Oracle® Communications Diameter Signaling Router Policy and Charging Application





Oracle Communications Diameter Signaling Router Policy and Charging Application, Release 8.5

F27855-01

Copyright © 2011, 2020, Oracle and/or its affiliates.

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

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

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

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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

Contents

Devicion History	1 1
Revision History Overview	1-1 1-1
Scope and Audience	1-1
Manual Organization	1-1
My Oracle Support	1-2
Policy and Charging Introduction	
Overview	2-1
The P-DRA Database	2-2
Bindings vs Sessions	2-3
Bindings	2-3
Sessions	2-4
The Binding Database	2-5
The Session Database	2-7
Binding Key Query Tool	2-8
The OC-DRA Database	2-8
Online Charging over Gy/Ro Reference Point	2-8
Binding-independent Interface	2-10
Deployment Topology	2-10
Policy DRA in Roaming Scenarios	2-12
PCA Configurable Components	2-13
Places	2-16
Place Associations	2-16
Server Groups	2-17
Resource Domains	2-18
Clients	2-19
PCRFs	2-20
OCSs	2-21
IPFE	2-21
PCA Scalability	2-23
MP Growth	2-24



Mated Pair Growth	2-24
Adding a Mated PCA DSR to an Existing PCA DSR	2-25
Adding a Mated Pair of PCA DSRs	2-25
Adding a Mated Triplet of PCA DSRs	2-26
Small System Support	2-26
Redundancy	2-29
MP Server Redundancy	2-29
Site Redundancy	2-30
Data Redundancy	2-31
OAM Server Redundancy	2-32
IP Networking	2-34
PCA Routing	2-35
Ingress Routing	2-36
Egress Routing	2-37
PCA Data Auditing	2-40
PCA and Application Chaining	2-43
The Communication Agent	2-44
Diameter Routing and Communication with PCA	2-46
PCA and IDIH Metadata	2-48
PCA Capacity Constraints	2-55
PCA Assumptions and Limitations	2-56
Policy DRA Overview	
The Policy DRA Function	3-1
PCRF Pools and Sub-Pools Concepts and Terminology	3-2
Policy DRA Functions	3-13
Diameter Request Message Processing	3-14
Query Subscriber's Binding Status	3-14
PCRF Selection and Routing	3-16
Topology Hiding Process	3-17
Diameter Answer Message Processing	3-17
Subscriber Session and Binding Database Management	3-17
Subscriber Identification and Binding	3-18
Binding-Capable Sessions	3-19
Binding-Dependent Sessions	3-25
In-Session Message Processing	3-27
Topology Hiding	3-27
Session Integrity	3-29
Suspect Binding Removal	3-33



3

4 Online Charging DRA Overview

	Online Charging DRA Functions	4-1
	OCS Selection and Routing	4-1
	Single OCS Pool Mode	4-1
	Multiple OCS Pools Mode	4-2
	Regionalized Routing	4-3
	Session State Maintenance	4-5
	Gy/Ro Diameter Request Message Processing	4-7
	Session Initiation Request Message Processing	4-9
	In-Session Request Message Processing	4-10
	Event Request Message Processing	4-11
	Gy/Ro Diameter Answer Message Processing	4-11
	Session Initiation Answer Message Processing	4-13
	In-Session Answer Message Processing	4-13
	Event Answer Message Processing	4-13
	DRL-Initiated Answer Message Processing	4-14
0	Configuration Policy and Charging Configuration Overview	5-1
	Policy and Charging Configuration Overview	5-1
	NOAM and SOAM Configuration	5-2
	Pre-Configuration Activities	5-3
	System Topology	5-4
	Networks	5-4
	Services	5-11
	Places	5-13
	Server Groups	5-15
	Devices	5-21
	PCA Topology	5-30
	Identifying Place and Place Association information	5-30
	Identifying Resource Domain information	5-31
	Diameter Network Check	5-31
	Diameter Network Check for Policy DRA	5-32
	Diameter Network Check for Online Charging DRA	5-33
	Health Check	5-34
	Verifying Server status	5-34
	Logging all current alarms	5-35
	PCA Configuration	5-36



Place and Place Association Configuration	5-36
Places	5-36
Place Associations	5-38
Resource Domain Configuration	5-40
Resource Domains	5-40
PCA Routing of Diameter Messages	5-42
Inter DSR Routing	5-44
Routing for Gx RAR Messages (PDRA Generated)	5-44
Diameter Configuration for PCA	5-44
Policy DRA Configuration	5-48
PCRFs	5-48
Binding Key Priority	5-52
Network-Wide Options	5-53
Policy Clients	5-65
Site Options	5-68
SBR Databases	5-70
PCRF Pools	5-70
PCRF Sub-Pool Selection Rules	5-77
PCRF Pool to PRT Mapping	5-84
Error Codes	5-88
Suspect Binding Removal Rules	5-101
Access Point Names	5-105
General Options	5-113
Online Charging DRA Configuration	5-115
OCSs	5-115
CTFs	5-118
SBR Databases	5-121
Access Point Names	5-121
OCS Session State	5-129
Realms	5-131
Network-Wide Options	5-133
Error Codes	5-138
General Options	5-150
Alarm Settings	5-152
Alarm Settings elements	5-153
Defining Alarm Settings	5-155
Congestion Options	5-156
Congestion Options elements	5-156
Setting Congestion Options	5-158
nfiguration of Policy DRA Function on a Running DSR PCA System	5-159
Configuring new Policy DRA Sites	5-159
• • •	



	Configuring Policy DRA in existing Sites	5-159
	Configuring Policy DRA in existing Sites with scaling	5-159
	Configuration of Online Charging Function on a Running DSR PCA System	5-159
	Configuring new Online Charging DRA Sites	5-160
	Configuring Online Charging DRA in existing Sites	5-160
	Configuring Online Charging DRA in existing Sites with scaling	5-160
	Unconfiguration of Policy DRA Function from a Running DSR PCA System	5-160
	Unconfiguring Policy DRA	5-160
	Unconfiguration of Online Charging DRA Function from a Running DSR PCA System	5-163
	Unconfiguring Online Charging DRA	5-163
	Diameter Common Configuration for PCA	5-165
	Post-Configuration Activities	5-165
	Enable the PCA Application	5-165
	Setting General Options	5-166
	Enable SBR Databases	5-166
	Restart Process	5-166
	Enable Connections	5-167
	Perform Health Check	5-167
	Bulk Import and Export	5-168
6	Maintenance	
	Introduction	6-1
	Policy and Charging Maintenance	6-2
	Policy Database Query	6-2
	Policy Database Query elements	6-2
	Alarms, KPIs, and Measurements	6-3
	Policy and Charging Alarms and Events	6-3
	PCA KPIs	6-3
	Policy and Charging Measurements	6-4
	Overload Management	6-4
	Overload Controls	6-4
	Shutdown	6-8
	Diameter Maintenance and Status Data for Components, DSR Applications, and DA-MPs	6-10
	Backup and Restore for Policy and Charging Configuration Data	6-11



List of Figures

2-1	Example of Gy/Ro Event-Based Charging	2-9
2-2	Example of Gy/Ro Session-Based Charging	2-10
2-3	Sites, Mated Pairs, and Region	2-11
2-4	Policy DRA in Roaming Scenarios	2-12
2-5	Example PCA Mated Pair - Hosting Binding SBRs	2-14
2-6	Example PCA Mated Pair - Not Hosting Binding Policy SBRs	2-15
2-7	Policy Client, PCRF, and Site Relationships	2-16
2-8	Smallest Supported PCA Field Deployment	2-27
2-9	Smallest Supported PCA Mated Pair	2-28
2-10	Smallest Supported PCA Mated Triplets	2-28
2-11	Comparing 2-Site and 3-Site Redundancy	2-31
2-12	Binding Table Partitioning Across Server Groups	2-32
2-13	Multi-Table Resources	2-32
2-14	Data Merging – Normal Case	2-33
2-15	Data Merging – Redundant Site Failure	2-34
2-16	Communication between ComAgents, Policy DRA, and SBR	2-45
2-17	Request Processing at the Diameter Routing Function and PCA	2-46
2-18	Answer Processing at the Diameter Routing Function and PCA	2-47
2-19	PCA Generated Answer Routing	2-47
2-20	PCA Generated Request Routing	2-48
2-21	Event Diagram Trace – CCR Example	2-53
2-22	Update Request Policy DRA Example	2-54
2-23	Event Diagram – Hover (Mouse-over) Example	2-55
3-1	Relationship between APNs and PCRF Pools	3-3
3-2	Relationship between IMSIs and PCRF Pools	3-4
3-3	Multiple PCRF Pools	3-5
3-4	Multiple PCRF Versions in a PCRF Pool	3-7
3-5	PCRF Pools and Sub-Pools Routing Scenarios	3-8
3-6	Subscriber Key Usage	3-18
3-7	Find or Create a Binding	3-24
3-8	Binding Dependent Session Initiation Request Processing Overview	3-26
4-1	Local OCS Server Selection	4-2
4-2	Local OCS Server Pool Selection	4-3
4-3	Regionalized OCS Server Pool Selection	4-4
5-1	PCA Configuration Flow	5-2





List of Tables

2-1	Example of a Binding Key Priority Configuration	2-6
2-2	Server Group Functions	2-17
2-3	SBR Server Group Configuration and Data Redundance	2-17
2-4	Client Connection Capability	2-19
2-5	IP Traffic-to-Service Mapping	2-34
2-6	P-DRA Application Routing Table Configuration	2-37
2-7	Communication between the Diameter Routing Function and the DAI	2-46
2-8	PCA Metadata-Generating Events	2-49
2-9	Engineered Capacity Constraints	2-55
3-1	Policy DRA Terminology	3-9
3-2	Example Key Priority Configuration	3-19
3-3	Topology Hiding Scope Configuration	3-28
3-4	Policy DRA Error Scenarios for Session Integrity	3-30
3-5	Session Integrity Conditions and Policy DRA Reaction	3-31
4-1	Session State Configuration Settings	4-5
4-2	Session State Information	4-7
4-3	Diameter AVPs used by OC-DRA for Request Message Processing	4-8
4-4	OC-DRA Header Information	4-12
4-5	Diameter AVPs used by OC-DRA for Answer Message Processing	4-12
5-1	Insert Network Elements Elements	5-5
5-2	Networks Insert Elements	5-7
5-3	Core Services	5-12
5-4	Places Insert Elements	5-13
5-5	Server Groups Insert Elements	5-16
5-6	Server Groups Edit Elements	5-17
5-7	Devices Elements	5-21
5-8	Devices General Options	5-23
5-9	Places Insert Elements	5-37
5-10	Place Association Insert Elements	5-39
5-11	Add New Resource Domain Elements	5-41
5-12	PCRFs Screen Elements	5-49
5-13	Binding Key Priority Elements	5-53
5-14	Policy DRA Network-Wide Options Elements	5-54
5-15	Policy Clients Elements	5-66
5-16	Site Options Elements	5-69



5-17	PCRF Pooling Concepts	5-72
5-18	PCRF Pooling Configuration Summary	5-73
5-19	PCRF Pools Elements	5-75
5-20	PCRF Sub-Pool Selection Rules Elements	5-80
5-21	PCRF Pools to PRT Mapping Elements	5-86
5-22	PCA Error Conditions	5-89
5-23	Interfaces Supported for Each Error Code	5-99
5-24	Error Codes Elements	5-99
5-25	Suspect Binding Removal Rules Elements	5-101
5-26	Access Point Names Elements	5-107
5-27	General Options Elements	5-114
5-28	OCSs Elements	5-116
5-29	CTFs Elements	5-119
5-30	Access Point Names Elements	5-123
5-31	OCS Session State Elements	5-129
5-32	Realms Elements	5-131
5-33	Online Charging DRA Network-Wide Options Elements	5-134
5-34	PCA Error Conditions	5-139
5-35	Interfaces Supported for Each Error Code	5-148
5-36	Error Codes Elements	5-148
5-37	General Options Elements	5-151
5-38	Alarm Settings Elements	5-153
5-39	Congestion Options Elements	5-156
6-1	Policy Database Query Elements	6-3
6-2	PCA Default Overload Control Thresholds	6-5
6-3	Diameter Routing Function Message Handling Based on PCA Operational Status	6-6
6-4	PCA-SBR Stack Event Priorities	6-7
6-5	Stack Event Load Shedding	6-8
6-6	PCA Operational Status	6-9



1

Introduction

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

Revision History

Date	Description
July 2016	Initial release
July 2016	Updated based on support for U-SBR
October 2020	Added Enable Reroute and Reroute Peer Route Table Name fields to Site Options.

Overview

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

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

Scope and Audience

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

Manual Organization

This manual is organized into chapters:

- Introduction contains general information about the DSR documentation, the organization of this manual, and how to get technical assistance.
- Policy and Charging Introduction describes the topology, architecture, components, and functions of the Policy and Charging application and the Session Binding Repository (SBR).
- Policy DRA Overview describes an overview of the Policy DRA feature and includes information about important fundamental concepts, as well as high-level functionality, including Pools and Sub-Pools.
- Online Charging DRA Overview describes an overview of the Online Charging DRA feature and includes information about important fundamental concepts, as well as high-level functionality, including OCSs and CTFs

- Configuration describes configuration of PCA application components.
- Maintenance describes PCA Maintenance functions, and Diameter Maintenance functions that provide maintenance and status information for PCA and the SBR.

My Oracle Support

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

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

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

You are connected to a live agent who can assist you with My Oracle Support registration and opening a support ticket.

My Oracle Support is available 24 hours a day, 7 days a week, 365 days a year.



2

Policy and Charging Introduction

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

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

Overview

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

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

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

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

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

A set of PCA mated pairs and NOAM server pair/triplet is described as a PCA network. All clients and PCRFs/OCSs are reachable for Diameter signaling from any PCA DSR in the PCA network.

PCA is also designed to do several things:

- Reduce Diameter signaling latency where possible by:
 - Limiting the need to access off-board databases



Off-board in this context means on a serverserver separate from the server handling the Diameter signaling

- Limiting to a single WAN traversal to route a diameter message within the PCA network
- Optimization of the most frequent sunny day scenarios, possibly at the expense of less common, or rainy day, scenarios
- Provide server redundancy by supporting clusters of active DA-MP servers
- Provides site redundancy by supporting mated pairs of P-DRA DSRs, as well as provide 3-site redundancy by supporting mated triplets of P-DRA DSRs
- Provide triple data redundancy for subscriber binding data by having geographically dispersed active, standby, and spare copies of each binding record for mated pair configuration
- Provide quadruple data redundancy for subscriber binding data by having geographically dispersed active, standby, spare, and spare copies of each binding record for mated triplet configuration
- Support scalability of each DSR by the addition of DA-MP servers, as well as support network scalability by the addition of PCA sites
- Limit network configuration complexity by making use of naming conventions for clients and PCRFs/OCSs
- Facilitate troubleshooting of network-wide database accesses and Diameter signaling by including correlation information in logs and traces

The P-DRA Database

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

A high-level view of SBR Binding and Session databases consists of several pieces:

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



Bindings vs Sessions

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

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

Bindings

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

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

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

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

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

P-DRA then checks to determine whether a Sub-Pool exists by locating the PCRF Pool and the Origin-Host from the session initiation request via **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **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 using the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pool To PRT Mapping** screen on the SOAM GUI. The P-DRA application instructs the Diameter Routing Layer to use the PRT table associated with the PCRF Pool or Sub-Pool to route the request.



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

A certain order is used to search for an existing binding:

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

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

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

Sessions

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



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**, and then **Configuration**, and then **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 several reasons:

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

The Binding Database

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

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

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

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



Anchor Key

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

See Error Codes.

Alternate Keys

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

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

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

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

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

Table 2-1 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 2-1 Example of a Binding Key Priority Configuration

Priority	Key
1	IMSI
2	IPv6
3	MSISDN
4	<not configured=""></not>



The Priority order defines the order in which P-DRA looks for a given key type in a binding-dependent session initiating message. In the example in Table 2-1, P-A will look for keys in order and AVP:

- IMSI: Subscription-Id AVP with Subscription-Id-Type of END_USER_IMSI
- IPv6 Address: Framed-IPv6-Prefix AVP (only high order 64 bits used)
- 3. MSISDN: Subscription-Id AVP with Subscription-Id-Type of END USER E164

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

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

The Session Database

The Session database consists of 2 tables:

- A Session table
- A SessionRef table

Session

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

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

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

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

Session Reference

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



removed. IP Address records are removed when the only binding-capable session that referenced them is removed.

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

Binding Key Query Tool

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

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



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

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

The OC-DRA Database

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

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

Online Charging over Gy/Ro Reference Point

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



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

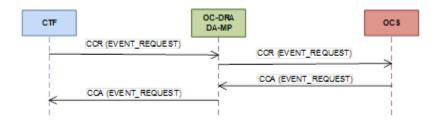
Basic scenarios are used for online charging:

- Event-Based Charging
- Session-Based Charging

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

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

Figure 2-1 Example of Gy/Ro Event-Based Charging



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



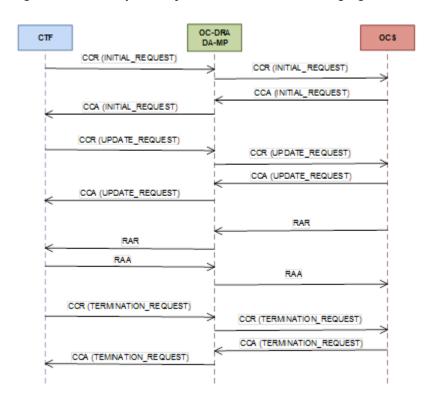


Figure 2-2 Example of Gy/Ro Session-Based Charging

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

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

Binding-independent Interface

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

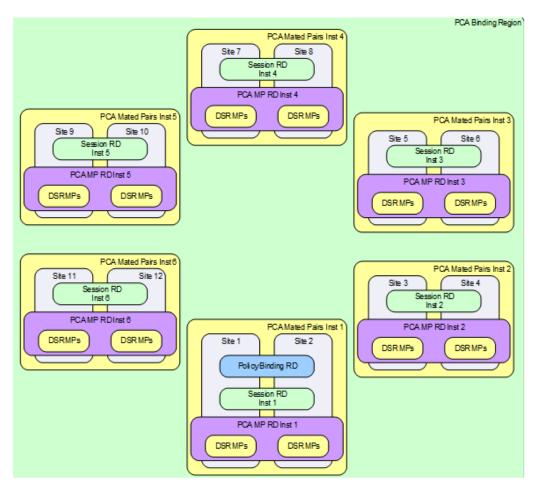
Deployment Topology

This section describes the makeup of a PCA network, regardless of its size. Figure 2-3 illustrates an example PCA network.



- A PCA Network can have up to 16 mated pairs or 32 sites, or can be as small as a single site. A PCA Network can also have up to 5 mated triplets.
- The PCA Binding Region provides the scope of the Policy Binding database.
 Binding records are accessible from every PCA DSR where the PDRA function is enabled in the Region.
- The Binding database need not be confined to a single mated pair or mated triplet.
- A PCA Mated Sites Place Association provides the scope for an instance of the Session database. Session records are accessible from each PCA DSR in the Mated Sites.
- Clients and PCRFs/OCSs have primary connections to their local PCA and secondary/tertiary connections to the mate(s) of their local PCA.
- PCA DSRs are connected to each other on the External Signaling Network. Each
 PCA site must be reachable from every other PCA site in the Region for Diameter
 signaling.
- The external signaling network handles stack events, database replication, and Diameter signaling. All three are required for the Diameter signaling to function correctly and with the required level of redundancy. Services (configured using the Configuration, and then Services screen) can be used to enforce separation of different types of traffic.

Figure 2-3 Sites, Mated Pairs, and Region





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

Policy DRA in Roaming Scenarios

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

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

Figure 2-4 illustrates an example Diameter network where the Policy DRAs are located in Home Access and Visited Access networks.

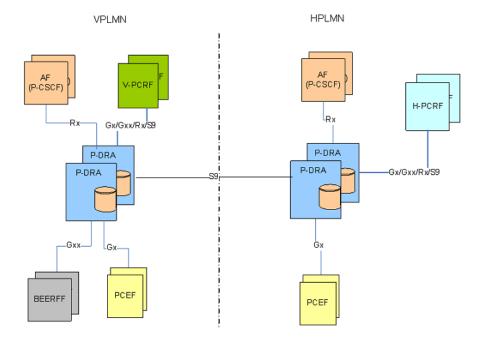


Figure 2-4 Policy DRA in Roaming Scenarios

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 QoC rules from the HPLMN to the VPLMN and transport the events occurred in the VPLMN to the HPLMN.

The S9 protocol makes use of exactly the same commands and messages as the Gx/Gxx protocols, except that a **V-PCRF** in VPLMN will provide an emergency treatment for any incoming CC-Request (INITIAL REQUEST) messages. This implies



that the Policy DRA does not check the existence of the Called-Station-ID AVP if the IMSI is missing in a CC-Request (INITIAL_REQUEST) over the S9 interface.

PCA Configurable Components

Key PCA configurable components include:

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

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

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

Figure 2-5 illustrates two PCA DSR Sites configured as a Mated Pair:

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



- Each Site has an SOAM Server Group consisting of 3 servers using the Active/ Standby/Spare redundancy model, allowing for Site redundancy.
- The PCA network has an NOAM Server Group consisting of 2 servers using the Active/Standby redundancy model. If NOAM site redundancy is desired, another pair of Disaster Recovery NOAM servers can be deployed at a different Site.
- Each Site has a number of DA-MP servers sufficient to carry the desired Diameter signaling load.
- Each Site has two pairs of IPFE servers one for use by Policy Clients and one for use by PCRFs. (IPFE is not required.)

Figure 2-5 Example PCA Mated Pair - Hosting Binding SBRs

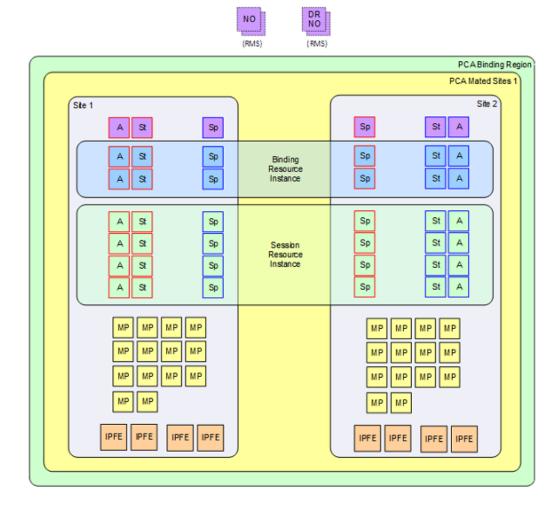


Figure 2-6 illustrates a possible configuration for additional mated pairs that do not host the Binding database:

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



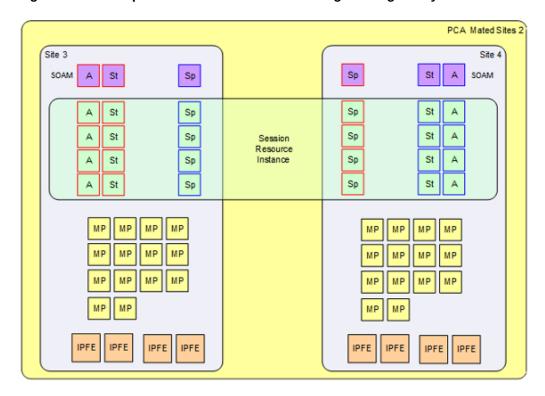


Figure 2-6 Example PCA Mated Pair - Not Hosting Binding Policy SBRs

Figure 2-7 illustrates example relationships between PCA DSR Sites and Policy Clients and PCRFs:

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



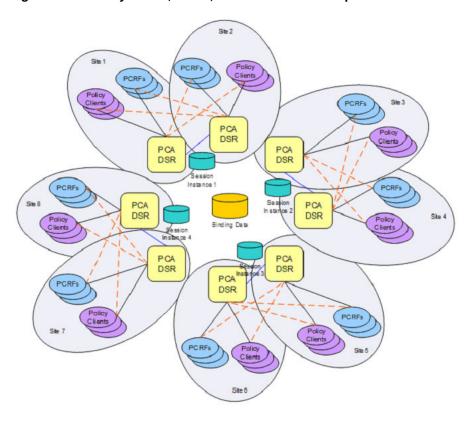


Figure 2-7 Policy Client, PCRF, and Site Relationships

Places

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

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

Place Associations

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

The PCA application defines two Place Association Types:

PCA Binding Region

As illustrated in Figure 2-3, the PCA application defines a Region to include all Sites that are part of the PCA network. This provides a scope for the Binding database, which is accessible to all PCA Sites in the PCA network.

PCA Mated Sites



As illustrated in Figure 2-3, pairs of PCA Sites are grouped together. Each Place Association with a PCA Mated Sites type includes exactly 2 sites. PCA Mated Sites has several attributes:

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

Server Groups

The PCA application makes use of several different types of Server Groups, as defined in Table 2-2.

Table 2-2 Server Group Functions

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

SBR Type

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

Each SBR Server Group consists of one to four servers, depending on the type of deployment. Table 2-3 describes the supported configurations for SBR Server Groups. See Redundancy for details on policy data redundancy.

Table 2-3 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.



Table 2-3 (Cont.) SBR Server Group Configuration and Data Redundance

# of Servers	Redundancy	Typical Use
3	Active/Standby/Spare	Mated Pair deployments to avoid a single-server failure from causing Session access requests to be routed to the mate Site. This is the target for large deployments. New sessions are equally distributed across all Session SBR Server Groups in the mated pair, meaning that ~50% of the Session accesses will be routed across the WAN.
		Note : SBR Server Groups must be configured with two WAN replication channels.
4	Active/Standby/Spare/Spare	Mated Triplet deployments where service is kept alive even two of the three sites with the service are down.

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

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

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

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

Resource Domains

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

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

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

Clients

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

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

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

How many connections a **Client**might initiate towards the PCA and how those connections are used are in customer control. The capabilities of the client, however, affect the functionality of the solution as shown in Table 2-4.

Table 2-4 Client Connection Capability

Number of Connections Supported by Policy and Charging Client

2

(per Diameter host) Effect on Solution Capability

- Site Redundancy cannot be supported.
 - Diameter signaling throughput is limited to the capacity of the connection.
 - Extra latency to reconnect in the event of a connection drop.
 - Site Redundancy supported if secondary connection is configured to connect to PCA mate site.
 - If both connections go to a single site and the policy and charging client has the capability to use both connections simultaneously, Diameter signaling throughput may be doubled vs. only one connection.
 - This configuration requires multiple Diameter connections to a single Diameter host something that is not supported by RFC 6733, but which many vendors support to allow capacity beyond what a single connection can support.
 - Extra latency is avoided in the event of a single connection drop because the other connection can be used without waiting for reconnect and Capabilities Exchange.
 - PCA Mated Triplet is not supported, since the client would become isolated if the two DSRs with which it has connections fail.



Table 2-4 (Cont.) Client Connection Capability

Number of Connections Supported by Policy and Charging Client (per Diameter host)	Effect on Solution Capability
3	 PCA Mated Triplet is supported if the client has one connection to each DSR in the triplet.
	 PCA Mated Pair is supported if and only if the client has at least one connection to each DSR in the pair. The extra connection to one of the two DSRs provides the opportunity for higher availability and throughput.
>3	There are many scenarios possible, depending on the capabilities of the policy and charging client. For example, there might be two connections to the primary PCA (for capacity) and two to each mate PCA.

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

PCRFs

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

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

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

- A PCRF that shares state across different Diameter hostnames.
 - Each Diameter hostname can all support Gx, Gxx, S9, Gx-Prime and Rx
 Diameter interfaces. This type of PCRF is supported by PCA.
 - Each hostname has a different connection for each different interface type.
 This type of PCRF is supported by PCA.
 - There is a different Diameter hostname for each connection for a specific Diameter interface. All of the Diameter hostnames share state. This type of PCRF is supported by PCA.
 - There are different Diameter hostnames for different policy client vendors.
 Policy state is shared across the Diameter hostnames, but origin based routing is required to select a set of PCRFs for distribution of the initial binding



depending on the policy client type. This type of PCRF is supported by PCA, but requires use of Diameter Routing Function PCRF selection as described in PCRF Selection for New Bindings.

- There is a different Diameter hostname for each connection. This type of PCRF is supported by PCA, but requires use of Diameter Routing Function PCRF selection based on the vendor type of the policy client as described in PCRF Selection for New Bindings.
- A PCRF that has one Diameter hostname, but supports a number of connections to that hostname using different IP addresses.
 Each connection can support Gx, Gxx, S9, Gx-Prime, and Rx Diameter interfaces.
 This type of PCRF is supported by PCA.

OCSs

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

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

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

IPFE

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

Various deployment scenarios involving IPFE are possible:

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



- An IPFE configuration in which Policy and Charging Clients are connected to a PCA mated pair, but PCRFs and/or OCSs are not capable of initiating connections to the PCA. Each Policy Client has a primary connection to one PCA site and a secondary connection to the mate site. For example:
 - Two PCA DSR sites, each with a pair of IPFE servers and 4 DA-MP servers.
 - Three Policy and Charging Clients with a primary connection to PCA DSR Site 1 and secondary connections to PCA DSR Site 2.
 - Three Policy and Charging Clients with a primary connection to PCA DSR Site
 2 and secondary connections to PCA DSR Site 1.
 - Two PCRFs or OCSs with primary connections to PCA DSR Site1 and secondary connections to PCA DSR Site 2.
 - Two PCRFs or OCSs with primary connections to PCA DSR Site2 and secondary connections to PCA DSR Site 1.
 - One IPFE at PCA DSR Site 1 is primary for TSA1. The other IPFE is standby for TSA1.
 - One FABR at PCA DSR Site 2 is primary for TSA2. The other IPFE is standby for TSA2.
- A single site PCA in which a single IPFE pair is used for both Policy and Charging Clients and PCRFs and/or OCSs. The use of IPFE for PCRFs is possible only if the PCRF can be configured to initiate connections towards the PCA. Some customers refer to an IPFE used by PCRFs as an IP Back-End, or IPBE, although there is no difference between an IPBE and an IPFE from a software or configuration perspective. For example:
 - One pair of IPFE servers, each server supporting two TSAs
 - Four Policy and Charging Clients connect to TSA1 with their secondary connection going to TSA3, or vice-versa.
 - The PCRFs or OCSs connect to TSA2 with their secondary connection going to TSA4, or vice-versa.
 - Six PCA MP servers, each capable of hosting connections from Policy and Charging Clients and PCRFs or OCSs
 - One IPFE server is primary for TSA1 and TSA2, and standby for TSA3 and TSA4.
 - The other IPFE server is primary for TSA3 and TSA4, and standby for TSA1 and TSA2.
- A single site PCA in which IPFE is used for both Policy Clients and PCRFs. In this
 case, two pairs of IPFE servers are deployed in order to support high Diameter
 signaling bandwidth. For example:
 - Two pairs of IPFEs, each supporting a two TSAs
 - The Policy and Charging Clients connect to either TSA1 or TSA2, with their secondary connection going to the other TSA.
 - The PCRFs or OCSs connect to either TSA3 or TSA4, with their secondary connection going to the other TSA.
 - Eight PCA DA-MPs, each capable of hosting connections from Policy and Charging Clients and PCRFs or OCSs



- One IPFE server on the Policy and Charging Client side is primary for TSA1 and standby for TSA2. The other IPFE server is primary for TSA2 and standby for TSA1.
- One IPFE server on the PCRF or OCS side is primary for TSA3 and standby for TSA4. The other IPFE server is primary for TSA4 and standby for TSA3.
- A PCA mated pair configured with an IPFE for Policy Clients and a separate IPFE for PCRFs. The Policy and Charging Clients and PCRFs have a primary connection to their local PCA DSR and a secondary connection to the mate PCA DSR. For example:
 - Two PCA DSR sites, each with a two pairs of IPFE servers and 6 DA-MP servers
 - Three Policy and Charging Clients with a primary connection to PCA DSR Site 1 and secondary connections to PCA DSR Site 2.
 - Three Policy and Charging Clients with a primary connection to PCA DSR Site
 2 and secondary connections to PCA DSR Site 1.
 - Two PCRFs or OCSs with primary connections to PCA DSR Site1 and secondary connections to PCA DSR Site 2.
 - Two PCRFs or OCSs with primary connections to PCA DSR Site2 and secondary connections to PCA DSR Site 1.
 - One IPFE on the Policy and Charging Client side at PCA DSR Site 1 is primary for TSA1. The other IPFE is standby for TSA1.
 - One IPFE on the Policy and Charging Client side at PCA DSR Site 2 is primary for TSA3. The other IPFE is standby for TSA3.
 - One IPFE on the PCRF or OCS side at PCA DSR Site 1 is primary for TSA2.
 The other IPFE is standby for TSA2.
 - One IPFE on the PCRF OCS side at PCA DSR Site 2 is primary for TSA4. The other IPFE is standby for TSA4.

PCA Scalability

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

For large systems, PCA can scale up as follows:

- 16 mated pairs of PCA DSRs (32 sites) or up to 5 mated triplets (15 sites) per network
- Three enclosures per PCA DSR site using half-height servers Each enclosure has 16 half-height slots.
- Two pairs of IPFE servers per PCA DSR
- Sixteen DA-MP servers per PCA DSR

Figure 2-3 illustrates a sample PCA network consisting of 6 mated pairs, or 12 sites with components that must be configured as follows:

 An instance of a Site (Place with type Site) is created for each physical location of a PCA DSR.



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

MP Growth

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

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

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

Mated Pair Growth

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

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

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



Adding a Mated PCA DSR to an Existing PCA DSR

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



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

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

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

Adding a Mated Pair of PCA DSRs

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

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



Adding a Mated Triplet of PCA DSRs

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

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

Small System Support

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

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

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

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

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



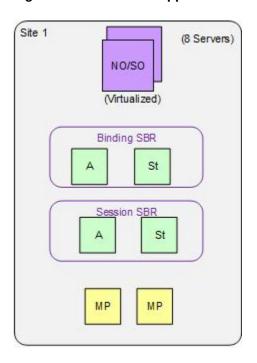


Figure 2-8 Smallest Supported PCA Field Deployment

The smallest supported Mated Pair of PCA DSRs, illustrated in Figure 2-9, has certain characteristics:

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

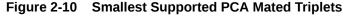


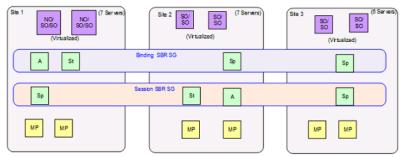
Site 1 (7 Servers) Site 2 NO (7 Servers) S0/ SO SO SO (Virtualized) (Virtualized) Binding SBRSG St Sp Session SBR SG St Sp ΜP MP MΡ

Figure 2-9 Smallest Supported PCA Mated Pair

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

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







Redundancy

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

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

MP Server Redundancy

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

- DA-MP Multi-Active Cluster
 - PCA MP servers are deployed using an Active/Active redundancy model. This means that every DA-MP actively processes Diameter signaling. In order to avoid single points of failure, a minimum of two DA-MPs must be deployed (except for lab and trial deployments, where one DA-MP is acceptable). DA-MPs at a given site must be configured such that loss of a single DA-MP will not cause the remaining DA-MP servers to go into signaling overload.
- SBR Active Only
 An SBR (either Session or Binding) can be deployed in simplex redundancy mode only for labs or trials. Otherwise this configuration represents a single point of failure for the SBR database being hosted by the Active-only Server Group. In this configuration, the SBR Server Groups consist of a single Server.
- SBR Active/Standby
 - The Active/Standby redundancy model should be used for single site PCA deployments, or for multi-site deployments when site redundancy is not important. In this configuration, the SBR Server Groups consist of two servers. On system initialization, one of the two servers in each SBR Server Group will be assigned the Active role and the other the Standby role. These roles will not change unless a failure or maintenance action causes a switch-over. For Active/Standby Server Groups, switch-overs are non-revertive, meaning that recovery of a formerly Active server will not cause a second switch-over to revert the Active role to that server.
- SBR Active/Spare
 - The Active/Spare redundancy model can be used for mated pair deployments in which it is acceptable for traffic to move from one site to the mate site on failure of a single server. In this configuration, the SBR Server Groups consist of two servers with one marked as Preferred Spare. On system initialization, the server not marked as Preferred Spare will be assigned the Active role and the other the Spare role. These roles will not change unless a failure or maintenance action causes a switch-over. For Active/Spare Server Groups, switch-overs are revertive, meaning that recovery of a formerly Active server will cause a second switch-over to revert the Active role to that server.
- SBR Active/Standby/Spare
 - The Active/Standby/Spare redundancy model should be used for PCA mated pair deployments in which site redundancy is desired. In this configuration, each SBR Server Group is configured with two servers at one site and the third at the mate site. The server at the mate site is designated in the Server Group configuration as Preferred Spare. On system initialization, one of the two servers that are located



at the same site will be assigned the Active role and the other the Standby role. The server at the mate site will be assigned the Spare role (as was preferred). If the Active server can no longer perform its function due to failure or maintenance, the Standby Server will be promoted to Active. Only if both Active and Standby servers at a site are unable to perform their function will the Spare server at the mate site be promoted to Active. Active and Standby role changes within a site are non-revertive, but if the server at the mate site is Active and one of the other servers recovers, a switch-over will occur to revert the Active role back to the site with two servers.

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

the standby server will be promoted to active. Only if both active and standby servers at a site are unable to perform their function will a spare server at a mate site be promoted to active. Active and standby role changes within a site are non-revertive, but the server at a mate site is active and one of the servers at the site with two servers recovers, a switch-over will occur to revert the active role back to the site with two servers.

Site Redundancy

2-Site Redundancy

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

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

3-Site Redundancy

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

- PCA is deployed on at least one mated triplet of PCA DSRs.
- Clients and PCRFs/OCSs connect to all sites in the mated triplet.
- SBR Server groups use the Active/Standby/Spare/Spare redundancy model.



- System OAM server groups use the Active/Standby/Spare/Spare redundancy model.
- Diameter Agent MP servers are recommended to be engineered at 26%%
 capacity across the mated triplet, for example, if two sites fail, then each remaining
 DA-MP operates at 80% capacity.

Comparing 2-Site and 3-Site Redundancy

To achieve the same redundant capacity of a mated pair in a mated triplet, multiple things need to be added:

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

PA1<Policy and Charging DRA Mated Sites PA1<Policy and Charging DRA Mated Sites: Place1<Site3 Place1<Site> Place2<Site3 Place3<Site3 Sp SOAM A St Sp (2) Sp Sp A St RD2 Policy Session RD2 Policy Session A St A St Sp Sp (4) Sp Sp A St

Figure 2-11 Comparing 2-Site and 3-Site Redundancy

Data Redundancy

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

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



Figure 2-12 illustrates how a given Policy Binding table might be partitioned across four SBR Server Groups in a Policy Binding resource domain.

Figure 2-12 Binding Table Partitioning Across Server Groups

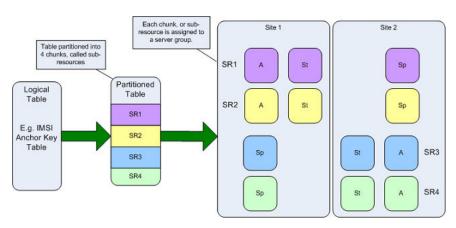
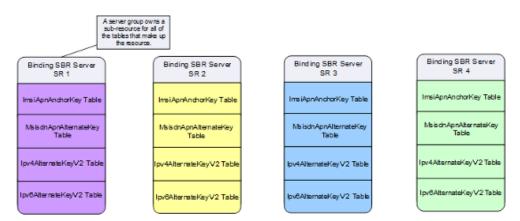


Figure 2-13 illustrates how each SBR Server Group hosts a partition of several tables. Only the Active server within each server group can write to the database. The Standby and Spare servers replicate only actions (adds, changes, and deletes) performed by the Active server.

Figure 2-13 Multi-Table Resources



OAM Server Redundancy

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

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

Active/Standby NOAM and Active/Standby DR NOAM
 The NOAM servers are deployed using the active/standby redundancy model at
 one of the sites in the PCA network. If site redundancy is desired, an optional pair

of Disaster Recovery (DR) NOAM servers can be deployed at a different site. The DR NOAM servers are used only if manually brought into service following loss of the site where the original NOAM pair was located.

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

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

Figure 2-14 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.

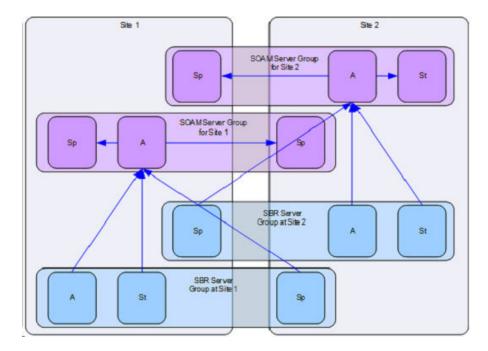


Figure 2-14 Data Merging - Normal Case

Figure 2-15 illustrates how a site failure affects merging of alarms, events, and measurements. When Site 2 fails, the servers at Site 1 that were marked as Preferred Spare are promoted to Active. The MP server that is now Active for the SBR Server Group for Site 2 will start merging its data to the SOAM server that is now Active for the SOAM Server Group for Site 2.



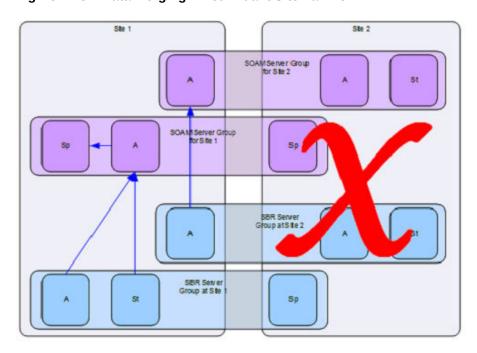


Figure 2-15 Data Merging – Redundant Site Failure

IP Networking

Diameter's flexibility 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 2-5 illustrates a possible Services configuration for enabling signaling traffic from OAM traffic. In Table 2-5, 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 2-5 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



Table 2-5 (Cont.) IP Traffic-to-Service Mapping

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

PCA Routing

Routing of Diameter messages in a DSR with PCA activated is divided into ingress routing and egress routing. Ingress routing includes routing of Diameter requests and Diameter answers to the PCA application from a remote peer. Egress routing includes

routing of Diameter requests and Diameter answers from PCA towards a remote peer. A remote peer can be a policy and charging client, a PCRF, an OCS or another DSR. Diameter requests can be initiated by both clients and PCRFs/OCSs.

Ingress Routing

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

Requests

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

- 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.
- Does the request match a rule in the Application Routing Table (ART)?If no rule is matched, the request is not routed to any DSR application. Otherwise the next condition is checked.
- 3. If the request matches an ART rule, is the application operational status for this DA-MP server set to Available? If the DSR Application is not Available, then the Unavailability action is performed by Diameter. For PCA, the Unavailability action for Request is Reject, which means PCA reject the Diameter Request messages by generating and sending an Answer message with a failure response with Result-Code AVP set to value configured for the error condition PCA Unavailable or Degraded.
 - If the DSR Application is Available, then Diameter routes the Request to the DSR Application specified in the matching Application Routing Rule.

Ingress Requests are examined by Diameter to determine whether they should be routed to a DSR Application. The rules for deciding how to route ingress requests are defined in an Application Routing Table, or ART. Table 2-6 describes the expected configuration of Application Routing Rules for PCA. These rules will cause every Request that includes one of these values in the Application-Id in the Diameter Header to be routed to the PCA application. Some of these rules can be omitted, depending on which interfaces are used for PCA.

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



The Application Name will always be PCA for the Policy and Charging application.
 PCA will show up in the Application Name drop-down only if the PCA application has been activated

Table 2-6 P-DRA Application Routing Table Configuration

Rule Name	Priority	Conditions	Application Name
PcaGxRule	5	Appld Equal 16777238 - 3GPP Gx	PCA
PcaGxxRule	5	Appld Equal 16777266 - 3GPP Gxx	PCA
PcaRxRule	5	Appld Equal 16777236 - 3GPP Rx	PCA
PcaGxPRule	5	Appld Equal 16777238 - 3GPP Gx- Prime	PCA
PcaS9Rule	5	Appld Equal 16777267 - 3GPP S9	PCA
PcaRoGyRule	5	Appld Equal 4 - 3GPP	PCA

Answers

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

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



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

Note:

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

Egress Routing

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

PCRF Selection for New Bindings

For the P-DRA function, when a binding capable session initiation message (CCR-I) arrives for an IMSI that is not already bound to a PCRF, the P-DRA function selects a PCRF from the list of adjacent PCRFs configured on **Policy and Charging**, and then **Configuration**, and then **PORFs** screen. This list of PCRFs



generally contains only PCRFs that are local to the site with the P-DRA node. PCRFs that are local to the PCA node's mate should generally not be included. The reason to include only local PCRFs is to avoid the extra latency associated with selection of a PCRF separated across a WAN from the policy client that initiated the session.

If the PCRF has different hostnames for different 3GPP interfaces (for example, Gx, Rx, Gxx, S9), only the binding capable hostnames should be configured on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRFs** screen.

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

PCRF Selection for Existing Bindings

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

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

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

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

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

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

Routing In-Session Messages Without Topology Hiding

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



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

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

Routing In-Session Message with Topology Hiding

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

Use of topology hiding is expensive in terms of the increased stack event processing required and the increased latency required to lookup the bound PCRF in the session record. For these reasons, topology hiding should be scoped as narrow as possible. For example, if topology should be hidden from only a few policy clients, choose the per policy client topology hiding scope instead of choosing to hide topology from all policy clients.

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

OCS Selection and Routing

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

- Single OCS Pool Mode
- Multiple OCS Pools Mode

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

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



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

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

Naming Conventions and Hierarchical Routing

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

PCA Data Auditing

P-DRA Binding/Session Database

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

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

To limit the effects of stale Binding and Session records, all SBRs that own an active part of the database continually audit each table to detect and remove stale records. The audit is constrained by both minimum and maximum audit rates. The actual rate



varies based on how busy the SBR server is. Audit has no impact on the engineered rate of signaling.

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

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

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

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

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

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

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



Note:

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

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

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

PCA Configuration Database

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

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

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

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



Note:

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

OC-DRA Session Database

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

Stale Gy/Ro sessions can occur for various reasons:

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

PCA and Application Chaining

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

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



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

P-DRA may receive the incoming Diameter request messages over binding capable interfaces (for example, Gx and Gxx) or over binding dependent interfaces (for example, Rx and Gx-Prime). The RBAR and PCA application chaining function is applicable ONLY on binding capable interfaces for the P-DRA feature, and binding independent interface (Ro/Gy) for the OC-DRA feature, but NOT on binding dependent interfaces. Specifically, the Diameter request messages over binding dependent interfaces (Rx or Gx-Prime) intending for being processed by P-DRA should never be routed to RBAR for address resolution.

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

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

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



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

The Communication Agent

The Communication Agent (ComAgent) enables reliable communication between Policy and Charging DRA and SBRs and among SBRs in a scalable and high

available PCA network. Figure 2-16 depicts the communication paths between the Policy and Charging DRA, the SBR, and their ComAgents, and the communication paths between the ComAgents.

Note:

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

Inter-ComAgent signaling PCA (SBR)-ComAgent PCA (SBR) -ComAgent Interface Interface SBR MP SBR MP PCA PCA SBR MF DA-MP SBR MP DA-MP SBR SBR MP Diameter signaling Signaling Network PCRF POPCEF

Figure 2-16 Communication between ComAgents, Policy DRA, and SBR

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

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

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

You can view and manage these connections using the ComAgent Connection Status GUI at the NOAM: **Communication Agent**, and then **Maintenance**, and then

Connection Status. There is also a ComAgent HA Service Status, which can be obtained from the **Communication Agent**, and then **Maintenance**, and then **HA Services Status** screen.

Diameter Routing and Communication with PCA

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

Table 2-7 describes two functions for communication between the Diameter Routing Function and the DAI.

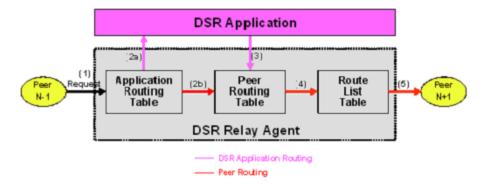
Table 2-7 Communication between the Diameter Routing Function and the DAI

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

Request Routing

Figure 2-17 shows the case where Diameter Request messages are routed from the Diameter Routing Function to the PCA based on the configured Application Routing Rule, and routed from the PCA to the Diameter Routing Function, all using the Application-Data function. The PCA will return the Request to the Diameter Routing Function for Peer Routing Rule processing and routing.

Figure 2-17 Request Processing at the Diameter Routing Function and PCA



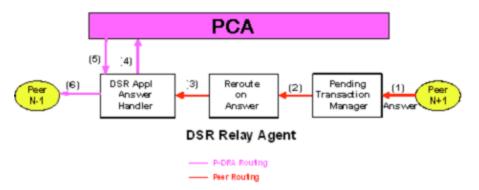
Answer Routing

When the PCA forwards a Request message to the Diameter Routing Function for routing, it must inform the Diameter Routing Function how to process the corresponding Answer. It can inform the Diameter Routing Function either to route the Answer to the PCA or to route the Answer to the downstream Peer without involving the PCA. Figure 2-18 shows the case where an Answer is transmitted back to the PCA. After the PCA completes processing of the Answer, it will send it to the Diameter



Routing Function for transmission to the Diameter Transport Function so that it can be routed to the downstream Peer.

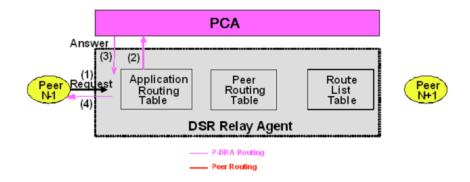
Figure 2-18 Answer Processing at the Diameter Routing Function and PCA



PCA Generated Answer

In some cases, the PCA needs to generate an Answer message in response to an incoming Request. For example, the Policy DRA function cannot find a PCRF to route the Request message to. Figure 2-19 shows the Diameter Routing Function routing for this scenario.

Figure 2-19 PCA Generated Answer Routing



PCA Generated Request

In some cases, the PCA needs to generate Diameter Requests. Figure 2-20 shows the Diameter Routing Function routing for this scenario.



DSR Application (1) Application Peer Route (2) Routing Routing List N+1 Table Table Table **DSR Relay Agent** DSR Application Routing Peer Routing

Figure 2-20 PCA Generated Request Routing

Policy DRA Function Use Cases

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

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

PCA and IDIH Metadata

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

PCA will record the Application-specific metadata events described in Table 2-8.



Table 2-8 PCA Metadata-Generating Events

Event	Instance Data	When Recorded	
PCA Function Invoked	PCA Function Name	When an application with multiple functionality receives an ingress Diameter message (including both requests and answers) and routes it to one of its functions for processing.	
		Note : This metadata is only recorded if the application function is enabled and available processing messages.	
PCRF Pool Selected	PCRF Pool NamePCRF Sub-Pool nameSub-Pool selection rule name	When P-DRA receives a binding- capable session initiation request	
Binding Query Request Sent	Anchor keyAPN namePCRF Pool nameSession reference	When P-DRA sends a Find or Create Binding stack event to a SBR	
Binding Query Result Received	 SBR IP Address (for example, 10.240.55.25) Result code Binding state Binding master Master session reference PCRF FQDN Suspect duration 	When P-DRA receives a Find or Create Binding Result stack event from a SBR	
Topology Hiding Applied	N/A	When P-DRA receives a bind- capable or binding-dependent session initiation/in-session answer destined for a peer for which topology hiding is configured	
Create Session Request Sent	 Session ID Session reference Anchor key MSISDN IPv4/IPv6 address key PCRF FQDN 	When P-DRA sends a Create Session stack event to a SBR	
Create Session Result Received	 SBR IP Address (for example, 10.240.55.25) Result code 	When P-DRA receives a Create Session Result stack event from a SBR	
Update Binding Request Sent	OperationAnchor keyFinal PCRF FQDNSession reference	When P-DRA sends an Update Binding stack event to a SBR	
Update Binding Result Received	 SBR IP Address (for example, 10.240.55.25) Result code 	When P-DRA receives an Update Binding Result stack event from a SBR	



Table 2-8 (Cont.) PCA Metadata-Generating Events

-	Luda and Buta	When Breeded
Event	Instance Data	When Recorded
Find Binding Request Sent	Key typeKey valueAPN name	When P-DRA sends a Find Binding stack event to a SBR
Find Binding Result Received	 SBR IP Address (for example, 10.240.55.25) Result code IMSI PCRF FQDN 	When P-DRA receives a Find Binding Result stack event from a SBR
Refresh Session Request Sent	Session ID	When P-DRA sends a Refresh Session stack event to a SBR
Refresh Session Result Received	SBR IP Address (for example, 10.240.55.25)Result code	When P-DRA receives a Refresh Session Result stack event from a SBR
Delete Session Request Sent	Session ID	When P-DRA sends a Remove Session stack event to a SBR
Delete Session Result Received	 SBR IP Address (for example, 10.240.55.25) Result code Session reference PCRF FQDN Anchor key MSISDN key IPv4/IPv6 key 	When P-DRA receives a Remove Session Result stack event from a SBR
Find Session Request Sent	Session ID	When P-DRA sends a Find Session stack event to a SBR
Find Session Result Received	 SBR IP Address (for example, 10.240.55.25) Result code Session reference PCRF FQDN 	When P-DRA receives a Find Session Result stack event from a SBR
Remove Suspect Binding Request Sent	Anchor keyPCRF FQDN	When P-DRA sends a remove Suspect Binding stack event to a SBR
Remove Suspect Binding Result Received	SBR IP Address (for example, 10.240.55.25)Result code	When P-DRA receives a Remove Suspect Binding Result stack event from a SBR
Session Release Initiated	Application Name	When an Update Binding request, a Create Session request or, a Create Alternate Key request fails
Session Query Initiated	Application Name	When a stale Gx session is detected by a SBR
Routing Exception	 Routing Exception Type (for example, SBR Congestion) Routing Exception Action (for example, Abandon Request) 	After any routing exception is encountered



Table 2-8 (Cont.) PCA Metadata-Generating Events

Event	Instance Data	When Recorded	
SBR Request Failure	 After any routing exception is encountered Resource name Sub-resource ID Failed Request Name 	When a PCA Function fails to send a request to the SBR	
SBR Response Timeout	Resource nameSub-resource ID	When a PCA Function times out waiting to receive a response from a SBR for a previous request	
Routing Error Indication Received	Routing Error	When a PCA Function initiates a Diameter request (Session Release RAR) that is rejected by DRL due to a routing error.	
		Note : The Routing Error recorded is the Error-Message AVP value of the Answer message initiated by DRL.	
Create OC Session Request Sent	 Session ID CTF Realm CTF FQDN OCS Realm OCS FQDN Subscriber ID APN Name 	When OC-DRA sends a Create OC Session stack event to the Session SBR	
Create OC Session Result Received	SBR IP AddressResult Code	When OC-DRA receives a Create OC Session Result stack event from the Session SBR	
Find and Refresh OC Session Request Sent	Session ID	When OC-DRA sends a Find and Refresh OC Session stack event to the Session SBR	
Find and Refresh OC Session Result Received	 Session ID Result Code CTF Realm CTF FQDN OCS Realm OCS FQDN Subscriber ID APN Name 	When OC-DRA receives the Find and Refresh OC Session Result stack event from the Session SBR	
Find and Remove OC Session Request Sent	Session ID	When OC-DRA sends a Find and Remove OC Session stack event to the Session SBR	
Find and Remove OC Session Request Received	 SBR IP Address Result Code CTF Realm CTF FQDN OCS Realm OCS FQDN Subscriber ID APN Name 	When OC-DRA receives the Find and Removed OC Session Result stack event from the Session SBR	



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

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

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

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

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



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



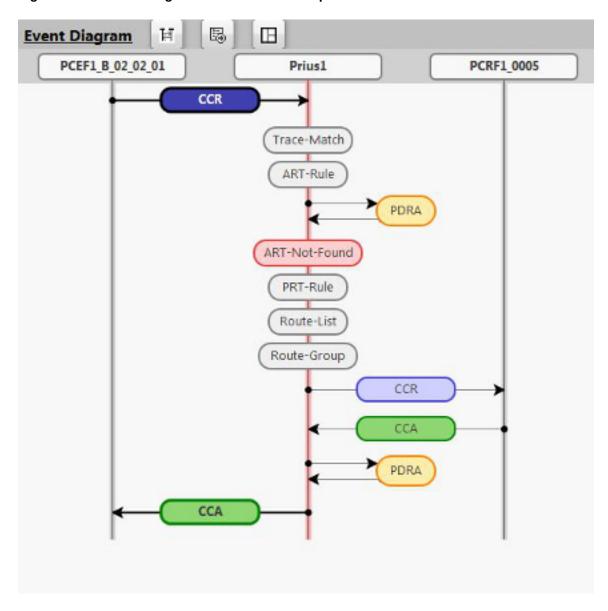


Figure 2-21 Event Diagram Trace – CCR Example

Figure 2-22 Update Request Policy DRA Example

	nt List 😩 🙈 💟	X H G D	
Rec#		Event Type	Event Data
1	03/07/2014 09:47:30:000		
2	•	Trace Match	INIT, TTP-IR
3		ART Rule Match	rule=Gx, table=GXRule1
4		Routing To App	PDRA
5	-	PCRF Paal Selected	PCRF Pool Name=Default PCRF Sub-Pool Name= Sub-Pool S
6		Binding Query Request Sent	Anchor Key=199754193097563 APN Name=APN1-B-02-02-01
7		Binding Query Result Received	p5BR IP Address=10.240.142.8 (IPv4) Binding State=Early Bind
В	-	App To Routing	requestRoutingAction=route via PRT rule overrideOriginRequ
9		ART Rule Not Found	GXRule1
10		PRT Rule Match	rule=Default_catch_all_table=Default
11		Route List Selected	Default_RL1
12	,	Route Group Selected	Default_RG1
13	03/07/2014 09:47:30:000	Message Sent	
14	03/07/2014 09:47:30:000	Message Received	2001 - DIAMETER_SUCCESS
15		Routing To App	PDRA
16		Update Binding Request Sent	Operation=Update PCRF Anchor Key=199754193097563 Fina
17	-	Create Session Request Sent	Session Reference-f25eb553805e22c368032d00 Anchor Key-
18		App To Routing	requestRoutingAction=unknown routing action overrideOrigin
19	03/07/2014 09:47:30:000	Message Sent	2001 - DIAMETER SUCCESS
20		Create Session Result Received	pSBR IP Address=172.18.24.14 (IPv4) Result Code=Success
21		Update Binding Result Received	p5BR IP Address=10.240.142.8 (IPv4) Result Code=Success



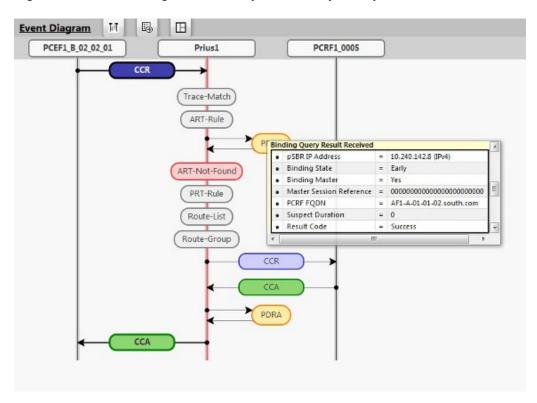


Figure 2-23 Event Diagram – Hover (Mouse-over) Example

PCA Capacity Constraints

PCA has engineered capacity constraints:

Table 2-9 Engineered Capacity Constraints

Constraint Name	Value
Maximum managed	objects per Network
Local PCRFs per Site	5000
APNs per Network	2500
PCRF Pools per Network	7
PCRF Sub-Pools per Network	14
Topo Hiding Policy Clients per Site	1000
Binding Server Groups per Network	0-8
Session Server Groups per Mated Pair	0-8
Binding Servers per Sg	8
Session Servers per Sg	8
Maximum Number of SBR Databases	65
Maximum Number of SBR Database Resizing Plans	33
Maximum Number of SBR Data Migration Plans	64
Maximum Number of Binding Region Place Associations	1



Table 2-9 (Cont.) Engineered Capacity Constraints

Constraint Name	Value
Maximum Number of Binding Resource Domains	2
Maximum Number of Mated Sites Place Associations	64
Maximum Number of PCA Sites	32
Maximum Number of Suspect Binding Rules Per Site	50
Cardinality relationships between managed obj	ects
Sub-Pool Rules per PCRF Pool	10
Maximum Number of Server Groups Per SBR Database	8
Maximum Number of Servers Per SBR Database	4

PCA Assumptions and Limitations

PCA has numerous assumptions and limitations:

Assumptions

- For the P-DRA function, the anchor key that identifies all subscribers in the PCA network is the IMSI, while MSISDN is the key that identifies the subscribers for the OC-DRA Function
- All Gx and Gxx session initiating Diameter messages will always include the IMSI.
 The only exception is emergency calls from devices with no SIM card (UICC-less).
- Messages sharing a common Diameter Session-Id will never arrive out of sequence.
- PCRF/OCS names and client names start with characters that can be used to identify which PCA DSR hosts the primary connection to that equipment. This greatly simplifies routing configuration for the PCA network. The network can be configured to work without such a naming convention, but routing setup and maintenance will be unnecessarily complex.
- The PCA Gx-Prime interface support feature is backward compatible and functions whether or not PCRF pooling is available.
- A complete suspect binding and session removal from the network (not only from the binding and session SBRs) requires the involvement of the policy clients with corresponding actions in response to PCA's RAR requests. PCA has no control on a policy client's behavior. When a PCEF receieves a Gx RAR with Session-Release-Cause AVP, it is expected to send:
 - A RAA (2001 Result Code AVP value) in response to the RAR
 - A GX CCR-T to terminate the GX session

Limitations

 When a PCRF is selected for a new subscriber binding, a simple round-robin selection mechanism is employed. PCA PCRF selections can be overridden



by DSR routing configuration. When PCRF selections are overridden by DSR, weighted load distribution can also be used.

- PCA does not support the 3GPP mechanism to redirect Policy Clients to a PCRF.
- Quota pooling is not supported. Quota pooling is a feature that would allow a
 number of subscribers to share a common pool of resources for policy decisions.
 For example, a family plan where all members of the family share access to
 resources such as bandwidth. PCA has no mechanism for identifying members of
 a quota pool such that their sessions could all be routed to the same PCRF.
- The P-DRA function supports only two of the Diameter Subscription-Id types: END_USER_IMSI (for IMSI) and END_USER_E164 (for MSISDN). Any other Subscription-Id type is ignored.
- PCA evenly distributes new sessions across the Session Policy SBR Server
 Groups at the mated pair, regardless of the physical location of the Active server.
 This results in ~50% of session accesses traversing the WAN between the mated
 pair sites. For mated triplet deployments having an even distribution of Session
 SBR server groups, ~66¾% of session accesses traverse the WAN between
 mated sites.
- In cases of Regionalized OCS deployments, where the Requests are routed to an OC-DRA on a remote DSR, RBAR may have to be invoked for subsequent Requests (CCR-U/Ts) as well. If MSISDN is not present in CCR-U/T messages, regionalized routing cannot be supported.
- Enabling/disabling the OC-DRA functionality or P-DRA functionality on a per site basis or on a per NE basis while enabling both at the NOAM is not supported.
- For customers upgrading from P-DRA to PCA, the customer team must ensure that there is enough spare capacity available in the session SBRs to support the additional online charging sessions.
- The P-DRA function will reject any binding capable session initiation request after the binding migration period if: (a) the message has no APN (for example, Called-Station-Id AVP), or has an APN that is not configured in PCA, and (b) PCRF Pooling is enabled. The specifications point out a case in which a Gxx session initiation request could be sent with no APN, which will not work for PCRF Pooling

Note:

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

RBAR currently cannot extract the MSISDN from the User-Name AVP, but can
extract it from other AVPs. If there is need to support regionalized routing for
CCRs with MSISDN stored in the User-Name AVP, the Diameter Mediation feature
will have to be used to extract the MSISDN from the User-Name AVP and include
it in an AVP that is supported by RBAR.



- OC-DRA extracts the subscriber's identity from the session initiation request (CCR-I) for the purpose of including it with the session state information stored at the Session SBR when session state is required to be maintained. OC-DRA does not extract the MSISDN or perform number conditioning when the subscriber's identity is retrieved in a format other than E.164 (for example, MSISDN) such as SIP URI, TEL URI or NAI. The subscriber's identity is stored in the format in which it is retrieved from the Request.
- The known error conditions could result in a split binding condition are:
 - 1. A binding sessionRef is removed as a result of the Suspect Binding mechanism, but the actual Diameter session survived the PCRF inaccessibility. This condition is expected to be rare because for a Diameter session to survive the PCRF inaccessibility, there would have to be no signaling attempted for the session during the outage and the PCRF would have to maintain session state over the outage.
 - 2. A binding sessionRef was removed due to being discovered in an Early state for longer than the Maximum Early Binding Lifetime, but the actual Diameter session was successfully established. This condition is expected to be rare because the binding record is explicitly updated to Final when the master session succeeds or slave polling succeeds. This condition should only result from software errors or SBR congestion causing database update requests to be discarded.
 - 3. An attempt was made to create a binding-capable session record, but the attempt failed, which triggered a Session Integrity session teardown. However, this mechanism cannot succeed if no session record exists and topology hiding was in use for the policy client that tried to create the session (for example, because the resulting CCR-T cannot be routed to a topology hidden PCRF). This condition is unlikely to cause a split binding because PCA will request that the policy client tear down the session. If the policy client complies, the PCRF will have a hung session that must be audited out. If the policy client declines to tear down the session, a split binding could occur.
- Site redundancy that occurs while an SBR Re-Mating migration is in progress is not supported.
- A binding is associated with a subscriber key (IMSI) that is used to route the
 request Diameter message to a specific PCRF. If a request message contains
 multiple subscriber keys (MSISDN, IPv4, and IPv6), then all associated keys are
 part of the suspect binding removal procedure.
- In some specific situations, a Binding Capable Session can be established without an IMSI value. In such scenarios, additional Sessions may be established and routed using the alternate keys from the original Binding Capable session without an IMSI. If a Diameter message fails and matches the Suspect Binding Removal Rules, the original Session is removed immediately, regardless of the Remove Immediately attribute value of the matching rule.



3

Policy DRA Overview

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

Details about the user interface, feature components, and specific tasks is included in the configuration sections. See Configuration.

The Policy DRA Function

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

Policy DRA provides various capabilities:

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



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

- Routing of binding capable and binding independent (e.g. Rx, Gx-Prime) policy signaling to the correct PCRF for subscribers with an existing binding
- 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.



APNs must be configured before enabling the PCRF Pooling feature.



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

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

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

PCRF Pools and Sub-Pools Concepts and Terminology

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

Related Topics

Configuration

PCRF Pools

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



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 3-1, 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.



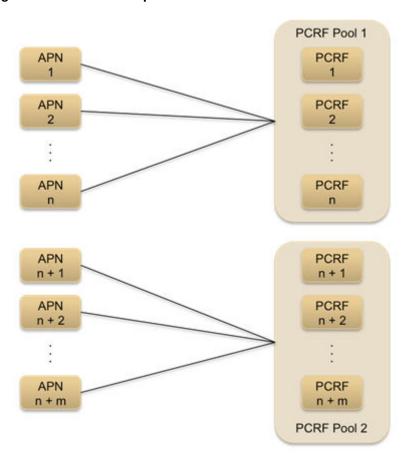


Figure 3-1 Relationship between APNs and PCRF Pools

Figure 3-2 illustrates the relationship between IMSIs and PCRF pool. The same IMSI must be able to have active bindings to multiple PCRF pools.

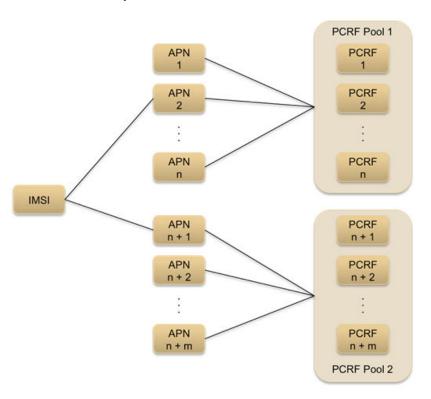


Figure 3-2 Relationship between IMSIs and PCRF Pools

Figure 3-3 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.



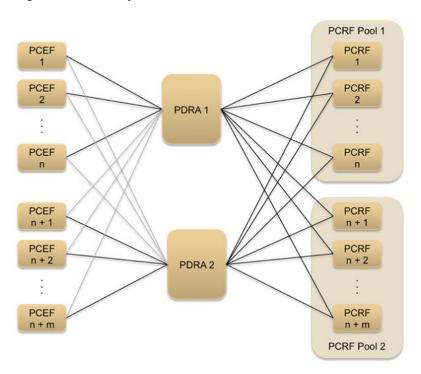


Figure 3-3 Multiple PCRF Pools

PCRF Pooling Modes

Routing used by PCRF Pooling depends on the Pool Mode. In Single Pool Mode, all binding capable session initiation request messages are routed to the Default PCRF Pool, regardless of the PCRF Pool mapped to the APN received in the request. The Default PCRF Pool is created automatically. This PCRF Pool is mapped to a Peer Route Table at every DSR site. In Single Pool Mode, all new binding capable session creation requests are routed to the default PCRF pool defined by a Peer Route Table at each DSR site.

In Multi Pool Mode, binding dependent session creation request messages, if required to correlate using MSISDN or IMSI keys, must include an APN. If an APN is not included in the requests, a Default APN is configured to be used to look up bindings using MSISDN or IMSI. The Default APN is used to operate PCRF Pooling in Multi Pool Mode but have a specific set of binding dependent interface equipment that initiate policy Diameter messages for subscribers for which the binding capable sessions were created using a single APN without including that APN in the binding dependent request message.

In Multi Pool Mode, binding dependent session creation request messages, if required to correlate using IPv4 or IPv6 addresses, need not include an APN in the request. The Default APN is also not used for binding correlation using IP addresses.

PCRF Sub-Pools

A PCRF Sub-Pool is a subset of a PCRF pool. The PCRF sub-pool is selected based on the Origin-Host of the PCEF sending a CCR-I.



Note:

Sub-pool rules are not applicable while operating in Single Pool mode.

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 3-4 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.



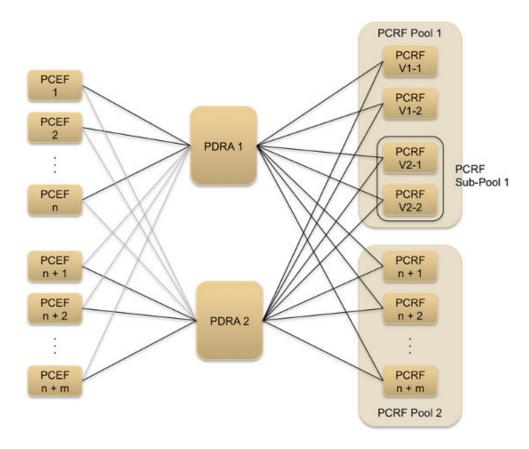


Figure 3-4 Multiple PCRF Versions in a PCRF Pool

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

- 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 3-5 shows example routing scenarios using PCRF Pools and Sub-Pools.



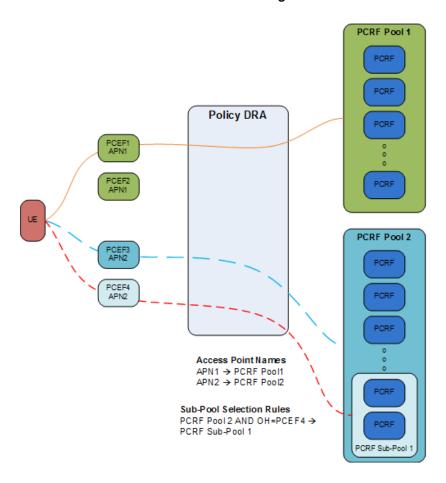


Figure 3-5 PCRF Pools and Sub-Pools Routing Scenarios

Planning for PCRF Sub-Pooling

To plan for PCRF Sub-Pooling, consider:

- 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. Use caution not to overwhelm the PCRFs



planned for the Sub-Pool by routing more signaling than they can reasonably support.

Session Binding

IMSI and MSISDN bindings always include an APN. IP address bindings do not include APN. In Multi Pool Mode, the APN is used to select the PCRF Pool for new bindings. In Single Pool Mode, the Default PCRF Pool is always used for new bindings.

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 session creation requests contain binding keys by which the subscriber's binding can be found. This process is called binding correlation. In Multi Pool Mode, if binding correlation is to be performed using IMSI or MSISDN keys, an APN must also be present in the binding dependent session creation request. In Single Pool Mode, binding correlation using IMSI or MSISDN need not include an APN. Binding correlation using IPv4 or IPv6 addresses never requires an APN in the binding dependent session creation request.

A binding capable session is a policy session that is allowed to cause a new binding to be created for a subscriber and APN. Binding capable session initiation requests must include both IMSI and APN. In Multi Pool Mode, new bindings for a subscriber are distributed across PCRFs in the PCRF Pool assigned to the APN provided in the session creation request. In Single Pool Mode, new bindings for a subscriber are distributed across PCRFs in the Default PCRF Pool, regardless of the PCRF Pool assigned to the APN in the session initiation request. Binding capable session initiation requests may also create alternate keys by which the subscriber may be identified. These alternate keys include MSISDN, IPv4 address, or IPv6 address. The binding of a subscriber and APN remains intact as long as the subscriber has at least one binding capable Diameter session for that binding.

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 3-1 shows a list of some Policy DRA terms and their meanings as they apply to this document.

Table 3-1 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.



Table 3-1 (Cont.) Policy DRA Terminology

Term	Meaning
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-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 in 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.



Table 3-1 (Cont.) Policy DRA Terminology

Term	Meaning
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 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.
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 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.)



Table 3-1 (Cont.) Policy DRA Terminology

Term	Meaning
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 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.



Table 3-1 (Cont.) Policy DRA Terminology

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

Policy DRA Functions

The Policy DRA functionality performs several major functions:

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



Managing subscriber Session and Binding databases

Diameter Request Message Processing

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

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

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

- The type of the Diameter Request: initiation, update, or termination
- The type of interface over which the Request message is carried and whether the
 session over this interface is binding-capable or binding-dependent.
 A session over a binding-capable interface will be eligible to establish a binding to
 a PCRF, while a session over a binding-dependent interface will rely on an existing
 binding to a PCRF but cannot create a new binding by itself.
- The subscriber's IDs from the appropriate AVPs (Subscription-ID AVP, Framed-IP-Address AVP, and Framed-IPv6-Prefix AVP)
- The Origin-Host and Realm AVPs, and Destination-Host and Realm AVPs.
- The access point name (APN) from which the request was received.
- Session-Id AVPs

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

Query Subscriber's Binding Status

Binding-capable Session Initiation Requests

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

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



Special Cases

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

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

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



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

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

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

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

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

- The APN in the binding-capable session initiation request (for example, CCR-I) is mapped to a PCRF Pool. This mapping is configured in **Policy DRA**, and then **Configuration**, and then **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 chosen PCRF Pool or Sub-Pool. The Diameter routing capabilities are used to load distribute the request across PCRFs in the specified pool.

✓ Note:

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

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

Early Binding

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

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

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

PCRF Selection and Routing

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

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



Note:

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

See Query Subscriber's Binding Status for a description of PCRF selection when PCRF Pooling is enabled.

Topology Hiding Process

See Network-Wide Options.

Diameter Answer Message Processing

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

- 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 an SBR, but no Answer is received after a configured time interval

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

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

Subscriber Session and Binding Database Management

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

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



Subscriber Identification and Binding

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

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

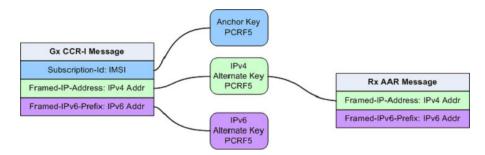
A session over a binding-capable interface will be eligible to establish a binding to a PCRF, while a session over a binding-dependent interface will rely on an existing binding to a PCRF but cannot create a new binding by itself.

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

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

In Figure 3-6, a Gx CCR-I message created 3 subscriber keys: one Anchor Key and two Alternate Keys, all bound to a PCRF called PCRF5. When a binding-dependent Rx session (AAR message) is created containing only IP addresses with no Anchor Key, the Policy DRA functionality looks up the IPv4 address of the subscriber and is able to relate it to the same PCRF because the Gx session had defined those IP addresses.

Figure 3-6 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 3-2 illustrates an example Binding Key configuration with priorities assigned to each key.

Table 3-2 Example Key Priority Configuration

Priority	Кеу Туре
1	IMSI
2	IPv4
3	MSISDN
4	IPv6
5	<not configured=""></not>

The example configuration in Table 3-2 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.
- If the Framed-IP-Address AVP is not found, a Subscription-Id AVP containing an MSISDN will be looked for.
- 4. If the Subscription-Id AVP with an MSISDN is found, look for a binding with that MSISDN.
- 5. If a Subscription-Id AVP containing an MSISDN is not found, then no Alternate Keys are present in the message and no Alternate Key records will be created by the application.

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

Binding-Capable Sessions

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

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

Binding-capable Session Initiation Request Processing Rules and Requirements

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



- Upon receipt of a binding-capable session initiation request containing no Called-Station-Id AVP (for example, no APN), Policy DRA generates and sends a binding-capable session initiation answer message using the Result Code configured for the Diameter interface for the Missing Or Unconfigured APN condition in the Error Codes GUI. The answer message shall include an Error-Message AVP with the 3-digit error code suffix of 500.
- Upon receipt of a binding-capable session initiation request containing no Called-Station-Id AVP (for example, no APN), Policy DRA asserts alarm-ID 22730 and increments measurement RxBindCapMissingApn by one.
- Upon receipt of a binding-capable session initiation request containing a Called-Station-Id AVP (for example, APN) that is not configured on the Access Point Names screen, 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 on the Error Codes screen. 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 screen, Policy DRA asserts Alarm-ID 22730 and increments measurement RxBindCapUnknownApn by one.
- Upon receipt of a binding-capable session initiation request containing a Called-Station-Id AVP (for example, APN) that is configured on the Access Point Names screen, the Policy DRA application performs PCRF Pool selection. Measurement RxBindCapApn2PcrfPool is incremented by one for the APN.
- If no PCRF Sub-Pool Selection rule matches, the suggested PCRF Pool is the PCRF Pool configured for the APN on the Access Point Names screen.
- If no PCRF Sub-Pool Selection Rule exists for the PCRF Pool that was assigned to the APN from the bindng-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 conditions 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 conditions 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 conditions 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.



The GUI prevents ambiguous Starts With and Ends With rules.

• If more than one PCRF Sub-Pool Selection Rule matches according to requirements, the Policy DRA application selects the match having the highest priority (for example, the lowest numeric priority value).

Note:

The GUI prevents creation of ambiguous, conflicting and duplicate rules.

- If a PCRF Sub-Pool Selection Rule matches according to requirements, Policy DRA application uses the PCRF Sub-Pool from the matching rule as the suggested PCRF Pool. Measurement RxBindCap2PcrfSubPool is incremented by one for the PCRF Sub-Pool Selection Rule that was matched.
- If a binding-capable session initiation request is received that would result in a
 new binding and no PCRFs are configured at the site, Policy DRA generates an
 error response with the 3002 Diameter Response-Code and Error-Message AVP
 including the string No PCRFs configured at this site.

Note:

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

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

Note:

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

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



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



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

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

Note:

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

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



PCRF Pool Selection

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

PCRF Pool Definition - Definition of the logical concept of a PCRF pool. This
includes configuring the 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

FODN (PCEF Origin-Host FODN)

An FQDN value or partial match.

PCRF Sub-Pool

One of the configured PCRF Pools.

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

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

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

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



Note:

The PCA GUI ensures no two rules with the same specificity have the same priority.

There is an order of precedence, from most specific to least:

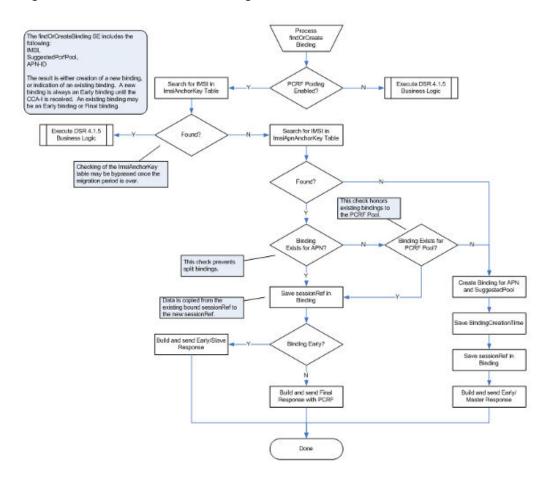
- Origin-Host full FQDN value
- 2. Origin-Host partial match

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 3-7 shows the logic used for this task.

Figure 3-7 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) certain actions occur:

- The answer message is relayed to the Policy Client that sent the request.
- A Session record is created with information related to the Diameter session.
- Alternate key binding records are created for the intersection of alternate keys configured in Policy and Charging, and then Configuration, and then Policy DRA, and then Binding Key Priority and alternate keys present in the bindingcapable 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) certain actions occur:

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

Related Topic

The P-DRA Database

Binding-Dependent Sessions

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

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

Figure 3-8 shows an overview of binding-dependent session initiation requests using IPv4 or IPv6 as correlation keys .



Try to find an MSISDN record If not found, set a failure indication for the findBindingResult stack event. Else, look for a binding with the specified APN. If a match is found, include the PCRF ID in the findBindingResult. If no match is found, set a failure indication in the findBindingResult stack event. Route the findBindingResult back to the caller (3). SBR(B) 2 findBinding 3 findBindingResult AF PCA DA-MP 4a AAR On receipt of the AAR (1) PCRF Find the highest priority alternate On receipt of the findBinding Result (3) key present in the AAR message. If the findBinding was successful, route the request Route a findBinding stack event (4a) via the routing layer (without specifying a PRT to a SBR(B) server (2). If the findBinding was not successful (and no other alternate keys are available), return an AAA to the AF (4b) using the Binding Not Found response code.

Figure 3-8 Binding Dependent Session Initiation Request Processing Overview

A specific 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 fount to PCRF X, which completes the process. If a match is not found, the result is binding not found, which competes the process.
- Binding-dependent session initiation requests using MSISDN as correlation key.
- Both MSISDN-Only and MSISDN+APN binding tables are audited.
- Both old and new IPv4 and IPv6 binding tables are audited.

Note:

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

Binding-dependent Session Initiation Request Processing Rules and Requirements

Binding-capable request processing uses the binding key priority table to determine which keys present in the message should have alternate keys created in the binding database. Binding-dependent processing uses the binding key priority configuration



to determine which keys to use and in what priority order when attempting binding correlation.

Related Topics

- The P-DRA Database
- In-Session Message Processing
- · Topology Hiding Process

In-Session Message Processing

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

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

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

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

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

Topology Hiding

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



Topology Hiding configuration is done on each Policy DRA DSR using the Policy and Charging section of the NOAM GUI. The configuration enables users to set the Topology Hiding function to be Enabled or Disabled for the Policy DRA node. After being enabled, the Topology Hiding function can be further configured to apply for a specific Topology Hiding Scope, as summarized in Table 3-3:

- 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 3-3 Topology Hiding Scope Configuration

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



Table 3-3 (Cont.) Topology Hiding Scope Configuration

Topology Hiding System Setting	Topology Hiding Scope Setting	Result
	All Foreign Realms + Specific Clients	Topology Hiding is performed if either All Foreign Realms or Specific Clients condition is met
	All Messages	Topology Hiding is performed for all messages destined to all policy clients

Session Integrity

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

Session Audit Premature Removal of Sessions

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

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

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

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

Incomplete Session Data

In order to reduce Diameter signaling latency for policy signaling, Policy DRA attempts to relay Diameter messages before updating its various database tables. Provided that all database updates are created successfully and in a timely manner, this works very well. There are scenarios in which records cannot be successfully updated and the Policy Client and the PCRF are not aware of any problem. Table 3-4 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 3-4 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 3-4 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 recreate 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.



Table 3-4 (Cont.) Policy DRA Error Scenarios for Session Integrity

Error Scenario	Policy DRA Behavior
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 3-5 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 3-5 Session Integrity Conditions and Policy DRA Reaction

Condition	RAR Type	Throttled	Comments
Session determined to be stale	Query	Y	Certain rules apply



Table 3-5 (Cont.) Session Integrity Conditions and Policy DRA Reaction

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

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

- Each session SBR performs RAR throttling independently
- The rate of RARs initiated by any given session SBR varies depending on the number of sessions in the list of sessions pending RAR treatment
- The rate of RARs initiated by a given session SBR is bounded to be no faster than on RAR every 20 ms and no slower than on RAR every 500 ms
- The fastest rate of RARs for a given session SBR is 1 per 20ms (50 per sec) when the list of sessions pending RAR treatment is 200 or more deep
- The slowest rate of RARs for a given session SBR is 1 per 500 ms (2 per sec) when the lsit of sessions pending RAR treatment is very small
- When the list of sessions pending RAR treatment is between 1 and 200 deep, the rate for sending RARs varies proportionally to the list depth between 2 per second and 50 per second
- If no response is received for an RAR for a given session, that RAR is not reattempted more frequently than every 2.5 seconds
- If no response is received from the RARs for a given session, the session is removed from the P-DRA database after 12 attempts

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

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

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



An RAA response with any other result code is ignored.

Suspect Binding Removal

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

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

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

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

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

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

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

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

 requesting DA-MP to send Query type RAR messages to policy clients for verifying session integrity of some session records stored on the Session SBR. Based on the RAA responses, the sessions in question will either be re-authorized or torn down in a later time by the policy clients



requesting DA-MP to send Release type RAR messages to policy clients for tearing down the sessions that the Session SBR has records. A Release RAR will request the policy clients to tear down the session in question by sending CCR-T messages to the bound PCRF.

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

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

For information on how to configure Suspect Binding Removal Rules, refer to Suspect Binding Removal Rules.

Per APN Subscriber Session Limiting

The Policy DRA functionality provides configurable options for setting a maximum per IMSI per APN session limit, as well as options if the session limit is exceeded. These fields for configuring these options are found on the **Policy and Charging**, and then **Configuration**, and then **Access Point Names** screen.

The **Maximum Allowed Sessions per IMSI** field sets the maximum number of bound sessions allowed per IMSI for a specific APN. The maximum number of simultaneous sessions allowed across all APNs combined is ten.

The **Per IMSI Session Exceeded Treatment** field determines what action Policy DRA takes if the maximum number of bound sessions for an IMSI for a specific APN is exceeded. If **Route** is selected, the CCR-I message will be routed and the oldest bound session will be replaced. If **Reject** is selected, the CCR-I message will be rejected using the Diameter response code configured for SBR Error.



If OCDRA Only is selected in the Function field on the Access Point Names screen, the Maximum Allowed Sessions per IMSI and Per IMSI Session Exceeded Treatment fields are disabled.

Refer to Access Point Names elements for further information.



4

Online Charging DRA Overview

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

Online Charging DRA Functions

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

- OCS Selection and Routing
- Session State Maintenance

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

OCS Selection and Routing

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

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

- Single PCS Pool
- Multiple OCS Pools

OCS Pool Selection modes are configurable using the NOAM Main Menu, and then Policy and Charging, and then Configuration, and then Online Charging DRA, and then Network-Wide Options, and then OCS Pool Selection Mode screen.

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

Single OCS Pool Mode

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

Note:

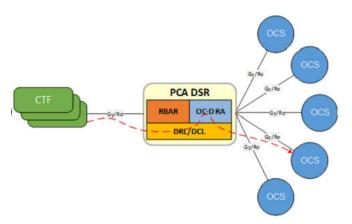
OC-DRA removing the Destination-Host AVP from the Online Charging Diameter Request when operating in the Single OCS Pool mode ensures that the request will not be accidentally rejected by the PRT or by the OCS Server if the Destination-Host AVP contained the PCA DSR's hostname.

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



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

Figure 4-1 Local OCS Server Selection



Multiple OCS Pools Mode

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



Note:

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

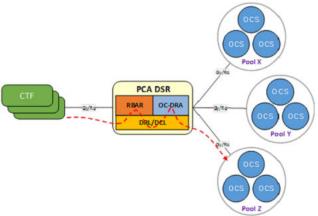
Note:

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

Note:

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

Figure 4-2 Local OCS Server Pool Selection



Regionalized Routing

Operators/service providers may have OCS deployments which are segmented based on ranges of subscriber identities (for example, MSISDNs) such that a given group of OCSs can only serve the subscriber range it has been assigned to serve. To support this architecture, the DSR RBAR application can be provisioned with higher priority

ART rules to be invoked prior to the invocation of PCA to perform regionalized routing based on subscriber's identity. If RBAR invocation fails, the DAL configuration should be provisioned such that RBAR is invoked on the mate DSR or an Answer response with a non-successful Result-Code/Experimental-Result AVP is generated and sent to the originator of the Diameter transaction.

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

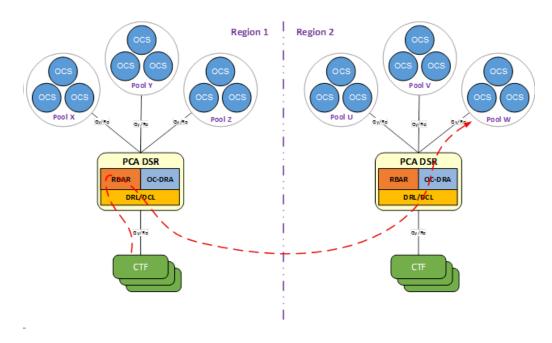


Figure 4-3 Regionalized OCS Server Pool Selection

Note:

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



Session State Maintenance

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

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

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

Session state is only applicable when Session-based charging is used and does not apply to Event-based online charging. As such, session state only applies to Diameter messages used for session based charging, which include CCR-I/U/T, CCA-I/U/T, RAR and RAA messages. Event-based charging Diameter messages, CCR-E and CCA-E do not create sessions. Thus, session state is not maintained for these messages.

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

Table 4-1 Session State Configuration Settings

Session State Setting	Scope	Description
No Messages	-	Session State is not maintained for any Gy/Ro sessions.
All Messages	-	Session State is maintained for all Gy/Ro sessions.



Table 4-1 (Cont.) Session State Configuration Settings

Session State Setting	Scope	Description
Specific Messages	Specific Realms	Session State is maintained for all sessions that are originated from an online charging client in a specific realm or being sent to an online charging server in the specific realm as configured in Policy and Charging, and then Configuration, and then Online Charging DRA, and then Realms.
	Specific Clients	Session State is maintained for all sessions that are originated from specific online charging clients as configured in Policy and Charging, and then Configuration, and then Online Charging DRA, and then CTFs.
	Specific Realms + Specific Client	Session State is maintained for all sessions that are originated from specific realms or specific online charging clients.
	Specific Servers	Session State is maintained for all sessions that are destined to specific online charging servers as configured in Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs.
	Specific Realms + Specific Servers	Session State is maintained for all sessions that are originated from specific realms or destined to specific online charging servers.
	Specific Realms + Specific Clients + Specific Servers	Session State is maintained for all sessions that are originated from specific realms or specific online charging clients or destined to specific online charging servers.

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



Table 4-2 Session State Information

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

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

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

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

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

Gy/Ro Diameter Request Message Processing

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

OC-DRA validates the Application ID and Command Code in the Diameter Request message for consistency. OC-DRA supports the Gy/Ro DCCA messages. If OC-DRA receives a Diameter Request message with a Command Code that is not supported, PCA will abandon message processing and send an Answer message with Result-



Code set to DIAMETER_COMMAND_UNSUPPORTED (3001) and Error-Message AVP to the downstream peer that initiated the Diameter transaction.

OC-DRA also makes Diameter Request message processing decision based on the small subset of AVPs for online charging. Those Diameter AVPs that are used specifically by OC-DRA for making routing decisions and maintaining session state. These AVPs are shown in Table 4-3 and marked M, O, $O_{\rm M}$, or - to indicate which ones are mandatory, optional, optional-mandatory, or not used for each of the supported Gy/Ro Diameter Credit Control Application Request messages.



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

Table 4-3 Diameter AVPs used by OC-DRA for Request Message Processing

	AVP	Used In			
AVP Name	Code	CCR	RAR	Value Type	Description
Session-Id	263	М	М	UTF8String	Contains the session identifier.
Origin-Host	264	M	М	DiamIdent	Contains the end point that originated the Diameter message.
Origin-Realm	296	M	М	DiamIdent	Contains the realm of the originator of the Diameter message.
Destination-Host	293	0	0	DiamIdent	Contains the end point to which the Diameter message is to be routed.
Destination-Realm	283	M	М	DiamIdent	Contains the realm to which the Diameter message is to be routed.
Auth-Application-Id	258	M	М	Unsigned32	Contains the application ID of the Diameter Credit Control Application which is 4.
CC-Request-Type	416	M	-	Enumerated	Contains the transfer type: event for event based charging and initial, update, terminate for session based charging.
User-Name	1	0	-	UTF8String	Contains the user name in the format of a NAI according to RFC 6733.
Subscription-Id	443	0	-	Grouped	Contains the identification of the user that is going to access the service in order to be identified by the OCS.
Subscription-Id- Type	450	O_{M}	-	Enumerated	Contains the type of the identifier, for example, value 0 is used for the international E.164 format according to ITU-T E.164 numbering plan. This AVP is a member of Subscription-Id Grouped AVP.
Subscription-Id- Data	444	O_{M}	-	UTF8String	Contains the user data content, for example, the MSISDN. This AVP is a member of Subscription-Id Grouped AVP.



Table 4-3 (Cont.) Diameter AVPs used by OC-DRA for Request Message Processing

	AVP	Used In		_	
AVP Name		CCR	RAR	Value Type	Description
Called-Station-Id	30	0	-	UTF8String	Contains the Access Point Name (APN) the user is connected to.

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

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

Session Initiation Request Message Processing

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

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

- Single OCS Pool
- Multiple OCS Pools

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

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

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



In-Session Request Message Processing

CCR

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

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

CCR-T

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

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

RAR

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

OC-DRA determines whether session state is maintained on the behalf of the online charging server based on session state configuration and message content. If session state is not maintained, OC-DRA routes the RAR without modification, expecting the online charging server to have set its Destination-Host AVP value to the hostname of the online charging client that originated the session initiation request. If session state is maintained, OC-DRA queries the Session SBR to retrieve and refresh the session state associated with the received Session-Id by sending a findAndRefreshOcSession



stack event to the Session SBR. If session state is found, OC-DRA replaces the Destination-Realm and Destination-Host in the RAR with the realm and hostname of the online charging client obtained from the session state and forwards it to DRL for routing. If session state is not found or an SBR error is encountered, OC-DRA will determine how to handle the message based upon the Session State Unavailable Action.

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

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

There are several user-configurable Session State Unavailable Actions are supported by OC-DRA:

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

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

Event Request Message Processing

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

Gy/Ro Diameter Answer Message Processing

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



PCA requests that DRL forward all Gy/Ro Answers back to PCA for Answer message processing in order to maintain measurements. PCA processing all Gy/Ro Answers for accounting purposes is not expensive since the OC-DRA and the routing layer are guaranteed to be on the same physical server (different threads in the same process) and avoids lots of explaining about why some measurements are not pegged.



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

Table 4-4 OC-DRA Header Information

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

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

OC-DRA makes Diameter Answer message processing decisions based on a small subset of AVPs defined in the Diameter protocol for online charging. Those Diameter AVPs that are used specifically by OC-DRA for making routing decisions. These AVPs are shown in Table 4-5 and marked M or - to indicate which ones are mandatory or not used for each of the supported Gy/Ro DCCA Answer messages.

Table 4-5 Diameter AVPs used by OC-DRA for Answer Message Processing

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

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

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



Session Initiation Answer Message Processing

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

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

In-Session Answer Message Processing

CCA-U

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

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

CCA-T

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

RAA

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

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

Event Answer Message Processing

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



DRL-Initiated Answer Message Processing

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

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



5

Configuration

This chapter defines the procedures required to configure the Policy and Charging Application (PCA) on a DSR system using the **Policy and Charging**, and then **Configuration** screens.

This chapter also contains an overview of information that is needed to configure and enable PCA, which includes configuring:

- Places and Place Associations
- Resource Domains
- Diameter Stack
- SBR Databases

Policy and Charging Configuration Overview

The **Policy and Charging**, and then **Configuration** screens for Policy and Charging components provide fields for entering the information needed to manage Policy and Charging configuration in the DSR.

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

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

 A pair of NOAM servers running in active/standby redundancy OAM configuration is performed on the NOAM.

Network-wide Policy and Charging configuration is performed on the NOAM.

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

Diameter protocol configuration is done on the SOAM.

Most DSR Application configuration is done on the SOAM.

Site-specific configuration for Policy and Charging is performed on the SOAM; all network-wide Policy and Charging configuration components are viewable on the SOAM.

 A set of MP servers, which can host signaling protocol stacks (such as DA-MPs) or in-memory database servers such as a Session Binding Repository (SBR) or IPFE servers.

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.

Three-tiered topology allows administrators to access all DSR GUI screens 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.

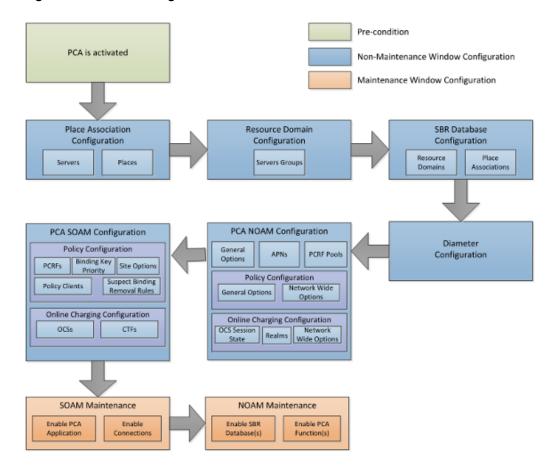


Figure 5-1 PCA Configuration Flow

NOAM and SOAM Configuration

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

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

In general, topology data like creation of sites, assignment of servers to sites, creation of server groups, and so on is A-scope data. DSR data configuration is generally site-scoped, or B-scope data.

Some Policy and Charging Application data must be configured at the A-scope level and some data must be configured at the B-scope level.

Policy related data configured at the NOAM include:

- Assignment of Servers to Site Places
- Assignment of Servers to SBR Server Groups



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

PCA-specific data configured at the NOAM include:

- Alarm Thresholds for:
 - PCA Application Ingress Message Rate
 - Session Database Capacity
 - Policy Binding Database Capacity
- · Policy DRA and Online Charging DRA function disabling/enabling
- Default Audit options
- OCS/Realm Session State
- Access Point Names (APN)
- Maximum Session Inactivity Time per APN
- PCRF Pools and PCRF Sub-Pool Selection Rules
- SBR Databases
- SBR Database Resizing Plans
- SBR Data Migration Plans

PCA-specific data configured at the SOAM include:

- PCRFs adjacent to the site
- Binding Key Priority for the site
- Topology Hiding configuration for the site
- OCSs and CTFs for OC-DRA
- Error response configuration for the site
- Congestion handling options
- Suspect Binding Removal Rules

For more information, see PCA Capacity Constraints.

Pre-Configuration Activities

Before PCA configuration can be performed, activities need to be performed in the system:

- Verify that the PCA application is activated in the system. (This is usually
 performed as part of the installation or upgrade activities.)
 Policy and Charging appears in the left-hand GUI menu on the NOAM and the
 SOAM after the application is activated.
- Verify that the NOAM and SOAM configuration is complete for PCA



- Gather component information that is required for Diameter, Diameter Common, and PCA configuration, including component item naming conventions and names, IP addresses, hostnames, and numbers of items to be configured.
 Naming Conventions and Hierarchical Routing illustrates the use of a naming convention.
- Configure Diameter Common components that are required for PCA configuration.
 See Diameter Common Configuration for PCA for PCA configuration information.
- Configure Diameter Configuration components that are required for PCA configuration. See Diameter Configuration for PCA.

System Topology

Networks

The Networks screen is used to create the networks used for internal, external, and signaling communications. The networks are grouped into logical buckets called network elements. Only after creating these buckets can the networks themselves be defined. One advantage of this architecture is simplified network device configuration and service mapping.

The workflow is to first create the network elements and then define the individual networks inside each element.

Network Elements

A network element is simply a collection of networks. In other words, a container of networks. Any servers belonging to a specific network element uses those networks exclusively to communicate internally and externally. A network element can contain multiple servers but a single server can only belong to one network element.

Using a three-tier DSR system as an example, a typical, regionally diverse, signaling network would have multiple network elements. Consider a system deployed across an east region and west region. The network element configuration might look like:

- NO East
- NO_West
- SO East
- SO West
- NO DR (Disaster Recovery Spare)



Caution:

Depending on the application, the workflow and provisioning instruction may differ from the direction provided here. Always follow the provisioning guidelines for your specific application and release.

There are two methods for creating network elements. The first method involves manual entry using the Networks [Insert network Element] screen. See Inserting a network element for more information on this method. The second method is more



encompassing and allows the user to simultaneously create the network element and associated networks. See Uploading a network element configuration file.

Networks Insert Network Element

Table 5-1 describes the elements of the Networks [Insert Network Element] screen.

Table 5-1 Insert Network Elements Elements

Field	Description	Data Input Notes
Network Element Name	The user-defined name for the	Must be unique.
	network element.	Format: String
		Range: 1-32 alphanumeric characters and underscore. Must contain at least one alphabetic character and must not start with a digit.
		Default: n/a
		A Value is required.

Inserting a network element

This procedure defines the manual process of inserting a network element. To view the procedure that involves the uploading of a network element configuration file see Uploading a network element configuration file.

Use this procedure to define and insert a network element:

- Click Configuration, and then Networking, and then Networks.
- 2. Click Insert Network Element.
- 3. Enter a unique name in the value field for **Network Element Name**.
- **4.** Enter a unique name across the network element table in **Network Element Name**.

See Network Insert elements for value limitations of the **Network Element Name** field.

Click OK to submit the information and return to the Networks screen or Cancel to discard the changes and return to the Networks screen.

The network element is added to the topology database tables, and the GUI displays the newly added network element in tab format on the Networks screen.

Uploading a network element configuration file

This procedure defines the automated process of uploading a network element configuration file to create the network element. To view the procedure that involves the manual process of inserting a network element see Inserting a network element.



Note:

Depending on the application, the workflow and provisioning instruction may differ from the direction provided here. Because applications differ, the format of the configuration file is not addressed here. Always follow the provisioning guidelines for your specific application and release.

Use this procedure to upload an XML file to configure a new network element:

- Click Configuration, and then Networking, and then Networks.
- Click Browse to locate the file you want to use to configure a new network element.

A file upload screen displays allowing you to navigate to and select the target configuration file.

Select the target file and click Open.

The screen disappears and the target file displays in the text box to the right of the **Browse** button.

4. Click Upload File.

The file is uploaded and data validation is performed.

Data validation is performed immediately. If the file is valid, a new network element is created and reflected in a new tab on the Networks screen. Alternately, a file that contains invalid parameters returns an error message, and no network element is created.

Viewing Network Elements

Use this procedure to view network elements:

- 1. Click Configuration, and then Networking, and then Networks.
- Network elements are presented in tabular form. If the target network element is not visible in the available screen space use the scroll right/left buttons located in the tool bar area to the right or left of the visible tabs.

Deleting a Network Element

Before deleting a network element the user must ensure that no servers are associated with the target network element. Attempting to delete a network element with at least one associated server results in an error message and the target network element is not deleted. If a network element contains networks, but is not associated with any servers, then deleting the network element is successful. The networks contained in the target network element are deleted along with the network element.

Use this procedure to delete a network element after confirming that no servers are associated with it:

- Click Configuration, and then Networking, and then Networks.
- Locate the target network element tab.



Network elements are presented in tabular form. If the target network element is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

- 3. Click the \mathbf{x} located on the tab of the target network element.
- 4. Click **OK** to delete the network element from the database tables.

A status message is presented stating the network element has successfully been deleted. Closing the status message returns you to the Networks screen.

The network element and related networks are deleted from the databases.

Exporting a network element configuration file

The network element **Export** button generates an installation script file used for configuration purposes. Use this procedure to export the configuration parameters of a network element:

- 1. Click Configuration, and then Networking, and then Networks.
- 2. Select the target network element tab.

Network elements are presented in tabular form. If the target network element is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

3. Click Export.

Open or save the configuration file from the screen. By default the name format of the output file is NE_<yyyymmdd>_<hhmmss>_<zone>.xml. You may change this as needed.

Network Insert elements

Table 5-2 describes the elements of the Networks Insert screen.

Table 5-2 Networks Insert Elements

Field	Description	Data Input Notes
Network Name	The name of the network.	Must be unique.
		Format: String
		Range: 1-31 alphanumeric characters. Must start with a letter. No special characters are allowed.
		Defaut: n/a
		A value is required.
Network Type	The type of network in the	Format: List
	context of the application.	Range: OAM or Signaling
		Default: OAM
VLAN ID	The VLAN ID of the Network	Format: Numeric
		Range: 1-4094
		A value is required.



Table 5-2 (Cont.) Networks Insert Elements

Field	Description	Data Input Notes
Network Address	The network address of the	Format: Valid network address
	Network	Range: Dotted decimal (IPv4) or colon hex (IPv6)
		Default: n/a
		A value is required.
Netmask	Subnetting to apply to servers within the Network	Format: Valid network netmask
		Range: Prefix length (IPv4 or IPv6) or dotted quad decimal (IPv4)
		Default: n/a
		A value is required.
Router IP	The IP address of a router on this network.	Format: Valid IP address
		Range: Dotted decimal (IPv4)
	Note: If this is a default	or colon hex (IPv6)
	network, this i used as the gateway address of the default route on servers with interfaces on this network. If customer router monitoring is enabled, this address is the one monitored.	Default: n/a
		Note: A value is not required. Networks without a router IP cannot be used as the default network. The default network selection defaults to No.
Default Network	Whether the network is the	Format: Option
	default gateway	Range: Yes or No
Routed	Whether the network is routed	Format: Option
	outside its network element.	Range: Yes or No
	Note: The network is automatically assigned to a network element when a server in a network element has an IP from the network assigned is to it.	Note : Select No to allow the same IMI network/IPs to be used at multiple Signaling sites.

Inserting a Network

Use the following procedure for manually inserting a network. Alternatively, you can use the automated process of uploading a network element configuration file to create both the network element and associated networks. See Uploading a network element configuration file.

- 1. Click **Configuration**, and then **Networking**, and then **Networks**.
- Locate and select the target network element tab where you want to create the network.

Network elements are presented in tabular form. If the target network element is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

3. Click Insert.

Enter a Network Name.

For more information about **Network Name**, or any field on this screen, see Network Insert elements.

- 5. Select a **Network Type** from the list.
- 6. Enter a VLAN ID.
- 7. Enter a Network Address.

This is a network address and not a host IP address.

- 8. Enter a Netmask.
- 9. Optional: Enter the Router IP

This is used as the gateway address of the default route if yes is chosen in step 10.

10. Choose whether this will be the network with a default gateway.

If yes is chosen, the gateway address entered in step 9 acts as the default route for servers with interfaces on this network.

- 11. Choose whether this network is routed outside its network element.
- 12. Click **OK** to submit the information and return to the Networks screen, or click **Apply** to submit the information and continue entering additional data. Clicking cancel discards your changes and returns you to the Networks screen.

The new network is added to the target network element.

Locking and Unlocking a Network

Any network on the system can be locked or unlocked. When a network is locked, no modifications may be made to any device or route that uses that network. To add a route or a device to a network, the network would have to be in an unlocked state.

- 1. Click Configuration, and then Networking, and then Networks.
- Locate and select the network element tab where the network you want to unlock exists.

Network elements are presented in tabular form. If the target network element is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

- 3. Identify the target network and determine the lock status. This can be accomplished by identifying the value of the Locked field for your target network. A value of yes indicates the network is currently locked; no indicates the network is not currently locked. Alternatively, you can select the target network and take note of the Lock/Unlock button. If the button transitions to Unlock then the network is currently locked; if the button transitions to Lock then the network is currently unlocked.
- To unlock a locked network, click Unlock and respond to the confirmation screen displayed. When unlocking you also have to confirm your decision using a checkbox.

The network is now unlocked.

5. To lock an unlocked network, click **Lock** and respond to the confirmation dialogue box that is presented.



The network is now locked.

The network is locked or unlocked.

Editing a Network

Not all networks can be edited. Pre-configured networks created during the install process, for example, cannot be edited. A network that cannot be edited is distinguished using italic font.



Before editing a network, generate a network report. The network report serves as a record of the network's original settings. Print or save the network report for your records. For more information about generating a network report, see Generating a Network Report.

- 1. Click Configuration, and then Networking, and then Networks.
- Locate and select the network element tab where the network you want to edit exists.

Network elements are presented in tabular form. If the target network element is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

3. Select the target network and determine the lock status. If the network is currently unlocked proceed to the next step. If the network is locked the Lock/Unlock button should be active and reflect Unlock. Click Unlock and respond to the confirmation dialogue box that is presented.

The network is now unlocked.

4. Navigate back to the target network elements tab and select the target network again. Click **Edit**.

If the network cannot be edited it means it is still locked or it is a pre-configured network.

5. Edit the available fields as necessary.

See Network Insert elements for details about the fields that display on this screen.



Fields that cannot be edited are disabled.

- 6. Click **OK** to submit the changes and return to the Networks screen, or click **Apply** to submit the information and continue editing additional data. Clicking cancel discards your changes and returns you to the Networks screen.
- 7. Return the target network to the desired lock status.

The network is changed.



Deleting a Network

Not all networks can be deleted. In-use networks and pre-configured networks created during the install process, for example, cannot be deleted. A network that cannot be deleted is distinguished using italic font.

Note:

Before deleting a network, generate a network report. The network report serves as a record of the network's original settings. Print or save the network report for your records. For more information about generating a network report, see Generating a Network Report.

- 1. Click Configuration, and then Networking, and then Networks.
- Locate and select the network element tab where the network you want to delete exists.

Network elements are presented in tabular form. If the target network element is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

- Click to select the network you want to delete. To delete multiple networks, press and hold Ctrl and click to select multiple networks.
 - If the network cannot be deleted, **Delete** is disabled. To delete multiple networks at one time, all selected networks must be deletable.
- Click Delete and respond to the confirmation dialogue box that is presented.
- 5. Click **OK** to delete the network.

The network has been removed from the database and it no longer displays in the network element tab.

Generating a Network Report

A network report provides a summary of the configuration of one or more networks. Reports can be printed or saved to a file.

- Click Configuration, and then Networking, and then Networks.
- 2. Locate and select the network element tab where the target networks exist.
 - Network elements are presented in tabular form. If the target network element is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.
- Click Report to generate a report for all networks. To generate a report for a single network, click to select the network and click Report. To generate a report for multiple networks, press and hold Ctrl as you click to select specific networks.
- 4. Click **Print** to print the report, or click **Save** to save the report to a file.

Services

This feature allows for flexible network deployment by allowing you to map an application service to a specific network. Additionally, this feature allows for the



differentiation of intra- and inter-networks on a per service basis. This means that traffic from different services can be segmented, which allows for service specific-networks and routes. This is predicated on the creation of network elements, networks, and routes to support the segmentation of service traffic.

Geo-redundant (spare) nodes and dual-path monitoring are special code on the node at the spare site that continually monitors the availability of the database instances at the primary site to determine if an automatic failover should occur due to loss of the active site servers. In the event of a network outage, it is possible that if the system is monitoring a single network path only and intra- and inter-networks are differentiated, an erroneous condition might occur where both sites try to assume activity. Inherent dual-path monitoring protects against this scenario.

The core services are:

- OAM
- Replication
- Signaling
- HA_Secondary
- HA_MP_Secondary
- Replication MP

For example, segregation of replication traffic might occur for inter-network (WAN) traffic only. Prerequisite configuration work would have included the creation of at least one LAN network and two WAN networks along with the related routes. For the purposed of this example, these could be named LAN1, WAN1, and WAN2. The services mapping might look similar to the settings in Table 5-3.

Table 5-3 Core Services

Name	Intra-NE Network	Inter-NE Network
OAM	Unspecified	Unspecified
	I AN1	WAN1
Replication		
Signaling	Unspecified	Unspecified
HA_Secondary	Unspecified	Unspecified
HA_MP_Secondary	Unspecified	Unspecified
Replication_MP	LAN1	WAN2



Services might vary depending on the application. For example, DSR adds a service known as ComAgent to the existing core services. Additionally, workflow and provisioning instruction might differ from the direction provided here. Always follow the provisioning guidelines for your specific application and release.

Editing Service information

Services are set during installation of the system. However, you can edit network characteristics of the services. Use this procedure to edit existing service information:



- 1. Click Configuration, and then Networking, and then Services.
- 2. Click Edit.
- 3. Select from the available choices to determine the Intra-NE Network.
- 4. Select from the available choices to determine the Inter-NE Network.
- 5. Click OK to submit the information and return to the Services screen, or click Apply to submit the information and remain in the Services [Edit] screen. Clicking cancel discards your changes and returns you to the Services screen.

Generating a Services Report

A services report provides a summary of the services configuration. This report can also be printed or saved to a file.

Use this procedure to generate a service reports:

- 1. Click Configuration, and then Networking, and then Services.
- 2. Click Report.
- 3. Click **Print** to print the report, or click **Save** to save a text file of the report. Clicking **Back** returns you to the Services screen.

Places

Places are used to build associations for groups of servers at a single geographic location. These places can then be grouped into place associations, which create relationships between one or more place.

Places Insert Elements

Table 5-4 describes the elements of the Places Insert screen.

Table 5-4 Places Insert Elements

Element	Description	Data Input Notes
Place Name	A unique name used to label the place.	Format: Alphanumeric characters and underscore (_) are allowed. A minimum of one alphabetic character is required.
		Range: Maximum length is 32 characters.
Parent	The parent of a place group	Format: List
		Note : This field is not used for PCA configuration. The only option is None.
Place Type	The place type.	Format: List
		Range: Site (default option).



Table 5-4 (Cont.) Places Insert Elements

Element	Description	Data Input Notes
Servers	List of the available servers in	Format: Checkbox
	the NO or SO	Note : Select all of the DA-MP and SBR servers that are physically located at this Site Place.

Inserting a Place

Use this procedure to configure a place:

- 1. Click Configuration, and then Places.
- 2. Click Insert.
- 3. Enter the Place Name.
- 4. Select a **Parent** from the list.



A Parent Place is not required for PCA Places and can be set as None.

- 5. Select a Place Type from the list.
- Select all of the available DA-MP and SBR Servers that are physically located at this Site Place.
- Click OK to submit the information and return to the Places screen, or click Apply to submit the information and continue adding additional data.

Editing a Place

Use this procedure to edit place information

- 1. Click Configuration, and then Places.
- 2. Select the place from the listing.
- 3. Click **Edit** at the bottom of the table.
- Modify one or more of the place information fields.
- Click OK to submit the information and return to the Places screen, or click Apply to submit the information and continue editing additional data.

The place information is updated in the network database and the changes take effect immediately.

Deleting a Place

Use this procedure to delete a place.

- 1. Click Configuration, and then Places.
- 2. Click to select the place you want to delete from the table.

A Place cannot be deleted if it includes servers. Before deleting, disassociate any servers.

- Click Delete.
- 4. Click **OK** to delete the place.

If you click **Cancel**, the place is not deleted, and you are returned to the Places screen.

Generating a Places Report

Use this procedure to generate a places report:

- 1. Click Configuration, and then Places.
- 2. Click to select the place for which you want to create a report.



To select multiple server groups, press and hold **Ctrl** as you click to select specific rows. Alternatively, if no servers are selected then all server groups appear in the report.

- 3. Click Report.
- 4. Click **Print** to print the report, or click **Save** to save a text file of the report.

Server Groups

The Server Groups feature allows the user to assign a function, parent relationships, and levels to a group of servers that share the same role, such as NOAM, SOAM, and MP servers. The purpose of this feature is to define database relationships to support the high availability architecture. This relates to replication, availability, status, and reporting at the server level.

From the Server Groups screen users can create new groups, edit groups, delete groups, and generate reports that contain server group data. Servers can be added or removed from existing groups using the edit function.

The Server Groups screen can be accessed from the main menu by navigating to **Configuration**, and then **Server Groups**. The screen displays a grid reflecting all currently configured server groups. A description of the elements displayed in the grid can be found in **Server Groups Edit Elements**.



Depending on the application configuration, the preferred HA role preference, or NE HA Pref, may not be displayed.



Server Groups Insert Elements

Table 5-5 describes the elements of the Server Groups [Insert] screen.

Table 5-5 Server Groups Insert Elements

Element	Description	Data Input Notes
Server Group Name	A unique name used to label	Format: String
	the server group.	Range: 1-32 characters. Alphanumeric and underscore are allowed. A minimum of one alphabetic character is required and must not start with a digit.
		Default: N/A
		A Value is required.
Level	The level of the servers	Format: List
	belonging to this group.	Range: Levels A, B, or C
		Note: Level A groups contain NOAMP and Query servers. Level B groups are optional and contain SOAM servers. Level C groups contain MP servers.
		A Value is required.
Parent	The parent server group that	Format: Pulldown menu
	functions as the replication parent of the selected server group	Note : If the level of the group being inserted is A, then the parent field is not editable and NONE is displayed in the list.
Function	The defined function for the server group.	Format: List
		Range: Functions supported by the system
WAN Replication Connection Count	Specifies the number of TCP connections that are used by replication over any WAN connection associated with this Server Group.	Format: Numeric
		Range = An integer between 1 and 8
		Default = 1

Inserting a Server Group

Use this procedure to configure a server group:



Servers are not added at this time. Only after the SG is created can servers be added using the edit function.

- 1. From the main menu select **Configuration**, and then **Server Groups**.
- 2. Click Insert.



3. Enter the Server Group Name.

For more information about **Server Group Name**, or any of the fields on this screen, see **Server Groups Insert Elements**.

- 4. Select a **Level** from the list.
- 5. Select a **Parent** from the list.
- 6. Select a **Function** from the list.
- 7. Enter a WAN Replication Connection Count.
- 8. Click **OK** to submit the information and return to the server groups screen or click **Apply** to submit the information and continue adding additional data. Clicking **Cancel** discards all changes and returns return you to the server groups screen.

Server Groups Edit Elements

The Server Groups [Edit] screen allows you to edit existing server groups. Table 5-6 describes the elements of the Edit Server Groups screen.

Table 5-6 Server Groups Edit Elements

Element	Description	Data Input Notes
Server Group Name	A unique name used to label	Format: String
	the server group.	Range: 1-32 characters. Alphanumeric and underscore are allowed. A minimum of one alphabetic character is required and must not start with a digit.
		Default: N/A
		A Value is required.
Level	The level of the servers	This field cannot be edited.
	belonging to this group.	Format: List
		Range: Levels A, B, or C
		Note: Level A groups contain NOAMP and Query servers. Level B groups are optional and contain SOAM servers. Level C groups contain MP servers.
		A Value is required.
Parent	The parent server group that functions as the replication parent of the selected server group.	Format: List
		Note : If the level of the group being inserted is A, then the parent field is not editable and NONE is displayed in the list.
Function	The defined function for the server group.	This field cannot be edited.
		Format: List
		Range: Functions supported by the system



Table 5-6 (Cont.) Server Groups Edit Elements

Element	Description	Data Input Notes
WAN Replication Connection Count	Specifies the number of TCP connections that are used by replication over any WAN connection associated with this Server Group.	Format: Numeric Range = An integer between 1 and 8 Default = 1
Prefer Network Element as spare	The Preferred HA Role Setting for the NE.	Format: Checkbox
	When marked as a preferred spare, the network element only assumes an active or standby role if all the other network elements are unavailable. This allows the user to isolate a dedicated disaster recovery element from normal operations.	
	Note : Depending on the application configuration, this selection may not be available.	
Server	The name of a server available for inclusion in the server group.	Automatically populated based on servers available for inclusion.
SG Inclusion	When checked, the server is included in the server group.	Checkbox
Preferred HA Role	The Preferred HA Role Setting for the server.	Checkbox
	When marked as a preferred spare, the server only assumes an active or standby role if all the other servers in the server group are unavailable. This allows the user to isolate a dedicated disaster recovery node from normal operations.	
VIP Assignment: VIP Address	A virtual IP address shared by the servers in this group that have networking interfaces on the same layer-2 network.	Format: Valid IP address Range: Four, 8-bit octets separated by periods [The first octet = 1-255; the last three octets = 0-255] Dotted quad decimal (IPv4) or colon hex (IPv6)

Editing a Server Group

Once a server group is created, certain values can be edited, and available servers can be added to or deleted from the server group. Additionally, the edit screen presents new fields and choices not present when initially creating the server group. For details regarding specific edit topics, select the appropriate link to display that information.

• For adding a server, see Adding a server to a server group.

- For deleting a server, see Deleting a server from a server group.
- For assigning a VIP to the server group, see Assigning a VIP to a server group.
- For removing a VIP from the server group, see Removing a VIP from a server group.

Use this procedure to edit a server group:

- 1. From the main menu select Configuration, and then Server Groups.
- 2. From the grid, click to select the server group you want to edit.
- 3. Click Edit.
- 4. Edit the values you want to change.
 - Fields that cannot be edited are grayed out. For more information about these fields, or any of the fields in this procedure, see Server Groups Edit Elements.
- 5. Click OK to submit the information and return to the Server Groups screen, or click Apply to submit the information and continue adding additional data. Clicking Cancel discards all changes and returns return you to the Server Groups screen.

Adding a server to a server group

Only after a server group is created can servers can be added. Use this procedure to add a server to a server group:

- 1. From the main menu select **Configuration**, and then **Server Groups**.
- 2. From the table, click to select the server group you want to edit.
- Click Edit.
 - The Servers Groups [Edit] screen displays the servers in the network element that are possible candidates for inclusion in the server group.
- 4. To add a server to the server group, select the checkbox for **SG Inclusion**. When checked, the server is included in the server group.
- 5. Click OK to submit the information and return to the Server Groups screen, or click Apply to submit the information and continue adding additional data. Clicking Cancel discards all changes and returns return you to the Server Groups screen.

Deleting a server from a server group

Use this procedure to delete a server from a server group:

- 1. From the main menu select **Configuration**, and then **Server Groups**.
- 2. From the table, click to select the server group you want to edit.
- Click Edit.
- 4. To delete a server from the server group, de-select the checkbox for **SG Inclusion**. When unchecked, the server is not included in the server group.
- 5. Click OK to submit the information and return to the Server Groups screen, or click Apply to submit the information and continue adding additional data. Clicking Cancel discards all changes and returns return you to the Server Groups screen.

Assigning a VIP to a server group

Use this procedure to assign a VIP to a server group.



Note:

This procedure is optional and is only supported if the system supports VIP.

- 1. From the main menu select Configuration, and then Server Groups.
- 2. From the table, click to select the server group you want to edit.
- 3. Click Edit.
- 4. Click **Add** to add a new VIP address to the server group.

Note:

Multiple VIP addresses can be added.

- 5. Insert the VIP address.
- 6. Click OK to submit the information and return to the Server Groups screen, or click Apply to submit the information and continue adding additional data. Clicking Cancel discards all changes and returns return you to the Server Groups screen.

Removing a VIP from a server group

Use this procedure to remove a VIP address from a server group:

- 1. From the main menu select **Configuration**, and then **Server Groups**.
- 2. From the table, click to select the server group you want to edit.
- 3. Click Edit.
- 4. Click **Remove** next to the VIP address you want to remove from the server group. The VIP address is removed from the server group.
- 5. Click **OK** to submit the information and return to the Server Groups screen, or click **Apply** to submit the information and continue adding additional data. Clicking **Cancel** discards all changes and returns return you to the Server Groups screen.

Deleting a Server Group

Use this procedure to delete a server group.

Note:

Only a server group with no existing servers in the group can be deleted. For information about how to delete a server from a server group, see Deleting a server from a server group.

- 1. From the main menu select Configuration, and then Server Groups.
- 2. Click to select the server group you want to delete from the table.
- 3. Click Delete.



4. Click **OK** to delete the server group.

If you click **Cancel**, the server group is not deleted, and you are returned to the Server Groups screen.

Generating a Server Group Report

Use this procedure to generate a server group report:



Depending on the application configuration, the NE HA Pref, or network element high availability preference, may not be displayed.

- 1. From the main menu select **Configuration**, and then **Server Groups**.
- 2. Click to select the server group for which you want to create a report.



To select multiple server groups, press and hold **Ctrl** as you click to select specific rows. Alternatively, if no servers are selected then all server groups appear in the report.

- 3. Click Report.
- 4. Click **Print** to print the report, or click **Save** to save a text file of the report.

Devices

The Devices screen is used to configure and manage additional interfaces other than what was configured during the initial installation.

Devices elements

Table 5-7 describes the elements of the Devices screen.

Table 5-7 Devices Elements

Tab/Field	Description	
Server	The server host name displayed in tabbed format at the top of the table	
Device Name	The name of the device (not user defined)	
Device Type	The device type. Supported types include: Bonding Vlan Alias Ethernet	
Device Options	A collection of keyword value pairs for the device options	



Table 5-7 (Cont.) Devices Elements

Tab/Field	Description
IP Interface (Network)	IP address and network name in the format: IP Address (network name)
Configuration Status	 The configuration status of the device. The possible states are: Discovered (provisioned directly on the server) Configured (provisioned through the GUI; server update is pending) Deployed (provisioned through the GUI; server update is complete) Pending (edit or delete update in progress) Deferred (server cannot be reached for updates) Error (specific error text is displayed in the Configuration Status field)
Is Locked?	Status of the lock state. The possible states are:
	 Locked (Not available for edit or delete)
	Unlocked (Available for edit or delete)

Viewing a Device

Devices are viewed on per server basis. The use of italics indicates a device that cannot be edited or deleted.

Use this procedure to view devices:

- 1. Click Configuration, and then Networking, and then Devices.
- 2. Locate and select the desired server tab.

Servers are presented in tabbed form. If the target server is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

The devices for that selected server are displayed.

Device Insert elements

Table 5-8 describes the elements of the Devices [Insert] screen.



Some fields are dynamic and only display when specific values are selected. Dynamic fields are noted in the description.



Table 5-8 Devices General Options

Field	Description	Data Input Notes
Device Type	The type of device. Note: A device type of Ethernet is system generated and not selectable from this screen.	Format: Options Range: Bonding, VLAN, Alias Default: N/A A value is required.
Start on Boot	When selected, this checkbox enables the device to start on boot.	Format: Checkbox Default: Enabled
Boot Protocol	The boot protocol.	Format: List Range: None, DHCP Default: None A value is required.
MTU Setting	The Maximum Transmission Unit (MTU) setting for the device (bytes per packet). Caution: Changing the MTU setting for an existing interface restarts the interface, which is service affecting.	Format: Numeric Range: 1280-65570 Default: 1500
Monitoring Type	The monitoring type to use with a bonding device. Note: This field is dynamic and only displays when bonding is selected as the device type.	Format: Options Range: MII, ARP Default: MI A Value is required.
Primary	The preferred primary interface. Note: This field is dynamic and only appears when bonding is selected as the device type and a monitoring type choice is selected.	Format: List Range: None - all available devices Default: None A value is required.
Monitoring Interval	MII monitoring interval in milliseconds. Note: A monitoring type is selected by default (MII).	Format: Numeric Range: A positive integer Default: 100ms A value is required.
Upstream Delay	MII monitoring upstream delay in milliseconds. Note: This field is dynamic and only appears when bonding is selected as the device type and MII is selected as the monitoring type.	Format: Numeric Range: A positive integer Default: 200ms A value is required.



Table 5-8 (Cont.) Devices General Options

Field	Description	Data Innut Nat
Field	Description	Data Input Notes
Downstream Delay	MII monitoring downstream delay in milliseconds.	Format: Numeric
	Note: This field is dynamic	Range: A positive integer Default: 200ms
	and only appears when	A value is required.
	bonding is selected as the device type and MII is selected as the monitoring type.	·
ARP Validation	The method to validate the	Format: List
	ARP probes and replies. Note: This field is dynamic	Range: None, Active, Backup, All
	and only appears when	Default: None
	bonding is selected as the device type and ARP is	A value is required.
	selected as the monitoring type.	
ARP Target IP(s)	Comma-separated ARP target	Format: Valid IP addresses
	IP address list.	Range: Dotted quad decimal (IPv4) or colon hex (IPv6)
	Note : This field is dynamic and only appears when	Default: None
	bonding is selected as the device type and ARP is	Range: Dotted quad decimal
	selected as the monitoring type.	(IPv4) or colon hex (IPv6)
Base Device(s)	The base device(s) for bond,	Format: Options
	alias, and VLAN device types. Note: Alias and VLAN devices	Range: Available base devices
	require one selection; bond devices require two selections. This cannot be changed after the device is created.	Default: N/A A Value is required.
IP Interfaces	and device to dicated.	
Add IP Interface	Presents a row with a single	Format: Button
	address box and network list.	At least one entry is required.
	Note: For each row, only one IP Address and network can be specified. To specify additional rows, click Add IP Interface.	
Remove	Removes the device interface IP Address on the selected row.	Format: Button
	Note: This is not a delete button. If Apply has already been selected, clicking Remove does not delete the interface. Deleting an interface that has already been defined takes place from the Devices screen.	
Submit Buttons		



Table 5-8 (Cont.) Devices General Options

Field	Description	Data Input Notes
OK	Submits the information to the database, and, if successful, returns you to the Devices screen.	
Apply	Submits the information to the database, and, if successful, remains on the Devices [Insert] screen so that you can enter additional data.	
Cancel	Discards the information and returns you to the Devices screen.	

Inserting a Device

The Devices [Insert] screen uses dynamic options. Depending on the selected value of a field, options may be added or removed from the screen. It is important to review and understand the elements associated with this screen by reviewing the Device Insert elements screen.



Devices cannot be created that use management networks (those configured after installation and designated in the Network listing in blue italic text). This ensures continued access to the GUI via the management networks. Additionally, device creation requires that the prerequisite networks are already configured. See Inserting a Network for more details.

- Click Configuration, and then Networking, and then Devices.
- 2. Locate and select the desired server tab.

Servers are presented in tabbed form. If the target server is not visible in the available screen space, use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

- Click Insert.
- 4. Select the **Device Type**. If the selected device type is **Bonding**, then continue with this step; otherwise, skip to 5.
 - a. By default, Start on Boot is enabled. Uncheck the checkbox if you want to disable Start on Boot.
 - b. Select the Boot Protocol.
 - c. Enter the MTU Setting if a default of 1500 is not desired.
 - Select the Monitoring Type.
 - e. Select the **Primary** interface.
 - Enter the Monitoring Interval.



- g. If MII was selected as the monitoring type, then enter the **Upstream Delay** in milliseconds; otherwise, skip to substep 4.j.
- **h.** Enter the **Downstream Delay** in milliseconds.
- i. Select Base Devices. Two must be selected.
- j. If ARP was selected as the monitoring type, then enter the ARP Validation method.
- **k.** Enter the **ARP Target IP(s)** using valid comma separated IP addresses.
- I. Skip to step 7 to continue.
- 5. If the selected device type is **VLAN**, then continue with this step; otherwise, skip to 6.
 - a. By default, Start on Boot is enabled. Uncheck the checkbox if you want to disable Start on Boot.
 - b. Select the **Boot Protocol**.
 - c. Enter the MTU Setting if a default of 1500 is not desired.
 - **d.** Select **Base Device**. Only one can be selected.
 - e. Skip to step 7 to continue.
- If the selected device type is Alias, then continue with this step; otherwise, skip to
 7.
 - a. By default, Start on Boot is enabled. Uncheck the checkbox if you want to disable Start on Boot.
 - b. Select the **Boot Protocol**.
 - **c.** Enter the **MTU Setting** if a default of 1500 is not desired. This is not an option for Alias device.
 - d. Select Base Device. Only one can be selected.
- Click Add IP Interface.

A new row is created with a textbox and list.

- 8. Enter an IP Address for the device.
- 9. Select a Network Name from the list.
- **10.** For each row, only one IP Address and Network Name can be specified. To specify additional interfaces, select **Add IP Interface** and complete steps 8 and 9.
- 11. When you are finished adding IP addresses, click OK to submit the information and return to the Devices screen, or click Apply to submit the information and continue entering additional data. Clicking Cancel discards your changes and returns you to the Devices screen.

Taking ownership of a device

Devices that have a configuration status of **Discovered** are devices that were configured during the initial install or extension process and not added manually. The user has limited abilities to modify these devices. When the need arrises to edit the attributes of these devices, the user must first take ownership of the device.

Before taking ownership of a device, the user should be familiar with the concept of locked/unlocked networks. Before editing or deleting any device that belongs to a



locked network, the network must be unlocked. See Locking and Unlocking a Network for more information.

Note:

Not all devices must belong to a network. For example, primary interfaces with a state of **Discovered** may not belong to a network.

The process of taking ownership of a device and then editing or deleting that device slightly differs depending on whether or not that device currently belongs to a locked network. See Editing a Device for more information.

Before taking ownership of a discovered device, the device has a configuration status of **Discovered**; **Locked**. **Edit** and **Delete** are disabled. Immediately after taking ownership of the device, the configuration status temporarily changes to **Configured** and then **Pending**. Within a few minutes, the device should transition to its final configuration status of **Deployed**. If the device belonged to a locked network before taking ownership, the status displays as **Deployed**; **Locked**, otherwise it displays as **Deployed**; **Unlocked**.



Before taking ownership of a device, generate a device report. The device report serves as a record of the device's original settings. Print or save the device report for your records. For more information about generating a device report, see Generating a Device Report.

Use the following procedure to take ownership of a device.

- 1. Click Configuration, and then Networking, and then Devices.
- Locate and select the desired server tab where the target device exists.
 - Servers are presented in tabbed form. If the target server is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.
- Click to select the device you want to take ownership of. Alternately, you can take ownership of multiple devices. Press and hold Ctrl and click to select more than one device.
 - If one or more selected devices have a configuration status of something other than **Discovered**, the **Take Ownership** button is disabled. To take ownership of multiple devices at one time, all selected devices must have a configuration status of **Discovered**.
- Click Take Ownership.

The configuration status temporarily displays **Configured**, then **Pending**, and finally **Deployed**.

The devices are now available for editing or deleting. Take note of the lock status. A device cannot be edited or deleted while in the **Locked** state. See Locking and Unlocking a Network for details on changing the lock status.



Editing a Device

Devices with a locked status cannot be edited without unlocking the network to which they belong. See Locking and Unlocking a Network for more information. Additionally, devices that have a configuration status of discovered cannot be unlocked until you take ownership of the device. See Taking ownership of a device for more information. Some discovered devices not belonging to a network are unlocked immediately after taking ownership. Other discovered devices require the extra step of unlocking the network after taking ownership.



Before editing a device, generate a device report. The device report serves as a record of the device's original settings. Print or save the device report for your records. For more information about generating a device report, see Generating a Device Report.

- 1. Click Configuration, and then Networking, and then Devices.
- Locate and select the desired server tab.

Servers are presented in tabbed form. If the target server is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

3. Click to select a device and click Edit.

If the device cannot be edited, **Edit** is disabled. Confirm the device is in a deployed and unlocked state. If the device can be edited, the Devices [Edit] screen appears.

4. Edit the available fields as necessary.

See Device Insert elements for details about the fields that appear on this screen.



Fields that cannot be edited are disabled.



Caution:

Changing the MTU setting for an existing interface restarts the interface, which affects service.

5. Click **OK** to submit the information and return to the Devices screen, or click **Apply** to submit the information and continue entering additional data. Clicking cancel discards your changes and returns you to the Devices screen.

Deleting a Device

Devices with a locked status cannot be deleted without unlocking the network to which they belong. See Locking and Unlocking a Network for more information. Additionally,

devices that have a configuration status of discovered cannot be unlocked until you take ownership of the device. See Taking ownership of a device for more information. Some discovered devices not belonging to a network are unlocked immediately after taking ownership. Other discovered devices require the extra step of unlocking the network after taking ownership.

Note:

Before deleting a device, generate a device report. The device report serves as a record of the device's original settings. Print or save the device report for your records. For more information about generating a device report, see Generating a Device Report.

- 1. Click Configuration, and then Networking, and then Devices.
- Locate and select the desired server tab.
 - Servers are presented in tabbed form. If the target server is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.
- Click to select the device you want to delete. Alternately, you can delete multiple
 devices. To delete multiple devices, press and hold Ctrl and click to select specific
 devices.
 - If the device cannot be deleted, **Delete** is disabled. Confirm the device is in a deployed and unlocked state. To delete multiple devices at one time, all selected devices must be deletable.
- 4. Click Delete.
- 5. Click OK.

Generating a Device Report

A device report can be generated on a single device, multiple devices within the same server, or all devices regardless of server.

- 1. Click Configuration, and then Networking, and then Devices.
- 2. Locate and select the desired server tab.

Servers are presented in tabbed form. If the target server is not visible in the available screen space use the scroll right/left buttons located below the tool bar area and to the right or left of the visible tabs.

The device data for the selected server displays.

- 3. To generate a device report, select one of the following procedures:
 - To generate a report for all devices under the current server tab, click Report.
 - To generate a report for a single device, click to select the device and click **Report**. Alternatively, you can select multiple devices. To generate a report for multiple devices, press and hold **Ctrl** as you click to select specific devices.
 - To generate a report for all devices regardless of server, click Report All.

The Device Report is generated.

4. Click **Print** to print the report.



5. Click **Save** to save the report to a file.

PCA Topology

Identifying Place and Place Association information

- 1. Identify and note the number of places and place names. A Place is used to identify the geographical location of a DSR node (SOAMs and subtending servers). There may be up to 32 places.
- 2. Identify the level of site redundancy to be deployed in the PCA system.
 - a. In case site redundancy is not required, the number of non-redundant PCA sites is the same as the number of Places
 - **b.** In case a two site redundancy model is chosen for some or all sites, identify and note the number of PCA mated pairs
 - c. In case a three site redundancy model is chosen for some or all sites, identify and note the number of PCA mated triplets
- 3. If the Policy DRA function is being configured, then identify and note the places that are associated to the Place Association with Policy Binding Region type. Generally, all stie places will be included in the Policy Binding Region Place Association.
- 4. Identify and note the places that are associated to the Place Association with Policy and Charging Mated Sites type.



The Policy and Charging Mated Sites Place Association type is used for all levels of site redundancy (no site redundancy, two-site redundancy, or 3-site redundancy).

Note:

There must be a Policy and Charging Mated Sites Place Association for each instance of a Policy and Charging Session Database. The number of site places in each place association depends on the level of site redundancy chosen.

- In the case of site redundancy, the Place Association will have one site
- In the case of two site redundancy, the Place Association will have two sites
- In the case of three site redundancy, the Place Association will have three sites

PCA Mated Sites - Identify and Log the site names for single sites, mated pairs or mated triplets



Identifying Resource Domain information

A Resource Domain is a collection of Server Groups that share a common purpose. For PCA, there are 3 types of Resource Domains: Policy Binding, Policy and Charging Session, and Policy and Charging DRA.

 Identify and log the number of Policy and Charging DRA resource domains and their server groups.

Note:

Depending on the redundancy model chosen, there can be up to three server groups in one Policy and Charging DRA resource domain. Each DSR (multi-active cluster) Server Group that will use a given instance of a Policy and Charging Session database should be included in a Policy and Charging DRA Resource Domain. A given DSR (multi-active cluster) Server Group may be included in only one Policy and Charging Session Resource Domain.

Identify and log the number of Policy Binding resource domains and its server groups.

Note:

This step is required for Policy DRA functionality only.

Note:

Depending on the capacity chosen, there can be up to eight server groups in one Policy Binding resource domain. All Session Binding Repository Server Groups that will host the Policy DRA binding database must be included in one Policy Binding Resource Domain.

3. Identify and log the number of Policy Session resource domains and their server groups.

Depending on the capacity chosen, there can be up to eight server groups in one Policy Session resource domain. All Session Binding Repository Server Groups that will host an instance of a PCA Session database must be included in the same Policy and Charging Session Resource Domain. Each Session database will have its own Policy and Charging Session Resource Domain.

Diameter Network Check

The Bulk Import/Export function of the Diameter Common application is useful for capturing configuration information for use by PCA.

For further details on how to use Bulk Import/Export, refer to Bulk Import and Export.



Diameter Network Check for Policy DRA

Identifying the Diameter network and properties for Policy DRA

- Identify and log the hardware profile type for each of the DA-MP Servers (PCA).
- 2. Identify and log the number of policy clients (PCEFs, BBERFs and AFs) and policy servers (PCRFs) in the network.
- 3. Identify and log the Diameter attributes for all the policy clients and policy servers in the network FQDN, Realm, IP address.
- Identify and log the type of Diameter Transport Protocol needed for all the policy clients and policy Servers - TCP/SCTP.
- Identify and log the type of Diameter connection mode needed for all the policy clients and policy server - Responder/Initiator/Responder-Initiator.
- Identify and log the Peer Node Identification for all the policy clients and policy servers - IP Address/FQDN.
- 7. Identify and log the route groups and route lists needed for Policy Servers and Policy Clients. Routing configuration is required for Policy Clients if the Policy Servers send Diameter request messages to be routed to the Policy Clients.
- 8. Identify and log the Policy Server configuration needed both that Gx and Rx are on same Policy Server or that they are on different servers.
- Identify and log the number of peer route tables needed for the Diameter configuration such as one for Rx Policy Servers and one for Gx Policy Servers.
- **10.** Identify and log the number of Application Route Table entries one for Gx Application and one for Rx Application message processing.
- 11. Identify and log the TSA used for local nodes if IPFE is used.

Identifying Diameter NOAM parameters for Policy DRA

- Identify and log the SBR Databases of Session and Binding types to be configured.
- Identify and log the Access Point Names used and the Stale Session Timeout for the same.
- 3. Identify and log the PCRF Pools and the Sub-Pool Selection Rules.



PCRF Sub Pool Selection Rules are optional.

- 4. Identify and log the General Options parameters for the Policy DRA network:
 - Default Stale Session Timeout
 - Binding Audit Session Query Rate
 - Audit Operation Rate
- 5. Identify and log the Network Wide Options parameters for the Policy DRA network:



- Early Binding Options
- Topology Hiding Options
- Suspect Binding Removal Options
- Session Integrity Options
- 6. Identify and log the Alarm Settings for the DSR Application Ingress Message Rate.
- 7. Identify and log the Congestion Alarm Thresholds and Message Throttling Rules.

Identifying Diameter SOAM parameters for Policy DRA

- 1. Identify and log the all the PCRFs handling the Policy Traffic for this site.
- 2. Identify and log the **Binding Key Priority** settings such as which binding keys will be used and in what order they will be used to correlate binding dependent session creation messages and route them to final bound PCRFs.
- 3. Identify and log the Policy Clients for which the topology hiding is needed.
- 4. Identify and log the PCRF Pool to PRT mapping configuration.
- Identify and log the error code configuration for each of the Error Conditions in the table per the policy client team request / interoperability requirements for the Policy Client vendor.
- 6. Identify and log the Suspect Binding Removal Rules.
- 7. Identify and log the Site Options for this site.

Diameter Network Check for Online Charging DRA

Identifying the Diameter network and properties for Online Charging DRA

- 1. Identify and log the hardware profile type for each of the DA-MP Servers (PCA).
- 2. Identify and log the number of Online Charging clients (CTFs) and Online Charging servers (OCSs) in the network.
- 3. Identify and log the diameter attributes for all the Online Charging clients and Online Charging servers in the network FQDN, Realm, IP address.
- 4. Identify and log the type of diameter Transport Protocol needed for all the Online Charging clients and Online Charging servers TCP/SCTP..
- Identify and log the type of diameter connection mode needed for all the Online Charging clients and Online Charging servers - Responder/Initiator/Responder-Initiator.
- 6. Identify and log the Peer Node Identification for all the Online Charging clients and Online Charging servers IP Address/FQDN.
- 7. Identify and log the route groups and route lists needed for Online Charging servers.
- **8.** Identify and log the number of peer route tables and peer route rules needed for the diameter configuration for Online Charging servers
- Identify and log the number of Application Route Table entries and for OC-DRA message processing.
- 10. Identify and log the TSA used for local nodes if IPFE is used.



Configuring Diameter NOAM parameters for Online Charging DRA

1. Identify and log the SBR Database of Session type to be configured.



Skip this step if Session type SBR Database was added during Policy DRA Function configuration.

- 2. Identify and log the Access Point Names used and the Stale Session Timeout for the same.
- Identify and log the General Options parameters for the Online Charging DRA network:
 - Default Stale Session Timeout
 - Audit Operation Rate
- 4. Identify and log the Online Charging Network Realms parameters for the Session State maintenance:
- 5. Identify and log the Network Wide Options for the Online Charging DRA network:
 - Session State Options
 - OCS Selection Options
- **6.** Identify and log the Alarm Settings for the DSR Application Ingress Message Rate.
- Identify and log the Congestion Alarm Thresholds and Message Throttling Rules.

Configuring Diameter SOAM parameters for Online Charging DRA

- 1. Identify and log the all the OCSs handling the Gy/Ro Traffic for this site.
- 2. Identify and log the all the CTFs to be configured for Session State maintenance.
- 3. Identify and log the error code configuration for each of the Error Conditions in the table for the Gy/Ro interface.

Health Check

Verifying Server status

1. Select Status & Manage, and then Server.

The Server Status screen is shown.

- 2. Verify all Server Statuses are Normal for Application Status, Alarms, Database, Collection, and Processes.
- 3. Do not proceed to PCA configuration if any of the statuses for Database, Collection, or Processes is not Normal. If any of these statuses are not Normal, corrective action should be taken to restore the non-Normal status to Normal before proceeding with the PCA configuration.



4. If the Alarm status is not Normal, but only Minor alarms are present, it is acceptable to proceed with the PCA configuration. If there are Major or Critical alarms present, analyze these alarms prior to proceeding with the PCA configuration. The activation can proceed in the presence of certain Major or Critical alarms.

Logging all current alarms

Viewing active alarms

Active alarms are displayed in a scrollable, optionally filterable table. By default, the active alarms are sorted by time stamp with the most recent alarm at the top.

Use this procedure to view active alarms.



The alarms and events that appear in View Active vary depending on whether you are logged in to an NOAM or **SOAM**. Alarm collection is handled solely by NOAM servers in systems that do not support SOAMs.

- 1. Click Alarms & Events, and then View Active.
- 2. If necessary, specify filter criteria and click Go.

The active alarms are displayed according to the specified criteria.

The active alarms table updates automatically. When new alarms are generated, the table is automatically updated, and the view returns to the top row of the table.

3. To suspend automatic updates, click any row in the table.

The following message appears: (Alarm updates are suspended.) If a new alarm is generated while automatic updates are suspended, a new message appears: (Alarm updates are suspended. Available updates pending.)

To resume automatic updates, press and hold **Ctrl** as you click to deselect the selected row.

Viewing alarm and event history

All historical alarms and events are displayed in a scrollable, optionally filterable table. The historical alarms and events are sorted, by default, by time stamp with the most recent one at the top. Use this procedure to view alarm and event history.



The alarms and events that appear in View History vary depending on whether you are logged in to an NOAM or SOAM. Alarm collection is handled solely by NOAM servers in systems that do not support SOAMs.

Click Alarms & Events, and then View History.



2. If necessary, specify filter criteria and click **Go**.



Some fields, such as **Additional Info**, truncate data to a limited number of characters. When this happens, a **More** link appears. Click **More** to view a report that displays all relevant data.

Historical alarms and events are displayed according to the specified criteria. The historical alarms table updates automatically. When new historical data is available, the table is automatically updated, and the view returns to the top row of the table.

3. To suspend automatic updates, click any row in the table.

The following message appears: (Alarm updates are suspended.) If a new alarm is generated while automatic updates are suspended, a new message appears: (Alarm updates are suspended. Available updates pending.)

To resume automatic updates, press and hold **Ctrl** as you click to deselect the selected row.

PCA Configuration

This section describes the **Policy and Charging**, and then **Configuration** screens.

Place and Place Association Configuration

The Policy and Charging Mated Sites type of Place Association is required for both Policy DRA and Online Charging DRA functions of PCA.

The Policy Binding Region type of Place Association is only required for Policy DRA function

Places

Places are used to build associations for groups of servers at a single geographic location. These places can then be grouped into place associations, which create relationships between one or more place.

Places Insert Elements

Table 5-9 describes the elements of the Places Insert screen.



Table 5-9 Places Insert Elements

Element	Description	Data Input Notes
Place Name A unique name used to label the place.	Format: Alphanumeric characters and underscore (_) are allowed. A minimum of one alphabetic character is required.	
		Range: Maximum length is 32 characters.
Parent	The parent of a place group	Format: List
		Note : This field is not used for PCA configuration. The only option is None.
Place Type	The place type.	Format: List
		Range: Site (default option).
Servers	List of the available servers in the NO or SO	Format: Checkbox
		Note : Select all of the DA-MP and SBR servers that are physically located at this Site Place.

Inserting a Place

Use this procedure to configure a place:

- 1. Click Configuration, and then Places.
- 2. Click Insert.
- 3. Enter the Place Name.
- 4. Select a Parent from the list.



A Parent Place is not required for PCA Places and can be set as **None**.

- 5. Select a **Place Type** from the list.
- **6.** Select all of the available DA-MP and SBR **Servers** that are physically located at this Site Place.
- 7. Click **OK** to submit the information and return to the Places screen, or click **Apply** to submit the information and continue adding additional data.

Editing a Place

Use this procedure to edit place information

- 1. Click Configuration, and then Places.
- 2. Select the place from the listing.
- 3. Click **Edit** at the bottom of the table.



- 4. Modify one or more of the place information fields.
- 5. Click **OK** to submit the information and return to the Places screen, or click **Apply** to submit the information and continue editing additional data.

The place information is updated in the network database and the changes take effect immediately.

Deleting a Place

Use this procedure to delete a place.

- 1. Click Configuration, and then Places.
- 2. Click to select the place you want to delete from the table.



A Place cannot be deleted if it includes servers. Before deleting, disassociate any servers.

- 3. Click Delete.
- 4. Click **OK** to delete the place.

If you click **Cancel**, the place is not deleted, and you are returned to the Places screen.

Generating a Places Report

Use this procedure to generate a places report:

- 1. Click Configuration, and then Places.
- 2. Click to select the place for which you want to create a report.



To select multiple server groups, press and hold **Ctrl** as you click to select specific rows. Alternatively, if no servers are selected then all server groups appear in the report.

- 3. Click Report.
- 4. Click **Print** to print the report, or click **Save** to save a text file of the report.

Place Associations

The Place Association function allows you to create relationships between places. Places are groups of servers at a single geographic location. For PCA, Place Associations are used to identify all sites that require access to the Policy DRA binding database, and to identify sites that share a PCA session database.

Place Association Insert Elements

Table 5-10 describes the elements of the Place Association Insert screen.



Table 5-10 Place Association Insert Elements

Element	Description	Data Input Notes
Place Association Name	A unique name used to label the place association.	Format: Alphanumeric characters and underscore (_) are allowed. A minimum of one alphabetic character is required.
		Range: Maximum length is 32 characters.
Place Association Type	The type of place association.	Format: List
		Range: defined by the application
Places	The places available to be grouped in this association.	Format: Option
		Range: list of places defined using Places function

Inserting a Place Association

Use this procedure to configure a place association:

- 1. Click Configuration, and then Place Association.
- 2. Click Insert.
- 3. Enter the Place Association Name.

For more information about **Place Association Name**, or any of the fields on this screen, see Place Association Insert Elements.

- 4. Optional: Select a **Place Association Type** from the list.
- Click OK to submit the information and return to the Place Associations screen, or click Apply to submit the information and continue adding additional data.

Editing a Place Associations

Use this procedure to edit place associations information

- 1. Click Configuration, and then Place Associations.
- 2. Select the place association from the listing.
- 3. Click **Edit** at the bottom of the table.
- **4.** Modify one or more of the place associations information fields.
- Click OK to submit the information and return to the Place Associations
 Configuration screen, or click Apply to submit the information and continue editing additional data.

The place association information is updated in the network database and the changes take effect immediately.

Deleting a Place Association

Use this procedure to delete a place association.



- 1. Click Configuration, and then Place Associations.
- 2. Click to select the place association you want to delete from the table.

Note:

You cannot delete a Place Association that includes Places. Before deleting the Place Association, disassociate the Places from the Place Association

- 3. Click Delete.
- 4. Click **OK** to delete the place association.

If you click **Cancel**, the place association are not deleted, and you are returned to the Place Association screen.

Generating a Place Associations Report

Use this procedure to generate a place associations report:

- 1. Click Configuration, and then Place Associations.
- 2. Click to select the place associations for which you want to create a report.
- Click Report.
- 4. Click **Print** to print the report, or click **Save** to save a text file of the report.

Resource Domain Configuration

The Policy and Charging DRA and Policy Session types of Resource Domains are required for both Policy DRA and Online Charging DRA functions of PCA.

A Resource Domain is a collection of Server Groups that share something in common. PCA uses the Policy Session Resource Domain to identify server groups that host an instance of a Policy and Charging session database. The Policy Binding Resource Domain identifies server groups that host the Policy DRA binding database. The Policy and Charging DRA Resource Domain identifies DA-MP server groups that require access to an instance of a Policy and Charging Session Database.

Resource Domains

The Resource Domains function allows you to assign servers to domains.

Add New Resource Domain Elements

Table 5-11 describes the elements for adding a resource domain element:



Table 5-11 Add New Resource Domain Elements

Element	Description	Data Input Notes
Resource Domain Name	The name for the resource domain.	Format: Alphanumeric (A-Z, a-z, 0-9) and underscore (_) characters.
		Range: Maximum length is 32 characters
		Default: N/A
Resource Domain Profile The profile as	The profile associated with the	Format: List
	resource domain.	Range: Policy Binding, Policy Session, and Policy and Charging DRA
Server Groups	The server groups associated with the resource domain	Format: Checkbox
		Range: MPSG, NOSG< SBRSG, SBSG2, SOSG

Inserting a Resource Domain

Use this procedure to insert a resource domain:

- 1. Click Configuration, and then Resource Domains.
- 2. Click **Insert** at the bottom of the table.
- 3. Enter a **Resource Domain Name**. This is a user-defined name for the domain. The domain name must be unique.
- 4. Select a Resource Domain Profile.
- 5. Select a Server Group.
- 6. Click **OK** to submit the information and return to the Resource Domains Configuration screen, or click **Apply** to submit the information and continue entering additional data.

The resource domain is added to the network database.

Editing a Resource Domain

Use this procedure to edit resource domain information

- 1. Click Configuration, and then Resource Domains.
- 2. Select the resource domain from the listing.
- 3. Click **Edit** at the bottom of the table.
- 4. Modify one or more of the resource domain information fields.
- Click OK to submit the information and return to the Resource Domains
 Configuration screen, or click Apply to submit the information and continue editing additional data.

The resource domain information is updated in the network database and the changes take effect immediately.



Deleting a Resource Domain

Use this procedure to delete a resource domain:

- 1. Click Configuration, and then Resource Domains.
- 2. Click to select the resource domain you want to delete.



To prevent large service disruptions, you cannot delete a Resource Domain with a profile type or Policy Binding or Policy Session, unless the Policy DRA feature is deactivated. However, resource domains with a profile type of Policy DRA can be deleted without deactivation of the Policy DRA feature.

3. Click Delete.

Click Yes to confirm.

The resource domain is deleted from the network database table.

Generating a Resource Domains Report

Use this procedure to generate a resource domains report:

- 1. Click Configuration, and then Resource Domains.
- 2. Click to select the resource domain for which you want to create a report.



To select multiple server groups, press and hold **Ctrl** as you click to select specific rows. Alternatively, if no servers are selected then all server groups appear in the report.

- 3. Click Report.
- 4. Click Print to print the report, or click Save to save a text file of the report.

PCA Routing of Diameter Messages

PCA routes Diameter messages depending on certain criteria:

- Answer message or Request message
- · New session Request or in-session Request
- New binding or existing binding new session Request

Peer Routing

PCA always attempts to route using Peer Route Tables. The Diameter Routing Function attempts to use Peer Route Tables in a predefined precedence:



- Peer Route Table configured for the originating Peer Node (Diameter, and then Configuration, and then Peer Nodes)
 If a match is found, the specified Peer Route Table is used.
- Peer Route Table configured for the Diameter Application-ID of the policy session initiation request being routed (Diameter, and then Configuration, and then Application Ids)
 If the ingress Peer Node is configured as Not Selected, that entry is skipped and
- 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, PCA will attempt to route the Request using Destination-Host routing (for example, to a connection or Alternate Implicit Route List associated with the destination Peer Node).

Routing of Session Initiation Requests for New Bindings

the Application Ids configuration is checked.

PCA allows a Peer Route Table to be configured for use when a new binding is created. This Peer Route Table can specify Peer Routing Rules to:

- Allow new bindings to be routed, for example, based on the Origin-Host or Origin-Realm of the PCEF
- Cause new bindings to be load-shared across all local PCRFs.

The Peer Route Table to use for new bindings is specified in the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Site Options** screen on the SOAM at each site.

If the Peer Route Table for new bindings is set to Not Selected, the Diameter Routing Function uses the precedence described in Peer Routing.

Routing of Session Initiation Requests for Existing Bindings

Sessions for subscribers that are already bound to a PCRF must be routed to the bound PCRF, or to a PCRF that shares state with the bound PCRF if the PCRF supports sharing of policy state. For existing bindings, no Peer Route Table is configured in the PCA application Site Options. Instead, the Diameter Routing Function uses the precedence described in Peer Routing.

Routing of Requests from PCRF to a Policy Client

In order to route Requests initiated by the PCRF, routing must be configured such that Requests from any PCRF can be routed to any Policy Client in the network. This type of routing is used to route RAR and ASR requests. For Requests from PCRFs to Policy Clients, no Peer Route Table is configured in the PCA application Site Options. Instead, the Diameter Routing Function uses the precedence described in Peer Routing.

Routing of In-Session Requests

In-session Requests are Requests within a Diameter session other than the Request that established the Diameter session. CCR-U, CCR-T, and STR are all examples of in-session Requests. In-session Requests are routed using the predefined precedence of Peer Route Tables described in Peer Routing.



Routing of Answer Messages

All Diameter Answer messages are routed over the same path on which the Request was routed, using hop-by-hop routing. No routing configuration is necessary to route Answer messages.

Inter DSR Routing

If Diameter messages need to be routed in between DSR sites (nodes), set up the routing as needed.

This is likely in 3-site redundancy deployments because many PCEFs likely only support primary and secondary connections. In such deployments routing can be set up between the three sites.

Routing for Gx RAR Messages (PDRA Generated)

Routing Rules can be configured to route a Gx RAR message generated at one site that is destined for a PCEF connected to another site.

This is likely in 3-site redundancy deployments because many PCEFs likely only support primary and secondary connections. In such deployments routing can be set up between the three sites.



Destination-Host based routing can be set up to route the Gx RAR messages to the appropriate site's DSR.

Diameter Configuration for PCA

The PCA application requires configuration of several Diameter Configuration components before the PCA configuration can be performed.

All Diameter Configuration components are configured using the SOAM GUI.

Use the explanations and procedures in the Diameter Configuration online help and the *Diameter User Guide* to complete the configuration of the Diameter Configuration components for the system, including the Diameter Configuration components for use with the PCA application.

1. Application Ids

Use the **Diameter**, and then **Configuration**, and then **Application Ids [Insert]** screen to define an Application Id for each Diameter interface that will be used by PCA in the system.

PCA supports values that can be selected in the **Application Id Value** pulldown list:

- 4 Diameter Credit Control
- 4 3GPP Gy/Ro
- 16777236 3GPP Rx



- 16777238 3GPP Gx
- 16777238 3GPP Gx-Prime
- 16777266 3GPP Gxx
- 16777267 3GPP S9
- 4294967295 Relay

Note:

Gx-Prime shares the same Application Id as Gx. To distinguish between them, the content of the Diameter message is checked against a configured Application Routing Table to determine if the message originates from a Gx or Gx-Prime interface.

PCA always attempts to route using Peer Route Tables. The Peer Route Table can be configured here for each Application Id, or can be configured for Peer Nodes. If neither is configured, the Default Peer Route Table will be used. See PCA Routing of Diameter Messages.

2. CEX Parameters

Use the **Diameter**, and then **Configuration**, and then **CEX Parameters [Insert]** screen to define the Capability Exchange parameters for each Application Id that was configured for use by PCA:

For each Application Id, select or enter:

- Application Id Type Authentication
- Vendor Specific Application Id, if the Application Id and Vendor Id will be grouped in a Vendor-specific Application Id AVP
- Vendor Id if Vendor Specific Application Id is selected
 The Vendor ID 10415 is defined in 3GPP:
 - 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)
 - Gy/Ro: 4 with Vendor-Id of 10415 (defined in 3GPP 32.299)

3. CEX Configuration Sets

Use the **Diameter**, and then **Configuration**, and then **Configuration Sets**, and then **CEX Configuration Sets [Insert]** screen 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 for connections with the PCRF nodes (Gx and Rx)



- A CEX Configuration Set to be used for connections with the OCS nodes (Gy/Ro)
- CEX Configuration Sets to be used with any other types of nodes, such as BBERF (Gxx)
- A CEX Configuration Set named Default is provided for the Relay Application Id; it can be edited if needed.

4. Transaction Configuration Sets

Use the **Diameter**, and then **Configuration**, and then **Configuration Sets**, and then **Transaction Configuration Sets** [Insert] screen to configure Transaction Configuration parameters.

The transaction configuration, once configured, needs to be set in the **Diameter**, and then **Configuration**, and then **Peer Nodes** screen.

5. Local Nodes

Use the **Diameter**, and then **Configuration**, and then **Local Nodes [Insert]** screen to configure the Diameter identity or identities by which the PCA DSR will be known to peer nodes.

The pulldown list of **IP Addresses** contains the XSI addresses configured on DSR MP Servers. If IPFE is being used, a local node should be created for each IPFE TSA.

6. Peer Nodes

Use the **Diameter**, and then **Configuration**, and then **Peer Nodes [Insert]** screen to configure PCEFs, AFs, BBERFs, and any other types of nodes as Peer Nodes to the PCA DA-MPs in the system. (PCA DA-MPs can also be Peer Nodes to each other at different sites). Additionally, select a **Transaction Configuration Set** from the pulldown list.



For PCA, the **Replace Dest Realm** and **Replace Dest Host** boxes should be checked for PCRF and OCS peer nodes.

See PCA Routing of Diameter Messages for details on routing of messages for PCA.

7. Connections

Use the **Diameter**, and then **Configuration**, and then **Connections [Insert]** screen to configure connections between the PCA DA-MPs and the Peer Nodes.

Any IPFE Target Set Address (TSA) that is used to configure a connection must use the same **Transport Protocol** (SCTP or TCP) that is selected to configure the connection.

8. Route Groups

Use the **Diameter**, and then **Configuration**, and then **Route Groups [Insert]** screen to configure Route Groups for use with PCA Peers.

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



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

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

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

9. Route Lists

Use the **Diameter**, and then **Configuration**, and then **Route Lists [Insert]** screen to configure Route Lists for use with the configured Route Groups.

For priority-based initial session binding, configure N Route Lists where N is the number of PCRF/OCSs in the system.

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

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

Enter 1 for the Minimum Route Group Availability Weight in all of the Route Lists.

10. Peer Route Table and Peer Routing Rules

Use the **Diameter**, and then **Configuration**, and then **Peer Route Tables [Insert]** screen to configure new Peer Route Tables if needed, and the Viewing Rules for Peer Route Table screen to configure Peer Routing Rules, such that DSR forwards messages based on the PCRF/OCS preference.

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

See PCA Routing of Diameter Messages for details on PRT routing of PCA messages.

The routing configuration will ensure that whenever PCA requests Diameter to route to a particular PCRF/OCS based on the PRT:

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

11. Application Route Tables and Application Routing Rules

Use the **Diameter**, and then **Configuration**, and then **Application Route Tables [Insert]** screen to configure new Application Routing Rules, if needed for each Diameter interface (such as GxGx-Prime, or Rx) that is configured in an Application Name, to be used for Diameter routing of messages to the PCA application. PCA must receive all Diameter Requests.

Use the Viewing Rules for Applicatin Route Table screen 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, and then Configuration, and then Application Ids (Filtered) screen for configured for PCA.
- Application Name PCA
- Gx-Prime
- Application Route table

Policy DRA Configuration

This section describes the **Policy and Charging**, and then **Configuration**, and then **Policy DRA** screens on the NOAM and the SOAM.

PCRFs

The **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRFs** screen contains the list of PCRF Peer Nodes that are to be used when a subscriber binding is created at this site. New bindings created at this Policy and Charging DSR are distributed evenly among the configured PCRFs.

PCRFs are responsible for authorizing and making policy decisions based on knowledge of subscriber resource usage and the capabilities allowed by the subscriber's account. All policy requests for a given subscriber must be routed to the same PCRF. Policy and Charging dynamically assigns subscribers to PCRFs using a load distribution algorithm, and maintains state about which subscribers are assigned to which PCRF. The relationship between a subscriber and a PCRF can change any time the subscriber transitions from having no Diameter policy sessions to having one or more Diameter policy sessions. After a policy session exists, all policy sessions for that subscriber are routed to the assigned PCRF.

The fields are described in PCRFs elements.



For details about configuring Peer Nodes, refer to the *Diameter User Guide* and Diameter online help.

On the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRFs** screen on the SOAM, you can perform a variety of 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 Insert.

The **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRFs [Insert]** screen opens and allows the user to add a PCRF. See **Inserting PCRFs**. If the maximum number of PCRFs (5000) already exists in the system, the **Policy and Charging**, and then **Configuration**, and then **Policy**



DRA, and then **PCRFs** [Insert] screen will not open, and an error message is displayed.

Select a PCRF in the list, and click Edit.

The **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRFs [Edit]** screen opens and allows the user to edit the selected PCRF. See Editing PCRFs.

 Select a PCRF in the list, and click **Delete** to remove the selected PCRF. See Deleting a PCRF.

PCRFs elements

Table 5-12 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRFs** screen on the Active SOAM.



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

Table 5-12 PCRFs Screen 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, and then Configuration, and then Peer Nodes (Filtered) screen where Diameter Peer Nodes are filtered by the PCRF Peer Node Name.	Format: Format: List
		Range: Configured Diameter Peer Nodes
		Note: The PCRF Peer Node Name cannot be changed on the [Edit] screen.
		trie (Euit) screen.
Comments	An optional comment to describe the PCRF Peer Node.	Format: Text box Range: 0-64 characters

Inserting PCRFs

Use this task to insert (create new) PCRFs.

The fields are described in PCRFs elements.

- 1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRFs**.
- 2. Click Insert.
- 3. Enter a unique PCRF Peer Node Name in the PCRF Peer Node Name field.



This name uniquely identifies the PCRF Peer Node to be included in the load distribution of new bindings to PCRFs.

4. Enter an optional comment in the **Comments** field.

5. Click:

- OK to save the new PCRF and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRFs screen.
- **Apply** to save the new PCRF and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRFs screen without saving any changes.

If **OK** or **Apply** is clicked and any of several conditions exist, an error message appears:

- The entered PCRF is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- Adding the new PCRF would cause the maximum number of PCRFs (5000) to be exceeded.

Editing PCRFs

Use this task to edit PCRF Comments.



The PCRF Pool Name cannot be edited.

1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRFs**.

The screen displays a list of the configured PCRF Peer Nodes that are used when a new subscriber binding is created.

2. Click in the **Comments** field of the row to select the PCRF to edit.

DO NOT click the blue PCRF Peer Node Name (except to see the configuration of the Peer Node). The blue color indicates a hyper-link that opens the **Diameter**, and then **Configuration**, and then **Peer Nodes [Filtered]** screen to display the configuration information for the Peer Node.

3. Edit the Comments field for the selected PCRF.

The PCRF Peer Node name cannot be changed.

- 4. Click:
 - OK to save the change and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRFs screen.
 - Apply to save the change and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRFs screen 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 this procedure to delete a PCRF.

This procedure describes the recommended steps for deleting a PCRF from a Policy and Charging configuration. In this procedure, PCRF refers to a Diameter peer of the PCA, which is sometimes referred to as a PCRF Front-end.

The PCRF procedure minimizes disruption to policy signaling by:

- Preventing sessions from creating new bindings to a PCRF that has been removed
- Allowing sessions with existing bindings to continue to use a PCRF that has been removed until those sessions terminate normally

The procedure describes the recommended steps for deletion of a PCRF from a Policy and Charging configuration. In this procedure, PCRF refers to a Diameter peer of the PCA, sometimes referred to as a PCRF Front-End.



The PCRF removal procedure is restricted to SOAM servers.

- Use Main Menu, and then Diameter, and then Configuration, and then Peer Nodes from the SOAM GUI screen to determine the Peer Node name of the PCRF(s) being removed.
- Use Main Menu, and then Diameter, and then Route Groups from the SOAM GUI screen, 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 screen, 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 screen, 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, and then Policy and Charging, and then Policy DRA, and then Configuration, and then PCRFs from the SOAM GUI screen 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**, and then **Policy and Charging**, and then **Policy DRA**, and then **Configuration**, and



then **PCRFs** on the SOAM GUI screen to delete the route group containing only the PCRF peer being removed.

- Use Main Menu, and then Diameter, and then Maintenance, and then Connections from the SOAM GUI screen to find the connection for the PCRF Peer Node and disable it
- Use Main Menu, and then Diameter, and then Maintenance, and then Connections from the SOAM GUI screen to delete the connection to the PCRF Peer Node.
- Use Main Menu, and then Diameter, and then Configuration, and then Peer Nodes from the SOAM GUI screen to delete the Diameter Peer Node for the PCRF being removed.

Binding Key Priority

The Binding Key Priority defines search priorities for Alternative Keys that can be used to locate a subscriber binding.

The Binding Key Priority controls:

- · Which keys are stored for binding correlation
- The order in which keys are searched for purposes of binding correlation

The priority determines the order used to find a binding for subsequent sessions. Alternative Keys with an assigned priority will be created with the binding if they are present in the session initiation message that created the binding. The Alternative Keys must be assigned a priority in order to be used to locate subscriber bindings. If any Alternative Keys are not assigned a priority, they will not be used to locate subscriber bindings even if the Alternative Key is present in the session initiation message.

The fields are described in Binding Key Priority elements.

On the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Binding Key Priority** screen on the Active SOAM, you can change the Binding Key Type for Binding Key Priority 2, 3, and 4.



Priority 1 for Binding Key Type IMSI is the highest priority and cannot be modified.

Enabling and disabling the binding key field depends on the value that you select for the Binding Key type.

Binding Key Priority elements

Table 5-13 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Binding Key Priority** screen.



Table 5-13 Binding Key Priority Elements

Field (* indicates a requried field)	Description	Data Input Notes
* Binding Key Type	The Binding Key Type which	Format: Pulldown list
	is assigned to a Binding Key Priority.	Range: MSISDN, IPv4, or IPv6 for Priority 2, 3, and 4
	GUID-7C4F863F-9E67-4B6C-89EE-B44E01292AF6The first row is Priority 1 and the corresponding Binding Key Type is IMSI. This row is readonly.	Default: -Select- (No Binding Key Type selected)

Setting Binding Key Priority

Use this task to set Binding Key Priority values.

The fields are described in Binding Key Priority elements.

- On the Active SOAM, click Policy and Charging, and then Configuration, and then Policy DRA, and then Binding Key Priority.
- 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 screen
 - Cancel to remain on the Policy and Charging, and then Configuration, and then Policy DRA, and then Binding Key Priority screen without saving any changes.

If **Apply** is clicked and any of several 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

Network-Wide Options

Click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Network-Wide Options** on an Active NOAM to configure Network-Wide Options.

The fields are described in Network-Wide Options elements.

General Options

- Select the PCRF Pooling Mode
- Select the Default APN for Non Specific Binding Key Correlation

Early Binding Options

- Set the Early Binding Polling Interval value
- Set the Maximum Early Binding Lifetime value



Topology Hiding Options

- Enable Topology Hiding
- Set the Topology Hiding Scope
- Set the Default Topology Hiding Virtual Name

Suspect Binding Removal Options

- Set the Suspect Binding Removal Events Ignore Interval
- Set the Suspect Binding Removal Events Reset Interval
- Set the Suspect Binding Removal Events Threshold

Session Integrity Options

- Indicate whether to use the Local Host Origin-Host and Origin-Realm or the PCRF Origin-Host and Origin-Realm as the Origin-Host and Origin-Realm in RAR messages that are constructed and sent by Policy DRA to the Policy Clients.
- Set the Maximum Query RAR Rate Per Session Server Group value
- Set the Maximum Attempts Per Query RAR value
- Set the Maximum Release RAR Rate Per Session Server Group value
- Set the Maximum Attempts Per Release RAR value
- Set the Query RAR Queue Capacity Per Session Server Group value
- Set the Release RAR Queue Capacity Per Session Server Group value

Network-Wide Options elements

Table 5-14 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Network-Wide Options** screen on the NOAM.

Table 5-14 Policy DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
General Options		
Enable PCRF Pooling	Indicates whether the PCRF Pooling feature is enabled.	Format: Checkbox Default: PCRF Pooling
	Check the box to allow a subscriber's policy sessions to be routed to different PCRFs depending on the Access Point Network the session originated from. this box must be checked following acceptance of upgrade or future upgrades will be disallowed.	Enabled (checked) for fresh installs, PCRF Pooling Disabled (unchecked) Range: Checked or Unchecked



Table 5-14 (Cont.) Policy DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
PCRF Pooling Mode	Indicates whether PCRF Pooling should operate in Single Pool or Multi Pool mode. Select Single Pool Mode if several items are true: IMSI or MSISDN are being used by at least one AF for binding correlation At least one AF does not include APN in session creation messages Otherwise, select Multi Pool Mode.	Format: Radio button Default: Multi Pool Mode Range: Single Pool Mode or Multi Pool Mode
	When Single Pool Mode is selected, all new binding capable session creation messages are routed using the Default PCRF Pool, regardless of the APN to PCRF Pool mapping.	
	When Multi Pool Mode is selected, new binding capable session creation messages are routed using the PCRF Pool associated with the APN in the session creation message.	



Table 5-14 (Cont.) Policy DRA Network-Wide Options Elements

Fields (* indicates a		
required field)	Description	Data Input Notes
Fields (* indicates a required field) Default APN for Non Specific Binding Key Correlation	APN to use for binding correlation if no APN is present in a binding dependent session creation request that needs to correlate using an IMSI or MSISDN key. If several items are true: IMSI or MSISDN are being used by at least one AF for binding correlation At least one AF does not include APN in session creation messages All bindings to be correlated using IMSI or MSISDN for AF(s) that do not include APN in the session creation messages are associated with a single APN Select the APN associated with all IMSI or MSISDN bindings to be used by AF(s) that do not include APN in the session creation messages. Otherwise, leave this field set toSelect If this field is enabled, no	Format: Drop-down list Default: N/A Range: List of configured APNs
	If this field is enabled, no APN is selected and a binding dependent session creation request is received that includes an IMSI or MSISDN key, neither IMSI nor MSISDN, will be used for binding correlation. Default APN for Non Specific Binding Correlation is disabled if Single Pool Mode is selected.	

Early Binding Options



Table 5-14 (Cont.) Policy DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
* Early Binding Polling Interval	The number of milliseconds between sending queries to the early binding master to determine which PCRF the master session was routed to so that the slave session can be routed to the same PCRF. The goal is to set this value such that the master session has time to receive an answer a high percentage of the	Format: Text box Default: 200 milliseconds Range: 50 to 10000 milliseconds
	time. Choosing a low value increases database queries, but may reduce latency. A high value does the opposite.	
	Note : This values is used only when PCRF Pooling is enabled.	
* Maximum Early Binding Lifetime	The maximum time that a binding is allowed to remain as an early binding. The ideal setting for this value	Format: Text box Default: 2500 milliseconds Range: 500 to 15000
	is 100 - 200 msec longer than the Diameter transaction timeout. This value prevents bindings from becoming stuck for long periods in the early binding state due to congestion or other error conditions. If a new Diameter request or polling attempt discovers a binding session that has been in the early state for longer than this time, the binding session is removed.	milliseconds
	Note : This value is used only when PCRF Pooling is Enabled.	
Topology Hiding Options		
Enable Topology Hiding	Enable or disable topology hiding using the check box. Once enabled or disabled here, the Topology Hiding is enabled or disabled at all SOAMs under this NOAM.	Format: Check box Default: Disabled (unchecked) Range: Enabled (checked), Disabled (unchecked)



Table 5-14 (Cont.) Policy DRA Network-Wide Options Elements

Fields (* indicates a		
required field)	Description	Data Input Notes
Topology Hiding Scope	messages where topology hiding will be applied. Select All Messages to perform topology hiding for all messages destined to policy clients. Select All Foreign Realms to perform topology	Range: All Messages, All Foreign Realms, Specific Clients, All Foreign Realms +
	hiding for messages destined to the policy clients whose realms are different from the realm of the PCRF to be bound. Select Specific Clients to perform topology hiding for the policy clients that are configured in on the SOAM GUI Main Menu:	
	Policy and Charging, and then Configuration, and then Policy DRA, and then Policy Clients screen. Select All Foreign Realms + Specific Clients to perform topology hiding if either condition (All Foreign Realms or Specific Clients) is met.	



Table 5-14 (Cont.) Policy DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
Default Topology Hiding	FQDN - This FQDN is	Format: Text box
Virtual name	used as a default value in the Origin-Host AVP	Default: N/A
	for answer messages routed from a PCRF to a policy client, or in the Destination-Host AVP for request messages routed from a PCRF to a policy client, only if Topology Hiding Virtual Name FQDN is not configured at a SOAM relevant to the policy client and PCRF. Realm - This Realm is used as a default value in the Origin-Realm AVP for answer messages routed from a PCRF to a policy client, or in the Destination-Realm AVP for request messages routed from a PCRF to a policy client, only if Topology Hiding Virtual Name Realm is not configured at a SOAM relevant to the policy client and PCRF.	Range: FQDN and Realm - a case-insensitive string consisting of a list of labels separated by dots, where a label may contain letters, digits, dashes (-) and underscore (_). A label must start with a letter, digit or underscore and must end with a letter or digit. Underscores may be used only as the first character. A label must be at most 63 characters long and a FQDN must be at most 255 characters long.
Suspect Binding Removal Option	ons	
* Suspect Binding Removal Events Ignore Interval	This value can be used to ignore Suspect Binding Removal Events that arrive in quick succession.	Format: Text box Default: 1 Range: 0-30 sec
	If a Suspect Binding Removal Event arrives for a given binding, but the time interval specified Suspect Binding Removal Events Ignore Interval has not yet elapsed, the event will not be counted against the Suspect Binding Removal Events Threshold. Setting the value to zero means that no events are ignored. It may be desirable to select a value such that retransmissions are ignored.	

Table 5-14 (Cont.) Policy DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
* Suspect Binding Removal Events Reset Interval	This value is used to separate occurrences of PCRF unavailability. The interval is started at the time when the last suspect binding event is counted. If no subsequent suspect binding event occurs after this interval, the suspect binding event count is reset and the suspect binding associated with the count is no longer considered as suspect until the next suspect binding event occurs.	Format: Text box Default: 60 min Range: 1-3600 min
* Suspect Binding Removal Events Threshold	The Suspect Binding Removal Events Threshold value can be used to avoid triggering suspect binding removal for transient events (for example, Diameter timeouts).	Format: Text box Default: 3 Range: 2-10
	This value specifies the number of Suspect Binding Removal Events that must occur for a given binding before the system will attempt to remove the binding by initiating a Session-Release RAR towards the Policy Client.	
	This value is used when the Remove Suspect Binding Immediately value in the matched rule on the Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules screen on the SOAM is set to No.	
Session Integrity Options		
Origin-Host and Origin-Realm for Policy DRA generated RAR messages	The selection option's Origin- Host and Origin-Realm will be used as the Origin-Host and Origin-Realm in the RAR messages constructed and sent by Policy DRA to the Policy Clients.	Format: Radio button Default: Local Host Range: Local Host or PCRF



Table 5-14 (Cont.) Policy DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
* Maximum Query RAR Rate Per Session Server Group	This value specifies the maximum rate in messages per second at which a given Session SBR Server Group can send RAR message to Policy Clients for the purpose of auditing to determine if the session still exists.	Format: Text box Default: 50 Range: 50-5000
* Maximum Attempts Per Query RAR	This value specifies the maximum number of times a given RAR will attempt to be sent to the Policy Client for purposes of querying for session existence when no response is received. If no response is received after the Maximum Attempts Per Query RAR has been reached, the Diameter session and associated binding keys are removed automatically.	Format: Text box Default: 12 Range: 1-12
* Maximum Release RAR Rate Per Session Server Group	This value specifies the maximum rate in messages per second at which a given Session SBR Server Group can send RAR message to Policy Clients for the purpose of requesting removal of a session. Session removal is requested if a session or its associated binding keys cannot be successfully stored in the SBR database, or when a Suspect Binding is to be removed due to PCRF inaccessibility.	Format: Text box Default: 50 Range: 50-5000
* Maximum Attempts Per Release RAR	This value specifies the maximum number of times a given RAR will attempt to be sent to the Policy Client for purposes of requesting removal of a session when no response is received. If no response is received after the Maximum Attempts Per Release RAR has been reached, the Diameter session and associated binding keys are removed automatically	Format: Text box Default: 12 Range: 1-12



Table 5-14 (Cont.) Policy DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
* Query RAR Queue Capacity Per Session Server Group	This value specifics the maximum number of RARs that can be queued in a given Session SBR Server Group for sending to Policy Clients for the purpose of querying for session existence. If a query RAR cannot be queued because the Pending Query RAR Capacity Per Session Server Group has been reached, the next pass of the session audit will attempt to queue the query RAR again.	Format: Text box Default: 1000 Range: 1000-50,000
* Release RAR Queue Capacity Per Session Server Group	This value specifies the maximum number of RARs that can be queued in a given Session SBR Server Group for sending to Policy Clients for the purpose of requesting removal of sessions. If a release RAR cannot be queued because the Pending Release RAR Capacity Per Session Server Groups has been reached, another attempt to queue the release RAR will occur the next time a Suspect Binding Removal Even occurs for that binding.	Format: Text box Default: 100,000 Range: 100,000-500,000

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 Policy DRA configuration options apply to the entire Policy DRA Network:

- Origin-Host and Origin-Realm for Policy DRA generated RAR messages
- PCRF Pooling Modes
- Early Binding
- Topology Hiding
- Click Policy and Charging, and then Configuration, and then Policy DRA, and then Network-Wide Options.
- 2. Select the PCRF Pooling Mode.

Indicates whether PCRF Pooling operates in Single Pool Mode or Multi Pool mode.



3. Select a Default APN for Non Specific Binding Key Correlation.

Select the APN to use for binding correlation if no APN is present in a binding dependent session creation request that needs to correlate using an IMSI or MSISDN key

4. Set the Early Binding Polling Interval.

The number of milliseconds between sending queries to the early binding master to determine which PCRF the master session was routed to so that the slave session can be routed to the same PCRF.

The goal is to set this value such that the master session has time to receive an answer a high percentage of the time. Choosing a low value increases database queries, but may reduce latency. A high value does the opposite.

Note:

This value is used only when PCRF Pooling is Enabled.

Set the Maximum Early Binding Lifetime.

The maximum time that a binding is allowed to remain as an early binding.

The ideal setting for this value is 100 - 200 msec longer than the Diameter transaction timeout. This value prevents bindings from becoming stuck for long periods in the early binding state due to congestion or other error conditions. If a new Diameter request or polling attempt discovers a binding session that has been in the early state for longer than this time, the binding session is removed.

Note:

This value is used only when PCRF Pooling is Enabled.

6. Select the **Enable Topology Hiding** check box.

Enable or disable topology hiding using the check box. Once enabled or disabled here, the Topology Hiding is enabled or disabled at all SOAMs under this NOAM.

7. Select a Topology Hiding Scope.

This sets the scope of messages where topology hiding will be applied. Select **All Messages** to perform topology hiding for all messages destined to policy clients. Select **All Foreign Realms** to perform topology hiding for messages destined to the policy clients whose realms are different from the realm of the PCRF to be bound. Select **Specific Clients** to perform topology hiding for the policy clients that are configured in one of SOAM GUI Main Menu: **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Policy Clients** screen. Select **All Foreign Realms + Specific Clients** to perform topology hiding if either condition (**All Foreign Realms** or **Specific Clients**) is met.

8. Enter a Default Topology Hiding Virtual Name.

The entered FQDN is used as a default value in the Origin-Host AVP for answer messages routed from a PCRF to a policy client, or in the Destination-Host AVP for request messages routed from a PCRF to a policy client, only if Topology Hiding Virtual Name FQDN is not configured at a SOAM relevant to the policy client and PCRF.



The entered Realm is used as a default value in the Origin-Realm AVP for answer messages routed from a PCRF to a policy client, or in the Destination-Realm AVP for request messages routed from a PCRF to a policy client, only if Topology Hiding Virtual Name Realm is not configured at a SOAM relevant to the policy client and PCRF.

9. Enter a Suspect Binding Removal Events Ignore Interval value.

The interval is used to ignore Suspect Binding Removal Events that arrive in quick succession

10. Enter a Suspect Binding Removal Events Reset Interval value.

The interval is used to separate occurrences of PCRF unavailability

11. Enter a Suspect Binding Removal Events Threshold value.

The threshold can be used to avoid triggering suspect binding removal for transient events (for example, Diameter timeouts).

12. Select the Local Host or PCRF radio button.

This sets the Origin-Host and Origin-Realm that will be used in the RAR messages constructed and sent by Policy DRA to policy clients.

13. Enter a Maximum Query RAR Rate Per Session Server Group value.

The value specifies the maximum rate in messages per second at which a given Session SBR Server Group can send RAR messages to Policy Clients for the purpose of auditing to determine if the session still exists.

14. Enter a Maximum Attempts Per Query RAR value.

The value specifies the maximum number of times a given RAR will be attempted to be sent to the Policy Client for purposes of querying for session existence when no response is received.

15. Enter a Maximum Release RAR Rate Per Session Server Group value.

The value specifies the maximum rate in messages per second at which a given Session SBR Server Group can send RAR messages to Policy Clients for the purpose of requesting removal of a session.

16. Enter a Maximum Attempts Per Release RAR value.

The value specifies the maximum number of times a given RAR will be attempted to be sent to the Policy Client for purposes of requesting removal of a session when no response is received.

17. Enter a Query RAR Queue Capacity Per Session Server Group value.

The value specifies the maximum number of RARs that can be queued in a given session SBR Server Group for sending to Policy Clients for the purpose of querying for session existence.

18. Enter a Release RAR Queue Capacity Per Session Server Group value.

The value specifies the maximum number of RARs that can be queued in a given session SBR Server Group for sending to Policy Clients for the purpose of requesting removal of sessions.

19. Click:

Apply to save the changes and remain on this screen.



 Cancel to discard changes and remain on the Policy and Charging, and then Configuration, and then Policy DRA, and then Network-Wide Options screen.

If **Apply** is clicked and the **Enable PCRF Pooling** checkbox transitioned from unchecked to checked, then an error message appears.

Policy Clients

Topology hiding configuration is performed at both the network level using the NOAM GUI and at the site level using the SOAM GUI.

The fields are described in Policy Clients elements.

SOAM Options

On the SOAM GUI, use the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Policy Clients** screen to define the list of Policy Client Peer Nodes from which the PCRF name is to be hidden. This screen can be used only if Topology Hiding is **Enabled** and the **Topology Hiding Scope** option is either **Specific Clients** or **All Foreign Realms + Specific Clients** on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Network-Wide Options** screen on the NOAM GUI. See **Site Options** for additional information.

- Filter the list of Policy Client Peer Node Names, to display only the desired Policy Client Peer Node Names.
- Sort the list entries in ascending or descending order by Policy Client Peer Node Names or by Comments, by clicking the column heading. By default, the list is sorted by Policy Client Peer Node Names in ascending numerical order.
- Click Insert.
 - You can add a Policy Client Peer Node Name and Comment. See Adding a New Policy Client for Topology Hiding. If the maximum number of Policy Client Peer Nodes (1000) already exists in the system, the Policy and Charging, and then Configuration, and then Policy DRA, and then Policy Clients [Insert] screen 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 Edit. (Clicking the blue Policy Client Peer Node Name will open the filtered Diameter, and then Configuration, and then Peer Nodes screen for the Peer Node.)
 - You can edit the **Comment** for the selected **Policy Client Peer Node Name**. (The Policy Client Peer Node Name cannot be changed). See Editing Policy Clients for Topology Hiding.
- Select the Comment in the row for a Policy Client Peer Node Name in the list, and click Delete to remove the selected Policy Client Peer Node Name. See Deleting a Topology Hiding Policy Client Peer Node.

Policy Clients elements

Table 5-15 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Policy Clients** screen. Data Input Notes apply to the Insert and Edit screens; the View screen is read-only.



Table 5-15 Policy Clients Elements

Elements (* indicates			
required field)	Description	Data Input Notes	
* Policy Client Peer Node	The name of a configured	Format: Pulldown list	
Name	Diameter Peer Node that identifies a Policy Client Peer Node.	Note: The Policy Client Peer Node Name cannot be changed on the Policy Clients	
	Selecting a Policy Client Peer	[Edit] screen.	
	Node name (blue hyperlink) displays the Diameter , and then Configuration , and ther Peer Nodes (Filtered) screer where Diameter Peer Nodes are filtered by the Policy Clien Peer Node Name.		
Topology Hiding Enabled	A read-only check box with	Format: Check box	
	default checked to indicate the Topology Hiding for the policy client peer node being enabled. It is the only option currently supported.	Range: N/A (Read Only)	
Comments	An optional comment that	Format: Text box	
	describes the Policy Client Peer Node.	Range 0-64 characters	

Adding a New Policy Client for Topology Hiding

Use this task to add a new Policy Client for Topology Hiding.



Topology Hiding is performed only if it is Enabled and the Topology Hiding Scope option is defined as Specific Clients or All Foreign Realms + Specific Clients in the Policy and Charging, and then Configuration, and then Policy DRA, and then Network-Wide Options screen on the NOAM.

The fields are described in Policy Clients elements.

- On the Active SOAM, click Policy and Charging, and then Configuration, and then Policy DRA, and then Policy Clients.
- 2. Click Insert.
- Select a Policy Client Peer Node Name from the Value dropdown list.
- Check Topology Hiding Enabled if Topology hiding is needed for the Policy Client.
- 5. Enter an optional comment in the **Comments** field.
- 6. Click:
 - OK to save the changes and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then Policy Clients screen.



- Apply to save the changes and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then Policy Clients screen without saving any changes.

If **OK** or **Apply** is clicked and any of several conditions exist, an error message appears:

- The entered comment exceeds 64 characters in length or contains something other than 7-bit ASCII characters.
- The Policy Client Peer Node Name is missing.
- The selected **Policy Peer Node Name** is already configured in the system.
- Any fields contain invalid input (for example, the wrong type of data was entered or a value is out of range).
- The maximum number (1000) of **Topology Hiding** records has already been configured.

Editing Policy Clients for Topology Hiding

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



Topology Hiding is performed only if it is Enabled and the Topology Hiding Scope option is defined as Specific Clients or All Foreign Realms + Specific Clients in the Policy and Charging, and then Configuration, and then Policy DRA, and then Network-Wide Options screen on the NOAM.

The fields are described in Policy Clients elements.

- 1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Policy Clients**.
- 2. Click Edit.

A read-only value is displayed in the Policy Client Peer Node Name Value field.

- Check Topology Hiding Enabled if Topology hiding is needed for the Policy Client.
- 4. Edit or enter an optional comment in the **Comments** field.
- 5. Click:
 - OK to save the edited Comment and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then Policy Clients screen.
 - Apply to save the edited Comment and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then Policy Clients screen without saving any changes.

If **OK** or **Apply** is clicked and 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, then an error message appears.



Deleting a Topology Hiding Policy Client Peer Node

Use this procedure to delete a Topology Hiding Policy Client Peer Node.

- 1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Policy Clients**.
- 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**, and then **Configuration**, and then **Peer Nodes** screen for the Peer Node.)
- 3. Click Delete.
- 4. Click:
 - OK to delete the Policy Client Peer Node Name.
 - Click Cancel to cancel the delete function and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then Topology Hiding screen.

If **OK** is clicked and the selected Policy Client Peer Node no longer exists (it was deleted by another user), an error message is displayed and the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Policy Clients** screen 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 Site Options can be configured on **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Site Options** screen on the active SOAM server:

- Topology Hiding Virtual Name FQDN and Realm. See Site Options elements
- Peer Route Table Name The name of the Peer Route Table to be used for routing new binding requests. This entry is no longer used when PCRF Pooling is enabled.
- Enable Reroute Indicates whether the Reroute is enabled or disabled.
- Reroute Peer Route Table Name Indicates the name of the Peer Route Table to be used for rerouting binding requests.

The fields are described in Site Options elements.

Site Options elements

Table 5-16 describes the elements on the SOAM Policy and Charging, and then Configuration, and then Policy DRA, and then Site Options screen. Data Input Notes apply to the Insert and Edit screens; the View screen is read-only.



Table 5-16 Site Options Elements

Field	Descriptions	Data Input Notes
Topology Hiding Virtual Name	FQDN	Format: Text box
	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.	Range: 1 - 255 characters. Valid characters are letters, digits, dots (.), and hyphens (-). At least one alpha character is required.
	Realm	Format: Text box
	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.	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	Format: Pulldown list
	Peer Route Table to be used for routing new binding requests.	Range: Not Selected, Default, configured Diameter Peer Route tables
	The Default PRT is always available, but must be selected from the list to be used.	Default: Not Selected
Enable Reroute	Indicates whether the Reroute is enabled or disabled.	Range: Checked or Unchecked
	Select this check box to configure Reroute and check whether subscriber binding is available or not.	Default: Reroute is disabled (unchecked)
Reroute Peer Route Table Name	Indicates the name of the Peer Route Table to be used for rerouting binding requests. This field is selected when Reroute is enabled.	Range: List of configured Diameter Peer Route Tables Default: Not Selected

Setting Site Options

Use this task to set Site Options on the Active SOAM server.

The fields are described in Site Options elements.

- 1. Click Policy and Charging, and then Configuration, and then Policy DRA, and then Site Options.
- 2. Enter an FDQN and Realm.



Note:

If no values are configured here when Topology Hiding is enabled, the FQDN and Realm values of the Default Topology Hiding Virtual Name configured in NOAM GUI Main Menu: Policy and Charging, and then Configuration, and then Policy DRA, and then Network-Wide Options will be used.

3. Select a Peer Route Table Name from the drop-down list.

This entry is no longer used once PCRF Pooling is Enabled.

- 4. Select the **Enable Reroute** check box to enable Reroute.
- Select a peer route table name from the Reroute Peer Route Table Name dropdown list.
- 6. Click:
 - Apply to save the changes and refresh this screen.
 - Cancel to discard the changes and remain on the Policy and ChargingConfigurationPolicy DRA screen.

If **Apply** is clicked and any entered value contains the wrong data type or is out of the allowed range, an error message appears.

SBR Databases

Refer to the SBR User's Guide for how to configure SBR Databases.

PCRF Pools

Policy DRA continues to support a single pool of PCRFs at each PCA site over which policy Diameter signaling is distributed using the subscriber's IMSI. This allows the incorporation of new services or new PCRF infrastructure without disturbing existing services. For example, one set of PCRF servers handle policy control for all consumer data accesses to their network and a second set of PCRF servers handle all enterprise data accesses for their network. The policy rules and/or PCRF implementations might be different enough to necessitate that these two services are segregated at the PCRF level.

This means that a given IMSI might concurrently have a binding to one PCRF for APN A and a binding to a different PCRF for APN B. Each APN is mapped to a set of PCRFs; this is called a PCRF Pool. In addition, if a binding to a PCRF Pool already and a new session is created that maps to that same PCRF Pool, the request must be routed to the same PCRF. When new bindings are created for different IMSIs and a given APN, the binding-capable session initiation requests are distributed across the PCRFs in the PCRF Pool assigned to that APN.

PCRF Pooling expands this capability for the creation of multiple pools of PCRFs, which are selected using the combination of IMSI and Access Point Name (APN). This allows you to route policy Diameter signaling initiating from a given APN to a designated subset of the PCRFs that can provide specialized policy treatment using knowledge of the APN.

PCRF Pooling modifies the logic in the Policy DRA to inspect the contents of binding generating Gx CCR-I messages to select the type of PCRF to which the CCR-I



messages are to be routed. In the initial P-DRA, it was assumed that all PCRFs could handle all Gx session bindings. PCRF Pooling provides service-specific sets of PCRFs. In this release, the APN used by the UE to connect to the network is used to determine the PCRF pool. The Origin-Host of the PCEF sending the CCR-I can then be used to select a PCRF sub-pool.

Multiple PCRF pools requires differentiation among the binding records in the binding SBR. It is possible for the same UE, as indicated by the IMSI, to have multiple active IPcan sessions spread across the different pools.

✓ Note:

Although the concept of a PCRF pool is a network-wide concept for a service provider, PCRF pools configuration is done on a PCA site-by-site basis. PCAs in different sites can support different PCRF Pool Selection configurations.

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

On the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pools** screen on the NOAM or SOAM, you can perform a variety of 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 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, a certain 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.

There are major differences between PCRF Pooling and non-pooling functionality:



Table 5-17 PCRF Pooling Concepts

Concept	Before PCRF Pooling	After PCRF Pooling
PCRF Pools	One PCRF Pool for all APNs.	Up to 7 PCRF Pools selected for new bindings using APN. More than one APN can be mapped to a given PCRF Pool, but a given APN can only be mapped to one PCRF Pool.
Subscriber Bindings	A binding is a simple mapping between an IMSI and a PCRF. Once a binding exists, all sessions for that IMSI are routed to the bound PCRF.	A binding is a mapping from an IMSI and APN to a PCRF, but with the caveat that before a new binding is created, the logic must check for existence of another binding to the same PCRF Pool for the IMSI. If such a binding exists, the new APN is bound to the same PCRF as an existing APN mapped to the same PCRF Pool. Once a binding exists, all sessions for that IMSI and APN are routed to the bound PCRF. Sessions for that IMSI and a different APN mapped to a different PCRF Pool can be routed to a different PCRF.
Number of Sessions per Binding	An IMSI may have up to 10 binding capable sessions.	An IMSI may have up to 10 binding capable sessions, which may be bound to different PCRFs based on APN.
Origin Based Routing	PRT table for new bindings specified in Site Options allows for selection of route list based on origin-host/realm.	After PCRF Pool selection, Sub-Pool rule matching is performed to select a PCRF Sub-Pool given the PCRF Pool and the origin-host of the PCEF.
PRT Table for New Bindings	Each site defines one PRT table to be used for all new bindings.	Each site can define a PRT table to be used for new bindings for each PCRF Pool.

Additionally, Pooling provides the ability to route to subsets of PCRFs in a PCRF Pool on the basis of the Diameter hostname of the PCEF that originated the binding capable session initiation request. These subsets are called PCRF Sub-Pools. This capability allows a controlled amount of policy Diameter signaling to be routed to one or more PCRFs within the PCRF Pool.

Figure 5-2 illustrates a sample PCA network configured for PCRF Pooling. The upper third of the figure shows data that is configured with the Policy and Charging GUI at the NOAM server. This data, including PCRF Pools, APN to PCRF Pool mapping, and PCRF Sub-Pool Selection Rules applies to all sites in the Policy DRA network.

The middle third of the figure shows data configured at the SOAM Policy and Charging GUI at each of two PCA sites. This data includes the PCRF Pool to PRT mappings,



PCRFs, PRT tables, Route Lists, Route Groups, Peer Nodes, and Connections. This data can differ at each PCA site.

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

Figure 5-2 PCRF Pooling Data

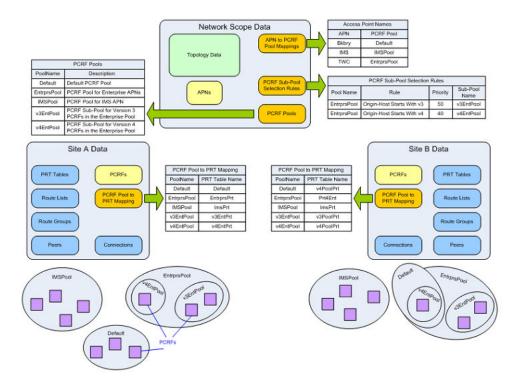


Table 5-18 describes each of the new PCRF Pooling configuration tables, including the order in which they should be configured.

Table 5-18 PCRF Pooling Configuration Summary

Configuration Order	GUI Screen	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.



Table 5-18 (Cont.) PCRF Pooling Configuration Summary

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

PCRF Pools elements

Table 5-19 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pools** screen.

Note:

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

The PCRF Pools table contains the list of configured PCRF Pools and Sub-Pools settings that you can use when selecting a set of PCRFs to host a new subscriber binding. The PCRF Pool to be used for a given subscriber binding attempt is determined based on the APN-to-PCRF Pool mappings configured in **Policy and Charging**, and then **Configuration**, and then **Access Point Names** and the PCRF Sub-Pool Selection Rules configured in **Policy and Charging**, and then **Configuration**, and then **PCRF Sub-Pool Selection Rules**.

Table 5-19 PCRF Pools Elements

Fields (* indicates required field)	Description	Data Input Notes
* PCRF Pool Name	A name that uniquely identifies the PCRF Pool A PCRF Pool names a set of PCRFs that should be used for policy request from a specified APN. THe mapping from APN to PCRF Pool is configured on the Policy and Charging, and then Configuration, and then Access Point Names screen.	Default: N/A Range: 1 to 32 characters,
Sub-Pool	Check this 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 on the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Sub-Pool Selection Rules screen. Note: If the check box on the	Format: Check box Default: No (Unchecked for Sub-Pool) Range: Yes (Checked for Sub-Pool), No (Unchecked for Sub-Pool)
Comments	PCRF Pools [Insert] screen is not checked, this PCRF Pool is a pool, not a sub-pool. An optional comment to	Format: Text box
	provide more information about the purpose of this PCRF Pool or Sub-Pool.	Default: N/A Range: 0-64 characters

Inserting PCRF Pools

Use this task to insert (create new) PCRF Pools.



- 1. On the Active NOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pools**.
- 2. Click Insert.
- 3. Enter a unique PCRF Pool Name in the PCRF Pool Name field.
- 4. Check the **Sub-Pool** check box if the PCRF Pool is to be used as a Sub-Pool.

A Sub-Pool is used if policy requests from specified origin-hosts should be routed to a different set of the PCRFs from those in the PCRF Pool selected by the APN. Sub-Pool Selection Rules are configured in **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **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).

- You can type an optional comment in the Comments field to describe the Pool or Sub-Pool. The entry must be characters in the range of 0 to 64, and the default is N/A.
- 6. Click:
 - OK to save the new PCRF Pool name and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools screen.
 - Apply to save the new PCRF Pool name and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools screen without saving any changes.

If **OK** or **Apply** is clicked and any of several conditions exist, an error message appears:

- The entered PCRF Pool name is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- Adding the new PCRF Pool would cause the maximum number of PCRF Pools (7) to be exceeded.

Editing PCRF Pools

Use this task to edit PCRF Pools comments. After a PCRF Pool is created, only the comment can be edited, and the Sub-Pool Indicator can only be changed by deleting the PCRF Pool and creating a new one.



The PCRF Pool Name cannot be edited.

1. On the Active NOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pools**.

The screen 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.
- Click in the Comments field.
- 5. Edit the **Comments** field for the selected PCRF Pool. The comment must be characters in the range of 0 to 64, and the default is N/A.
- 6. Click:
 - OK to save the change and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools screen.
 - **Apply** to save the change and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools screen 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 screens at any of the sites.

- 1. On the Active NOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pools**.
- 2. Select the PCRF Pool Name or PCRF Sub-Pool Name to be deleted.
- 3. Click Delete.
- 4. Click:
 - OK to delete the PCRF Pool or PCRF Sub-Pool.
 - Cancel to cancel the delete function and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools screen.

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 screen is refreshed. The row that was selected is no longer displayed in the list.

PCRF Sub-Pool Selection Rules

The PCRF Sub-Pool Selection table contains rules for selection of a PCRF Sub-Pool for a given PCRF Pool and Origin-Host value.



It is sometimes necessary to subdivide a PCRF Pool into sub-pools (such as a need 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 certain attributes:

- The Default PCRF Pool can have sub-pools.
- The PCRF Pool Name column contains hyperlinks to the PCRF Pools screen 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 a pair of conditions exist:

- An APN is mapped to the PCRF Pool using the Access Point Names screen.
- 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:

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 certain 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 several conditions 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 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 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 5-20 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Sub-Pool Selection Rules** screen.



Table 5-20 PCRF Sub-Pool Selection Rules Elements

Fields (* indicates		
required field)	Description	Data Input Notes
* PCRF Sub-Pool Selection Rule Name	A name that uniquely identifies the PCRF Sub-Pool Selection Rule.	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 Default: N/A Range: Valid name
* Priority	A priority value.	Format: Text box
•	A low value indicates a higher	Default: 50
	priority.	Range: 1-99, inclusive
	Note: 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.	
* PCRF Pool Name	The name of the PCRF Pool for which a Sub-Pool is being defined.	Format: Dropdown menu Default: N/A
	A dropdown 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.	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	Format: Dropdown menu and Text box
	a given Diameter AVP using	Default: N/A
	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.	Range: Equals, Starts With, and Ends With.
* PCRF Sub-Pool	The PCRF Sub-Pool that is	Format: Dropdown menu
Name	to be used for Gx and Gxx	Default: N/A
	session initiation request messages matching this Rule.	Range: Assigned PCRF Sub-Pool Name



Table 5-20 (Cont.) PCRF Sub-Pool Selection Rules Elements

Fields (* indicates required field)	Description	Data Input Notes
Last Updated	This read-only field displays the date and time that this rule was created, or the last time the rule was changed, whichever is most recent.	Format: Text box (read-only)
	This field records the time and date of changes that may affect routing of binding capable session initiation requests. This date and time can be compared against binding creation times when troubleshooting using Policy and Charging, and then Maintenance, and then Policy Database Query.	

Inserting PCRF Sub-Pool Selection Rules

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

- 1. On the Active NOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Sub-Pool Selection Rules**.
- 2. Click Insert.
- 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 dropdown 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.

Select a condition from the Operator dropdown 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 **PRCF Sub-Pool Name** from the dropdown menu. Choices include all the qualified PCRF Sub-Pool configured from the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pools** screen. A qualified PCRF Sub-Pool is a PCRF Pool that is non-retired and has been marked as Sub-Pool. A retired PCRF Sub-Pool entry can be created by first adding a new PCRF Sub-Pool and then deleting it.

This is the PCRF Sub-Pool that is to be used for Gx and Gxx session initiation request messages that match this Rule. The default is N/A and the range is the choice of configured PCRF Pools.

9. The Last Updated field is a read-only field that displays the date and time that this rule was created, or the last time the rule was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding-capable session initiation requests. This date and time can be compared against binding creation times when troubleshooting using the Binding Key Query Tool.

10. Click:

- OK to save the new PCRF Pool name and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools screen.
- Apply to save the new PCRF Sub-Pool Selection Rule and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Sub-Pool Selection Rule screen 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 screen 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, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **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.

2. Select a PCRF Sub-Pool Selection Rule to edit.

DO NOT click the blue PCRF Pool Name or the PCRF Sub-Pool Name(except to see the configuration of the PCRF Pool Name or PCRF Sub-Pool Name). The blue color indicates a hyper-link that opens the **Diameter**, and then **Configuration**, and then **Peer Nodes [Filtered]** screen to display the configuration information for the Peer Node.

Click Edit.

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** from the dropdown menu.

This PCRF Sub-Pool that will be used for Gx and Gxx session initiation request messages matching this Rule The default is N/A, and the range is the choice of configured PCRF Sub-Pool Selection Rules.

8. Last Updated is a read-only field that displays the date and time that this rule was created, or the last time the rule was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding capable session initiation requests. This date and time can be compared against binding creation times when troubleshooting using the Binding Key Query Tool.

9. Click:

- OK to save the change and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Sub-Pool Selection Rules screen.
- **Apply** to save the change and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Sub-Pool Selection Rules screen 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 this procedure to delete a PCRF.

A PCRF Sub-Pool Selection Rule can be deleted at any time.

- 1. Click Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Sub-Pool Selection Rules.
- 2. Select the PCRF Sub-Pool Selection Rule Name to be deleted.
- 3. Click Delete.
- 4. Click:
 - OK to delete the PCRF Sub-Pool Selection Rule Name.



 Cancel to cancel the delete function and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Sub-Pool Selection Rules screen.

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 screen is refreshed. The row that was selected is no longer displayed in the list.

PCRF Pool to PRT Mapping

In initial DSR release installations, PCRF Pools and PRT tables must be configured as part of configuring the PCA application. For initial installs and upgrades from releases in which PCA was not activated, the Default PCRF Pool is created and mapped to the Not Selected PRT.

When a PCRF Pool or Sub-Pool is added at the NOAMP, the data is replicated on the SOAM servers at each site. When a user opens the PCRF Pool to PRT Mapping screen, 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.



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 screen 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 screen, and it will have the same PRT choice as was previously configured, provided the PRT table still exists.

A Peer Route Table cannot be deleted from a site if that Peer Route Table is referenced by a current PCRF Pool to PRT Mapping entry. Entries for retired PCRF Pools or Sub-Pools are not included in this restriction. As a result, if a PCRF Pool A had a mapping to PRT table X, then PCRF Pool A was deleted at the NOAMP, it is possible to delete PRT X (provided no other active PCRF Pool to PRT Mappings referenced PRT X). If PCRF Pool A was added back at the NOAM after the deletion of PRT X, PCRF Pool A would appear on the PCRF Pool to PRT Mapping GUI with its PRT entry set to the default of Not Selected.

If a PCRF Pool or Sub-Pool is changed from being mapped to a PRT table to the -Select- value in the PRT pulldown menu, you might see a confirmation window that includes a warning if a certain condition applies:



- If an APN is mapped to the PCRF Pool being changed, a confirmation window is
 displayed on the PCRF Pool to PRT Mapping screen 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 screen 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 5-21 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pools to PRT Mapping** screen.



Data Input Notes apply to the Insert and Edit screens; the View screen 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 used when a new binding is created for the PCRF Pool. The PCRF Pool or Sub-Pool to be used for a given subscriber binding attempt is determined based on Access point Name to PCRF Pool mappings, or by rules configured at the NOAMP in **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **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 5-21 PCRF Pools to PRT Mapping Elements

Field (* indicates required		
field)	Description	Data Input Notes
* PCRF Pool Name	The name of the PCRF Pool or Sub-Pool defined for the network on the PCRF Pools screen.	Format: Text box; string of 1-32 alphanumeric characters, must contain at least one alpha character, must not start
	When a PCRF Pool or PCRF Sub-Pool is configured at	with a digit, and can contain underscores
	the NOAMP, it automatically appears on the PCRF Pool to PRT Mappings screen so that a PRT can be defined for it if needed. This field is a hyper-link to the PCRF Pools (Filtered) view screen, filtered by the PCRF Pool or Sub-Pool name.	Range: Valid name
Peer Route Table Name	The name of a configured Peer Route table used to route new binding requests destined to the PCRF Pool or PCRF Sub-Pool.	Format: String Range: All Peer Route Tables configured at this site Default: Not Selected
	This field is a hyperlink to the Diameter , and then Configuration , and then Peer Route Tables view screen, filtered by the PRT name.	

Editing PCRF Pool to PRT Mapping

Use this task to edit PCRF Pool to PRT Mapping settings.

1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pool to PRT Mapping**.

The screen displays a list of PCRF Pools or Sub-Pools configured at the NOAMP.

2. Select a row to edit (click in the row, but do not click on a specific element within the row).

DO NOT click the blue PCRF Pool Name or the Peer Route Table Name (except to view the PCRF Pools (Filtered) screen or the Peer Routes Table (Filtered) screen. The blue color indicates a hyper-link. The PCRF Pool Name hyper-link opens the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pools (Filtered)** screen and the Peer Route Table hyper-link opens the **Diameter**, and then **Configuration**, and then **Peer Routes Table (Filtered)** screen.

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.

(optional) Click Pause updates to suppress the automatic screen 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 screen the next time an update occurs.

4. Click Edit.

The Peer Route Table Name dropdown menu initially displays the Peer Route Table from the row being edited and contains all configured Peer Route Tables and Not Selected. Not Selected provides backwards compatibility for users who had the Site Options Peer Route Table Name set to Not Selected. When Not Selected is chosen, PCA does not instruct DRL to use an application specified PRT, but enables DRL use its normal PRT precedence for PRT selection instead. If Edit is clicked and the PCRF Pool Name of the selected row has been deleted, an error is displayed and this row is no longer displayed. If Edit is clicked and the PCRF Pool Name of the selected row still exists (has not been retired), the PCRF Pool To PRT Mapping [Edit] screen is displayed with data populated from the selected row.

5. Select an item from the **Peer Route Table Name** dropdown menu. The default is Not Selected, and the range is All Peer Route Tables configured at this site.

6. Click:

- OK to save the selection and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pool to PRT Mapping screen.
- Apply to save the selection and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pool to PRT Mapping screen without saving any changes.

Additionally, multiples issues 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] screen 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 screen is automatically refreshed every *N* seconds to show the latest PCRF Pools configured at the NOAMP **Policy and Charging**, and then **Configuration**, and then **POICY DRA**, and then **PCRF Pools** screen.

Pausing update applies to all rows in the table on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Pool to PRT Mapping** screen. Selecting this check box pause the automatic update function for all items in the table.

- On the Active SOAM, click Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pool to PRT Mapping.
 - The screen displays a list of the configured PCRF Pool Names and corresponding Peer Route Table Names.
- (optional) Click Pause updates to suppress the automatic screen 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 screen the next time an update occurs.

Error Codes

For each Policy and Charging Site, the Diameter Error Result Code value to send to the Request initiator for policy related errors can be configured according to which condition applies. Each condition can be mapped to a different Result Code for each supported interface. Result Codes can be Diameter IANA defined or experimental.

When PCRF Pooling is enabled, new binding cannot be created unless the binding-capable session initiation request contains a configured APN. If the binding-capable session initiation request arrives with either no APN, or an APN that is not configured in the Access Point Names table, the request is answered by Policy DRA using a configurable error response code. To configure the Diameter response code for this scenario, a new Missing Or Unconfigured APN error condition has been added to the existing SOAM error. This error response applies to all binding capable interfaces (for example, Gx, Gxx, and S9) and can be configured with either an IANA Diameter result code, or an experimental result code and vendor-id.

Three-digit error codes in Diameter Error-Message AVPs indicate exactly why a slave session could not be routed. This provides more robust troubleshooting using Diameter capture tools.

A 3-digit error code is an identifier to uniquely identify a specific error scenario (not error category) encountered in a Diameter Answer message generated by PCA. 3-digit codes are unique across all DSR layers (DSR connection layer, routing layer and application layer) and all DSR applications (such as PCA, RBAR, FABR, and IDIH) for errors they represent. The ranges of 500-549 and 850-899 are for the PCA application, while the DSR connection layer, routing layer and other DSR applications uses other non-overlapping ranges. Multiple errors may belong to a same error category and are associated with a same Result-Code. It is the 3-digit code that can distinguish an error from others. Users should search for the 3-digit code when identifying an error if possible and available.



Note:

The error conditions in this table are GUI-configurable.

Table 5-22 PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Missing or Unconfigured APN	PDRA	Policy-related binding capable session initiation request messages	3002	500	Missing or Unconfigured APN. Binding capable session initiation request is received with no APN
Missing or Unconfigured APN	PDRA	Policy-related binding capable session initiation request messages	3002	501	Missing or Unconfigured APN. Binding capable session initiation request is received with APN, but the APN is not configured in the APN configuration.
Unable To Route	PDRA	Policy-related binding capable and dependent session initiation request messages	3002	502	Unable To Route. Request message is received and a binding with a PCRF was found. P-DRA cannot route the request to PCRF due to DSR queue full error.



Table 5-22 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Unable To Route	PDRA	Policy-related binding capable and dependent session initiation request messages	3002	3-digit error code from DRL	Unable To Route. Request message is received and a binding with a PCRF was found. P-DRA cannot route the request to PCRF due to PCRF being unreachable. DRL Text string.
No Usable Keys In Binding Dependent Message	PDRA	Policy-related binding dependent session initiation request messages	3002	503	No Usable Keys In Binding Dependent Message. No binding key in Binding Key Priority GUI can be matched, or no key is included in the binding dependent message.
Binding Not Found	PDRA	Policy-related binding dependent session initiation request messages	3002	505	Binding Not Found. Binding record is not found after examining all configured binding keys in Diameter message.
SBR Error	PDRA	Policy-related binding capable and dependent session initiation, update or terminate answer messages,	3002	507	SBR Error. ComAgent timeout.



Table 5-22 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
SBR Error	PDRA	Policy-related binding capable and dependent session initiation, update or terminate answer messages	3002	508	SBR Error. SBR database error prevents SBR from reading, writing or deleting a record,
Session Not Found	PDRA	Policy-related binding capable and dependent session update or terminate request messages	3002	509	Session Not Found. Session record does not exist for given session ID
PCA Unavailable Or Degraded	PDRA/ OCDRA	Any Diameter Requests forwarded to PCA	3002	305	PCA Unavailable or Degraded
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	520	SBR Error. Binding capable session initiation request received, but no PCRFs are configured at the site, or PCRF ID is not found in PCRF table.
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	521	SBR Error. Maximum number of sessions per binding is exceeded that fails the binding creation for given subscriber's key.



Table 5-22 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Session ID is missing from Request	PDRA	Any Policy- related Diameter Requests forwarded to P-DRA	5005	522	Session ID is missing from Request
CC - Request - Type AVP is missing from CCR message	PDRA	Policy-related binding capable session initiation, update or terminate request messages	5005	523	CC-Request- Type AVP is missing from CCR message
Not In Use		J			
Invalid AVP value in request message	PDRA	Any Policy- related Diameter Requests forwarded to P-DRA	5004	525	Invalid AVP value in request message
Destination - Host AVP is missing in in- session request	PDRA	Policy-related binding capable update and terminate request and dependent session initiation update or terminate request messages	5012	506	Destination- Host AVP is missing in in- session request
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	510	Unable To Route. A slave session could not be routed because on polling the slave sessionRef was no longer in the binding database.



Table 5-22 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	511	Unable To Route. A slave session could not be routed because on polling the master sessionRef was no longer in the binding database.
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	512	Unable To Route. A slave session could not be routed because on polling the master sessionRef was early too long.
SBR Error	PDRA	Policy-related Requests and Answers	3002	504	SBR Error. ComAgent unavailable when sending stack event to SBR
Unsupported Application ID	PDRA/ OCDRA	Diameter Requests	3007	530	Application ID unsupported by PCA
Command Code and App ID no match	PDRA	Policy-related Requests and Answers	5019	531	Command Code does not match App ID or not exist
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	513	Unable To Route. A slave session could not routed because on polling the master session and internal error occurred.



Table 5-22 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
PCA Functionality Unavailable or Disabled	PDRA	Policy related binding capable and dependent session update or terminate request messages	3002	532	PCA Functionality Unavailable or Disabled. Policy DRA Function Disabled.
PCA Functionality Unavailable or Disabled	PDRA	Policy related binding capable and dependent session update or terminate request messages	3002	533	PCA Functionality Unavailable or Disabled. Policy DRA Function Unavailable.
PCA Functionality Unavailable or Disabled	OCDRA	Online Charging related binding independent session request messages	3002	534	PCA Functionality Unavailable or Disabled. Online Charging DRA Function Disabled.
PCA Functionality Unavailable or Disabled	OCDRA	Online Charging related binding independent session request messages	3002	535	PCA Functionality Unavailable or Disabled. Online Charging DRA Function Unavailable
Session ID is missing from Request	OCDRA	Any Online Charging - related Diameter Requests forwarded to OC-DRA	5005	536	Session ID is missing from Request
CC - Request - Type AVP is missing from CCR message	OCDRA	Any Online Charging- related Diameter Requests forwarded to OC-DRA	5005	537	CC-Request- Type AVP is missing from CCR message



Table 5-22 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Invalid AVP value in request message	OCDRA	Any Online Charging- related Diameter Requests forwarded to OC-DRA	5004	538	Invalid AVP value in request message
Not In Use					
Unable To Route	OCDRA	Online Charging- related binding independent session request messages	3002	540	Unable To Route. Request message is received, OC- DRA cannot route the request to OCS due to DSR queue full error.
Unable To Route	OCDRA	Online Charging- related binding independent session initiation request messages	3002	539	Unable To Route. Request message can not be routed to peer node. DIAM-ERR- XXXX- XXX:DRL Text string.
SBR Error	OCDRA	Online Charging- related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	541	SBR Error. ComAgent timeout.



Table 5-22 (Cont.) PCA Error Conditions

		Applied			
Error Category	Functionality	Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
SBR Error	OCDRA	Online Charging- related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	542	SBR Error. SBR database error prevents SBR from reading, writing or deleting a record,
SBR Error	OCDRA	Online Charging- related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	543	SBR Error . ComAgent unavailable when sending stack event to SBR,
Session Not Found	OCDRA	Online Charging- related session update or terminate request messages, if session state or topology hiding applies	5002	544	Session Not Found. Session record does not exist for given session ID
Command Code and App ID no match	OCDRA	Online Charging- related Requests.	3001	545	Command Code does not match App ID or not exist



Table 5-22 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	546	SBR Error. A binding capable session initiation request could not be routed, maximum sessions per IMSI per APN limit is exceeded.
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	547	SBR Error. A binding capable session initiation request could not be routed, maximum sessions per IMSI per APN limit is exceeded and existing sessions could not be replaced because binding is in early state.



Table 5-22 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	548	SBR Error. A binding capable session initiation request could not be routed, maximum sessions per IMSI per APN limit is exceeded and existing sessions could not be replaced because Maximum Early Binding lifetime is not elapsed for existing sessions.

On the **Policy and Charging**, and then **Configuration**, and then **Error Codes** screen on the SOAM:

Select an Error Condition in the list and click Edit.
 You can edit the selected Error Code. See Editing Error Codes.

The fields are described in Error Codes elements.

Error Codes elements

Table 5-24 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Error Codes** screens. Data Input Notes apply to the Error Codes [Edit] screen; the View screen is read-only.

The Error Codes define the Result Codes to be returned for various Policy and Charging Error Conditions. Each Error Condition will return the Result Code configured for each applicable Diameter interface.

Table 5-23 indicates the Diameter interfaces that are supported for each Error Code.

The default Result Code is 3002-DIAMETER_UNABLE_TO_DELIVER.



Table 5-23 Interfaces Supported for Each Error Code

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

Table 5-24 Error Codes Elements

Fields (* indicates required		
field)	Description	Data Input Notes
* Error Condition	The name of the selected Policy and Charging Error Condition.	View only; cannot be edited
* Gx/Gxx, Result Code	The Result Code to be returned on the Gx and Gxx interfaces	Format: Text box Range: 1-9999 Default: 3002
Gx/Gxx Vendor ID	The Vendor ID that corresponds with the Gx and Gxx interfaces.	Format: Text box Range: 1-4294967295
	The Vendor ID means the RFC standard error code will be sent.	
* Rx Result Code	The Result Code to be returned to the Rx interface.	Format: Text box Range: 1-9999 Default: 3002
Rx Vendor ID	The Vendor ID that corresponds with the Rx interface.	Format: Text box Range: 1-4294967295
	The Vendor ID means the RFC standard error code will be sent.	
* S9 Result Code	The Result Code to be returned to the S9 interface.	Format: Text box Range: 1-9999 Default: 3002



Table 5-24 (Cont.) Error Codes Elements

Fields (* indicates required field)	Description	Data Input Notes
S9 Vendor ID	The Vendor ID that corresponds the S9 interface. The Vendor ID means the RFC standard error code will be sent.	Format: Text box Range: 1-4294967295
* Gx-Prime Result Code	The Result Code to be returned on the Gx-Prime interface	Format: Text Box Range: 1-9999 Default: 3002
Gx-Prime Vendor ID	The Vendor ID that corresponds with the Gx-Prime interface. The Vendor ID means the RFC standard error code will be sent.	Format: Text Box Range: 1-4294967295
* Gy/Ro Result Code	The Result code to be returned on the Gy/Ro interface.	Format: Text Box Range: 1-9999
Gy/Ro Vendor ID	The Vendor ID that corresponds with the experimental code for the Gy/Ro interface.	Format: Text Box Range: 1-4294967295

Editing Error Codes

Use this task to edit Error Codes on the Active SOAM.

The fields are described in Error Codes elements.

- 1. Click **Policy and Charging**, and then **Configuration**, and then **Error Codes**.
- 2. Select the **Error Condition** that you want to edit.
- 3. Click Edit.

The fields that appear on the **Policy and Charging**, and then **Configuration**, and then **Error Codes [Edit]** screen are dependent on the Error Condition that was selected.

- 4. Edit the fields to define the selected Error Condition.
- 5. Click:
 - OK to save the changes and return to the Policy and Charging, and then Configuration, and then Error Codes screen
 - Apply to save the changes and remain on this screen
 - Cancel to discard the changes and return to the Policy and Charging, and then Configuration, and then Error Codes screen

If **OK** or **Apply** is clicked and any of several 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).

Suspect Binding Removal Rules

On the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Suspect Binding Removal Rules** screen on the SO GUI, Suspect Binding Removal Rules are listed. During Diameter messaging, when a Suspect Binding Removal Rule is matched, the rule configuration determines if the Subscriber's bindings should be considered suspect or be removed immediately.

When configuring Suspect Binding Removal Rules:

- The (Experimental) Result Code attribute is applicable only if the Error Scenario Category is set to External Result.
- Vendor ID is applicable only if the Result Code is an experimental result code.
- If Remove Suspect Binding Immediately is set to Yes, a single match on this rule will cause the removal of all sessions for the IMSI and PCRF.
- If Remove Suspect Binding Immediately is set to No, a rule match causes all sessions for the IMSI and PCRF to be considered suspect and the number of rule matches is incremented.



If the number of rule matches reaches or exceeds the configured in the Suspect Binding Removal Events Threshold field on the **Policy and Charging**, and then **Configuration**, and then **Policy** screen, all sessions for the IMSI and PCRF will be removed.

The fields are described in Suspect Binding Removal Rules elements.

Suspect Binding Removal Rules elements

Table 5-25 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Suspect Binding Removal Rules** screen on the SO GUI.

Table 5-25 Suspect Binding Removal Rules Elements

Field (* indicates a required field)	Description	Data Input Notes
* Rule Name	A name that uniquely identifies	Format: Text box
	the Suspect Binding Removal	Default: N/A
	Rule	Range: A 32-character string. Valid characters are alphanumeric and underscore. Must contain at least one alphabetic character and must not start with a digit.



Table 5-25 (Cont.) Suspect Binding Removal Rules Elements

Field (* indicates a required	Description	Date leavet Nation
field)	Description	Data Input Notes
* Application Name	The Diameter Application name and Id to which this Suspect Binding Removal Rule applies. Session initiation answer messages including the Application-Id are candidates to match this rule.	Format: Pull-down list
		Default: N/A
		Range: Supported P-DRA Application-Ids
* Command Code	The Diameter Command Code or Extended Command Code name and value to which	Format: Pull-down list
		Default: N/A
	this Suspect Binding Removal Rule applies. Session initiation answers messages including this Command Code are candidates to match this rule.	Range: Supported P-DRA session initiation answer messages
* Error Scenario Category	The error category to which	Format: Pull-down list
	the Suspect Binding Removal	Default: N/A
	Rule applies. Category Unable to Route is for when no session initiation answer is received from the PCRF (possibly because the request could not be routed). If Unable to Route is chosen, the (Experimental) Result Code sent to the policy clients is the one configured on the Policy and Charging, and then Configuration, and then Error Codes screen for the specific interface. Category External Result is for when a specified session imitation error answer is received from the PCRF. If External Result is chosen, a Result Code must be specified, otherwise no Result Code is necessary.	Range: External Result, Unable to Route
Result Code	The session initiation error answer (Experimental) Result Code to which this Suspect Binding Removal rule applies if the Error Scenario Category is External Result. This field is not applicable when Error Scenario Category is set to Unable to Route.	Format: Text box Default: N/A Range: 1-9999



Table 5-25 (Cont.) Suspect Binding Removal Rules Elements

Field (* indicates a required field)	Description	Data Input Notes
Vendor ID	If a Result Code is entered in the Result Code field and that Result Code is an experimental result code, enter a Vendor-Id in this field. Otherwise, leave the field blank.	Format: Text box Default: N/A Range: 1-4294967295
Remove Suspect Binding Immediately	Check this box if a single occurrence of this rule match means that the binding should be removed. Uncheck this box if multiple occurrences of this rule match are required before the binding should be removed. If this box is unchecked, the Suspect Binding Removal Events Threshold field in Policy and Charging, and then Configuration, and then Policy DRA, and then Network-Wide Options on the NOAM controls how many Suspect Binding Removal Events must occur before a Session-Release RAR will be sent to the policy client to request removal of the binding.	Format: Check box Default: No (Unchecked) Range: Yes (Checked), No (Unchecked)
Comments	An optional comment to describe this suspect binding removal rule	Format: Text box Default: N/A Range: 0-64 characters

Inserting Suspect Binding Removal Rules

Use this task to insert a Suspect Binding Removal Rule.

The fields are described in Suspect Binding Removal Rules elements.

- 1. Click Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules.
- 2. Click Insert.
- 3. Enter a unique Rule Name in the **Rule Name** field.
- 4. Select an **Application Name** from the pull-down list.
- Select a Command Code from the pull-down list.
- 6. Select an Error Scenario Category from the pull-down list.
- 7. Enter a Result Code unless Error Scenario Category is Unable to Route.



- **8.** Enter an optional Vendor ID in the **Vendor ID** field only if a Result Code is found in the **Result Code** field.
- 9. Check or uncheck the **Remove Suspect Binding Immediately** box as needed.
- 10. Add an optional comment in the **Comments** box.

11. Click:

- Ok to save the new Suspect Binding Removal Rule to the Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules screen.
- Apply to save the new Suspect Binding Removal Rule and remain on the screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules screen without any changes.

If **Ok** or **Apply** is clicked and any of several conditions exist, an error message appear:

- A Field value is missing
- The syntax is invalid
- A Field value must be unique
- The maximum number of suspect binding removal rules allowed to be configured
- An operation failed because the suspect binding removal rule already exists
- An operation failed because the new suspect binding removal rule conflicts with another existing suspect binding removal rule
- An operation failed because inconsistent field values are configured
- An operation failed because the Rule Name cannot be Total because Total is a reserved Rule Name

Editing Suspect Binding Removal Rules

Use this task to edit a Suspect Binding Removal Rule.

The fields are described in Suspect Binding Removal Rules elements.

- Click Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules.
- 2. Click Edit.
- 3. Select an **Application Name** from the pull-down list.
- 4. Select a **Command Code** from the pull-down list.
- 5. Select an Error Scenario Category from the pull-down list.
- 6. Enter a Result Code unless Error Scenario Category is Unable to Route.
- 7. Enter an optional Vendor ID in the **Vendor ID** field only if a Result Code is found in the **Result Code** field.
- 8. Check or uncheck the **Remove Suspect Binding Immediately** box as needed.
- **9.** Add an optional comment in the **Comments** box.



10. Click:

- Ok to save the edited Suspect Binding Removal Rule to the Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules screen.
- Apply to save the edited Suspect Binding Removal Rule and remain on the screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules screen without any changes.

If **Ok** or **Apply** is clicked and any of several conditions exist, an error message appear:

- A Field value is missing
- The syntax is invalid
- A Field value must be unique
- The maximum number of suspect binding removal rules allowed to be configured
- An operation failed because the suspect binding removal rule already exists
- An operation failed because the edited suspect binding removal rule conflicts with another existing suspect binding removal rule
- An operation failed because inconsistent field values are configured
- An operation failed because the Rule Name cannot be Total because Total is a reserved Rule Name

Deleting a Suspect Binding Removal Rule

Use this task to delete a Suspect Binding Removal Rule.

- Click Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules.
- 2. Select the **Suspect Binding Removal Rule** to be deleted.
- 3. Click Delete.
- 4. Click:
 - Ok to delete the Suspect Binding Removal Rule.
 - Cancel to cancel the delete function and return to the Policy and Charging, and then Configuration, and then Policy DRA, and then Suspect Binding Removal Rules screen.

Access Point Names

An Access Point Name (APN) is a unique Packet Data network identifier. The PCA uses configured Access Point Names to validate APN entries received in Diameter signaling, and to apply appropriate Stale Session Timeout values during database audits.

PCRF pool selection allows the APN used by the UE to connect to the network is used to determine the PCRF pool. This allows multiple bindings to exist for a single IMSI, one for each PCRF pool. The Origin-Host of the PCEF sending the CCR-I can then be



used to select a PCRF sub-pool. Each APN is mapped to a PCRF Pool designated to manage policy bindings originated from that APN. In addition, a stale session timeout is assigned to the APN to control how long a session from the APN can remain idle before being subject to audit.

When an APN entry is added, new bindings from that APN are routed to a PCRF in the specified PCRF Pool (or a Sub-Pool if a matching PCRF Sub-Pool Selection Rule also exists). When an APN is mapped to a PCRF Pool using the Access Point Names screen, a check is performed to determine if the selected PCRF Pool is configured with a PRT mapping at each site. If at least one site does not have a mapping for the selected PCRF Pool, a confirmation dialog displays a warning as follows:

- If a PCRF Pool is not mapped to a PRT table for a site, a confirmation dialog is displayed on the APN GUI warning that Site *X* does not have a mapping defined for this PCRF Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN that is mapped to that PCRF Pool.
- If a site cannot be reached due to network errors, a confirmation dialog is displayed on the APN GUI warning that it cannot be determined whether Site *X* has a mapping defined for this PCRF Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN that is mapped to that PCRF Pool.

Single PCRF pool support is achieved by using the default pool, with all APNs mapped to that pool. This results in all bindings pointing to a single PCRF Pool.

If an APN is successfully deleted from the NOAMP GUI, the entry is internally marked as retired. Retired entries are not displayed on the GUI, but cannot be removed from the internal tables because that APN could still be referenced by any number of bindings. If you add a new APN with the same name as one that has been retired, the record comes out of retirement, but with the PCRF Pool and Stale Session Lifetime configured when the record was re-added.

The fields are described in Access Point Names elements.

On the **Policy and Charging**, and then **Configuration**, and then **Access Point Names** screen on the Active NOAM, you can perform a variety of actions:

- Filter the list of Access Point Names, to display only the desired Access Point Names.
- Sort the list entries in ascending or descending order by Access Point Names or by Stale Session Timeout, by clicking the column heading. By default, the list is sorted by Access Point Names in ascending numerical order.
- Work with PCRF Pool Names and Sub-Pools for PDRA APNs.
- Configure settings for Per IMSI Session Limiting.
- Click Insert.
 - You can add a new Access Point Name. See Inserting Access Point Names. If the maximum number of Access Point Names (2500) already exists in the system, the Policy and Charging, and then Configuration, and then Access Point Names [Insert] screen will not open, and an error message is displayed.
- Select an Access Point Name in the list, and click Edit.
 You can edit the selected Access Point Name. See Editing Access Point Names.
- Select an Access Point Name in the list, and click Delete to remove the selected Access Point Name. See Deleting an Access Point Name.



On the **Policy and Charging**, and then **Configuration**, and then **Access Point Names** screen on the SOAM, you can view the configured Access Point Names, and perform the several 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 5-26 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Access Point Names** screen.

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

Table 5-26 Access Point Names Elements

Elements (* indicates required field)	Description	Data Input Notes
* Access Point Name	The unique network identifier of a Packet Data Access Point.	Format: Text box; valid characters are alphabetic characters (A-Z and a-z), digits (0-9), hyphen (-), and period (.). Must begin and end with an alphabetic character or a digit.
		Default: N/A
		Range: 1 to 100
* Function	The PCA function that uses this Access Point. PCRF Pool	Format: Radio button
		Default: PDRA Only
•	is required to be configured for PDRA only.	Range: PDRA Only, OCDRA Only, and PDRA and OCDRA
PCRF Pool Name	The PCRF Pool associated with the Access Point Name.	Format: Dropdown Menu Default: Default PCRF Pool
	PCRF Pool Names in the row are hyperlinks to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools (Filtered) view screen filtered by the PCRF Pool Name.	Range: Configured PCRF Pools



Table 5-26 (Cont.) Access Point Names Elements

Elements (* indicates required field)	Description	Data Input Notes
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, and then Configuration, and then PCRF Sub-Pool Selection Rules screen. If the value is not zero, each Sub-Pool in the row is a hyperlink to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Sub-Pool Selection Rules (Filtered) view screen filtered by the PCRF Sub-Pool Selection Rule.	Format: List Range: N/A
	If the number of Sub-Pools is zero, this is not a hyperlink field.	
Maximum Allowed Sessions per IMSI	This setting is the maximum number of bound sessions allowed per IMSI for this APN	Format: Text box Default: 2 Range: 1-10
Per IMSI Session Exceeded Treatment	This setting defines the treatment of new binding capable session initiation	Format: Radio button Default: Route
attempts when the maximum number of bound sessions for an IMSI for this APN is exceed.	Range: Route, Reject	
	If Route is selected, the CCR-I message will be routed and the oldest bound session will be replaced. If Reject is selected, the CCR-I message will be rejected using the Diameter response code configured for SBR Error.	



Table 5-26 (Cont.) Access Point Names Elements

Elements (* indicates		
required field)	Description	Data Input Notes
* Stale Session Timeout (Hrs)	This setting is a time value (in hours), after which a session is considered to be stale. For PDRA, a session is considered stale only if no RAR/RAA messages are received in longer than this configured time. For OCDRA, a session is considered stale if no any in session messages are received in longer than this configured time. If a session's age exceeds this value, that session is eligible to be audited out of the database.	Format: Text box. Value must be numeric. Default: 168 hours (7 days) Range: 1-2400 (1 hour to 100 days)
	This value is used for sessions associated with this Access Point Name. For sessions which are not associated with any configured Access Point Names, the Default Stale Session Timeout value in the Policy and Charging Configuration General Options table is used.	
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.	Format: Text box Range: N/A
	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.	

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.

- Click Policy and Charging, and then Configuration, and then Access Point Names.
- Click Insert.
- 3. Enter a unique Access Point Name in the Access Point Name Value field.
- 4. Select the **Function**.
- 5. Select a PCRF Pool Name from the dropdown menu. This field contains all the qualified PCRF Pools configured from Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools. A qualified PCRF Pool is non-retired and has not been marked as Sub-Pool.

Note:

This step is only valid for PDRA Only or PDRA and OCDRA.

This identifies the PCRF Pool to which new bindings initiated from the Access Point Network are to be routed.

Note:

A retired PCRF Pool entry can be created by first adding a new PCRF Pool and then deleting it.

The Number of Sub-Pools field is a read-only field that displays the number of PCRF Sub-Pools associated with the selected PCRF Pool. The mapping between PCRF Pool and PCRF Sub-Pool is configured from the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Sub-Pool Selection Rules** screen.

- 6. Enter the Maximum Allowed Sessions per IMSI.
- 7. Select a Per IMSI Session Exceeded Treatment.
- 8. If a value other than the default Stale Session Timeout value is desired, enter the desired length of time in hours in the Stale Session Timeout (Hrs) **Value** field.

For sessions that are not associated with any configured Access Point Names, the default Stale Session Timeout value in the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **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.

9. Click:

- OK to save the new Access Point Name and return to the Policy and Charging, and then Configuration, and then Access Point Names screen.
- Apply to save the new Access Point Name and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Access Point Names screen without saving any changes.

If **OK** or **Apply** is clicked and any of several possible conditions exist, an error message appears:

- The entered Access Point Name is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range
- Any required field is empty (not entered)
- Adding the new Access Point Name would cause the maximum number of Access Point Names (2500) to be exceeded

Editing Access Point Names

Use this task to edit Access Point Stale Session Timeout values.



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.

- Click Policy and Charging, and then Configuration, and then Access Point Names.
- 2. Click Edit.
- 3. Select the Function.
- 4. Select a PCRF Pool Name from the PCRF Pool Name pulldown 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**, and then **Configuration**, and then **PCRF Sub-Pool Selection Rules** screen.



This step is only valid for PDRA Only or PDRA and OCDRA.

- 5. Edit the Maximum Allowed Sessions per IMSI.
- 6. Select a Per IMSI Session Exceeded Treatment.
- Enter the desired length of time in hours in the Stale Session Timeout (Hrs) Value field.

For sessions that are not associated with any configured Access Point Names, the default Stale Session Timeout value in the **Policy and Charging**, and then **Configuration**, and then **Network-Wide Options** table is used. The default is 168 hours (7 days), and the range is 1-2400 hours (1 hour to 100 days).

The Last Updated field is a read-only field that displays the date and time that this APN was created, or the last time the PCRF Pool Name was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding-capable session initiation requests. You can compare this date and time to the binding creation times when troubleshooting using the Policy Database Query Tool.

8. Click:

- OK to save the changes and return to the Policy and Charging, and then Configuration, and then Access Point Names screen.
- Apply to save the edited Access Point Name and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Access Point Names screen without saving any changes.

If **OK** or **Apply** is clicked and 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, then an error message appears.

Deleting an Access Point Name

Use this task to delete an Access Point Name.



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

- Click Policy and Charging, and then Configuration, and then Access Point Names.
- 2. Select the Access Point Name to be deleted.
- 3. Click Delete.
- 4. Click:



- OK to delete the Access Point Name.
- Click Cancel to cancel the delete function and return to the Policy and Charging, and then Configuration, and then Access Point Names screen.

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 screen.

General Options

On the **Policy and Charging**, and then **Configuration**, and then **General Options** screen on an Active NOAM, the General Options can be configured:



General Options is also available to be viewed on the SOAM. However, these options are only able to be sorted and filtered on the SOAM. Modifying these options is only permissible on the NOAM.

The fields are described in General Options elements.

General Options

- Indicate whether or not the Policy DRA function of PCA is enabled.
- Indicate whether or not the Online Charging DRA Function of PCA is enabled.

Audit Options

 Change the **Default Stale Session Timeout** value to a value other than the default value in the field.

This setting is a time value (in hours), after which a session is considered to be stale. For PDRA, a session is considered stale only if no RAR/RAA messages are received in longer than this configured time. For OCDRA, a session is considered stale if no any in session messages are received in longer than this configured time. If a session's age exceeds this value, that session is eligible to be audited out of the database.

This value is only used if a session is not associated with a configured Access Point Name in the Access Point Names configuration table. For sessions that are associated with a configured Access Point Name, the appropriate Stale Session value in the Access Point Name configuration table is used.

- Change the Binding Audit Session Query Rate, which is the maximum rate at which a binding SBR can send query messages to session servers to verify that sessions are still valid.
- Change the Audit Operation Rate
 For session SBRs the maximum rate at which Diameter sessions are checked for staleness.

For binding SBRs – the maximum rate at which binding session references are examined, if not already throttled by the Binding Audit Session Query Rate.



General Options elements

Table 5-27 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **General Options** screen on the NOAM.

Table 5-27 General Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
General Options		
Policy DRA Enabled	Indicates whether the Policy DRA Function of PCA is enabled	Format: Check box Range: Enabled (Checked) or Disabled (Unchecked) Default: Disabled (Unchecked)
Online Charging DRA Enabled	Indicates whether the Online	Format: Check box
	Charging DRA Function of PCA is enabled	Range: Enabled (Checked) or Disabled (Unchecked)
		Default: Disabled (Unchecked)
Audit Options		
* Default Stale Session	This setting is a time value	Format: Text box
Timeout	(in hours), after which a session is considered to be	Range: 1-2400 hours (1 hour to 100 days)
	stale For PDRA a session	Default: 168 hours (7 days)
	This value is only used if a session is not associated with a configured Access Point Name in the Access Point Names configuration table. For sessions that are associated with a configured Access Point Name, the appropriate Stale Session value in the Access Point Name configuration table is used.	
* Binding Audit Session Query Rate	The maximum rate at which a binding SBR can send query messages to session servers to verify that sessions are still valid.	Format: Text box Range: 5000-25000 records per second Default: 12000 per second



Table 5-27 (Cont.) General Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
* Audit Operation Rate	For session SBRs: the maximum rate at which Diameter sessions are checked for staleness. For binding SBRs: the maximum rate at which binding session references	Format: Text box Range: 25000-50000 per second Default: 50000 per second
	are examined, if not already throttled by the Binding Audit Session Query Rate.	

Online Charging DRA Configuration

This section describes the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA** screens on the NOAM and the SOAM.

OCSs

On an Active SOAM, the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCSs** screen lists the Online Charging Servers (OCS) Peer Nodes configured on a site.

The list of OCS Peer Nodes is updated by inserting, editing, or deleting an OCS Peer Node from the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCSs** screen at each site's SOAM.

OCSs elements

Table 5-28 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCSs** screen on the Active SOAM.



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



Table 5-28 OCSs Elements

Fields (* indicates required field)	Description	Data Input Notes
* OCS Peer Node Name	the OCS Peer Node to be included in the load distribution of new session initiation diameter request messages to OCSs.	
		Range: Configured Diameter Peer Nodes
		Note : The OCS Peer Node Name cannot be changed on the [Edit] screen.
Selecting a OCS Peer Node name (blue hyperlink) displays the Diameter , and then Configuration , and then Peer Nodes (Filtered) screen where Diameter Peer Nodes are filtered by the OCS Peer Node Name.	the [Edit] Screen.	
	Peer Nodes (Filtered) screen where Diameter Peer Nodes are filtered by the OCS Peer	
•	An optional comment to	Format: Text box
	describe the OCS Peer Node.	Range: 0-64 characters

Inserting OCSs

Use this task to insert (create new) OCSs.

The fields are described in OCSs elements.

- 1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCSs**.
- 2. Click Insert.
- 3. Select an OCS Peer Node Name from the dropdown menu.
- 4. Enter an optional comment in the **Comments** field.
- 5. Click:
 - OK to save the new OCS and return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs screen.
 - Apply to save the new OCS and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs screen without saving any changes.

If **OK** or **Apply** is clicked and any of several conditions exist, an error message appears:

- The entered OCS is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).
- Adding the new OCS would cause the maximum number of OCSs (2500) to be exceeded.



Editing OCSs

Use this task to edit OCSs.

1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCSs**.

The screen displays a list of the configured OCS Peer Nodes that are used when a new subscriber binding is created.

2. Click in the **Comments** field of the row to select the OCS to edit.

DO NOT click the blue OCS Peer Node Name (except to see the configuration of the Peer Node). The blue color indicates a hyper-link that opens the **Diameter**, and then **Configuration**, and then **Peer Nodes [Filtered]** screen to display the configuration information for the Peer Node.

3. Edit the **Comments** field for the selected OCS.

The OCS Peer Node name cannot be changed.

- 4. Click:
 - OK to save the change and return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs screen.
 - Apply to save the change and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs screen without saving any changes.

If **Apply** or **OK** is clicked and the selected **OCS Peer Node Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting an OCS

Use this procedure to delete an OCS.

This procedure describes the recommended steps for deleting an OCS from a Policy and Charging configuration. In this procedure, OCS refers to a Diameter peer of the PCA, which is sometimes referred to as an OCS Front-end.

The OCS procedure minimizes disruption to policy signaling by:

- Preventing sessions from creating new bindings to an OCS that has been removed
- Allowing sessions with existing bindings to continue to use an OCS that has been removed until those sessions terminate normally

The procedure describes the recommended steps for deletion of an OCS from a Policy and Charging configuration. In this procedure, OCS refers to a Diameter peer of the PCA, sometimes referred to as an OCS Front-End.



The PCRF removal procedure is restricted to SOAM servers.



- 1. Use **Diameter**, and then **Configuration**, and then **Peer Nodes** from the SOAM GUI screen to determine the Peer Node name of the OCS(s) being removed.
- 2. Use Diameter, and then Configuration, and then Route Groups from the SOAM GUI screen to filter by Peer Node with the corresponding Peer Node name of the OCS. This will display only the Route Groups that are associated with the OCS.
- 3. From the same GUI screen, determine if there are any Route Groups that contain other Peer Nodes in addition to the OCS to be removed.
 - There are generally at least two Route Groups for each OCS. One Route Group with only the specified OCS peer, and one or more Route Groups with the specified OCS peer plus other OCS peers. The goal is to leave the route group with only the specified OCS peer, but delete the OCS peer from the other route groups. This allows routing for existing bindings to the OCS peer, but prevents alternate routing to the OCS peer.
- 4. From the same GUI screen, edit each of the determined Route Groups and remove the OCS/OCS Front-End Peer Nodes from the Route Group.
 - This prevents alternate routing selection of the OCS peer being removed.
- 5. Use Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs from the SOAM GUI screen to delete the OCS.
 - This prevents new Bindings from using the OCS peer being removed.
- 6. After enough time has elapsed such that all Diameter sessions that could be bound to the OCS peer should have terminated normally, use Main Menu, and then Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs on the SOAM GUI screen to delete the route group containing only the OCS peer being removed.
- 7. Use **Diameter**, and then **Maintenance**, and then **Connections** from the SOAM GUI screen to find the connection for the OCS Peer Node and disable it
- **8.** Use **Diameter**, and then **Maintenance**, and then **Connections** from the SOAM GUI screen to delete the connection to the OCS Peer Node.
- Use Diameter, and then Configuration, and then Peer Nodes from the SOAM GUI screen to delete the Diameter Peer Node for the OCS being removed.

CTFs

On an Active SOAM, the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **CTFs** screen lists the CTF Peer Nodes for which the Session state is to be stored. This screen is only used if Session State Scope is set to Specific Messages in the Network-Wide Options Configuration on the NOAM

The list of CTF Peer Nodes is updated by inserting, editing, deleting a CTF Peer Node from the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **CTFs** screen at each site's SOAM.

CTFs elements

Table 5-29 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **CTFs** screen on the Active SOAM.



Note:

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

Table 5-29 CTFs Elements

Fields (* indicates required field)	Description	Data Input Notes
* CTF Peer Node Name	A name that uniquely identifies the CTF Peer Node.	Format: List
		Range: Configured Diameter
	Selecting a CTF Peer Node name (blue hyperlink) displays the Diameter , and then Configuration , and then Peer Nodes (Filtered) screen where Diameter Peer Nodes are filtered by the CTF Peer Node Name.	Peer Nodes
		Note : The CTF Peer Node Name cannot be changed on the [Edit] screen.
Comments	An optional comment to describe the CTF Peer Node.	Format: Text box Range:0-64 characters

Inserting CTFs

Use this task to insert(create new) CTFs.

The fields are described in CTFs elements.

- 1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **CTFs**.
- 2. Click Insert.
- 3. Enter a unique CTF Peer Node Name in the CTF Peer Node Name field.
 - This name uniquely identifies the CTF Peer Node.
- 4. Enter an optional comment in the **Comments** field.
- 5. Click:
 - OK to save the new CTF and return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then CTFs screen.
 - Apply to save the new CTF and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then CTFs screen without saving any changes.

If **OK** or **Apply** is clicked and any of several conditions exist, an error message appears:

- The entered CTF is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Any required field is empty (not entered).



 Adding the new CTF would cause the maximum number of CTFs (2500) to be exceeded.

Editing CTFs

Use this task to edit CTFs.

1. On the Active SOAM, click **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **CTFs**.

The screen displays a list of the configured CTF Peer Nodes that are used when a new subscriber binding is created.

- 2. Click in the **Comments** field of the row to select the CTF to edit.
- 3. Edit the **Comments** field for the selected CTF.

The CTF Peer Node name cannot be changed.

- 4. Click:
 - OK to save the change and return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then CTFs screen.
 - Apply to save the change and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then CTFs screen without saving any changes.

If **Apply** or **OK** is clicked and the selected **CTF Peer Node Name** entry no longer exists (was deleted by another user), an error message appears.

Deleting a CTF

Use this procedure to delete an CTF.

This procedure describes the recommended steps for deleting an CTF from a Policy and Charging configuration. In this procedure, CTF refers to a Diameter peer of the PCA, which is sometimes referred to as an CTF Front-end.

The CTF procedure minimizes disruption to policy signaling by:

- Preventing sessions from creating new bindings to an CTF that has been removed
- Allowing sessions with existing bindings to continue to use an CTF that has been removed until those sessions terminate normally

The procedure describes the recommended steps for deletion of an CTF from a Policy and Charging configuration. In this procedure, CTF refers to a Diameter peer of the PCA, sometimes referred to as an CTF Front-End.



The PCRF removal procedure is restricted to SOAM servers.

1. Use **Diameter**, and then **Configuration**, and then **Peer Nodes** from the SOAM GUI screen to determine the Peer Node name of the CTF(s) being removed.



- Use Diameter, and then Configuration, and then Route Groups from the SOAM GUI screen to filter by Peer Node with the corresponding Peer Node name of the CTF. This will display only the Route Groups that are associated with the CTF.
- From the same GUI screen, determine if there are any Route Groups that contain other Peer Nodes in addition to the CTF to be removed.
 - There are generally at least two Route Groups for each CTF. One Route Group with only the specified CTF peer, and one or more Route Groups with the specified CTF peer plus other CTF peers. The goal is to leave the route group with only the specified CTF peer, but delete the CTF peer from the other route groups. This allows routing for existing bindings to the OCS peer, but prevents alternate routing to the OCS peer.
- 4. From the same GUI screen, edit each of the determined Route Groups and remove the OCS/OCS Front-End Peer Nodes from the Route Group.
 - This prevents alternate routing selection of the OCS peer being removed.
- Use Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs from the SOAM GUI screen to delete the OCS.
 - This prevents new Bindings from using the OCS peer being removed.
- 6. After enough time has elapsed such that all Diameter sessions that could be bound to the OCS peer should have terminated normally, use Main Menu, and then Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCSs on the SOAM GUI screen to delete the route group containing only the OCS peer being removed.
- 7. Use **Diameter**, and then **Maintenance**, and then **Connections** from the SOAM GUI screen to find the connection for the OCS Peer Node and disable it
- **8.** Use **Diameter**, and then **Maintenance**, and then **Connections** from the SOAM GUI screen to delete the connection to the OCS Peer Node.
- Use Diameter, and then Configuration, and then Peer Nodes from the SOAM GUI screen to delete the Diameter Peer Node for the OCS being removed.

SBR Databases

Refer to the SBR User's Guide for how to configure SBR Databases.

Access Point Names

An Access Point Name (APN) is a unique Packet Data network identifier. The PCA uses configured Access Point Names to validate APN entries received in Diameter signaling, and to apply appropriate Stale Session Timeout values during database audits.

PCRF pool selection allows the APN used by the UE to connect to the network is used to determine the PCRF pool. This allows multiple bindings to exist for a single IMSI, one for each PCRF pool. The Origin-Host of the PCEF sending the CCR-I can then be used to select a PCRF sub-pool. Each APN is mapped to a PCRF Pool designated to manage policy bindings originated from that APN. In addition, a stale session timeout is assigned to the APN to control how long a session from the APN can remain idle before being subject to audit.

When an APN entry is added, new bindings from that APN are routed to a PCRF in the specified PCRF Pool (or a Sub-Pool if a matching PCRF Sub-Pool Selection Rule



also exists). When an APN is mapped to a PCRF Pool using the Access Point Names screen, a check is performed to determine if the selected PCRF Pool is configured with a PRT mapping at each site. If at least one site does not have a mapping for the selected PCRF Pool, a confirmation dialog displays a warning as follows:

- If a PCRF Pool is not mapped to a PRT table for a site, a confirmation dialog is displayed on the APN GUI warning that Site *X* does not have a mapping defined for this PCRF Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN that is mapped to that PCRF Pool.
- If a site cannot be reached due to network errors, a confirmation dialog is displayed on the APN GUI warning that it cannot be determined whether Site *X* has a mapping defined for this PCRF Pool. You can choose to continue, but with the knowledge that a call might fail at that site if a binding-capable session initiation request arrives with an APN that is mapped to that PCRF Pool.

Single PCRF pool support is achieved by using the default pool, with all APNs mapped to that pool. This results in all bindings pointing to a single PCRF Pool.

If an APN is successfully deleted from the NOAMP GUI, the entry is internally marked as retired. Retired entries are not displayed on the GUI, but cannot be removed from the internal tables because that APN could still be referenced by any number of bindings. If you add a new APN with the same name as one that has been retired, the record comes out of retirement, but with the PCRF Pool and Stale Session Lifetime configured when the record was re-added.

The fields are described in Access Point Names elements.

On the **Policy and Charging**, and then **Configuration**, and then **Access Point Names** screen on the Active NOAM, you can perform a variety of actions:

- Filter the list of Access Point Names, to display only the desired Access Point Names.
- Sort the list entries in ascending or descending order by Access Point Names or by Stale Session Timeout, by clicking the column heading. By default, the list is sorted by Access Point Names in ascending numerical order.
- Work with PCRF Pool Names and Sub-Pools for PDRA APNs.
- Configure settings for Per IMSI Session Limiting.
- · Click Insert.

You can add a new Access Point Name. See Inserting Access Point Names. If the maximum number of Access Point Names (2500) already exists in the system, the Policy and Charging, and then Configuration, and then Access Point Names [Insert] screen will not open, and an error message is displayed.

- Select an Access Point Name in the list, and click Edit.
 You can edit the selected Access Point Name. See Editing Access Point Names.
- Select an Access Point Name in the list, and click Delete to remove the selected Access Point Name. See Deleting an Access Point Name.

On the **Policy and Charging**, and then **Configuration**, and then **Access Point Names** screen on the SOAM, you can view the configured Access Point Names, and perform the several 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 5-30 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Access Point Names** screen.

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

Table 5-30 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
* Function	The PCA function that uses this Access Point. PCRF Pool is required to be configured for PDRA only.	Format: Radio button
		Default: PDRA Only
		Range: PDRA Only, OCDRA Only, and PDRA and OCDRA
PCRF Pool Name	The PCRF Pool associated with the Access Point Name.	Format: Dropdown Menu Default: Default PCRF Pool
	PCRF Pool Names in the row are hyperlinks to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools (Filtered) view screen filtered by the PCRF Pool Name.	Range: Configured PCRF Pools



Table 5-30 (Cont.) Access Point Names Elements

Elements (* indicates required field)	Description	Data Input Notes
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, and then Configuration, and then PCRF Sub-Pool Selection Rules screen. If the value is not zero, each Sub-Pool in the row is a hyperlink to the Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Sub-Pool Selection Rules (Filtered) view screen filtered by the PCRF Sub-Pool Selection Rule.	Format: List Range: N/A
	If the number of Sub-Pools is zero, this is not a hyperlink field.	
Maximum Allowed Sessions per IMSI	This setting is the maximum number of bound sessions allowed per IMSI for this APN	Format: Text box Default: 2 Range: 1-10
Per IMSI Session Exceeded Treatment	This setting defines the treatment of new binding capable session initiation attempts when the maximum number of bound sessions for an IMSI for this APN is exceed.	Format: Radio button Default: Route
		Range: Route, Reject
	If Route is selected, the CCR-I message will be routed and the oldest bound session will be replaced. If Reject is selected, the CCR-I message will be rejected using the Diameter response code configured for SBR Error.	



Table 5-30 (Cont.) Access Point Names Elements

Elements (* indicates		
required field)	Description	Data Input Notes
* Stale Session Timeout (Hrs)	This setting is a time value (in hours), after which a session is considered to be stale. For PDRA, a session is considered stale only if no RAR/RAA messages are received in longer than this configured time. For OCDRA, a session is considered stale if no any in session messages are received in longer than this configured time. If a session's age exceeds this value, that session is eligible to be audited out of the database.	Format: Text box. Value must be numeric. Default: 168 hours (7 days) Range: 1-2400 (1 hour to 100 days)
	This value is used for sessions associated with this Access Point Name. For sessions which are not associated with any configured Access Point Names, the Default Stale Session Timeout value in the Policy and Charging Configuration General Options table is used.	
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.	Format: Text box Range: N/A
	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.	

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.

- Click Policy and Charging, and then Configuration, and then Access Point Names.
- Click Insert.
- 3. Enter a unique Access Point Name in the Access Point Name Value field.
- 4. Select the **Function**.
- 5. Select a PCRF Pool Name from the dropdown menu. This field contains all the qualified PCRF Pools configured from Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools. A qualified PCRF Pool is non-retired and has not been marked as Sub-Pool.



This step is only valid for PDRA Only or PDRA and OCDRA.

This identifies the PCRF Pool to which new bindings initiated from the Access Point Network are to be routed.



A retired PCRF Pool entry can be created by first adding a new PCRF Pool and then deleting it.

The Number of Sub-Pools field is a read-only field that displays the number of PCRF Sub-Pools associated with the selected PCRF Pool. The mapping between PCRF Pool and PCRF Sub-Pool is configured from the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Sub-Pool Selection Rules** screen.

- 6. Enter the Maximum Allowed Sessions per IMSI.
- 7. Select a Per IMSI Session Exceeded Treatment.
- 8. If a value other than the default Stale Session Timeout value is desired, enter the desired length of time in hours in the Stale Session Timeout (Hrs) **Value** field.

For sessions that are not associated with any configured Access Point Names, the default Stale Session Timeout value in the **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **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.

9. Click:

- OK to save the new Access Point Name and return to the Policy and Charging, and then Configuration, and then Access Point Names screen.
- Apply to save the new Access Point Name and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Access Point Names screen without saving any changes.

If **OK** or **Apply** is clicked and any of several possible conditions exist, an error message appears:

- The entered Access Point Name is not unique (already exists).
- Any fields contain a value that contains invalid characters or is out of the allowed range
- Any required field is empty (not entered)
- Adding the new Access Point Name would cause the maximum number of Access Point Names (2500) to be exceeded

Editing Access Point Names

Use this task to edit Access Point Stale Session Timeout values.



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.

- Click Policy and Charging, and then Configuration, and then Access Point Names.
- 2. Click Edit.
- 3. Select the Function.
- 4. Select a PCRF Pool Name from the PCRF Pool Name pulldown 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**, and then **Configuration**, and then **PCRF Sub-Pool Selection Rules** screen.



This step is only valid for PDRA Only or PDRA and OCDRA.

- 5. Edit the Maximum Allowed Sessions per IMSI.
- 6. Select a Per IMSI Session Exceeded Treatment.
- Enter the desired length of time in hours in the Stale Session Timeout (Hrs) Value field.

For sessions that are not associated with any configured Access Point Names, the default Stale Session Timeout value in the **Policy and Charging**, and then **Configuration**, and then **Network-Wide Options** table is used. The default is 168 hours (7 days), and the range is 1-2400 hours (1 hour to 100 days).

The Last Updated field is a read-only field that displays the date and time that this APN was created, or the last time the PCRF Pool Name was changed, whichever is most recent. This field records the time and date of changes that might affect routing of binding-capable session initiation requests. You can compare this date and time to the binding creation times when troubleshooting using the Policy Database Query Tool.

8. Click:

- OK to save the changes and return to the Policy and Charging, and then Configuration, and then Access Point Names screen.
- Apply to save the edited Access Point Name and remain on this screen.
- Cancel to return to the Policy and Charging, and then Configuration, and then Access Point Names screen without saving any changes.

If **OK** or **Apply** is clicked and 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, then an error message appears.

Deleting an Access Point Name

Use this task to delete an Access Point Name.



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

- Click Policy and Charging, and then Configuration, and then Access Point Names.
- 2. Select the Access Point Name to be deleted.
- 3. Click Delete.
- 4. Click:



- OK to delete the Access Point Name.
- Click Cancel to cancel the delete function and return to the Policy and Charging, and then Configuration, and then Access Point Names screen.

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 screen.

OCS Session State

On an Active NOAM, the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCS Session State** screen lists the network-wide list of Online Charging Servers (OCSs), listed by their Realm and FQDN. It is used to configure the Session State setting for OCSs.

The list of OCSs is updated by inserting or deleting an OCS from the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCSs** screen at each site's SOAM. Additionally, the Realm and FQDN are configured from each site's **Diameter**, and then **Configuration**, and then **Peer Nodes** screen on the SOAM.

Once the list of OCSs is populated, the options become available:

- Editing whether or not OCS Session State is enabled
- Pausing the updating of the OCS list

OCS Session State elements

Table 5-31 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCS Session State** screen.



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

Table 5-31 OCS Session State Elements

Fields (* indicates required field)	Description	Data Input Notes
* Realm	Realm of this Peer Node.	Format: Text box
	Realm is a case-insensitive string consisting of	Default: N/A
	a list of labels separated by dots, where a label may contain letters, digits, dashes (-) and underscore (_). A label must start with a letter, digit or underscore and must end with a letter or digit. Underscores may be used only as the first character. A label must be at most 63 characters long and a Realm must be at most 255 characters long.	Range: A valid Realm



Table 5-31 (Cont.) OCS Session State Elements

Fields (* indicates	Describetion	Data Innuit Nation
required field)	Description	Data Input Notes
* FQDN	Fully Qualified Domain Name of this Peer Node.	Format: Text box
	FQDN is a case-insensitive string consisting of	Default: N/A
	a list of labels separated by dots, where a label may contain letters, digits, dashes (-) and underscore (_). A label must start with a letter, digit or underscore and must end with a letter or digit. Underscores may be used only as the first character. A label must be at most 63 characters long and a FQDN must be at most 255 characters long.	Range: A valid FQDN
Session State	Setting to enable Session State for OCSs.	Format: Check box
Enable	Check this box if the sessions are to be maintained for this OCS. The Sessions shall be maintained if the Session State Scope is set to Policy and Charging , and then Configuration , and then Online Charging DRA , and then Network-Wide Options configuration or if Session State Scope is set to Specific Messages and this Session State Enabled setting is checked.	Default: No (unchecked) - Do not maintain session states
		Range: Yes (checked) - Maintain session states, or No (unchecked) - Do not maintain sessions states

Editing OCS Session State

Use this task to edit an OCS Session State.

- 1. On the Active NOAM, click **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCS Session State**.
- 2. Click Edit.
- 3. The **Realm** and **FQDN** fields are disabled and cannot be edited from this screen.
- 4. Check or uncheck the box to Enable or Disable OCS Session State.
- 5. Click:
 - OK to save the edited OCS Session State and return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCS Session State screen.
 - Apply to save the edited OCS Session State and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then OCS Session State screen without saving any changes.

If **OK** or **Apply** is clicked and a required field no longer exists, an error message appears.



Realms

The **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **Realms** screen on an NOAM contains the list of Online Charging network realms for which the Session state is stored.



This screen is only use if **Session State Scope** is set to Specific Messages on the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **Network-Wide Options** screen.

Realms elements

Table 5-32 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **Realms** screen.



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

The Realms table lists the Online Charging network realms for which the Session state is to be stored. This table is only used if Session State Scope is set to Specific Messages in the Network-Wide Options Configuration.

Table 5-32 Realms Elements

Fields (* indicates required field)	Description	Data Input Notes
* Realm Name	Realm name is a case- insensitive string consisting of a list of lables separated by dots, where a label may contain letters, digits, dashes(-), and underscore(_). A label must start with a letter, digit, or underscore and must end with a letter or digit. Underscores may be used only as the first character. A label can be at most 63 characters long and a Realm can be at most 255 characters long.	Format: text box Default: N/A Range: 1-1000 entries
Comments	An optional comment to provide more information about the purpose of this PCRF Pool or Sub-Pool.	Format: Text box Default: N/A Range: 0-64 characters



Inserting Realms

Use this task to insert (create new) Realms.

- 1. On the Active NOAM, click **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **Realms**.
- Click Insert.
- 3. Enter a unique Realm Name in the **Realm Name** field.
- 4. If desired, enter an optional comment in the **Comments** field to describe the Realm. The entry must be characters in the range of 0 to 64, and the default is N/A.
- 5. Click:
 - OK to save the new Realm name and return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then Realms screen.
 - Apply to save the new Realm name and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then Realms screen without saving any changes.

If **OK** or **Apply** is clicked and any of several conditions exist, an error message appears:

- Any required field is empty (not entered).
- Any fields contain a value that contains invalid characters or is out of the allowed range.
- Adding a new Realm would cause the maximum number of Realms (1000) to be exceeded.
- The entered Realm name is not unique (already exists).

Editing Realms

Use this task to edit a Realm.

- 1. On the Active NOAM, click **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **Realms**.
- 2. Click Edit.
- 3. Edit the unique Realm Name in the **Realm Name** field.
- 4. If desired, edit an optional comment in the **Comments** field.
- 5. Click:
 - OK to save the edited Realm name and return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then Realms screen.
 - **Apply** to save the edited Realm name and remain on this screen.
 - Cancel to return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then Realms screen without saving any changes.



If **OK** or **Apply** is clicked and any of several conditions exist, an error message appears:

- Any required field no longer exists
- Any fields contain a value that contains invalid characters or is out of the allowed range.

Deleting Realms

Use this task to delete a Realm.

- 1. On the Active NOAM, click **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **Realms**.
- 2. Select the **Realm** to be deleted.
- 3. Click Delete.
- 4. Click:
 - OK to delete the Realm.
 - Cancel to cancel the delete function and return to the Policy and Charging, and then Configuration, and then Online Charging DRA, and then Realms screen.

If **OK** is clicked and the selected Realm no longer exists (it was deleted by another user), an error message is displayed, and the Realms screen is refreshed. The row that was selected is no longer displayed in the list.

Network-Wide Options

On the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **Network-Wide Options** screen on an Active NOAM, the Network-Wide Options can be configured:

The fields are described in Network-Wide Options elements.

Session Options

- Set the scope of messages for which Session State will be stored.
- Set the action to be performed if an in-session Request message cannot be successfully processed due to the inability to retrieve session state associated with the received Session-Id from the Session SBR (for example, session state is not found or an SBR error is encountered).

OCS Selection Options

 Set the operating mode for selecting the OCS Server for routing the Session Initiation Request messages.

Network-Wide Options elements

Table 5-33 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **Network-Wide Options** screen on the NOAM.



 Table 5-33
 Online Charging DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
Session Options		
Session State Scope	This sets the scope of messages for which Session State will be stored. Format: Pulldown menu Range: None, All Message Specific Messages	
	Select All Messages to store Session State for all messages. Select None to disable Session State for all messages. Select Specific Messages to store Session State only if the CTF client is configured in the CTFs configuration or OCS is configured with Session State as enabled in OCSs configuration or realm is configured in Realms configuration.	Default: None
Session State Unavailable Action	Sets the action to be performed if an in-session Request message cannot be successfully processed due to the inability to retrieve session state associated with the received Session-Id from the Session SBR (for example, session state is not found or an SBR error is encountered). Route to Peer will route the message to a peer using the Peer Routing Table. Send Answer will abandon message processing and send an Answer response containing Answer Result-Code value configured for Session Not Found or SBR	Format: Pulldown menu Range: Send Answer, Route To Peer Default: Send Answer
	Error.	



Table 5-33 (Cont.) Online Charging DRA Network-Wide Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
OCS Pool Selection Mode	mode for selecting the OCS Server for routing the Session Initiation Request messages. Range: Single Pools	
When Single Pool m selected, the Session Requests are distribut weighted round-robin among all available C servers connected to Node. When Multiple Pools selected, the Session Requests are routed OCS server identified	When Single Pool mode is selected, the Session Initiation Requests are distributed in a weighted round-robin scheme among all available OCS servers connected to this Node.	Default: Single Pool
	When Multiple Pools mode is selected, the Session Initiation Requests are routed to an OCS server identified by RBAR in a specific pool of OCS servers.	



Note:

Keep these consideration in mind when working with network-wide options:

- If **Apply** is clicked and the Session State Scope transitioned from None to All Messages, a confirmation dialog with a checkbox shall be displayed containing the text: IMPORTANT! The Session State for all the messages will be enabled. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected. Check the checkbox and click OK to continue, otherwise click **Cancel**.
- If Apply is clicked and the Session State Scope transitioned from None
 to Specific Messages, a confirmation dialog with a checkbox shall be
 displayed containing the text: IMPORTANT! The Session State Scope
 Specific Messages requires the OCSs, CTFs or Realms to be configured
 for maintaining Session State. The Session State may not be found for
 already established sessions and the subsequent requests for already
 established sessions may be rejected.
- If Apply is clicked and the Session State Scope transitioned from All Messages to Specific Messages, a confirmation dialog with a checkbox shall be displayed containing the text: IMPORTANT! The Session State Scope Specific Messages requires the OCSs, CTFs or Realms to be configured for maintaining Session State. Some or all subsequent in-session messages may be rejected if Destination-Host AVP is not present in them. Check the checkbox and click OK to continue, otherwise click Cancel.
- If Apply is clicked and the Session State Scope transitioned from Specific Messages to All Messages, a confirmation dialog with a checkbox shall be displayed containing the text: IMPORTANT! The Session State for all the messages will be enabled. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected. Check the checkbox and click OK to continue, otherwise click Cancel.
- If the confirmation dialog for Session State Scope is cancelled by clicking Cancel, control is returned to the Network-Wide Options screen with no data committed.
- If the confirmation dialog for Session State Scope is confirmed by checking the checkbox and clicking OK and no OCS or CTF or Realm is configured for session state maintenance, a Warning Box is displayed on the Network-Wide Options screen containing the text: Session State Scope is configured as Specific Messages but no OCS, CTF or Realm is configured for Session State maintenance. The configured data is saved in the configuration database.

Setting Network-Wide Options

Use this task to set Online Charging DRA Network-Wide Options on the NOAM.

The fields are described in Network-Wide Options elements.

The Network-Wide Options associated with Online Charging DRA can be set:



- Setting the Session State Scope
- Setting the action to be taken if the Session State is Unavailable
- Setting the OCS Pool Selection Mode
- Click Policy and Charging, and then Configuration, and then Online Charging DRA, and then Network-Wide Options.
- 2. Select a **Session State Scope** from the pulldown list.
- 3. Select a Session State Unavailable Action from the pulldown list.
- 4. Select an OCS Pool Selection Mode from the pulldown list.
- 5. Click:
 - Apply to save the changes and remain on this screen.
 - Cancel to discard changes and remain on the Policy and Charging, and then Configuration, and then Online Charging DRA, and then Network-Wide Options screen.

If **Apply** is clicked and if a problematic condition exists, a warning message appears:

- If Apply is clicked and the Session State Scope transitioned from None to All Messages, a confirmation dialog with a checkbox shall be displayed containing the text: IMPORTANT! The Session State for all the messages will be enabled. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected. Check the checkbox and click OK to continue, otherwise click Cancel..
- If Apply is clicked and the Session State Scope transitioned from None to Specific Messages, a confirmation dialog with a checkbox shall be displayed containing the text: IMPORTANT! The Session State Scope Specific Messages requires the OCSs, CTFs or Realms to be configured for maintaining Session State. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected..
- If Apply is clicked and the Session State Scope transitioned from All Messages to Specific Messages, a confirmation dialog with a checkbox shall be displayed containing the text: IMPORTANT! The Session State Scope Specific Messages requires the OCSs, CTFs or Realms to be configured for maintaining Session State. Some or all subsequent in-session messages may be rejected if Destination-Host AVP is not present in them. Check the checkbox and click OK to continue, otherwise click Cancel..
- If Apply is clicked and the Session State Scope transitioned from Specific Messages to All Messages, a confirmation dialog with a checkbox shall be displayed containing the text: IMPORTANT! The Session State for all the messages will be enabled. The Session State may not be found for already established sessions and the subsequent requests for already established sessions may be rejected. Check the checkbox and click OK to continue, otherwise click Cancel.
- If the confirmation dialog for Session State Scope is cancelled by clicking Cancel, control is returned to the Network-Wide Options screen with no data committed.
- If the confirmation dialog for Session State Scope is confirmed by checking the checkbox and clicking OK and no OCS or CTF or Realm is configured for



session state maintenance, a Warning Box is displayed on the Network-Wide Options screen containing the text: Session State Scope is configured as Specific Messages but no OCS, CTF or Realm is configured for Session State maintenance. The configured data is saved in the configuration database.

Error Codes

For each Policy and Charging Site, the Diameter Error Result Code value to send to the Request initiator for policy related errors can be configured according to which condition applies. Each condition can be mapped to a different Result Code for each supported interface. Result Codes can be Diameter IANA defined or experimental.

When PCRF Pooling is enabled, new binding cannot be created unless the binding-capable session initiation request contains a configured APN. If the binding-capable session initiation request arrives with either no APN, or an APN that is not configured in the Access Point Names table, the request is answered by Policy DRA using a configurable error response code. To configure the Diameter response code for this scenario, a new Missing Or Unconfigured APN error condition has been added to the existing SOAM error. This error response applies to all binding capable interfaces (for example, Gx, Gxx, and S9) and can be configured with either an IANA Diameter result code, or an experimental result code and vendor-id.

Three-digit error codes in Diameter Error-Message AVPs indicate exactly why a slave session could not be routed. This provides more robust troubleshooting using Diameter capture tools.

A 3-digit error code is an identifier to uniquely identify a specific error scenario (not error category) encountered in a Diameter Answer message generated by PCA. 3-digit codes are unique across all DSR layers (DSR connection layer, routing layer and application layer) and all DSR applications (such as PCA, RBAR, FABR, and IDIH) for errors they represent. The ranges of 500-549 and 850-899 are for the PCA application, while the DSR connection layer, routing layer and other DSR applications uses other non-overlapping ranges. Multiple errors may belong to a same error category and are associated with a same Result-Code. It is the 3-digit code that can distinguish an error from others. Users should search for the 3-digit code when identifying an error if possible and available.



The error conditions in this table are GUI-configurable.



Table 5-34 PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Missing or Unconfigured APN	PDRA	Policy-related binding capable session initiation request messages	3002	500	Missing or Unconfigured APN. Binding capable session initiation request is received with no APN
Missing or Unconfigured APN	PDRA	Policy-related binding capable session initiation request messages	3002	501	Missing or Unconfigured APN. Binding capable session initiation request is received with APN, but the APN is not configured in the APN configuration.
Unable To Route	PDRA	Policy-related binding capable and dependent session initiation request messages	3002	502	Unable To Route. Request message is received and a binding with a PCRF was found. P-DRA cannot route the request to PCRF due to DSR queue full error.
Unable To Route	PDRA	Policy-related binding capable and dependent session initiation request messages	3002	3-digit error code from DRL	Unable To Route. Request message is received and a binding with a PCRF was found. P-DRA cannot route the request to PCRF due to PCRF being unreachable. DRL Text string.



Table 5-34 (Cont.) PCA Error Conditions

Error		Applied Diameter Interface/	Default	Error- Message	
Category	Functionality		Result-Code	Suffix	Error Text
No Usable Keys In Binding Dependent Message	PDRA	Policy-related binding dependent session initiation request messages	3002	503	No Usable Keys In Binding Dependent Message. No binding key in Binding Key Priority GUI can be matched, or no key is included in the binding dependent message.
Binding Not Found	PDRA	Policy-related binding dependent session initiation request messages	3002	505	Binding Not Found. Binding record is not found after examining all configured binding keys in Diameter message.
SBR Error	PDRA	Policy-related binding capable and dependent session initiation, update or terminate answer messages,	3002	507	SBR Error. ComAgent timeout.
SBR Error	PDRA	Policy-related binding capable and dependent session initiation, update or terminate answer messages	3002	508	SBR Error. SBR database error prevents SBR from reading, writing or deleting a record,



Table 5-34 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Session Not Found	PDRA	Policy-related binding capable and dependent session update or terminate request messages	3002	509	Session Not Found. Session record does not exist for given session ID
PCA Unavailable Or Degraded	PDRA/ OCDRA	Any Diameter Requests forwarded to PCA	3002	305	PCA Unavailable or Degraded
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	520	SBR Error. Binding capable session initiation request received, but no PCRFs are configured at the site, or PCRF ID is not found in PCRF table.
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	521	SBR Error. Maximum number of sessions per binding is exceeded that fails the binding creation for given subscriber's key.
Session ID is missing from Request	PDRA	Any Policy- related Diameter Requests forwarded to P-DRA	5005	522	Session ID is missing from Request



Table 5-34 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
CC - Request - Type AVP is missing from CCR message	PDRA	Policy-related binding capable session initiation, update or terminate request messages	5005	523	CC-Request- Type AVP is missing from CCR message
Not In Use					
Invalid AVP value in request message	PDRA	Any Policy- related Diameter Requests forwarded to P-DRA	5004	525	Invalid AVP value in request message
Destination - Host AVP is missing in in- session request	PDRA	Policy-related binding capable update and terminate request and dependent session initiation update or terminate request messages	5012	506	Destination- Host AVP is missing in in- session request
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	510	Unable To Route. A slave session could not be routed because on polling the slave sessionRef was no longer in the binding database.
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	511	Unable To Route. A slave session could not be routed because on polling the master sessionRef was no longer in the binding database.



Table 5-34 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	512	Unable To Route. A slave session could not be routed because on polling the master sessionRef was early too long.
SBR Error	PDRA	Policy-related Requests and Answers	3002	504	SBR Error. ComAgent unavailable when sending stack event to SBR
Unsupported Application ID	PDRA/ OCDRA	Diameter Requests	3007	530	Application ID unsupported by PCA
Command Code and App ID no match	PDRA	Policy-related Requests and Answers	5019	531	Command Code does not match App ID or not exist
Unable To Route	PDRA	Policy-related binding capable session initiation request	3002	513	Unable To Route. A slave session could not routed because on polling the master session and internal error occurred.
PCA Functionality Unavailable or Disabled	PDRA	Policy related binding capable and dependent session update or terminate request messages	3002	532	PCA Functionality Unavailable or Disabled. Policy DRA Function Disabled.
PCA Functionality Unavailable or Disabled	PDRA	Policy related binding capable and dependent session update or terminate request messages	3002	533	PCA Functionality Unavailable or Disabled. Policy DRA Function Unavailable.



Table 5-34 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
PCA Functionality Unavailable or Disabled	OCDRA	Online Charging related binding independent session request messages	3002	534	PCA Functionality Unavailable or Disabled. Online Charging DRA Function Disabled.
PCA Functionality Unavailable or Disabled	OCDRA	Online Charging related binding independent session request messages	3002	535	PCA Functionality Unavailable or Disabled. Online Charging DRA Function Unavailable
Session ID is missing from Request	OCDRA	Any Online Charging - related Diameter Requests forwarded to OC-DRA	5005	536	Session ID is missing from Request
CC - Request - Type AVP is missing from CCR message	OCDRA	Any Online Charging- related Diameter Requests forwarded to OC-DRA	5005	537	CC-Request- Type AVP is missing from CCR message
Invalid AVP value in request message	OCDRA	Any Online Charging- related Diameter Requests forwarded to OC-DRA	5004	538	Invalid AVP value in request message
Not In Use	OCDDA	Online	2002	F 40	Unabla Ta
Unable To Route	OCDRA	Online Charging- related binding independent session request messages	3002	540	Unable To Route. Request message is received, OC- DRA cannot route the request to OCS due to DSR queue full error.



Table 5-34 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Unable To Route	OCDRA	Online Charging- related binding independent session initiation request messages	3002	539	Unable To Route. Request message can not be routed to peer node. DIAM-ERR- XXXX- XXX:DRL Text string.
SBR Error	OCDRA	Online Charging- related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	541	SBR Error. ComAgent timeout.
SBR Error	OCDRA	Online Charging- related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	542	SBR Error. SBR database error prevents SBR from reading, writing or deleting a record,
SBR Error	OCDRA	Online Charging- related binding independent session update or terminate answer messages, if session state or topology hiding applies	5012	543	SBR Error . ComAgent unavailable when sending stack event to SBR,

Table 5-34 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
Session Not Found	OCDRA	Online Charging- related session update or terminate request messages, if session state or topology hiding applies	5002	544	Session Not Found. Session record does not exist for given session ID
Command Code and App ID no match	OCDRA	Online Charging- related Requests.	3001	545	Command Code does not match App ID or not exist
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	546	SBR Error. A binding capable session initiation request could not be routed, maximum sessions per IMSI per APN limit is exceeded.
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	547	SBR Error. A binding capable session initiation request could not be routed, maximum sessions per IMSI per APN limit is exceeded and existing sessions could not be replaced because binding is in early state.



Table 5-34 (Cont.) PCA Error Conditions

Error Category	Functionality	Applied Diameter Interface/ Message	Default Result-Code	Error- Message Suffix	Error Text
SBR Error	PDRA	Policy-related binding capable session initiation request messages	3002	548	SBR Error. A binding capable session initiation request could not be routed, maximum sessions per IMSI per APN limit is exceeded and existing sessions could not be replaced because Maximum Early Binding lifetime is not elapsed for existing sessions.

On the **Policy and Charging**, and then **Configuration**, and then **Error Codes** screen on the SOAM:

Select an Error Condition in the list and click Edit.
 You can edit the selected Error Code. See Editing Error Codes.

The fields are described in Error Codes elements.

Error Codes elements

Table 5-36 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Error Codes** screens. Data Input Notes apply to the Error Codes [Edit] screen; the View screen is read-only.

The Error Codes define the Result Codes to be returned for various Policy and Charging Error Conditions. Each Error Condition will return the Result Code configured for each applicable Diameter interface.

Table 5-35 indicates the Diameter interfaces that are supported for each Error Code.

The default Result Code is 3002-DIAMETER_UNABLE_TO_DELIVER.



Table 5-35 Interfaces Supported for Each Error Code

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

Table 5-36 Error Codes Elements

Fields (* indicates required field)	Description	Data Input Notes
* Error Condition	The name of the selected Policy and Charging Error Condition.	View only; cannot be edited
* Gx/Gxx, Result Code	The Result Code to be returned on the Gx and Gxx interfaces	Format: Text box Range: 1-9999 Default: 3002
Gx/Gxx Vendor ID	The Vendor ID that corresponds with the Gx and Gxx interfaces.	Format: Text box Range: 1-4294967295
	The Vendor ID means the RFC standard error code will be sent.	
* Rx Result Code	The Result Code to be returned to the Rx interface.	Format: Text box Range: 1-9999 Default: 3002
Rx Vendor ID	The Vendor ID that corresponds with the Rx interface.	Format: Text box Range: 1-4294967295
	The Vendor ID means the RFC standard error code will be sent.	
* S9 Result Code	The Result Code to be returned to the S9 interface.	Format: Text box Range: 1-9999 Default: 3002



Table 5-36 (Cont.) Error Codes Elements

Fields (* indicates required	Donasis di su	Data Invest Nata
field)	Description	Data Input Notes
S9 Vendor ID	The Vendor ID that corresponds the S9 interface. The Vendor ID means the RFC standard error code will be sent.	Format: Text box Range: 1-4294967295
* Gx-Prime Result Code	The Result Code to be returned on the Gx-Prime interface	Format: Text Box Range: 1-9999 Default: 3002
Gx-Prime Vendor ID	The Vendor ID that corresponds with the Gx-Prime interface.	Format: Text Box Range: 1-4294967295
	The Vendor ID means the RFC standard error code will be sent.	
* Gy/Ro Result Code	The Result code to be	Format: Text Box
	returned on the Gy/Ro interface.	Range: 1-9999
Gy/Ro Vendor ID	The Vendor ID that	Format: Text Box
	corresponds with the experimental code for the Gy/Ro interface.	Range: 1-4294967295

Editing Error Codes

Use this task to edit Error Codes on the Active SOAM.

The fields are described in Error Codes elements.

- 1. Click Policy and Charging, and then Configuration, and then Error Codes.
- 2. Select the **Error Condition** that you want to edit.
- 3. Click Edit.

The fields that appear on the **Policy and Charging**, and then **Configuration**, and then **Error Codes [Edit]** screen are dependent on the Error Condition that was selected.

- 4. Edit the fields to define the selected Error Condition.
- 5. Click:
 - OK to save the changes and return to the Policy and Charging, and then Configuration, and then Error Codes screen
 - Apply to save the changes and remain on this screen
 - Cancel to discard the changes and return to the Policy and Charging, and then Configuration, and then Error Codes screen

If **OK** or **Apply** is clicked and any of several 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).

General Options

On the **Policy and Charging**, and then **Configuration**, and then **General Options** screen on an Active NOAM, the General Options can be configured:



General Options is also available to be viewed on the SOAM. However, these options are only able to be sorted and filtered on the SOAM. Modifying these options is only permissible on the NOAM.

The fields are described in General Options elements.

General Options

- Indicate whether or not the Policy DRA function of PCA is enabled.
- Indicate whether or not the Online Charging DRA Function of PCA is enabled.

Audit Options

 Change the **Default Stale Session Timeout** value to a value other than the default value in the field.

This setting is a time value (in hours), after which a session is considered to be stale. For PDRA, a session is considered stale only if no RAR/RAA messages are received in longer than this configured time. For OCDRA, a session is considered stale if no any in session messages are received in longer than this configured time. If a session's age exceeds this value, that session is eligible to be audited out of the database.

This value is only used if a session is not associated with a configured Access Point Name in the Access Point Names configuration table. For sessions that are associated with a configured Access Point Name, the appropriate Stale Session value in the Access Point Name configuration table is used.

- Change the Binding Audit Session Query Rate, which is the maximum rate at which a binding SBR can send query messages to session servers to verify that sessions are still valid.
- Change the Audit Operation Rate
 For session SBRs the maximum rate at which Diameter sessions are checked for staleness.

For binding SBRs – the maximum rate at which binding session references are examined, if not already throttled by the Binding Audit Session Query Rate.

General Options elements

Table 5-37 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **General Options** screen on the NOAM.



Table 5-37 General Options Elements

Fields (* indicates a	Description	Data Input Notes
required field)	Description	Data Input Notes
General Options Policy DRA Enabled	Indicates whether the Policy DRA Function of PCA is enabled	Format: Check box Range: Enabled (Checked) or Disabled (Unchecked) Default: Disabled (Unchecked)
Online Charging DRA Enabled		Format: Check box
	Charging DRA Function of PCA is enabled	Range: Enabled (Checked) or Disabled (Unchecked)
		Default: Disabled (Unchecked)
Audit Options		
* Default Stale Session Timeout	This setting is a time value (in hours), after which a session is considered to be	Format: Text box Range: 1-2400 hours (1 hour to 100 days)
	stale. For PDRA, a session is considered stale only if no RAR/RAA messages are received in longer than this configured time. For OCDRA, a session is considered stale if no any in session messages are received in longer than this configured time. If a session's age exceeds this value, that session is eligible to be audited out of the database.	Default: 168 hours (7 days)
	This value is only used if a session is not associated with a configured Access Point Name in the Access Point Names configuration table. For sessions that are associated with a configured Access Point Name, the appropriate Stale Session value in the Access Point Name configuration table is used.	
* Binding Audit Session Query Rate	The maximum rate at which a binding SBR can send query messages to session servers to verify that sessions are still valid.	Format: Text box Range: 5000-25000 records per second Default: 12000 per second



Table 5-37 (Cont.) General Options Elements

Fields (* indicates a required field)	Description	Data Input Notes
* Audit Operation Rate For session SBRs: the maximum rate at which Diameter sessions are checked for staleness. For binding SBRs: the maximum rate at which binding session references are examined, if not already throttled by the Binding Audit Session Query Rate.	maximum rate at which Diameter sessions are	Format: Text box Range: 25000-50000 per second
	Default: 50000 per second	

Alarm Settings



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

On the **Policy and Charging**, and then **Configuration**, and then **Alarm Settings** screen on an SOAM, the user can view the configured Alarm Thresholds and Suppress indications.

Each alarm can be configured with Minor, Major, and Critical threshold percentages.

The fields are described in Alarm Settings elements.

On the **Policy and Charging**, and then **Configuration**, and then **Alarm Settings** screen on the NOAM, you can change the Alarm Thresholds and the Suppress indications for the alarms:

- DSR Application Ingress Message Rate
 The DSR Application Ingress Message Rate alarm is raised when the average
 Policy and Charging ingress messages rate exceeds the configured Alarm
 Threshold. The thresholds are based on the engineered system value for Ingress
 Message Capacity.
- SBR Sessions Threshold Exceeded
 The SBR Sessions Threshold Exceeded alarm percent full is based on the number of Session records compared to an engineered maximum that varies according to the number of session SBR Server Groups per mated pair.
 - The SBR Sessions Threshold Exceeded alarm is raised when number of concurrent policy and Charging sessions exceeds the configured threshold.
- SBR Bindings Threshold Exceeded
 The SBR Bindings Threshold Exceeded alarm measures the number of IMSI Anchor Key records against an engineered maximum value that varies according to the number of binding SBR Server Groups.



The Policy SBR Bindings Threshold Exceeded alarm works similarly to the session capacity alarm except that the scope of the binding capacity alarm is networkwide.

Alarm Settings elements

Table 5-38 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Alarm Settings** screen. 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 screens; the View screen is read-only.

The screen contains three sets of input fields for the alarms:

- DSR Application Ingress Message Rate
- SBR Sessions Threshold Exceeded
- SBR Bindings Threshold Exceeded

The element labels are the same for each input field set, but some serve different purposes and have different values. These distinctions are noted in the table.

Table 5-38 Alarm Settings Elements

Elements (* indicates required field)	Description	Data Input Notes
DSR Application Ingress Mess	sage Rate	
* Alarm Name	This alarm is raised when average Policy and Charging ingress messages rate exceeds the configured threshold. The thresholds are based on the engineered system value for Ingress Message Capacity.	Format: Non-editable text box Range: DSR Application Ingress Message Rate
* Critical Alarm Threshold (Percent)	The Policy and Charging ingress message rate threshold for this alarm to be raised as Critical. The threshold is a percentage of the Ingress Capacity Capability.	Format: Text box Range: 100-200 Default: 160
Suppress Critical	Controls whether this alarm is raised as Critical.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)
* Major Alarm Threshold (Percent)	The Policy and Charging ingress message rate threshold for this alarm to be raised as Major. The threshold is a percentage of the Ingress Capacity Capability.	Format: Text box Range: 100-200 Default: 140
Suppress Major	Controls whether this alarm is raised as Major.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)



Table 5-38 (Cont.) Alarm Settings Elements

Elements (* indicates		
required field)	Description	Data Input Notes
* Minor Alarm Threshold	The Policy and Charging	Format: Text box
(Percent)	ingress message rate threshold for this alarm to be	Range: 100-200
	raised as Minor. The threshold is a percentage of the Ingress Capacity Capability.	Default: 110
Suppress Minor	Controls whether this alarm is	Format: Check box
	raised as Minor.	Range: Unchecked (No) or Checked (Yes)
		Default: Unchecked (No)
SBR Sessions Threshold Exc	ceeded	
* Alarm Name	This alarm is raised when	Format: Non-editable text box
	the number of concurrent Policy and Online Charging SBR sessions exceeds the configured threshold.	Range: Policy SBR Sessions Threshold Exceeded
* Critical Alarm Threshold	The concurrent sessions	Format: Text box
(Percent)	threshold for this alarm to be raised as Critical. The	Range: 1-99
	threshold is a percentage of the Maximum SBR Sessions.	Default: 95
Suppress Critical	Controls whether this alarm is	Format: Check box
	raised as Critical.	Range: Unchecked (No) or Checked (Yes)
		Default: Unchecked (No)
* Major Alarm Threshold (Percent)	The concurrent sessions	Format: Text box
	threshold for this alarm to	Range: 1-99
	be raised as Major. The threshold is a percentage of the Maximum SBR Sessions.	Default: 90
Suppress Major	Controls whether this alarm is	Format: Check box
	raised as Major.	Range: Unchecked (No) or Checked (Yes)
		Default: Unchecked (No)
* Minor Alarm Threshold	The concurrent sessions	Format: Text box
(Percent)	threshold for this alarm to	Range: 1-99
	be raised as Minor. The threshold is a percentage of the Maximum SBR Sessions.	Default: 80
Suppress Minor	Controls whether this alarm is	Format: Check box
	raised as Minor.	Range: Unchecked (No) or Checked (Yes)
		Default: Unchecked (No)
SBR Bindings Threshold Exc	eeded	
* Alarm Name	This alarm is raised when the	Format: Non-editable text box
	number of concurrent Policy SBR bindings exceeds the configured threshold.	Range: Policy SBR Bindings Threshold Exceeded



Table 5-38 (Cont.) Alarm Settings Elements

Elements (* indicates required field)	Description	Data Input Notes
* 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
Suppress Minor	Controls whether this alarm is raised as Minor.	Format: Check box Range: Unchecked (No) or Checked (Yes) Default: Unchecked (No)

Defining Alarm Settings

Use this task to define Alarm Settings on an Active NOAM.



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. Click Policy and Charging, and then Configuration, and then Alarm Settings.
- 2. Enter values in the editable fields to define the alarm settings.
- 3. Click:



- Apply to save the changes and remain on this screen.
- Cancel to discard the changes and remain on the Policy and Charging, and then Configuration, and then Alarm Settings screen.

If **Apply** is clicked and any of several 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 Congestion Options can be configured:

- Alarm Thresholds, which are used to:
 - Set the percentage of the Policy and Charging ingress message rate capacity at which an alarm is raised with Critical, Major, or Minor severity.
 - Set the percentage of the Policy and Charging ingress message rate capacity at which a Critical, Major, or Minor severity alarm is cleared.

The percentages control the onset and abatement of the corresponding Congestion Levels.

Default thresholds are based n the engineered system value for Ingress Policy and Charging Request Message Capacity.

 Message Throttling Rules, which determine the percentage of Session Creation, Update, and Terminate Request messages that are discarded when Congestion Levels 1, 2, and 3 exist.

The fields are described in Congestion Options elements.

Congestion Options elements

Table 5-39 describes the elements on the **Policy and Charging**, and then **Configuration**, and then **Congestion Options** screen. The elements can be configured and viewed on the NOAM.

The screen contains two sets of input fields:

- Alarm Thresholds
- Message Throttling Rules

Table 5-39 Congestion Options Elements

Fields (* indicates required		
field)	Description	Data Input Notes
Alarm Thresholds		



Table 5-39 (Cont.) Congestion Options Elements

Fields (* indicates required	D	Data Innut Nation
field)	Description	Data Input Notes
Alarm Name	Alarm is raised when average Policy and Charging ingress request messages rate exceeds the configured threshold. The thresholds are based on the engineered system value for Ingress Policy and Charging Request Message Capacity.	Format: Non-editable text box Range: Policy and Charging Server in Congestion
* Critical Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 3.	Format: Text box Range: 100-200 Default: 160
* Critical Alarm Abatement Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 3.	Format: Text box Range: 100-200 Default: 150
* Major Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 2.	Format: Text box Range: 100-200 Default: 140
* Major Alarm Abatement Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 2.	Format: Text box Range: 100-200 Default: 130
* Minor Alarm Onset Threshold	Percentage of Policy and Charging Ingress Request Message Rate capacity at which this alarm gets raised with Critical severity. This implies that the system is at Congestion Level 1.	Format: Text box Range: 100-200 Default: 110



Table 5-39 (Cont.) Congestion Options Elements

Fields (* indicates required			
field)	Description	Data Input Notes	
* Minor Alarm Abatement	Percentage of Policy and	Format: Text box	
Threshold	Charging Ingress Request Message Rate capacity at	Range: 100-200	
	which this alarm with Critical severity is cleared. This implies that the system has come out of Congestion Level 1.	Default: 100	
Message Throttling Rules			
Tabs for Congestion Level 1, C	Congestion Level 2, and Congesti	ion Level 3	
* Discard Session Creation	Percentage of Request	Format: Text box	
Requests	messages that result in new session creation, to	Range: 0-100	
	be discarded when this	Default:	
	congestion level exists.	Level 1 - 25	
		Level 2 - 50	
		Level 3 - 100	
* Discard Session Update	Percentage of Request messages that result in updating existing sessions, to be discarded when this	Format: Text box	
Requests		Range: 0-100	
		Default:	
	congestion level exists.	Level 1 - 0	
		Level 2 - 25	
		Level 3 - 50	
* Discard Session Terminate	Percentage of Request	Format: Text box	
Requests	messages that result in terminating existing sessions,	Range: 0-100	
	to be discarded when this	Default:	
	congestion level exists.	Level 1 - 0	
		Level 2 - 0	
		Level 3 - 0	

Setting Congestion Options

Use this task to set the Congestion Options on the Active NOAM:

- Alarm Thresholds for the Policy and Charging Server in Congestion onset and abatement alarm for Critical, Major, and Minor severities.
- Message Throttling Rules for discarding Session Creation, Update, and Terminate Requests for Congestion Levels 1, 2, and 3.
- 1. Click Policy and Charging, and then Configuration, and then Congestion Options.
- 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 screen to show the changes.
- Cancel to discard the changes and refresh the screen.

If **Apply** is clicked and any of several 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.

Configuration of Policy DRA Function on a Running DSR PCA System

Configuring new Policy DRA Sites

- Execute the procedures in the DSR Hardware and Software Installation Procedure and DSR Software Installation and Configuration Procedure to add new site(s) in the DSR network.
- 2. Configure the PCA Policy DRA function using the procedures in Policy DRA Configuration.

Configuring Policy DRA in existing Sites

 Configure the PCA Policy DRA function using the procedures in Policy DRA Configuration.

Configuring Policy DRA in existing Sites with scaling

 If the need arises to scale Policy DRA on a running PCA system, use the information from PCA Scalability, MP Growth, Mated Pair Growth, and Small System Support.

Configuration of Online Charging Function on a Running DSR PCA System



Configuring new Online Charging DRA Sites

- Execute the procedures in the DSR Hardware and Software Installation Procedure and DSR Software Installation and Configuration Procedure to add new site(s) in the DSR network.
- 2. Configure the PCA Online Charging DRA function using the procedures in Online Charging DRA Configuration.

Configuring Online Charging DRA in existing Sites

 Configure the PCA Online Charging DRA function using the procedures in Online Charging DRA Configuration.

Configuring Online Charging DRA in existing Sites with scaling

 If the need arises to scale Online Charging DRA on a running PCA system, use the information from PCA Scalability, MP Growth, Mated Pair Growth, and Small System Support.

Unconfiguration of Policy DRA Function from a Running DSR PCA System

Unconfiguring Policy DRA

This procedure unconfigures the Policy DRA function of the PCA application.

Disable the Policy DRA function.



Executing this step irretrievably deletes all the subscriber binding and Policy session records from the SBR Databases.

Note:

This procedure should be performed in a maintenance window. After clicking **Apply**, several instances (depending on the number of redundant SBR servers) of one or more alarms are expected. They are expected to clear in an interval of 15 minutes of less. The alarms that may arise are 31101 (DB Replication To Slave Failure), 19800 (Communication Agent Connection Down), and 31201 (Process Not Running).

 On the NOAM, navigate to Policy and Charging, and then Configuration, and then General Options.



- b. Uncheck the Policy DRA Enabled box.
- c. Click Apply.
- Disable the Policy DRA specific Binding SBR Database.
 - a. Navigate to SBR, and then Maintenance, and then SBR Database Status.
 - b. Select the desired Binding type SBR Database.
 - c. Click Disable.
- 3. Disable the Policy DRA Session Database.
 - **a.** If the Online Charging DRA function is not enabled, disable all the Session Database(s).
 - b. Navigate to SBR, and then Maintenance, and then SBR Database Status.
 - **c.** One by one, select the Session type SBR Database and click **Disable**.
- 4. Delete the Policy DRA specific Binding SBR Database.
 - a. Navigate to SBR, and then Maintenance, and then SBR Databases.
 - b. Select the desired Binding type SBR Database.
 - c. Click **Delete**.
- 5. Delete the Policy DRA Session SBR Databases.
 - a. If the Online Charging DRA function is not enabled, disable all the Session Database(s).
 - **b.** Navigate to **SBR**, and then **Maintenance**, and then **SBR Databases**.
 - c. Delete the Session type SBR Databases.
- 6. Delete the Policy DRA specific APNs.

Note:

This step is optional. This step can be skipped if you are going to enable Policy DRA again on this system and you want to reuse the APN configuration data after reenable.

- a. Navigate to **Poicy and Charging**, and then **Configuration**, and then **Access Point Names**.
- **b.** Delete the Policy DRA specific configuration data from this screen.
- 7. De-reference all the PRTs from PCRF Pools.

Note:

This step is optional. This step can be skipped if you are going to enable Policy DRA again on this system and you want to reuse the PCRF Pool configuration data after reenable.

 On the SOAM, navigate to Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pool to PRT Mapping.



- Edit all the PCRF Pool Name entries and set the Peer Route Table Name to Not Selected.
- 8. Delete all the PCRFs.

Note:

This step is optional. This step can be skipped if you are going to enable Policy DRA again on this system and you want to reuse the PCRF configuration data after reenable.

- Navigate to Policy and Charging, and then Configuration, and then Policy DRA, and then PCRFs.
- **b.** Delete the complete configuration data from this screen.
- 9. Delete all the Policy Clients configuration.
 - Navigate to Policy and Charging, and then Configuration, and then Policy DRA, and then Policy Clients.
 - **b.** Delete the complete configuration data from this screen.
- **10.** Unconfigure the Site Options.
 - Navigate to Policy and Charging, and then Configuration, and then Policy DRA, and then Site Options.
 - **b.** Delete the configuration data from this screen.

Note:

If no value for **Topology Hiding Virtual Name** is configured on this screen, navigate to **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **Network-Wide Options** on the NOAM to delete the **Default Topology Hiding Virtual Name** value.

11. Restore default values of Error Codes

Note:

This step is optional.

- Navigate to Policy and Charging, and then Configuration, and then Error Codes.
- b. Edit all Error Conditions and set the Result Code as 3002 for all Policy DRA application interfaces (Gx/Gxx, Rx, S9, Gx-Prime etc.).
- 12. Delete all the Sub-Pool Delection Rules
 - a. Only the NOAM, navigate to **Policy and Charging**, and then **Configuration**, and then **Policy DRA**, and then **PCRF Sub-Pool Selection Rules**.
 - **b.** Delete the complete configuration data from this screen.



13. Delete all the PCRF Pools

- Navigate to Policy and Charging, and then Configuration, and then Policy DRA, and then PCRF Pools.
- b. Delete the complete configuration data from this screen.

Unconfiguration of Online Charging DRA Function from a Running DSR PCA System

Unconfiguring Online Charging DRA

This procedure unconfigures the Online Charging DRA function of the PCA application.

- 1. Disable the Online Charging DRA function.
 - a. On the NOAM, navigate to Policy and Charging, and then Configuration, and then General Options.



This procedure should be performed in a maintenance window. After clicking **Apply**, several instances (depending on the number of redundant SBR servers) of one or more alarms are expected. They are expected to clear in an interval of 15 minutes of less. The alarms that may arise are 31101 (DB Replication To Slave Failure), 19800 (Communication Agent Connection Down), and 31201 (Process Not Running).

- b. Uncheck the Online Charging DRA Enabled box.
- c. Click Apply.
- 2. Disable the Online Charging DRA Session SBR Database.
 - a. If the Policy DRA function is not enabled, disable all the Session Database(s).
 - Navigate to SBR, and then Maintenance, and then SBR Database Status.
 - c. One by one, select the Session type SBR Database and click **Disable**.
- 3. Delete the Online Charging DRA Session SBR Databases.
 - If the Policy DRA function is not enabled, delete all the Session Database(s).
 - Navigate to SBR, and then Maintenance, and then SBR Databases.
 - c. Delete the Session type SBR Databases.
- 4. Delete the configured Realms.

Note:

This step is optional. This step can be skipped if you are going to enable Online Charging DRA again on this system and you want to reuse the Online Charging Realms configuration data after reenable.



- a. Navigate to Poicy and Charging, and then Configuration, and then Online Charging DRA, and then Realms.
- **b.** Delete the complete configuration data from this screen.
- Delete the Online Charging DRA specific APNs.

Note:

This step is optional. This step can be skipped if you are going to enable Online Charging DRA again on this system and you want to reuse the APN configuration data after reenable.

- Navigate to Poicy and Charging, and then Configuration, and then Access Point Names.
- b. Delete the Online Charging DRA specific configuration data from this screen.
- **6.** Delete all the Online Charging Servers..

Note:

This step is optional. This step can be skipped if you are going to enable Online Charging DRA again on this system and you want to reuse the OCS configuration data after reenable.

- a. On the SOAM, navigate to **Online Charging and Charging**, and then **Configuration**, and then **Online Charging DRA**, and then **OCSs**.
- **b.** Delete the complete configuration data from this screen.
- 7. Delete all the Online Charging Clients.
 - a. Navigate to Online Charging and Charging, and then Configuration, and then Online Charging DRA, and then CTFs.
 - **b.** Delete the complete configuration data from this screen.
- 8. Restore default values of Error Codes

Note:

This step is optional.

- a. Navigate to Online Charging and Charging, and then Configuration, and then Error Codes.
- b. Edit the SBR Error error condition and set the Gy/Ro Result Code as 5012.
- **c.** Edit the Session Not Found error condition and set the Gy/Ro Result Code as 5002.
- d. Edit all other Error Conditions and set the Gy/Ro Result Code as 3002.



Diameter Common Configuration for PCA

Diameter Common configuration must be done before PCA configuration can be performed.

Use the explanations and procedures in the Diameter Common configuration help and the *Diameter Common User's Guide* to complete the Diameter Common configuration, including the Diameter Common components needed for use with PCA.

SOAM Diameter Common Configuration

Diameter Common configuration for MP Profile assignment for PCA is done from the SOAM GUI **Main Menu: Diameter Common**, and then **MPs**, and then **Profile Assignments**.

Click **Diameter Common**, and then **MPs**, and then **Profile Assignments** and verify that the correct Session MP Profiles have been assigned for PCA DA-MPs. If assignments need to be made or changed:

- Use the Diameter Common, and then MPs, and then Profile Assignments screen to assign an MP Profile for each configured PCA DA-MP shown in the DA-MP list.
- From the pulldown list, select the MP Profile that is for the correct blade type and for a Session application (such as **G6 Session** or **G8 Session**).

Post-Configuration Activities

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



It is recommended to perform these procedures in a maintenance window.

- Enable the PCA application
- Enable SBR Databases
- Restart Process
- Enable Diameter Connections with Peer Nodes
- Perform Health Check

Enable the PCA Application

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

- 1. Click **Diameter**, and then **Maintenance**, and then **Applications**.
- Under DSR Application Name, select each PCA row.
 To select more than one row, press and hold Ctrl while you click each row.
- 3. Click Enable.



4. Verify the application status on the screen.

The Admin State, Operational Status, Operational Reason, and Congestion Level in each of the selected rows should change respectively to Enabled, Available, Normal, Normal.

Setting General Options

Use this task to set General Options on the NOAM.

The fields are described in General Options elements.

The general options can apply to the configuration of Policy and Charging:

- Policy DRA Enabled
- Online Charging DRA Enabled
- 1. Click Policy and Charging, and then Configuration, and then General Options.
- The Policy DRA Enabled check box allows the user to enable or disable Policy DRA.
- The Online Charging DRA Enabled check box allows the user to enable or disable Online Charging DRA.
- 4. Click:
 - Apply to save the changes and remain on this screen.
 - Cancel to discard changes and remain on the Policy and Charging, and then Configuration, and then General Options screen.

If **Apply** is clicked and the entered Default Stale Session Timeout value contains invalid characters, is out of the allowed range, or the field is empty, then an error message appears.

Enable SBR Databases

Refer to the SBR User's Guide for how to enable SBR Databases.

Restart Process

Use this task to restart the DSR and Policy and Charging SBR process

- 1. On the NOAM, navigate to **Status & Manage**, and then **Server**.
- Select the MP servers with a DSR (multi-active cluster) function that are or will be handling PCA traffic and all MP servers with a Policy and Charging SBR function.

Click Restart.



The function of an MP Server is the same as the function assigned to its Server Group on **Configuration**, and then **Server Groups**.



Note:

If the DSR system is processing traffic other than PCA, then do not restart all DA-MP servers simultaneously because doing so will cause a network-wide outage. Restart the DA-MP servers in a controlled order to minimize traffic loss.

Enable Connections

Use this task to enable one or more connections to Peer Nodes.

- At the Active SOAM, click Diameter, and then Maintenance, and then Connections.
- 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.

- Click Enable.
- 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.

Verify 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.

Perform Health Check

Use this task to perform a health check.

- 1. Verify SBR Database Status
 - a. On the NOAM, navigate to SBR, and then Maintenance, and then SBR Database Status.
 - **b.** Verify that the status for all the SBR Database rows have values for:
 - Administrative State = Enabled
 - Operational Status = Normal
 - Resource User Operational Reason = X of X available
 - Resource Provider Operational Reason = Y of Y available
- 2. Verify Policy and Charging SBR Status.



- a. On the NOAM, click SBR, and then Maintenance, and then 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 tabs.

The Resource HA Role of Standby applies if there is server level redundancy configured in the DSR system. The Resource HA Role of Spare applies if there is site level redundancy configured in the DSR system.

If all the verifications are successful, then proceed with signaling call flow execution. Otherwise, stop the process and contact My Oracle Support.

- 3. Verify there are no PCA alarms raised.
 - a. On the NOAM, navigate to Alarms & Events, and then View Active.
 - **b.** Verify that there are no Alarms raised with Product **PCA/SBR**.

If all the verifications are successful, then proceed with signaling call flow execution. Otherwise, stop the process and contact My Oracle Support.

Bulk Import and Export

The *Diameter Common User's Guide* describes the use and operation of Bulk Import and Export functions:

- Help, and then Diameter Common, and then Bulk Import
- Help, and then Diameter Common, and then Bulk Export

The Bulk Import and Export functions can be used to export Diameter, IPFE, and Application configuration data in CSV files to a location outside the system, and to import the files (usually edited) into the system where the Import function is executed.

Bulk Import

The Bulk Import operations use configuration data in ASCII Comma-Separated Values (CSV) files (.csv), to insert new data into, update existing data in, or delete existing data from the configuration data in the system.



Some configuration data can be imported only with the Update operation, and other data can be imported with Insert and Delete operations but not Update. Refer to the *Diameter Common User's Guide* or the **Diameter Common**, and then **Import** Help for valid Import operations.

Import CSV files can be created by using a Bulk Export operation, or can be manually created using a text editor.

Note:

The format of each Import CSV file record must be compatible with the configuration data in the release used to import the file. Across different release versions, column counts may not be compatible, and the import fails.



Files that are created using the Bulk Export operation can be exported either to the local Status & Manage File Management Directory (**Status & Manage**, and then **Files** screen), or to the local Export Server Directory.

CSV files that are in the local File Management area can be used for Bulk Import operations on the local system.

Files can be created manually using a text editor; the files must be uploaded to the File Management area of the local system before they can be used for Import operations on the local system.

Multiple Import operations can be performed:

- Insert new configuration data records that do not currently exist in the system
- Update existing configuration data in the system
- Delete existing configuration data from the system

Each Import operation creates a log file. If errors occur, a Failures CSV file is created that appears in the File Management area. Failures files can be downloaded, edited to correct the errors, and imported to successfully process the records that failed. Failures files that are unchanged for more than 14 days and log files that are older than 14 days are automatically deleted from the File Management area.

Bulk Export

The Bulk Export operation creates ASCII Comma-Separated Values (CSV) files (.csv) containing Diameter, IPFE, and Application configuration data. Exported configuration data can be edited and used with the Bulk Import operations to change the configuration data in the local system without the use of GUI screens. The exported files can be transferred to and used to configure another system.

Each exported CSV file contains one or more records for the configuration data that was selected for the Export operation. The selected configuration data can be exported once immediately, or exports can be scheduled to periodically occur automatically at configured times.

Configuration data can be exported in one Export operation:

- All exportable configuration data in the system
- All exportable configuration data from the selected Application, IPFE, or Diameter (each component's data is in a separate file)
- Exportable configuration data from a selected configuration component for the selected Application, IPFE, or Diameter

Exported files can be written to the File Management Directory in the local File Management area (**Status & Manage**, and then **Files** screen), 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 are not logged.



6

Maintenance

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

- Maintenance and status information that is maintained by the Policy and Charging Configuration and Maintenance components and displayed on the Policy and Charging, and then Maintenance screens.
- Maintenance and status data that is maintained by Diameter for Diameter Configuration components, DSR Applications, and DA-MPs and displayed on the Diameter, and then Maintenance screens.
- Descriptions of Policy and Charging alarms, KPIs, and measurements
- Auditing of the Session and Binding databases
- Policy and Charging overload management
- Database Backup and Restore of Policy and Charging configuration data

Introduction

This chapter describes:

 Policy and Charging Maintenance describes maintenance and status data that is maintained by the Policy and Charging application and by Policy and Charging DA-MPs.

On the **Policy and Charging**, and then **Maintenance** screens, the user can:

- Define and execute a Policy Database 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**, and then **Maintenance** screens include status information for:

- Peer Nodes
- Connections
- DSR Applications (including Policy and Charging)
- DA-MPs
- Alarms, KPIs, and Measurements describes Policy and Charging-specific database alarms, and indicates the location of descriptions of PCA and SBR alarms, KPIs, and measurements.
- PCA Data Auditing describes the auditing of the Session and Binding databases.
- Overload Management describes overload controls and load shedding and for PCA and SBR.



 Backup and Restore for Policy and Charging Configuration Data describes the OAM database backup and restore of Policy and Charging configuration data.

Policy and Charging Maintenance

The **Policy and Charging**, and then **Maintenance** screen on the NOAM provides access to the Policy Database Query tool.

Policy Database Query

Use the **Policy and Charging**, and then **Maintenance**, and then **Policy Database Query** screen 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 gueried binding key does not exist, an error message is displayed...



The Policy Database Query tool can be used only with Gx sessions. It is not applicable to Rx sessions.

The fields are described in Policy Database Query elements.

To use the Policy Database Query tool,

- On the Active NOAM, click Policy and Charging, and then Maintenance, and then Policy Database Query.
- 2. Select the **Binding Key Type** in the pulldown list.
- 3. Enter the **Binding Key** value to search for.
- 4. Click Search.

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

Policy Database Query elements

Table 6-1 describes the elements on the **Policy and Charging**, and then **Maintenance**, and then **Policy Database Query** screen.



Table 6-1 Policy Database Query Elements

Elements (* indicates a required field)	Description	Data Input Notes
* Binding Key Type	Select the type of binding key data entered in the Binding	Format: Pulldown list Range: IMSI, MSISDN, IPv4
	Key field.	Address, IPv6 Address Default: N/A
* Binding Key	Enter the binding key string to search for. Note: If Binding Key Type field is set toSelect, 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:23 c::

Alarms, KPIs, and Measurements

This section describes the type of alarm, KPI, and measurements information that is available for the Policy and Charging Application's combination of Policy DRA, Online Charging DRA, and SBR, as well as how to access the information in the DSR GUI.

Policy and Charging Alarms and Events

The Policy and Charging Application alarms and events are described in the *Alarms* and *KPIs Reference* and the DSR online help for alarms and events.

Active alarms and events and alarm and event history can be displayed on the **Alarms & Events**, and then **View Active** and **Alarms & Events**, and then **View History** screens.

PCA KPIs

Key Performance Indicators, or KPIs, provide a means to convey performance information to the user in near real-time. All the KPIs for the Policy and Charging Application are displayed on the **Status & Manage**, and then **KPIs** screen. Selecting the tab for a server and a label under the tab displays the KPI information for the selected server.



The PCA KPIs are described in the *Alarms and KPIs Reference* and the DSR Alarms and KPIs online help.

Policy and Charging Measurements

Measurements for the Policy and Charging Application 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**, and then **Reports** screen.

The *Measurements Reference* and online help explain the report selection criteria, and describe each measurement in each measurement group.

Overload Management

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

Overload Controls

The SBRs that implement the Session and Binding databases must protect themselves from becoming so overloaded that they cease to perform their function. There are two parts to achieving this goal:

- Detecting the overload condition and severity
- Shedding work to reduce load.

Overload Control in PCA

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

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

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

Three PCA congestion levels (CL1, CL2 and CL3) are defined, each of them is associated with onset and abatement threshold values. The onset and abatement



values are configurable (see Congestion Options). When PCA is in congestion, a PCA congestion alarm will be raised at the severity (Minor, Major or Critical) corresponding to the congestion level (CL1, CL2 or CL3).

When congestion is detected, DA-MP overload control throttles a portion of incoming messages to keep PCA from being severely impacted. The type and percentage of the messages to be throttled are configurable through the PCA GUI as displayed in Table 6-2:

Table 6-2 PCA Default Overload Control Thresholds

	Alarm ID 22721				
PCA Operation al Status	Severit y	Onset Threshol d	Abatemen t Threshold	PCA Congestio n Level	PCA Message Throttling Rules
Available	N/A	N/A	N/A	CL0	No messages are discarded (Accept and process 100% Request and Answer messages)
Available	Minor	110%	100%	CL1	 Discard 25% of requests for creating new sessions Discard 0% of requests for updating existing sessions Discard 0% of requests for terminating existing sessions Discard 0% of answer messages
Available	Major	140%	130%	CL2	 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
Degraded	Critical	160%	150%	CL3	 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

The PCA's internal congestion state contributes to PCA's Operational Status directly, along with its Admin state and Shutdown state. Consequently, the congestion state of the PCA impacts the Diameter Routing Function message transferring decision. Depending on the PCA's Operational Status (Unavailable, Degraded, Available), the Diameter Routing Function forwards all the ingress messages to the PCA when the PCA's Operational Status is Available, or discard some or all of the ingress messages when the Operational Status is Degraded or Unavailable. Table 6-3 describes the Diameter Routing Function handling of the messages to the PCA.



Table 6-3 Diameter Routing Function Message Handling Based on PCA Operational Status

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

PCA verifies if an ingress message has the priority or not. Priority messages are inviolable and are not discarded by the DA-MP overload control functionality, regardless of the congestion state of the PCA MP. Such messages are assigned with minimum inviolable priority.

PCA also assigns message priority to Gx Re-Authorization Request (RAR) messages that originate in PCA. The priority of PCA-generated RAR messages is determined by the intent of the RAR message, such as querying the status of a session or removing an existing session. PCA distinguishes between the two different types of RAR message by inclusion or exclusion of the Session-Release-Clause AVP in the generated RAR. If the Session-Release-Clause AVP is included, the RAR is intended to remove an existing session. Otherwise, the RAR is intended to query the status of a session. The priority for an RAR without the Session-Release-Clause AVP is set to a lower priority, while an RAR with the Session-Release Clause AVP is set to a higher priority.

The DA-MP overload control function's message priority detection checks the priority of an ingress message. If the message priority is greater than or equal to the minimum inviolable priority, the message is not throttled by the DA-MP overload control function, regardless of the congestion level of the PCA MP. However, if the message priority is smaller than the minimum inviolable priority, the DA-MP overload control function discards the message based on the congestion level thresholds shown in Table 6-2.

Overload Control in SBR

SBR relies on ComAgent for resource monitoring and overload control. The ComAgent Resource Monitoring and Overload Framework monitors local MP's resource utilizations, defines MP congestion based on one or multiple resource utilizations, communicates the MP congestion levels to Peers, and reports local MP congestion level to the local application (SBR).

Messages called stack events are used for communication to and from ComAgent.

ComAgent defines MP congestion levels based on a CPU utilization metric and ingress stack event rate (number of stack events received per second at local ComAgent), whichever is higher than the pre-defined congestion threshold, and broadcasts the MP congestion state to all its Peers. ComAgent provides APIs that the local SBR can call for receiving congestion level notifications.

SBR congestion is measured based on the SBR CPU utilization level. There are four SBR congestion levels: CL0 (normal), CL1 (Minor), CL2 (Major) and CL3 (Critical). There are related Onset and Abatement threshold values, and Abatement time delays.

The SBR congestion state (CPU utilization) is managed and controlled by the ComAgents on both PCA and SBR MPs based on the ComAgent MP Overload



Management Framework. Messages to a SBR from a PCA are handled based on the congestion state of the SBR. A SBR congestion alarm will be raised when MP congestion notification is received from ComAgent. The appropriate alarm severity information will be included in the notification. The alarm will be cleared if the congestion level is changed to Normal, also indicated in the notification from ComAgent.

To manage the overload situation on a SBR, all stack event messages are associated with pre-defined priorities. Before a stack event message is sent, its priority is compared with the congestion level of the SBR to which the stack event is sent. If the priority is higher than or equal to the SBR current congestion level, the message will be forwarded. Otherwise, it will be discarded.

Table 6-4 PCA-SBR Stack Event Priorities

Stack Event Category	Priority	Reasoning
Audit stack events	0	Audit get lowest priority in the presence of overload.
Response stack events	3	Responses get the highest priority since the request has already been made.
Remove stack events	0 (Audit) 3 (Call Processing)	If done for auditing, Remove gets lowest priority. If part of call processing, Remove gets highest priority because it is cleaning up data.
Update stack events	2	Falls under the category of in- session processing. Existing sessions/bindings are more important than new sessions/ bindings.
Find stack events	2	Falls under the category of in- session processing. Existing sessions/bindings are more important than new sessions/ bindings.
Create stack events	1	New sessions/bindings are lower priority than existing sessions, but higher priority than audit.
Query stack events	1	Query stack events are used for troubleshooting, so they are higher priority than audit, but still lower priority than most of the call processing stack events.
MITM RAR events	0 (Query) 3 (Terminate)	If used for query RAR, priority 0 is used. If used for terminate RAR, priority 3 is used.

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



Stack events that are triggered by Diameter messages with inviolable priorities have the highest priority among all the stack events to ensure the Diameter messages and are more favorably processed by SBR or PCA.

Four priority levels (P0, P1, P2, and P3) are used for the stack event priority setting. PCA determines if a stack event to be sent to an active SBR is a priority message. If it is, the stack event is assigned the highest priority (P3). Otherwise, the stack event's priority level is assigned based on the values shown in Table 6-4.

Load Shedding

After the SBR has determined that it is in overload (CL1-CL3), it informs ComAgent that its resources and sub-resources are in congestion. ComAgent then broadcasts this information to all of the resource users for the specified resources and sub-resources. The resource users now begin to shed load by sending only certain requests for database updates. The resource users determine which database requests to discard based on the current congestion level of the resource provider.

Database requests are delivered to SBRs using ComAgent stack events. Each stack event has a priority. The resource user software (on either DA-MPs or SBRs) sets the stack event priority for every Stack Event it sends, depending on the type of stack event and the circumstances under which the Stack Event is being used. For example, the same stack event may be used for signaling and for audit, but may have a different priority in each circumstance. The Stack Event priority is compared with the congestion level of the server that is the target of the stack event to determine whether stack event should be sent, as shown in Table 6-5.

Table 6-5 Stack Event Load Shedding

Congestion Level	Description
CLO	The resource provider is not congested. No load shedding occurs. Send all Stack Events.
CL1	Minor congestion. Auditing is suspended. Send all Stack Events not related to auditing.
CL2	Major congestion. No new bindings or sessions are created. Existing bindings and sessions are unaffected. Send only Stack Events related to existing sessions.
CL3	Critical congestion. Send only Stack Events already started and Stack Events that remove sessions or bindings.

Shutdown

DA-MP - The Policy and Charging Application running on DA-MPs supports the DSR Application Infrastructure graceful shutdown with 5 seconds grace period. This means that when PCA is Disabled (using the **Diameter**, and then **Maintenance**, and then **Applications** screen), the application will transition to the Degraded Operational Status for 5 seconds to allow in-flight messages to be processed without accepting any new Requests before transitioning to the Unavailable Operational Status. In the Unavailable status, neither Requests nor Answers are processed by the PCA.

SBR - Because SBR servers use the Active/Standby/Spare redundancy model, and ComAgent supports reliable transactions, there is no need for a graceful shutdown mode. Shutdown of a SBR server will cause a failover to another server in the same



Server Group. (The exception is if the Server Group only has one server, as might be the case in a small demo system.)

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

When the PCA is not processing requests (in Operational Status of Degraded or Unavailable), the Diameter Routing Function will attempt to route new Requests using the rules in the Peer Routing Tables. If the Request has no Destination-Host AVP, as would be the case for session-initiating Requests, the routing will fail and the Diameter Routing Function will respond with a 3002 DIAMETER_UNABLE_TO_DELIVER Answer.

When a Server is Stopped using the Stop function on the **Status & Manage**, and then **Server** screen, Diameter will terminate all Diameter connections by sending a DPR and waiting for the DPA. If all DPAs have not been received within 15 seconds, Diameter begins termination of its layers and queues. If Diameter is still not shut down after another 15 seconds, the process is abruptly terminated.

To properly shut down a PCA DA-MP server,

- Go to the **Diameter**, and then **Maintenance**, and then **Applications** screen and Disable the PCA application.
 The Operational Status of the application will transition to Unavailable
- Go to the Status & Manage, and then Server screen and Stop the Server's application processes.After 30 seconds maintenance can proceed as necessary.

Table 6-6 shows an example of the PCA Operational Status determination where the Shutdown mode is Graceful Shutdown. The Shut down and Shutting down in the Operational Reason column indicate the states where the (Graceful) shutdown process has been completed (Shut down) and is in progress (Shutting down) respectively. While the Graceful Shutdown is in progress, the PCA continues to process the messages in its queue for a time interval that is engineering configurable.

Table 6-6 PCA Operational Status

Admin State	Congestion Level	Shutdown State	Operational Status	Operational Reason
N/A	N/A	N/A	Unavailable	Not initialized
Disabled	0 ,1, 2, 3	False	Unavailable	Shut down
Disabled	0 ,1, 2, 3	True	Degraded	Shutting down
Enabled	0	N/A	Available	Normal
	1			Available with
	2			CL_1
				Available with CL_2
Enabled	3	N/A	Degraded	Congested with CL_3

SBR - Because SBR servers use the Active/Standby/Spare redundancy model, and ComAgent supports reliable transactions, there is no need for a graceful shutdown mode. Shutdown of a SBR server will cause a failover to another server in the same



Server Group. (The exception is if the Server Group only has one server, as might be the case in a small demo system.)

Diameter Maintenance and Status Data for Components, DSR Applications, and DA-MPs

Maintenance and status data is maintained and displayed on the **Diameter**, and then **Maintenance** screens for Diameter Configuration components, DSR Applications including Policy and Charging, and DA-MPs including those that run the Policy and Charging application:

- Route Lists Maintenance The Diameter, and then Maintenance, and then Route Lists screen 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, and then Maintenance, and then Route Groups screen displays the configured and available capacity for Route Groups and displays information about Peer Nodes or Connections assigned to a Route Group.

This information can be used to determine if changes need to be made to the Peer Node or Connection assignments in a Route Group in order to better facilitate Diameter message routing. Additionally, this information is useful for troubleshooting alarms.

Note:

Policy and Charging will create and add one metadata record to the TTR for each event that occurs while any Diameter message in the transaction is being processed. This function is achieved through Policy and Charging support of IDIH.

- Peer Nodes Maintenance The Diameter, and then Maintenance, and then Peer Nodes screen provides the Operational Status of Peer Node connections, including a Reason for the status.
- Connections Maintenance The Diameter, and then Maintenance, and then Connections screen displays information about existing connections, including the Operational Status of each connection.
 - The **Diameter**, and then **Maintenance**, and then **SCTP Statistics** screen displays statistics about paths within an SCTP connection. Each line on the screen represents a path within an SCTP connection.
- Applications Maintenance The Diameter, and then Maintenance, and then Applications screen displays status, state, and congestion information about activated DSR Applications. The data is refreshed every 10 seconds.
 On the Diameter, and then Maintenance, and then Applications screen, you can change the Admin State of the selected DSR Application to Enabled or Disabled.



DA-MPs Maintenance – The Diameter, and then Maintenance, and then DA-MPs screen provides state and congestion information about Diameter Agent Message Processors.

On the **Diameter**, and then **Maintenance**, and then **DA-MPs** screen,

- 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**, and then **Reports**, and then **MP (SCTP) Reports** screen displays the Message Processor (MP) SCTP statistics per MP, for all MPs or for a selected set of MPs. Each row shows the statistics for one MP.

Diameter Maintenance is described in more detail in the *Diameter User Guide* and in the Diameter Help.

Backup and Restore for Policy and Charging Configuration Data

Because the Policy and Charging Application is required to run on a 3-tier OAM topology where some data is mastered at the NOAM and some data is mastered at SOAMs at each site, backup and restore must be performed on the NOAM and on the SOAMs at each site.

Only configured data is backed up and restored. Dynamic data such as policy and policy charging sessions and policy bindings that is mastered on SBR MP servers is not backed up or restored.

The PCA feature uses the capabilities of the Backup and Restore functions provided by the OAM **Status & Manage**, and then **Database** screen, as described in the Database Backups and Restores chapter of the *DSR Administration Guide*.

