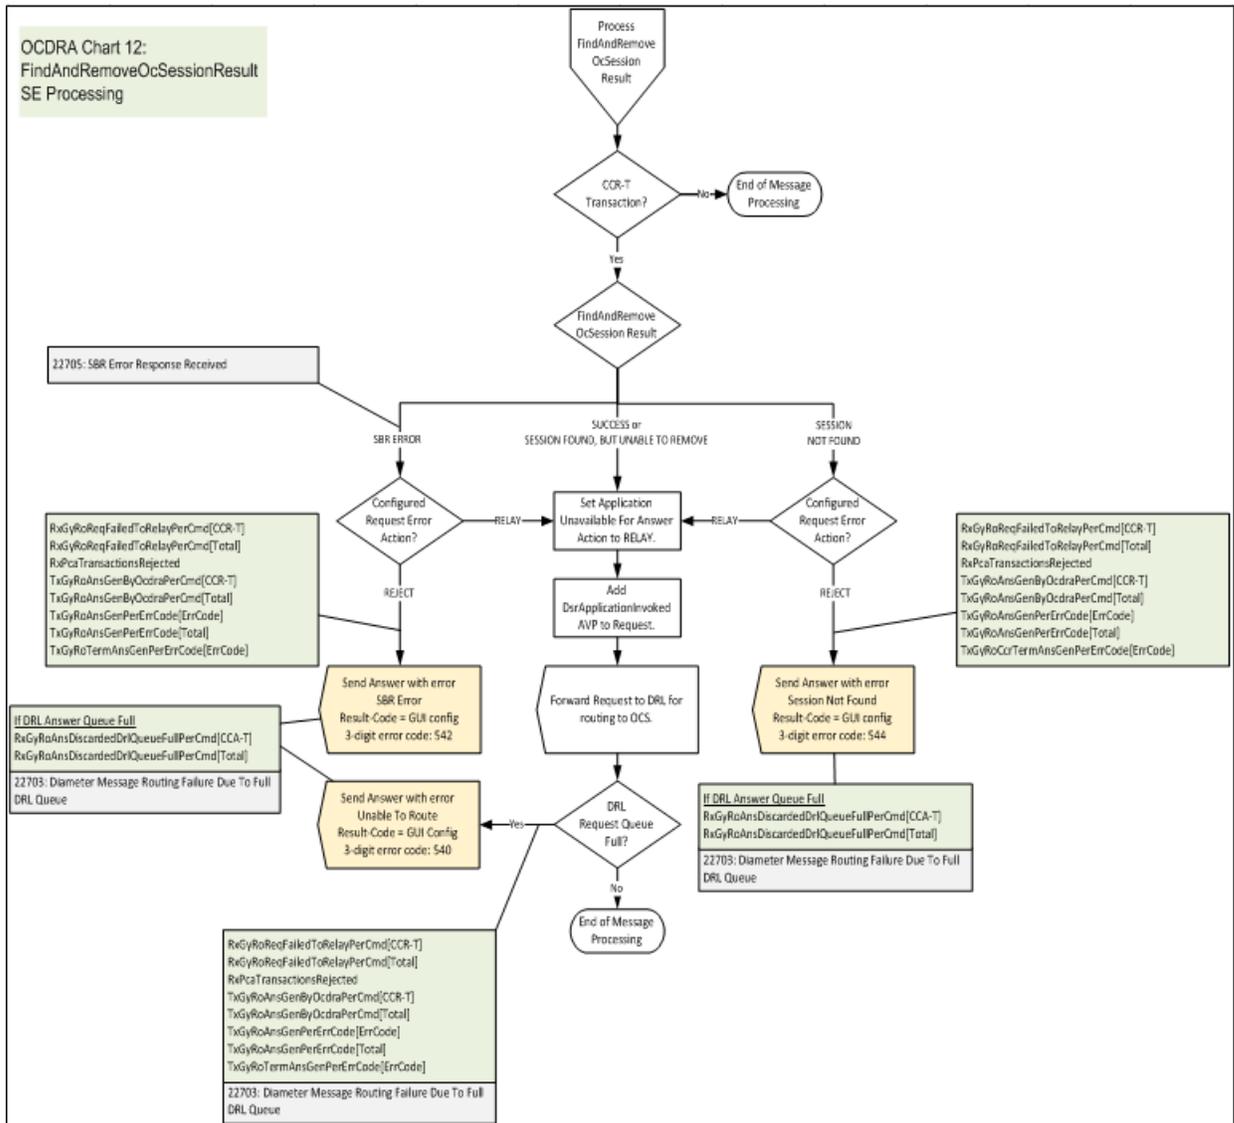


Figure 84: FindAndRemoveOcSessionResult Processing

CCA Processing

Figure 85: CCA Processing shows an error resolution flowchart that illustrates where CCA Processing error.

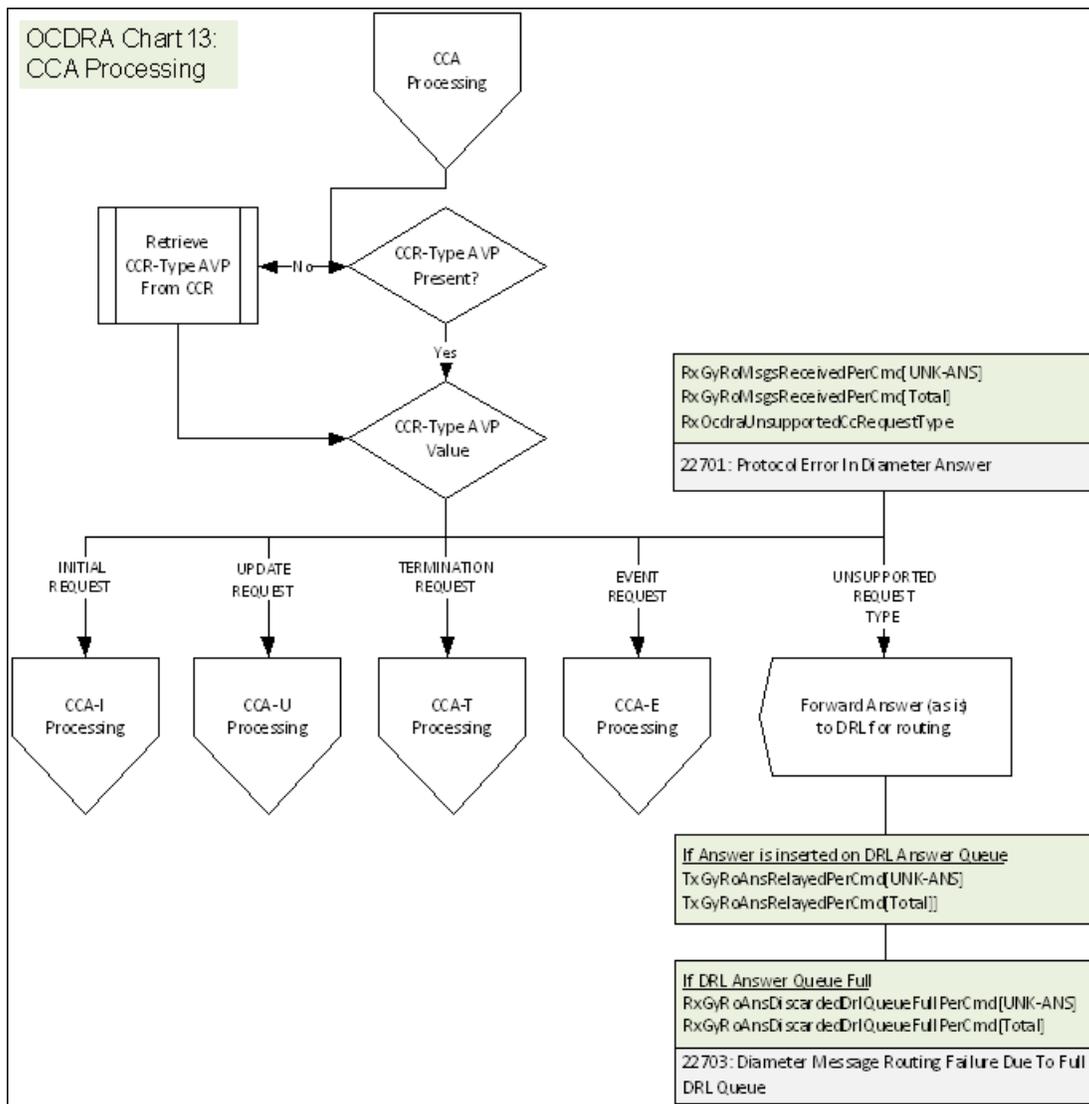


Figure 85: CCA Processing

CCA-I Processing

Figure 86: CCA-I Processing shows an error resolution flowchart that illustrates where CCA-I Processing error.

OCDRA Chart 14:
CCA-I Processing

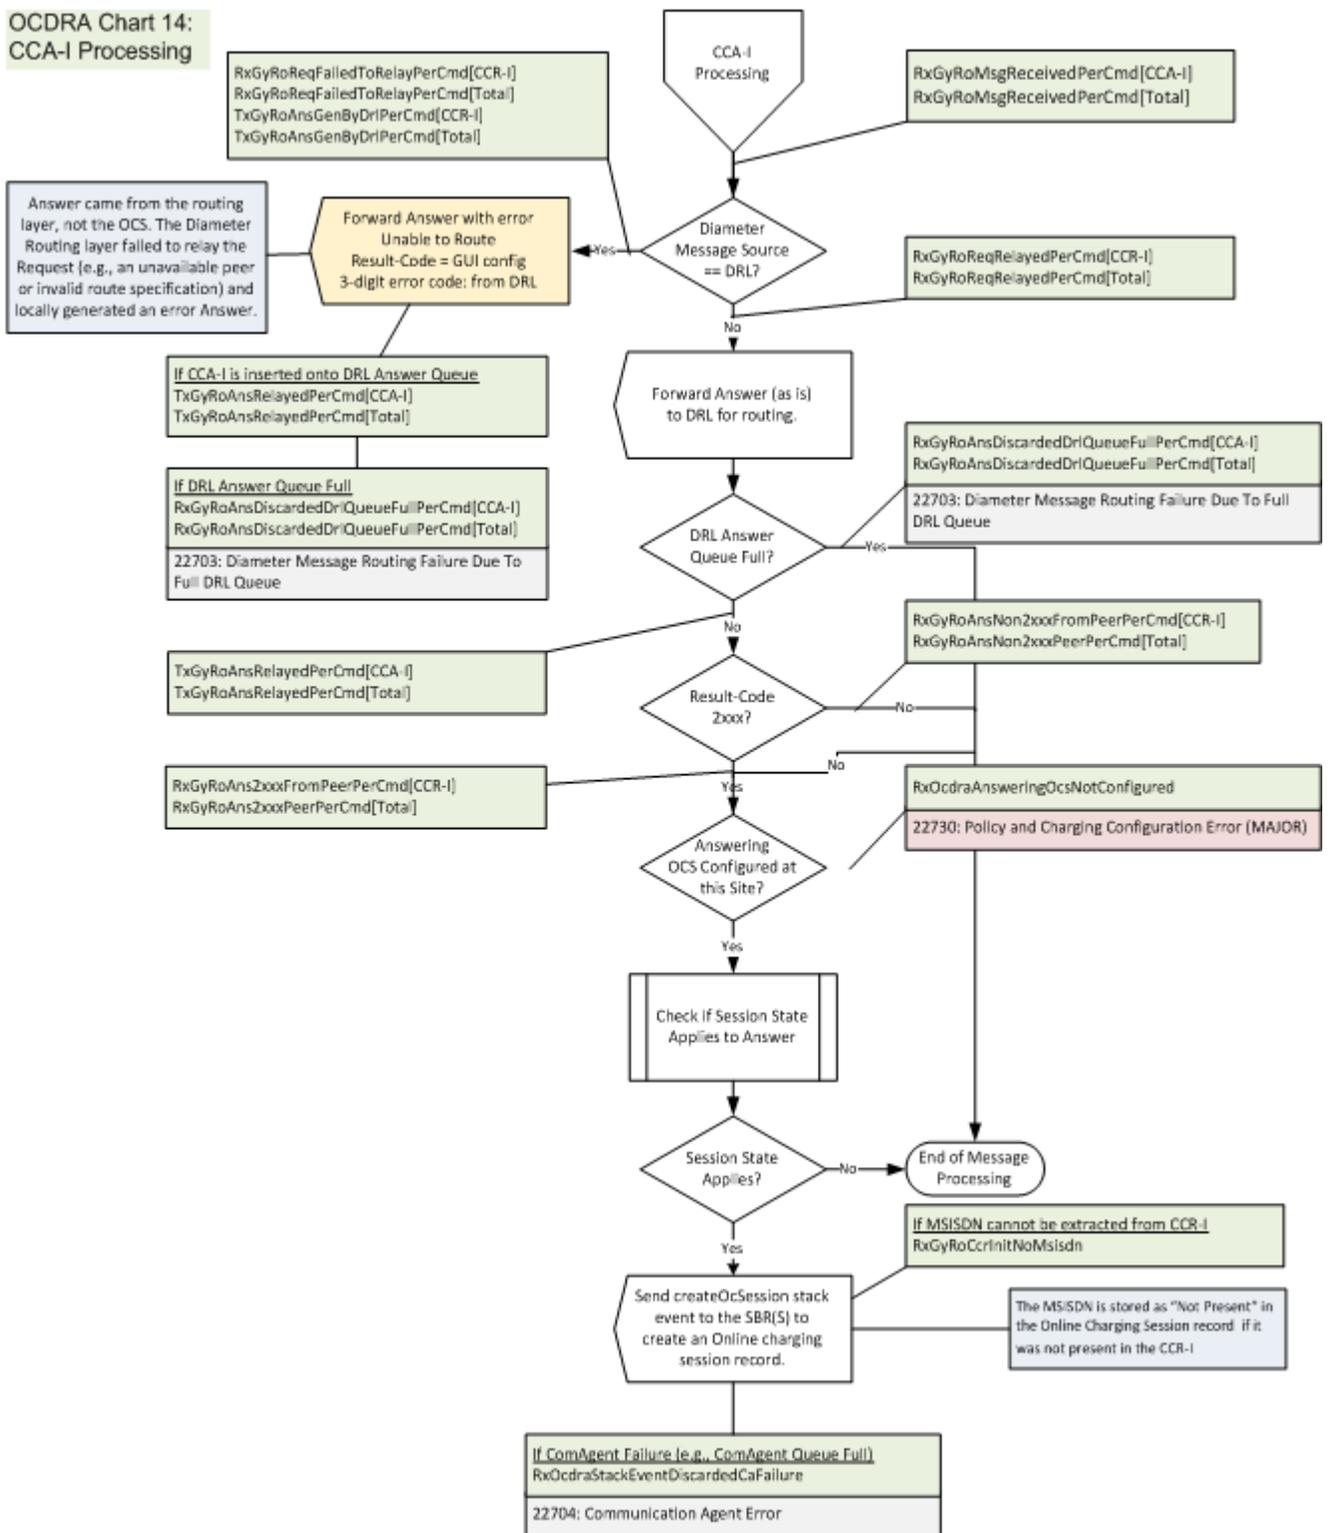


Figure 86: CCA-I Processing

CCA-U Processing

Figure 87: CCA-U Processing shows an error resolution flowchart that illustrates where CCA-U Processing error.

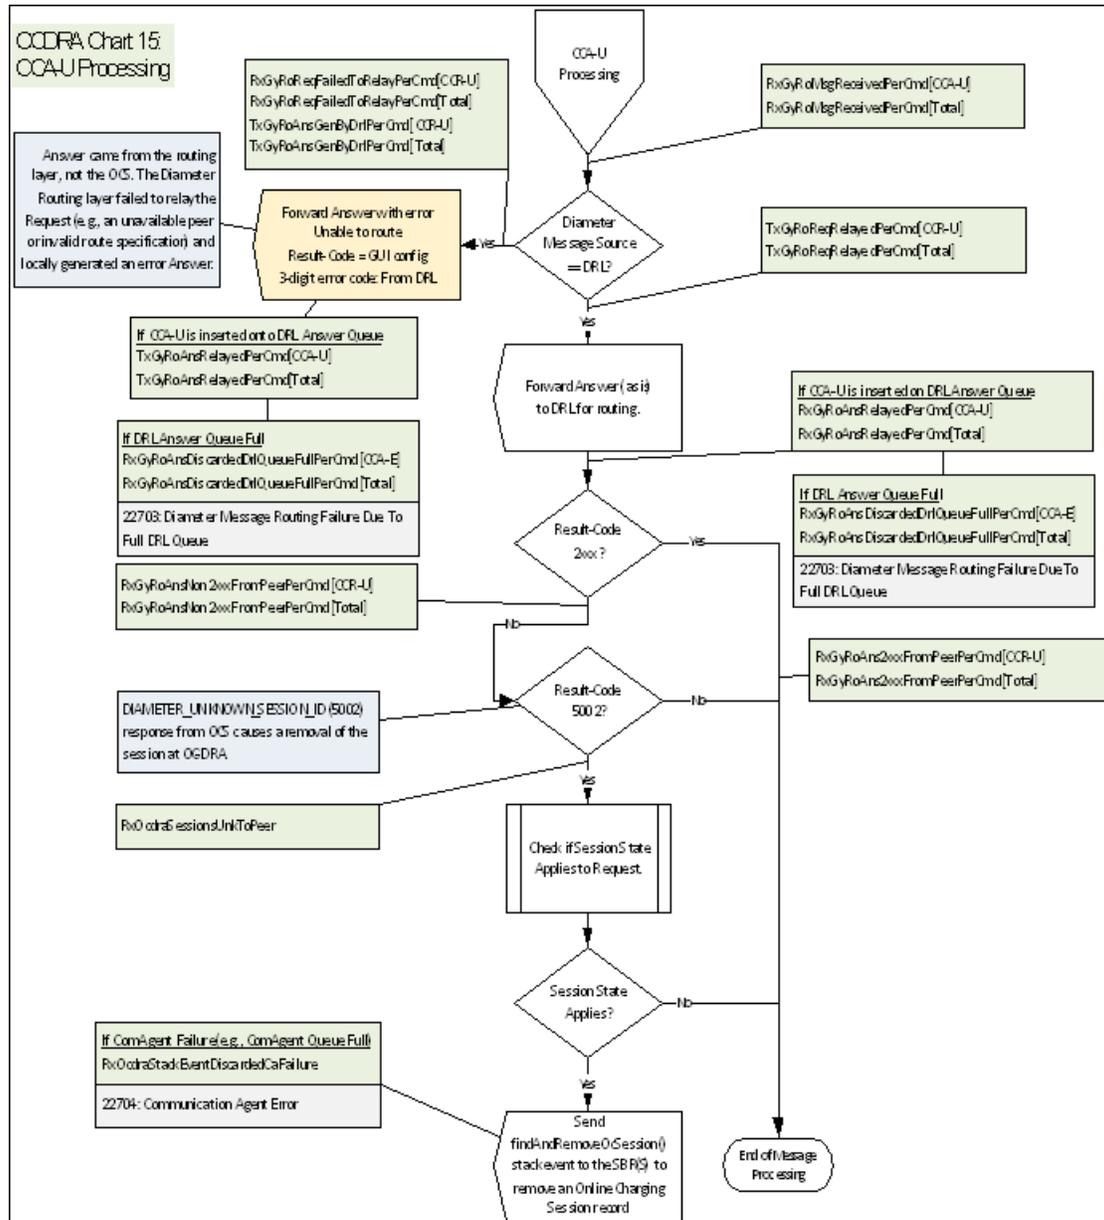


Figure 87: CCA-U Processing

CCA-T Processing

CCA-T Processing shows an error resolution flowchart that illustrates where CCA-T Processing error.

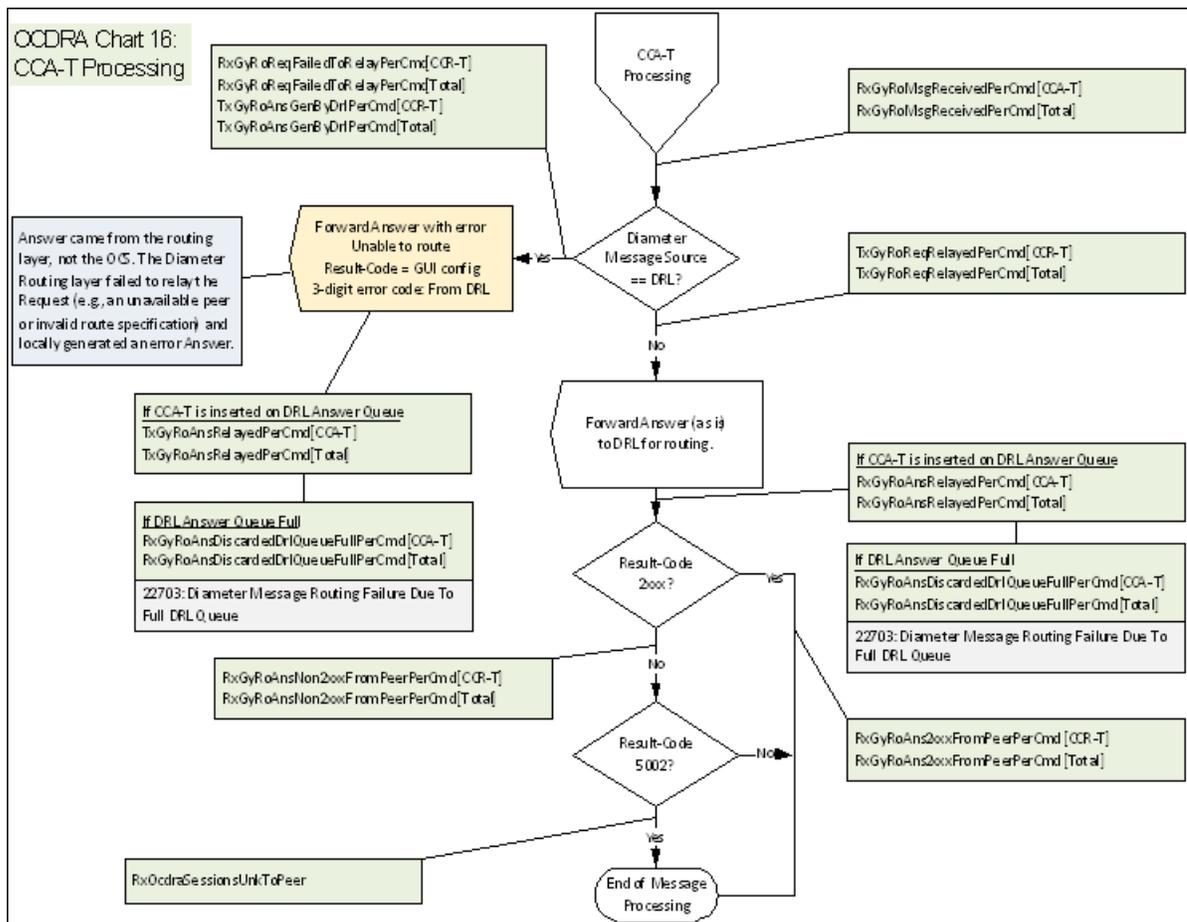


Figure 88: CCA-T Processing

CCA-E Processing

Figure 89: CCA-E Processing shows an error resolution flowchart that illustrates where CCA-E Processing error.

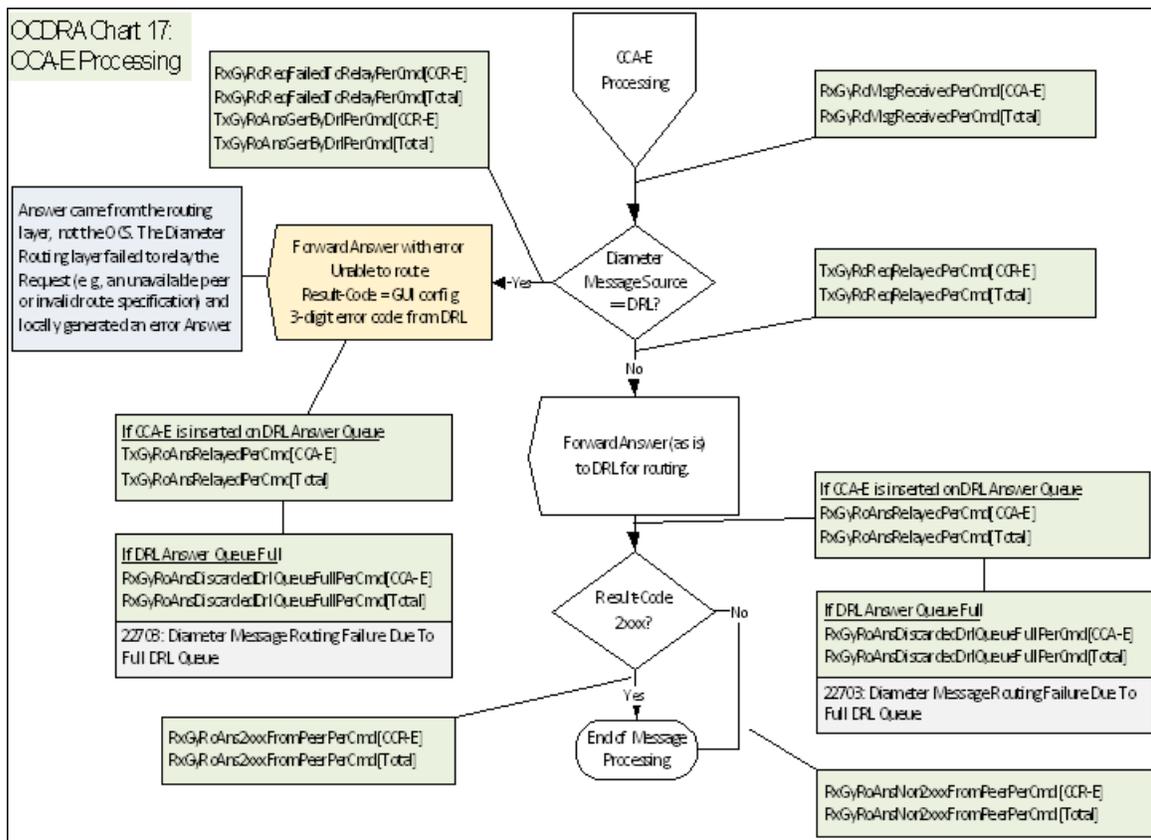


Figure 89: CCA-E Processing

RAA Processing

RAA Processing shows an error resolution flowchart that illustrates where RAA Processing error.

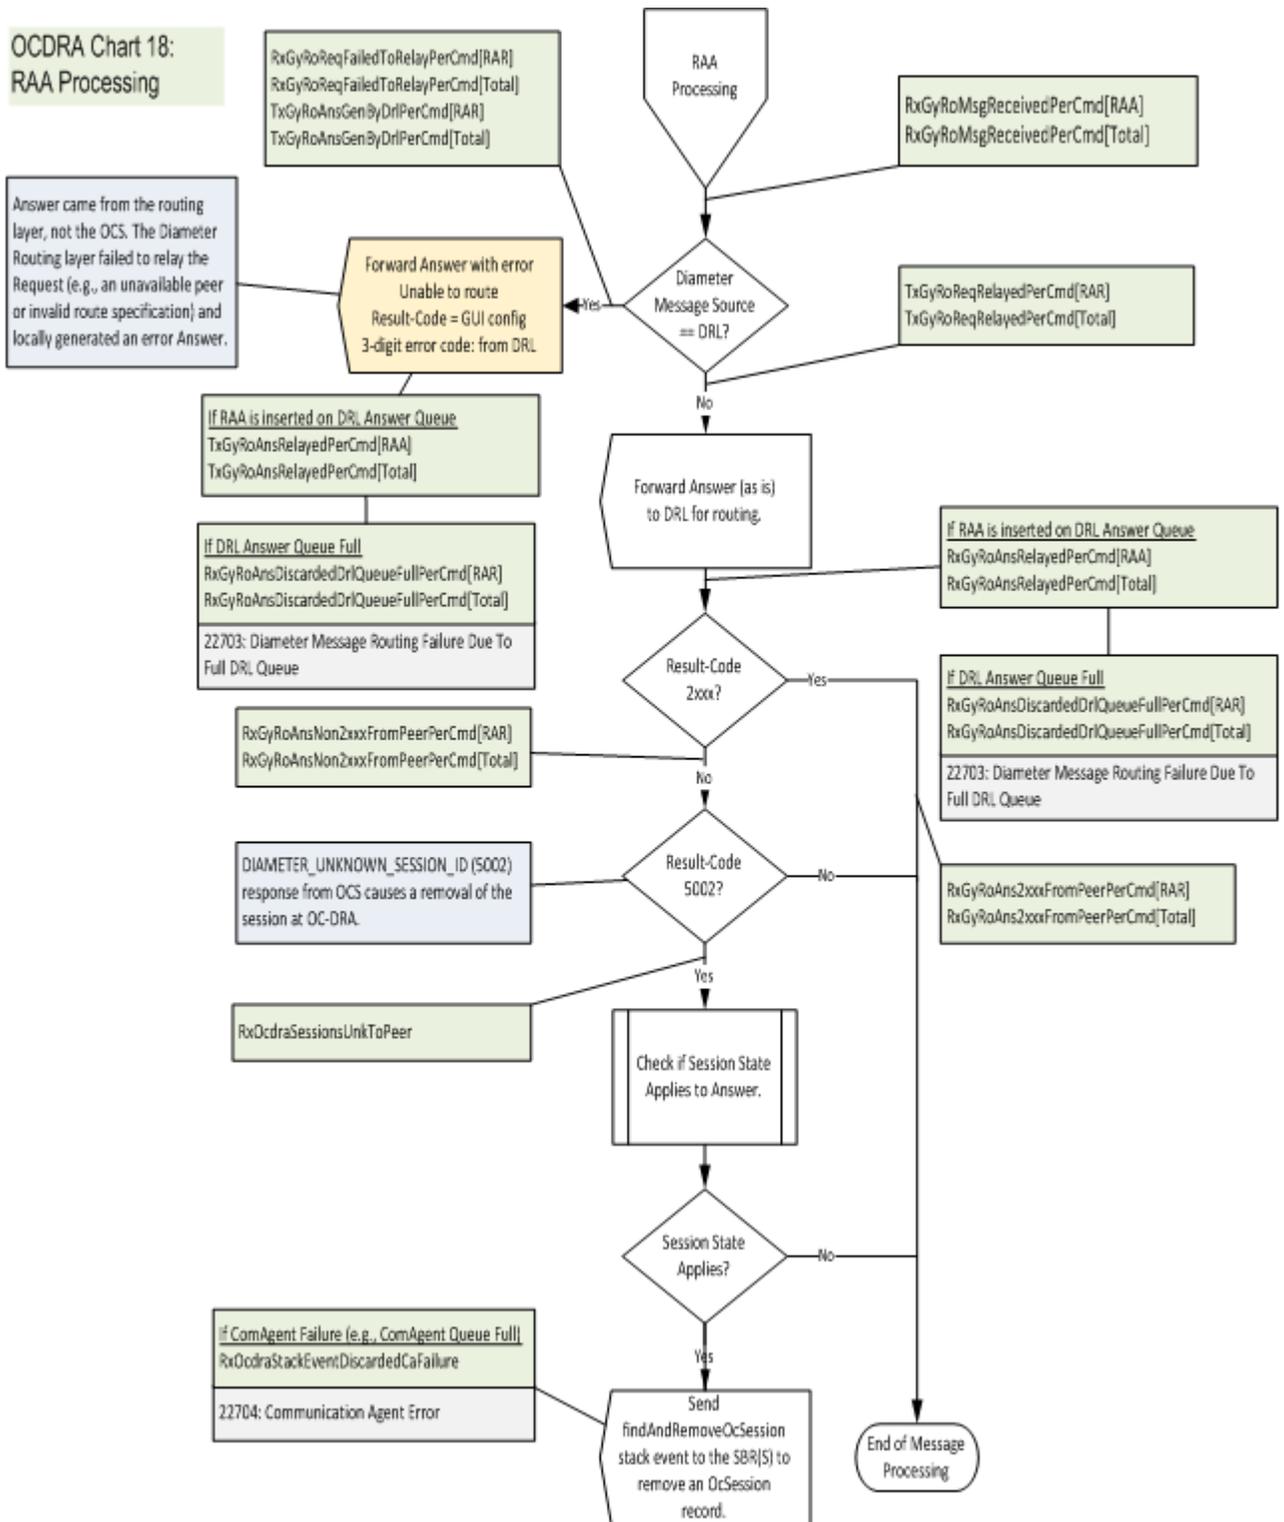


Figure 90: RAA Processing

A

AAA	Authentication, Authorization, and Accounting (Rx Diameter command)
AF	Application Function (such as P-CSCF)
Alternate Key	A subscriber key other than the anchor subscriber key, i.e., IP addresses or MSISDNs. Binding capable interfaces can include alternate subscriber keys. Binding dependent interfaces (Rx) cannot add alternate subscriber keys, but they can use them to find a binding.
APN	Access Point Name The name identifying a general packet radio service (GPRS) bearer service in a GSM mobile network. See also GSM.

B

BBERF	Bearer Binding and Event Reporting Function: A type of Policy Client used to control access to the bearer network (AN).
Binding Capable Interface	Gx and Gxx interfaces are capable of creating a binding if no binding exists for a subscriber. The CCR-I message must include the anchor subscriber key and may include alternate subscriber keys.

B

Binding database	Policy SBR database that holds network-wide subscriber binding information. Maps subscriber keys to the PCRF that hosts the subscriber's policy rules. A given binding record is maintained by 3 servers in the network: an Active server, a Standby server, and a Spare server.
------------------	--

C

CCA-I	Credit Control Answer – Initial
-------	---------------------------------

CCR-I	CCR Initial
-------	-------------

CEX Configuration Set	A mechanism for assigning Application IDs and supported Vendor IDs to a Local Node or to a Connection.
-----------------------	--

CTF	Charging Trigger Function
-----	---------------------------

D

DA-MP	Diameter Agent Message Processor A DSR MP (Server Role = MP, Server Group Function = Diameter Signaling Router). A local application such as CPA can optionally be activated on the DA-MP. A computer or blade that is hosting a Diameter Signaling Router Application.
-------	--

DIH	Diameter Intelligence Hub A troubleshooting solution for LTE, IMS, and 3G Diameter traffic processed by the DSR. DIH does not require separate probes or taps.
-----	---

D

DPI	Diameter Plug-In is a reusable Diameter stack consisting of DCL, DRL, and an application interface.
DRL	Diameter Routing Layer - The software layer of the stack that implements Diameter routing.
DSR	Diameter Signaling Router A set of co-located Message Processors which share common Diameter routing tables and are supported by a pair of OAM servers. A DSR Network Element may consist of one or more Diameter nodes.

F

FQDN	Fully Qualified Domain Name The complete domain name for a specific computer on the Internet (i.e., www.oracle.com). A domain name that specifies its exact location in the tree hierarchy of the DNS.
------	--

G

GUI	Graphical User Interface The term given to that set of items and facilities which provides you with a graphic means for manipulating screen data rather than being limited to character based commands.
Gx	The Diameter credit control based interface between a PCRF and a PCEF as defined by 3GPP. The interface is used to convey session information from the PCEF to the

G

PCRF, and in reply the PCRF provides rule information for the PCEF to enforce.

Gxx

Short for Gxa and Gxc. The Diameter credit control based interface between a BBERF and a PCRF, as defined by 3GPP.

I

IANA

Internet Assigned Numbers Authority

An organization that provides criteria regarding registration of values related to the Diameter protocol.

IDIH

Integrated Diameter Intelligence Hub

IMSI

International Mobile Subscriber Identity

IPFE

IP Front End

A traffic distributor that routes TCP traffic sent to a target set address by application clients across a set of application servers. The IPFE minimizes the number of externally routable IP addresses required for application clients to contact application servers.

M

MOS

Media Optimization Server

MSISDN

Mobile Station International Subscriber Directory Number

M

The MSISDN is the network specific subscriber number of a mobile communications subscriber. This is normally the phone number that is used to reach the subscriber.

N

NOAM

Network Operations, Administration, and Maintenance

O

OAM

Operations, Administration, and Maintenance

The application that operates the Maintenance and Administration Subsystem that controls the operation of many products.

P

PCEF

Policy and charging enforcement function

A system responsible for enforcing policies on network subscriber authentication, authorization, accounting, and mobility. A PCEF device, such as a CMTS or GGSN, communicates with a PCRF device, such as a policy server.

PCRF

Policy and Charging Rules Function

The ability to dynamically control access, services, network capacity, and charges in a network.

Maintains rules regarding a subscriber's use of network resources. Responds to CCR and AAR messages. Periodically sends RAR messages. All policy sessions for a given subscriber, originating

P

anywhere in the network, must be processed by the same PCRF.

P-DRA

Policy DRA

Peer Route Table

A set of prioritized Peer Routing Rules that define routing to Peer Nodes based on message content.

Place

An OAM configured component that defines physical locations. The Site Place groups the servers at a physical location. Each server is associated with exactly one Site Place.

Place Association

An OAM configured component used by P-DRA to group Site Places into Policy DRA Mated Pairs and Policy DRA Binding Regions.

PRT

Peer Route Table or Peer Routing Table

R

RAA

Re-Authorization Answer (Gx or Rx Diameter command)

RAR

Re-Authorization Request (Gx or Rx Diameter command)

RBAR

Range Based Address Resolution

A DSR enhanced routing application which allows you to route Diameter end-to-end transactions based on Application ID, Command Code, "Routing Entity" Type, and Routing Entity address ranges.

R

Resource Domain

A list of Server Groups that support a logical resource.

S

S9

The S9 Diameter interface includes Rx, Gx, and Gxx messages, but when these messages are used between a visited PCRF and the home PCRF, the interfaces are collectively referred to as S9. Defined by 3GPP 29.215 as the interface between a visited PCRF and a home PCRF. There is no difference in processing of Rx over S9 versus. Rx not over S9. The S9 interface is binding capable for Gx and Gxx only. Rx over S9 is binding dependent.

SBR

Session Binding Repository

A highly available, distributed database for storing Diameter session binding data.

SOAM

System Operations,
Administration, and Maintenance

Subscriber Key

One of several possible keys that can be used to uniquely identify a subscriber. Subscriber Keys are delivered in the Subscriber-Id Diameter AVP of a CCR-I message. One of the Subscriber Keys is designated as an Anchor Key.

Suggested PCRF

PCRF that will be used for the binding unless an error causes alternate routing. Avoids the need to update the binding if the suggested PCRF successfully answers the CCR-I.

S

Suspect Binding

A Policy DRA IMSI Anchor Key binding record is considered to be “suspect” if the last attempt to route a CCR-I message to the bound PCRF failed with a 3002 Error Code response. The concept of Suspect Binding allows bindings to be removed after a short period of time (called the Suspect Binding Interval) from a PCRF that has become unreachable.

T

TSA

Target Set Address

An externally routable IP address that the IPFE presents to application clients. The IPFE distributes traffic sent to a target set address across a set of application servers.

TTR

Trace Transaction Record - A record describing a Diameter transaction, including all of the Diameter messages that were part of the transaction, plus the operations performed by DSR while processing those messages.

U

UE

User Equipment

V

V-PCRF

Visited PCRF

X

XSI

External Signaling Interface