

**Oracle® Communications
Diameter Signaling Router**

Policy DRA User Guide

910-6921-001 Revision A

May 2014

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, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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 on 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. 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.

Table of Contents

Chapter 1: Document Introduction.....	12
Purpose of this Manual.....	13
Manual Organization.....	13
Scope and Audience.....	13
Documentation Admonishments.....	13
Related Publications.....	14
Customer Care Center.....	15
Emergency Response.....	18
Locate Product Documentation on the Customer Support Site.....	18
 Chapter 2: Policy DRA Introduction.....	 19
The Policy DRA Application.....	20
The Policy DRA Database.....	20
The Binding Database.....	23
The Session Database.....	25
Policy DRA Functions.....	25
Diameter Request Message Processing.....	26
Query Subscriber's Binding Status.....	26
PCRF Selection and Routing.....	28
Topology Hiding Process.....	28
Diameter Answer Message Processing.....	28
Subscriber Session and Binding Database Management.....	29
The Communication Agent.....	29
Subscriber Identification and Binding.....	30
Diameter Routing and Communication with Policy DRA.....	31
Ingress Routing.....	34
Egress Routing.....	35
 Chapter 3: Policy DRA Overview.....	 39
PCRF Pools and Sub-Pools Concepts and Terminology.....	40
Binding-capable Sessions.....	51
Binding-dependent Sessions.....	56
In-session Message Processing.....	58

Topology Hiding.....	58
Session Integrity.....	59
Policy Data Auditing.....	62
Policy DRA Assumptions and Limitations.....	67
Policy DRA Capacity Constraints.....	69
Chapter 4: Policy DRA Deployment.....	70
High Level Deployment Description.....	71
Deployment Topology.....	71
Policy DRA in Roaming Scenarios.....	72
Policy DRA Configurable Components.....	73
IPFE.....	83
Redundancy.....	85
MP Server Redundancy.....	85
Site Redundancy.....	86
Data Redundancy.....	86
OAM Server Redundancy.....	87
Policy DRA Scalability.....	89
MP Growth.....	90
Database Growth.....	90
Mated Pair Growth.....	91
Small System Support.....	92
IP Networking.....	93
Chapter 5: Policy DRA Configuration.....	95
Policy DRA Configuration Overview.....	96
Pre-Configuration Activities.....	98
Initial Installation for PCRF Pooling.....	99
Diameter Configuration for Policy DRA.....	100
Policy DRA Configuration on the NOAM and the SOAM.....	104
Access Point Names.....	104
PCRF Pools.....	110
PCRF Pool to PRT Mapping.....	116
PCRF Sub-Pool Selection Rules.....	120
Network-Wide Options.....	126
Alarm Settings.....	130
Congestion Options.....	134
PCRFs.....	137
Binding Key Priority.....	141
Topology Hiding.....	142

Site Options.....	145
Error Codes.....	148
Post-Configuration Activities.....	153
Enable the Policy DRA Application.....	153
Enable Connections.....	153
Status Verification.....	154
DSR Bulk Import and Export.....	154
Chapter 6: Policy DRA Maintenance.....	156
Introduction.....	157
Policy DRA Maintenance Pages.....	157
Policy SBR Status.....	157
The Binding Key Query.....	158
Alarms, KPIs, and Measurements.....	161
Policy DRA and Policy SBR Alarms and Events.....	161
Policy DRA and Policy SBR KPIs.....	162
Policy DRA and Policy SBR Measurements.....	163
Binding and Session Database Auditing.....	163
Overload Management.....	163
Overload Controls.....	163
Shutdown.....	166
Diameter Maintenance and Status Data for Components, DSR Applications, and DA-MPs.....	167
Backup and Restore for Policy DRA Configuration Data.....	168
Appendix A: Policy DRA PCRF Pooling Upgrade.....	169
Upgrade Paths.....	170
Configuration After Upgrade.....	172
Concepts and Terminology.....	173
Configuring PCRF Pooling.....	176
Processing Phases.....	179
Binding Migration.....	180
Appendix B: Policy DRA Error Resolution.....	183
Introduction.....	185
Policy DRA Error Resolution Flowchart Summary.....	187
Diameter Message Validation Error Resolution Flowchart.....	189
Generic CCR Processing Error Resolution Flowchart.....	190
CCR-I Processing without PCRF Pool Error Resolution Flowchart.....	191

findSessionRef Processing Error Resolution Flowchart.....	193
findOrCreBindResShort Processing Error Resolution Flowchart.....	193
CCR-I Processing with PCRF Pool Error Resolution Flowchart.....	196
findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart.....	197
Early Bind Pool Error Resolution Flowchart.....	198
CCA-I Processing Error Resolution Flowchart.....	199
findSession Response Processing Error Resolution Flowchart.....	200
CCR-U Processing Error Resolution Flowchart.....	201
CCR-T Processing Error Resolution Flowchart.....	202
CCA-U/T Processing Error Resolution Flowchart.....	203
RAR Processing Error Resolution Flowchart.....	204
RAA Processing Error Resolution Flowchart.....	205
AAR Processing Error Resolution Flowchart.....	206
AAA Processing Error Resolution Flowchart.....	207
STR Processing Error Resolution Flowchart.....	208
STA Processing Error Resolution Flowchart.....	209
ASR/ASA Processing Error Resolution Flowchart.....	210
Glossary.....	212

List of Figures

Figure 1: Communication between ComAgents, Policy DRA, and Policy SBR.....	29
Figure 2: Subscriber Key Usage.....	31
Figure 3: Request Processing at the Diameter Routing Function and Policy DRA	32
Figure 4: Answer Processing at the Diameter Routing Function and Policy DRA.....	32
Figure 5: Policy DRA Generated Answer Routing.....	33
Figure 6: Policy DRA Generated Request Routing.....	33
Figure 7: Relationship between APNs and PCRF Pools.....	41
Figure 8: Relationship between IMSIs and PCRF Pools.....	42
Figure 9: Multiple PCRF Pools.....	43
Figure 10: Multiple PCRF Versions in a PCRF Pool.....	44
Figure 11: PCRF Pools and Sub-Pools Routing Scenarios.....	45
Figure 12: Find or Create a Binding.....	55
Figure 13: Binding Dependent Session Initiation Request Processing Overview.....	57
Figure 14: Sites, Mated Pairs, and Region.....	72
Figure 15: Policy DRA in Roaming Scenarios.....	73
Figure 16: Policy DRA Component Relationships.....	74
Figure 17: Example Policy DRA Mated Pair - Hosting Binding Policy SBRs.....	76
Figure 18: Example Policy DRA Mated Pair - Not Hosting Binding Policy SBRs.....	77
Figure 19: Policy Client, PCRF, and Site Relationships.....	78
Figure 20: Resource Domains.....	81
Figure 21: Binding Table Partitioning Across Server Groups.....	87
Figure 22: Multi-Table Resources.....	87

Figure 23: Data Merging - Normal Case.....	88
Figure 24: Data merging - Redundant Site Failure.....	89
Figure 25: Smallest Supported Policy DRA Mated Pair.....	93
Figure 26: GUI Structure for 3-tiered DSR Topology with Policy DRA.....	97
Figure 27: PCRF Pooling Data.....	112
Figure 28: Policy DRA Default Overload Control Thresholds.....	164
Figure 29: PCRF Pooling Effects on Policy DRA.....	175
Figure 30: Error Resolution Flowchart Summary.....	189
Figure 31: Diameter Message Validation Error Resolution Flowchart.....	190
Figure 32: Generic CCR Processing Error Resolution Flowchart.....	191
Figure 33: CCR-I Processing without PCRF Pool Error Resolution Flowchart.....	192
Figure 34: findSessionRef Processing Error Resolution Flowchart.....	193
Figure 35: findOrCreBindResShort Processing Error Resolution Flowchart.....	195
Figure 36: CCR-I Processing with PCRF Pool Error Resolution Flowchart.....	196
Figure 37: findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart.....	197
Figure 38: Early Bind Pool Error Resolution Flowchart.....	199
Figure 39: CCA-I Processing Error Resolution Flowchart.....	200
Figure 40: findSession Response Processing Error Resolution Flowchart.....	201
Figure 41: CCR-U Processing Error Resolution Flowchart.....	202
Figure 42: CCR-T Processing Error Resolution Flowchart.....	203
Figure 43: CCA-U/T Processing Error Resolution Flowchart.....	204
Figure 44: RAR Processing Error Resolution Flowchart.....	205
Figure 45: RAA Processing Error Resolution Flowchart.....	206
Figure 46: AAR Processing Error Resolution Flowchart.....	207

Figure 47: AAA Processing Error Resolution Flowchart.....	208
Figure 48: STR Processing Error Resolution Flowchart.....	209
Figure 49: STA Processing Error Resolution Flowchart.....	210
Figure 50: ASR/ASA Processing Error Resolution Flowchart.....	211

List of Tables

Table 1: Admonishments.....	14
Table 2: Example of a Binding Key Priority Configuration.....	24
Table 3: Example Key Priority Configuration.....	31
Table 4: Communication between the Diameter Routing Function and the DAI.....	32
Table 5: Policy DRA Application Routing Rule Configuration.....	34
Table 6: Answer Processing by Policy DRA.....	35
Table 7: Policy DRA Terminology.....	46
Table 8: Topology Hiding Scope Configuration.....	59
Table 9: Policy DRA Error Scenarios for Session Integrity.....	60
Table 10: Session Integrity Conditions and Policy DRA Reaction.....	62
Table 11: Effects of Stale Binding and Session Records.....	63
Table 12: Audit Report Formats.....	65
Table 13: Server Group Functions.....	79
Table 14: Policy SBR Server Group Configuration and Data Redundance.....	79
Table 15: Policy Client Connection Capability.....	81
Table 16: IP Traffic-to-Service Mapping.....	93
Table 17: Access Point Names elements.....	106
Table 18: PCRF Pooling Concepts.....	111
Table 19: PCRF Pooling Configuration Summary.....	113
Table 20: PCRF Pools elements.....	114
Table 21: PCRF Pools to PRT Mapping elements.....	118
Table 22: PCRF Sub-Pool Selection Rules elements.....	122

Table 23: Network-Wide Options elements.....	127
Table 24: Alarm Settings elements.....	131
Table 25: Congestion Options elements.....	135
Table 26: PCRFs page elements.....	138
Table 27: Binding Key Priority elements.....	141
Table 28: Topology Hiding elements.....	143
Table 29: Site Options elements.....	146
Table 30: Topology Hiding Scope Configuration.....	147
Table 31: Policy DRA Error Conditions.....	149
Table 32: Interfaces Supported for Each Error Code.....	151
Table 33: Error Codes elements.....	151
Table 34: Policy SBR Status elements.....	158
Table 35: Binding Key Query elements.....	160
Table 36: Diameter Routing Function Message Handling Based on Policy DRA Operational Status.....	164
Table 37: Stack Event Load Shedding.....	166
Table 38: Policy DRA Operational Status.....	167
Table 39: Processing During Transition Period.....	171
Table 40: Upgrading Policy DRA Terminology.....	175
Table 41: Error Resolution Attributes.....	185

Chapter 1

Document Introduction

Topics:

- *Purpose of this Manual.....13*
- *Manual Organization.....13*
- *Scope and Audience.....13*
- *Documentation Admonishments.....13*
- *Related Publications.....14*
- *Customer Care Center.....15*
- *Emergency Response.....18*
- *Locate Product Documentation on the Customer Support Site.....18*

This chapter contains a brief description of the Policy Diameter Routing Agent (Policy DRA) feature. The contents include sections about the document scope, audience, and organization; how to find related publications; and how to contact customer assistance.

Purpose of this Manual

This content:

- Gives a conceptual overview of the application's purpose, architecture, and functionality
- Describes the pages and elements on the application GUI (Graphical User Interface)
- Provides procedures for using the application interface
- Explains the organization of, and how to use, the documentation

Manual Organization

This manual is organized into the following chapters:

- [Document Introduction](#) contains general information about the DSR documentation, the organization of this manual , and how to get technical assistance.
- [Policy DRA Introduction](#) describes the topology, architecture, components, and functions of the Policy DRA application and the Policy Session Binding Repository (Policy 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.
- [Policy DRA Deployment](#) describes Policy DRA and Policy SBR deployment in a DSR system.
- [Policy DRA Configuration](#) describes configuration of Policy DRA application components.
- [Policy DRA Maintenance](#) describes Policy DRA Maintenance functions, and Diameter Maintenance functions that provide maintenance and status information for Policy DRA and the Policy SBR.
- [Policy DRA PCRF Pooling Upgrade](#) describes how to upgrade for PCRF Pooling from Policy DRA release 4.1.5 or 5.0.
- [Policy DRA Error Resolution](#) describes information to help users diagnose and resolve Policy DRA errors encountered while processing Diameter messages in the Policy DRA application.

Scope and Audience

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

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

The Diameter Signaling Router (DSR) documentation set includes the following publications, which provide information for the configuration and use of DSR and related applications.

Getting Started includes a product overview, system architecture, and functions. It also explains the DSR GUI features including user interface elements, main menu options, supported browsers, and common user interface widgets.

Feature Notice describes new features in the current release, provides the hardware baseline for this release, and explains how to find customer documentation on the Oracle Customer Support Site.

Roadmap to Hardware Documentation provides links to access manufacturer online documentation for hardware related to the DSR.

Operation, Administration, and Maintenance (OAM) Guide provides information on system-level configuration and administration tasks for the advanced functions of the DSR, both for initial setup and maintenance.

Communication Agent User's Guide explains how to use the Communication Agent GUI pages to configure Remote Servers, Connection Groups, and Routed Servers, and to maintain configured connections.

Diameter User's Guide explains how to use the Diameter GUI pages to manage the configuration and maintenance of Diameter Configuration components, including Local and Peer Nodes, Connections, Configuration Sets, Peer Routing Rules, Application Route Tables, System Options, and DNS options; describes the functions of Diameter Message Copy; and describes DSR capacity and congestion controls.

Diameter Mediation User's Guide describes the functions of Diameter Mediation, and explains how to use the Diameter Mediation GUI pages (nested inside the Diameter GUI folder) to configure and test Rule Templates, how to use the Formatting Value Wizard, and how to configure Rule Sets.

IP Front End (IPFE) User's Guide explains how to use the IPFE GUI pages to configure IPFE to distribute IPv4 and IPv6 connections from multiple clients to multiple nodes.

Range-Based Address Resolution (RBAR) User's Guide explains how to use the RBAR GUI pages to configure RBAR to route Diameter end-to-end transactions based on Diameter Application ID, Command Code, Routing Entity Type, and Routing Entity address ranges and individual addresses.

Full-Address Based Resolution (FABR) User's Guide explains how to use the FABR GUI pages to configure FABR to resolve designated Diameter server addresses based on Diameter Application ID, Command Code, Routing Entity Type, and Routing Entity addresses.

Charging Proxy Application (CPA) and Offline Charging Solution User's Guide describes the Offline Charging Solution and explains how to use the CPA GUI pages to set System Options for CPA, configure the CPA's Message Copy capability, and configure the Session Binding Repository for CPA.

Policy DRA User's Guide describes the topology and functions of the Policy Diameter Routing Agent (Policy DRA) DSR Application and the Policy Session Binding Repository, and explains how to use the GUI pages to configure Policy DRA.

Gateway Location Application (GLA) User's Guide describes the functions of retrieving subscriber data stored in Policy Session Binding Repository (pSBR) provided by Policy DRA and explains how to use the GUI pages to configure GLA.

DSR Alarms, KPIs, and Measurements Reference provides detailed descriptions of alarms, events, Key Performance Indicators (KPIs), and measurements; indicates actions to take to resolve an alarm, event, or unusual Diameter measurement value; and explains how to generate reports containing current alarm, event, KPI, and measurement information.

DSR Administration Guide describes DSR architecture, functions, configuration, and tools and utilities (IPsec, Import/Export, DIH, and database backups); and provides references to other publications for more detailed information.

Customer Care Center

Oracle's Tekelec Customer Care Center is your initial point of contact for all product support needs. A representative takes your call or email, creates a Customer Service Request (CSR) and directs your requests to the Technical Assistance Center (TAC). Each CSR includes an individual tracking number. Together with TAC Engineers, the representative will help you resolve your request.

The Customer Care Center is available 24 hours a day, 7 days a week, 365 days a year, and is linked to TAC Engineers around the globe.

TAC Engineers are available to provide solutions to your technical questions and issues 7 days a week, 24 hours a day. After a CSR is issued, the TAC Engineer determines the classification of the trouble. If a critical problem exists, emergency procedures are initiated. If the problem is not critical, normal support procedures apply. A primary Technical Engineer is assigned to work on the CSR and provide a solution to the problem. The CSR is closed when the problem is resolved.

Technical Assistance Centers are located around the globe in the following locations:

Related - Global

Email (All Regions): support@tekelec.com

- **USA and Canada**

Phone:

1-888-367-8552 (toll-free, within continental USA and Canada)

1-919-460-2150 (outside continental USA and Canada)

TAC Regional Support Office Hours:

8:00 a.m. through 5:00 p.m. (GMT minus 5 hours), Monday through Friday, excluding holidays

- **Caribbean and Latin America (CALA)**

Phone:

+1-919-460-2150

TAC Regional Support Office Hours (except Brazil):

10:00 a.m. through 7:00 p.m. (GMT minus 6 hours), Monday through Friday, excluding holidays

- **Argentina**

Phone:

0-800-555-5246 (toll-free)

- **Brazil**

Phone:

0-800-891-4341 (toll-free)

TAC Regional Support Office Hours:

8:00 a.m. through 5:48 p.m. (GMT minus 3 hours), Monday through Friday, excluding holidays

- **Chile**

Phone:

1230-020-555-5468

- **Colombia**

Phone:

01-800-912-0537

- **Dominican Republic**

Phone:

1-888-367-8552

- **Mexico**

Phone:

001-888-367-8552

- **Peru**

Phone:

0800-53-087

- **Puerto Rico**

Phone:

1-888-367-8552

- **Venezuela**

Phone:

0800-176-6497

- **Europe, Middle East, and Africa**

Regional Office Hours:

8:30 a.m. through 5:00 p.m. (GMT), Monday through Friday, excluding holidays

- **Signaling**

Phone:

+44 1784 467 804 (within UK)

- **Software Solutions**

Phone:

+33 3 89 33 54 00

- **Asia**

- **India**

Phone:

+91-124-465-5098 or +1-919-460-2150

TAC Regional Support Office Hours:

10:00 a.m. through 7:00 p.m. (GMT plus 5 1/2 hours), Monday through Saturday, excluding holidays

- **Singapore**

Phone:

+65 6796 2288

TAC Regional Support Office Hours:

9:00 a.m. through 6:00 p.m. (GMT plus 8 hours), Monday through Friday, excluding holidays

Emergency Response

In the event of a critical service situation, emergency response is offered by Oracle's Tekelec Customer Care Center 24 hours a day, 7 days a week. 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's Tekelec Customer Care Center.

Locate Product Documentation on the Customer Support Site

Oracle customer documentation is available on the web at the Oracle Technology Network (OTN) 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 www.adobe.com.

1. Log into the Oracle Customer Support site at <http://docs.oracle.com>.
2. Under **Applications**, click the link for **Communications**.
The **Oracle Communications Documentation** window opens with Tekelec shown near the top.
3. Click **Oracle Communications Documentation for Tekelec Products**.
4. Navigate to your Product and then the Release Number, and click the **View** link (the **Download** link will retrieve the entire documentation set).
5. To download a file to your location, right-click the PDF link and select **Save Target As**.

Chapter 2

Policy DRA Introduction

Topics:

- [The Policy DRA Application.....20](#)
- [The Policy DRA Database.....20](#)
- [Policy DRA Functions.....25](#)
- [The Communication Agent.....29](#)
- [Subscriber Identification and Binding.....30](#)
- [Diameter Routing and Communication with Policy DRA.....31](#)

This section introduces Policy DRA-related applications, key concepts, and basic functionality.

The Policy Diameter Routing Agent (Policy DRA) is a feature of the Diameter Signaling Router (DSR) product, which is part of the Oracle product line of signaling products. Policy DRA runs as a DSR Application, to solve Diameter routing problems that are specific to the policy management domain.

Note: If you are upgrading from a Policy DRA release prior to 5.1 and Policy DRA was activated on that release, see [Policy DRA PCRF Pooling Upgrade](#). The differences between Policy DRA prior to DSR 5.1 and Policy DRA in 5.1 and later releases are described there.

The Policy DRA Application

Policy DRA offers a scalable, geo-diverse Diameter application 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 Policy Clients.

Policy DRA provides the following capabilities:

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

The Policy DRA Database

The Policy DRA application uses the Session and Binding databases in the Policy Session Binding Repository. 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 Policy DRA Binding and Session databases:

- There is one instance of the Binding database in the entire Policy DRA network.
- There is one instance of the Session database per Policy DRA Mated Pair.

- 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, MSISDN, and IPv6 Alternate Keys are optional. They represent alternate ways, other than the IMSI, to identify a subscriber.

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

Policy Session Binding Repository (pSBR) servers host the Policy Session and Policy Binding databases for use by the Policy DRA 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.

Policy 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 PCEF 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 Policy DRA must be accessible from anywhere in the Policy DRA network. This means that the information in the Binding database must be accessible from all Policy 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 Policy DRA DSR in the Policy DRA network.

PCRF Pooling now 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 DRA > Configuration > Access Point Names**.

Policy 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 DRA > Configuration > 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 DRA > Configuration > PCRF Pool To PRT Mapping**. 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. Policy 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). The Policy DRA application 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.

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. An anchor key is required for binding-capable session initiation requests. This allows a binding to be established (with the exception of emergency calls from a phone without a SIM card). Serialization is accomplished by requiring that binding-capable session origination messages (CCR-I) always contain an IMSI and that the IMSI is always used for creation of new bindings. IMSI is the anchor key for Gx, Gxx, and S9.

See [Error Codes](#).

Alternate Keys

Alternate Keys provide different ways to identify a subscriber for binding-dependent interfaces (Rx and Gx-Prime). 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 Policy 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.

Policy DRA allows the priority 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 2: 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 2: 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 Policy DRA looks for a given key type in a binding-dependent session initiating message. In the example in [Table 2: Example of a Binding Key Priority Configuration](#), Policy DRA 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, Policy DRA will attempt to find a Binding record in the corresponding Binding database table. If a key is not present, Policy DRA 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 Table

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.

SessionRef Table

SessionRef records are used to tie Binding records to Diameter sessions. This allows Policy 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 Policy DRA database audit. See [Binding and Session Database Auditing](#) for more information about the database audit.

Policy DRA Functions

The Policy DRA application 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 application 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 Policy 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 Policy 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 Policy SBR interface. A response to the request from the Active Policy 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 Policy SBR for the binding status for this session. Depending on the output from the interactions with the Policy 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 Policy 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.

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 application no longer performs a round-robin selection among all configured PCRFs. The Policy DRA application 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 [Topology Hiding](#).

Diameter Answer Message Processing

After the Policy DRA routes a Diameter Request message to a selected PCRF, and updates the Policy 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 Policy 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

The Communication Agent

The Communication Agent (ComAgent) enables reliable communication between Policy DRA and Policy SBRs and among Policy SBRs in a scalable and high available Policy DRA network. [Figure 1: Communication between ComAgents, Policy DRA, and Policy SBR](#) depicts the communication paths between the Policy DRA, the Policy SBR, and their ComAgents, and the communication paths between the ComAgents.




Figure 1: Communication between ComAgents, Policy DRA, and Policy SBR

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

Policy 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 pSBRs in the network.

- All session pSBRs in the network connect to all binding pSBRs in the network
- All DA-MPs in a mated pair connect to all session pSBRs 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 DRA > Maintenance > Policy SBR Status** page.

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 Policy 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 2: 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 application 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 2: 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 3: Example Key Priority Configuration](#) illustrates an example Binding Key configuration with priorities assigned to each key. IMSI, IPv4, IPv6, or MSISDN

Table 3: Example Key Priority Configuration

Priority	Key Type
1	IMSI
2	IPv4
3	MSISDN
4	IPv6
5	<Not configured>

The example configuration in [Table 3: 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 application will NOT look for MSISDN or use it in the Binding database.

Diameter Routing and Communication with Policy DRA

The Policy DRA, as a DSR 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 4: Communication between the Diameter Routing Function and the DAI describes two functions for communication between the Diameter Routing Function and the DAI.

Table 4: Communication between the Diameter Routing Function and the DAI

Function	Communication Direction	Description
Application Data	Policy DRA <-> Diameter Routing Function	Either a Request or an Answer with supporting information
Application Status	Policy DRA <-> Diameter Routing Function	The Policy DRA 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 Policy DRA based on the configured Application Routing Rule, and routed from the Policy DRA to the Diameter Routing Function, all using the Application-Data function. The Policy DRA will return the Request to the Diameter Routing Function for Peer Routing Rule processing and routing.




Figure 3: Request Processing at the Diameter Routing Function and Policy DRA

Answer Routing

When the Policy DRA 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 Policy DRA or to route the Answer to the downstream Peer without involving the Policy DRA. *Figure 4: Answer Processing at the Diameter Routing Function and Policy DRA* shows the case where an Answer is transmitted back to the Policy DRA. After the Policy DRA 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 4: Answer Processing at the Diameter Routing Function and Policy DRA

Policy DRA Generated Answer

In some cases, the Policy DRA needs to generate an Answer message in response to an incoming Request. For example, the Policy DRA cannot find a PCRF to route the Request message to. [Figure 5: Policy DRA Generated Answer Routing](#) shows the Diameter Routing Function routing for this scenario.




Figure 5: Policy DRA Generated Answer Routing

Policy DRA Generated Request

In some cases, the Policy DRA needs to generate Diameter Requests. [Figure 6: Policy DRA Generated Request Routing](#) shows the Diameter Routing Function routing for this scenario.




Figure 6: Policy DRA Generated Request Routing

Policy DRA Application Use Cases

The following typical Policy DRA application signaling use cases demonstrate the Policy DRA and Policy 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](#).

Ingress Routing

This section describes how Diameter Request and Answer messages are routed to Policy DRA.

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 an Application Routing Rule?

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 Application Routing Rule, is the DSR Application Operational Status for this DA-MP set to Available?

If the DSR Application is not Available, then the "Unavailability action" is performed by Diameter. For Policy DRA, the Unavailability action is "Continue Routing", which means to route the Request using PRT Peer Routing Rules.

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 Diameter Configuration Application Routing Rules. [Table 5: Policy DRA Application Routing Rule Configuration](#) describes the expected configuration of Application Routing Rules for Policy DRA. 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 Policy DRA application. Some of these rules can be omitted, depending on which interfaces are used for Policy DRA.

- 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 (one rule is more specific than another) can use the Priority field to remove ambiguity about which rule should be chosen. ("Best Match" semantics is not supported for Application Routing Rules.)
- 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 is always PDRA for the Policy DRA application. PDRA appears in the Application Name pulldown list only if the Policy DRA feature has been activated.

Table 5: Policy DRA Application Routing Rule Configuration

Rule Name	Priority	Conditions	Application Name
P-DRA Gx	1	AppId Equal 16777238 - 3GPP Gx	PDRA

Rule Name	Priority	Conditions	Application Name
P-DRA Gx-Prime	1	AppId Equal 16777238 - 3GPP Gx-Prime	PDRA
P-DRA Rx	1	AppId Equal 16777236 - 3GPP Rx	PDRA
P-DRA S9	1	AppId Equal 16777267 - 3GPP S9	PDRA
P-DRA Gxx	1	AppId Equal 16777266 - 3GPP Gxx	PDRA

Answers

Diameter Answer messages can be routed to a DSR Application, or relayed by Diameter using Peer Routing Rules. [Table 6: Answer Processing by Policy DRA](#) lists all of the supported Answer messages and indicates which ones are processed by Policy DRA under what conditions.

If the Policy DRA application has requested to receive an Answer, but the Policy DRA application has Operational Status of Unavailable, the Diameter Routing Function will relay the Answer message directly to the remote Peer.

Note: Relaying an Answer while the Policy DRA application is Unavailable might result in exposing a PCRF name that was supposed to be topology hidden. This is because Diameter Routing does not support configuration of whether to relay or discard Answers when a DSR Application requested receipt of the Answer, but became Unavailable before the Answer was received.

Table 6: Answer Processing by Policy DRA

Answer	Requested by Policy DRA	Condition
CCA-I	Always	N/A
CCA-U	Conditional	CCA-U is processed by Policy DRA only if Policy DRA is configured to update the Session Last Touched Time on CCA-U in addition to RAA.
CCA-T	Always	N/A
RAA (Gx)	Always	N/A
AAA	Always	N/A
ASA	Never	N/A
RAA (Rx)	Always	N/A
STA	Always	N/A

Egress Routing

This section describes how Diameter Request messages are routed from Policy DRA. Diameter Request messages are routed from the Diameter Routing Function to the Policy DRA based on the configured Application Routing Rule, and routed from the Policy DRA to the Diameter Routing Function. The Policy DRA will return the Request to the Diameter Routing Function for PRT processing and routing.

When the Policy DRA 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 to either route the Answer to the Policy DRA or route the Answer to the downstream Peer without involving the Policy DRA. After the Policy DRA 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. 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

When a binding capable session initiation message (CCR-I) arrives for an IMSI that is not already bound to a PCRF, the Policy DRA application selects a PCRF from the list of adjacent PCRFs that are configured using the **Policy DRA -> Configuration -> PCRFs** GUI page. This list of PCRFs generally contains only PCRFs that are local to the Site with the Policy DRA node. PCRFs that are local to the Policy DRA node's mate are generally not 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 (Gx, Rx, Gxx, Gx-Prime, S9), only the binding capable Hostnames are configured in the **Policy DRA -> Configuration -> PCRFs** GUI.

Policy DRA distributes new bindings across the set of configured PCRFs. The distribution occurs independently on each DA-MP server; predicting the next PCRF that will be used is difficult on a Policy DRA node that has Policy Client connections to multiple DA-MP servers. In addition, the distribution is determined 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.

It is also possible to support more complicated PCRF selection by pushing the PCRF selection into Diameter Routing and out of the Policy DRA application. This can be accomplished by configuring a separate Peer Routing Table to be used for new binding creations using the **Policy DRA -> Configuration -> Site Options** GUI. The Peer Routing Rules can be configured to cause selection of different Route Lists. In this mode, Policy DRA can support weighted PCRF selection and different PCRF pools based on the origin of the Request.

The Policy DRA APN-based PCRF Pool Selection function 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.

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 creating two Peer Routing Rules 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

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 (CCA-I or AAA). This PCRF name is 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 Policy 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 number of Policy 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 include the Destination-Host AVP with the Policy DRA Diameter hostname. In order to support such Policy Clients, Policy DRA 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. 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. Refer to Policy DRA online help for Topology Hiding Scope options. If the PCRF name is hidden from the Policy Client, 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, Policy DRA 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 requires increased stack event processing and increased latency to look up the bound PCRF in the Session record. For these reasons, Topology Hiding should have the narrowest possible Scope. 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 (that is not proxy-compatible).

Naming Conventions and Hierarchical Routing

When Policy DRA is deployed in large networks with multiple Policy DRA mated pairs, the Diameter Routing configuration can be greatly simplified by employing some simple naming conventions. For example, naming all Policy Clients and PCRFs local to a particular Policy DRA node such that they start with a common prefix allows Peer Routing Rules like “Destination-Host Starts-With xxx”, where xxx is the Site prefix for that Policy DRA node. The “Starts-With” rule will point to a Route List that routes to the Policy DRA node where the equipment is located. Then if a new Policy Client or PCRF is added at a given Policy DRA node, routing changes are needed only at that node and that node’s mate, which have Peer Node entries and Diameter connections (that is, are adjacent) to the new Policy Client or PCRF. Policy DRA nodes that are non-adjacent do not require any routing updates.

Chapter 3

Policy DRA Overview

Topics:

- [PCRF Pools and Sub-Pools Concepts and Terminology.....40](#)
- [Binding-capable Sessions.....51](#)
- [Binding-dependent Sessions.....56](#)
- [In-session Message Processing.....58](#)
- [Topology Hiding.....58](#)
- [Session Integrity.....59](#)
- [Policy Data Auditing.....62](#)
- [Policy DRA Assumptions and Limitations.....67](#)
- [Policy DRA Capacity Constraints.....69](#)

This section gives an overview of Policy DRA, 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 DRA Configuration](#).

Note: If you are upgrading from a Policy DRA release prior to 5.1 and Policy DRA was activated on that release, see [Policy DRA PCRF Pooling Upgrade](#). The differences between Policy DRA prior to DSR 5.1 and Policy DRA in 5.1 and later releases are described there.

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 DRA Configuration](#)
- [Policy DRA PCRF Pooling Upgrade](#)

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 7: 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 7: Relationship between APNs and PCRF Pools

Figure 8: 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 8: Relationship between IMSIs and PCRF Pools

Figure 9: Multiple PCRF Pools illustrates multiple PCRF pools, each supporting a different service. In this example, PCRF pool 1 might be dedicated to policy control over the usage of enterprise data services and PCRF 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 PCRF pools.




Figure 9: 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 10: 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 10: 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 11: PCRF Pools and Sub-Pools Routing Scenarios shows example routing scenarios using PCRF Pools and Sub-Pools.




Figure 11: 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 DRA > 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 7: Policy DRA Terminology](#) shows a list of some Policy DRA terms and their meanings as they apply to this document.

Table 7: 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. A binding ceases to exist when the last Diameter

Term	Meaning
	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 Binding Slave for that binding. There may be

Term	Meaning
	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	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. 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.
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 session initiation request is considered to be a

Term	Meaning
	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	The suspect binding mechanism allows a binding to be removed if the PCRF that the subscriber is bound to becomes unreachable. A binding is marked suspect if after being successfully established, a subsequent binding-capable session initiation request for that same binding receives a 3002 response (unable to route) from the routing layer. If another binding-capable session initiation request for the binding arrives after the suspect binding interval and also receives a 3002 response, the suspect binding is removed, allowing the next request to be routed to another PCRF.

Binding-capable Sessions

A binding is a relationship stored in the pSBR-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 pSBR-S 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 (11345) 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 (11344) 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 (11340) 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 (11341) 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 PsbrFinalBindingsFollowed (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 application 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 application 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 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.

- 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 (11342) 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 (11343) is incremented by one.

Note: When the Policy 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 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

There following configuration data 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 Policy DRA 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 12: Find or Create a Binding shows the logic used for this task.




Figure 12: 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 DRA -> Configuration -> 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 Policy DRA Database.](#)

Binding-dependent Sessions

A binding is a relationship stored in the pSBR-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 pSBR-S that associate additional sessions with a binding.

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

Figure 13: Binding Dependent Session Initiation Request Processing Overview shows an overview of binding-dependent session initiation requests using IPv4 or IPv6 as correlation keys .




Figure 13: 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. See

Related Topics

- [The Policy 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 pSBR Session Database holds session information that is used for routing in-session messages. A given session record is accessible on every pSBR 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 DRA 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 8: 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 8: Topology Hiding Scope Configuration

Topology Hiding System Setting	Topology Hiding Scope Setting	Result
Disabled	N/A	No Topology Hiding is performed
Enabled	Specific Hosts	Topology Hiding is performed for messages destined for the Policy Clients only if the Policy Clients' FQDNs are configured for Topology Hiding
	All Foreign Realms	Topology Hiding is performed for messages destined for the Policy Clients if the realms of the Policy Clients are different from the Realm of the PCRF that sends the messages
	All Foreign Realms + Specific Hosts	Superset of All Foreign Realms and Specific Hosts options
	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:

1. Session Audit Premature Removal of Sessions

Policy DRA uses the mechanism of the Session Audit (see [Binding and Session Database 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 [Binding and Session Database 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 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.

2. 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 9: 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 9: 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 9: Policy DRA Error Scenarios for Session Integrity

Error Scenario	Policy DRA Behavior
----------------	---------------------

Failed to create IMSI Anchor Key for new binding	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. The Policy Client will attempt to re-establish the binding capable session.
Failed to create binding capable session	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). 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.
Failed to create an alternate key	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). If Alternate Keys are not created, a binding dependent session that needs to use those keys will fail.
Failed to update a new binding with the answering PCRF	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). 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.

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 10: 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 10: 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 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.

Policy Data Auditing

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.

- Policy 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 Policy 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 Policy SBR server is. Audit has no impact on the engineered rate of signaling.

Table 11: Effects of Stale Binding and Session Records describes the possible effects of a stale record, according to the type of stale record.

Table 11: Effects of Stale Binding and Session Records

Record Type	Effect of a Stale Record
IMSI	A stale IMSI anchor key record will cause all sessions for that IMSI to be handled by whatever PCRF was assigned to the IMSI when the IMSI anchor key record was created. Having one or more subscribers tied indefinitely to a given PCRF may hinder the Policy DRA load distribution algorithm from keeping PCRFs evenly loaded.
MSISDN	A stale MSISDN alternate key record will cause all binding dependent sessions that rely on the MSISDN for subscriber identity to be routed to the PCRF that was assigned to the MSISDN record when it was created. If some binding capable sessions include MSISDN and some do not, it is possible that a subscriber's sessions could be routed to two different PCRFs (one used by the stale MSISDN record, and one used by a new IMSI anchor key record). This situation is expected to be rare and is further mitigated by the Policy DRA software overwriting MSISDN records if the same MSISDN is later assigned to a different PCRF.
IPv4	A stale IPv4 alternate key record could result in mis-routing of a Diameter session containing an IP address that was reassigned (such as by a DHCP server) to another subscriber. This situation is mitigated by the Policy DRA software overwriting IPv4 Address records if the same IP address is later assigned to a different PCRF.
IPv6	A stale IPv6 Alternate key record could result in mis-routing of a Diameter session containing an IP address that was reassigned (such as by a DHCP server) to another subscriber. This situation is mitigated by the Policy DRA software overwriting IPv6 Address records if the same IP address is later assigned to a different PCRF.
Session	The main problem caused by a stale Session record is that removal of a binding capable session is what triggers removal of Binding records. If a stale Session record exists, the associated and correspondingly stale Binding records probably also exist.

Record Type	Effect of a Stale Record
SessionRef	SessionRef records are written in lock-step with Session records such that if one fails, the other is removed. As a result, if a SessionRef record is stale, the corresponding Session record is probably also stale. Because Binding database records contain SessionRef instances, a stale SessionRef record prevents those Binding records from being removed when they should be. Binding records are removed only if they are not associated with a valid SessionRef.

Binding table audits are confined to confirming with the Session Policy 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.

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:

- The default duration can be configured using the **NOAM Policy DRA > Configuration > Network-Wide Options** GUI page.
- A session lifetime can be configured per Access Point Name using the **NOAM Policy DRA > Configuration > Access Point Names** GUI page.

If the session initiating message (CCR-I) contains a Called-Station-Id AVP (an Access Point Name) and the Access Point Name is configured in the Access Point Names GUI, the session will use the value associated with that Access Point Name for the session lifetime value. If the session initiating message contains no Called-Station-Id Access Point Name, or contains a Called-Station-Id Access Point Name that is not configured in the Access Point Names GUI, the default session lifetime from Network-Wide Options will be used.

If the audit discovers a session record for which the current time minus the last touched time (either when the session was created, or when the last RAA was processed, whichever is more recent) exceeds the applicable session lifetime, the record is considered to be stale. Stale records are scheduled for Policy DRA initiated RAR messages to query the policy client that created the session to ask if the session is still valid.

Generally, Policy 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.

The maximum audit rate is configurable (with a default of 12,000) so that the audit maximum rate can be tuned according to the customer's traffic levels. For example, if the Policy SBR servers are using only 50% capacity for signaling, a higher rate could be made available to audit.

If the Policy SBR signaling load plus the audit load cause a Policy SBR server to exceed 100% capacity, that Policy SBR server will report congestion, which will cause an automatic suspension of auditing. Audit will continue to be suspended until no Policy SBR server is reporting congestion. Any Policy SBR on which audit is suspended will have minor alarm 22715 to report the suspension. The alarm is cleared only when congestion abates.

A Policy SBR server determines that it is in congestion by examining the rate of incoming stack events.

- Local congestion refers to congestion at the Policy SBR server that is walking through Binding or Session table records.

- Remote congestion refers to congestion at one of the Session Policy SBR servers that a Binding Policy SBR server is querying for the existence of session data (using sessionRef).

A Binding Policy SBR server will suspend audit processing if the server on which it is running is congested (local congestion), or if any of the Session Policy SBR servers to which it is connected through ComAgent connections have reported congestion (remote congestion). Audit processing will remain suspended until both local congestion and all instances of remote congestion have abated.

A Session Policy SBR server will suspend audit processing if the server on which it is running is congested (local congestion). The Session Policy SBR does not have to worry about remote congestion because it does not rely on binding data to perform its auditing function. Recall that session records are removed by audit if they are determined to be stale and the policy client that created the session indicates that the session is no longer needed (or if the session integrity feature has exhausted all attempts to communicate with a policy client that created a session). Session auditing will remain suspended until the local congestion abates.

When a Policy SBR server starts up (i.e. Policy SBR process starts), or when a Policy 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.

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 a Policy 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 "Policy 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. The audit reports for each table type are formatted as described in [Table 12: Audit Report Formats](#).

Table 12: Audit Report Formats

Data Type	Audit Report Format
IMSI Binding Records	Policy SBR Table Audit Pass Statistics Table: ImsiAnchorKey Records Visited: N Session References Audited: N Session References Removed: N Records Removed: N Audit Pass Duration: N seconds Suspended Duration: N seconds
IPv4 Alternate Key Binding Records	Policy SBR Table Audit Pass Statistics Table: Ipv4AlternateKey Records Visited: N Records Removed: N Audit Pass Duration: N seconds

	Suspended Duration: N seconds
IPv6 Alternate Key Binding Records	Policy SBR Table Audit Pass Statistics Table: Ipv6AlternateKey Records Visited: N Records Removed: N Audit Pass Duration: N seconds Suspended Duration: N seconds
MSISDN Alternate Key Binding Records	Policy SBR Table Audit Pass Statistics Table: MsisdnAlternateKey Records Visited: N Session References Audited: N Session References Removed: N Records Removed: N Audit Pass Duration: N seconds Suspended Duration: N seconds
Sessions Records	Policy SBR Table Audit Pass Statistics Table: Session Records Visited: N Records Requiring a Policy Client Query: N Records Removed due to Policy Client Query Results: N Stale Binding Dependent Records Removed: N Audit Pass Duration: N seconds Suspended Duration: N seconds
Session Reference Records	Policy SBR Table Audit Pass Statistics Table: SessionRef Records Visited: N Records Removed: N Audit Pass Duration: N seconds Suspended Duration: N seconds
pSBR	Policy SBR Audit Statistics Report Number of Query PendingRar records: 0 Number of Release PendingRar records: 0 Number of Query PendingRar records added : 0

	Number of Release PendingRar records added: 0 Number of Query PendingRar records deleted: 0 Number of Release PendingRar records deleted: 0 Number of Query PendingRar records not added due to max capacity: 0 Number of Release PendingRar records not added due to max capacity: 0 Average number of times Rar attempted before Query PendingRar records deleted: 0.00 Average number of times Rar attempted before Release PendingRar records deleted : 0.00 Maximum number of times Rar attempted before Query PendingRar records deleted: 0 Maximum number of times Rar attempted before Release PendingRar records deleted: 0
--	--

Policy DRA Assumptions and Limitations

Policy DRA has the following assumptions and limitations:

Assumptions

- The Anchor Key that identifies all subscribers in the Policy DRA network is the IMSI.
- 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 names and Policy Client names start with characters that can be used to identify which Policy DRA DSR hosts the primary connection to that equipment. This greatly simplifies routing configuration for the Policy DRA network. The network can be configured to work without such a naming convention, but routing setup and maintenance will be unnecessarily complex.
- The Policy DRA Gx-Prime interface support feature is backward compatible and functions whether or not PCRF pooling is available.

Limitations

Note: These limitations are not necessarily specific to PCRF Pooling.

- Policy DRA does not support adding, changing, or removing IP addresses using binding-capable update messages (for example, CCR-U/CCA-U). Alternate keys are currently created on receipt of CCA-I for keys present in the corresponding CCR-I.
- All binding-capable session initiation requests without an APN (for example, Called-Station-Id AVP), or with an APN that is not configured in P-DRA, received with PCRF Pooling enabled and after the binding migration period are rejected.

Note: This condition is not a typical limitation, but it clarifies 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 bind IPv4 address a.b.c.d to PCRF B.

- The following known error conditions exist 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 pSBR 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 Policy DRA 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.
- In an early binding scenario, if the master session is established and torn down successfully in rapid succession (for example, in the interval between slave polling attempts, typically about 250 ms), the slave sessions will also fail. This occurs because the master sessionRef is gone when the polling occurs. This condition should be rare because the policy client should generally not set up sessions lasting less than 1 second.
- Policy DRA PCRF selections can be overridden by DSR routing configuration. When PCRF selections are overridden by DSR, weighted load distribution can also be used.
- 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. Policy DRA has no mechanism for identifying members of a quota pool such that their sessions could all be routed to the same PCRF.
- Policy data Binding records are guaranteed to survive only a single site failure.
- Policy DRA does not support the 3GPP mechanism to redirect Policy Clients to a PCRF.
- Policy DRA does not support growth of Policy SBR resources. A Policy DRA system can be configured at activation time to be as small as 3 servers, or as large as 8 Policy DRA mated pairs of 3 enclosures each, but once the number of Policy SBR(B) Server Groups per network and the number of Policy SBR(S) Server Groups per mated pair is chosen at feature activation time, neither growth nor de-growth is supported without first deactivating the feature. Feature deactivation requires a total network-wide outage for Policy DRA. (Additional mated pairs can, however, be added to grow the Policy DRA network provided that the new mated pairs have the same number of session Policy SBR Server Groups as the existing mated pairs.)
- Policy DRA 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.

- Policy DRA evenly distributes new sessions across the Policy 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.
- If a Gx-Prime session is detected as being stale, it will be removed immediately by the pSBR audit function. It will not be subject to MITM handling like Gx sessions are:
 - Consistent session audit handling behavior as for other binding dependent interface such as Rx.
 - There should be no strong case for Gx-Prime sessions to be topology hidden for DPI clients. Additionally, MITM handling for Gx-Prime sessions is not needed because the session states will not be stored if no topology hiding.

Note: Gx-Prime sessions may still have Topology Hiding and may be audited in the same way Gx sessions are.

Policy DRA Capacity Constraints

Policy DRA has the following engineered capacity constraints:

Constraint Name	Value
Maximum managed objects per Network	
Local PCRS 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	1-8 (Configurable during Feature Activation)
Session Server Groups per Mated Pair	1-8 (Configurable during Feature Activation)
Binding Servers per Sg	8
Session Servers per Sg	8
Cardinality relationships between managed objects	
Sub-Pool Rules per PCRF Pool	10

Chapter 4

Policy DRA Deployment

Topics:

- *High Level Deployment Description.....71*
- *Deployment Topology.....71*
- *Redundancy.....85*
- *Policy DRA Scalability.....89*
- *IP Networking.....93*

Policy DRA can be deployed in customer networks to solve Policy Diameter signaling routing issues.

A Policy DRA DSR consists of a number of Policy DRA DA-MPs, a number of Policy SBRs, OAM servers, and IPFE servers.

High Level Deployment Description

A Policy DRA DSR consists of a number of Policy DRA DA-MPs, a number of Policy SBRs, OAM servers, and optional IPFE servers.

The Policy DRA DA-MPs are responsible for handling Diameter signaling the Policy DRA application.

Policy SBRs are special purpose MP blades that provide an off-board database for use by the Policy DRA feature hosted on the Policy DRA DA-MPs. Policy SBRs host the Policy session and Policy Binding databases.

Each Policy DRA DSR hosts connections from Policy Clients and PCRFs. Policy Clients are devices that request authorization for access to network resources on behalf of user equipment (such as mobile phones) from the PCRF. 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 are routed to the same PCRF.

Policy DRA DSRs can be deployed in mated pairs such that Policy session state is not lost even if an entire Policy DRA DSR fails or becomes inaccessible. When Policy DRA mated pairs are deployed, Policy Clients and PCRFs are typically cross-connected such that both Policy DRA DSRs have connections to all Policy Clients and all PCRFs at both mated sites.

“Policy DRA network” is the term used to describe a set of Policy DRA mated pairs and NOAM server pair. All Policy Clients and PCRFs are reachable for Diameter signaling from any Policy DRA DSR in the Policy DRA network.

Deployment Topology

This section describes the makeup of a Policy DRA network, regardless of its size. [Figure 14: Sites, Mated Pairs, and Region](#) illustrates an example Policy DRA network.

- A Policy DRA Network can have up to 8 mated pairs or 16 sites, or can be as small as a single site.
- The Policy DRA Binding Region provides the scope of the Policy Binding database. There is one instance of the Binding database in the entire Policy DRA network. Binding records are accessible from every Policy DRA DSR in the Region.

The Binding database need not be confined to a single mated pair, but can be deployed across multiple Policy DRA DSRs. All Policy Binding Server Groups must be deployed before the Policy DRA network can be used.

- Mated Pair provides the scope for an instance of the Policy Session database.

There is one instance of the Session database per Policy DRA Mated Pair. Session records are accessible from each Policy DRA DSR in the Mated Pair.

- Policy Clients and PCRFs have primary connections to their local Policy DRA and secondary connections to the mate of their local Policy DRA.
- Policy DRA DSRs are connected to each other on the External Signaling Network. Each Policy DRA Site must be reachable from every other Policy DRA 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 14: Sites, Mated Pairs, and Region

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

If the deployment includes more than one mated pair, all mated pairs that host the Binding database must be deployed before the Policy DRA network can be functional. Subsequent mated pairs can be deployed as needed, but will host only instances of the Session database.

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 15: Policy DRA in Roaming Scenarios](#) illustrates an example Diameter network where the Policy DRAs are located in Home Access and Visited Access networks.




Figure 15: 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.

Policy DRA Configurable Components

Figure 16: Policy DRA Component Relationships illustrates the relationships between the following key Policy DRA configurable components:

- Policy DRA Binding Region - consisting of all Policy DRA Sites

Policy DRA Deployment

- Policy DRA Mated Pairs - consisting of pairs of Policy DRA Sites
- Policy DRA Sites
- Policy Session Resource Domains – one per Policy DRA Mated Site consisting of all Session Policy SBR Server Groups at the mated pair
- Policy Binding Resource Domain – one per Policy DRA Binding Region consisting of all Binding Policy SBR Server Groups
- Policy DRA Resource Domains – one per Policy DRA Mated Site consisting of all DSR (multi-active cluster) Server Groups at the mated pair.
- Policy SBR Server Groups – enough to handle the load in Stack Events per second
- Diameter Signaling Router (multi-active cluster) Server Groups – one per Policy DRA Site




Figure 16: Policy DRA Component Relationships


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

- One mated pair that hosts the Policy Binding database and an instance of the Policy Session database
- N mated pairs that each host only an instance of the Policy Session database

Figure 17: Example Policy DRA Mated Pair - Hosting Binding Policy SBRs illustrates two Policy DRA DSR Sites configured as a Mated Pair that is hosting the Binding database:

- This Mated Pair hosts the Policy Binding database and an instance of the Policy Session database.
- The Policy Binding database is represented by a Policy Binding Resource Domain consisting of a number of Policy SBR Server Groups.
- The Policy Session database instance is represented by a Policy Session Resource Domain consisting of a number of Policy SBR Server Groups.
- Each Policy SBR Server Group consists of 3 servers using the Active/Standby/Spare redundancy model, allowing for Site redundancy.
- The number of Policy 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.
- Each Site has an SOAM Server Group consisting of 3 servers using the Active/Standby/Spare redundancy model, allowing for Site redundancy.
- The Policy DRA 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 blades – one for use by Policy Clients and one for use by PCRFs. (IPFE is not required.)

Policy DRA Deployment




There is 1 Type 1 P-DRA Mated Pair per P-DRA Network
 SOAM SGs are Act/Sby/Sp
 Policy SBR SGs are Act/Sby/Sp
 DSR SGs are N:K Act – 1 SG per Site
 IPFE SGs are Act – 1 SG per IPFE server

Figure 17: Example Policy DRA Mated Pair - Hosting Binding Policy SBRs

Figure 18: Example Policy DRA 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.



There are 5 Type 2 P-DRA Mated Pairs per P-DRA Network
 SOAM SGs are Act/Sby/Sp
 Policy SBR SGs are Act/Sby/Sp
 DSR SGs are N:K Act – 1 SG per Site
 IPFE SGs are Act – 1 SG per IPFE server

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

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

- Each Policy DRA DSR Site has a set of Policy Clients whose primary connection is directed to that Policy DRA.
- Each Policy DRA DSR Site has a set of PCRFs to which it distributes new bindings. Each PCRF at this Site has a primary connection to the Policy DRA DSR at that Site.
- Each Policy Client should have a secondary connection to the mate of the Policy DRA DSR for which it has a primary connection. (Without this “cross-connect”, Policy DRA 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 Policy DRA DSR for which it has a primary connection. (Without this “cross-connect”, Policy DRA site failure would leave the PCRF inaccessible.)
- Each Mated Pair of Policy DRA DSRs shares an instance of the Policy Session database.
- All Policy DRA 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 Policy DRA DSR has no connection, the message must be routed to a Policy DRA DSR that does have a connection. This routing is configured using the DSR routing tables.

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




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 Policy DRA application defines two Place Association Types:

- Policy DRA Binding Region

As illustrated in [Figure 14: Sites, Mated Pairs, and Region](#), the Policy DRA application defines a Region to include all Sites that are part of the Policy DRA network. This provides a scope for the Binding database, which is accessible to all Policy DRA Sites in the Policy DRA network.

- Policy DRA Mated Pair

As illustrated in [Figure 14: Sites, Mated Pairs, and Region](#), pairs of Policy DRA Sites are grouped together as Mated Pairs. Each Place Association with Type of Policy DRA Mated Pair includes exactly 2 sites. A Policy DRA Mated Pair has the following attributes:

- Hosts an instance of the Policy DRA Session database
- Hosts Policy Client Diameter connections for Policy Clients at both Sites in the Mated Pair
- Hosts PCRF Diameter connections for PCRFs at both Sites in the Mated Pair

Server Groups

The Policy DRA application makes use of several different types of Server Groups, as defined in [Table 13: Server Group Functions](#).

Table 13: Server Group Functions

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

- Policy SBR Type

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

Each Policy SBR Server Group consists of one, two, or three servers, depending on the type of deployment. [Table 14: Policy SBR Server Group Configuration and Data Redundance](#) describes the supported configurations for Policy SBR Server Groups. See [Redundancy](#) for details on Policy data redundancy.

Table 14: Policy 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 Policy SBR Server Groups in the mated pair, meaning that ~50% of the Session accesses will be routed across the WAN.

# of Servers	Redundancy	Typical Use
		Note: PSBR Server Groups must be configured with two WAN replication channels.

Because only the active server in a Policy SBR Server Group is actually processing Stack Events, a Policy 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 Policy DRA, 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 Policy Clients and PCRFs connected to every DA-MP.

The DA-MPs in the Server Group will be treated as a cluster of active servers. There should be at least two DA- MPs in the Server Group in order to support in-service maintenance or upgrade. The DA- MPs 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 Policy DRA is being deployed in mated pairs, the DA- MPs 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 Policy DRA application defines three Resource Domain Profiles: Policy Session, Policy Binding, and Policy DRA.

After Policy 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. A Policy SBR Server Group must be associated with either a Policy Session or Policy Binding Resource Domain, or with both Policy Session and Policy Binding Resource Domains. The latter configuration is expected to be used only for small deployments.

DA- MPs are configured in a single Server Group per Policy DRA DSR with a Server Group function type 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 Policy DRA Resource Domain. For a non-mated deployment, the DSR (multi-active cluster) Server Group must be in its own Policy DRA Resource Domain.

Figure 20: Resource Domains illustrates the possible relationships between a single Policy SBR Server Group and the Policy Resource Domains. Although not shown in the figure, each Resource Domain will probably contain a number of Server Groups.




Figure 20: Resource Domains

Policy Clients

Policy Clients act on behalf of the user equipment (UEs) to request Policy authorization and enforce Policy rules received from the PCRFs. Policy Clients send Policy requests to the Policy DRA, which ensures that the Policy request are sent to the PCRF in charge of Policy for the subscriber associated with the UE.

Policy DRA supports three different types of Policy Clients, referred to by 3GPP as AF, PCEF, BBERF, and DPI/MOS:

- The AF uses the Rx Diameter interface.
- The PCEF uses the Gx Diameter interface.
- The BBERF uses the Gxx Diameter interface.
- The DPI/MOS uses the Gx-Prime Diameter interface

How many connections a Policy Client might initiate towards the Policy DRA and how those connections are used are in customer control. The capabilities of the Policy Client, however, affect the functionality of the solution; as shown in [Table 15: Policy Client Connection Capability](#).

Table 15: Policy Client Connection Capability

Number of Connections Supported by Policy Client (per Diameter host)	Effect on Solution Capability
1	• Site Redundancy cannot be taken advantage of. • Diameter signaling throughput is limited to the capacity of the connection. • Extra latency to reconnect in the event of a connection drop.

Number of Connections Supported by Policy Client (per Diameter host)	Effect on Solution Capability
2	• Site Redundancy supported if secondary connection is configured to connect to Policy DRA mate site. • If both connections go to a single site and the Policy Client has the capability to use both connections simultaneously, Diameter signaling throughput may be doubled vs. only one connection. This configuration requires multiple Diameter connections to a single Diameter host – something that is not supported by RFC 3588, but which many vendors support to allow capacity beyond what a single connection can support. • 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.
>2	There are many scenarios possible, depending on the capabilities of the Policy Client. For example, there might be two connections to the primary Policy DRA (for capacity) and two to the mate Policy DRA (for Site redundancy).

Any Diameter Request can be sent to either Policy DRA 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 Policy DRA 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 Policy DRA 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, however, all Policy sessions for that subscriber are routed to the assigned PCRF.

Policy DRA 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 Policy DRA. 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 Policy DRA.
 - Each hostname has a different connection for each different interface type. This type of PCRF is supported by Policy DRA.

- 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 Policy DRA.
- 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 Policy DRA, 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 Policy DRA, 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 Policy DRA.

IPFE

In order to simplify network connectivity, Policy DRA will typically be deployed with one or two pairs of IPFEs per Policy DRA 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 Policy DRA in which the PCRFs are not capable of initiating connections to the Policy DRA. For example:
 - A Policy DRA DSR Site with a pair of IPFE blades, 8 DA-MP blades, and some Policy SBR blades
 - Four Policy Clients connected to two IPFE TSAs, with primary connections and secondary connections
 - The DA-MP blades are split into two groups that host connections to TSA1 and TSA2 respectively. This is necessary to ensure that a Policy Client's primary and secondary connections do not end up being connected to the same DA-MP.
 - One IPFE blade is primary for TSA1 and standby for TSA2; the other IPFE blade is primary for TSA2 and standby for TSA1.
 - Policy DRA MPs-to-PCRFs connectivity need not be fully meshed.
- An IPFE configuration in which Policy Clients are connected to a Policy DRA mated pair, but PCRFs are not capable of initiating connections to the Policy DRA. Each Policy Client has a primary connection to one Policy DRA site and a secondary connection to the mate site. For example:
 - Two Policy DRA DSR sites, each with a pair of IPFE blades and 4 DA-MP blades
 - Three Policy Clients with a primary connection to Policy DRA DSR Site 1 and secondary connections to Policy DRA DSR Site 2.
 - Three Policy Clients with a primary connection to Policy DRA DSR Site 2 and secondary connections to Policy DRA DSR Site 1.
 - Two PCRFs with primary connections to Policy DRA DSR Site1 and secondary connections to Policy DRA DSR Site 2.
 - Two PCRFs with primary connections to Policy DRA DSR Site2 and secondary connections to Policy DRA DSR Site 1.
 - One IPFE at Policy DRA DSR Site 1 is primary for TSA1. The other IPFE is standby for TSA1.

- One IPFE at Policy DRA DSR Site 2 is primary for TSA2. The other IPFE is standby for TSA2.
- A single site Policy DRA in which a single IPFE pair is used for both Policy Clients and PCRFs. The use of IPFE for PCRFs is possible only if the PCRF can be configured to initiate connections towards the Policy DRA. 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 blades, each blade supporting two TSAs
 - Four Policy Clients connect to TSA1 with their secondary connection going to TSA3, or vice-versa.
 - The PCRFs connect to TSA2 with their secondary connection going to TSA4, or vice-versa.
 - Six Policy DRA MP servers, each capable of hosting connections from Policy Clients and PCRFs
 - One IPFE blade is primary for TSA1 and TSA2, and standby for TSA3 and TSA4.
 - The other IPFE blade is primary for TSA3 and TSA4, and standby for TSA1 and TSA2.
- A single site Policy DRA in which IPFE is used for both Policy Clients and PCRFs. In this case, two pairs of IPFE blades are deployed in order to support high Diameter signaling bandwidth. For example:
 - Two pairs of IPFEs, each supporting a two TSAs
 - The Policy Clients connect to either TSA1 or TSA2, with their secondary connection going to the other TSA.
 - The PCRFs connect to either TSA3 or TSA4, with their secondary connection going to the other TSA.
 - Eight Policy DRA DA-MPs, each capable of hosting connections from Policy Clients and PCRFs
 - One IPFE blade on the Policy Client side is primary for TSA1 and standby for TSA2. The other IPFE blade is primary for TSA2 and standby for TSA1.
 - One IPFE blade on the PCRF side is primary for TSA3 and standby for TSA4. The other IPFE blade is primary for TSA4 and standby for TSA3.
- A Policy DRA mated pair configured with an IPFE for Policy Clients and a separate IPFE for PCRFs. The Policy Clients and PCRFs have a primary connection to their local Policy DRA DSR and a secondary connection to the mate Policy DRA DSR. For example:
 - Two Policy DRA DSR sites, each with a two pairs of IPFE blades and 6 DA-MP blades
 - Three Policy Clients with a primary connection to Policy DRA DSR Site 1 and secondary connections to Policy DRA DSR Site 2.
 - Three Policy Clients with a primary connection to Policy DRA DSR Site 2 and secondary connections to Policy DRA DSR Site 1.
 - Two PCRFs with primary connections to Policy DRA DSR Site1 and secondary connections to Policy DRA DSR Site 2.
 - Two PCRFs with primary connections to Policy DRA DSR Site2 and secondary connections to Policy DRA DSR Site 1.
 - One IPFE on the Policy Client side at Policy DRA DSR Site 1 is primary for TSA1. The other IPFE is standby for TSA1.
 - One IPFE on the Policy Client side at Policy DRA DSR Site 2 is primary for TSA3. The other IPFE is standby for TSA3.
 - One IPFE on the PCRF side at Policy DRA DSR Site 1 is primary for TSA2. The other IPFE is standby for TSA2.
 - One IPFE on the PCRF side at Policy DRA DSR Site 2 is primary for TSA4. The other IPFE is standby for TSA4.

Redundancy

Making the Policy DRA feature 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 Policy DRA feature 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 Policy SBR MPs:

- DA-MP Multi-Active Cluster

Policy DRA DA-MPs 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.

- Policy SBR Active Only

A Policy 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 Policy SBR database being hosted by the Active-only Server Group. In this configuration, the Policy SBR Server Groups consist of a single Server.

- Policy SBR Active/Standby

The Active/Standby redundancy model should be used for single site Policy DRA deployments, or for multi-site deployments when site redundancy is not important. In this configuration, the Policy SBR Server Groups consist of two servers. On system initialization, one of the two servers in each Policy 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.

- Policy 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 Policy 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.

- Policy SBR Active/Standby/Spare

The Active/Standby/Spare redundancy model should be used for Policy DRA mated pair deployments in which site redundancy is desired. In this configuration, each Policy 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.

Site Redundancy

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 Policy DRA this means no loss of Policy Binding or Policy Session data. In order to achieve site redundancy, the following configuration applies:

- Policy DRA is deployed on at least one mated pair of Policy DRA DSRs.
- Policy Clients and PCRFs are able to connect to both sites in the mated pair.
- Policy 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.
- DA-MPs are recommended to be engineered at 40% capacity across the mated pair.

Data Redundancy

The Policy Session and Policy Binding databases are partitioned such that each Server Group in a Policy Session or Policy 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 Policy SBR Server Group all maintain exact replicas of the data for the partition that the Server Group is responsible for. 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 21: Binding Table Partitioning Across Server Groups illustrates how a given Policy Binding table might be partitioned across four Policy SBR Server Groups in a Policy Binding Resource Domain.




Figure 21: Binding Table Partitioning Across Server Groups

Figure 22: Multi-Table Resources illustrates how each Policy 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 22: Multi-Table Resources

OAM Server Redundancy

The Policy DRA application can be deployed with varying degrees of redundancy on the NOAM and SOAM servers. Like the Policy 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 Policy DRA 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 Policy DRA 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 Policy 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 Policy SBR servers would have no parent OAM server and would not be able to report measurements, events, and alarms.

Policy SBR servers in a given Policy 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 23: 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 23: Data Merging - Normal Case

Figure 24: 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 Policy 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 24: Data merging - Redundant Site Failure

Policy DRA Scalability

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

For large systems, Policy DRA can scale up as follows:

- Eight mated pairs of Policy DRA DSRs (16 sites)
- Three enclosures per Policy DRA DSR site using half-height blades

Each enclosure has 16 half-height slots.

- Two pairs of IPFE blades per Policy DRA DSR
- Sixteen DA-MP blades per Policy DRA DSR

Figure 14: Sites, Mated Pairs, and Region illustrates a sample Policy DRA 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 Policy DRA DSR.
- All MP servers (both Policy SBRs and DA-MPs) are assigned to the Site where they are physically located.
- An instance of a Policy DRA Mated Pair (Place Association with type Policy DRA Mated Pair) is created for each pair of sites that are mates.
- A pre-determined number of Policy Binding Server Groups are created on the Policy DRA 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.
 - Policy Binding Server Groups can exist on more than 2 sites, but the Policy DRA network is not operational until all sites hosting Policy Binding Server Groups are operational.

- A Policy Binding Resource Domain is created including all Policy Binding Server Groups.
- A pre-determined number of Policy Session Server Groups are created at each mated pair.

Each Policy 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 Policy 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 DRA Resource Domain is created including the DSR Server Group at each of the mated Sites.
- A Policy DRA Binding Region (Place Association with type Policy DRA Binding Region) is created containing all Sites.

The Mated Pair of Policy DRA DSR sites illustrated in Figure 5 could support approximately 336,000 Diameter MPS with site redundancy (with DA-MPs engineered at 40%).

The single site Policy DRA DSR illustrated in Figure 6 could support approximately 384,000 Diameter MPS (with DA-MPs engineered at 80%).

MP Growth

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

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

Database Growth

The Policy DRA feature does not support growth of the Policy Session or Policy Binding databases after feature activation.

Note: The percentages of different types of Policy Diameter messages in the overall Policy Diameter traffic load is referred to as the call model.

This has the following implications:

- The number of Server Groups that will host the Policy Session database for each mated pair (or single site if no mated pair is planned) must be determined prior to feature activation.

The number of Policy Session Server Groups required depends on the expected Diameter traffic rate in MPS for Policy signaling and the ratio of Diameter MPS to Session stack events determined by the call model.

- The number of Policy Binding database Server Groups for the entire planned Policy DRA network must be determined prior to feature activation.

The number of Policy Binding Server Groups required depends on the number of Policy subscribers and the expected Diameter traffic rate in MPS for Policy signaling and the ratio of Diameter MPS to Binding stack events determined by the call model.

- After the number of Policy Binding and Policy Session Server Groups has been configured at Policy DRA feature activation time, these numbers cannot be changed without deactivating the feature.

Deactivation of the Policy DRA feature results in an outage for all Policy signaling that traverses all Policy DRA DSRs in the Policy DRA network.

Mated Pair Growth

A mate Policy DRA DSR can be added to a single-site Policy DRA DSR.

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

Adding a Mate Policy DRA DSR to an Existing Policy DRA DSR

Because Policy SBR growth is not supported, a Policy DRA DSR deployed without a mate must host all of the Policy SBR Server Groups that are planned for deployment across the mated pair when the mate is added. This requires planning ahead for the eventual mate.

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

A Policy DRA 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 Policy 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 Policy 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 Policy 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 Policy 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 Policy 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 Policy 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 Policy DRA DSRs

Policy DRA network capacity can be expanded by adding mated pairs of Policy DRA DSRs. Policy DRA mated pairs added after the Policy DRA network is up and running cannot include additional Policy Binding Policy SBR Servers.

The number of Policy Session Policy SBR Servers must be the same for each of the new Policy DRA mates, and must be determined at Policy DRA feature activation. Every Policy DRA mated pair must have the same number of Policy Session Policy SBR Server Groups. After the number is selected the value cannot change until a software upgrade becomes available that supports Policy SBR growth.

While Policy SBR growth (adding Policy SBR Server Groups) is not supported, Policy DRA 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 Policy DRA feature can scale down to run on a small hardware footprint. This section describes the smallest supported Policy DRA 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 Policy DRA 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 Policy SBR Server is also running in simplex mode and is configured to host both the Policy Binding and Policy Session databases. 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 Policy DRA DSR 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.
- The Policy SBR Server pair uses 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.

The smallest supported Mated Pair of Policy DRA DSRs, illustrated in [Figure 25: Smallest Supported Policy DRA 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 single combined Session and Binding Policy SBR triplet is deployed with two servers at Site 1 and one server at Site 2.
- Two DA-MPs are deployed at each site to support server redundancy at each site.




Figure 25: Smallest Supported Policy DRA Mated Pair

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 16: IP Traffic-to-Service Mapping](#) illustrates a possible Services configuration for enabling signaling traffic from OAM traffic. In [Table 16: 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 16: 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 Policy SBRs, and between Policy SBRs	ComAgent	Signaling VLAN 4	Signaling VLAN 4
Replication of data among DA-MPs	Replication_MP	Signaling VLAN 4	Signaling VLAN 4
Replication of data among Policy SBRs	Replication_MP	Signaling VLAN 4	Signaling VLAN 4

Traffic Type	Service Name	Intra-NE Network	Inter-NE Network
HA Heartbeating among Policy 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 Policy 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

Chapter 5

Policy DRA Configuration

Topics:

- [*Policy DRA Configuration Overview.....96*](#)
- [*Pre-Configuration Activities.....98*](#)
- [*Policy DRA Configuration on the NOAM and the SOAM.....104*](#)
- [*Post-Configuration Activities.....153*](#)

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

Policy DRA Configuration Overview

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

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

The DSR 3-tiered Operations, Administration, and Maintenance (OAM) topology is required for the Policy DRA 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 26: GUI Structure for 3-tiered DSR Topology with Policy DRA](#), network-wide Policy DRA configuration is performed on the NOAM.

- A pair or triplet of SOAM servers at each site running in active/standby, or active/standby/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 26: GUI Structure for 3-tiered DSR Topology with Policy DRA](#), site-specific configuration for Policy DRA is performed on the SOAM; some network-wide Policy DRA 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, Policy 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 26: GUI Structure for 3-tiered DSR Topology with Policy DRA

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 DRA 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 Policy SBR Server Groups
- Assignment of Policy SBR Server Groups to Policy Session and/or Policy Binding Resource Domains
- Assignment of DSR Multi-active Cluster Server Groups to Policy DRA Resource Domains
- Assignment of Site Places to Policy DRA Mated Sites Place Associations
- Assignment of Site Places to Policy DRA Binding Region Place Associations

Policy DRA-specific data configured at the NOAM include:

- Alarm Thresholds for:
 - Policy DRA Application Ingress Message Rate
 - Policy Session Database Capacity
 - Policy Binding Database Capacity
- Access Point Names (APN)
- Maximum Session Inactivity Time per APN
- PCRF Pools and PCRF Sub-Pool Selection Rules

Policy DRA-specific data configured at the SOAM include:

- PCRFs, PCRF Pools, PCRF Pool to PRT Mapping, PCRF Sub-Pool Selection Rules, and local to the site
- Binding Key Priority for the site
- Topology Hiding configuration for the site
- Error response configuration for the site

For more information, see [Policy DRA Capacity Constraints](#).

Pre-Configuration Activities

Before Policy DRA configuration can be performed, the following activities need to be performed in the system:

- Verify that the Policy DRA application is activated in the system. (This is usually performed as part of the installation or upgrade activities.)

Policy DRA 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 Policy DRA:

- 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 Policy Binding or Policy Session profile, unless the P-DRA feature is deactivated. Resource Domains that are part of Policy DRA profiles can be deleted when the Policy DRA application is activated.

- Gather component information that is required for Diameter and Policy DRA 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 Configuration components that are required for Policy DRA configuration. See *Diameter Configuration for Policy DRA*.

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 DRA feature must be activated.
3. Policy DRA must be configured.
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 Policy DRA is activated on a DSR that was upgraded to a release that supports PCRF Pooling and the Policy DRA activation occurs after the upgrade is completed and accepted, the considerations listed above apply to the initial install. Activation of Policy DRA on a network where the upgrade is not completed and accepted on all servers is prohibited by the activation script.

Diameter Configuration for Policy DRA

The Policy DRA application requires configuration of several Diameter Configuration components before the Policy DRA 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 and Mediation User Guide* to complete the configuration of the Diameter Configuration components for the system, including the following Diameter Configuration components for use with the Policy DRA application.

1. MP Profiles

Select **Diameter > DA-MPs > Profile Assignments**, and verify that the correct Session MP Profiles have been assigned for Policy DRA 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 Policy DRA 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**).

2. Application Ids

Use the **Diameter > Configuration > Application Ids [Insert]** page to define an Application Id for each Diameter interface that will be used by Policy DRA in the system.

Policy DRA supports the following values that can be selected in the **Application Id Value** pulldown list:

- 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.

Policy DRA 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 [Policy DRA Routing of Diameter Messages](#).

3. 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 Policy DRA:

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)

4. 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 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.

5. Local Nodes (Policy DRA DA-MPs)

Use the **Diameter > Configuration > Local Nodes [Insert]** page to configure the Policy DRA 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 Policy DRA DA-MPs in the system. (Policy DRA DA-MPS can also be Peer Nodes to each other at different sites.)

See [Policy DRA Routing of Diameter Messages](#) for details on routing of messages for Policy DRA.

7. Connections

Use the **Diameter > Configuration > Peer Nodes [Insert]** page to configure connections between the Policy DRA 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 Policy DRA Peers.

For priority-based initial CCR-I routing, configure N+1 Route Groups where N is the number of PCRFs in the system. The first N Route Groups contain one corresponding PCRF Peer Node in each one, and the last Route Group contains all PCRFs.

The goal is to setup a routing configuration such that if there is no route available to the suggested PCRF in an initial (binding capable) session Request, Diameter automatically sends the Request messages to any other available PCRF.

Define a Route Group for each PCRF; enter the **Route Group Name**, select the **Peer Node** name (PCRF name) and enter the **Provisioned Capacity** as 1.

Define a last Route Group for all PCRFs; enter the **Route Group Name**, then add a **Peer Node, Connection and Capacity** entry for every PCRF. Select the **Peer Node** (PCRF) and enter the **Provisioned Capacity** as 1 for each PCRF 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 PCRFs in the system.

All Route Lists must contain at least two Route Groups, one for a single PCRF and one for all PCRFs.

Assign **Priority** value 1 to each Route Group for a single PCRF; assign **Priority** value 2 to the Route Group containing all the PCRFs.

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 preference.

Peer Routing Rules can be added to the Default Peer Route Table (PRT) or to new Peer Route Tables.

See [Policy DRA Routing of Diameter Messages](#) for details on PRT routing of Policy DRA messages.

The routing configuration will ensure that whenever Policy DRA requests Diameter to route to a particular PCRF based on the PRT:

- If the PCRF is available, Diameter will route to it.
- If the PCRF is not available, Diameter will route the message to any other available PCRF.

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 Policy DRA application. Policy DRA must receive all Policy 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 Policy DRA.
- **Application Name - PDRA**
- **Gx-Prime**
- **Application Route table**

Policy DRA Routing of Diameter Messages

Policy DRA 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

Policy DRA 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.
2. 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.
3. Default Peer Route Table
If no match is found in the Application-Ids configuration, the Default Peer Route Table is used.
4. Destination-Host Routing
If no Peer Routing Rule matches in the Default Peer Route Table, Policy DRA 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

Policy DRA 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 DRA->Configuration->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 Policy DRA 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 Policy DRA 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.

Policy DRA Configuration on the NOAM and the SOAM

This section describes the **Policy DRA > Configuration** GUI pages on the NOAM and the SOAM.

Access Point Names

An Access Point Name (APN) is a unique Packet Data network identifier. The Policy DRA 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.

Note: Access Point Names are configurable only on Active NOAM servers, and are viewable on SOAM and NOAM servers.

The fields are described in [Access Point Names elements](#).

On the **Policy DRA > 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.
- Click the **Insert** button.

The **Policy DRA > 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 (200) already exists in the system, the **Policy DRA > 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 DRA > 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 DRA > 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 17: Access Point Names elements](#) describes the elements on the **Policy DRA > Configuration > Access Point Names** page.

Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 17: 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 alphabetic character or a digit. Default: N/A Range: 1 to 100
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)	The time value (in hours) after which a session is considered to be stale. A session is considered stale only if no RAR/RAA messages are received in longer than the configured time. This value is used for sessions associated with this Access Point Name. For sessions that are not associated with any configured Access Point Names, the Default Stale Session Timeout value configured on the NOAM Policy DRA > Configuration > Network-Wide Options page is used.	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
	If a session's age exceeds this value, that session is eligible to be audited out of the database.	
Last Updated	This read-only field displays a timestamp of the time the Access Point Name was created or last updated, whichever occurred most recently. 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. 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.	Format: List Range: N/A

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 DRA > Configuration > Access Point Names**.

The **Policy DRA > Configuration > Access Point Names** page appears.

2. Click **Insert**.

The **Policy DRA > Configuration > Access Point Names [Insert]** page appears.

3. Enter a unique Access Point Name in the Access Point Name **Value** field.

4. Select a PCRF Pool Name from the **PCRF Pool Name** dropdown menu. This field contains all the qualified PCRF Pools configured from **Policy DRA -> Configuration -> PCRF Pools**. A qualified PCRF Pool is non-retired and has not been marked as Sub-Pool.

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 DRA -> Configuration -> PCRF Sub-Pool Selection Rules** page.

5. 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 DRA > 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 Binding Key Query Tool.

6. Click:

- **OK** to save the new Access Point Name and return to the **Policy DRA > Configuration > Access Point Names** page.
- **Apply** to save the new Access Point Name and remain on this page.
- **Cancel** to return to the **Policy DRA > 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 (200) 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 DRA > Configuration > Access Point Names**.

The **Policy DRA Configuration Access Point Names** page appears.

2. Click **Edit**.

The **Policy DRA > Configuration > Access Point Names [Edit]** page appears.

3. 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.

4. 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 DRA > 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 Binding Key Query Tool.
5. Click:
 - **OK** to save the changes and return to the **Policy DRA > Configuration > Access Point Names** page.
 - **Apply** to save the edited Access Point Name and remain on this page.
 - **Cancel** to return to the **Policy DRA > 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 DRA > Configuration > Access Point Names**.
The **Policy DRA > 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.

PCRF Pools

Policy DRA continues to support a single pool of PCRFs at each P-DRA 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.

Note: This includes the ability to upgrade from previous releases of Policy DRA, but new binding sessions automatically default to PCRF Pooling.

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 application 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 application, 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 Policy DRA site-by-site basis. Policy DRAs 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 Policy DRA peers that are a part of the pool.

On the **Policy DRA > Configuration > 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 18: 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.

Concept	Before PCRf Pooling	After PCRf Pooling
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 Policy DRA network configured for PCRf Pooling. The upper third of the figure shows data that is configured with the Policy DRA 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 DRA GUI at each of two Policy DRA 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 Policy DRA site.

The bottom third of the figure shows the PCRfs logically grouped into PCRf Pools as defined by the network operator.




Figure 27: PCRf Pooling Data

Table 19: PCRf Pooling Configuration Summary describes each of the new PCRf Pooling configuration tables, including the order in which they should be configured.

Table 19: PCRF Pooling Configuration Summary

Configuration Order	GUI Page	Purpose
1	PCRF Pools	Define the names of the PCRF Pools and Sub-Pools that are needed for grouping PCRFs to handle policy signaling for the various APNs.
2	PCRF Pool to PRT Mapping	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. Note: Before this step can be performed, PRT tables must be defined in the Diameter folder.
3	PCRF Sub-Pool Selection Rules	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. 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.
4	Access Point Names	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

Configuration Order	GUI Page	Purpose
		creation of a new binding are routed using the PCRF Pool.

PCRF Pools elements

[Table 20: PCRF Pools elements](#) describes the elements on the **Policy DRA > Configuration > PCRF Pools** page.

Note: Data Input Notes apply to the Edit page; 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 DRA > Configuration > Access Point Names** and the PCRF Sub-Pool Selection Rules configured in **Policy DRA > Configuration > PCRF Sub-Pool Selection Rules**.

Table 20: 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 DRA -> 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 DRA > Configuration > PCRF Pools**.

The **Policy DRA > Configuration > PCRF Pools** page appears.

2. Click **Insert**.

The **Policy DRA > Configuration > 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 DRA -> Configuration -> 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 DRA > Configuration > PCRF Pools** page.
 - **Apply** to save the new PCRF Pool name and remain on this page.
 - **Cancel** to return to the **Policy DRA > Configuration > 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 (2500) 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 DRA > Configuration > PCRF Pools**.

The **Policy DRA > Configuration > 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 DRA > Configuration > 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 DRA > Configuration > PCRF Pools** page.
- **Apply** to save the change and remain on this page.
- **Cancel** to return to the **Policy DRA > Configuration > 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 DRA > Configuration > PCRF Pools**.

The **Policy DRA > Configuration > 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 DRA > Configuration > 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 Policy DRA application. For initial installs and upgrades from releases in which Policy DRA was not activated, the Default PCRF Pool is created and mapped to the Not Selected PRT.

Note: If upgrading to a release of Policy DRA that supports PCRF Pools or PCRF Sub-Pools, see [Policy DRA PCRF Pooling Upgrade](#).

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 NOAMP 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 21: PCRF Pools to PRT Mapping elements describes the elements on the **Policy DRA > Configuration > PCRF Pools to PRT Mapping** page.

Note: Data Input Notes apply to the Edit page; 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 DRA -> Configuration -> 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 21: 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 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.	Format: Text box; string of 1-32 alphanumeric characters, must contain at least one alpha character, must not start with a 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 DRA > Configuration > PCRF Pool to PRT Mapping**.

The **Policy DRA > Configuration > 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 (unless you want 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 DRA > Configuration > 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, Policy DRA 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 DRA > Configuration > PCRF Pool to PRT Mapping** page.
 - **Apply** to save the selection and remain on this page.
 - **Cancel** to return to the **Policy DRA > Configuration > 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 DRA -> Configuration -> PCRF Pools** page.

Pausing update applies to all rows in the table on the **Policy DRA > 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 DRA > Configuration > PCRF Pool to PRT Mapping**.
The **Policy DRA > Configuration > 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 22: PCRF Sub-Pool Selection Rules elements describes the elements on the **Policy DRA > Configuration > PCRF Sub-Pool Selection Rules** page.

Table 22: 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	Format: Textbox Range: Equals, Starts With, and Ends With.

Fields (* indicates required field)	Description	Data Input Notes
	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.	
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 [Insert]

Use this task to insert (create new) PCRF Sub-Pool Selection Rules.

1. On the Active NOAM, select **Policy DRA > Configuration > PCRF Sub-Pool Selection Rules**.

The **Policy DRA > Configuration > PCRF Sub-Pool Selection Rules** page appears.

2. Click **Insert**.

The **Policy DRA > Configuration > 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 DRA -> Configuration -> 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 DRA > Configuration > PCRF Pools** page.
- **Apply** to save the new PCRF Sub-Pool Selection Rule and remain on this page.
- **Cancel** to return to the **Policy DRA > Configuration > 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 DRA > Configuration > PCRF Sub-Pool Selection Rules**.
The **Policy DRA > Configuration > PCRF Sub-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 (unless you want 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 DRA > Configuration > 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 DRA > Configuration > PCRF Sub-Pool Selection Rules** page.

- **Apply** to save the change and remain on this page.
- **Cancel** to return to the **Policy DRA > Configuration > 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 DRA > Configuration > PCRF Sub-Pool Selection Rules**.

The **Policy DRA > Configuration > 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 DRA > Configuration > 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 DRA > Configuration > Network-Wide Options** page on an Active NOAM, the following **Network-Wide Options** can be configured:

- **General Options**
 - Indicate whether to relay or discard an Answer message when the Policy DRA is Unavailable (for answers).
 - 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.
- **Audit Options**
 - Change the **Default Stale Session Timeout** value to a value other than the default value in the field.

This setting is a length of time in hours after which a session is considered to be stale. A session is considered stale only if no RAR/RAA messages are received in a length of time 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 only if a session is not associated with a configured Access Point Name. For sessions that are associated with a configured Access Point Name, the **Stale Session Timeout** value configured for the Access Point Name is used.

- Change the **Maximum Audit Frequency** default value to a different number of records per second for auditing the Policy SBR database.
- **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).

The fields are described in [Network-Wide Options elements](#).

Network-Wide Options elements

[Table 23: Network-Wide Options elements](#) describes the elements on the **Policy DRA > Configuration > Network-Wide Options** page on the NOAM.

Table 23: Network-Wide Options elements

Fields (* indicates a required field)	Description	Data Input Notes
General Options		
Policy DRA Unavailable (for answers)	A choice to relay or discard an Answer message when Policy DRA is Unavailable.	Format: Radio buttons Range: Relay or Discard Default: Relay
Origin-Host and Origin-Realm for Policy DRA generated RAR messages	A radio button choice to control the selected option's Origin-Host and Origin-Realm use as the Origin-Host and Origin-Realm in the RAR messages built and sent by Policy DRA to Policy Clients.	Format: Radio buttons Range: Local Host or PCRF Default: Local Host
Enable PCRF Pooling	A checkbox choice to control the PCRF Pooling feature. Check the box to allow a subscriber's policy sessions to be routed to different PCRFs depending on the originating point of the Access Point Network. Note: If upgrading from an activated pre-5.1 release, check this box following acceptance of the upgrade in order to allow future upgrades.	Format: Checkbox Range: Yes (Checked) or No (Unchecked) Default: PCRF Pooling Enabled (checked) for initial installs; PCRF Pooling Disabled (Unchecked) for upgrades from activated pre-5.1 releases.
Audit Options		

Fields (* indicates a required field)	Description	Data Input Notes
* Default Stale Session Timeout	The time (in hours) after which a session is considered to be stale. A session is considered stale only if no RAR/RAA messages are received in longer than the 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 only if a session is not associated with a configured Access Point Name. For sessions that are associated with a configured Access Point Name, the Stale Session Timeout value configured for the Access Point Name is used.	Format: Text box Range: 1-2400 hours (1 hour to 100 days) Default: 168 hours (7 days)
* Maximum Audit Frequency	The maximum records per seconds for auditing the Policy SBR database.	Format: Text box Range: 1000-25000 Default: 12000
Early Binding Options		
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. Set a 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 might reduce latency. A high value does the opposite. Note: This values is used only when PCRF Pooling is enabled.	Format: Text box Range: 1000 to 25000 milliseconds Default: 200 milliseconds
Maximum Early Binding Lifetime	The maximum time that a binding is allowed to remain as an early binding. The suggested setting for this value is 100 to 200 milliseconds longer than the Diameter transaction timeout. This value prevents bindings from becoming stagnate 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 values is used only when PCRF Pooling is enabled.	Format: Text box Range: 500 to 15000 milliseconds Default: 2500 milliseconds

Note: Keep these consideration in mind when working with net-work options:

- If **Apply** is clicked and the **Enable PCRF Pooling** checkbox transitioned from unchecked to checked, a confirmation window with a checkbox is displayed containing the text: "IMPORTANT! Enabling PCRF Pooling causes all new bindings to be stored using both IMS and AN. All ANS. MUST be configured in **Policy DRA -> Configuration -> Access Point Names** or session initiation requests might be rejected. After confirming this change, it is not possible to return to the IMS Only binding storage. Check the checkbox and click OK to enable PCRF Pooling; otherwise, click **Cancel** to continue using IMS Only."
- If the confirmation dialog for enabling PCRF Pooling is confirmed by checking the checkbox and clicking **OK** and any server in the Policy DRA network is not upgraded to the PCRF Pooling release with upgrade Accepted, no data is committed, and an error Box is displayed.
- If PCRF Pooling has been Enabled successfully, and then PCRF Pooling is disabled (through BE), no data will be committed and an error message is logged.
- If an attempt is made to Enable PCRF Pooling via BE, no data will be committed and an error message is logged.
- If the confirmation dialog for enabling PCRF Pooling is confirmed by checking the checkbox and clicking **OK** and all data is valid, the data is written to the database and the **Network-Wide Options** page is redrawn with all fields populated as applied and the **Enable PCRF Pooling** field is 'disabled' and 'checked'.
- If **Apply** is clicked and no change was made to the **Enable PCRF Pooling** value and all data is valid, the data is written to the database and the **Network-Wide Options** page is redrawn with all fields populated as applied.

Viewing Network-Wide Options

Use this task to view configured Network-Wide Options on the NOAM.

Select **Policy DRA > Configuration > Network-Wide Options**.

The **Policy DRA > Configuration > 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:

- Policy DRA Unavailable (for answers)
 - Enable PCRF Pooling
 - Default Stale Session Timeout (in hours)
 - Origin-Host and Origin-Realm for Policy DRA generated RAR messages
1. Select **Policy DRA > Configuration > Network-Wide Options**.
The **Policy DRA Network-Wide Options** page appears.
 2. The **Relay** or **Discard** radio button setting is an engineered system value (uneditable) that controls the network option for the Policy DRA Unavailable (for answers) field.
 3. Enter a number in the Default Stale Session Timeout **Value** field.

4. Select the **Enable PCRF Pooling** 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.

5. Select the **Local Host** or **PCRF** radio button.

PCRF Pooling can only be enabled after all of the servers in the network have been successfully upgraded to the release supporting PCRF Pooling and the upgrade has been accepted on all servers. You cannot enable PCRF Pooling until this state has been achieved.

6. Set the **Early Binding Polling Interval**.

This sets 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. Ideally, set the 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 might reduce latency.

7. Click:

- **Apply** to save the changes and remain on this page.
- **Cancel** to discard changes and remain on the **Policy DRA > Configuration > Network-Wide Options** page.

If **Apply** is clicked and the following condition exists, an error message appears:

- The entered Default State Session Timeout value contains invalid characters, is out of the allowed range, or the field is empty.

Alarm Settings

Note: Alarm Settings are configurable only on Active NOAM servers, and are viewable on NOAM and SOAM servers.

On the **Policy DRA > Configuration > Alarm Settings** page on an SOAM, you 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 DRA > 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 DRA ingress messages rate exceeds the configured Alarm Threshold. The thresholds are based on the engineered system value for Ingress Message Capacity.

- Policy SBR Sessions Threshold Exceeded

The Policy 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 Policy SBR Server Groups per mated pair chosen during Policy DRA feature activation.

The Policy SBR Sessions Threshold Exceeded alarm is raised when number of concurrent Policy SBR sessions exceeds the configured threshold.

- Policy SBR Bindings Threshold Exceeded

The Policy 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 Policy SBR Server Groups specified at Policy DRA feature activation.

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 24: Alarm Settings elements](#) describes the elements on the **Policy DRA > 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
- Policy SBR Sessions Threshold Exceeded
- Policy 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 24: 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 DRA 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 DRA 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)
* Major Alarm Threshold (Percent)	The Policy DRA 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

Policy DRA 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 Policy DRA 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)
Policy SBR Sessions Threshold Exceeded		
* Alarm Name	This alarm is raised when the number of concurrent Policy 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 Policy 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 Policy SBR Sessions.	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)

Policy DRA Configuration

Elements (* indicates required field)	Description	Data Input Notes
* Minor Alarm Threshold (Percent)	The concurrent sessions threshold for this alarm to be raised as Minor. The threshold is a percentage of the Maximum Policy 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)
Policy 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)
* 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

Elements (* indicates required field)	Description	Data Input Notes
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 DRA > Configuration > Alarm Settings**.

The **Policy DRA > 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 DRA > Configuration > Alarm Settings**.

The **Policy DRA > 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 DRA > 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 DRA ingress message rate capacity at which an alarm is raised with Critical, Major, or Minor severity.
 - Set the percentage of the Policy DRA 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 DRA 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 25: Congestion Options elements](#) describes the elements on the **Policy DRA > 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 25: Congestion Options elements

Fields (* indicates required field)	Description	Data Input Notes
Alarm Thresholds		
Alarm Name	The Policy DRA Server in Congestion alarm is raised when average Policy DRA ingress request messages rate exceeds the configured threshold. The thresholds are based on the engineered system value for Ingress Policy DRA Request Message Capacity.	Format: Non-editable text box Range: Policy DRA Server in Congestion
* Critical Alarm Onset Threshold	Percentage of Policy DRA 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 DRA 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
* Major Alarm Onset Threshold	Percentage of Policy DRA 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

Policy DRA Configuration

Fields (* indicates required field)	Description	Data Input Notes
* Major Alarm Abatement Threshold	Percentage of Policy DRA 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 DRA 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 DRA 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
Message Throttling Rules Tabs for Congestion Level 1, Congestion Level 2, and Congestion Level 3		
* 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 Level 3 - 0

Viewing Congestion Options

Use this task to view configured Congestion Options on the NOAM.

Select **Policy DRA > Configuration > Congestion Options**.

The **Policy DRA > 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 DRA 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 DRA > Configuration > Congestion Options**.

The **Policy DRA > 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.

PCRFs

The **Policy DRA > Configuration > 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 DRA 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. The Policy DRA 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 and Mediation User Guide* and Diameter online help.

On the **Policy DRA > Configuration > 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 DRA > Configuration > PCRFs [Insert]** page opens. You can add a PCRF. See [Inserting PCRFs](#). If the maximum number of PCRFs (2500) already exists in the system, the **Policy DRA > Configuration > 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 DRA > Configuration > PCRFs [Edit]** page opens. You can 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 26: PCRFs page elements](#) describes the elements on the **Policy DRA > Configuration > PCRFs** page on the Active SOAM.

Note: Data Input Notes apply to the Edit page; the View page is read-only.

Table 26: 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 DRA > Configuration > PCRFs**.

The **Policy DRA > Configuration > 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 DRA > Configuration > PCRFs**.

The **Policy DRA > Configuration > PCRFs** page appears.

2. Click **Insert**.

The **Policy DRA > Configuration > 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 DRA > Configuration > PCRFs** page.
- **Apply** to save the new PCRF and remain on this page.
- **Cancel** to return to the **Policy DRA > Configuration > 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 (2500) 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 DRA > Configuration > PCRFs**.

The **Policy DRA > Configuration > 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 (unless you want 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 DRA > Configuration > PCRFs** page.
- **Apply** to save the change and remain on this page.
- **Cancel** to return to the **Policy DRA > Configuration > 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 DRA configuration. In this procedure, PCRF refers to a Diameter peer of the Policy DRA, 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 DRAs configuration. In this procedure, PCRF refers to a Diameter peer of the Policy DRA, 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 > Configuration > 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 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 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 DRA > Configuration > 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.

Binding Key Priority elements

[Table 27: Binding Key Priority elements](#) describes the elements on the **Policy DRA > Configuration > Binding Key Priority** page.

Table 27: 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 DRA > Configuration > Binding Key Priority**.

The **Policy DRA > Configuration > 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 DRA > Configuration > Binding Key Priority**.
The **Policy DRA > Configuration > 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 DRA > Configuration > 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

Topology Hiding

Use the **Policy DRA > Configuration > Topology Hiding** page 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 Hosts** or **All Foreign Realms + Specific Hosts** on the **Policy DRA > Configuration > Site Options** page. See [Site Options](#).

The fields are described in [Topology Hiding elements](#).

On the **Policy DRA > Configuration > Topology Hiding** page, you can:

- 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 DRA > Configuration > Topology Hiding [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 DRA > Configuration > Topology Hiding [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 DRA > Configuration > Topology Hiding [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.)

- 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](#).

Topology Hiding elements

[Table 28: Topology Hiding elements](#) describes the elements on the **Policy DRA > Configuration > Topology Hiding** page. Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 28: Topology Hiding 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
Comments	An optional comment that describes the Policy Client Peer Node.	Format: Text box Range 0-64 characters

Viewing Topology Hiding

Use this task to view all configured Topology Hiding settings on the SOAM.

Select **Policy DRA > Configuration > Topology Hiding**.

The **Policy DRA > Configuration > Topology Hiding** page appears.

The fields are described in [Topology Hiding 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 Hosts** or **All Foreign Realms + Specific Hosts** in the Policy DRA > Configuration > Site Options page.

The fields are described in [Topology Hiding elements](#).

1. On the Active SOAM, select **Policy DRA > Configuration > Topology Hiding**.
The **Policy DRA > Configuration > Topology Hiding** page appears.
2. Click **Insert**.
The **Policy DRA > Configuration > Topology Hiding [Insert]** page appears.
3. Select a Policy Client Peer Node Name from the **Value** pulldown list.
4. Enter an optional comment in the **Comments** field.
5. Click:
 - **OK** to save the changes and return to the **Policy DRA > Configuration > Topology Hiding** page.
 - **Apply** to save the changes and remain on this page.
 - **Cancel** to return to the **Policy DRA > Configuration > Topology Hiding** 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 Topology Hiding

Use this task to edit a Policy Client for Topology Hiding.

Note: Topology Hiding is performed only if it is also activated and the Topology Hiding **Scope** option is defined as **Specific Hosts** or **All Foreign Realms + Specific Hosts** in the **Policy DRA > Configuration > Site Options** page.

The fields are described in [Topology Hiding elements](#).

1. On the Active SOAM, select **Policy DRA > Configuration > Topology Hiding**.
The **Policy DRA > Configuration > Topology Hiding** page appears.
2. Click **Edit**.
The **Policy DRA > Configuration > Topology Hiding [Edit]** page appears.
A read-only value is displayed in the Policy Client Peer Node Name **Value** field.
3. Edit or enter an optional comment in the **Comments** field.
4. Click:
 - **OK** to save the edited Comment and return to the **Policy DRA > Configuration > Topology Hiding** page.
 - **Apply** to save the edited Comment and remain on this page.
 - **Cancel** to return to the **Policy DRA > Configuration > Topology Hiding** 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. Select **Policy DRA > Configuration > Topology Hiding**.
The **Policy DRA > Configuration > Topology Hiding** 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 DRA > Configuration > 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 DRA > Configuration > Topology Hiding** page is refreshed. The row that was selected is no longer displayed in the list.

Site Options

The Policy DRA Site Options apply independently to each Policy DRA site. The following Site Options can be configured on **P-DRA > Configuration > Site Options** page on the active SOAM server:

- Topology Hiding Options - Enable/Disable, Scope, FQDN, and Realm. See [Site Options elements](#)
- Peer Route Table Name - The name of the configured Diameter Peer Route Table that contains the Peer Routing Rules to be used for routing new binding Requests.
- Ingress Message Capacity (read only) - Ingress message capacity for a single MP server; the value is the same as the **Engineered Ingress MPS** value in the Session **MP Profile** assigned for the blade type on which the Policy DRA application is running

The fields are described in [Site Options elements](#).

Site Options elements

[Table 29: Site Options elements](#) describes the elements on the SOAM **Policy DRA > Configuration > Site Options** page. Data Input Notes apply to the Insert and Edit pages; the View page is read-only.

Table 29: Site Options elements

Field (* indicates field is required)	Descriptions	Data Input Notes
Topology Hiding Options	Enabled Check or uncheck the box to Enable (checked) or Disable (unchecked) Topology Hiding. When the box is checked, select the Scope and enter FQDN and Realm that apply for Topology Hiding.	Format: Check box Range: Checked (Enabled), Unchecked (Disabled) Default: Unchecked (Disabled)
	Scope The scope of messages where Topology Hiding will be applied.	Format: Pulldown list Range: • All Messages -Perform Topology Hiding for all messages destined for Policy Clients. • All Foreign Realms - Perform Topology Hiding if the Realm of the Policy Client is different from the Realm of the PCRF that originated the message. • Specific Hosts - Perform Topology Hiding only if the Policy Client is configured on the Policy DRA > Configuration > Topology Hiding page. • All Foreign Realms +Specific Hosts - Perform Topology Hiding if either the 'All Foreign Realms' or 'Specific Hosts' condition is met.
	FQDN 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.	Format: Text box Range: 1 - 255 characters. Valid characters are letters, digits, dots (.), and hyphens (-). At least one alpha character is required.

Field (* indicates field is required)	Descriptions	Data Input Notes
	Realm 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.	Format: Text box Range: 1 - 255 characters. Valid characters are letters, digits, dots (.), and hyphens (-). At least one alpha character is required.
Peer Route Table Name	The name of the Diameter Peer Route Table to be used for routing new binding requests. The Default PRT is always available, but must be selected from the list to be used.	Format: Pulldown list Range: Not Selected, Default, configured Diameter Peer Route Tables Default: Not Selected

Table 30: Topology Hiding Scope Configuration shows the available Topology Hiding settings and corresponding results.

Note: Topology Hiding must be performed at the originating P-DRA.

Table 30: Topology Hiding Scope Configuration

Topology Hiding System Setting	Topology Hiding Scope Setting	Result
Disabled	N/A	No Topology Hiding is performed
Enabled	Specific Hosts	Topology Hiding is performed for messages for the Policy Clients only if the Policy Clients' FQDNs are configured for Topology Hiding.
	All Foreign Realms	Topology Hiding is performed for messages for the Policy Clients if the Realms of the Policy Clients are different from the Realm of the PCRF that sends the messages.
	All Foreign Realms + Specific Hosts	Superset of All Foreign Realms
	All	Topology Hiding is performed for all messages to all Policy Clients

Viewing Site Options

Use this task to view all configured Site Options on an SOAM.

Select **Policy DRA > Configuration > Site Options**.

The **Policy DRA > Configuration > 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](#).

Note: The **Ingress Message Capacity** field is read-only; it cannot be changed.

1. Select **Policy DRA > Configuration > Site Options**.

The **Policy DRA > Configuration > Site Options** page appears.

2. Select a mate DSR Peer Node from the **Policy DRA Mate DSR Peer Node Name** pulldown list.

This defines the Peer Node to which binding capable session initiating messages are routed if no local PCRFs are configured.

3. Check the Topology Hiding Options **Enabled** check box to enable or disable Topology Hiding.

Note: If Enabled, select the **Scope**, **FDQN** and **Realm** to apply topology hiding.

4. If the Topology Hiding Options **Enabled** box is checked, enter or select the **Scope**, **FDQN**, and **Realm** values to apply for Topology Hiding.

5. Select a Peer Route Table from the **Peer Route Table Name** pulldown list.

This identifies the Peer Route Table that contains the Peer Routing Rules that are used for routing new binding Requests.

6. Click:

- **Apply** to save the changes and refresh this page.
- **Cancel** to discard the changes and remain on the **Policy DRA > Configuration > 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.

Error Codes

For each Policy DRA 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 P-DRA. 3-digit codes are unique across all DSR layers (DSR connection layer, routing layer and application layer) and all DSR applications (P-DRA, RBAR, FABR etc.) for errors they represent. The ranges of 500-549 and 850-899 are for P-DRA application, while 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 31: Policy DRA Error Conditions

Error Conditions	Error Codes	Description	Applies to
P-DRA Unavailable Or Degraded	305	Returned if 1) The Policy DRA application Operational Status=Unavailable due to disabling the application on the Diameter > Maintenance > Applications GUI page (Admin State=Disabled), or 2) The Policy DRA application is in a Degraded stat due to congestion	Gx, Gxx, Gx-Prime, Rx, S9 messages
Binding Not Found	505	Returned if an Rx or Gx-Prime session is created or updated with an AAR message and the Binding Key in the Rx or Gx-Prime message cannot be found in the database	Rx and Gx-Prime sessions
Binding Found, But Unable To Route	502, 3-digit error code from DRL, 510 - 5513	Returned if a binding is found or created and the Policy DRA is unable to route the message to the PCRF	Gx, Gxx, Gx-Prime, Rx, S9 session creation messages (CCR-I and AAR)
Policy SBR Error	507, 508, 520, 521, 504	Returned if the Policy DRA receives an unexpected error while executing a database	Gx, Gxx, Gx-Prime, Rx, S9 messages

Error Conditions	Error Codes	Description	Applies to
		operation such as a lookup, insertion, or deletion of records	
No Usable Keys In Binding Dependent Message	503	Returned if an AAR message does not contain a key or keys that match the keys configured in the Policy DRA > Configuration > Binding Key Priority GUI page	Rx and Gx-Prime sessions
Session Not Found	509	Returned if the Policy DRA is unable to find a session record matching the in-session message	In-session Gx, Gxx, Gx-Prime, Rx, S9 messages, only when Topology Hiding applies to the message
Missing or Unconfigured APN	500, 501	Returned if 1) Binding capable session initiation request received with no APN, or 2) Binding capable session initiation request received with an APN, but the APN is not configured in the APN configuration.	Gx, Gxx, S9 messages

On the **Policy DRA > 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 DRA > 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 33: Error Codes elements](#) describes the elements on the **Policy DRA > 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 DRA Error Conditions. Each Error Condition will return the Result Code configured for each applicable Diameter interface.

[Table 32: 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 32: Interfaces Supported for Each Error Code

Error Code	Result Code	Vendor ID
Policy DRA Unavailable Or Degraded	Gx/Gxx, Gx-Prime, Rx, S9	Gx/Gxx, Gx-Prime, Rx, S9
Binding Not Found	Rx, Gx-Prime	Rx, Gx-Prime
Binding Found, But Unable To Route	Gx/Gxx,Gx-Prime, Rx, S9	Gx/Gxx,Gx-Prime, Rx, S9
Policy SBR Error	Gx/Gxx,Gx-Prime, Rx, S9	Gx/Gxx,Gx-Prime, Rx, S9
No Usable Keys In Binding Dependent Message	Rx,Gx-Prime	Rx,Gx-Prime
Session Not Found	Gx/Gxx,Gx-Prime, Rx, S9	Gx/Gxx,Gx-Prime, Rx, S9
Missing or Unconfigured APN	Gx/Gxx,S9	Gx/Gxx,S9

Table 33: Error Codes elements

Fields (* indicates required field)	Description	Data Input Notes
Error Condition	The name of the selected Policy DRA 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
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
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
* Rx Result Code	The Result Code to be returned to the Rx interface.	Format: Text box Range: 1-9999 Default: 3002

Fields (* indicates required field)	Description	Data Input Notes
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

Viewing Error Codes

Use this task to view configured Error Codes on the SOAM.

Select **Policy DRA > Configuration > Error Codes**.

The **Policy DRA > 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 DRA > Configuration > Error Codes**.

The **Policy DRA > Configuration > Error Codes** page appears

2. Select the **Error Condition** that you want to edit.

3. Click **Edit**.

The **Policy DRA > Configuration > Error Codes [Edit]** page opens

The fields that appear on the **Policy DRA > 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 DRA > Configuration > Error Codes** page
- **Apply** to save the changes and remain on this page
- **Cancel** to discard the changes and return to the **Policy DRA > 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).

Post-Configuration Activities

After Policy DRA configuration is complete, the following activities need to be performed to make the Policy DRA application fully operational in the system:

- Enable the Policy DRA application
- Enable Diameter Connections with Peer Nodes
- Status Verification

Enable the Policy DRA Application

Use this task to enable the Policy DRA application. For each Active SOAM,

1. Select **Diameter > Maintenance > Applications**.

The **Diameter > Maintenance > Applications** page appears.

2. Under **DSR Application Name**, select each **PDRA** 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 Policy DRA and Policy 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 Policy SBR Status.
 - a) At the Active NOAM, select **Policy DRA > Maintenance > Policy 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.

DSR Bulk Import and Export

The following documents describe the use and operation of DSR Bulk Import and Export functions:

- *Diameter User's Guide*, "Diameter Configuration", "DSR Bulk Import", "DSR Bulk Export"
- **Help > Diameter > Configuration > DSR Bulk Import**
- **Help > Diameter > Configuration > DSR Bulk Export**

The DSR Bulk Import and Export functions can be used to export Diameter, IPFE, and DSR 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.

Configuration data refers to any data that is configured for one of the Export **Export Application** types (FABR, RBAR, Policy DRA, GLA , or CPA and SBR DSR Applications; IPFE; and the Diameter Configuration components).

DSR Bulk Export

The DSR Bulk Export operation creates ASCII Comma-Separated Values (CSV) files (.csv) containing Diameter , IPFE, and DSR Application configuration data. Exported configuration data can be edited and used with the DSR 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 DSR 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 DSR Application, IPFE, or Diameter (each component's data is in a separate file)
- Exportable configuration data from a selected configuration component for the selected DSR Application, IPFE, or Diameter

Exported files can be written to the File Management Directory in the local File Management area (**Status & Manage > File** 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.

DSR Bulk Import

The DSR 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 Diameter Configuration, IPFE Configuration, or DSR Applications (FABR, RBAR, Policy DRA, GLA, and CPA/SBR) 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 "DSR Bulk Import" section of the *Diameter User's Guide* or the **Diameter > Configuration > Import** Help for valid Import operations.

Import CSV files can be created by using a DSR 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 DSR release that is used to import the file.

Files that are created using the DSR 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 on a computer; 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.

Chapter 6

Policy DRA Maintenance

Topics:

- [Introduction.....157](#)
- [Policy DRA Maintenance Pages.....157](#)
- [Alarms, KPIs, and Measurements.....161](#)
- [Binding and Session Database Auditing.....163](#)
- [Overload Management.....163](#)
- [Shutdown.....166](#)
- [Diameter Maintenance and Status Data for Components, DSR Applications, and DA-MPs.....167](#)
- [Backup and Restore for Policy DRA Configuration Data.....168](#)

This chapter describes or indicates where to find the following information that can be used for the Policy DRA application and Policy SBR:

- Maintenance and status information that is maintained by the Policy DRA Configuration and Maintenance components and displayed on the **Policy DRA > 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 DRA and Policy SBR alarms, KPIs, and measurements
- Auditing of the Session and Binding databases
- Policy DRA and Policy SBR overload management
- Database Backup and Restore of Policy DRA configuration data

Introduction

This chapter describes:

- [Policy DRA Maintenance Pages](#) describes maintenance and status data that is maintained by the Policy DRA application and by Policy DRA DA-MPs.

On the **P-DRA > Maintenance** pages you can:

- View Policy SBR Status
- Define and execute a Binding Key Query
- [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 DRA)
- DA-MPs
- [Alarms, KPIs, and Measurements](#) describes Policy DRA-specific database alarms, and indicates the location of descriptions of Policy DRA and Policy SBR alarms, KPIs, and measurements.
- [Binding and Session Database Auditing](#) describes the auditing of the Session and Binding databases.
- [Overload Management](#) describes overload controls and load shedding and for Policy DRA and Policy SBR.
- [Backup and Restore for Policy DRA Configuration Data](#) describes the OAM database backup and restore of Policy DRA configuration data.

Policy DRA Maintenance Pages

The Policy DRA > Maintenance GUI pages on the NOAM display Policy SBR status information and provide access to the Binding :Key Query tool.

Policy SBR Status

The **Policy DRA > Maintenance > Policy SBR Status** page displays a collapsed or expanded detailed report for Policy SBR. The data is displayed within Server Groups by configured Place Associations.

Fields are described in [Policy SBR Status elements](#).

Policy SBR Status elements

Table 34: Policy SBR Status elements describes the elements on the **Policy SBR Status** page, which displays Policy 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 34: Policy SBR Status elements

Elements	Description	Data Input Notes
One <Binding Region> tab	A list of all configured Server Groups that are assigned to the Binding Region Place Association. The Resource Domain Name and the Resource Domain Profile of each Server Group is shown. The 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.	The page is view-only. 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.
A tab for each Policy DRA Mated Site in the system	Each tab displays a list of all configured Server Groups that are assigned to that Mated Pair Place Association. The Resource Domain Name and the Resource Domain Profile of each Server Group are shown. The 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.	

The Binding Key Query

Use the **Policy DRA > Maintenance > Binding Key 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 Binding Key Query tool can be used only with Gx sessions. It is not applicable to Rx sessions.

The fields are described in *Binding Key Query elements*.

To use the Binding Key Query tool,

1. On the Active NOAM, select **Policy DRA > Maintenance > Binding Key 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. The following is an example of a report:

```
=====
B i n d i n g      K e y      Q u e r y      R e p o r t      -      I M S I
=====
Report Generated: Thu Oct 24 2:21:16 2013 EDT
From: Network OAM&P on host E10301Blade01N002
Report Version: 0.45740
User: guiadmin

-----

Results for IMSI Binding Key = 111111180000001
  Binding 1 Data:
    Access Point Name: tekelecLte.com
    PCRF Pool Name: tekelecLtePcrfSubPool1
    PCRF FQDN: pcrf1.tekelec.com
    Binding State: Final
    Binding Creation Date/Time: Tue Oct 22 04:45:42 2013 America/New_York
    pSBR(S) Site Name: site1
    pSBR(S) Server Group: BPSBRMP_SG
    pSBR(S) Server Name: E10301Blade04Bpsbr1
    Binding Sub Resource Id:6

    Session 1 Data:
      Diameter Session-Id: OAM;1096298391;1
      IMSI: 111111180000001
      MSISDN Key: 9999999999
      IPv4 Address: 10.39.223.126
      IPv6 Address: 2606:ae00:4ecc:e901::
      Binding Capable: Yes
      PCEF FQDN:
      Creation Date/Time: Tue Oct 22 04:45:42 2013 America/New_York
      Last Touched Date/Time: Tue Oct 22 04:45:42 2013 America/New_York
      pSBR(S) Site Name: site1
      pSBR(S) Server Group: SPSBRMP_SG
      pSBR(S) Server Name: E10301Blade07Spsbr2
      Session Reference: 363b6652c0060200b20a2000
      Session Sub Resource Id: 1

  Binding 2 Data:
    Access Point Name: tekelecVolte.com
    PCRF Pool Name: tekelecVoltePcrfPool
    PCRF FQDN: pcrf2.tekelec.com
    Binding State: Final
    Binding Creation Date/Time: Tue Oct 22 07:28:27 2013 America/New_York
    pSBR(S) Site Name: site1
    pSBR(S) Server Group: BPSBRMP_SG
    pSBR(S) Server Name: E10301Blade04Bpsbr1
    Binding Sub Resource Id:6

    Session 1 Data:
      Diameter Session-Id: OAM;1096298391;1
      IMSI: 111111180000001
      MSISDN Key: 9999999999
      IPv4 Address: 10.39.223.126
      IPv6 Address: 2606:ae00:4ecc:e901::
```

```

Binding Capable: Yes
PCEF FQDN:
Creation Date/Time: Tue Oct 22 04:45:42 2013 America/New_York
Last Touched Date/Time: Tue Oct 22 04:45:42 2013 America/New_York
pSBR(S) Site Name: site1
pSBR(S) Server Group: SPSBRMP_SG
pSBR(S) Server Name: E10301Blade07Spsbr2
Session Reference: 363b6652c0060200b20a2001
Session Sub Resource Id: 1

```

```

-----
E n d   O f   B i n d i n g   K e y   Q u e r y   R e p o r t   -   I M
S I
=====

```

To enter another query, click **Clear**, and select and enter the values for the new search.

Binding Key Query elements

[Table 35: Binding Key Query elements](#) describes the elements on the **Policy DRA > Maintenance > Binding Key Query** page.

Table 35: Binding Key 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. • 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::

Alarms, KPIs, and Measurements

This section describes the type of alarm, KPI, and measurements information that is available for Policy DRA and Policy SBR, and how to access the information in the DSR GUI.

Policy DRA and Policy SBR Alarms and Events

Policy DRA application and Policy SBR alarms and events are described in the *Alarms, KPIs, and Measurements Reference* and the DSR Alarms, KPIs, and Measurements online help.

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.

Database Alarms

The Policy DRA application supports two Policy SBR alarms related to database capacity:

- A **Binding Capacity alarm**: "Policy SBR Bindings Threshold Exceeded"

The Binding Capacity alarm scope is network-wide. The Binding Capacity alarm is raised and cleared based on the percentage full of the Binding database.

The assertion threshold values are specified as percentages and can be configured at any time using the **Policy DRA > Configuration > Alarm Settings** GUI page on the active NOAM. Each alarm severity can be suppressed if desired by checking a box on the **Policy DRA > Configuration > Alarm Settings** GUI page.

The Binding Capacity alarm measures the number of binding (IMSI) records against an engineered maximum value that varies according to the number of Binding Policy SBR Server Groups that are specified at feature activation.

Because no single Binding Policy SBR server holds the entire Binding database (except in the case of small systems with only one Binding Policy SBR Server Group), each Binding Policy SBR reports the size of its portion of the database to the NOAM server. A mechanism on the NOAM aggregates the reported database size records such that only the records from active servers in each Server Group are counted. This summation is then converted into a percent-full of the maximum database size and compared against the assertion and abatement thresholds, causing alarms to be raised and cleared accordingly.

Alarm abatement occurs at 5% below the assertion threshold for each alarm severity. For example, if the minor alarm threshold is configured as 70%, a minor alarm will clear only after the database size drops below 65% full.

- A **Session Capacity alarm**: "Policy SBR Sessions Threshold Exceeded"

The Session Capacity alarm is scoped to a mated pair of Policy DRA DSRs because each mated pair has its own instance of the Session database. The Session Capacity alarm is raised and cleared based on the percentage full of an instance of the Session database.

The assertion threshold values are specified as percentages and can be configured any time, using the **Policy DRA > Configuration > Alarm Settings** GUI page on the active NOAM. Each alarm severity can be suppressed if desired by checking a box on the **Policy DRA > Configuration > Alarm Settings** GUI page.

The Session Capacity alarm percent full is based on the number of Session records compared to an engineered maximum, which varies according to the number of Session Policy SBR Server Groups per mated pair that are chosen at Policy DRA feature activation.

Because no single Session Policy SBR server holds the entire Session database (except in the case of small systems with only one Session Policy SBR Server Group), each session Policy SBR reports the size of its portion of the database to the NOAM server. A mechanism on the NOAM aggregates the reported database size records such that only the records from active servers in each Server Group in an instance of the Session database are counted. This summation is then converted into a percent-full of the maximum database size and compared against the assertion and abatement thresholds, causing alarms to be raised and cleared accordingly.

Alarm abatement occurs at 5% below the assertion threshold for each alarm severity. For example, if the minor alarm threshold is configured as 70%, a minor alarm will clear only after the database size drops below 65% full.

If a Policy SBR Session Capacity alarm is asserted, the “instance” field of the alarm is set to the name of the Policy DRA Mated Pair **Place Association** that identifies the Policy DRA mated pair.

DSR Application Ingress Message Rate Alarm

The number of ingress messages (both Requests from PCEF and Answers from PCRF) per second received by Policy DRA is counted as input to Policy DRA ingress message processing capacity. The capacity is an engineering system value for the number of ingress messages per second processed by Policy DRA for a single MP server.

Thresholds (in percentages) associated with the Policy DRA ingress message capacity can be configured on the **Policy DRA > Configuration > Alarm Settings** GUI page on the active NOAM.

If the ingress message rate received at Policy DRA exceeds the configured percentage of the maximum capacity, this alarm is raised at the appropriate severity (Minor, Major, Critical).

This alarm is cleared when the Ingress Message rate drops below the configured percentage of the ingress message capacity for the alarm severity (Minor, Major, Critical).

Policy SBR Audit Report Event 22716

To limit the effects of stale Binding and Session records, all Policy DRA MPs that own an active part of the database continually audit each table to detect and remove stale records.

In order to have some visibility into what the audit is doing, the audit generates Event 22716 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. The audit reports for each table type are formatted as described in [Table 12: Audit Report Formats](#).

Policy DRA and Policy 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 Policy DRA and Policy SBR are displayed on the **Status & Manage > KPIs** GUI page. Selecting the tab for a server and either **P-DRA** or **pSBR** under the tab displays the KPI information for the selected server.

The Policy DRA and Policy SBR KPIs are described in the *DSR Alarms, KPIs, and Measurements Reference* and the DSR Alarms, KPIs, and Measurements online help.

Policy DRA and Policy SBR Measurements

Measurements for Policy DRA and Policy 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 Alarms, KPIs, and Measurements* online help and PDF explain the report selection criteria, and describe each measurement in each measurement group.

Binding and Session Database Auditing

See [Policy Data Auditing](#).

Overload Management

The Policy DRA application provides mechanisms to manage the overload and congestion that can occur on the Policy DRA and Policy SBR. The Policy DRA might receive ingress messages at a rate higher than the engineered capacity. The internal queues on the Policy DRA might experience higher utilization level than configured. The same might happen on the Policy SBR servers, directly or indirectly resulting from the overloaded traffic from the network or from the Policy DRA.

Overload Controls

The Policy 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 DRA DA-MP Overload Control

The number of ingress messages (both Requests and Answers) per second received by Policy DRA is counted as input to Policy DRA ingress message processing capacity. The capacity is an engineering number of ingress messages per second processed by Policy DRA. The number of Request messages received at Policy DRA per second is also measured separately.

Policy DRA defines alarms on the queue utilization levels based on configured threshold values. Thresholds (in percentage) are configured in association with the Policy DRA ingress message capacity. If the ingress message rate received at Policy DRA exceeds the configured percentage of the maximum capacity, alarms will be raised. Policy DRA ingress Request capacity can be engineering configured to provide the value based on which thresholds (in percentage) are configured. See [Alarm Settings](#).

The Policy DRA congestion is then defined by the ingress Request messages capacity and the configured threshold values. Policy DRA will be considered in congestion if the ingress Request rate at Policy DRA exceeds the configured percentages (thresholds) of Policy DRA ingress Request capacity.

Three Policy DRA 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 Policy DRA is in congestion, a Policy DRA 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, Policy DRA will perform overload control by throttling a portion of incoming messages to keep Policy DRA from being severely impacted. The type and percentage of the messages to be throttled will be configurable through the Policy DRA GUI as displayed in [Figure 28: Policy DRA Default Overload Control Thresholds](#):

Congestion Levels	Alarm-ID 22721 Severity	P-DRA Operation Status	Message Throttling Rules
CL_0	N/A	Available	No messages are discarded (Accept 100% Request and Answer messages)
CL_1	Minor	Available	- 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
CL_2	Major	Available	- 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
CL_3	Critical	Degraded	- Discard 100% of requests for creating new sessions - Discard 50% of requests for updating existing sessions - Discard 0% of requests for terminating existing sessions - Discard 0% of answer messages

Figure 28: Policy DRA Default Overload Control Thresholds

The Policy DRA's internal congestion state contributes to Policy DRA's Operational Status directly, along with its Admin state and Shutdown state. Consequently, the congestion state of the Policy DRA impacts the Diameter Routing Function message transferring decision. Depending on the Policy DRA's Operational Status (Unavailable, Degraded, Available), the Diameter Routing Function will forward all the ingress messages to the Policy DRA when the Policy DRA's Operational Status is Available, or discard some or all of the ingress messages when the Operational Status is Degraded or Unavailable. [Table 36: Diameter Routing Function Message Handling Based on Policy DRA Operational Status](#) describes the Diameter Routing Function handling of the messages to the Policy DRA.

Table 36: Diameter Routing Function Message Handling Based on Policy DRA Operational Status

Policy DRA Operational Status	Diameter Routing Function Message Handling
Available	Forward all Request and Answer messages to Policy DRA
Degraded	Forward all Answer messages only to Policy DRA
Unavailable	Discard all messages intended for Policy DRA

Policy SBR Congestion

Policy 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 (Policy 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 Policy SBR can call for receiving congestion level notifications.

Policy SBR congestion is measured based on the Policy SBR CPU utilization level. There are four Policy 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 Policy SBR congestion state (CPU utilization) is managed and controlled by the ComAgents on both Policy DRA and Policy SBR MPs based on the ComAgent MP Overload Management Framework. Messages to a Policy SBR from a Policy DRA are handled based on the congestion state of the Policy SBR. A Policy 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 Policy 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 Policy SBR to which the stack event is sent. If the priority is higher than or equal to the Policy SBR current congestion level, the message will be forwarded. Otherwise, it will be discarded.

The stack events may also be routed from a Policy SBR to another Policy SBR in some scenarios. The congestion control in this case should be conducted based on the congestion state of the receiving Policy SBR, i.e. the ComAgent on the sending Policy SBR is responsible to compare the stack event priority with the congestion level of the receiving Policy SBR and make the routing decision accordingly.

Load Shedding

After the Policy 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 Policy SBRs using ComAgent stack events. Each stack event has a priority. The resource user software (on either DA-MPs or Policy 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 37: Stack Event Load Shedding](#).

Table 37: 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 DRA application running on DA-MPs supports the DSR Application Infrastructure graceful shutdown with 5 seconds grace period. This means that when Policy DRA 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 Policy DRA application.

Policy SBR - Because Policy 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 Policy 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 Policy DRA Operational Status (Unavailable, Degraded and Available) is determined by its Admin State, Congestion Level, and the Shutdown State. The Policy DRA application calculates and maintains its own operational status and reports it to the Diameter Routing Function.

When the Policy DRA application 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 Policy DRA DA-MP server,

1. Go to the Diameter -> Maintenance -> Applications GUI page and Disable the Policy DRA 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 38: Policy DRA Operational Status shows an example of the Policy DRA 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 Policy DRA continues to process the messages in its queue for a time interval that is engineering configurable.

Table 38: Policy DRA 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	N/A	Available	Normal
	1			Available with CL_1
	2			Available with CL_2
Enabled	3	N/A	Degraded	Congested with CL_3

Policy SBR - Because Policy 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 Policy 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 DRA, and DA-MPs including those that run the Policy DRA 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.

- **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 and Mediation User Guide* and in the Diameter Help.

Backup and Restore for Policy DRA Configuration Data

Because Policy DRA 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 sessions and policy bindings that is mastered on Policy SBR MP servers is not backed up or restored.

The Policy DRA 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*.

Appendix

A

Policy DRA PCRF Pooling Upgrade

Topics:

- [Upgrade Paths.....170](#)
- [Configuration After Upgrade.....172](#)
- [Concepts and Terminology.....173](#)
- [Configuring PCRF Pooling.....176](#)
- [Processing Phases.....179](#)
- [Binding Migration.....180](#)

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 Policy DRA 4.1.5 or 5.0. 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.

Upgrade Paths

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.

Upgrading on Previously Activated Policy DRA Releases

If Policy DRA is already activated on a DSR that has been upgraded to the release that supports PCRF Pooling and the Policy DRA activation occurs after the upgrade is completed and accepted, the following conditions apply:

- Diameter must be configured according to the appropriate release documentation.
- Policy DRA feature must be activated.
- 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 P-DRA 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

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-I's.
- 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 39: Processing During Transition Period

Request Type	Processing during transition period
CCR-I	If an existing binding matches for the IMSI+APN combination, then route to the PCRF indicated in the binding. Note: Any binding for the IMSI that existed prior to enabling PCRF Pooling will match any IMSI+APN combination for that IMSI. The following logic applies: • 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)	If an existing secondary key matches for the MSISDN+APN combination, then route to the PCRF indicated by the secondary key. Note: Any binding for the MSISDN that existed prior to enabling PCRF Pooling will match any MSISDN+APN combination for that MSISDN. Else, existing behavior for invalid request (binding not found).
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.

Related Topic

[*Policy DRA Configuration on the NOAM and the SOAM*](#)

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.

Enabling PCRF Pooling

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 will be routed according to the PCRF Pooling behavior.

The following rules apply to PCRF Pooling enablement:

- When upgrading the Policy DRA application, all existing binding and session database entries are maintained.
- When upgrading to a version of the Policy DRA application that supports PCRF Pools, the default PCRF pool must be defined.
- When upgrading to a version of the Policy DRA application that supports PCRF Pools, all existing APNs must be mapped to the default pool.
- When upgrading to a version of the Policy DRA application that supports PCRF Pools, the default pool must point to the PRT table used for routing of new-binding requests prior to the upgrade.
- When upgrading to a version of the Policy DRA application that supports the PCRF pool feature prior to the PCRF Pool feature being enabled, the PCRF pool feature shall not be enabled as a result of the upgrade procedure.
- Upgrade procedures from the Policy DRA version 4.1.5 to version 6.0 or later must ensure that the migration from the pre PCRF Pools pSBR database to the PCRF Pools pSBR database has finished prior to starting the upgrade.

Policy DRA Before and After PCRF Pooling Upgrade

The following figure illustrates the main differences between P-DRA before and after PCRF Pooling.




Figure 29: PCRF Pooling Effects on Policy DRA

See [PCRF Pools](#) for a description of the major differences between PCRF Pooling and non-pooling functionality.

Terminology for Upgrading Policy DRA

The following table shows a list of some common Policy DRA acronyms related to upgrading.

Table 40: 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.

Acronym/Term	Description
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:

1. Use the alarm history at **Alarms & Events > View History** with filter setting as follows:
 - a. Scope set to the NOAM Network Element
 - b. Display Filter set to Event ID=22730
 - c. Collection Interval set to N Days, where N is long enough to cover the period after the upgrade, or something shorter if required.
2. 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.
3. 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.

1. Determine a set of PCRF Peer Nodes that will be grouped into a PCRF Pool.
2. At each site's SOAM, for the PCRF Peer Nodes from item 1, create a Route Group containing those PCRFs.
3. For the Route Group in item 2 in this step, create a Route List that uses that Route Group.
4. For the Route List in item 3 in this step, create a Peer Route Table.
5. 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.
6. 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.

1. At **Policy DRA -> Configuration -> PCRF Pools**, insert a PCRF Pool using a descriptive name.
2. 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.

1. 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.
2. 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.

1. Use Policy DRA -> Configuration -> Error Codes.
2. Select the row for Missing or Unconfigured APN and click **Edit**.
3. 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.

1. 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

Policy DRA Error Resolution

Topics:

- [Introduction.....185](#)
- [Policy DRA Error Resolution Flowchart Summary.....187](#)
- [Diameter Message Validation Error Resolution Flowchart.....189](#)
- [Generic CCR Processing Error Resolution Flowchart.....190](#)
- [CCR-I Processing without PCRF Pool Error Resolution Flowchart.....191](#)
- [findSessionRef Processing Error Resolution Flowchart.....193](#)
- [findOrCreBindResShort Processing Error Resolution Flowchart.....193](#)
- [CCR-I Processing with PCRF Pool Error Resolution Flowchart.....196](#)
- [findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart.....197](#)
- [Early Bind Pool Error Resolution Flowchart....198](#)
- [CCA-I Processing Error Resolution Flowchart.....199](#)
- [findSession Response Processing Error Resolution Flowchart.....200](#)
- [CCR-U Processing Error Resolution Flowchart.....201](#)
- [CCR-T Processing Error Resolution Flowchart.....202](#)
- [CCA-U/T Processing Error Resolution Flowchart.....203](#)
- [RAR Processing Error Resolution Flowchart...204](#)
- [RAA Processing Error Resolution Flowchart...205](#)
- [AAR Processing Error Resolution Flowchart...206](#)
- [AAA Processing Error Resolution Flowchart..207](#)

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.

- *STR Processing Error Resolution Flowchart....208*
- *STA Processing Error Resolution Flowchart...209*
- *ASR/ASA Processing Error Resolution Flowchart.....210*

Introduction

Error resolution flowcharts illustrate Policy DRA 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 Policy DRA application and the error location.

You should use the information in the flowcharts with the additional error resolution steps included in *Alarms, KPIs, and 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 Policy DRA GUI during error resolution efforts.

Note: See [Error Codes](#) for more information about Policy DRA error conditions and error code numbers.

Measurement tags in the documentation are named to associate with receive and transmit. For example, the Rx in RxBindCapMissingApn indicates receive and the Tx in TxPdraErrAnsGeneratedCaFailure indicates transmit of send. Thus, the message name Diameter-interface-Rx-Binding-Cap-Missing-APN actually means Received-Binding-Cap-Missing-APN. Also, the Rx measurement tags are valid in the Gx interface, but a misinterpretation of Rx is possible and might lead a reader to ask how can an Rx (read as Diameter-interface-Rx) exist in a Gx diameter scenario.

Understanding Error Resolution Procedure Attributes

[Table 41: Error Resolution Attributes](#) lists the attributes in the resolution procedures that you should be familiar with before using this appendix. See *Alarms, KPIs, and 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 41: Error Resolution Attributes

Attribute	Procedure
Error Categories / Names	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 DRA SOAM GUI (Policy DRA -> Configuration -> Error

Attribute	Procedure
	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. In a single error scenario, if the value in the corresponding GUI Configurable attribute is No, 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	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. 3-digit codes are unique across all DSR layers (DSR connection layer, routing layer, and application layer) and all DSR applications (Policy DRA, RBAR, FABR, and so on) for errors the codes represent. The ranges of 500-549 and 850-899 are for Policy DRA application, while the DSR connection layer, routing layer, and other DSR applications use other non-overlapping ranges. 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.
Result Code in Policy DRA Generated Answer	If the error is not GUI configurable, this attribute is the value of the Result Code AVP in the Answer generated by Policy DRA. If the error is GUI configurable, the Result-Code value can be found in Policy DRA SOAM GUI (Policy DRA -> 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	This attribute provides information about the applicable Diameter interfaces or Diameter message types on which the corresponding error could occur. The Diameter interfaces are categorized as binding-capable (Gx/Gxx) or binding-dependent (Rx or Gx-Prime). Together with the P-DRA business logic flowcharts in this appendix, the

Attribute	Procedure
	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.
Direct Policy DRA Alarm/Event	The Policy DRA 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 Policy DRA business logic flowcharts indicating specifically where and when in the business logic that the corresponding error might occur. Policy DRA business logic flowcharts are included in this appendix.
Troubleshooting Steps/Customer Actions	See <i>Alarms, KPIs, and Measurements Reference</i> .

Policy DRA Error Resolution Flowchart Summary

Figure 30: 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 Policy 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 31: Diameter Message Validation Error Resolution Flowchart
Chart 2	Generic CCR Processing	Figure 32: Generic CCR Processing Error Resolution Flowchart
Chart 3	CCR-I Processing without PCRF	Figure 33: CCR-I Processing without PCRF Pool Error Resolution Flowchart
Chart 4	findSessionRef Processing	Figure 34: findSessionRef Processing Error Resolution Flowchart

Chart number	Flowchart Name	Link
Chart 5	findOrCreBindResShort Processing	Figure 35: findOrCreBindResShort Processing Error Resolution Flowchart
Chart 6	CCR-I Processing with PCRF Pool	Figure 36: CCR-I Processing with PCRF Pool Error Resolution Flowchart
Chart 7	findOrCreateBinding Response Processing with PCRF Pool	Figure 37: findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart
Chart 8	Early Bind Pool	Figure 38: Early Bind Pool Error Resolution Flowchart
Chart 9	CCA-I Processing	Figure 39: CCA-I Processing Error Resolution Flowchart
Chart 10	findSession Response Processing	Figure 40: findSession Response Processing Error Resolution Flowchart
Chart 11	CCR-U Processing	Figure 41: CCR-U Processing Error Resolution Flowchart
Chart 12	CCR-T Processing	Figure 42: CCR-T Processing Error Resolution Flowchart
Chart 13	CCA-U/T Processing	Figure 43: CCA-U/T Processing Error Resolution Flowchart
Chart 14	RAR Processing	Figure 44: RAR Processing Error Resolution Flowchart
Chart 15	RAA Processing	Figure 45: RAA Processing Error Resolution Flowchart
Chart 16	AAR Processing	Figure 46: AAR Processing Error Resolution Flowchart
Chart 17	AAA Processing	Figure 47: AAA Processing Error Resolution Flowchart
Chart 18	STR Processing	Figure 48: STR Processing Error Resolution Flowchart
Chart 19	STA Processing	Figure 49: STA Processing Error Resolution Flowchart
Chart 20	ASR/ASA Processing	Figure 50: ASR/ASA Processing Error Resolution Flowchart




Figure 30: Error Resolution Flowchart Summary

Diameter Message Validation Error Resolution Flowchart

Figure 31: 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 31: Diameter Message Validation Error Resolution Flowchart

Generic CCR Processing Error Resolution Flowchart

Figure 32: 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 32: Generic CCR Processing Error Resolution Flowchart

CCR-I Processing without PCRF Pool Error Resolution Flowchart

Figure 33: 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 33: CCR-I Processing without PCRF Pool Error Resolution Flowchart

findSessionRef Processing Error Resolution Flowchart

Figure 34: *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 34: findSessionRef Processing Error Resolution Flowchart

findOrCreBindResShort Processing Error Resolution Flowchart

Figure 35: *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 35: findOrCreBindResShort Processing Error Resolution Flowchart

CCR-I Processing with PCRF Pool Error Resolution Flowchart

Figure 36: *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 36: CCR-I Processing with PCRF Pool Error Resolution Flowchart

findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart

Figure 37: *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 37: findOrCreateBinding Response Processing with PCRF Pool Error Resolution Flowchart

Early Bind Pool Error Resolution Flowchart

Figure 38: 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 38: Early Bind Pool Error Resolution Flowchart

CCA-I Processing Error Resolution Flowchart

Figure 39: 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 39: CCA-I Processing Error Resolution Flowchart

findSession Response Processing Error Resolution Flowchart

Figure 40: 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 40: findSession Response Processing Error Resolution Flowchart

CCR-U Processing Error Resolution Flowchart

Figure 41: 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 41: CCR-U Processing Error Resolution Flowchart

CCR-T Processing Error Resolution Flowchart

Figure 42: CCR-T Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where CCR-T Processing errors can occur.




Figure 42: CCR-T Processing Error Resolution Flowchart

CCA-U/T Processing Error Resolution Flowchart

Figure 43: 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 43: CCA-U/T Processing Error Resolution Flowchart

RAR Processing Error Resolution Flowchart

Figure 44: 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 44: RAR Processing Error Resolution Flowchart

RAA Processing Error Resolution Flowchart

Figure 45: 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 45: RAA Processing Error Resolution Flowchart

AAR Processing Error Resolution Flowchart

Figure 46: 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 46: AAR Processing Error Resolution Flowchart

AAA Processing Error Resolution Flowchart

Figure 47: 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 47: AAA Processing Error Resolution Flowchart

STR Processing Error Resolution Flowchart

Figure 48: 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 48: STR Processing Error Resolution Flowchart

STA Processing Error Resolution Flowchart

Figure 49: STA Processing Error Resolution Flowchart shows an error resolution flowchart that illustrates where STA Processing errors can occur.




Figure 49: STA Processing Error Resolution Flowchart

ASR/ASA Processing Error Resolution Flowchart

Figure 50: 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 50: ASR/ASA Processing Error Resolution Flowchart

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; for example, 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.
ASA	Abort-Session-Answer

B

BBERF	Bearer Binding and Event Reporting Function: A type of Policy Client used to control access to the bearer network (AN).
Binding	A binding between a subscriber identifier (e.g. IMSI, MSISDN, IP Address) and an MPE. The MRA maintains bindings, and there is

B

one binding per subscriber even if the subscriber has multiple active sessions.

See Policy 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. A binding ceases to exist when the last Diameter session for that IMSI and APN is terminated. See also PCRF Pool Binding.

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.

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

CCA-T

Credit Control Answer - Terminate

CCA-U

Credit Control Answer - Update

CCR-I

CCR Initial

C

CEX Configuration Set

A mechanism for assigning Application IDs and supported Vendor IDs to a Local Node or to a Connection.

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.

DPI

Diameter Plug-In is a reusable Diameter stack consisting of DCL, DRL, and an application interface.
Deep Packet Inspection is a form of packet filtering that examines the data and/or header part of a packet as it passes an inspection point. The MPE device uses DPI to recognize the application for establishing QoS or managing quota. See also packet inspection.

DRL

Diameter Routing Layer
The software layer of the Eagle XG Diameter 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.

G

G

GUI	Graphical User Interface The term given to that set of items and facilities which provide the user 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 PCRF, and in reply the PCRF provides rule information for the PCEF to enforce.
Gx-Prime	A vendor specific Gx like interface with minor variations as the protocol for DPI and PCRF communications before the standardized Sd reference point/protocol was available. Gx-Prime uses the same Application Id (16777238) as Gx does and the same command code set (Credit Control Request/ Answer and Re-Auth Request/ Answer) as well.
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.
------	---

I

IMSI	International Mobile Subscriber Identity International Mobile Station 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

MITM	Man in the Middle
MOS	Media Optimization Server
MSISDN	Mobile Station International Subscriber Directory Number 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. Mobile Subscriber Integrated Services Digital Network [Number] Mobile Station International Subscriber Directory Number. The unique, network-specific subscriber number of a mobile communications subscriber. MSISDN follows the E.164 numbering plan; that is, normally the MSISDN is the phone number that is used to reach the subscriber.

O

O**OAM**

Operations, Administration, and Maintenance

The application that operates the Maintenance and Administration Subsystem which controls the operation of many products.

P**PCEF**

Policy and Charging Enforcement Function

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 anywhere in the network, must be processed by the same PCRF.

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 anywhere in the network, must be processed by the same PCRF.

Peer

A Diameter node to which a given Diameter node has a direct transport connection.

Peer Route Table

A set of prioritized Peer Routing Rules that define routing to Peer Nodes based on message content.

P

Peer Routing 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.
Policy DRA Binding Region	A type of Place Association that defines the scope of an instance of the P-DRA Binding database. In the context of the P-DRA network, a region is all of the sites in the P-DRA network. P-DRA supports only one instance of the Policy Binding Region, meaning that there is only one Binding database for the entire P-DRA Network.
Policy DRA Mated Pair	A type of Place Association. In the context of a P-DRA network, a Mated Pair is two P-DRA DSRs that are paired for redundancy such that if one site fails, the other site can take over the failed site's entire load. A Mated Pair sets the scope of an instance of the Policy Session database.
PRT	Peer Route Table or Peer Routing Table
pSBR	Policy SBR

R

RAA	Re-Authorization Answer (Gx or Rx Diameter command)
RAR	Re-Authorization Request (Gx or Rx Diameter command)
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
STA	Session-Termination-Answer Session Termination Answer (Rx Diameter command)
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.

S

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.

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.

The suspect binding mechanism allows a binding to be removed if the PCRF that the subscriber is bound to becomes unreachable. A binding is marked suspect if after being successfully established, a subsequent binding capable session initiation request for that same binding receives a 3002 response (unable to route) from the routing layer. If another binding capable session initiation request for the binding arrives after the suspect binding interval and also receives a 3002 response, the suspect binding is removed, allowing the next request to be routed to another PCRF.

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.

U

UE User Equipment

V

V-PCRF Visited PCRF

x

XSI	External Signaling IP Address
-----	-------------------------------

XSI External Signaling Interface