

Oracle® Communications Diameter Signaling Router Feature Guide



Release 9.1.0.1.0
G54512-01
April 2026

ORACLE®

Copyright © 2014, 2026, 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, software documentation, data (as defined in the Federal Acquisition Regulation), 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," "commercial computer software documentation," or "limited rights data" 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®, Java, MySQL, and NetSuite 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

1	Introduction to Diameter Signaling Router	
1.1	Diameter Routing Challenges	1
1.2	Diameter Signaling Router Solution	3
2	DSR Features and Functions	
2.1	Overview	1
2.2	DSR system architecture	1
2.2.1	Operations, Administration and Maintenance	3
2.2.2	Diameter Agent Message Processor (DA MP)	3
2.2.3	STP Message Processor	4
2.2.4	IP Front End	4
2.2.5	Session or Subscriber Binding Repository	4
2.2.6	Subscriber Data Server	5
2.2.7	Database Processor	5
2.2.8	Query Server	6
2.2.9	Integrated Diameter Intelligence Hub	6
2.3	DSR OAMP	6
2.3.1	Network Interfaces	7
2.3.2	Web-Based GUI	7
2.3.3	Operations and Provisioning	7
2.3.4	Maintenance	8
2.3.5	DSR Dashboard	15
2.3.6	Automatic Performance Data Export	17
2.3.7	Administration	18
2.3.7.1	Database Management	18
2.3.7.2	File Management	19
2.3.8	Security	19
2.3.9	Machine/Machine Interface	20
2.4	DSR Nodes	20
2.5	Diameter Core Routing	21
2.5.1	Extended Command Codes	24
2.5.2	Redirect Agent Support	24
2.6	Routing and Transaction Related Parameters in the DSR	25

2.6.1	Peer Routing Table	26
2.6.2	Application Routing Table	28
2.6.3	Routing Option Sets	28
2.6.4	Pending Answer Timer	29
2.6.5	Transport	30
2.6.6	Message Prioritization	32
2.6.7	Diameter Routing Message Priority	33
2.6.8	TLS / DTLS	33
2.6.9	Configurable Disable of CEx Peer IP Validation	33
2.6.10	Diameter Peer Discovery	34
2.6.11	Implicit Realm Routing	35
2.6.12	DNS Support	37
2.6.13	Signaling Firewall	37
2.6.14	Support Answer on Any Connection	38
2.6.15	Congestion Control	39
2.6.15.1	Per Connection Ingress MPS Control	40
2.6.15.2	MP Overload Control	43
2.6.15.3	Internal Resource Management	44
2.6.15.4	Egress Transport Congestion	45
2.6.15.5	Per Connection Egress MPS Control	45
2.6.15.6	Egress Throttle Group (ETG) Limiting	46
2.6.15.7	Coordinated Egress Throttling Across Multiple DSRs	49
2.6.15.8	Remote Busy Congestion	50
2.6.15.9	Remote Transport Congestion Control	53
2.6.15.10	Diameter Overload Indication Conveyance	53
2.7	Next Generation Network Priority Service	72
2.8	IP Front End	73
2.8.1	Traffic Distribution	73
2.8.2	High availability	75
2.9	RADIUS Signaling Router	75
2.9.1	RADIUS Routing	75
2.9.2	RADIUS Overload Control	76
2.9.3	RADIUS Message Format	76
2.9.4	Authenticator	77
2.9.5	Message Authenticator	77
2.9.6	Connections and Peers	77
2.9.7	Routing and Load-balancing	78
2.9.8	Duplicate Detection	79
2.9.9	Message / Traffic Control	79
2.10	RADIUS-Diameter IWF for Authentication	79
2.11	Diameter Mediation	81
2.11.1	Rule Templates and Rules	81

2.11.2	States of a Rule Template	83
2.11.3	Trigger Points	83
2.11.4	Measurements Associated with Rule	84
2.11.5	AVP Dictionaries	84
2.12	Topology Hiding	84
2.12.1	S6a/S6d Topology Hiding	85
2.12.2	Path Topology Hiding	90
2.12.3	S9 PCRF Topology Hiding	93
2.12.4	S9 AF/pCSCF Topology Hiding	94
2.13	DSR Applications	94
2.13.1	Range Based Address Resolution (RBAR)	95
2.13.2	Full Address Based Resolution	95
2.13.3	Policy and Charging Application (PCA)	100
2.13.4	Gateway Location Application (GLA)	110
2.13.5	Diameter Security Application (DSA)	112
2.13.6	DSA – Cross Protocol Security	114
2.13.7	DSA – Common Visualization Framework	115
2.13.8	Service Based Interface Support	115
2.14	Diameter Message Copy	115
2.15	Virtualized network functions manager VNF	116
2.16	Custom Application Framework (CAF)	117
2.16.1	CAF Application Life Cycle Management	117
2.16.2	Reference CAF Applications	118
2.17	Security Assertion Markup Language (SAML)	119
2.17.1	SAML Authentication Flow	119
2.17.2	SAML Feature Description	120
2.17.3	Enabling SAML Authentication functionality	121
2.17.4	Disabling SAML Authentication functionality	121
2.17.5	Viewing SAML Authentication functionality	122
2.17.6	Uploading IDP Metadata	122
2.17.7	Deleting IDP Metadata	123
2.17.8	DSR/SDS Metafile	123
2.17.9	Measurement Data Streaming (MDS)	123
2.17.9.1	Troubleshooting	129
2.17.9.2	Limitations	129
2.18	Signaling Transfer Point (STP) Virtual Network Function (VNF)	129
2.18.1	General	129
2.18.2	Signaling Protocols	130
2.18.3	SCCP -GLOBAL TITLE TRANSLATIONS (GTT) -ANSI/ITU	131
2.18.4	Equipment Identity Register	134
2.18.5	FLOBR (Flexible Linkset Optional Based Routing)	136
2.18.6	MBR (MAP Based Routing)	137

2.18.7	TOBR (TCAP Opcode Based Routing)	137
2.18.8	IDP Query	138
2.18.9	Loop Detection	138
2.18.10	Multiple PC	139
2.18.11	MTP Screening	140
2.18.12	ANSI ITU Conversion	140

A Appendix A: Supported Diameter Interfaces

What's New in This Guide

Release 9.1.0.1.0 - G54512-01, April 2026

Updated the maximum value of the Application Routing Rules from 50000 to 65000 in the [Application Routing Table](#) section.

Reference

For more information, see *Diameter Signaling Router Operations, Administration and Maintenance Guide (OAM) Guide*.


1

Introduction to Diameter Signaling Router

Mobile data traffic is growing exponentially, fueled by the introduction of smart phones, laptop dongles, flat-rate plans, social networking, and applications like mobile video. Operators are looking at Internet protocol (IP) networks such as Long Term Evolution (LTE) and IP multimedia subsystem (IMS) to provide the bandwidth required to support high data traffic requirements and applications. The CSPs need cost-effective solutions to address the growing gap between traffic and revenue growth.

The 3GPP Evolved Packet core (EPC) and IP Multimedia Subsystem (IMS) network architectures have specified the use of Diameter over Stream Control Transmission Protocol (SCTP) or Transmission Control Protocol (TCP) for many network interfaces such as for policy, charging, authentication, and mobility management. Many of interfaces are illustrated in the figure below. 3GPP and ETSI defines the Diameter protocol, it is the foundation Authentication, Authorization, and Accounting (AAA) functions in the Next Generation Network (NGN).

Figure 1-1 Selected Diameter Interfaces in LTE and IMS



1.1 Diameter Routing Challenges

For years operators have employed signaling system 7 (SS7) as the international, standardized protocol to communicate globally between operator networks. In LTE and IMS networks, many of the functions performed by SS7-based signaling in current networks are replaced by equivalent functions based on the Diameter protocol. Operators expect the same network behavior and robustness currently provided by SS7 networks.

Without a separate Diameter signaling infrastructure at the network core to facilitate signaling between network elements, endpoints such as Mobility Management Entities (MMEs) and Home Subscriber Servers (HSSs) must utilize direct signaling connections to each other, forming a mesh-like network architecture. Network endpoints must handle all session related tasks such as routing, traffic management, redundancy and service implementation. Implementing an IMS or LTE network without a signaling framework may be sufficient initially, but as traffic levels grow, the lack of a capable signaling infrastructure poses a number of challenges:

- **Scalability and load balancing:** Each endpoint must maintain a separate SCTP association or TCP connection with each of its Diameter peers and keep track of the status of each association. This network arrangement increases the overheads on the endpoints as the number of nodes increases, and the endpoints have the additional responsibility of load balancing. This burden is made more complex with the responsibility of load balancing placed on each end point.
 - **Congestion control:** Diameter lacks well-defined congestion control mechanisms found in other protocols such as SS7. For example, if an HSS has multiple Diameter front ends, the lack of sufficient congestion control increases the risk of a cascading HSS failure.
 - **Secure Network interconnect:** A fully meshed network is completely unworkable when dealing with connections to other networks because there is no central interconnect point, which also exposes the operator's network topology to other operators and can lead to security breaches.
 - **Interoperability:** Protocol inter working becomes unmanageable as the number of devices supplied by multiple vendors increases. With no separate signaling or session framework, interoperability testing (IOT) must be performed at every existing node when a new node or software load is placed in service. IOT activities consume a considerable amount of operator time and resources, with costs increasing in proportion to the number of tests that must be performed.
 - **Support for legacy EIR:** A need for MAP to Diameter inter working is required as transitions are made and LTEs quickly introduced into a network while still needing to support legacy HLRs.
 - **Support for both SCTP and TCP implementations:** SCTP elements cannot communicate with TCP elements. Without a central conversion element, operators will either have to upgrade TCP elements or require all elements in the network to support both stacks.
 - **Subscriber to HSS mapping:** When there are multiple HSS in the network, subscribers may be homed on different HSS. Therefore, there must be some function in the network that maps subscriber identities to HSSs. With no separate Diameter signaling infrastructure, that task must be handled by a standalone Subscription Locator Function (SLF), or by the HSS itself. Either approach wastes MME (or call session control function[CSCF]) processing and can add unnecessary delays. The HSS approach wastes HSS resources and may even result in the need for more HSSs than would otherwise be necessary.
 - **Policy and charging rules function (PCRF) binding:** When multiple PCRFs are required in the network, there must be a way to ensure that all messages associated with a user's particular IP connectivity access network(IP-CAN) session are processed by the same PCRF. This requires an element in the network that maintains session binding dynamically.

In recognition of Diameter routing issues, 3GPP has defined the need for a Diameter signaling infrastructure and a Diameter border infrastructure as shown below which is taken from TR

29.909. In addition, the GSMA has specified the need for a Diameter Proxy Agent as shown below which is taken from PRD IR.88.

Figure 1-2 3GPP Inter/Intra-operator Diameter infrastructure





Figure 2 - 3GPP Inter/Intra-operator Diameter infrastructure

Figure 1-3 GSMA roaming implementation architecture



1.2 Diameter Signaling Router Solution

Oracle Communication's Diameter Signaling Router (DSR) creates a centralized core Diameter signaling layer that relieves LTE, IMS and 3G Diameter endpoints of routing, traffic management and load balancing tasks and provides a single interconnect point to other networks. Each endpoint only needs one connection to a DSR to gain access to all other Diameter destinations reachable by the DSR. This approach eliminates the Diameter or SCTP(or TCP) mesh that is created by having direct signaling connections between each network element. Having one or more connection hubs that centralize the Diameter traffic to all end nodes simplifies interoperability between different network elements and enhances network scalability.


Centralizing Diameter routing with a DSR creates a signaling architecture that reduces the cost and complexity of the core network and enables core networks to grow incrementally to support increasing service and traffic demands. It also facilitates network monitoring by providing a centralized vantage point in the signaling network.

Advantages of a centralized signaling architecture as listed below:

- Improves signaling performance and scalability by alleviating issues related to the limited signaling capacity of MMEs, HSSs, CSCFs and other Diameter endpoints.
- Provides a centralized point from which to implement load balancing.
- Simplifies network expansion because routing configuration changes for new endpoints are performed only on the DSR.
- Increases reliability by providing geographic redundancy.
- Provides mediation point for Diameter variants to support interoperability between multi-vendor endpoints.
- Creates a gateway to other networks to support roaming, security and topology hiding.
- Reduces provisioning, maintenance and IOT costs associated with adding new network nodes.
- Enables HSS routing flexibility by providing a central point to perform HSS address resolution.
- Creates a centralized monitoring and network intelligence data collection point to isolate problems and track Key Performance Indicators (KPIs).
- Provides network wide PCRF binding to ensure that all messages associated with a user's particular IP-CAN session are processed by the same PCRF.

The DSR can be deployed as a core router routing traffic between Diameter elements in the home network and as a gateway router routing traffic between Diameter elements in the visited network and the home network. Refer to the figure below for a representation of an operator's EPC or IMS core network with DSR.

Figure 1-4 Example of Operator's EPC/IMS Core network with DSR



The resulting architecture enables IP networks to grow incrementally and systematically to support increasing service and traffic demands. A centralized Diameter router is the ideal place to add other advanced network functionalities like network performance intelligence via centralized monitoring, address resolution, Diameter interworking and traffic steering.

2

DSR Features and Functions

Primary function of the DSR is as a Diameter relay per RFC 6733 to route Diameter traffic based on provisioned routing data. As a result, the DSR reduces the complexity and cost of maintaining a large number of SCTP connections in LTE, IMS, and 3G networks. It simplifies the Diameter network and streamlines the provisioning of Diameter interfaces. The DSR supports flexible traffic load sharing and redundancy schemes and offloads Diameter clients and servers from having to perform many of these tasks, thereby reducing cost and time to market and freeing up valuable resources in the end points. For a full list of all supported Diameter interfaces please see, [Appendix A: Supported Diameter Interfaces](#).

DSR network elements are deployed in geographically diverse mated pairs with each NE servicing signaling traffic to form a collection of Diameter clients, servers, and agents. The DSR Message Processor (MP) provides the Diameter message handling function and each DSR MP supports connections to all Diameter peers (defined as an element to which the DSR has a direct transport connection).

2.1 Overview


One primary function of the DSR is as a Diameter relay per RFC 6733 to route Diameter traffic based on provisioned routing data. As a result, the DSR reduces the complexity and cost of maintaining a large number of SCTP connections in LTE, IMS and 3G networks, simplifies the Diameter network and streamlines the provisioning of Diameter interfaces. The DSR supports flexible traffic load sharing and redundancy schemes and offloads Diameter clients and servers from having to perform many of these tasks, thereby reducing cost and time to market and freeing up valuable resources in the end points. For a full list of all supported Diameter interfaces please see Appendix A: Supported Diameter Interfaces.

DSR network elements are deployed in geographically diverse mated pairs with each NE servicing signaling traffic to/from a collection of Diameter clients, servers, and agents. The DSR Message Processor (MP) provides the Diameter message handling function and each DSR MP supports connections to all Diameter peers (defined as an element to which the DSR has a direct transport connection).

2.2 DSR system architecture

The following image shows an overview of a DSR system architecture. Only single elements are shown for simplicity.

Figure 2-1 DSR 7.x Architecture



The key components of the solution are:

- Operations, Administration, Maintenance, and Provisioning (OAMP)
 - System OAM per signaling node
 - Network OAMP
- Diameter Agent Message Processor (DA MP) (handles Diameter and Radius)
- SS7 Message Processor
- STP Message Processor
- IP Front End (IPFE)
- Session Binding Repository (SBR)
- Database Processor (DP) / Subscriber Data Server (SDS)
- Query Server (QS)
- Integrated Diameter Intelligence Hub (IDIH)
- User Data Repository (UDR NO)–Optional
- VNF Manager (VNFM)–Optional
- Service Proxy Function (SPF) and NRF (Network Repository Function)–Optional

List of non-compatible features in DSR 8.3:

- Diameter Intelligence Hub (DIH)
- Gateway Location Application (GLA)
- Radius

In rest of this document any reference to these non-complaint features shouldn't be considered for feature enabling.

These components are described at a higher level in the following subsections. Although each component plays a key role, the OAM and DA MP components are the mandatory components of the system.

For details on the licensing of the various DSR Features see, DSR Licensing Information User Manual, available on Oracle Help Center (OHC).

2.2.1 Operations, Administration and Maintenance

The Operations, Administration, Maintenance and Provisioning components of the DSR include the System OAM located at each signaling node and the Network OAMP (NOAMP).

Key characteristics of the Network OAMP are as follows:

- Centralized OAMP for the DSR network.
- Central location for network wide data configuration, like “topology hiding.”
- Supports SNMP northbound interface to operations support systems for fault management.
- Runs on a pair of servers in active/standby configuration or can be virtualized on the System OAM blades at one signaling site (for small systems with two DSR signaling nodes only).
- Optionally supports disaster recovery site for geographic redundancy.
- Provides configuration and management of topology data.
- Maintains event and security logs.
- Centralizes collection, access to measurements, and reports.
- Centralized view of key operational metrics which identifies potential operational issues.
- Centralized architecture for the configuration and management of geo-redundant state DBs for policy and charging proxy.

Key characteristics of the System OAM at each signaling node are as follows:

- Centralized OAM interface for the node.
- Provides mechanism to configure the diameter data (routing tables, mediation, so on).
- Maintains local copy of the configuration database.
- Supports SNMP northbound interface to operations support systems for fault management.
- Provides mechanism to create user groups with various access levels.
- Maintains event and security logs.
- Centralizes collection, access to measurements, and reports.
- Centralized view of key operational metrics which identifies potential operational issues.

2.2.2 Diameter Agent Message Processor (DA MP)

The DA MP hosts Proxy applications such as address resolution, policy, charging Application, charging Proxy and scales by adding blades or instances.

Key characteristics of a DA MP are as follows:

- Provides application specific handling of real-time Diameter and/or RADIUS messages.
- Accesses DPs for real-time version of the subscriber DB, as needed.
- Accesses session and subscriber binding from SBRs as needed.

- Interfaces with System OAM or IDIH.

2.2.3 STP Message Processor

The STP Message Processor provides the functionality of Signaling Transfer Point (STP).

Key characteristics of an STP MP are as follows:

- Supports M3UA in signaling gateway mode.
- Supports M2PA in client and server mode.
- Supports SCCP routing with enhanced GTT capabilities .
- Provides flow control at SCTP to manage traffic rates on each link.
- Interfaces with System OAM and supports configurations through RESTful MMI.

2.2.4 IP Front End

The DSR IP Front End provides TCP/SCTP connection based load balancing to hide the internal DSR hardware architecture and IP addresses from the customer network. The IPFE is typically deployed in sets of Active-Active pairs and it distributes connections to DA MPs. IPFE provides load balancing of connections to DA MPs. The connections are active/active with TSAs (Target Set Addresses) and they provide TCP and SCTP connectivity.

Key characteristics of an IPFE are as follows:

- Optional component of the DSR.
- Supports up to two active or standby pairs with 3.2 Gbps bandwidth per active/standby pair.
- Supported with SCTP Multi-homing.

2.2.5 Session or Subscriber Binding Repository

The SBR stores diameter sessions and subscriber bindings for stateful applications. The Policy Charging Application (PCA) supports Policy DRA (P-DRA) and Online Charging DRA (OC-DRA) functionalities. OC-DRA uses session database SBRs (SBR(s)) and policy DRA uses both session database SBRs (SBR(s)) and subscriber binding database SBR's (SBR(b)). Throughout this document the SBRs are referred to individually when there are significant differences discussed, and referred as SBR, without distinguishing the application, when the attribute applies to all types.

Key characteristics of an SBR are as follows:

- Optional component of the DSR.
- Provides repository for subscriber and session state data.
- Provides DSRs with network-wide access to bindings.
- Provides procedures for in-service augmentation of the DSR signaling node-to-session SBR database relationships.

A number of capabilities are available to allow the SBR to be reconfigured once deployed including:

- Binding SBR capacity Growth/Degrowth: Allows in-service growth and degrowth of the Binding SBR database capacity in an existing P-DRA deployment, to include augmenting the physical location of the Binding SBR servers.

- Session SBR Capacity Growth/Degrowth: Allows in-service growth and degrowth of the Session SBR database capacity in an existing P-DRA / OC-DRA deployment, to include augmenting the physical location of the Session SBR servers.
- SBR Data Migration of a Session SBR Database: Allows reconfiguring an SBR Database topology by moving data from one data base to another such as: Mating/Un-Mating/Re-Mating. SBR data migration plan is used to move from an initial SBD DB to a target SBR DB without affecting traffic.
- Per mated pair sizing of Session SBR: Supports independent sizing of the Session SBR databases in a P-DRA / OC-DRA network managed by a common DSR NOAM.
- P-DRA support for 2.1M network wide MPS on P-DRA: Provides world-class scaling of policy network traffic, supporting up to 2.1 M network wide MPS of P-DRA traffic, including network-wide stateful Gx/Rx correlation to support VoLTE.

2.2.6 Subscriber Data Server

The SDS provides a centralized provisioning system for distributed subscriber data repository. The SDS is a highly-scalable database with flexible schema.

Key characteristics of the SDS are as follows:

- Interfaces with provisioning systems to provision subscriber related data.
- Interfaces with DPs at each DSR network element.
- Replicates data to multiple sites.
- Stores and maintains the master copy of the subscriber database.
- Supports bulk import of subscriber data.
- Correlates records belonging to a single subscriber.
- Provides web based GUI for provisioning, configuration, and administration of the data.
- Supports SNMP v2c northbound interface to operations support systems for fault management.
- Provides mechanism to create user groups with various access levels.
- Provides continuous automated audit to maintain integrity of the database.
- Supports backup and restore of the subscriber database.
- Runs on a pair of servers in active / hot standby, and can provide geographic redundancy by deploying two SDS pairs at diverse locations.
- Disaster Recovery site capabilities.

2.2.7 Database Processor

The database processor is the repository of subscriber data on the individual DSR node elements. The database processor hosts the full address resolution database and scales by adding blades.

Key characteristics of a DP are as follows:

- provides high capacity real-time database query capability to DA MPs.
- Interfaces with DP-SOAM (application hosted on the same blades as the DSR SOAM) for provisioning of subscriber data and for measurements reporting across all DPs.
- Maintains synchronization of data across all database processor.

- Hosting other Oracle SDS based applications.

2.2.8 Query Server

The Query Server contains a replicated copy of the local SDS database and supports a northbound MySQL interface for free-form verification queries of the SDS Provisioning Database. The query server's northbound MySQL interface is accessible via its local server IP.

Key characteristics of the QS are as follows:

- optional component that contains a real-time, replicated instance of the subscriber DB.
- provides LDAP, XML and SQL access.


2.2.9 Integrated Diameter Intelligence Hub

Refer to *Integrated Diameter Intelligence Hub User Guide*.

2.3 DSR OAMP

The DSR has a 3-tiered topology as described in the following image:

Figure 2-2 DSR 3-tiered Topology Architecture



Key services provided by the OAM components include:

- Centralized operational interface.
- Distribution of provisioned and configuration data to all message processors in all sites.
- Event collection and administration from all message processors.
- User and access administration.
- Supports northbound SNMP interface towards an operator EMS/NMS.

- Supports a web based GUI for configuration.

The DSR MPs host the Diameter and RADIUS Signaling Router applications, process Diameter and RADIUS messages.

2.3.1 Network Interfaces

Three types of network interfaces are used in the DSR:

- XMI – External Management Interface: Interface to the operator's management network. XMI can be found on the OAM servers. All OAM&P functions are available to the user through the XMI.
- IMI – Internal Management Interface: DSRs internal management network interface. All DSR nodes have this interface and use the IMI for exchange of crucial internal data. The user does not have access to the internal management network.
- XSI – Signaling Interface: Interface to the operator's signaling network. Only the Message Processors (MPs) have this interface. The XSI is used exclusively by the application and is not used by OAM&P for any purpose.

2.3.2 Web-Based GUI

The DSR provides a web-based graphical user interface as the primary interface that administrators and operators use to configure and maintain the network. GUI access is user id and password protected.

2.3.3 Operations and Provisioning

Operations and Provisioning of the DSR can be accomplished through one of the ten GUI sessions that are made available to the user through internal web server(s). Through the GUI, the User is able to make all operations and provisioning changes to the DSR, including:

- Network Information (does not include switch configuration)
- Network Element
- Servers
- Routing and Configuration Databases
- Status and Manage for:
 - Network Elements
 - Servers
 - Replication
 - Collection
 - High Availability
 - Database
 - KPIs
 - Processes
 - Files

Network Information

The network information defines the network name, the layout or shape of the network elements and their components. It defines the interlinking and the intercommunicating of the components. The network information represents all server relationships within the application. The server relationships are then used to control data replication, data collection, and define high availability relationships. Switch configuration is not defined by the network information.

Network Elements

The DSR application is a collection of servers linked by standardized interfaces. Network Elements (NE) are containers that group and create relationships among servers in the network. A network element can contain multiple servers but a single server is part of only one network element. The DSR solution is comprised of a Network OAMP network element, at least one signaling node, and an optional database provisioning node (SDS).

2.3.4 Maintenance


The DSR provides the following maintenance capabilities:

- Alarms and Events
- Measurements
- Key Performance Indicators
- Bulk Import/Export

Alarms and Events

The platform and DSR software raise minor, major, critical alarms, and events for a wide variety of conditions. These are immediately sent to the OAM system also sent to the operator's network management system using SNMP. Alarm or event logs at the OAM are stored up to seven days. The OAM provides a dashboard view of all alarms on the downstream MPs. This information is maintained locally up to three days.

Figure 2-3 Flow of Alarms



Following are some of the alarms and events supported by DSR:

- Connection to peer failed/ restored
- Peer unavailable/available
- Connection to peer congested/not-congested
- Route list available/unavailable
- OAM server failed/ restored
- MP failed/ restored
- MP entered/exited/changed local congestion

A detailed list of all alarms supported in DSR can be found in Platform Feature Guide.

Key Performance Indicators

Key Performance Indicators (KPIs) allow the user to monitor system performance data, including CPU, memory, swap space, and uptime per server. This performance data is collected from all servers within the defined topology. Key Performance Indicators supported by the platform and DSR software are in the following tables.

Table 2-1 DSR KPI Summary

KPI Category	KPI Examples
Server Element KPIs	A group of KPIs that appear regardless of server role such as CPU and Network Element.
CAPM KPIs	Counters related to computer-aided policy making such as active templates and test templates.
Charging Proxy Application KPIs	KPIs related to the CPA feature such as CPA Answer Message Rate, CPA Ingress Message Rate, and cSBR Query Error Rate.
Communications Agent KPIs	KPIs related to the communication agent such as user data ingress message rate.
Connection Maintenance KPIs	KPIs pertaining to connection maintenance such as RxConnAvgMPS.
DIAM KPIs	Basic Diameter KPIs such as Avg Rsp time and ingress trans success rate.
IPFE KPIs	KPIs associated with IPFE such as CPU % and IPFE Mbytes/Sec.
MP KPIs	KPIs relating to the message processor such as Avg Diameter Process CPU Util and average routing message rate.
FABR KPIs	KPIs related to the full address based resolution feature such as Ingress Message Rate and DP Response Time Average.
RBAR KPIs	KPIs related to the Range Based Address Resolution feature such as Average Resolved Message Rate and Ingress Message Rate.
SBR KPIs	KPIs related to Session Binding Repository such as Current Session Bindings and Request Rate.

Table 2-2 Platform KPI Summary

KPI Name	KPI Description
System.CPU_UtilPct	Reflects current CPU usage, from 0-100%. (100% means all CPU Cores are completely busy).
System.RAM_UtilPct	Reflects the current committed RAM usage as a percentage of total physical RAM. Based on the Committed_AS measurement from Linux /proc/meminfo. This metric can exceed 100% if the kernel has committed more resources than provided by physical RAM, in which case swapping will occur.
System.Swap_UtilPct	Reflects the current usage of Swap space as a percentage of total configured swap space. This metric will be 0-100%.
System.Uptime_Srv	Length of time since the last server reboot.

A detailed list of all KPIs supported in DSR can be found in the *Platform Feature Guide* found on the Oracle Help Center (OHC).

Measurements

All components of the DSR solution measure the amount and type of messages sent and received. Measurement data collected from all components of the solution can be used for multiple purposes, including discerning traffic patterns, user behavior, traffic modeling, size traffic sensitive resources, and troubleshooting.

The measurements framework allows applications to define, update, and produce reports for various measurements:

- Measurements are ordinary counters that count occurrences of different events within the system, for example, the number of messages received. Measurement counters are also called pegs.
- Applications simply peg (increment) measurements upon the occurrence of the event that needs to be measured.
- Measurements are collected and merged at the OAM servers.
- The GUI allows reports to be generated from measurements.

A subset of the measurements supported in DSR are listed in the following table. A detailed list of all KPIs supported in DSR can be found in the *Platform Feature Guide* found on the Oracle Help Center (OHC).

Table 2-3 DSR Measurements

Measurement Category	Description
Application Routing Rules	A set of measurements associated with the usage of application routing rules. These allow the user to determine which application routing rules are most commonly used and the percentage of times that messages were successfully or unsuccessfully routed.
Charging Proxy Application (CPA) Performance	This group contains measurements that provide performance information that is specific to the CPA application.

Table 2-3 (Cont.) DSR Measurements

Measurement Category	Description
Charging Proxy Application Exception	These measurements provide information about exceptions and unexpected messages and events that are specific to the CPA application.
Charging Proxy Application Session DB	These measurements provide information about events that occur when the CPA queries the SBR.
Computer Aided Policy Making (CAPM)	A set of measurements containing usage-based measurements related to the Diameter Mediation feature.
Communication Agent Performance	This group is a set of measurements that provide performance information that is specific to the ComAgent protocol. They allow the user to determine how many messages are successfully forwarded and received to and from each DSR application.
Communication Agent Exception	This group is a set of measurements that provide information about exceptions and unexpected messages and events that are specific to the ComAgent protocol.
Connection Congestion	These measurements contain per-connection measurements related to Diameter connection congestion states.
Connection Exception	These measurements provide information about exceptions and unexpected messages and events for individual SCTP/TCP connections that are not specific to the Diameter protocol.
Connection Performance	This group contains measurements that provide performance information for individual SCTP/TCP connections that are not specific to the Diameter protocol.
DSR Application Exception	A set of measurements that provide information about exceptions and unexpected messages and events that are specific to the DSR protocol.
DSR Application Performance	A set of measurements that provide performance information that is specific to the DSR protocol. These allow the user to determine how many messages are successfully forwarded and received to and from each DSR application.
Diameter Egress Transaction	These are measurements providing information about Diameter peer-to-peer transactions forwarded to upstream peers.
Diameter Exception	A set of measurements that provide information about exceptions and unexpected messages and events that are specific to the Diameter protocol.
Diameter Ingress Transaction Exception	These measurements provide information about exceptions associate with the routing of Diameter transactions received from downstream peers.
Diameter Ingress Transaction Performance	A set of measurements providing information about the outcome of Diameter transactions received from downstream peers.
Diameter Performance	Measurements that provide performance information that is specific to the Diameter protocol.

Table 2-3 (Cont.) DSR Measurements

Measurement Category	Description
Diameter Rerouting	These measurements allow the user to evaluate the amount of message rerouting attempts which are occurring, the reasons for why message rerouting is occurring, and the success rate of message rerouting attempts.
Full Address Based Resolution (FABR) Application Performance	A set of measurements that provide performance information that is specific to the FABR feature. They allow the user to determine how many messages are successfully forwarded and received to and from the FABR application.
Full Address Based Resolution (FABR) Application Exception	A set of measurements that provide information about exceptions and unexpected messages and events that are specific to the FABR feature.
IP Front End (IPFE) Exception	This group is a set of measurements that provide information about exceptions and unexpected messages and events specific to the IPFE application.
IP Front End (IPFE) Performance	This group contains measurements that provide performance information that is specific to the IPFE application. Counts for various expected/normal messages and events are included in this group.
Message Copy	These measurements from the Diameter Application Server reflect the message copy performance. They allow the user to monitor the amount of traffic being copied and the percentage of times that messages were successfully or unsuccessfully copied.
Message Priority	This group contains measurements that provide information on message priority assigned to ingress Diameter messages.
Message Processor (MP) Performance	These measurements provide performance information for an MP server.
OAM Alarm	General measurements about the alarm system such as number of critical, major, and minor alarms.
OAM System	General measurements about the overall OAM system
Peer Node Performance	Measurements that provide performance information that is specific to a Peer Node. These measurements allow users to determine how many messages are successfully forwarded and received to/from each peer node.
Peer Routing Rules	These are measurements associated with the usage of peer routing rules. They allow the user to determine which peer routing rules are most commonly used and the percentage of times that messages were successfully or unsuccessfully routed using the route list.
Range Based Address Resolution (RBAR) Application Performance	A set of measurements that provide performance information that is specific to the RBAR application. They allow the user to determine how many messages are successfully forwarded and received to/from each RBAR application.

Table 2-3 (Cont.) DSR Measurements

Measurement Category	Description
Range Based Address Resolution (RBAR) Exception	A set of measurements that provide information about exceptions and unexpected messages and events that are specific to the RBAR feature
Route List	A set of measurements associated with the usage of route lists. They allow the user to determine which route lists are most commonly used and the percentage of times that messages were successfully or unsuccessfully routed using the route list.
Routing Usage	This report allows the user to evaluate how ingress request messages are being routed internally within the relay agent.
Session Binding Repository (SBR) Exception	A set of measurements that provide information about exceptions and unexpected messages and events specific to the SBR application.
Session Binding Repository (SBR) Performance	This group contains measurements that provide performance information that is specific to the SBR application. Counts for various expected/normal messages and events are included in this group.

Table 2-3 (Cont.) DSR Measurements

Measurement Category	Description
Interface Level Connection	This group contains the set of measurements that allows the user to determine the number of total, successful, and rejected messages forwarded and received to or from DSR over a connection on a diameter interface. Note - Interface Level Connection measurement group is not visible in the Measurements Report of SOAM GUI, hence it cannot be exported. - The measurement statistics of the measurement pegs present in this group can be collected through MDS feature. For more information, refer to <i>Oracle Communications Measurement Data Streaming Guide</i>. - The measurement statistics of these pegs will not be visible over Kafka server before the first initial traffic run. After traffic is initiated and indexes are created, pegs will be visible on Kafka screen.

Bulk Import/Export


DSR supports bulk import and export of provisioning and configuration data using comma separated values (csv) file format. The import and export operations can be initiated from the DSR GUI. The import operation supports insertion, updating, and deletion of provisioned data. Both the import and export operations generate log files.

2.3.5 DSR Dashboard

This GUI display is an operational tool allowing customers to easily identify the potential for or existence of a DSR Node or Diameter Network outage. This dashboard is accessible via the SOAM or NOAM GUI and provides the following high-level capabilities:

- Centralized view: Allows operators to view a high level summary of key operational metrics.
- Identifies potential operational issues: Assists operators in identifying problems through visual enhancements such as colorization and highlighting.
- Centralized Launch-Point: Allows operators to drill-down to the next level of status information to assist in pinpointing the source of a potential problem.

Figure 2-4 DSR Dashboard on the NOAM



The Dashboard is comprised of the following concepts and components:

Dashboard Metrics:

- Metrics are the core component of the DSR Dashboard. The operator can determine which Metrics can be viewed on their Dashboard display through configuration.
- Server metrics are maintained by each MP. Per-Server metric values are periodically pushed to their local SOAM which can be displayed on the SOAM Dashboard display.
- "Server Type" metrics allow the operator see to a roll-up of Server metrics by Server type. The formula for calculating a Server Type metric value is identical to that for calculating the per-NE metric for that metric.
- Network Element (NE) metrics are derived from per-server metrics. A "Network Element" is the set of servers managed by a SOAM. The formula for calculating a per-NE metric value is metric-specific although, in general, most NE metrics are the sum of the per-Server metrics.

- Per Network metrics are derived from per-NE summary metrics. A “Network” is the set of DSR NEs managed by a NOAM. The formula for calculating a Network metric value is identical to that for calculating the per-NE metric for that metric.

Metric Groups:

- A Metric Group allows the operator to physically group Metrics onto the Dashboard display and for creating an aggregation status for a group of metrics.
- The “status” of a Metric Group is the worst-case status of the metrics within that group.

Server Type:

- A Server Type physically groups Metrics associated with a particular type of Server (e.g., DA-MP) onto the Dashboard display and for creating summary metrics for Servers of a similar type.
- The following Server Types are supported: DA-MP, SS7-MP, IPFE, SBR, cSBR, SOAM.

Network Element (NE):

- A “Network Element” is a set of Servers which are managed by a SOAM.
- The set of servers which are managed by a SOAM is determined through standard NOAM configuration and cannot be modified via Dashboard configuration.
- A NOAM can manage up to 32 NEs.

Dashboard Network Element (NE):

- A “Dashboard Network Element” is a logical representation of a Network Element which can be assigned a set of Metrics, NE Metric Thresholds and Server Metric Thresholds via configuration that defines the content and thresholds of a SOAM Dashboard display.
- Up to 32 Dashboard NEs are supported.

Dashboard Network:

- A “Dashboard Network” is a set of Dashboard Network Elements, Metrics and associated Network Metric Thresholds that is created by configuration that defines the content and thresholds of a NOAM Dashboard display.
- The set of Dashboard Network Elements assigned to a Dashboard Network is determined from configuration.
- One Dashboard Network is supported.

Visualization Enhancements:

- Visualization enhancements such as coloring are used on the Dashboard to attract the operator’s attention to a potential problem.
- Visualization enhancements are enabled through metric thresholds.
- Visualization enhancements can be applied independently to Server Type, NE and Network summary metrics and Server metrics.
- Visualization enhancements are applied to Dashboard row and columns headers to ensure that any metric value which has exceeded a threshold but cannot be physically viewed on a single physical monitor is not totally hidden from the operator’s view.

Metric Thresholds:

- Metric thresholds allow the operator to enable visualization enhancements on the Dashboard.

- Up to three separate threshold values (e.g., thresh-1, thresh-2, thresh-3) can be assigned to each metric.
- Dashboard Network summary, Dashboard NE summary and Server metric thresholds are supported.
- Dashboard Network summary and Dashboard NE summary metric threshold values can be assigned by the operator.
- Metric thresholds are used for Dashboard visualization enhancements.
- Most (but not necessarily all) metrics have thresholds.
- Whether a Metric can be assigned thresholds is determined from configuration.

Dashboard GUI Display:

- The Dashboard GUI display allows an operator to view a set of metric values used for monitoring the status of a Network or NE.
- The NOAM Dashboard allows the operator to view both Network summary and NE summary metrics.
- The SOAM Dashboard allows the operator to view the NE's summary metrics, its per-Server Type summary metrics and its per-Server metrics.
- Metric values are displayed as text.
- Sets of Metrics associated with network components are displayed vertically on the Dashboard in network hierarchical order. For example, on the NOAM Dashboard, Network metrics are displayed first followed by per-NE metrics.
- Each column on the Dashboard contains the set of values for a particular Metric.
- The operator can control which metrics are displayed on the Dashboard via configuration.
- The order that Metric Groups are displayed on the Dashboard is determined from configuration.
- The order that Metrics are displayed within a Metric Group on the Dashboard display is determined from configuration.
- Metrics selected for display on the Dashboard via configuration are hidden/viewed via a Dashboard GUI control based on "threshold level" filters (for example, only display metrics having at least one value exceeding its threshold-3 value).

Drill-down via hyperlinks:

- A Dashboard provides high level metrics providing an overall view of the health of one or more Network Elements of the customer's network.
- When a visual enhancement on the Dashboard is enabled when a user-defined threshold is exceeded, the operator may want to investigate the potential problem by inspection of additional information.
- The Dashboard facilitates operator trouble-shooting via context-sensitive hyperlinks on the Dashboard to assist in viewing more detailed information via existing DSR status and maintenance screens.
- The linkage between content on the Dashboard to DSR status and maintenance screens is determined from configuration.

2.3.6 Automatic Performance Data Export

The Automatic Performance Data Export feature provides the following capabilities:

- Periodic generation and remote copy of filtered performance data.
- Proper management of the file space associated with the exported data.

Specifically, Automatic PDE provides the ability to create custom queries of performance data and to schedule periodic remote copy operations to export the performance data to remote export systems.

2.3.7 Administration

Administration functions are tasks that are supported at the system level. Administration functions of the DSR include:

- User Administration
- Passwords
- Group Administration
- User's Session Administration
- Authorized IPs
- System Level Options
- SNMP Administration
- ISO Administration
- Upgrade Administration
- Software Versions

For more details on platform related features see, *Platform Feature Guide*.

2.3.7.1 Database Management

Database Management for DSR provides 4 major functions:

- Database Status - maintains status information on each database image in the DSR network and makes the information accessible through the OAM server GUI.
- Backup and Restore - Backup function captures and preserves snapshot images of Configuration and Provisioning database tables. Restore function allows user to restore the preserved databases images. The DSR supports interface to and/or integration with 3rd party backup systems (that is Symantec NetBackup).
- Replication Control - allows the User to selectively enable and disable replication of Configuration and Provisioning data to servers.

Note

This function is provided for use during an upgrade and should be used by Oracle Personnel only.

- Provisioning Control - provides the User the ability to lockout Provisioning and Configuration updates to the database.

Note

This function is provided for use during an upgrade and should be used by Oracle Personnel only.

2.3.7.2 File Management

The File Management function includes a File Management Area, which is a designated storage area for any file the user requests the system to generate. The list of possible files includes, but is not limited to: database backups, alarms logs, measurement reports and security logs. The File Management function also provides secure access for file transfer on and off the servers. The easy-to-use web pages give the user the ability to export any file in the File Management Area off to an external element for long term storage. It also allows the user to import a file from an external element, such as an archived database backup image.

2.3.8 Security

Oracle addresses Product Security with a comprehensive strategy that covers the design, deployment, and support phases of the product life-cycle. Drawing from industry standards and security references, Oracle hardens the platform and application to minimize security risks. Security hardening includes minimizing the attack surface by removing or disabling unnecessary software modules and processes, restricting port usage, consistent use of secure protocols, and enforcement of strong authentication policies. Vulnerability management ensures that new application releases include recent security updates. In addition, a continuous tracking and assessment process identifies emerging vulnerabilities that may impact fielded systems. Security updates are delivered to the field as fully tested Maintenance Releases.

Networking topologies provide separation of signaling and administrative traffic to provide additional security. Firewalls can be established at each server with IP Table rules to establish White List and/or Black List access control. The DSR supports transporting Diameter messages over IPSec thereby ensuring data confidentiality and data integrity of Diameter messages traversing the DSR.

Oracle realizes the importance of having distinct interfaces at the Network-Network Interface layer. To maintain the separation of traffic between internal and external Diameter elements, the DSR supports separate network interfaces towards the internal and external traffic. The routing tables in DSR support the implementation of a Diameter Access Control List which make it possible to reject requests arriving from certain origin-hosts or origin-realms or for certain command codes.

Oracle recommends that Layer 2 and Layer 3 ACLs be implemented at the Border Gateway. However, Professional Services available from the Oracle Consulting team can implement Layer 2 and Layer 3 ACLs at the aggregation switch which serves as the demarcation point or at the individual MPs that serve the Diameter traffic.

In addition to supporting security at the transport and network layers, Oracle's solution provides Access Control Lists based on IP addresses to restrict user access to the database on IP interfaces used for querying the database. These interfaces support SSL.

DSR maintains a record of all system users' interactions in its Security Logs. Security Logs are maintained on OAM servers. Each OAM server is capable of storing up to seven days' worth of Security Logs. Log files can be exported to an external network device for long term storage. The security logs include:

- Successful logins

- Failed login attempts
- User actions (for example, configure a new OAM, initiate a backup, view alarm log).

Please see the *Diameter Signaling Router (DSR) 8.6.0.0.0 Security Guide* – Available at Oracle.com docIPSec

The DSR optionally supports IPSec encryption per Diameter connection or association. Use of IPSec reduces MPS throughput by up to 40%. IPSec is supported for SCTP over IPv6 connections. The DSR IPSec implementation is based on 3GPP TS 33.210 version 9.0.0 and supports the following:

- Encapsulating Security Payload (ESP).
- Internet Key Exchange (IKE) v1 and v2.
- Tunnel Mode (entire IP packet is encrypted and/or authenticated).
- Up to 100 tunnels.
- Encryption transforms/ciphers supported: ESP_3DES (default) and AES-CBC (128 bit key length).
- Authentication transform supported: ESP_HMAC_SHA-1.
- Configurable Security Policy Database with backup and restore capability.

2.3.9 Machine/Machine Interface

DSR REST MMI's provides Application Programming Interface allows EMS, OSS, or NMS systems at customer's network to interface directly with the DSR in order to access, store, change and delete OAM&P. Use of the MMI will allow real time changes in the DSR that is initiated by a configuration change in north-bound customer management systems.

Benefits of REST MMI's include:

- Industry moving towards automation of network operations.
- Provides a consistent interface for all OAM&P data.
- Advancement toward automated installations.

Administration, Configuration, Alarms & Events, Status & Manage, Measurements, Diameter, and IPFE managed objects of DSR can be managed using REST MMIs.

2.4 DSR Nodes

Each DSR message processor (MP) can host up to 48 Diameter Nodes (also called Diameter Identities). Hosting more than one node/identity allows a DSR deployment at the Network Edge where DSR acts as the single point of contact for all Diameter elements external to the operator network and similarly all internal Diameter elements use it as the point of contact when reaching Diameter servers external to the operator network. Another use case for hosting multiple Diameter nodes on each MP is to support multiple connections from an external Diameter element to the DSR.

Each Diameter Node has the following attributes:


- Diameter Realm that may be unique or shared across the nodes.
- Up to 128 local IP addresses - IPv4 or IPv6 addresses or a combination of IPv4 and IPv6 addresses. (Each DA-MP supports up to 8 local IP addresses and 32 DA-MPs are supported).

- A unique Fully Qualified Domain Name (FQDN).

DSR allows an IP address to be shared across nodes provided the combination of IP address, port and transport are unique across nodes.

See the following figure for a sample configuration:

Figure 2-5 Multiple Nodes per Message Processor



2.5 Diameter Core Routing

The DSR application provides a Diameter Routing Agent to forward messages to the appropriate destination based on information contained within the message including header information and applicable Attribute Value Pairs (AVP). As per the core Diameter specification, the DSR provides the capability to route Diameter messages based on any combination, or presence/absence, of Destination-Host, Destination-Realm, and Application-ID. In addition DSR optionally provides the capability to look at Command-Code and origination information, namely Origin-Realm and Origin-Host for advanced routing functionality. The average diameter message size supported is 2K bytes with a maximum message size of 60K bytes.

DSR high level message processing and routing is shown below. The numbers show the message flow through the system.

Figure 2-6 High Level Message Processing and Routing in DSR



DSR supports the following routing functions:

- Message routing to Diameter peers based upon user-defined message content rules.
- Message routing to Diameter peers based upon user-defined priorities and weights.
- Message routing to Diameter peers with multiple transport connections.
- Alternate routing on connection failures.

- Alternate routing on Answer timeouts.
- Alternate routing on user-defined Answer responses.
- Route management based on peer transport connection status changes.
- Route management based on OAM configuration changes.

Routing rules and rule actions are used to implement the routing behavior required by the operator. Routing rules are defined using combinations of the following data elements:

- Destination-Realm (leading, trailing characters, exact match, contains, not equal or always true).
- Destination-Host (leading, trailing characters, exact match, contains, always true, present and not equal, or presence/absence).
- Application-ID (exact match, not equal, or always true).
- Command-Code (exact match, not equal or always true).
- Origin-Realm (leading, trailing characters, exact match, contains, not equal or always true).
- Origin-Host (leading, trailing characters, exact match, contains, not equal or always true).

A set of configurable timers (100 – 180,000 milliseconds) control the length of time the DSR waits to receive an answer to an outstanding request. The maximum number of times a request can be rerouted upon connection failure or timeout is configurable from 0 – 4 retries.

DSR supports the concepts of routes, peer route tables, peer route groups, connection route groups, route lists, and peer node groups to provide a very powerful and flexible load balancing solution. A Route Group is comprised of a prioritized list of peers or connections used for routing messages. A route list is comprised of multiple route groups – only one of which is designated as active at any one time. Each route list supports the following configurable information:

- Route List ID.
- Up to five Route Groups with associated Route Group Priority level (1-5).
- Minimum Route Group Availability Weight to control which Route Group in the Route List is actively used for routing requests.
- 0-10 optional Traffic Throttle Groups with associated Max Loss % Threshold for use with IETF Diameter Overload Indicator Conveyance (DOIC) feature.

Each Route Group supports the following configurable information:

- Route Group ID.
- Up to 160 Peer IDs -OR- 512 Connection IDs.
- Weight (1-64K) for each Peer ID or Connection ID.

When peers or connections have the same priority level a weight is assigned to each peer/ connection which defines the weighted distribution of messages amongst the peers/ connections. For example, if two peers with equal priority have weights 100 and 150 respectively then 40% of the messages will be forward to peer-1 ($100/(100+150)$) and 60% of the messages will be forward to peer-2 ($150/(100+150)$).

Peer Route Tables can be assigned to Peer Nodes or Application IDs. Each Peer Route Table has its own set of Peer Route Rules.

A set of peers with equal priority within a Route List is called a “Peer Route Group”. Multiple connections to the same peer can be assigned to a Connection Route Group (CRG). The use of CRGs allows for prioritized routing between connections to the same peer. An example use

case would be connecting to Peers across different sites which share the same hostname. The peer within the site would be contacted for any traffic originated within the site and the remote peer should be contacted only if the local peer is unavailable.

Figure 2-7 Connection Route Group



When multiple Route Groups are assigned to a Route List, only one of the Route Groups is designated as the "Active Route Group" for routing messages for that Route List. The remaining Route Groups within the Route List are referred to as "Standby Route Groups". DSR designates the "Active Route Group" within each Route List based on the Route Group's priority and available capacity relative to the provisioned minimum capacity (described below) of the Route List. When the "Operational Status" of peers change or the configuration of either the Route List or Route Groups within the Route List change, then DSR may need to change the designated "Active Route Group" for the Route List. An example of Route List and Route Group relationships is shown below.

Figure 2-8 Route List, Route Group, Peer Relationship Example



Showing a different set of route lists and route groups, an example of peer routing based on route groups with a route list is shown in the figure below.

Figure 2-9 Load Balancing Based on Route Groups and Peer Weights



DSR supports provisioning up to 160 routes in a route group (same priority) and allows for provisioning of 3 route groups per route list.

To further enhance the load balancing scheme, the DSR allows the operator to provision a "minimum route list capacity" threshold for each route list. This provisioned "minimum route list capacity" is compared against the route group capacity. The route group capacity is dynamically computed based on the availability status of each route within the route group and is the sum of all the weights of "available" routes in a route group. If the route group capacity is higher than the threshold, the route group is considered "available" for routing messages. If the route group capacity is lower (due to one or more failures on certain routes in the route group), the route group is not considered "available" for routing messages. DSR uses the highest priority (lowest value) "available" route group within a route list when routing messages over the route list. If none of the route groups in the route list are "available", DSR will use the route group with the most "available" capacity, also honoring route group priority, when routing messages over the route list.

A peer node group is a configuration managed object that provides a container for a collection of DSR peer nodes with like attributes (Example: same network element or same capacity)

requirement). The user configures DSR peer nodes with their IP addresses in the peer node group container. Applications can use this IP address grouping for various functions such as IPFE for a distribution algorithm.

2.5.1 Extended Command Codes

Routing attributes by extended command code broadens the definition of a Diameter command code to include additional application specific single Diameter or 3GPP AVP content per command code. ECC are used for advanced routing selection and are comprised of the following attributes:

- ECC Name
- CC value
- AVP code value
- AVP data value

For example, there are four types of Credit-Control-Request (CCR) transactions which are uniquely identified by the content of the CCR's "CC-Request-Type" AVP: (For a complete list of ECCs please see the DSR Documentation set available at Oracle.com on the Oracle Technology Network (OTN).)


1. Initial_Request (typically called CCR-I).
2. Update_Request (typically called CCR-U).
3. Termination_Request (typically called CCR-T).
4. Event_Request (typically called CCR-E).

Extended command codes can be used in Routing Option Sets (ROS), Pending Answer Timer (PAT), and Message Priority Configuration Set (MPCS) (see Message).

2.5.2 Redirect Agent Support

The DSR supports the processing of notifications sent by a Redirect Agent. The DSR processes the redirect notification (DIAMETER_REDIRECT_INDICATION response) and continues routing the original request upstream using the Redirect-Host in the response (RFC7633). In addition, the DSR processes realm redirect notification and continues routing the original request upstream using the Redirect-Realm in the response (RFC7075). Finally, an optional re-evaluation of the application routing table and peer routing table is supported for routing the redirected request.

Figure 2-10 Redirect Agent



2.6 Routing and Transaction Related Parameters in the DSR

The DSR has a hierarchical configuration and selection criteria for routing and transaction related (ART, PRT, ROS and PAT) parameters. Customers can configure DSR and choose per ingress peer node scoped additional transaction-specific granularity in routing and transaction parameters selection process.

Customers can create Transaction Configuration Groups which are composed of Transaction Configuration Sets. The Transaction Configuration Sets are composed of individual Diameter Transactions (represented by Appl-id+Extended Command Codes) with each transaction optionally specifying an ART, PRT, ROS and PAT. Once a Transaction Configuration Group is associated with an ingress peer, any Requests from the peer that match a Transaction Configuration Set within the assigned Transaction Configuration Group uses the associated ART, PRT, ROS and PAT if specified. The following table provides the precedence order for routing and transaction related parameter selection.

Table 2-4 Modified Routing and Transaction Parameter Selection Precedence Order

Parameter Selection Criteria	Parameter Selection Precedence Order			
	ROS (Note 3)	PAT	ART (Note 1)	PRT (Note 2)
Ingress Peer Node Selected Transaction Configuration Group	1	1	1	1
Ingress Peer Node	2	2	2	2
Egress Peer Node	NA	3	NA	NA

Table 2-4 (Cont.) Modified Routing and Transaction Parameter Selection Precedence Order

Parameter Selection Criteria	Parameter Selection Precedence Order			
Default Transaction Configuration Group	3	4	4	3
System Default	4	5	4	4

Note

- For multiple DRA Application invocation on the same message, the applications can select a different ART and override the core routing ART precedence.
- Local DSR applications can select a different PRT and override this core routing PRT precedence
- Existing OAM configuration rule: A Routing Option Set with a configured Pending Answer Timer cannot be associated with an application-ID.

DSR supports configuring of up to 100 Transaction Configuration Groups, where each group instance can contain up to 1000 transaction configuration set entries. The maximum transaction set entries per DSR system cannot be greater than 1000.

2.6.1 Peer Routing Table

A peer route table is a set of prioritized peer routing rules that define routing to peer nodes based on message content. Peer routing rules are prioritized lists of user-configured rules that define where to route a message to upstream peer nodes. Routing is based on message content matching a peer routing rule's conditions. There are six peer routing rule parameters:

- Destination-Realm
- Destination-Host
- Application-ID
- Command-Code
- Origin-Realm
- Origin-Host

When a diameter message matches the condition of peer routing rules then the action specified for the rule occurs. If you choose to route the diameter message to a peer node, the message is sent to a peer node in the selected route list based on the route group priority and peer node configured capacity settings. If you choose to send an answer, then the message is not routed and the specified diameter answer code is returned to the sender.

Peer routing rules are assigned a priority in relation to other peer routing rules. A message is handled based on the highest priority routing rule that it matches. The lower the number a peer routing rule is assigned the higher priority it has. (1 is the highest priority and 1000 is the lowest priority.)

If a message does not match any of the peer routing rules and the destination-host parameter contains a Fully Qualified Domain Name (FQDN) matching a peer node, then the message is directly routed to that peer node if it has an available connection. If there is not an available connection, the message is routed using the alternate implicit route configured for the peer node.

PRT Partitioning

The DSR supports up to 500 PRTs on the DSR. Any one of the PRTs can be optionally associated with either the (ingress) peer or Ingress Peer Node selected Transaction Configuration Group or Default Transaction Configuration Group. A local application can also specify the PRT that needs to be used for routing a request. Each of these PRTs have no more than 2000 rules and the total number of rules across all PRTs cannot exceed 50,000. Which means, if one PRT has 2000 rules then DSR can have maximum 25 such PRTs, so that the total number of rules will remain within 50,000. Peer routing rules can be prioritized (1 – 2000) for cases where an inbound Diameter request may match multiple user-defined routing rules. A system wide PRT is also present by default and is used if a PRT has not been assigned.

The PRT can be associated with the ingress peer node which can be useful to separate routing tables for example for LTE domain, IMS domain, or routing partners.

Rule Action defines the action to perform when a routing rule is invoked. Actions supported are:

- Route to Peer - use Route List Table.
- Send Answer Response - an Answer response is sent with a configurable Result-Code and no further message processing occurs.
- Abandon With No Answer - discard the message and no Answer is sent to the originating Peer Node.

Forward to Peer Route Table - forward the message to the specified Peer Route Table.

The table below is used to determine the PRT instance to be used:

Table 2-5 PRT Precedence

PRT Used	PRT specified by local app (if supported)	PRT associated with Ingress Peer Node Selected Transaction Configuration Group	PRT associated with an Ingress Peer	PRT associated with Default Transaction Configuration Group	Default PRT
Default PRT	No	No	No	No	Yes
Default Transaction Configuration Group PRT	No	No	No	Yes	Yes
Peer PRT	No	No	Yes	Don't Care	Yes
PRT associated with Ingress Peer Node Selected Transaction Configuration Group	No	Yes	Don't Care	Don't Care	Yes
Local App PRT	Yes	Don't Care	Don't Care	Yes	Yes

2.6.2 Application Routing Table

An application route table contains one or more application routing rules that can be used for routing request messages to DSR applications. Up to 1,500 Application Route Tables can be configured per DSR network element, a total of upto 65,000 application routing rules across all ARTs can be configured across the application route tables per network element.

An application routing rule defines message routing to a DSR application based on message content matching the application routing rule's conditions. There are six application routing rule parameters:

- Destination-Realm
- Destination-Host
- Application-Id
- Command-Code
- Origin-Realm-
- Origin-Host

When a diameter message matches the conditions of an application routing rule the message is routed to the DSR application specified in the rule.

Rule Action defines the action to perform when a routing rule is invoked. Actions supported are:

- Route to Application - route the message to the local Application associated with this Rule.
- Forward to Egress Routing - ART search stops and moves on to PRT.
- Send Answer Response – ART generates an Answer. This Answer unwinds any previously encountered DSR Applications that want to process the Answer. Normal controls for Answer are given (Result-Code vs Experimental Result Code, Result-Code value, Vendor-ID, and ErrorMessage string).
- Forward to Application Route Table - ART forwards the request message to the specified Application Route Table.
- Forward to Peer Route Table - ART will forward the request message to the specified Peer Route Table.

Abandon With No Answer - discard the message and no Answer is sent to the originating Peer Node. Application routing rules are assigned a priority in relation to other application routing rules. A message is handled based on the highest priority routing rule that it matches. The lower the number an application routing rule is assigned the higher priority it has. (1 is highest priority and 1000 is lowest priority.)

One or more DSR applications must be activated before application routing rules can be configured.

2.6.3 Routing Option Sets

A Routing Option Set defines the request attempt timeout and/or the routing actions the DSR takes in response to a connection failure, no-peer-response or connection congestion conditions. These are assigned per App ID, or Ingress Peer Node. This feature allows for the creation of up to 50 routing option sets (ROS) (including default) which can then be optionally associated to a diameter transaction in several ways (in precedence order):

- If the Transaction Configuration Group is selected on the ingress peer node configuration object, then the Transaction Configuration Group is used and the longest/strongest match search criteria is applied.
- The Routing Option Set is assigned to the ingress peer node.
- The Routing Option Set is assigned to the default TCG.
- The system default ROS is used.

Some items included in the Routing Option Set are:

- Resource Exhausted Action
- No Peer Response Action
- Connection Failure
- Connection Congestion Action
- Maximum Forwarding
- Transaction LifeTime
- Pending Answer Timer (PAT)

Alternate routing is supported in cases of transport failure, message response timeout and upon receipt of user defined answer responses.

Alternate Routing on Answer

- User defines which Result Codes trigger alternate routing.
- User defines which Application IDs are associated with each Result Code.

Alternate routing on transport failure

- Connection failure occurs after message has been sent.
- T-bit set on re-routed message to warn of possible duplicate.

Alternate routing on timeout

- No response received for message.
- T-bit set on re-routed message to warn of possible duplicate.

2.6.4 Pending Answer Timer

Pending Answer Timers specify the amount of time the DSR waits for an Answer after sending a Request to a Peer Node. DSR allows for the specification of up to 16 pending answer timers that can be associated with the transactions/peers. This allows for different peers to respond to answers with different response times.

This feature addresses the ability to configure the Pending Answer Timer in the DSR which can then be optionally associated to a diameter transaction in several ways (in precedence order):

- If the Transaction Configuration Group is selected on the ingress peer node configuration object, then the transaction configuration group is used and the longest/strongest match criteria is applied for request message parameters to compare and if a match is found, then the PAT assigned to the transaction set defined under this group.
- The PAT from the ROS assigned to the ingress peer node is used.
- The PAT assigned to the egress peer node is used.
- The PAT assigned to the default TCG is used.

- The System default PAT is used.

2.6.5 Transport

The DSR supports SCTP and TCP transport simultaneously including support for both protocols to the same Diameter peer. The DSR supports UDP transport for Radius. The DSR supports up to 64 connections per single Diameter peer which can either be uni-homed via TCP or SCTP or multi-homed via SCTP. The DSR maintains the availability status of each Diameter peer. Supported values are *available*, *unavailable* and *degraded*.

The following information are some of the configurable items for each connection:

- Peer Host FQDN, Realm ID and optionally IPv4 or IPv6 address.
- Local Host and Realm ID (defined as part of the Diameter node).
- Message Priority Configuration Set.
- Egress Throttling Configuration Set.
- Remote Busy Usage / Remote Busy Abatement Timer.
- Transport Congestion Abatement Time-out.
- DSR Local Node status as the connection initiator only, initiator & responder (default) or responder-only.
- Other connection characteristics such as timer values detailed below.
- For SCTP connections:
 - RTO.Initial
 - RTO.Min
 - RTO.Max
 - RTO.Max.Init
 - Association.Max.Retrans
 - Path.Max.Retrans
 - Max.Init.Retrans
 - HB.Interval
 - SACK Delay
 - Maximum number of Inbound and Outbound Streams
 - Partial Reliability Lifetime
 - Socket Send/Rx Buffer
 - Max Burst
 - Datagram Bundling
 - Maximum Segment Size
 - Fragmentation Flag
 - Data Chunk Delivery Flag
- For TCP connections:
 - Nagle Algorithm ON/OFF indicator.
 - Socket Send/Rx Buffer.

- Maximum Segment Size (bytes).
- TCP Keep Alive.
- TCP Idle Time For Keep Alive.
- TCP Probe Interval For Keep Alive.
- TCP Keep Alive Max Count.
- Diameter Connect Timer (Tc as per RFC6733).
- Diameter Watchdog Timer Initial value (as per RFC3539).
- Diameter Capabilities Exchange Timer (Oracle extension to RFC6733).
- Diameter Disconnect Timer (Oracle extension to RFC6733).
- Diameter Proving Mode (Oracle extension to RFC3539).
- Diameter Proving Timer (Oracle extension to RFC3539).
- Diameter Proving Times (Oracle extension to RFC3539).

DSR supports multiple SCTP streams as follows:

- DSR negotiates the number of SCTP inbound and outbound streams with peers per RFC4960 during connection establishment using the number of streams configured for the connection.
- DSR sends CER, CEA, DPR, and DPA messages on outbound stream 0/
- If stream negotiation results in more than 1 outbound stream toward a peer, DSR evenly distributes DWR, DWA, Request, and Answer messages across non-zero outbound streams.
- DSR accepts and processes messages from the peer on any valid inbound stream.

The DSR supports SCTP multi-homing as an option which provides a level of fault tolerance against IP network failures. By implementing multi-homing the DSR can establish an alternate path to the Diameter peers it connects to through the IP network using SCTP protocol. Failure of the primary network path will result in the DSR re-routing Diameter messages through the configured alternate IP path. Multi-homed associations can be created through multiple IP interfaces on a single MP blade. This is independent of any port bonding existing on the Ethernet interfaces. Multi-homing is supported for both IPv4 and IPv6 networks but IPv4 and IPv6 cannot co-exist on the same connection.

Figure 2-11 SCTP Multi-Homing





Figure 2-12 SCTP Multi-Homing via Port Bonding

2.6.6 Message Prioritization

This feature provides a method for DSR administrators to assign message priorities to incoming Diameter requests. This priority configuration can be associated with a connection, peer node, application routing rule, or a peer routing rule. As messages arrive they are marked with a message priority. Once the message priority is set it can be used as input into decisions around load shedding and message throttling.

The Message Priority Configuration Set (MPCS) table is used for this configuration. The following are some of the defined methods used for setting message priority:

- Based on the connection upon which a message arrives.
- Based on the peer from which a message is sent.
- Based on an Application Routing Rule .
- Based on a Peer Routing Rule.

Each MPCS contains the following information:

MPCS ID – The ID is used when associating the configuration set with a connection.

Set of Application-ID, Command-code, priority tuples, also called message priority rules.

- **Application-ID** – The Diameter application-ID. The application-id can be a wildcard indicating that all application-ids match this message priority rule.
- **Command-code** – The Diameter command-code. The command-code can be a wildcard indicating that all command-codes within the specified application match this message priority rule.

Note

If multiple command-codes with the same appl-id are to get the same message priority then there will be a separate message priority rule tuple for each command-code.

- **Priority** – The priority applied to all request messages that match the Application-ID, Command-Code combination.

2.6.7 Diameter Routing Message Priority

RFC 7944 Diameter Routing Message Priority (DRMP) is the IETF standard which defines a mechanism to allow Diameter endpoints to indicate the relative priority of Diameter transactions. With this information, Diameter nodes can factor that priority into routing, resource allocation and overload abatement decisions. Message priority is embedded into IETF defined DRMP AVP of diameter messages with priority value ranging from 0 through 15 where 0 is the highest priority value and 15 is the lowest priority value. DRMP allows message priority assignment based on Diameter transactions that is request and answer message shall have same message priority. DSR uses the DRMP AVP based message priorities for message throttling decisions during congestion conditions similar to message priorities defined using Message Priority Configuration Set (MPCS) at DSR.

DSR provides the system configuration option to enable the support for 16 message priorities or legacy 5 message priorities. DRMP feature can be used at DSR only if the support for 16 message priorities has been enabled. DRMP feature can be enabled for individual Diameter Application Ids which allows DSR to assign message priorities to ingress diameter messages based on DRMP AVP only for configured Diameter Application Id's. If no DRMP AVP is present in the ingress diameter message then message priority shall be assigned based on MPCS configurations at DSR. The operator is also provided a configuration option called "Answer Priority Mode", a System Options attribute, for selecting which method to use for assigning priority to Answer messages - via the DSR legacy method of reserving the highest priorities for normal Answers (Highest Priority Mode) or the DRMP method of making the Answer priority the same as the Request priority (Request Priority Mode). When Highest Priority Mode is set, DSR ignores DRMP AVPs in Answer messages because the operator has chosen to ignore the DRMP method of assigning priority to Answers.

2.6.8 TLS / DTLS

The DSR optionally supports TLS for TCP connections and DTLS for SCTP associations in the DSR. This provides RFC compliant support for security protocol enabled certificate and key exchange. TLS/DTLS can be independently enabled on each DSR diameter connection. TLS/DTLS encrypts packets within a segment of network TCP connections or SCTP associations at the application layer using asymmetric cryptography for key exchange, symmetric encryption for privacy, and message authentication codes for message integrity. TLS/DTLS provides tighter encryption via handshake mechanisms. This feature uses the certificate management component from platform. Please see DSR for more information on the certificate management feature Capability Exchanges.

The Capability Exchanges on the DSR provide flexibility to inter-op with other Diameter nodes. These enhancements include:

- Support of any Application –Id.
- Configurable list of Application-Ids (up to 20 maximum) that can be advertised to the peer on a per connection basis.
- Authentication of minimum mandatory Application-Ids in the advertised list.
- Support for more than one Vendor specific Application-Id.

2.6.9 Configurable Disable of CEx Peer IP Validation

The DSR provides a mechanism to enable or disable the validation of Host-IP-Address AVPs in the CEx message against the actual peer connection IP address on a per connection configuration set basis.

2.6.10 Diameter Peer Discovery

The base Diameter protocol specification RFC6733 mandates that both dynamic Diameter agent discovery and manual configuration mechanisms be supported by all Diameter implementations; and either or both may be used in the network deployment.


From the DSR signaling point of view there are three basic use-cases for dynamic Diameter peer discovery.

- **initiator mode:** DSR discovering the last-hop Diameter peers.
- **initiator+responder mode:** DSR discovering Diameter (Edge) Agent for further handling of a Diameter operation. It is combination of the above 2 uses cases between two end-points.

The DSR supports the above listed deployment use-cases. The support for Dynamic Peer Discovery provides:

- the capability to configure realms that are dynamically discovered using RFC 6733 extended NAPTR methods.
- For a DNS Client Application instance that performs dynamic discovery.
- OAM functions that update/create the managed objects that are used for Diameter signaling.
- The ability to accept connections from configured realms.

Figure 2-13 Dynamic Diameter Peer Discovery: Example



In the above example for 'initiator' mode, each DSR node does the following:


- Monitors configuration changes.

- Creates tags required for Diameter extended NAPTR (S-NAPTR) query.
- Invokes DNS Client Application Interface for query resolution towards configured DNS Servers.
- Provides DNS Client Application Interface, processes the DNS responses and resolves NAPTR, SRV, A, AAAA lookups
- Performs target server resolution mapping to Diameter peer attributes for specified realm.
- Invokes OAM interface to update discovered Diameter peer attributes in DSR configuration managed objects.
- Replicates DSR configuration managed objects to DA-MPs. The signaling functions become aware of the required peer attributes and initiates connection establishment and Diameter capabilities exchange.

In the above example for 'initiator+responder' mode, 'Initiator+responder' mode for peer discovery is possible using one 'initiator+responder' connection.

2.6.11 Implicit Realm Routing

Implicit Realm Routing provides realm routing using DNS SRV load balancing information. The figure below illustrates the high level flow of Diameter Request forwarding/routing decision points on DA-MP blades. Note that Destination-REALM and Application-Id based implicit realm routing is added after the Destination-Host based implicit routing. Implicit realm routing is only performed for routing messages to dynamically discovered peers.

Figure 2-14 High Level DSR Routing Flow – Fall through to Dest-Realm Based Implicit Routing

2.6.12 DNS Support

The DSR supports DNS lookups for resolving peer host names to an IP address. The operator can configure up to two DNS server addresses designated as primary and secondary servers. The wait time for DNS queries for connections initiated by the DSR is configurable between 100 to 5000 milliseconds with a default of 500 milliseconds. This process is used for both dynamic peer discovery and A/AAAA lookups.

The DSR supports both A (IPv4) and AAAA (IPv6) DNS queries. If the configured local IP address of the connection is IPv4 the DSR will perform an "A" lookup and if it is IPv6 the DSR will perform an "AAAA" lookup. If the IP address of the connection is undefined by the operator, the DSR will resolve the host name using both A and AAAA DNS queries when initiating the connection. The DSR can either use the peer's FQDN or an FQDN specified for the connection as a hostname for the DNS lookup.


2.6.13 Signaling Firewall

Signaling Firewall feature is the network security feature of DSR which configures native Linux 'iptables' rules in the Linux firewall on each DA-MP server to allow only essential network traffic pertaining to the active signaling configuration. The in-bound signaling traffic is accepted by the DSR application only over the administratively enabled Diameter and Radius connections configured at DSR SOAM.

Signaling Firewall feature provides the following capabilities at DSR:

- Capability to automatically configure the Linux firewall to allow desired signaling network traffic on DA-MPs.
- Capability to dynamically update the Linux firewall configuration on DA-MPs to allow or disallow signaling traffic.
- Capability to administer (Enable and Disable) the DSR Signaling Firewall on the Signaling Node via System OAM configuration user interfaces.

Figure 2-15 DSR Signaling Firewall




20

Note

This feature does not apply to IPFE servers and hence there is no impact on the IPFE function.


2.6.14 Support Answer on Any Connection

DSR supports processing of answer messages from connections that are different to the connections used to send the request to the upstream Peer node. This feature can be enabled for individual Peers configured at DSR. Upstream Peer nodes can respond back answers to request received from DSR on any connection without the need to follow the same path as the received request.

Figure 2-16 Answer processing on connection different from the connection used to send Request message

2.6.15 Congestion Control

The DSR supports local and remote congestion control via the use of congestion levels. Congestion levels are defined for which only a percentage of Request messages are processed during the congestion period. The DSR supports a method for limiting the volume of Diameter Request traffic that DSR is willing to receive from DSR peers. In addition, the DSR provides a method for partitioning the MPS capacity among DSR peer connections, providing some user-configurable prioritization of DSR traffic handling. Congestion levels correspond to minor, major and critical alarms associated with resource utilization. The percentage of Request messages to be processed for each level is shown below. The DSR may return a user configurable Answer message when a Request message is not successfully routed during congestion. Under severe congestion conditions, the DSR may not return an Answer message. Request messages that are not processed will be discarded. An OAM event will be raised upon entering and exiting congestion levels. If the Next Generation Network Priority Service (NGN-PS) feature is enabled, these DSR Congestion Control mechanisms do not affect processing of NGN-PS messages. Refer to Next Generation Network Priority Service (NGN-PS) for more information.

Figure 2-17 Congestion Control

2.6.15.1 Per Connection Ingress MPS Control

The Per-Connection Ingress MPS Control feature provides the following:

- A method to reserve/guarantee a user-configured minimum ingress message capacity for each peer connection.
- A method for limiting the ingress message capacity for a peer connection to a user-configured maximum.
- A method for multiple peer connections to have a 'shared' ingress message capacity.
- A method to prevent the total reserved ingress message capacity of all active peer connections on a DA MP from exceeding the DA MP's capacity.
- A method for limiting the overall rate at which a DA MP attempts to process messages from all peer connections.
- A method for coloring (Green or Yellow) messages ingressing a DSR.

There are six user-configurable capacity configuration set parameters for DSR Connections: Ingress MPS Minor Alarm Threshold, Ingress MPS Major Alarm Threshold, Abatement Time, Reserved Ingress MPS, Maximum Ingress MPS and Convergence Time. Additional details on some of these follow.

- Reserved Ingress MPS:
 - Ingress capacity (in Messages per Second) reserved for use by the peer connection. It is not available for use by other connections on the same DA MP.
 - Min value: 0
 - Max value: Minimum (Connection engineered capacity, DA MP's licensed MPS capacity)
 - Default: 0

When a DSR Connection's ingress message rate is equal to or below its configured Reserved Ingress MPS, all messages ingressing the connection are colored Green. When a DSR Connection's ingress message rate is above its configured Reserved Ingress MPS, messages below the reserved capacity are colored green and messages above the reserved capacity are colored yellow.

- Maximum Ingress MPS:
 - Maximum ingress capacity (in Messages per Second) allowed on this connection. Capacity beyond “reserved” and up to “max” is shared by all connections on the DA MP and comes from DA MP capacity leftover after all connections’ “reserved” capacities have been deducted from the DA MP capacity.
 - Min value: 10.
 - Max value: Minimum (Connection engineered capacity, DA MP’s licensed MPS capacity).
 - Default: Minimum (Connection engineered capacity, DA MP’s licensed MPS capacity).

A fundamental principal of Per-Connection Ingress MPS Control is to allocate a DA-MP’s ingress message processing capacity among the Diameter peer connections that it hosts. Each peer connection is allocated, via user-configuration, a reserved and a maximum ingress message processing capacity. The reserved capacity for a connection is available for exclusive use by the connection. The capacity between a connection’s reserved and maximum is shared with other connections hosted by the DA-MP. The DA-MP reads messages arriving from a peer connection and attempts to process them as long as reserved or shared ingress message capacity is available for the connection.

- Convergence Time:
 - Convergence time in ms used by an algorithm for Ingress MPS rate computation by PCIMC.
 - Min value: 250ms
 - Max value: 4000ms
 - Default: 1000ms

Rate Convergence Time is used in message rate computation done using an algorithm called Sliding-Historic Metric where traffic history (that is message count) is stored for a configured time. The DSR maintains a sliding history using an array of traffic counts (ex: one second history), where each element in the array represents time in ms (for example: 50 ms) of elapsed time. Message rate is calculated as message per second (MPS). Rate Convergence Time allows the user to control the sensitivity of the request traffic bursts and allows them to tune accordingly.

When the ingress messages are above the connection Maximum ingress MPS rate the DA-MP enforces a short discard period, during which time ingress messages over the defined maximum Ingress MPS are read from the connection and discarded based on the message priority. This approach provides some user-configurable bounding of the DSR application memory and compute resources that are allocated for each peer connection, reducing the likelihood that a subset of DSR downstream peers which are offering an excessive/unexpected Request load can cause DSR congestion or congestion of DSR upstream peers. The discarding of ingress messages by the DSR results in the DSR Peer experiencing Request timeouts (when DSR discards Request messages) and/or receiving duplicate Requests (when DSR discards Answer messages).


It should be noted that the DSR is enforcing ingress message rate independent of the type (that is Request or Answer) or size of the ingress messages.

The figure below depicts a DSR DA MP hosting 3 connections with the attributes shown in the following table:

Table 2-6 DSR Ingress MPS Configuration Example 1

Connection	Reserved Ingress MPS	Maximum Ingress MPS	MPS shared with other connections
Connection1	100	500	400
Connection 2	0	5000	5000
Connection 3	500	500	0

Figure 2-18 DSR Ingress MPS Configuration Example 1 – Normal Case




The DSR prevents the total Reserved Ingress MPS of all connections hosted by a DA MP from exceeding the DA MP's maximum ingress MPS. The enforced limit for this is the DA MP's licensed MPS capacity, which defaults to the DA MP's maximum engineered capacity. The enforcement of this requirement on 'configured' connections versus 'Enabled' or 'Active' connections is a design decision.

This feature addresses the functionality to assist DSR overload and throttling algorithms in differentiating messages ingressing a DSR connection whose ingress message rate is above (vs equal to or below) its configured reserved ingress MPS.

When a DSR connection's ingress message rate is equal to or below its configured reserved ingress MPS, all messages ingressing the connection are colored green. When a DSR connection's ingress message rate is above its configured reserved ingress MPS, messages below the reserved capacity are colored green and messages above the reserved capacity are colored yellow. Message color is used as a means for differentiating diameter connections that are under-utilized versus those that are over-utilized with respect to ingress traffic. Traffic from under-utilized connections are marked "green" by the per-connection ingress MPS control

(PCIMC) feature. As stated above, when a DSR connection's ingress message rate is above its configured reserved ingress MPS, messages below the reserved capacity are colored green and messages above the reserved capacity are colored yellow. In the event of danger of congestion or of CPU congestion and based on the specified discard policy, traffic from over-utilized connections is considered for discard before traffic from under-utilized connections. Traffic discarded by PCIMC due to capacity exhaustion (per-connection or shared) is marked "red" and is not considered for any subsequent processing.

Figure 2-19 Message Coloring and Priority or Color based DA-MP Overload Control



2.6.15.2 MP Overload Control

DSR MP Overload Control utilizes proven platform infrastructure to monitor the CPU utilization of each DSR MP and implement incremental load-shedding algorithms as engineered CPU utilization thresholds are exceeded. MP overload control provides DSR stability in the presence of extremely deteriorated network conditions, message loads that exceed the engineered capacity of a DSR MP, or improper configurations. It is important to note that MP overload control algorithm only monitors and acts on the CPU utilization of the DSR MP software functions (that is message & event handling), allowing a sufficient CPU budget for other non-critical (that is best effort) DSR MP functions. In this way, the load-shedding algorithms are not invoked when non-critical DSR MP functions consume more than their budgeted CPU when it has no impact on critical DSR MP functions. Message priority and Message color are used as input to the DSR's message throttling and shedding decisions. In addition, exponential smoothing is applied to the CPU utilization samples in order to prevent the load-shedding algorithms from introducing more instability to an already degraded system

A DA-MP Danger of Congestion (DOC) threshold is less than the threshold set for DA-MP congestion level "1". There is a DOC onset threshold, a DOC abatement threshold, and a DOC warning event.

When it has been determined that a system is actually in congestion, the request messages discarded are based on the priority of the message, the color of the message, and the user-

configurable DA-MP Danger of Congestion discard policy. There are three user-configurable options:

- Discard by color within priority (Y-P0, G-P0, Y-P1, G-P1, Y-P2, G-P2,....., Y-P15, G-P15).
- Discard by priority within color (Y-P0, Y-P1, Y-P2,.....,Y-P15, G-P0, G-P1, G-P2,.....,G-P15).
- Discard by priority only (P0, P1, P2,....., P15).

The following elements are configurable for the DA-MP Overload Control feature:

- Congestion Level 1 Discard Percentage – The percent below the DA-MP engineered ingress MPS that DA-MP overload control polices the total DA-MP ingress MPS when the DA-MP is in congestion level 1.
- Congestion Level 2 Discard Percentage – The percent below the DA-MP engineered ingress MPS that DA-MP overload control polices the total DA-MP ingress MPS to when the DA-MP is in congestion level 2.
- Congestion Level 3 Discard Percentage – The percent below the DA-MP engineered ingress MPS that DA-MP overload control polices the total DA-MP ingress MPS to when the DA-MP is in congestion level 3.
- Congestion Discard Policy – The order of message priority and color-based traffic segments to consider when determining discard candidates for the application of treatment during DA-MP congestion processing.
- Danger of Congestion Discard Percentage – The percent of total DA-MP ingress MPS above the DA-MP Engineered Ingress MPS that DA-MP Overload Control discards when the DA-MP is in danger of congestion.
- Danger of Congestion Discard Policy – The order of Message Priority and Color-based traffic segments to consider when determining discard candidates for the application of treatment during DA-MP Danger of Congestion (DOC) processing. The following order is considered: Color within Priority, Priority within Color, and Priority Only.

As the DSR MP CPU utilization exceeds the engineered thresholds, the MP congestion level is updated and message load-shedding is performed by the DSR.

2.6.15.3 Internal Resource Management

DSR utilizes proven platform infrastructure to monitor, alarm, and manage the resources used by internal message queues and protocol data unit (PDU) buffer pools to prevent loss of critical events and monitor and manage PDU pool exhaustion.

Message Queue Management

- Enforces a maximum queue depth for non-critical events; non-critical events are never allowed to overflow a queue's maximum capacity.
- The system attempts to always queue critical events even when the queue's maximum capacity is reached.
- Measurements and informational alarms are maintained for discards of all events.

PDU Buffer Pool Management

- Similar to message queues, the DSR monitors the size of each PDU Buffer Pool, alarms when the utilization crosses configured thresholds, and discards messages when the PDU Buffer pool is exhausted.
- Measurements are maintained for all discards.


2.6.15.4 Egress Transport Congestion

When a DSR peer connection becomes blocked due to transport layer congestion the DSR acts in the following manner:

- When a DSR peer connection becomes blocked, the DSR sets the connection's congestion level to CL-4 (Requests nor Answers can be sent on the connection).
- The DSR waits for the connection to unblock and then abate a connection's egress transport congestion using a time-based step-wise abatement algorithm similar to Remote BUSY Congestion.
- A user-configurable Egress Transport Abatement Timer exists for each DSR Peer Connection. The abatement timer defines the time spent abating each congestion level during abatement and is not started until the socket unblocks and becomes writable.
- Messages already committed to the connection by the DSR routing layer when a connection initially becomes transport congested will be discarded.

The above can be summarized using the chart below.

Figure 2-20 Example Congestion level abatement



2.6.15.5 Per Connection Egress MPS Control

The Egress Message Throttling feature provides a mechanism that assists with the prevention of Diameter peer overload. It does so by allowing the user to configure the max Egress Message Rate (EMR) on a per connection basis and shedding messages as the offered message rate gets closer to the max EMR. The feature works in conjunction with the message prioritization infrastructure and provides intelligent load shedding based on the volume of the offered load. The load shedding is performed by dropping requests based on priority and the offered Message Rate. It should be noted that message priorities are assigned using DRMP AVP when DRMP feature is enabled or through MPCS configurations.

The connection egress message throttling behavior is governed by user-configurable Egress Message Throttling Configuration Sets. Each Egress Message Throttling Configuration Set contains:

- A maximum allowed EMR.
- A minimum of one and up to a maximum of three pairs of user-configurable EMR Throttle and Abatement Thresholds (TT & AT) expressed as % of max EMR.
- Convergence Rate: The time the algorithm takes for the measured rate to converge on the actual rate. Useful for bursty traffic.
- Abatement Time

The “maximum allowed EMR” dictates the maximum volume of traffic that can be served over a particular connection. Each EMR throttle & abatement threshold pair are then expressed as percentages of the maximum allowed EMR and dictate how the connection congestion state will be updated.

The DSR allows for egress message throttling to be enabled for at least 500 peer connections in a single DSR NE. To enable egress message throttling on a connection, the user creates an Egress Message Throttling Configuration Set and assigns it to one or more DSR peer connections that are to be throttled using the configuration set settings. The DSR supports at least 50 user-configurable Egress Message Throttling Configuration Sets.

2.6.15.6 Egress Throttle Group (ETG) Limiting

Network operators cannot control the ingress load-shedding behavior of all nodes in their networks and many become unstable and fail when offered excessive ingress traffic loads. Therefore, DSR can be utilized to enforce maximum egress traffic rates and maximum pending transaction counts on a connection, a peer, or an aggregate group of connections/peers.

- Egress Throttle Group Rate Limiting: A method to control the total egress Request traffic rate that DSR can route to a user-defined group of connections or peers.
- Egress Throttle Group Pending Transaction Limiting: A method to control the total number of transactions that DSR can allow to be pending for a user-defined group of connections or peers.

DSR supports two modes of ETG limiting:

- Threshold Throttling Mode: DSR limits the diameter requests or pending transactions for a given ETG based on user-defined onset and abatement thresholds set for ETG's rate congestion level (ETG-RCL), values: CL1, CL2 and CL3. This mode can only be used when 16 Message Priorities is disabled.
- Limit Throttling Mode: DSR performs throttling by measuring the rate of Request messages offered to each ETG and divert the traffic based on configured Congestion Discard Policy when request rate or pending transaction exceeds the user-defined maximum traffic allowed. This method eliminates congestion levels and the need for the user-defined congestion level onset and abatement thresholds. This mode can be used with both legacy 5 message priorities and 16 message priorities.

These features provide DSR egress throttling capability that allows the user to:

- Configure an ETG with a max of 128 entries, each peer/connection can be in only 1 ETG.
- Identify a group of peers and/or connections and associate them with an Egress Throttle Group.
- Set the ETG's maximum egress Request rate.


- Configure throttling and abatement thresholds or maximum allowed ETG request rate and pending transactions based on ETG limiting mode.
- Set convergence rate and abatement timer.
- Set the ETG's maximum pending transaction limit.

2.6.15.6.1 Example: DSR Connects to a Single Server Node with Multiple Connections

DSR typically connects to a single server node with more than 1 connection for redundancy (and sometimes for capacity). DSR per-connection egress throttling functionality may result in underutilization of a server node's capacity when a subset of the DSR connections to the server node fail and the remaining connections are capable of carrying the full capacity of the server node. For example, consider the scenario depicted in the figure below where:

- Constraint 1: Server 1 has a total capacity of X TPS.
- Constraint 2: Server 1 can process as much as 50% of its total capacity on a single connection.
- DSR throttles each connection to Server 1 to $X/3$ (addresses constraint 1 only).

Figure 2-21 DSR Per-Connection Egress Throttling




In the above example, the per-connection egress throttling is used to limit the aggregate egress traffic rate to Server 1 (constraint 1). As a result, each of the 3 connections to Server 1 must be throttled at $1/3$ of Server 1's capacity to prevent DSR from offering a load greater than X when all 3 connections are in-service. However, if one of the connections to Server 1 fails DSR will restrict egress traffic to $2/3$ of Server 1's capacity even though the remaining two connections are capable of carrying the entire capacity of Server 1.

The ability for DSR to throttle the aggregate egress traffic across all 3 DSR connections to Server 1 while also throttling the egress traffic on individual connections to Server 1 reduces the limitations described above. This is shown in the figure above where:

- Constraint 1: Server 1 has a total capacity of X TPS.
- Constraint 2: Server 1 can process as much as 50% of its total capacity on a single connection.

- DSR throttles the aggregate egress traffic over all connections to Server 1 to X (addresses constraint 1).
- DSR throttles each connection to Server 1 to X/2 (addresses constraint 2).

Figure 2-22 DSR Aggregate and Per-Connection Egress Throttling



In the figure above, use of aggregate egress traffic rate limiting to address constraint 1 allows the per-connection egress throttling limits to be relaxed as it is being used appropriately to address the connection constraint (constraint 2).

The DSR can aggregate and distribute information about the ETG across all DA-MPs for use in routing decisions.

During Request routing, if the DSR selects a peer/connection that is a member of an ETG and determines that either the rate or pending transaction cumulative limit for that ETG has already been reached, then the DSR does not route to that peer/connection and continues to search for an acceptable peer/connection via standard DSR routing operations.

DSR utilizes the existing user-configurable response behavior in the Routing Option Set for Requests that are throttled and cannot be routed via other connections.

DSR uses standard alarming capabilities against the ETG to alert the user when limits are exceeded.

2.6.15.6.2 Connection Pending Transaction Limiting

This feature makes the connection Pending Transaction Limiting attribute user configurable and tunable on a per connection basis. The primary use of Connection Pending Transaction Limits on a DSR DA-MP is to prevent a small number of connections on a DA-MP from consuming a disproportionate number of the available Pending Transaction Records on the DA-MP, which could result in limited Pending Transaction Record availability for the remaining connections.

DSR peer nodes have differing requirements regarding the maximum number of pending transactions required on the DSR


- DSR-to-Server connections typically carry higher traffic volumes than DSR-to-Client connections due to DSR aggregation of traffic from many client connections to few server connections.
- A high percentage of the traffic on DSR-to-Server connections requires Pending Transaction Records in the DSR since the majority of the traffic egressing the DSR on these connections are Requests.
- A low percentage of the traffic on DSR-to-Client connections requires Pending Transaction Records in the DSR since the majority of the traffic egressing the DSR on these connections are Answers.
- DSR-to-Server connections may encounter significant increases in offered load in a very short time immediately following network events such as MME failures or failures of redundant Servers providing the service. 'Riding through' these types of sudden increases in traffic volume may require higher Pending Transaction Limits on the connections.

In order to support customization of the distribution of the available Pending Transaction Records on a DA-MP based on the varying deployment requirements, this feature provides user-configuration of the Connection Pending Transaction Limit for each DSR peer connection. The limit configured is enforced independently by all DA-MPs in the DSR.

2.6.15.7 Coordinated Egress Throttling Across Multiple DSRs

When multiple DSRs (mated pair or triplet) connect to common servers, there is a need for the DSRs to share egress throttling information to avoid under-utilization or overload of the common servers in load share or failure scenarios. This feature allows multiple DSRs to share real-time Egress Throttle Group Rate and Pending Transaction information in order to maximize utilization of servers common to the DSRs while also protecting the common servers from overload.

To address communication failure amongst the contributing DSRs when under coordinated egress throttling, DSR supports a user configuration option that specifies how much the coordinated ETGs Rate and/or Pending Transaction Limit should be reduced from the coordinated maximum egress rate and pending transaction value. This user configurable option 'Coordination Failure (% Reduction)' affects egress Request rate and pending transaction maximum value proportional to the number of peer DSR communication failures. Also, please note that this 'Coordination Failure (% Reduction)' parameter does not apply when a DSR is providing SOAM managed single DSR scoped egress throttling.

Figure 2-23 2 DSR Sites- Coordinated Egress Throttling Example

2.6.15.8 Remote Busy Congestion

The intent of this feature is to provide remedial measures if it is determined that a connection to a DSR peer node is unable to process messages as fast as they are sent to it on a given DSR connection to the peer node. A connection is considered congested (BUSY) if an Answer message containing 'DIAMETER_TOO_BUSY' result code is received on the connection and was originated by the peer node.

Remote BUSY Congestion is determined by analyzing Diameter Answer from a connected peer. The result code 'DIAMETER TOO BUSY' in a Diameter Answer from a connected peer indicates the connection is congested or BUSY.

When this feature is configured, DSR sets the status of a connection to 'BUSY' in the following conditions:

- The result code of Diameter Answer is 'DIAMETER TOO BUSY'.
- Origin-Host of the Answer messages is same as the connection's Peer FQDN.

The DSR sets the status 'BUSY' only to the connection of a peer on which 'DIAMETER TOO BUSY' is received. The other connections between the DSR and the peer may or may not be BUSY.

Typically, if a connection is BUSY, it is not selected for routing of Diameter Request messages. However, based on the configuration, this behavior may be overridden and a BUSY connection may be selected to route the Request when the message is addressed to the connection's peer FQDN.

A BUSY connection becomes uncongested after a certain minimum time has elapsed in 'BUSY' state. DSR provides a configurable timer to set this value.

Note


Diameter Protocol does not provide any mechanism for a node to signal to its peers that its busy condition has abated.

The figure below shows the message flow diagram for determination of congestion in a normal case.

- DSR receives a Diameter Request Message.
- DSR selects a connection and forwards it to a connected peer (Server).
- The peer replies with 'DIAMETER_TOO_BUSY' result code in the Answer.
- DSR sets the Connection Status to 'BUSY' and starts 'Connection Busy Abatement Timer'.
- DSR forward the DIAMETER_TOO_BUSY to client.

If 'Reroute on Answer' feature is configured, the DSR may attempt to perform alternate routing of Request based on DSR routing configuration.

Figure 2-24 Connection Busy



Request Priority for which a remote busy was received	Associated Connection Congestion Level	Message Priorities Allowed when support for 5 message priorities is enabled	Messages Priorities Not Allowed when support for 4 message priorities is enabled	Comment
2	CL-3	3	0,1,2	Only allow Answers to be sent on connection.

Request Priority for which a remote busy was received	Associated Connection Congestion Level	Message Priorities Allowed when support for 5 message priorities is enabled	Messages Priorities Not Allowed when support for 4 message priorities is enabled	Comment
1	CL-2	3,2	0,1	Only allow Answers and Priority=2 Requests to be sent on connection.
0	CL-1	3,2,1	0	Only allow Answers and Priority=2, 1 Requests to be sent on connection.


Note

When support for 16 message priorities is enabled then message priorities to be allowed in each of the Connection Congestion Level is decided based on 'Connection Configuration Set' MO assigned to individual diameter connection. The operator must specify a "Minimum Request Priority Allowed" attribute for each of CL1, CL2 and CPL3 in Connection Configuration Set. The "Minimum Request Priority Allowed" attribute assigned to a CL value defines the minimum Request message priority criteria for forwarding Request messages to that connection. For example, if the "Minimum Request Priority Allowed" attribute assigned to CL3 is 12, then DSR will forward Request messages to that connection with a priority of 12, 13, 14 or 15.

When the abatement timer expires, the congestion level is decremented by one thereby allowing Requests with the next lower priority and the abatement timer is restarted. For the example above, after the abatement timer expires, priority 2 and above Requests will be allowed over the connection. This process continues until the congestion level of the connection drops back to zero. This behavior is illustrated in the figure below:

Note

Diameter Protocol does not provide any mechanism for a node to signal to its peers that its busy condition has abated.

Figure 2-25 Congestion level abatement over time for Remote Busy

2.6.15.9 Remote Transport Congestion Control

Egress transport congestion control occurs when a DSR diameter peer connection's TCP/SCTP send buffer is exhausted, as indicated by the TCP/SCTP socket becoming 'blocked'.

When this occurs the DSR sets the connection's priority level to CPL-4. This means that no requests or answers can be sent on the connection. A user configurable abatement timer is used to control the period a connection stays in the CPL mode. The receive or transmit buffer sizes are user-configurable for the system.

2.6.15.10 Diameter Overload Indication Conveyance

Diameter Overload Indication Conveyance (DOIC) is a new IETF standard for supporting dynamic overload controls between Diameter servers and Diameter clients. It allows for Diameter servers to send overload reports requesting that Diameter clients reduce the traffic that they are sending to the server. It also allows for Diameter Agents such as the DSR to act as a proxy for any clients, by reducing traffic as requested by the servers, or as a proxy for the servers by requesting that traffic be reduced by the client.


There are two major interactions defined between the DOIC Reporting Node and the DOIC Reacting Node:

- The DOIC Capabilities Announcement (DCA) function allows the DOIC Reacting Node to send a list of its supported DOIC capabilities to the DOIC Reporting Node, and the DOIC Reporting Node to respond with its selected options if the DOIC Reacting Node gave it multiple options.
- The DOIC Reporting Node sends DOIC Overload Reports (OLR) to the DOIC Reacting Node requesting a reduction in traffic. The defined loss algorithm is an "abatement algorithm" which tells the DOIC Reacting Node to reduce the amount of traffic being sent by a given percentage.

DOIC layers on top of existing congestion controls in the DSR. Therefore, all of the current static controls such as ETLs/ETGs and the connection level congestion controls work as previously described.

The DSR supports the Reacting Node role in DOIC.

Figure 2-26 Reacting Node Role for DOIC




2.6.15.10.1 DOIC Capabilities Announcement

The DOIC solution supports the ability for Diameter nodes to determine if other nodes in the path of a request support the DOIC solution. The DOIC Capabilities Announcement (DCA) mechanism uses the OC-Supported-Features AVPs to indicate the Diameter overload features supported. This AVP is added by the DSR to all requests that are routed to a HostID/AppID defined in a Traffic Throttling Point (TTP).


The figure below shows the basic flow of DCA. This example assumes that a TTP has been created for a combination of HostID/AppID.

- A request is received by the DSR that AFTER ROUTING contains a HostID/AppID combination that matches the key of a TTP indicating that the DSR is functioning as a Reacting Node for that HostID/AppID. The evaluation has to be after routing in the DSR, since the request might have been Realm routed to DSR, or the HostID might have been changed by routing in the DSR (for instance from a Pseudo HostID to a real HostID).
- The DSR inserts the OC-Supported-Features AVP into every request message sent to that HostID/AppID. This AVP includes the list of all of the supported Abatement Algorithms on the DSR.
- The host returns in every answer message an OC-Supported-Features AVP indicating which of the abatement algorithms the DSR said it supported that the Host wants to use. While the DSR can include multiple supported abatement algorithms, the Reporting Node can only return one, the one it wants to use. The Host sends back the OC-Supported-Features AVP and optionally includes the OC-Feature-Vector that specifies the abatement algorithm. If the answer doesn't include the OC-Supported-Features AVP then the abatement algorithm defaults to the Loss algorithm.
- At this point the Host can start sending requests (Overload Reports) that causes the DSR to reduce the request traffic sent to the Host. The lifetime of an OC-Supported-Features exchange lasts for a single request and answer and so these steps are repeated for each request.

Figure 2-27 DOIC Capabilities Announcement

2.6.15.10.2 DOIC Overload Reports

Once the DSR starts sending OC-Supported-Features AVPs to the Reporting Node, the Reporting Node can start sending back DOIC Overload Reports (OLRs) requesting traffic abatement. The figure below gives a simple example of the Overload Report (OLR) mechanism.

Figure 2-28 DOIC Overload Reports Example

When the Host decides that it needs to reduce the traffic being sent to it, it includes (piggybacks) an OC-OLR AVP in an answer to a request message that included the OC-Supported-Features AVP. The OC-OLR AVP includes:

- Type of report (host, realm).
- Report id (Sequence Number).
- Length of time the report is valid.
- Abatement algorithm specific AVPs.

The DSR abates traffic based on the data in the OLR for the duration given in the report, or until the report is effectively cancelled by the DOIC Reporting Node sending a report with 0 for time the report is valid.


2.6.15.10.3 DOIC AVP Blocking

There are potential security issues with DOIC, for instance, an unauthorized third party or unauthorized node might inject an overload report into the network to throttle 100% of the traffic as a form of a Denial-of-Service (DoS) attack. OLRs also include potentially sensitive information such as network topology, current network status. And DOIC overload reports could contain sensitive information about the status of a vendor's network if they were allowed to transit to a roaming partner.

The DSR enforces a Hop by Hope trust model at the peer level to address this potential security issue. For every peer configured on a DSR it is possible to either allow (pass through)

or block (strip the DOIC AVPs) on both requests and answers. The figure below shows a logical picture of a network with both a “first hop” DSR 1, and a “last hop” DSR 2.

Figure 2-29 DOIC Security Setting Example



The following three scenarios are supported:

- All DOIC AVPs are stripped on all requests and answers sent and received on a connection to a given peer.
- All OC-Supported-Features AVPs are stripped from requests sent to the peer, and all OC-Supported-Features AVPs and OLR AVPs are stripped from answers received from the peer. The OC-Supported-Features AVP and the OLR AVP are allowed on answers sent to the Peer from the DSR. This mode allows the DOIC Reporting Node function to be done by either DSR or by a downstream peer. But it blocks the DSR or downstream peer from doing the DOIC Reacting Node function.
- All OC-Supported-Features AVPs and OLR reports are stripped on answers sent to the Peer. The OC-Supported-Features AVP is allowed on requests sent to the peer. This mode allows the DOIC Reacting Node function to be done on the DSR or on a downstream Peer.

2.6.15.10.4 Loss Abatement Algorithm

The supported DOIC abatement algorithm is the “Loss” abatement algorithm. It specifies a percentage of traffic that is abated for a given TTP. The Loss algorithm is stateless. It specifies a percentage of traffic that is abated of the traffic that would have been sent without the abatement, not a percentage of the previous traffic that caused the abatement request. Thus the reacting node does not guarantee that there is an absolute reduction in traffic sent, since the offered traffic may have increased since the Overload Report was sent. Rather, it guarantees that the requested percentage of new requests are given abatement treatment.

2.6.15.10.4.1 Abatement by Color/Priority


Since the DOIC Loss abatement algorithm is rate based, it is necessary for the DSR’s DOIC abatement algorithm to use rate based algorithm rather than the “threshold” (CL0-CL3) mechanisms currently used for ETGs/ETLs. The DOIC implementation of the loss abatement algorithm uses a Priority/Color mechanism similar to the one used for the DA-MP overload controls. It supports throttling by:

- Message priority only
- Message Priority first, then color
- Color first then Message Priority

For instance, if the abatement request is for a 10% reduction, rather than reducing all of the traffic by that rate, the DSR rejects from the lowest rank to the highest rank the requests to hit that target rate. The figure below shows an example of throttling by Priority. Based on the


requested reduction in an OLR, the DSR calculates the Max ETR for that TTP. The DSR then start discarding the lower priority messages (in this case P0 messages) until it hits the required rate.

Figure 2-30 Throttle by Message Priority Only when support for 5 message priorities is enabled



If the customer has selected the “Discard by Priority within Color” option, then the DSR makes the same calculation about the Max ETR, but as shown in the figure below it starts discarding “yellow” messages from lowest to highest priority before it discards any “green” messages of any priority.

Figure 2-31 Throttle by Color then Message Priority when support for 5 message priorities is enabled



Supporting Color and Priority throttling requires the rate of each Color and Priority be tracked per TTP. These message rates track messages arriving at the TTP, and do not include ingress messages discarded by the Per-Connection Ingress MPS Control component or by DA-MP overload. These rates are tracked using the existing Sliding Historic Metric (SHM) that is used for other overload controls within the DSR. This SHM mechanism allows the user to set a “convergence” time for how quickly the rate metric reacts to changes in the rate. A small Rate Convergence Time causes the calculated rate to react to short term bursts, where a larger number “smooths” bursty traffic.

The discard policy for the TTP is inherited from the Congestion Discard Policy parameter set in the DA-MP profile for each DA-MP.

2.6.15.10.5 Coming Out Of Overload

The DSR recovers at a fixed rate configurable by the user on a TTP basis, with the default being a relatively small number such as 5-10% increase per second. This approach allows the DSR to react quickly to small changes and slowly to large ones. It also provides a behavior which is easy to predict especially if the traffic is ramping up.

2.6.15.10.6 Traffic Throttle Point

A logical Traffic Throttle Point (TTP) is required to manage the DOIC relationship between the DSR and the Reporting Nodes. Some of the major items in the TTP include:

- DOIC scope for the TTP: HostID/AppID pair.
- Configuration parameters for the TTP.
- Tracks the rate information per color/priority.
- Tracks the administrative, operational and a throttling status (enabled/disabled, current abatement requests, so on).

The table below lists the data the TTP contains. The configuration data that is common between TTPs has been split out into a separate table shown below:

Table 2-7 TTP Data Components by Type

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
Scope	Entity Type	HostID	Y	Y	N/A	
	Entity Name	HostID	Y	Y	N/A	
	AppID	AppID name	Y	Y	N/A	
Configuration	See the table below	Choice list of Congestion Configuration Sets	Y	N	System Default	

Table 2-7 (Cont.) TTP Data Components by Type

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
	Maximum ETR	ETR value in Messages per Second	Y	Y	N/A	The maximum ETR allowed for the TTP in the absence of DOIC abatement. This field is mandatory since it is used as part of the calculation for TTG loss %. Note that this is ETR (requests only), not EMR (requests and answers).
	Alternate Implicit Route	Valid Route List on the DSR	Y	N	Blank	An optional Route List which specifies an alternate route (list) to use when "implicit routing" is invoked and the primary route to the Host is unavailable. The TTP Alternate Implicit Route List is higher priority (that is is used instead of) any Alternate Implicit Route List defined at the Peer level.
	Maximum Loss % Threshold	0-100%	Y	Y	100%	If the current loss rate for the TTP is greater than or equal to this number, routing should "skip" this TTP, and take whatever the next routing action is (that is treat it just like it didn't meet the "minimum weight" requirements for a Route List). A default of 100% mimics the current DSR behavior (that is ignores DOIC loss data).

Table 2-7 (Cont.) TTP Data Components by Type

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
Status	Throttling Admin State	Enabled/ Disabled	Y		Disabled	This admin state controls the overall throttling status of the TTP. When it is disabled, no throttling is done. When it is enabled, the TTP will at least do static throttling if the EMR value is defined. Whether the TTP also does "DOIC" dynamic throttling is set by the Dynamic Throttling Admin State.
	Dynamic Throttling Admin State					This admin state controls whether the TTP also performing DOIC throttling. When it is enabled the TTP will send DCA AVPs to the peer, and look for OLR AVPs in answers. It will also comply with any loss requests. When it is disabled the TTP will only do static throttling.
	Operational State	Available, Degraded, Inactive	Y	N/A	N/A	This state is driven by a number of factors such as the current loss % (that is the TTP degraded), and the operational status of the underlying peer.

Table 2-7 (Cont.) TTP Data Components by Type

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
	Operational State Reason	(similar to the existing ETG states)	Y	N/A	N/A	This is the reason for the operations state. For instance, when in the degraded mode the operational reason could be Peer Overload, or static rate Exceeded. When in the Inactive state the reason could be TTP disabled or SMS service degraded.
	Current Abatement Algorithm	Loss, Rate, or NA	Y	N/A	N/A	Note that this is only set while the DSR is in an overload state. Otherwise it will be set to Not Applicable (NA).
Current DOIC Status	Normal, Overload, Recovering	Y	N/A	N/A	- Normal means no overload condition. - Overload means that the Current Time to Expire is greater than 0 - Recovering means that the DSR is ramping up the traffic after an overload state has ended.	Current DOIC Status.

Table 2-7 (Cont.) TTP Data Components by Type

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
Current Time to Expire	Time in seconds	Y	N/A	N/A	If this is non-zero then the DSR is in an active DOIC Overload Control State (OCS) requested by the DOIC Reporting Node. The DSR moves from the "Overload" to the "Recovering" state when either this timer expires, or the Host sends a time of "0". Time is shown in seconds since it can only be set in seconds.	Current Time to Expire.
Current Loss Rate	0-100% loss	Y	N/A	N/A	From the OC-OLR when using the Loss abatement algorithm.	Current Loss Rate.
Priority 1/Color 1 OTR	OTR in messages per second	N	N/A	N/A	Offered Transaction Rate (OTR) not EMR since these are requests only. There's no need to display the breakdown of rate by color/priority.	Priority 1/Color 1 OTR.
.....		N	N/A	N/A		
Priority X/Color X OTR	OTR in Messages per Second	N	N/A	N/A		Priority X/Color X OTR.
Total OTR	OTR in Messages per Second	Y	N/A	N/A	Only the total across the different Colors or Priorities needs to be displayed.	Total OTR

Table 2-7 (Cont.) TTP Data Components by Type

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
Target ETR	ETR in Messages per Second	Y	N/A	N/A	The current Max Egress Target Transaction Rate. Normally this is the configured Max ETR for this TTP. But this number is reduced to reflect the current loss rate during overload.	Target ETR
Percentage of Transactions Diverted	0-100%	Y	N/A	N/A	The percentage of the OTR for this TTP that's being diverted due to overload.	Percentage of Transactions Diverted.

The following table shows the items in the TTP configuration set:

Table 2-8 TTP Configuration Set Components

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
Scope	Configuration Set Name	Customer Defined	Y	Y	N/A	The text string name for this configuration set.
Configuration	Abatement Recovery Rate	1-100%/second	Y	Y	5%/Sec	The rate at which the DSR goes from the requested loss to zero abatement after an OLR expires. If the current requested loss is -20%, then the DSR decreases the loss linearly from -20 % to 0 at this rate.

Table 2-8 (Cont.) TTP Configuration Set Components

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
	Supported Abatement Algorithms	Loss	Y	Y	Loss	This is the list of abatement algorithms sent in the OC-feature-vector to the Reporting Node. It is configured at the TTP configuration set level since the customer may not want to allow all of the supported algorithms on a given TTP.
	Default OC-Validity-Duration	0-86,400 seconds	Y	Y	30 seconds	This is the default time for the OC-Validity if a time isn't specified in an OLR. Note that 0 in an OLR means stop abating. The suggested default here of 30 seconds is from the DOIC spec.

Table 2-8 (Cont.) TTP Configuration Set Components

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
	Rate Convergence Time	250-2000ms	Y	Y	1000ms	This parameter controls the sensitivity of the calculated rate to bursts of traffic on the TTP. The ETR calculated by the Sliding Historic Metric is always normalized to 1 second (as per the DOIC spec), but the DOIC specification specifically allows for the rate to be higher within that second as long as the per-second average is maintained.
	Dynamic Throttling Override Message Priority Threshold	Priority 1-2	N	N	N/A	Messages with this priority or higher will be routed at the TTP level even if routing them will cause the TTP to exceed a requested abatement level, as long as the message rate is below the TTP Static Throttle Rate. A priority of 0 is not allowed since then the flag effectively disables DOIC dynamic throttling.

2.6.15.10.7 DOIC Interaction with Routing

As mentioned above, abatement includes both throttling and diversion. The additional data collected at the TTP level for the DOIC throttling is also used to improve routing decisions, this includes:

- If all Hosts on a DSR for a given AppID are congested, do not send traffic to the mate DSR if it's also congested.
- Between Peers in a Route Group distribute traffic by both static capacity (also reflecting availability), and by DOIC loss level.
- When selecting Route Groups in a Route List, "bypass" Route Groups that exceed a loss threshold (equivalent to "minimum Route Group weight").
- When selecting a connection or a peer within a Route Group, take into account the overload condition of the underlying TTPs.
- Provide distinct Diameter error leg (user configurable) for requests that failed due to DOIC congestion.

In general all of the existing DSR routing capabilities remain unchanged, but the following optional capabilities are added:

- At the Route List it is possible to "skip" Route Groups that have too high of a DOIC loss rate.
- Within Peer Route Groups traffic is balanced between peers by both their static weighting (existing functionality) and their DOIC loss rates.
- Within Connection Route Groups traffic is balanced between connections by both their static weighting (existing functionality) and their DOIC loss rates.
- Implicit Routing first looks for a matching TTP (more specific, since it's both HostID and an AppID) before it looks for a matching Peer. There is also a new Alternate Implicit Route List associated with a TTP.
- A new error leg for DOIC congestion is defined for when all of the Route Groups in a Route List are skipped due to not meeting the DOIC loss rate cutoff, or when a request message is rejected at the TTP level due to DOIC throttling. This same error code is also used for requests blocked at the TTP level by the Priority/Color algorithm rejecting requests to meet a DOIC abatement request.


Note

The DOIC capabilities layer on top of the existing routing functions. Unless specifically noted, all of the current DSR routing functions continue to operate.

2.6.15.10.7.1 Traffic Throttling Group

To make routing decisions at the Route List level it's necessary to aggregate some of the individual TTP-level data into data that represents the entire Route Group. This summary data is called a Traffic Throttling Group (TTG). There is one parameter summarized in the TTG:

The calculated Loss % of the TTG. This is the weighted (by Max ETR) average of the % loss in the available TTPs.

Figure 2-32 TTP to TTG Rollup Logic

A TTG is a B-scoped item, containing only TTPs from the same DSR. However, it is useful to be able to share TTGs between DSRs. For instance, for the local DSR to decide whether it's a good idea to send traffic it can't handle to the mate DSR, the local DSR needs to know the congestion status of the mate for that particular AppID. To prevent more split-scoped data, the DSR allows the user to define at the B-level which TTGs should be shared between DSRs. The NOAM is then responsible for distributing that list of TTGs to the other DSRs.

2.6.15.10.7.2 TTG Configuration and Status Data

The table below lists the data required in the TTG:

Table 2-9 TTG Data Components by Type

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
Scope	Site Name	DSR Node Name	Y	Y	N/A	Since TTGs can be shared across DSRs the DSR that owns this TTG is part of the key to the record.
	TTG Name	Text field	Y	Y	N/A	Name of the TTG.

Table 2-9 (Cont.) TTG Data Components by Type

Type of Data	Data	Values	Displayed?	Mandatory	Default	Comments
	AppID	AppID name	Y	Y	N/A	The application ID associated with this TTG. This field is used by the DRL to determine whether the TTG is applicable to a request message being routing, and by the GUI to determine which TTPs can be assigned to the TTG.
State	Admin State	Enabled, Disabled	Y	N/A	N/A	Whether the TTG is active and can be used for routing.
	Current Loss %	0-100G	Y	Y	N/A	The current % loss of the TTG calculated.
Configuration	TTP List	List of TTPs	Y	Y	N/A	List of the TTPs assigned to the TTG. Note that all of the TTGs assigned to a TTP must match the Application ID assigned to the TTG.

2.6.15.10.7.3 Congestion-Aware Route Lists

The DSR can use the congestion information in the TTGs to skip Route Groups in the Route List that do not meet threshold criteria for their congestion status. A typical use for skipping congested Route Groups is to prevent a DSR that can't handle traffic itself due to congestion from sending that traffic to a mate DSR that is just as overloaded already.


When defining a Route List two new optional parameters are part of each Route Group:

- The TTG data associated with that Route Group.
- A threshold for the maximum acceptable loss before that Route Group is skipped.


Logically the new threshold functions just like the current "Minimum Route Group Availability Weight", in that it causes the Route List to skip to the priority Route Group.

The figure below shows an example of the Route Group provisioning within a Route List, and the logic used to evaluate that data. There is an action associated with the Routing Option Set called the Routing Congestion Action. This is the action that is taken if all of the Route Groups in the Route List are skipped because they didn't meet the congestion thresholds. Note that the existing action is taken if the Route List successfully selects any of the Route Groups, but then fails to route the request anyway. Like all ROS actions, the Routing Congestion action allows the user to specify whether the answer should be abandoned, or if an error answer is set, what the error should be. This new congestion-related error leg can be used to return an experimental error number indicating a congestion failure.


Figure 2-33 Congestion-Aware Route List Logic



The figure below gives an example of using these thresholds. In DSR 1 there are two Route Groups that can handle S6b initial requests, a primary one (RG1) and a secondary one (RG2). If neither of those Route Groups can handle the request, then the customer wants to try sending the requests to the mate DSR, DSR 2. Since RG1 and RG2 are local to the DSR, they have corresponding TTGs, DSR1_S6b and DSR1_2_S6b. However, to send S6b traffic to the mate DSR there isn't a dedicated S6b Route Group, just RG3 which represents the connections to the mate. The TTG associated with RG3 in the Route List is then the TTG of the target resources on the mate, in this case TTG DSR2_S6b. This works because when the DSR is deciding whether or not to route traffic to the mate, it doesn't care about the congestion status of the route group to get the DSR (that will be handled by ETGs for instance), it cares whether the S6b handling resources on the mate are congested or not.

Figure 2-34 Congestion-Aware Route List Example 1

The next figure shows the evaluation logic for this example. Since the TTG currently has a higher loss (90%) than the threshold for RG1 (Max Loss % = 80%), the Route List skips RG1 and goes directly to evaluating RG2. Since the loss and rate thresholds for RG2 are acceptable, the RL sends the request to RG2.

Figure 2-35 Congestion-Aware Route List Logic Example 1a

The next figure shows what happens when none of the Route Groups meets their minimum threshold. In this case all three Route Groups are skipped, and the action defined in the new “Routing Congestion Action” in the Routing Option Set is executed. Like the other actions defined in the ROS, this new action can either abandon the answer, or return a user-configured

error number, error text and vendor ID. This action is only taken when the request skips all of the available Route Groups due to not meeting the Max % Loss threshold. If no Route Group is found due to other reasons, such as a Route Group was selected, but it couldn't handle the request, then all of the existing error legs are used as appropriate.

Figure 2-36 Congestion-Aware Route List Logic Example 1b



2.6.15.10.7.4 Interaction of DOIC within Route Groups

When a DSR receives a DOIC Overload Report with an abatement request for a given TTP, the DSR follows the abatement percentage regardless of how traffic is routed to that TTP: through a peer route group, a connection route group, or via implicit routing.

2.6.15.10.8 DOIC Override Flag

The DOIC override flag is an attribute of the TTP configuration set. It is a single message priority threshold so that all request messages of that priority and above are given the DOIC Override handling. For instance, if this threshold is set for Priority 2 messages, then all Priority 2 messages are given this treatment. If the threshold is set for Priority 1 messages, then both Priority 1 and Priority 2 messages receive this treatment.

2.6.15.10.8.1 Handling of the DOIC Override Flag

- Priorities still apply within the flagged messages. For instance, if the flagged level is set to priority 1, then priority 1 messages continue to be discarded before Priority 2 messages.
- The DOIC flag overrides color. Thus if the algorithm is set to “discard by priority within color”, but have the “DOIC Override” threshold set for priority 2 messages, then the sort order (low priority to high priority) would be:
 - Yellow/priority 0
 - Yellow/priority 1
 - Green/priority 0
 - Green/priority 1

- Yellow/priority 2/DOIC Override
- Green/priority 2/DOIC Override

2.7 Next Generation Network Priority Service


New Generation Network Priority-Service (NGN-PS) is a US Presidentially-directed program to define and deploy priority communications in commercial service provider next generation packet-switched networks. First Responders, National Security, and Emergency Preparedness need priority access and priority services to better serve the public interest when the network is facing severe overload conditions.

Despite the original directive of NGN-PS as mentioned above, this feature can be used by any operator to prioritize VoLTE traffic based on subscription information stored in the HSS or SPR.

The DSR assigns all NGN-PS messages higher priority over non NGN-PS messages (request and answers). The DSR then allows control for the NGN-PS feature on a per Diameter interface basis and for any combination of the following LTE and IMS diameter interfaces: Gx, Rx, Cx/Dx and Dh/Sh. The DSR supports optional NGN-PS/Advance-Priority HSS and NGN-PS/Advance-Priority SPR modes for Gx interface as specified in LTE_GIR and IMS_GIR specifications. Additionally, DSR also allows to tag the ingress messages as NGN-PS messages using DRMP AVP value "0" for any diameter interface. This allows operators to support NGN-PS for diameter interfaces which do not have 3GPP defined AVP's to identify the message as NGN-PS message.

If required the DSR throttles normal traffic to an engineered capacity such that internal resources (related to message queuing, resource pools, and message control functions) are always available for a small amount of additional NGN-PS traffic. Specific alarms and measurements allow evaluation of NGN-PS transactions.

Figure 2-37 NGN-PS Support



2.8 IP Front End

The presence of IPFE does not prevent a system from having DA MPs directly connected to clients using for example SCTP Multi-homing connections.

The IP Front End (IPFE) is a traffic distributor that transparently does the following:

- Presents a routable IP address representing a set of up to 16 application servers to application clients. This reduces the number of addresses with which the clients need to be configured.
- Routes packets from the clients that establish new TCP or SCTP connections to selected application servers.
- Routes packets in existing TCP or SCTP connections to the correct servers for the connection.

2.8.1 Traffic Distribution

The IPFE presents one or more externally routable IP addresses to accept TCP or SCTP traffic from clients. These externally visible addresses are known as Target Set Addresses (TSAs). Each TSA has an associated set of IP addresses for application servers, up to 16 addresses, known as a Target Set. The IP addresses in a given Target Set are of the same IP version (that is, IPv4 or IPv6) as the associated TSA.

A typical client is configured to send TCP or SCTP traffic to one or more of the TSAs, rather than directly to an application server. When the IPFE receives a packet at a TSA, it first checks to see if it has a transaction state that associates the packet's source address and port to a particular application server.

This state is known as an "association." If no such association exists (that is, the packet was an "initial" packet), the IPFE runs a selection function (which has been configured by the user selecting a method such as hash, least load, peer node aware least load, so on.) to choose an application server address from the eligible addresses in the Target Set. The selection function uses a configurable weighting factor when selecting the target address from the list of eligible addresses. The IPFE routes the packet to the selected address, and creates an association mapping the source address and port to the selected address. When future packets arrive with the same source address and port, the IPFE routes them to the same selected address according to the association.

Because the IPFE has no visibility into the transaction state between client and application server, it cannot know if an association no longer represents an active connection. The IPFE makes available a per Target Set configuration parameter, known as delete age, that specifies the elapse of time after which an association is to be deleted. The IPFE treats packets that had their associations deleted as new packets and runs the application server selection function for them. The IPFE sees only packets sent from client to server. Return traffic from server to client bypasses the IPFE for performance reasons. However, the client's TCP or SCTP stack "sees" only one address for the TSA; that is, it sends all traffic to the TSA, and perceives all return traffic as coming from the TSA.

The IPFE neither interprets nor modifies anything in the TCP or SCTP payload. The IPFE also does not maintain TCP or SCTP state, per se, but keeps sufficient state to route all packets for a particular session to the same application server.

In high-availability configurations, four IPFEs may be deployed as two mated pairs, with each pair sharing TSAs and Target Sets. The mated pairs share sufficient state so that they may identically route any client packet sent to a given TSA.

The IPFE supports the following types of DSR Diameter connections:


- Responder Only
- Initiator Only
- Initiator and Responder

Support for the IPFE initiator or responder connections removes the need for roaming partners to negotiate Initiator / Responder responsibilities. DSR initiates and listens for Diameter connections on a single connection using shared IPFE signaling IP addresses. The DSR provides a system wide distributed connection election algorithm to resolve race conditions between IPFE initiator and responder state machine instances.

The DSR currently allows up to 1 IPFE 'initiator+responder' per TSA per peer node. If there are more than 1 TSA per DSR, each TSA can be associated with 1 'initiator+responder' connection. Please note that this can co-exist 'initiator only' or 'responder only' connections to the same Peer node. In the case of an election, one of the two connections shuts down.

- Local Node FQDN > Peer Node FQDN = responder connection survives.
- Local Node FQDN < Peer Node FQDN = initiator connection survives.
- All subsequent messages are sent on the surviving connection.

Figure 2-38 IPFE Initiator or Responder Support



Connection Balancing

Under normal operation, the IPFE distributes connections among application servers according to the weighting factors defined in the Target Sets. However, certain failure and recovery scenarios can result in an application server having significantly more or fewer connections than is intended by its weighting factor. The IPFE considers the system to be "out of balance" if this discrepancy is so large that the overall system cannot reach its rated capacity even though individual application servers still have capacity to spare, or so that a second failure is likely to cause one of the remaining servers to become overloaded. The IPFE determines this by measuring the number of packets sent to each server and applying a "balance" heuristic.

When the IPFE detects that the system is out of balance, it sets an alarm and directs any new connections to under loaded application servers to relieve the imbalance. There are a few types of connection distribution algorithms that can be used: hash, least load, and peer node group aware least load distribution.

2.8.2 High availability

When paired with another IPFE instance and configured with at least two Target Set Addresses, the IPFE supports high availability. In the case of an IPFE pair and two Target Set Addresses, each IPFE is configured to handle one Target Set Address. Each IPFE is automatically aware of the ruleset for the secondary Target Set Address. If one IPFE should become unavailable, the other IPFE becomes active for the failed IPFE's Target Set Address while continuing to handle its own.

In the case of an IPFE pair, but only one Target Set Address, then one IPFE is active for the Target Set Address and the other is standby.

2.9 RADIUS Signaling Router

The RADIUS Signaling Router feature of the DSR covers RADIUS message routing (RADIUS message in and RADIUS message out) without Diameter interworking. RSR supports RADIUS over UDP transport. RADIUS signaling is handled by a DA-MP instance. Both Diameter and RADIUS connections can be hosted on a DA-MP. Some examples where RSR may be used are:


- WLAN authentication and authorization, both in non-roaming and roaming cases.
- For a GGSN to authenticate a user and provide (accounting) information to an AAA server on a per APN basis.
- Policy control in broadband networks via dynamic authorization mechanisms.

Base RADIUS support on the DSR, aka RSR, is provided in a way that the system may or may not also be handling Diameter traffic. On a DSR with RSR, a message could come in as either RADIUS or Diameter and egress in the same protocol as it ingressed.

2.9.1 RADIUS Routing

An example of the RADIUS proxy use case is for WiFi roaming. The figure below shows the interfaces between a visited and home wi-fi service provider network. As seen in the figure, RADIUS authentication and authorization is used between these networks.

Figure 2-39 RADIUS Interfaces in WLAN Roaming Architecture



2.9.2 RADIUS Overload Control

RADIUS may be used as an option, on a per APN basis, for a GGSN to use RADIUS authentication to authenticate a user and RADIUS accounting to provide information to a AAA server. There have been cases where RADIUS servers have become overloaded due to excessive traffic load and AAA networks have gone down.


With RADIUS there is no message defined for a server to tell a client that it is currently experiencing overload. Messages simply time out, and clients need to retransmit, which has a tendency to make the problem in the network worse and potentially spread to other servers.

To address this problem, a RADIUS router or proxy can be used between the clients and servers to control the amount of traffic presented to the servers and prevent server outages. Additionally, in the event there was an outage of one or more servers, such a centralized point of traffic control could help ensure a smooth re-introduction of active servers to the network.

2.9.3 RADIUS Message Format

The figure below shows the basic RADIUS message data format.

Figure 2-40 RADIUS Data Format



2.9.4 Authenticator

RADIUS clients and servers share a secret (password). This means RSR shares a secret with every peer node. This same secret is used for multiple peer nodes. Various fields and methods are used for authentication, accounting, COA and disconnect messages to pass along and verify the secret.

2.9.5 Message Authenticator

A Message-Authenticator attribute (different from the Authenticator in the RADIUS packet header) is used to authenticate and integrity-protect RADIUS packets in order to prevent spoofing.

A server or client receiving a message with a Message Authenticator attribute present must calculate the expected value of the message authenticator and silently discard the packet if it does not match the value sent.

2.9.6 Connections and Peers

Despite the fact that RADIUS uses the connectionless UDP transport, the concept of a RADIUS connection is helpful to facilitate understanding RSR operation. From the point of view of RADIUS peers connecting to DSR, the DSR can act either as a RADIUS server (to RADIUS clients) or as a RADIUS client (towards RADIUS servers). We can generally think of a RADIUS peer as defined by its IP address and (optional) port and a RADIUS connection as an association of source IP address + (optional) ports and recipient IP address + port, where either the source or the recipient would be represented by DSR.

In RADIUS, specific recipient ports are typically associated with specific services - for example, Authentication, Accounting, and Change of Authorization would each have their own distinct ports. This means that on a given connection, requests always flow in one direction and the responses in the other. There are two types of DSR RADIUS connections:

- Client connection: remote IP + port combined with local (DSR) IP + port range. These are connections towards servers.

- Server connection: remote IP combined with local (DSR) IP + port. These are the connections towards clients.

2.9.7 Routing and Load-balancing

Base RADIUS routing utilizes the same mechanisms as Diameter by means of encapsulating the RADIUS messages in a Diameter wrapper. Specific Diameter AVPs in the Diameter wrapper are created based on information from the corresponding RADIUS message and configuration information. The Diameter wrapper contains the following AVPs and Diameter header information which is then used to route the Diameter wrapper using existing Diameter routing mechanisms.

Table 2-10 RADIUS Message Mapping

Diameter AVP	RADIUS Message Type	Value
Application-ID	Any	Derived based on configured mapping from RADIUS command code.
Command-Code	Any	Derived based on configured mapping from RADIUS command code.
Origin-Realm	CoA or Disconnect	Derived via configured mapping from ingress peer.
	Not CoA or Disconnect	Derived via configured mapping from NAS-Identifier, NAS-IP(v6)-Address, or ingress peer.
Origin-Host	COA or Disconnect	Derived via configured mapping from ingress peer.
	Not CoA or Disconnect	Derived via configured mapping from NAS-Identifier, NAS-IP(v6)-Address, or ingress peer.
Destination-Realm	CoA or Disconnect	Derived via configured mapping from NAS-Identifier, NAS-IP(v6)-Address, or ingress peer.
	Not CoA or Disconnect	Domain part of User-Name, if possible, else same value as Origin-Realm (see above).
Destination-Host	CoA or Disconnect	Derived via configured mapping from NAS-Identifier or NAS-IP(v6)-Address.
	Not CoA or Disconnect	Destination-Host omitted.

Load-balancing in RADIUS is the same as for Diameter with respect to route groups, weighted loadsharing, so on. Only Peer Route Groups are supported for RADIUS peers – Connection Route Groups are not applicable.

As RSR uses the Diameter routing mechanisms, all Diameter routing capabilities can be used to route RADIUS messages. For instance, if a response is not received in a timely manner after a RADIUS request is forwarded to a RADIUS server, RSR can resend the request a configurable number of times, and if a response is still not received, the request can be routed to an alternate server. RSR does not support the Diameter "alternate routing on answer" capability. RSR supports receipt of Status-Server message from RADIUS clients. RSR can be configured to respond to Status-Server with either Access-Accept or Accounting-Response.

2.9.8 Duplicate Detection

According to the base RADIUS specification, any message received within a short span of time with the same client source IP address, source UDP port, Authenticator, and Identifier is considered a duplicate request.

DSR detects duplicate requests received from clients and in such cases, avoids sending duplicate requests to servers. DSR supports retransmission of a request to the same connection a user configurable number of times. Such retransmissions contain the same source IP address, source UDP port, Authenticator, and Identifier value to allow the server to detect retransmitted requests.

2.9.9 Message / Traffic Control


The following features work for RADIUS the same as the do for Diameter:

- Per-connection ingress message control.
- Egress throttle groups (Egress message rate limiting and Global egress request window limiting are supported).
- Per-connection egress message throttling (Egress message rate limiting and Egress request window limiting are supported).

2.10 RADIUS-Diameter IWF for Authentication

The RADIUS-Diameter Interworking (R-D IWF) for Authentication feature provides message conversion and interworking between a RADIUS based client (server) and a Diameter based server (client). An example is shown below where RADIUS authentication and accounting is used by a WLAN AP, but the AAA server is Diameter based.

Figure 2-41 RADIUS-Diameter IWF for WLAN Authentication




The Figure below shows a generic case for deployment of a mated pair of DSRs with RSR and/or R-D IWF capability. This figure shows the case where a single pair of DSR IWFs is serving a many-to-many relationship between RADIUS and Diameter networks. The blue lines in the figure depict RADIUS connections and the redlines depict Diameter connections/ connection sets. The following routing options are supported:

- RADIUS net to same RADIUS net.
- RADIUS net to different RADIUS net.


- RADIUS net to Diameter net.
- Diameter net to same Diameter net.
- Diameter net to different Diameter net.
- Diameter net to RADIUS net.

Figure 2-42 RSR and R-D IWF Deployment



As shown in the above figure, Diameter transport is planned for the 'c-links' between DSR mates, even for RADIUS messages. This is possible because RADIUS messages are encapsulated within a Diameter shell for internal routing within a DSR.

Figure 2-43 RADIUS-Diameter IWF for Authentication



Supported mappings include:

- RADIUS Access-Request - Diameter DER.
- Diameter DEA - RADIUS Access-Challenge.
- Diameter DEA - RADIUS Access-Accept.
- Diameter DEA - RADIUS Access-Reject.

2.11 Diameter Mediation

The Diameter Protocol has been designed with extensibility in mind. Standards bodies have defined quite a few applications on top of the base Diameter protocol for use in 3G, LTE and IMS networks. Over time, the standards bodies will continue to extend these applications by adding, altering or deleting AVPs or modifying the header to meet new market needs.

In an effort to differentiate themselves, Vendors often include additional functionality into the protocol by adding proprietary AVPs or overloading existing AVPs. Such additions do not pose an interoperability issue where all the equipment is provided by a single vendor, but that is rarely the case. As most operators rely on equipment from multiple vendors, interoperability issues are almost guaranteed. To make matters worse, vendors continue to extend their proprietary versions of the protocol making them incompatible with other elements that communicate using the previous version of the proprietary protocol.

Even in the absence of vendor-specific extensions, it is possible that two vendors interpret the standard in slightly different ways which could then lead to interoperability issues. The operator can mitigate this by forcing the two vendors to perform interoperability testing prior to deployment. However, in certain scenarios, such as the S9interface (HPCRF-VPCRF), where two operator networks have to exchange Diameter traffic between each other, performing interoperability exercises with all other operator networks is not practical.

Operators may choose to deploy components of a solution in a phased manner. For example, an operator can start with just the charging and billing systems and roll in the policy control parts of the solution at a later time. As new components are added to the solution, operators will have to ensure that these new components work seamlessly with the existing setup. In such situations, operators often see a need for performing activities such as Digit Manipulation or mapping of Result-Codes.

Therefore, as Diameter networks get more complex, inter-operability issues in a multi-vendor environment or interoperator Diameter traffic exchange could pose challenges. Also as new components are added to the solution, operators will have to ensure that these new components work seamlessly with the existing setup.

The Diameter Mediation feature offers an intuitive GUI that can be used by the operator to build mediation rules to resolve inter-operability issues. This logic can be seamlessly applied to all messages transiting the DSR. As an example, the mediation feature can be utilized by the customer for topology hiding. Operators often desire to hide the topology details of their network for protection purposes and for seamless interworking functionality. The customer is able to use the provided mediation framework to create the necessary rules that would implement topology hiding in their network. In addition mediation enables the DSR to route based on session-id. This is done by using the hashing mechanism to identify messages with matching session-ids that are then all configured to go to the same host.

2.11.1 Rule Templates and Rules

Upon identifying the need for message mediation, an operator begins by creating a "Rule Template". A Rule Template includes the logic required to perform a specific mediation.

Conditions and Actions are defined as part of the template and then the rule template is associated with one or more Trigger Points (defined below). Once the definition is complete, the operator provisions the data (Rules) needed for the conditions and the actions. An operator can provision up to 2000 Rules per Rule Template and 2000 counters that can be used for measurements based on message content. To ensure system performance is not impacted, the total number of rules across all rule sets combined should not exceed 3750.

The Rule Template allows for up to 5 conditions and 5 actions to be defined in a template. When multiple conditions are present in a Rule Template, the framework allows the conditions to be combined using the logical operators (AND, OR) and also the order in which the actions must be executed.

Some examples of the conditions supported are:

- Checking for the presence or absence of well-known or proprietary AVPs.
- Checking for the value of AVP header components or data part of well-known or proprietary AVPs.
- Checking for any other component of an AVP such as AVP flag.
- Checking for any component of the Diameter message (flags, appl-id, cmd-code, so on).
- Checking for ranges.
- Checking for peer and connection names/ids.
- Checking for message priority.
- Checking for bit set/reset.
- Checking if a message has been redirected.

Some examples of the actions supported are:

- Adding or deleting AVPs.
- Modifying parts of AVP header.
- Modifying the Diameter header.
- Set a message priority.
- Activate message copy.
- Set alarm/event.
- User defined measurements associated with the use of measurement rules.
- Redirect a message.
- Parse decorated NAI.
- Peg a mediation framework counter.

Both actions and conditions can be applied to Grouped AVPs. A max depth of 8 is supported for the Grouped AVPs.

Rule Templates and their associated Rules can be independently exported on one system (such as a lab system) and then imported into another system (such as a production system). This capability is useful when the Rule Templates and Rules are being tested in a lab environment and for moving the Rule templates and Rules to production system upon successful completion of testing. The import and export all comes in handy when a Rule Template has to be updated and replaced with a newer version of the Rule Template but the older Rules need to be preserved.

2.11.2 States of a Rule Template


A Rule Template is in one of three states at any point in time. These states are Development, Test and Active. Each Mediation Template begins in the “Development” state when created. Once the template definition is complete the State can be changed to “Test” or “Active”. An operator can provision rules (data) against the Template only after a Template is in the “Test” or “Active” states. In the “Test” state, the template logic is executed for Requests arriving on “test” connections. (See connections GUI to designate a connection as a “test” connection). However, only Requests (not answers) can be processed in this state and so it is recommended to test the Templates by placing them in an “Active” state but on a lab system prior to moving into production. Upon successful execution of tests in the lab system, the templates and the associated rules (if applicable) can be imported to the production system and the state of the Mediation Template can be changed to “Active” by the operator.. If the execution of tests is unsuccessful, the Mediation Template can be transitioned back into the “Development” state where it can be altered and the process is repeated. It should be noted that rules cannot be associated with a template in “development” state and hence it is recommended to export the rules associated with the template prior to this operation to avoid the need of manually configuring the rules again.

Mediation provides the KPI's to monitor the CPU utilization of threads/processes that are impacted by the use of Mediation rules. These KPI's can be used to assess the performance impact of using Mediation rules before and after Mediation rules are deployed. Mediation user's guide explains the recommended “safe” ranges for these KPIs under which mediation rules can be deployed without adversely impacting the performance of the DSR.

2.11.3 Trigger Points

Trigger points are specific points in call processing where the Rule Templates along with their associated Rules can be executed. The trigger points supported in the DSR are:

- Upon receipt of a Request (including a Redirected Request) (including CER, DWR, DPR).
- Prior to relay or proxy or sending of the Request (including CER, DWR, DPR).
- Prior to forwarding a re-routed Request.
- Upon receipt of an Answer (including CEA, DWA, DPA).
- Just prior to forwarding/sending the Answer downstream. (including CEA, DWA, DPA)
- Just prior to the invocation of an application.
- Immediately after the Request exits the application.
- Just prior to the Answer being routed to the application.
- Immediately after the answer exits the application.

Figure 2-44 Mediation Trigger Points

The mediation framework also supports defining multiple mediation rules at a single trigger point or invoking the same mediation rule at multiple trigger points.

2.11.4 Measurements Associated with Rule

In order to allow an operator to see how many times a rule is invoked for debugging purposes or for fine tuning purposes, rule counts are maintained for the rules in a rule set. These counts can be enabled/disabled as a property of the template and once enabled the counters appear against the individual rules in the rule set (that is there is one counter per each rule in the rule set.) These counters track the number of times a rule is successfully matched on all the conditions in the template. The counters are based on conditions only and the outcomes of the actions do not impact the counters. They are incremented sequentially until they are disabled.

2.11.5 AVP Dictionaries

The GUI driven definition is much simplified by using AVP names instead of AVP codes wherever possible. The Diameter Mediation Framework includes a Base AVP Dictionary where well known AVPs are defined. This dictionary includes AVPs defined in the base Diameter Protocol and AVPs defined by popular applications such as Diameter Credit Control Application, and S6a interface. Any additions made by the operator are included into the Custom AVP Dictionary. Once defined, these AVPs are available for use by their name during rule template definition.

A grouped or non-grouped AVP defined in the base dictionary or in the custom dictionary can be cloned, modified and saved into the customer dictionary. An AVP cannot be saved if the combination of the same AVP code and/or AVP name already exists in the custom dictionary. If the user clones an AVP that is referred from some template/rule, then the GUI only allows adding new sub AVPs to the grouped AVP, no other changes are allowed. If the AVP is not used by any template/rule, the user can do other modifications.

2.12 Topology Hiding

In various interworking scenarios LTE service providers need to protect their networks. The Topology Hiding features remove or hide all Diameter addresses from messages being routed out of the home network on connections with this feature enabled. This feature also re-inserts the appropriate addresses in messages coming back into the home network on these connections. In addition, peer networks are prevented from determining the topology of the


home service provider's network by obscuring the number of host names in the network. As a result of this, the peer network service provider is not able to determine how many MME/SGSNs, HSSs, PCRF, AFs, and pCSCFs are deployed. Nor can the peer service providers derive any deployment architecture information through inspection of host names.

2.12.1 S6a/S6d Topology Hiding

In S6a/S6d transactions, a host name sent by the MME/SGSN in the Origin-Host AVP in a ULR message is saved by the HSS and used in the Destination-Host AVP for requests, such as the CLR, sent by the HSS. The figure below shows this linking of host names across Diameter transactions. As a result of this, it is necessary to ensure that a DSR receiving a CLR request from an untrusted peer network HSS can determine which MME/SGSN host is the target of the request.

With this approach, there is a configured mapping of real MME/SGSN host names to MME/SGSN pseudo-host names. When a request or answer associated with a protected network is forwarded towards an untrusted peer network, the MME/SGSN host name in the message is replaced by a MME/SGSN pseudo-host name. When a request or answer is received by a DSR with TH enabled on the ingress Peer Node and it contains a MME/SGSN pseudo-host name, the MME/SGSN pseudo-host name is replaced by the real MME/SGSN host name.

Figure 2-45 MME/SGSN Topology Hiding



The MME/SGSN topology hiding feature also hides the number of MME/SGSNs in the protected network. To achieve this requirement the MME/SGSN Topology Hiding feature allows for the mapping of a variable number of MME/SGSN pseudo-host names per real MME/SGSN host name. For details on the configuration of the host names, see S6a/S6d Configuration.

The algorithm for selection of the MME/SGSN pseudo-host name ensures that the same MME/SGSN pseudo-host name is always selected for the same IMSI from the same MME/SGSN. This is to ensure that the HSS receiving a ULR doesn't mistakenly think that the request is from a new MME/SGSN, triggering a CLR transaction. The MME/SGSN topology hiding feature also hides the host names included as part of the Session-Id AVP.

S6a/S6d HSS Topology Hiding

The S6a/S6d HSS topology hiding feature applies to all Diameter S6a/S6d messages between a protected network HSS and an untrusted peer network MME/SGSN. The HSS topology hiding feature also hides the number of HSSs in the protected network. To achieve this requirement the HSS Topology Hiding feature allows for the mapping of a variable number of HSS pseudo-host names per real HSS host name. For details on the configuration of the host names, see S6a/S6d Configuration.

For Diameter transactions originated by an MME/SGSN in an untrusted peer network, the following actions are taken for S6a/S6d HSS Topology Hiding:

- Request Messages – If the request message contains the Destination-Host address of S6a/S6d HSS and if HSS pseudo-name was selected from a list of HSS pseudo-names in previous S6a/S6d HSS Answer, then S6a/S6d HSS Topology Hiding restores the original S6a/S6d HSS addresses in the Destination-Host AVP. Restoration of Protected S6a/S6d HSS original host name is not done if single pseudo-name is used in S6a/S6d HSS Topology Hiding. Instead this replacement is done by HSS Address resolution application such as DSR's FABR or RBAR application.
- Answer Messages – The answer message contains the HSS real host name in the Origin-Host AVP. This real host name is replaced based on one of the following 2 methods for HSS pseudo host name selection.
 - A single HSS pseudo-host name which has been defined for all the network HSS real host names in the protected Network.
 - A HSS pseudo-host name selected from a list of HSS pseudo-host names that have been defined for each real HSS host name in the Protected Network (this approach is similar to the one described for MME/SGSN Topology Hiding).
For Diameter transactions originated by the protected network HSS and targeted for an untrusted peer network MME/SGSN the following actions must be taken for S6a/S6d HSS Topology Hiding.
- Request Messages –
 - The request message contains the HSS real host name in the Origin-Host AVP. Based on which HSS pseudo-host name selection method has been selected (as described above), this host name is replaced with either the single HSS pseudo-host name defined for all HSS real host names in the protected network, or by a HSS pseudo-host name from the list of HSS pseudo host names defined for each of the Protected Network real HSS host names.
 - The request message also contains a Session-Id AVP that contains the HSS's Diameter-ID. Based on which HSS pseudo-host name selection method has been selected (as described above), this HSS real host name is also replaced with either the single HSS pseudo-host name defined for all HSS real hostnames in the protected network, or by a HSS pseudo-host name from the list of HSS pseudo host names defined for each of the Protected Network real HSS host names.
- Answer Messages –
 - The answer message also contains a Session-Id AVP that contains a HSS pseudo host name in the Diameter-ID portion. This is replaced with the HSS real host name stored in the transaction state.

The figures below show message flows illustrating S6a/S6d HSS TH for requests originating at an untrusted peer network MME/SGSN as well as the protected network HSS.

Figure 2-46 S6a/S6d HSS Topology Hiding - ULR Message Flow





Figure 2-47 S6a-S6d HSS Topology Hiding CLR Message Flow

S6a/S6d Configuration

When configuring the Topology Hiding features a GUI is used to input the necessary data. The real host names of the network elements (MME/SGSN/HSSs) are entered. A pattern is entered that is used to generate the pseudo-host names. The DSR then generates from one to three pseudo-host names per entered MME/SGSN/HSS.

The following example is based on an MME/SGSN perspective, but the same configuration applies for HSS topology hiding as well. This example assumes that a carrier has five MME/SGSNs with the following real names:

- mme1.westregion.example.com
- mme2.westregion.example.com
- mme1.eastregion.example.com
- mme2.eastregion.example.com
- mme1.texasregion.example.com

When configuring the topology hiding, the carrier enters these five real MME/SGSN host names. The carrier also enters the pattern to be used in generating the MME/SGSN pseudo-host names. The pattern is in the form:

prefix|digits|suffix

where the variable portion of the name is the digits field. For example, assume the carrier enters the following pattern:

prefix = “mme”

digits = “nnn”

suffix = “.example.com”

The resulting generated names look as follows:

mme|nnn|.example.co

In this case, the nnn portion of the MME/SGSN pseudo-host name contains three digits used to differentiate the MME/SGSN pseudo-host names.

The DSR then generates the mapping between real and pseudo-host names. The following table is an example mapping that could result from this example:

Table 2-11 MME/SGSN PSEUDO-HOST NAME MAPPING EXAMPLE

MME/SGSN Real HostName	MME/SGSN Pseudo-Host Name(s)	
mme1.westregion.example.com	mme042.example.com	m m e 1 2 3 e x a m p l e c o m
mme2.westregion.example.com	mme533.example.com	
mme1.eastregion.example.com	mme922.example.com	
mme2.eastregion.example.com	mme411.example.com	mm mm ee 23 13 81 ee xx aa mm pp ee cc pp mm

Table 2-11 (Cont.) MME/SGSN PSEUDO-HOST NAME MAPPING EXAMPLE

MME/SGSN Real HostName	MME/SGSN Pseudo-Host Name(s)	
mme1.texasregion.example.com	mme776.example.com	mm mm ee 23 98 53 ee xx aa mm pp ee cc pp mm

This mapping is then used for replacing MME/SGSN real host names with MME/SGSN pseudo-host names for messages directed toward the untrusted peer network HSS and for replacing MME/SGSN pseudo-host names with real host names for messages from the untrusted peer network HSS targeted for a protected network MME/SGSN. These same steps are used to create the pseudo-host names for HSSs to support S6a/S6d HSS topology hiding.

2.12.2 Path Topology Hiding

Path Topology Hiding is the most generic form of topology hiding. It is required for Topology Hiding on any Diameter interface type. Path Topology Hiding involves removing Diameter host names from the Route-Record AVPs included in request messages. This feature does more than just Path Topology Hiding. It might be better called Diameter Topology Hiding, as there are host names that are hidden that are beyond just the path recorded in Route-Record AVPs. This feature hides all of the host names included by the base Diameter protocol, with the exception of the Session-Id header, which is left to the TH feature for the specific interface to handle.

Path Topology Hiding also hides addresses in other AVPs that are part of the base Diameter specification. This includes the following:

- The Error-Reporting-Host AVP contains the name of the host that generated an error response. When present, this host name needs to be obscured in answer messages.
- The Proxy-Host which is an embedded AVP within the grouped Proxy-Info AVP contains the name of a proxy that handled a request. This is used as a way for the proxy to insert state into a request message and receive the state back in the answer message. As such, the method for hiding the name of the Proxy-Host name must allow for reconstruction of the name when the answer message is received.

Route-Record Hiding

The Route-Record AVP has two uses in Diameter signaling:

- The primary purpose is to detect loops in the routing of Diameter Request. In this case, a Diameter Relay or Proxy looks at Route-Record AVPs to determine if a message loop has or will occur. This is detected either by the relay or proxy (the DSR in our case) finding its own host-id in the Route-Record message or by the DSR determining

that the host to which the request is to be routed in the Route-Record AVP(referred to as forward loop detection). Note that not all Diameter Relays/Proxies do forward loop detection. The DSR, however, does.

Note

For the purposes of this feature, the definition of a loop is modified slightly to include any time that a Request leaves the home or interworking network and then returns to the home or interworking network. This is independent of the DEA or DIA at which request returns to the home or interworking network. This means that a Request leaving the network on one DEA/DIA and returning to the network on a different DEA/DIA is considered a loop.

- The other defined purpose of the Route-Record AVP is for authorization of the request. A Diameter service might not want to accept a request if it has traveled through a suspect realm. While the DSR does not support such an authorization feature, the Path TH feature does not remove the ability for other Diameter agents or servers to use the Route-Record AVPs to authorize the request.

Each Route-Record AVP contains a Host-Id of a Diameter node that has handled the request. A Relay/ProxyAgent inserts a Route-Record AVP into the message containing the Host-Id of the Diameter node from which it received the request.

It is the Protected Network's Host-Ids included in the Route-Record AVPs that need to be hidden.

For Request messages leaving a protected network, the Path TH feature handles Route-Record AVPs by stripping the protected network's Route-Record AVPs and replacing them with a single Route-Record AVP containing a Route-Record pseudo-host name.

For example, the following request:

```
xxR
...
Route-Record: host1.protectednetwork1.net
Route-Record: host2.protectednetwork1.net
...
```

Would be modified to the following:

```
xxR
...
Route-Record: pseudohost.protectednetwork1.net
...
```

Route-Record AVPs for network other than the Protected Network are preserved. As such, the following request:

```
xxR
...
Route-Record: host.foreign1.net
Route-Record: host.foreign2.net
```

Route-Record: host.protectednetwork1.net

...

Would be modified to the following:

xxR

...

Route-Record: host.foreign1.net

Route-Record: host.foreign2.net

Route-Record: pseudohost.protectednetwork1.net

...

For requests ingressing into a protected network, the Path TH feature examines the Route-Record headers in the request. If any of the Route-Record AVPs contains a host name matching a protected network's Route Record pseudo-host name then the DSR considers it a loop and returns an answer message with Result-Code AVP value 3005 (DIAMETER_LOOP_DETECTED).

It is also necessary to hide the names of hosts that occur in the other base Diameter AVPs listed here:


- Proxy-Host AVP (embedded in the grouped Proxy-Info AVP).
- Error-Reporting-Host AVP.

Proxy-Host Hiding

The handling of the Proxy-Host AVP can be achieved using a pseudo-host name. In this case, the real name is stored in the pending transaction record. The pseudo-host name found in the answer message is replaced by the real host name stored in the pending transaction record. The figure below shows a simple message flow illustrating this functionality.

This handles the instance that multiple proxies are in the path of the request. As a result, a single Proxy-Host pseudo-host name is not sufficient, as the original name is restored when the answer returns. To address this, the DEA/DIA is able to insert a different Proxy-Host pseudo-host name per Proxy-Host AVP. These Proxy-Host pseudo-host names are also generated in a fashion that does not expose the number of proxies in the protected network. In order to achieve this, the Proxy-Host pseudo-host name consists of two components, the user-defined Proxy-Host pseudo-host name string and a random set of 3-digits prefixed to that name. If the user-defined Proxy-Host pseudo-host name string is proxy.example.com, then the value inserted into a Proxy-Host AVP would then be of the form nnnproxy.example.com, where "nnn" is a randomly generated set of digits.

Figure 2-48 Proxy-Host Topology Hiding Message Flow



Error-Reporting-Host Hiding

When obscuring the Error-Reporting-Host AVP the real host name is recovered in case it is needed for troubleshooting activities. Encryption is used for obscuring the Error-Reporting-Host AVP. This allows for troubleshooters in the protected network to decrypt the AVP to determine the original value. The encryption algorithm used only requires the operator to know the key for decrypting this value in a common troubleshooting tool such as Wireshark.

2.12.3 S9 PCRF Topology Hiding

S9 PCRF topology hiding is concerned with hiding the identity of a Protected Network's PCRFs, as well as the number of PCRF's in the network, when it exchanges messages with Untrusted Networks. A PCRF's identity is embedded in the Origin-Host and Session-Id AVPs sent in Request messages and the Origin-Host AVP sent in Answer messages. This capability is associated with the Diameter S9 and Rx application messages over the S9Reference Point. This S9 PCRF Topology Hiding feature encompasses:

- PCRF Topology Hiding in inbound and outbound roaming use cases – Hiding of PCRF host names in S9messages over the S9 Reference Point in Local Breakout (LBO) roaming architecture with the AF in the Visited Network or with the AF in the Home Network. Also hiding of PCRF host names in S9 messages over the S9Reference Point in the Home Routed Access roaming architecture.
- PCRF Topology Hiding in outbound roaming use case: Hiding of PCRF host names in Rx messages over the S9Reference Point in Local Breakout (LBO) roaming architecture with the AF in the Visited Network.

- PCRF Topology Hiding in inbound roaming use case – Hiding of PCRF host names in Rx messages over the S9Reference Point in Local Breakout (LBO) roaming architecture with the AF in the Visited Network where the Visited PCRF is implemented as a client/server of Rx messages to/from the Home PCRF.

The technique to hide and restore PCRF identities is similar to as described in S6a/S6d MME Topology Hiding.

2.12.4 S9 AF/pCSCF Topology Hiding

S9 AF/pCSCF topology hiding is concerned with hiding the identity of a Protected Home Network's AF/pCSCFs, as well as the number of AF/pCSCF's in the network, when it exchanges messages with Untrusted Networks. An AF/pCSCF identity is embedded in the Origin-Host and Session-Id AVPs sent in Request messages and the Origin-Host AVP sent in Answer messages. This is associated with the Diameter Rx application messages over the S9 Reference point. This AF/pCSCF Topology Hiding feature encompasses:

- 9 AF/ pCSCF Topology Hiding (inbound roaming use case) – Hiding of AF/pCSCF host names in Rx messages over the S9 Reference Point in Local Breakout (LBO) roaming architecture with the AF in the Visited Network where the Visited PCRF is implemented as a Proxy of Rx messages to/from the Home PCRF.

The technique to hide and restore S9 AF/pCSCF identities is similar to as described in S6a/S6d MME Topology Hiding.

2.13 DSR Applications

Certain functionality on the DSR is deemed important or complicated enough to be called an application and the details on those items can be found in this section. In general, the DSR is positioned as a flexible multi-functional router that can provide any or all of the applications listed below, and would evolve to support additional applications.

- Range Based Address Resolution (RBAR): a DSR enhanced routing application which allows the user to route Diameter end-to-end transactions based on Application ID, Command Code, "Routing Entity" Type, and Routing Entity address ranges.
- Full Address Based Resolution (FABR): a DSR enhanced routing application which allows the user to route Diameter end-to-end transactions based on Application ID, Command Code, "Routing Entity" Type, and individual Routing Entity.
- Policy and Charging Application (PCA): a DSR application providing two functions: 1) Online Charging Proxy or Online Charging Diameter Routing Agent (OC-DRA) and 2) Policy Proxy or Policy Diameter Routing Agent (P-DRA)
- Gateway Location Application (GLA): manages state information required to route Gx, Rx and other policy related Diameter sessions.
- RADIUS-Diameter IWF (R-D IWF): This feature provides message conversion and interworking between a RADIUS based client (server) and a Diameter based server (client).
Support for multiple applications and application chaining is supported with some restrictions. The following application limitations exist:
- The following applications are mutually exclusive on the same DSR Signaling node:
 - GLA is only supported on nodes with PCA
- The following application combinations are not supported on the same Diameter Agent Server
 - All three of FABR, RBAR and PCA

- The following application and function chaining combinations are supported. The priority of the routing rules in the ART determine the chaining order of these applications:
 - RBAR to P-DRBAR to P-DRA
 - RBAR to OC-DRA
 - RBAR to MAP IWF
 - FABR to MAP IWF
 - RBAR to FABR
 - RBAR to RADIUS IWF
 - FABR to RBAR
 - FABR to P-DRA
 - FABR to OC-DRA
 - FABR to RADIUS IWF

2.13.1 Range Based Address Resolution (RBAR)

Range based address resolution is a DSR enhanced routing application which allows the user to route Diameter end-to-end transactions based on Application ID, Command Code, "Routing Entity" Type, and Routing Entity address ranges. A Routing Entity can be a User Identity (IMSI, MSISDN, External Identifier, IMPI or IMPU) or an IP Address associated with the User Equipment (IPv4 or IPv6-prefix address). Charging characteristics are supported for the "Routing Entity" Type as well. Routing resolves to a "Destination" which can be configured with any combination of a Realm and FQDN (Realm-only, FQDN-only, or Realm and FQDN). Prefix filtering is provided with the creation of a user-configurable table filled with invalid IMSI MCC values that is used during IMSI validation prior to using the IMSI value for address resolution. The address resolution application checks against ranges of MCC values which are then used to invalidate an IMSI. The RBAR supports third (tertiary) routing entity search in the priority list for performing Destination lookups for a given application-id and command code. That is, customers can use up to 3 routing entities to perform address Resolution, for e.g. IMSI, MSISDN and External-Identifier in the preferred order of priority.

The RBAR application routes all messages as a Diameter Proxy Agent. Each Routing Entity supports up to two prioritized AVPs that are searched for in the ingress Diameter message resolving to configure Destination node. When a message successfully resolves to a Destination, RBAR replaces the Destination-Host and possibly Destination-Realm AVP in the ingress message, with the corresponding values assigned to the resolved Destination, and forwards the message to the DSR Relay Agent for egress routing into the network. A GUI is provided allowing the operator to provision MCC-MNC combinations of all network operators in the world which includes the country and network name. A list of all the well-known MCC-MNC combinations are pre-populated at installation time but these can be modified/deleted at a later time.

2.13.2 Full Address Based Resolution

Full address based resolution is a DSR enhanced routing application which allows the user to route Diameter end-to-end transactions based on Application ID, Command Code, "Routing Entity" Type, and individual Routing Entity. For FABR a Routing Entity can be a User Identity (IMSI, MSISDN, URI, wild carded NAI, External-Identifier, IMPI or IMPU). The FABR also supports third (tertiary) routing entity search in the priority list for performing Destination lookups for a given application-id and command code. That is, customers can use up to 3 routing entities to perform address Resolution, for e.g. IMSI, MSISDN and External-Identifier in the preferred order of priority. Each Routing Entity supports up to two prioritized AVPs that are

searched for in the ingress Diameter message resolving to configure Destination node. As in RBAR, routing resolves to a "Destination" which can be configured with any combination of a Realm and FQDN (Realm-only, FQDN-only, or Realm and FQDN). Prefix filtering is provided with the creation of a user-configurable table filled with invalid IMSI MCC values that is used during IMSI validation prior to using the IMSI value for address resolution. The address resolution application checks against ranges of MCC values which are then used to invalidate an IMSI.

The FABR application routes all messages as a Diameter Proxy Agent. When a message successfully resolves to a Destination, FABR replaces the Destination-Host and possibly Destination-Realm AVP in the ingress message, with the corresponding values assigned to the resolved Destination, and forwards the message to the DSR RelayAgent for egress routing into the network. FABR uses the remote database storage called DSR Data Repository(DDR) to store subscriber data. DDR is hosted on the Database Processor blades at each node.

A GUI is provided allowing the operator to provision MCC-MNC combinations of all network operators in the world including the country and network name. A list of all the well-known MCC-MNC combinations are pre-populated at installation time but these can be modified/deleted at a later time.


Subscriber Data Server (SDS) Integration

Oracle Communication's Subscriber Data Server (SDS) integrates with the DSR to provide the following functions:

- Provisioning and storage of large amounts of database information required for the Full Address Based Resolution (FABR) feature.
- Replication of information across multiple sites so that the data may be queried at the DSR sites.
- Support for querying by backend Operating systems to maintain reports and audit information.

The central provisioning capability is provided by the SDS component. The SDS is deployed optionally geo-redundantly at a Primary and Disaster recovery site. A Query Server component that processes queries from backend customer operations systems is deployed optionally geo-redundantly at the Primary and Disaster Recovery SDS site. FABR data along with any other future DSR specific subscriber data is termed DSR Data. The application hosting the DSR Data is termed the DSR Data Repository (DDR). The SDS supports a SOAP/XML interface for provisioning. This interface supports Insert, Update & Delete functions on the Subscriber profile.

Figure 2-49 Subscriber Data Server Architecture



The SDS also supports Split NPA data. When a service provider exhausts all MSISDNs within a Numbering Plan Area (NPA), the service provider commonly adds another NPA to the region. The result of assigning a new NPA is called a NPA Split. As new NXXs are defined in the new NPA, existing exchanges (NXXs) may be assigned to the newly created NXXs from the old NPA. The new and the old NXX have the same value.

When an NPA split occurs, a period of time is set aside during which a subscriber can be reached via phone number using old NPA-NXX and via phone number using new NPA-NXX. This period is called Permissive Dialing Period (PDP).

NPA splits apply to MSISDNs. During the NPA Split process, the SDS will automatically create duplicate MSISDN records at the start of Permissive Dialing Period (PDP) time (activation) and delete old MSISDN records at the end of PDP time (completion).

The SDS Subscriber Identity Grouping (Subscribers page) allows users to group optional customer-specified account IDs, multiple MSISDNs routing entities, and/or multiple IMSI routing entities together into one Subscriber. After a Subscriber (a group of related routing entities and an optional Account ID value) is created, the destinations for all of the related routing entities can be updated, all data from the subscriber can be read, and the subscriber can be deleted or its addresses modified by using any of the subscriber's addresses (account ID, MSISDN, or IMSI).

In order to help maintenance personnel with trouble shooting at the Query Server, records belonging to a single subscriber are now correlated at the SDS and the Query Server.


User Data Repository (UDR) Integration

Oracle Communication's User Data Repository (UDR) integrates with the DSR to provide the following functions:

- Provisioning and storage of large amounts of database information required for the Full Address Based Resolution (FABR) feature.
- Replication of information across multiple sites so that the data may be queried at the DSR sites.
- Support for querying by backend Operating systems to maintain reports and audit information.

The central provisioning capability is provided by the UDR component. The UDR is deployed optionally geo-redundantly at a Primary and Disaster recovery site. A Query Server component that processes queries from backend customer operations systems is deployed optionally geo-redundantly at the Primary and Disaster Recovery UDR site. FABR data along with any other future DSR specific subscriber data is termed DSR Data. The application hosting the DSR Data is termed the DSR Data Repository (DDR). The UDR supports a SOAP/XML interface for provisioning. This interface supports Insert, Update & Delete functions on the Subscriber profile.

Figure 2-50 User Data Repository Architecture



The UDR also supports Split NPA data. When a service provider exhausts all MSISDNs within a Numbering Plan Area (NPA), the service provider commonly adds another NPA to the region. The result of assigning a new NPA is called a NPA Split. As new NXXs are defined in the new NPA, existing exchanges (NXXs) may be assigned to the newly created NXXs from the old NPA. The new and the old NXX have the same value.

When an NPA split occurs, a period of time is set aside during which a subscriber can be reached via phone number using old NPA-NXX and via phone number using new NPA-NXX. This period is called Permissive Dialing Period (PDP).

NPA splits apply to MSISDNs. During the NPA Split process, the UDR will automatically create duplicate MSISDN records at the start of Permissive Dialing Period (PDP) time (activation) and delete old MSISDN records at the end of PDP time (completion).

The UDR Subscriber Identity Grouping (Subscribers page) allows users to group optional customer-specified account IDs, multiple MSISDNs routing entities, and/or multiple IMSI routing entities together into one Subscriber. After a Subscriber (a group of related routing entities and an optional Account ID value) is created, the destinations for all of the related routing entities can be updated, all data from the subscriber can be read, and the subscriber can be deleted or its addresses modified by using any of the subscriber's addresses (account ID, MSISDN, or IMSI).

In order to help maintenance personnel with trouble shooting at the Query Server, records belonging to a single subscriber are now correlated at the UDR and the Query Server.

Limitations of UDR

As DSR can now send FABR queries to UDR, unlike SDS the following features are currently not supported by UDR:

- Prefix Search for IMSI and MSISDN.
- External Identifier match partial (Domain identifier).
- 16 priority Support.
- NGN PS priority Support.
- Wild Card Nai User.

FABR Blacklist

The FABR application also supports the rejection of Diameter requests which carry a blacklisted IMSI/MSISDN. A blacklist search is performed prior to the Full address search. This search can be enabled for a combination of Application-Id, Command-Code, and Routing Entity. If a match is found during the blacklist search, the operator is able to configure FABR, on a per Application-Id basis, to either respond to the Diameter request with a configurable Result-Code/ Experimental Result-Code, or Forward the Request to a default destination or forward the Request unchanged.

A total of 1 Million IMSIs and 1 Million MSISDNs (not prefixes) are supported for blacklisting. The IMSIs are of fixed length (15 digits long) and the MSISDNs are provisioned as E.164 numbers (includes the Country code but with out the + sign). The blacklisted IMSIs and MSISDNs are provisioned via the SDS/UDR GUI or via bulk import using a CSV file.

IMSI/MSISDN Prefix Lookups

Operators use FABR to resolve individual subscriber IMSIs or MSISDNs to specific end points such as a HSS. This ability to resolve the address on an individual subscriber basis provides the highest degree of freedom and flexibility to the operator and allows for subscribers to be assigned to an HSS based on a criteria that fits the operator's needs.

The prefix lookups allow an operator to manage routing based on IMSI prefixes/ranges. All the IMSIs that fall under a particular IMSI prefix/range resolve to the same end point. For example, a block of IMSIs for Machine-to-Machine(M2M) communication could be used and the operator wishes to route all registration requests arising from these IMSIs to a specific HSS (or a set of HSSs) that is dedicated for M2M. Providing the ability to provision ranges results in significant operational savings from a provisioning point of view.

Prefix based lookups are performed after the full address lookup. The prefix based lookup is only performed if the full address lookup does not find a match and can be enabled by the operator for a combination of Application-Id, Command-Code and Routing Entity Type. For example, an operator can choose to perform the prefix lookup only on the S6a-AIR request but not on the other S6a requests. The Routing Entity Type provides additional granularity when the same request carries multiple subscriber identities and the prefix lookup is performed only for one of those identities but not both. For example, certain Cx Requests are known to carry

both an IMSI and an MSISDN and this feature allows an operator to perform a prefix lookup for the IMSI but not for the MSISDN.⁹² | ORACLE COMMUNICATIONS DIAMETER SIGNALING ROUTER RELEASE 9.0.0.0.0 FEATURE GUIDE MSISDN prefixes are supported as well. This allows an operator to route a Diameter Request such as the Cx-LIR based on a prefix if the individual entry is not found.

2.13.3 Policy and Charging Application (PCA)

The Policy and Charging Application provides two functions on the DSR:

- Online Charging Proxy (also known as Online Charging Diameter Routing Agent (OC-DRA)).
- Policy Proxy (also known as Policy Diameter Routing Agent (P-DRA)).

A PCA DSR can be deployed in a Diameter network with either P-DRA function or OC-DRA function enabled or with both P-DRA and OC-DRA functions enabled on a network-wide basis.


Online Charging Proxy (OC-DRA – Online Charging Diameter Routing Agent).

Mobile Operators are increasingly using Diameter based infrastructure for subscriber charging. 3G operators use a mix of CAMEL and Diameter for charging voice and data sessions respectively while LTE/VoLTE standards call for using Diameter exclusively for the transport of charging messages between charging servers and charging clients.

Online Charging and Offline Charging mechanisms were originally put in place by the standards bodies to address prepaid and postpaid subscribers, but lately, operators seem to be migrating towards convergent charging systems that use Online Charging mechanisms for both prepaid and postpaid subscribers. In the DSR, the Online Charging Proxy provides the Online Charging Diameter Routing Agent (OC-DRA) function.

The figure below shows the Online Charging Architecture as per 3GPP. The architecture does not mandate a DRA and thus does not depict a DRA but shows the various CTFs that can initiate Online Charging messages via Ro or CAP. The figure also shows the components within the Online Charging System (blue box) which typically maps to the Online Charging Server.


Figure 2-51 Online Charging System and Architecture



The following features are supported as part of the Online Charging Proxy.

- Support Gy/Ro interfaces for online charging sessions between Charging Trigger Function (CTF) and Online Charging System (OCS).
- Selection of an OCS or OCS cluster for a specific user based on subscriber's ID and/or APN.
- Creation and maintaining of session state info for some online charging sessions, if configured.
- Stateful session-base routing of online charging messages to available OCSs.
- High Availability within the site using N+1 DA MP deployment model.
- Geo-Redundancy by sharing session state across mated sites where needed.

The OC-DRA solution retrieves the subscriber's identity from any of the above mentioned AVPs and stores them as part of subscriber state if needed and used for debugging/tracing customer sessions.

Figure 2-52 A typical Online Charging Session**Policy Proxy (PDRA – Policy Diameter Routing Agent)**

With the advent of LTE and high-speed wireless networks, network providers have a need to manage subscriber resource usage across their entire network. To accomplish network-wide resource monitoring and control requires identification of subscriber resource usage using multiple keys (e.g. IMSI, MSISDN, IP addresses) in a network with large numbers of policy enforcement clients and policy rules servers (PCRFs). Subscriber requests for access to network resources must be routed to a single PCRF in the network so that policy decisions can be made with knowledge of all the resources being used by all of that subscriber's policy sessions. Rather than creating a provisioned relationship between subscribers and PCRFs, which would be difficult and expensive to manage, subscribers are dynamically assigned to a PCRF when the initial bearer session (Gx or Gxx interface) is created. All subscriber policy sessions from anywhere in the network are routed to the assigned PCRF until that subscriber's last Gx or Gxx session ends, at which point the next Gx or Gxx session may be routed to a different PCRF. This dynamic mapping of subscribers to PCRFs provides automatic load distribution to available PCRFs, while still mapping all of a subscriber's sessions to a single PCRF.

Operators are relying on PDRA for its session binding/correlation abilities to enable VoLTE in their networks. In the VoLTE scenarios, Rx Requests initiated by the AS (P-CSCF) are correlated by the PDRA and routed to the PCRF serving the corresponding Gx session. PDRA creates bindings as policy sessions are established and this binding information is then used to route subsequent sessions initiated by the subscriber. In certain situations, such as the failure or the reboot of a PCRF, the binding information in the PDRA becomes invalid and must be deleted as soon as possible. In the case of a PCRF failure the subscriber's Gx session is torn down. This cleanup action forces the subscriber to re-initiate the IP-CAN session and the Gx session so that it may be routed to a functioning PCRF.

This feature allows the removal of any binding capable interface supported by PDRA which can be triggered off Diameter based failures. The DSR monitors the type and the number of error responses originated by the PCRF. (In some situations, the error responses maybe

generated by the DSR on behalf of the PCRF.) The PDRA marks a binding as suspect upon seeing certain error responses (also called as session removal events) and tears down the subscriber's Gx session when the number of such error responses exceed a pre-configured value. This forces the subscriber to re-initiate the Gx session which can then be routed to a functioning PCRF. Furthermore, the feature removes all of the subscriber's Gx sessions (or other binding capable sessions) associated with the failed PCRF. The subscriber's Gx sessions (or other binding capable sessions) associated with other PCRFs are not impacted.


In addition to managing a subscriber's resource usage across the network, network providers may have a need to perform topology hiding of the PCRF from some policy clients. This topology hiding prevents the policy client from obtaining knowledge of the PCRF identity (host name or IP address), or indeed knowledge of the number or location of PCRFs deployed in the network.

In summary, the Policy DRA function provides the following capabilities:

- Distribution of Gx, Gxx, and S9 policy sessions (that is binding capable sessions) to available PCRFs.
- Binding of subscriber keys such as IMSI, MSISDN, and IP addresses to the PCRF selected when the initial Gx, Gxx, or S9 session was established.
- Providing network-wide correlation of subscriber sessions such that a policy session initiated anywhere in the network will be routed to the PCRF that is serving the subscriber.
- Providing multiple binding keys by which a subscriber can be identified so that policy clients that use different keys can still be routed to the PCRF assigned to the subscriber.
- Efficient routing of Diameter messages such that any policy client in the network can signal to any PCRF in the network, and vice-versa, without requiring full-mesh Diameter connectivity.
- Hiding of PCRF topology information from specified policy clients.


The figure below illustrates an example policy network with P-DRA DSRs deployed:

Figure 2-53 Network View of P-DRA Mated Pairs



The primary Diameter interfaces to/from the PCRF in a non-roaming environment are Gx (PCEF-PCRF), Gxx(BBERF-PCRF), Gx'/Gx-Lite and Rx (AF-PCRF). These are highlighted in the figure below.. All of these may not be, and often are not, present in all networks. In addition, variants of these interfaces are sometimes used, for example from systems which perform DPI (Deep Packet Inspection) and augment other PCEFs such as GGSNs and PGWs.

Figure 2-54 Overall PCC logical architecture (non-roam)



The DRA first provides distribution of subscribers' initial Gx sessions, which correspond to their data (IP-CAN) sessions, to PCRFs. This can be done in dynamic (e.g. round-robin) or static (e.g. range-based routing) fashion. Via PCRF binding, the DRA then remembers the PCRF that has been assigned for a subscriber's data session(s) and makes sure that all policy related messages associated with that user's active data session(s) are routed to the same PCRF. Via session correlation, the DRA associates multiple simultaneous Gx/Gxx and Rx sessions for the same user to the same PCRF.

For various reasons, there may be the need to hide the specific Diameter identities of PCRFs from other devices or networks. The DRA is the logical place to perform such topology hiding.

The primary purposes of the DSR Policy DRA function are:

- Distributing initial Gx, Gxx and S9 sessions across available PCRFs.
- Providing network wide subscriber binding by storing the relationship between various subscriber data session identities, such as MSISDN / IP address(es) / IMSI, and the assigned PCRF. All P-DRA in the defined P-DRA pool must work together as a single logical P-DRA.
- Providing network wide session correlation by using the stored binding data to associate other Diameter sessions with the initial session for the subscriber and route messages to the assigned PCRF.
- Performing topology hiding to hide the true identities of the PCRFs from other elements in the network.

Support for Gx' / Gx Lite

The PCRF's primary enforcement point today in the mobile networks is the PGW and is achieved over the Gx interface. This control is based on the subscriber's profile which is provisioned by the operator and provides a certain amount of control over the subscriber's voice and data sessions.

Lately, operators are seeing the need for a finer level of control that is based on the data being exchanged between a user and the internet. This can be for reasons such as video optimization, parental controls, content filtering and traffic/bandwidth management. To help with this, several vendors have built products (generally called as DPI/MOS servers) that reside in the data path and can inspect the data being exchanged at much finer granularity and provide feedback to the PCRF servers. The PCRF servers can then use this information to influence the PGW via the Gx session (in a manner similar to how the Rx interface influences the Gx session).


3GPP has defined the Sd interface in 3GPP release 11 and beyond, for use between the DPI and PCRF servers. However, some of the DPI vendors have produced these boxes before the Sd interface was standardized, adopted Gx with minor variations as the protocol between DPI and PCRF servers. These Gx variations are referred to by some as Gx` and by others as Gx-Lite. It should be noted that Gx` interface does not carry the IMSI which is usually present on the Gx interface. The same is true for Sd interface as well.

The DSR based Policy DRA application manages state required to route Gx, Gxx, Rx and S9 Diameter sessions that belong to a single subscriber to the same PCRF. Given the introduction of DPI/MOS servers into the mobile networks, the Policy DRA must be enhanced to support the interfaces used by these servers (Gx`) so that these sessions are routed to the same PCRF that is hosting the corresponding Gx/Gxx session.

Supporting the Gx`/Gx Lite interface involves identifying these sessions, extracting the subscriber keys from , performing a binding lookup and finally routing these requests to the appropriate PCRF. The lookup is typically done on the session initiating the request with subsequent requests performing destination-host based routing but if PCRF topology hiding is enabled, the session information has to be stored in the session data base and a lookup is required for subsequent requests in the session.

PCRF Topology Hiding

The P-DRA also supports PCRF topology hiding, which can optionally be enabled on a per-destination basis. If enabled for a destination, topology hiding means the PCRF appears as a single large PCRF to that destination. An example where the peer is a PCEF is shown in the figure below, which shows the message flow for a CCR message. This same flow applies to all CCR messages, with the exception that the Initial message might not contain a Destination-Host, in which case the P-DRA adds a Destination-Host to the message before sending to the PCRF. The P-DRA distributes CCR-Initial messages for a user's first session over the Diameter connections to a pool of PCRF connections. The P-DRA, absent of failures, sends all messages of a Diameter session to the same PCRF for the duration of the session.

Figure 2-55 PCRF topology hiding

In the CCR-I, the PCEF optionally includes the Destination-Host of P-DRA and upon receiving an initial CCA from the P-DRA, populates the Destination-Host AVP with the P-DRA ID for subsequent messages (CCR-U and CCR-T). This is based on the Origin-Host AVP received in the initial CCA from the P-DRA.

Topology hiding also applies to Request messages sent from a PCRF to the affected destination.

APN Based PCRF Pooling

Service providers require flexibility in the deployment of new policy-controlled services. They need the ability to roll in new services or new PCRF infrastructure without disturbing existing services. For instance, a carrier might want to have 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 needs to have these two services segregated at the PCRF level.

The introduction of multiple PCRF pools also introduces the requirement to differentiate the binding records in the binding SBR. It is possible for the same UE, as indicated by the IMSI, to have multiple active IP can sessions spread across the different pools.

The contents of binding generating Gx CCR-I messages are inspected to select the type of PCRF to which the CCR-I messages are to be routed. This feature allows sets of PCRFs to be service specific. The IMSI and/or 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.


A PCRF pool is a set of PCRF's able to handle a set of policy-based services. Multiple pools are supported requiring the PDRA to allow the selection to which a new-binding CCR-I belongs.

Note


While the concept of a PCRF pool might be a network wide concept for a service provider, the configuration of PCRF pools is done on a PDRA site-by-site basis. It is a requirement that PDRA sites be able to have different PCRF Pool Selection configuration.

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 PDRA peers that are a part of the pool.

Figure 2-56 Relationship between APNs and PCRF Pools



The figure below illustrates the relationship between IMSI and PCRF pool. The same IMSI is able to have active bindings to multiple PCRF pools.

Figure 2-57 Relationship between IMSIs and PCRF pools

PCA Deployment

A PCA DSR consists of a number of PCA DA-MP servers, a number of SBR servers, OAM server, and optionally, IPFE servers. The PCA DA-MP servers are responsible for handling Diameter signaling and implementing the Policy DRA and Online Charging DRA feature business logics. PCA DA-MP servers run the PCA application in the same process with the Oracle Diameter stack.

SBR servers host the policy session and policy binding databases for P-DRA function, and online charging session database for OC-DRA function respectively. These are special purpose MP blades that provide an off-board database for use by the PCA application business logic hosted on the PCA DA-MP servers. The P-DRA functional ways maintains session records for binding capable sessions (Gx, Gxx, and the S9 versions of Gx and Gxx), and binding dependent sessions (Rx and Gx-Prime) for which topology hiding is in effect. The OC-DRA function maintains session records for binding independent sessions (Gy and Ro) based on configuration and Diameter message content.


Each PCA DSR hosts connections to clients and to policy/charging servers such as OCSs and PCRFs. Clients are devices (not provided by Oracle) that request authorization for access to network resources on behalf of user equipment (e.g. mobile phones) from the 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 P-DRA. P-DRA makes sure that all policy authorization requests for a given subscriber are routed to the same PCRF. Charging clients (CTF) generates charging events based on the observation of network resource usage and collect 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.

PCA DSRs can be deployed in mated pairs such that policy 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 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/OCSSs are cross-connected such that all three PCA DSRs have connections to all policy clients and all PCRFs/OCSSs associated with the mated triplet.

PCA network is the term used to describe a set of PCA mated pairs and network OAM&P server pair/triplet. All clients and PCRFs/OCSSs are reachable for Diameter signaling from any PCA DSR in the PCA network.

Figure 2-58 PCA Example Deployment




2.13.4 Gateway Location Application (GLA)

The DSR based PCA PDRA function manages state required to route Gx, Rx and other policy related Diameter sessions. The Policy DRA SBR-B is a network wide repository for that state.

Customers are recognizing the value of having a centralized, network wide repository for binding state and are identifying additional ways to leverage the Policy DRA managed state.


The Gateway Location Application (GLA) provides a Diameter signaling approach for accessing that binding state. The GLA gives the ability to retrieve the Diameter identity that initiated Gx sessions for a given IMSI or MsISDN.

A use case for this application is an IMSI query with a single matching Gx session. The figure below shows this use case where the GGR message includes a query that has IMSI as the query key. In this example a single Gx session matches the query.

Figure 2-59 IMSI Query with Single Matching Gx Session Use Case

The steps for this use case are as follows:

- Existing Policy DRA handling of a Gx CCR-I session. This session is the first for the IMSI and results in a new binding.
- The Policy DRA application stores the gateway state associated with the Gx session. This includes the APN for the session and the Origin-Host received in the CCR-I message. The Origin-Host contains the Diameter Identity of the PCEF that originates the CCR-I and will generally be the FQDN of the PCEF.
- The GQC generates a GGR message with IMSI as the query key.
- The GLA queries the SBR-B to get the gateway state for the Gx session or sessions associated with the IMSI combination.
- The SBR-B returns the gateway state for all sessions associated with the IMSI. In this case there is one Gx session, the one that resulted in the binding. The state returned included the Origin-Host and APN associated with the session. A time stamp for when the session was initiated is also included.
- The GLA returns the Gx session state in a GGA message. If no matching sessions are included in the GW State Response then the GLA returns a response.
- The GLA application's role is to provide access to state generated by the PCA PDRA function. As a result, the GLA application must be deployed in a network that includes the PCA. The implication of this is that the PCA and the GLA application must be managed by the same NOAM. This is illustrated in the figure below:

Figure 2-60 PCA and GLA NOAM Architecture

Within a single DSR Network Element, there are three alternatives for deploying the GLA application


Dedicated GLA DA-MPs – The GLA application is deployed in a DSR NE that also supports the PCA but is deployed on dedicated DA-MPs. The benefit of this deployment architecture is that it isolates the GLA Diameter traffic from the Policy DRA Diameter traffic. The GLA traffic can vary greatly and at times can spike to a high traffic rate. This deployment alternative helps to minimize the impact of those traffic spikes on the mainline PCA. Note that the full impact of the traffic cannot be isolated as the GLA queries result in interactions with the SBR-B database.

- **Shared GLA DA-MPs** – The GLA application is deployed in a DSR NE that also supports the PCA. The GLA application and PCA are both enabled on common DA-MPs.
- **Dedicated GLA Network Element** – The GLA application is deployed as a separate set of DSR NEs. This must be in a network that includes DSR NEs running the PC.


When deployed using separate sets of MPs and when using IPFE to distribute client-initiated connections, it is necessary to configure separate target sets for each application. One IPFE target set contains the PCR MPs and a second IPFE target set contains the GLA MPs.

2.13.5 Diameter Security Application (DSA)

Network security is a key concern for service providers as they interconnect networks to provide universal services for their subscribers. Signaling networks at interconnect are secured through business arrangements rather than firewalls between other MNO's where there is trust between MNO's when communicating information. Most of the Diameter Security vulnerabilities happen in the interconnect from roaming networks through IPX or directly from roaming partner networks.

Figure 2-61 Interconnect Network and Security Vulnerabilities

GSMA has published the interconnect diameter network security recommendations in FS.19 and IR.88 standard specifications. It includes detailing the different security threats associated with both Diameter networks, describing the possible attacks which can be implemented, evaluating the risks and defining a set of best practice countermeasures. Best practice to secure Diameter network is to use single point of interconnect and controlling access permission for all incoming requests at Diameter Edge Agent (DEA). DEA is the only point of contact into and out of an operator's network at the Diameter application level. DSR provides GSMA FS.19 and IR.88 compliant integrated Diameter Firewall functionality through Diameter Security Application (DSA) when DSR acting as DEA. Diameter Security application (DSA) implements all GSMA FS.19 standard defined security message category filters and other security countermeasures to mitigate different Diameter attacks. Diameter Security application(DSA) shall apply different countermeasures for ingress messages received from external foreign network and for egress messages sent to external foreign network.

Figure 2-62 Diameter Security Application (DSA)

DSA security countermeasures are applied to two types of subscribers:

- Inbound roaming subscribers: Countermeasures are applied for visited network subscribers roaming inhome network.
- Outbound roaming subscribers: Countermeasures are applied for home network subscribers roaming invisited network.

Different Security countermeasure profiles can be created for different IPX or roaming partners by enabling and disabling countermeasures individually for different IPX provider or roaming partner Diameter Peers. DSA shall provide two modes of operation for individual countermeasure:

- Detection Only: DSA shall monitor Diameter Traffic and report different Diameter Vulnerabilities.
- Detection and Correction:
 - By Dropping message: DSA shall drop the vulnerable Diameter messages.
 - By Rejecting message: DSA shall reject the vulnerable Diameter request messages with error answer messages. If the vulnerable message is answer message then answer message shall be dropped by DSA.

Detection Only' shall be default operation mode set for all counter measures.

DSA security countermeasures can be classified as:

- Stateless Security Countermeasures: Counter measures which can be applied without maintaining the state information at DSA. Some of the important countermeasures are:
 - Application-Id whitelist screening.
 - Application-Id and Command Code consistency Check.
 - AVP Multiple Instance Check.
 - Avp Whitelist Screening
 - Destination-Realm and Origin-Realm match Check.
 - Origin Realm and Destination Realm whitelist screening.
 - Realm and IMSI consistency check.
 - Specific AVP Screening.
 - Session-Id validation check.
 - Subscriber Identity validation.
 - Origin-Host and Origin-Realm format check.
 - Origin host and Origin Realm consistency check.
 - Visited-PLMN-ID and Origin-Realm consistency check.
- Stateful countermeasures: Countermeasures that requires maintaining the state information at DSA for correlating different Diameter Transactions in call flows. Some of the important countermeasures are:
 - Previous Location Check
 - Time-Distance Check
 - Source Host validation – MME Validation
 - Source Host validation – HSS Validation
 - Message rate monitoring

2.13.6 DSA – Cross Protocol Security

Multiple protocol technologies are available in the network today, for example, 2G, 3G, and 4G. Each Protocol class is solving security issues independently. There is a class of threats that may affect the networks supporting 2G/3G and 4G protocols together. The hacker could potentially imitate an MME, sending messages for an IMSI for a Subscriber who is the 2G/3G coverage Area. Similarly, a Hacker can simulate VLR for an IMSI of a Subscriber in the 4G Coverage Area. This solution intends to provide a way to implement security across protocol classes. This solution covers Velocity Check/Time Distance Check.

2.13.7 DSA – Common Visualization Framework

DSA has introduced Common Visualization Framework using industry-leading tools, such as Elastic, Logstash, and Kibana. This framework allows customers to proactively detect, reconfigure, add, and delete any countermeasure. Active SOAM receives logs from DA-MP, which are enriched, and passed on to the configured visualization framework. This feature enhancement allows customers to use a combination of parameters, such as Category, Counter Measure name, Action Applied, IMSI, App_ID, Command Code, VPLMN Id, Origin-Host, Origin Realm, Destination Host and Destination Realm, Country Code, and design dashboards within the capabilities provided by Kibana.

2.13.8 Service Based Interface Support

When DSR receives an Rx-AAR message, the message is processed in the following ways:


- For 4G users, Policy DRA (P-DRA) queries the SBR database.
- For 5G users, P-DRA forwards the Rx-AAR message to the configured diameter peer (OCIWF) because SBR does not contain 5G information. When the OC-IWF sends the Rx AAA – 3006 (Diameter Redirect)Redirect-Host – PCF FQDN, the Redirect Agent functionality in DSR redirects the Authentication,Authorization Request (AAR) to the PCF FQDN.

2.14 Diameter Message Copy

The DSR is able to copy certain Diameter Requests or Requests and Answers that transit the system. The copied messages can be used for book keeping/verification or for offering additional services such as sending a welcome SMS. The copied messages are sent towards Diameter Application Servers (DAS) which behave like RFC6733compliant standard Diameter servers.

The figure below provides a high level overview and shows the message processing sequence followed by DSR when performing Message Copy. It should be noted that the Message Copy is performed after the completion of the original transaction. In cases where a copy of the Answer message is to be copied, the Answer message is embedded into a Proprietary AVP and included in the copied message.

Figure 2-63 Message Copy Overview



The Message Copy function can be triggered by the following mechanisms:


- PRT based triggering.
- Using DSR's mediation rules.
- DSR application triggering (for example: FABR).

2.15 Virtualized network functions manager VNF

Network Function Virtualization-Management and Network Orchestration (NFV-MANO) architecture is the framework created for performing the necessary management and orchestration by ETSI standards. DSR 9.0.0.0 includes a VNFM for automating life cycle management operation of DSR application.

DSR VNFM is a standalone VM deployed separately from the DSR VNFCs. The DSR VNFM provides a REST API where in subset of VNF specific function provided in ETSI NFV SOL 003 can be supported. DSR VNFM provides supports ETSI direct model to interface with Open Stack and providing HEAT templates to deploy DSR/vSTP applications.

Figure 2-64 DSR VNFM



DSR application automatic instantiation process includes two operations specified by the ETSI NFV SOL 003. Those operations supported by DSR VNFM currently are as follows:

- Create VNF Identifier
- Instantiate VNF
- VNFC Scale Out capabilities
- Query Life Cycle Management
- Auto Discover capabilities

- Terminate VNF

2.16 Custom Application Framework (CAF)

Custom Application Framework (CAF) is the call processing application development framework which allows CSP's to rapidly build and deploy value added services on DSR. Applications built using CAF will run and interface with the DSR in a manner similar to the native DSR applications like RBAR, FABR, and PCA.

Figure 2-65 Custom Application Framework (CAF)




Figure 70 – Custom Application Framework (CAF)

CAF provides the GUI based IDE at NOAM to build the applications using Perl language. Developers can define and access configuration and maintenance data tables required for custom application business logic. CAF allows to define and access custom application state data using the DSR's high-capacity, highly-available, geo-redundant Subscriber Data Repository (UDR) subsystem. CAF also includes meta-data in DSR's Integrated Diameter Intelligence Hub (IDIH) trace records for advanced troubleshooting using IDIH.

CAF includes the following set of inbuilt Perl API's which can be used to build value added call processing applications at DSR:

- Message manipulation API's to encode and decode diameter messages.
- API's to perform basic routing functions like: set ART, set PRT, forward/drop a Diameter message or respond back to request with Diameter Answer message.
- Database access API's to perform create/read/update/delete operations on UDR tables and read operations on configuration tables used by CAF applications.
- Utilities API's to peg the measurement counters, raise the alarms and events, log the debug events, so on.


2.16.1 CAF Application Life Cycle Management

CAF Application Life Cycle Management

CAF application life cycle starts with activation of CAF application at NOAM and takes effect after the CAF has been administratively enabled. CAF provides the following application life cycle capabilities:

- Management of multiple versions per CAF applications.
- Supports importing/exporting of application business logic and configuration data across deployments.
- Each CAF application can be in one of the following states:
 - Development (initial state) – Allows to creation of new application or update of existing version of CAF application.
 - Trial (at most one at a time) – Allows to run the CAF application only on configured Trial DA-MPs before actual deployment into Production environment.
 - Production (at most one at a time; if missing, the CAF app goes into „Unavailable“ operational state) – Allows CAF application to be deployed into production environment across DA-MP.
 - Archived – Allows archiving of older versions CAF application for later reuse.
- Provides capability to roll back the system to a previous state with minimum service interruption, in case the new production version experiences unexpected runtime errors or performance issues.

Figure 2-66 Custom Application Framework (CAF) Application State Transition



2.16.2 Reference CAF Applications

DSR delivers the following reference CAF applications which can be used by customers:

- **Steering of Roaming (SoR):** The SoR application lets home network operators control and distribute registration traffic of their outbound roamers. Use SoR to define static distribution roaming steering policies for each group of roaming partners that are part of the same country. Refer SoR User's Guide for more details.
- **Zero Balance Application (ZBA):** When a mobile data service subscriber initiates a data session, a network Element (PCEF or CTF) may attempt to setup a data session for the subscriber and query the OCS for service units. In the case where the subscriber has no credit in their account, the OCS will reject the request, usually with an indication that the subscriber has no credit. However, in most cases the PCEF, or other CTFs (Charging Trigger Functions), will continue to attempt to query the OCS. This can create overload on the OCS to handle requests for subscribers that will always be denied during the period in which the subscriber has no credit. The purpose of the ZBA application is to detect when a subscriber, identified by a MSISDN, has a zero balance in his/her account, the ZBA application under certain configurable conditions, may respond on behalf of the OCS and reduce the number of queries the OCS must handle for subscribers who temporarily have no credit. Refer Zero Balance Application User's Guide for more details.

Enhanced DCA Framework to create and send Diameter Message

- DCA Framework has been enhanced to create and send a diameter message to another node. DCA framework can create diameter messages and once it is created successfully it shall be able to send the message to the diameter peer node.

2.17 Security Assertion Markup Language (SAML)

SAML authentication is an additional authentication mechanism in DSR/SDS to existing local and LDAP authentication mechanisms for authenticating user. SAML is an XML-based open standard for exchanging authentication data between a Service Provider (providing a service to the user) and an Identity Provider (providing user identity verification for the Service Provider). Web applications leverage SAML via the IDP to authenticate the user. Service Provider does not need to store passwords and not having to address forgotten password issues. IDP and SP exchange their metadata containing information required for interaction between them.

Table 2-12 Terms and definition

Terms	Definition
Identity Provider (IDP)	Entity that verifies the identity of the user, in response to a request by the Service Provider. The Identity Provider is responsible for maintaining and authenticating the user's identity.
Service Provider(SP)	Service Provider (SP) offers a service to the user and allows the user to sign in by using SAML.
SAML Metadata	SAML metadata is an XML document containing necessary information for communication between identity provider and service provider.
SAML Assertion	An XML document returned by the Identity Provider to the Service Provider after authentication of the user.
Assertion Consumer Service(ACS)endpoint	The endpoint where the Service Provider will receive SAML assertions issued by the Identity Provider.
Entity ID	A unique identifier for a SAML entity. A SAML entity can be a Service Provider or an Identity Provider.
Bindings	SAML requestors and responders communicate by exchanging messages. The mechanism to transport these messages is called a SAML binding.
Metadata	A set of information supplied by the IdP to the SP, and/or vice versa, in XML format.

2.17.1 SAML Authentication Flow

SAML Authentication Flow:

1. Enable SAML authentication functionality from General Options screen.
2. Upload IDP Metadata file in DSR/SDS.
3. Customer needs to upload DSR/SDS Metadata file on their IDP.
4. Once configuration is completed user can login via saml using url: <ipaddress>?auth=SAML.

2.17.2 SAML Feature Description

SAML Feature Description


- DSR and SDS will act as a service provider for SAML authentication.
- DSR/SDS SAML authentication and configuration is supported through GUI only. There is no MMI support of it.
- DSR / SDS GUI can be accessed using XMI or VIP address for SAML version 2.0 based authentication.
- DP provided XML metadata to be uploaded from Active NOAM screen “Administration -> Remote Servers -> SAML Authentication”. Same IDP metadata will be applicable to other OAM servers in topology.
- DSR / SDS supports HTTP-POST and HTTP-Redirect binding when SAML request is send by SP to IDP via user agent. This information should be available in IDP provided metadata.
- DSR / SDS supports HTTP-POST binding to be used when IDP sends SAML response to SP via user agent. This information will be available in DSR / SDS metadata.
- No notification is sent to IDP when user accessing DSR / SDS GUI is voluntarily or forcefully logged out from DSR / SDS GUI.
- DSR / SDS will not report any error if SAML response / assertion is not received. Few instances are:
 - User agent does not redirect SAML authentication request to IDP.
 - IDP does not send SAML response / assertion.
 - User agent does not redirect SAML response / assertion to DSR / SDS.
 - No communication between IDP and user agent.
- Few instances of SAML authentication login failure are given below. Note this is not the complete list of potential failures.
 - SAML feature is not enabled.
 - IDP metadata is not uploaded.
 - IDP metadata is not syntactically correct.
 - Required parameters to send SAML authentication request are missing in IDP metadata.
 - SAML response does not contain required parameters.
 - Failed SAML response is received from IDP.
 - Authenticated user is not created on DSR / SDS.
 - Authenticated user account is disabled on DSR / SDS.
- Below URL is used to access GUI for SAML based authentication. Each NOAM/SOAM has different IP and will be authenticated separately.
 - `https://<XMI OR VIP IP>?auth=SAML`

DSR / SDS has no control over communication between IDP and user agent.

Note

IDP configuration OR configuring DSR or SDS metadata on IDP is out of scope of DSR and SDS.

Figure 2-67 SAML Call Flow



2.17.3 Enabling SAML Authentication functionality

Enabling SAML Authentication functionality will allow SAML authentication of users. Use this procedure to enable the SAML functionality.

Note

This procedure pertains to GUI access only.

- Click Administration, and then General Options.
- Set SAML Enabled parameter to 1 to enable SAML functionality.

Note

By default, SAML Enabled parameter is 0 (i.e. disabled).

2.17.4 Disabling SAML Authentication functionality

Use this procedure to disable the SAML functionality.

Note

This procedure pertains to GUI access only.

- Click Administration, and then General Options.
- Set SAML Enabled parameter to 0 to disable SAML functionality.

2.17.5 Viewing SAML Authentication functionality

Viewing SAML Authentication functionality

Use this procedure to view SAML authentication page.

1. Click on Administration
2. Then click on Remote Servers
3. Navigate to SAML Authentication
4. The SAML Authentication page lists Entity Id of an IDP Server and allows to insert Metadata of IDP

2.17.6 Uploading IDP Metadata

Uploading IDP Metadata

This procedure defines the automated process of uploading a IDP Metadata configuration file for SAML authentication. Use this procedure to upload an XML file to configure SAML authentication.

Note

This procedure pertains to GUI access only. User can insert IDP Metadata using upload functionality only. Click Administration, and then Remote Servers, and then SAML Authentication.

- Click Browse to locate the file you want to use to configure a IDP Metdata. A file upload screen displays allowing you to navigate to and select the target configuration file.
- Select the target file and click Open.
 - Only XML file will be supported.
 - Ensure that the file name length including .extension is restricted to 255 characters.

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

- Click Upload File.
- The file is uploaded and data validation is performed.
- Metadata validation is performed immediately. If the file is valid, then IDP Metadata will be inserted and Entity Id will be displayed on the page. Alternately, a file that contains invalid parameters returns an error message, and IDP Metadata will not be inserted.

Note

The maximum number of IDP Metadata that can be inserted is 1. If user wants to insert new IDP Metadata, then first he needs to delete the existing one and after that he can upload a new Metadata.

2.17.7 Deleting IDP Metadata

Deleting IDP Metadata

Use this procedure to delete a IDP Metadata:

- Click Administration, and then Remote Servers, and then SAML Authentication.
- Select the appropriate Entity Id from the table listing.
- Click Delete.
- Click OK to delete IDP metadata.

The IDP Metadata is deleted from the database and the Entity Id entry will no longer appear in the table listing.

2.17.8 DSR/SDS Metafile

DSR/SDS Metafile

DSR/SDS needs to provide the metadata file for configuration on the IDP.

- Sample data file needs to be added in the user guide document in Appendix.

Figure 2-68 DSR Metadata

```

DSR Metadata
<EntityDescriptor ID="SM38148aa4977a48e7cc446a01f6ba0c02f97179aea5a" entityID="https://oracle.com"
xmlns="urn:oasis:names:tc:SAML:2.0:metadata">

  <SPSSODescriptor AuthnRequestsSigned="false" WantAssertionsSigned="true"
ID="SM6d552c1f7fb6ed52383838e24696ebe501724ae5936" protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">

    <NameIDFormat>urn:oasis:names:tc:SAML:2.0:nameid-format:persistent</NameIDFormat>

    <AssertionConsumerService isDefault="true" index="0" Location="https://10.75.236.49"
Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"/>

  </SPSSODescriptor>
</EntityDescriptor>


```

2.17.9 Measurement Data Streaming (MDS)

Measurements Data streaming solution provides DSR's statistics to remote server communicating through Apache Kafka servers in the network. Using Kafka producer, MDS feature sends selected DSR measurement statistics to Kafka server. Measurements collected for all the DSR components (A, B, C levels) at NOAM and SOAM are sent to remote Kafka server in JSON format. Hence, avoiding users to enter directly into the DSR system and

interpret measurements. Statistics, which are pushed over to Kafka server can be used for DSR system analysis.

Figure 2-69 DSR Mediation with MDS



The MDS collects user-enabled measurement statistics data and sends the collected data to a user-configured Kafka cluster at user configured intervals. Additionally, MDS stores any unsent data on disk in the event of a connection loss with the Kafka cluster. It attempts to resend this unsent data once the connection to the Kafka cluster is reestablished, provided it occurs before the expiration of the user-configured retention period. Kafka connection properties and DSR properties can be configured at NOAM, but measurements can be enabled at both NOAM and SOAM.

Selection of measurements

There are many measurements for DSR nodes, together, which produce a large quantity of data. Since it takes more time to generate this data and high CPU power to process it, MDS feature provides an option to stream only the required measurements by enabling them. Using NOAM GUI, the user can enable measurements whose scope is A(A level server). Using SOAM GUI, the user can enable measurements whose scope is DSR component B(B level server).

Interval

This property defines the regular interval at which MDS collects, parse, and sends the statistics data of measurements to the remote Kafka server. The user can configure this property from NOAM GUI. If the transfer of measurements data to the Kafka cluster is not completed within the user configured interval, an event will be triggered.

The different Interval Options are:

Table 2-13 Interval

Interval	Default
5 minutes	Yes
15 minutes	-

Table 2-13 (Cont.) Interval

Interval	Default
30 minutes	-

Retention

MDS feature ensures that there is no data loss when there is a disconnection or unable to establish a connection to both primary and backup Kafka clusters. It stores the unsent data into the MdsPendingRecordsA or MdsPendingRecordsB tables.

MDS sends this unsent data to Kafka server if connection to Kafka reestablishes before the expiry of retention period. MDS deletes the old unsent data from MdsPendingRecordsA or MdsPendingRecordsB after the expiry of configured retention period. MDS provides the following options for maximum retention time:

Table 2-14 Retention

Data	Table Name
Unsent data of A-level metrics	MdsPendingRecordsA
Unsent data of B, C –level metrics	MdsPendingRecordsB

Table 2-15 Retention time

Retention Time	Default
No Retention	-
5 minutes	-
15 minutes	-
30 minutes	Yes
1 hour	-

Note

“No Retention” option disables data retention in case of send failure.

Security

MDS feature supports the users to send statistics data to remote Kafka server securely. It allows the users to configure any one of the following protocols:

- Additional configurable properties required by the security protocols are listed in the MDS Configuration.
- Kafka properties, refer to *Diameter Signaling Router Operations, Administration and Maintenance Guide*.

Data Size and Data Format

The dimension of measurements can be single or arrayed. Since the arrayed measurements can contain the data in large quantity, MDS provides "Max Records Size" option to configure the maximum size of the message. MDS divides the large data into segments based on the configured "Max Record Size" value and adds segment information to each segment.

Table 2-16 Segment Field

Info	Description
"segmentInfo": "F0"	Indicates the first segment.
"segmentInfo": "In"	'I' indicates the intermediate segment and 'n' is the segment number.
"segmentInfo": "Ln"	'L' indicates the last segment and 'n' is the segment number.

Note

If message is not segmented it contains only one segment, then segmentInfo header will contain only L0 segmentInfo: "L0" indicates that the message contains only one segment.

The user can configure the "Batch Size" Kafka property, based on which the Kafka producer collects multiple messages together until the batch size is reached and sends it through the network as one batch.

Every full and split message will contain a header with the following characteristics:

Table 2-17 Header characteristics

Name	Description
hostname	Server name defined by the server.
node	Contains a detailed server type, example: - MP_DA-MP - SYSTEM_OAM_SOAM - MP_IPFE
date	Indicates when the measurement was taken.
interval	Indicates the interval in which the measurements are collected.
metric	Name of the measurement.
group	Name of the group, which the measurement belongs.
type	Indicates the measurement type, example: - Single - Average - Max - Min
dim	Indicates the measurements dimension, example: - Single - Arrayed
segmentInfo	Indicates the segment number.
Total	It contains actual statistics data of measurement. This is only used when dim is single.
Sets	It contains actual statistics data of measurement. This is only used when dim is arrayed and contains all elements and their values.

Example of Single dimensioned measurement record:

```
[{
  "hostname": "MDS-NOAM",
  "node": "NETWORK_OAMP",
  "date": "04/03/2023 13:30 ",
  "interval": "13:30:00-13:34:59",
  "dim": "Single",
  "metric": "System.Disk_UtilPct_Peak",
  "group": "OAM.SYSTEM",
  "type": "max",
  "segmentInfo": "L0",
  "TOTAL": "0",
}]
```

Example of Arrayed measurement record without applying "Max Record Size" limit:

```
[{
  "hostname": "DSRTESTS0-DS000",
  "node": "SYSTEM_OAM",
  "date": "04/03/2023 13:30",
  "interval": "13:30:00-13:34:59",
  "metric": "TxRequestSuccessAllConn",
  "group": "Diameter Egress Transaction",
  "type": "simple",
  "dim": "Arrayed",
  "segmentInfo": "L0",

  "sets":
  [
    {"id": "Connection_1", "TOTAL": "3"},
    {"id": "Connection_2", "TOTAL": "5"}
    -----
    {"id": "Connection_20", "TOTAL": "6"},
    {"id": "Connection_21", "TOTAL": "2"}
    -----
    {"id": "Connection_30", "TOTAL": "1"},
    {"id": "Connection_31", "TOTAL": "51"}
    -----
    {"id": "Connection_n", "TOTAL": "10"}

  ]
}]
```

Example of Arrayed measurement record after applying 2KB record size limit

The default value for the *Max Record Size* is 2KB for the parameter Record 1.

Record 1:

```
[{
  "hostname": "DSRTESTS0-DS000",
  "node": "SYSTEM_OAM",
  "date": " 04/03/2023 13:30",
```

```

    "interval": "13:30:00-13:34:59",
    "metric": "TxRequestSuccessAllConn",
    "group": "Diameter Egress Transaction",
    "type": "simple",
    "dim": "Arrayed",
    "segmentInfo": "F0",
    "sets": [
      { "id": "Connection_1", "TOTAL": "3" },
      { "id": "Connection_2", "TOTAL": "5" }
      -----
      { "id": "Connection_20", "TOTAL": "6" },
    ]
  }
}

```

Record 2:

```

[ {
  "hostname": "DSRTESTS0-DS000",
  "node": "SYSTEM_OAM",
  "date": "04/03/2023 13:30",
  "interval": "13:30:00-13:34:59",
  "metric": "TxRequestSuccessAllConn",
  "group": "Diameter Egress Transaction",
  "type": "simple",
  "dim": "Arrayed",
  "segmentInfo": "I1",
  "sets": [
    { "id": "Connection_21", "TOTAL": "2" }
    -----
    { "id": "Connection_30", "TOTAL": "1" },
  ]
}
]

```

Record n:

```

[ {
  "hostname": "DSRTESTS0-DS000",
  "node": "SYSTEM_OAM",
  "date": "04/03/2023 13:30",
  "interval": "13:30:00-13:34:59",
  "metric": "TxRequestSuccessAllConn",
  "group": "Diameter Egress Transaction",
  "type": "simple",
  "dim": "Arrayed",
  "segmentInfo": "Ln",
  "sets": [
    -----
    -----
    { "id": "Connection_n", "TOTAL": "10" }
  ]
}
]

```

2.17.9.1 Troubleshooting

1. In case the Measurement Data Streaming does not work, the alarms and events can help in investigating the problems. For more information, see *Alarms and KPIs*.
2. Check Network or Route Configuration on DSR NOAM or SOAM and Kafka servers, if Measurement Data Streaming is unable to establish a connection to Kafka cluster.
3. Check Kafka properties on Active DSR NOAM GUI, from **Main Menu** select **Measurement Streaming** and click **Stream Options**.
4. Verify configuration of the `/etc/hosts` file on CLI of Active DSR NOAM and DSR SOAM if configured.

2.17.9.2 Limitations

- In case of Kafka connection failure, if the pending records go beyond 2GB within the retention period, any new unsent records will not be stored in `MdsPendingRecordsA` or `MdsPendingRecordsB` tables.
- Pending records from `MdsPendingRecordsA` or `MdsPendingRecordsB` are attempted to be sent in the next interval once the latest records for the corresponding interval will be sent. Hence the order of records at Kafka server will be different.

Note

Measurement Data Streaming is tested with 50 measurement streams. It is recommended for the user to select up to 50 measurements for streaming.

Measurement Data Streaming Feature Activation guide captures following details:

- Procedures to activate or deactivate the MDS Feature.
- Enabling or disabling Measurements on DSR NOAM & SOAM.

Measurement Data Streaming user guide captures details of Kafka properties to be configured from NOAM GUI.

2.18 Signaling Transfer Point (STP) Virtual Network Function (VNF)

Signaling Transfer Point (STP) Virtual Network Function (VNF)


2.18.1 General

General

Signaling Transfer Point (STP) VNF aka Virtual Signaling Transfer Point (vSTP) allows operators to have integrated SS7 signaling platform supporting key functionalities of signaling gateway and advanced SS7 routing. Service providers can optimize the use of network resources, manage subscribers and also interoperate between networks with disparate technologies. STP functionality at DSR provides scalability, reliability, security, and flexibility, while also providing investment protection by enabling migration to next generation technologies based on diameter signaling.

vSTP functionality is implemented through dedicated message processor in the DSR network element as shown below:

Figure 2-70 DSR Architecture with STP VNF Message Processor



DSR supports STP VNF only for virtual deployment model. Representational State Transfer (RESTful) MMI's shall be used for all STP configuration. Administration, Measurements and Alarms menus are available for STP feature via SOAM and NOAM GUI.

2.18.2 Signaling Protocols

STP VNF shall support following Sigtran protocols:

M2PA

The MTP2 User Peer-to-Peer Adaptation Layer Protocol (M2PA) is designed for communications between peers(e.g., STP to STP), as well as IP enabled end points. The STP VNF supports the RFC version of the M2PA protocol, as well as draft version 6. M2PA is specifically designed to fully support the transport of SS7 Message Transfer Part (MTP) Layer 3 signaling messages over IP using the services of SCTP. This includes full MTP3 message handling and network management capabilities between any two SS7 nodes, communicating over an IP network. The M2PA protocol provides the following advantages over the M3UA and SUA protocol.

- Greater support for MTP2 features such as link proving, processor outage, and blocking.
- Support for full retrieval due to link failure or processor outage (reduced message loss on failure).
- Simple conversions of all SS7 messages into M2PA messages and vice versa (better gateway, higher performance).

M3UA Protocol

M3UA seamlessly transports SS7 MTP3 user part signaling messages over IP using SCTP. M3UA-connected IP endpoints do not have to conform to standard SS7 topology, because each M3UA association does not require an SS7 link. Each M3UA-connected IP endpoint can be addressed by an SS7 point code unique from the signaling gateway's point code. STP VNF provides M3UA without routing keys.

2.18.3 SCCP -GLOBAL TITLE TRANSLATIONS (GTT) -ANSI/ITU

SCCP -GLOBAL TITLE TRANSLATIONS (GTT) -ANSI/ITU

General

The SCCP Global Title Translations (GTT) feature uses the signaling connection control part (SCCP) to translate addresses (Global Titles) from signaling messages that do not contain explicit information allowing the message transfer part (MTP) to route the message.

Global Title Translation Provisioning

The STP VNF uses tables for performing global title translations. Each table points to another table. The following tables are used for SCCP routing and management.

- Translation Type (TT) Table
- Global Title Translation (GTT) Table
- MAP Table

Translation Type Table

The Translation Type (TT) table is used to direct the translation process to the proper GTT tables for translation and further routing or processing. The Translation Type table supports translation values from 0 to 255.

Global Title Translation Table

The Global Title Translation (GTT) Table contains the digits or ranges of digits that are used to translate the in bound MSU to either another node for additional global title translation (intermediate GTT) or the MSU's final destination (final GTT). The EAGLE GTT table allows up to 1,000,000 total GTT entries with a performance restriction of up to 200,000 GTT entries per TT. Each entry may be a single value or a range of values

For example, an inbound MSU that arrives with a translation type of 253 and the digits 3038258000 in the CdPA would be translated by the range 3038258-3038259

Table 2-18 Global Title Translation

TT	GTA	EGTA	PC	XLAT	RI
253	3137070	3137080	1-1-3	DPC	GT
253	3137254	3137258	-1-1-4	DPC	GT
253	3038258	3038259	1-1-1	DPCSSN	SSN

There are five possible results to a global title translation:

- DPC only, route on GT –This result indicates that the DPC should be altered with the new translated point code, and the message will be routed to that node for further translation. The address indicator routing flag remains set to “route on GT.” If the called party address contains a point code then it is also replaced with the new point code.
- DPC only, route on SSN –This result indicates that the final destination SSN is already in the called party address and, with the addition of the new translated DPC, the final destination of the message is known. The address indicator routing flag is set to “route on SSN.” The new point code becomes the DPC of the message. If the called party address contains a point code then it is also replaced with the new point code.

- DPC and SSN, route on SSN –This result indicates that the final destination SSN and DPC should be determined by translation. No further translation is required, and the message can be routed to its final destination. The translated point code is placed into the DPC, and the SSN should be placed in the called party address. The address indicator routing flag is set to “route on SSN.”
- New GT –This result indicates that the translation type in the called party address should be replaced with the new translation type from the translation. This also indicates that the DPC in the message should be altered with the new translated point code and routed to that node for further translation. The routing indicator flag should remain set at “route on GT.”

MAP Table

The MAP table provides the set of remote subsystems associated to a particular remote point code. Each table contains up to ten subsystems assigned to a particular point code. This table also provides timers used for the subsystem status test (SST) procedure and information for locating the replicated point code and subsystem for any particular SSN. An option is provided on a per point code basis to send an SST upon receipt of an MTP-RESUME to ensure the subsystem is indeed available.

Global Title Translations may result in a choice of up to eight node/subsystems (replicated subsystems). Routing between the replicated pairs is based upon the global title translation results, which are provisioned in the database

There are four routing possibilities:

- Solitary –the GTT has a single node subsystem
- Dominant –all traffic is routed to the primary node/subsystem(s) if it is available. If the primary becomes unavailable, the traffic is routed to the backup subsystem(s). If the primary subsequently becomes available again, traffic is then routed back to the primary node/subsystem(s).
- Load sharing –the load is shared equally between replicated subsystem.
- Load sharing –the load is shared equally between replicated subsystem.
- Combination Load Share/Dominant -allows a group of primary node/subsystem(s) to loadshare as a dominant group while having the flexibility to form secondary, tertiary, so on, node or subsystem groups

Advanced Global Title Translation Functionality

The DSR STP VNF provides advanced Global Title Translation functionality to meet specific network needs. The STP VNF provides the following advanced Global Title Translation capabilities:

Flexible GTT Load sharing

The Flexible GTT Load sharing feature enables STP operators to create multiple load-sharing groups using the same destination point codes and/or SubSystem Number. This typically helps applying different load-sharing rules for different traffic types involving the same network elements. Flexible Intermediate GTT Load Sharing.

- Thanks to the use of Mated Relay Node tables (MRN), the STP will allow multiple load-sharing groups for GT routed traffic (Routing Indicator set to '0')

Flexible Final GTT Load Sharing

- In case of Final Global Title Translation, Load-sharing rules can differ not only based on the Destination Point Codes but also based on the SubSystem Number of the incoming

MSU. Here, the use of multiple Mated Application (MAP) tables will enable full flexibility to the STP operator.

Weighted GTT Load Sharing

This feature allows unequal traffic load-sharing for either Final or Intermediate Global Title Translations. The MAP and MRN tables will offer additional parameters to be provisioned in order to define new rules:

- Individual weighting for each entity in a relative cost (RC) group.
- In-service threshold for each RC group.

An RC group is a group of entries in either the MAP or the MRN groups that have the same relative Sucha group can also be referred to as an entity set.

Transaction-Based GTT Load Sharing

Transaction-Based GTT Load Sharing feature allows messages with the same transaction parameters (TCAP, SCCP, MTP...) to be routed to the same destination within an entity set. An entity set is a group of entities that are used to determine the proper destination of a post-GTT message. This group of entities can be one of the following:

- A mated application (MAP) group.
- A mated relay node (MRN) group.
- A mated application set (MAPSET) if the Flexible GTT Load Sharing feature is enabled.
- A mated relay node (MRNSET) if the Flexible GTT Load Sharing feature is enabled.

The feature applies to the following types of SCCP messages: Class 0/Class 1 UDT/UDTS/XUDT/XUDTS.

Different keys derived from the signaling messages can then be used for the load-sharing:

- MTP Parameters – the first 3 bytes of the incoming OPC and 1 byte of the SLS.
- SCCP Parameters – the last 4 bytes of the global title address field of the Called Party Address.
- TCAP Parameter – the TCAP transaction Id in the messages.
- Enhanced MTP Parameters – a combination of the SLS and incoming OPC values.

Flexible Linkset Optional Based Routing

Flexible Linkset Optional Based Routing allows the STP-MP to route GTT traffic based on the incoming link set and to route GTT traffic based on a variety of parameters (MTP, SCCP and TCAP depending on features that are enabled and turned on) in a flexible order on a per-translation basis. This feature enables routing/screening to be performed based on Calling Party information as well as based on the origin linkset from which the message came in the STP. Typically used to differentiate routing/screening rules for international/national gateways from the intra-network linksets.

Flexible Linkset Optional based routing enables new GTTSET Types as follows:

- GTT Set types for CdPA GTT Selectors:
 - CdPA GTA,
 - CdPA SSDN,
 - DPC
- GTT Set types for CgPA GTT Selectors:

- CgPA GTA,
- CgPA SSN,
- CgPA PC,
- OPC

TCAP Opcode Based Routing

TCAP Opcode Based Routing allows the EAGLE to route messages based on their operation codes. The TCAP Opcode Based Routing feature uses the information contained in the TCAP portion of messages (Operation code) is used for performing global title translation. This feature supports both ITU and ANSI messages.

GTT Actions

GTT actions allows ability to do more than just routing and screening the messages in the GTT framework. This enables both black and whitelisting at SCCP layer based on CdPA/CgPA information.


Following actions are supported in STP-MP:

- DISCARD: Allows silent discard of incoming message. No response is generated.
- UDTs: Allows to discard the message and generate a UDTs with a customizable SCCP Error Code.
- TCAP Error: Allows to discard the message and generate response with a customizable TCAP Error Code.
- FORWARD: Allows to “intercept” signaling traffic and forward it to a different destination than the normal one (if this destination is available).
- DUPLICATE: Allows message copy over SS7 link done together (or not) with another action. Duplicated messages consume 1 additional ingress MPS.

Please refer Virtual Signaling Transfer Point (vSTP) User's Guide for more details on GTT capabilities.

2.18.4 Equipment Identity Register

To reduce the number of GSM mobile handset thefts EIR functionality providing a mechanism that will allow the network operators to prevent stolen or disallowed handsets from accessing the network. This control is done by using the International Mobile Equipment Identity (IMEI) provided during handset registration and comparing it against a set of lists provided by the network operator. Application may also provide an option to validate IMEI based on the IMSI (International Mobile Subscriber Identity). IMSI Range to be used for special handling for certain subscribers.

Figure 2-71 EIR solution architecture.

EIR is a DSR Application on DA MP and STP MP. Screening supported for IMEI Individual or Range (with possible IMSI mapping), TAC (provisioned as IMEI range), SVN, IMSI range, Status override possible with IMSI association. EIR only provides the device status information.

The UDR NO provides the functionality of the Equipment Identity Register (EIR) database to the DSR. The database stores white, gray, and black lists of IMEI numbers.

Key characteristics of a UDR NO are as follows:


- Supports an internal query interface towards DA-MP and STP MP for vEIR.
- Provides options of provisioning GUI as well as Bulk provisioning.
- Provides FTP and SFTP based provisioning. Provisioning interfaces supported are REST/ SOAP/XML/LDAP.

Data Types that are supported at EIR data base:

IMEI


- 100M Individual entries.
- 10M IMEI range/TAC.
- 1 IMEI up to 10 IMSI.
- SVN supported
- IMSI Range (1000)
 - To support whitelisting special subscribers.

Figure 2-72 SS7 EIR call flow



- MAP version 1, 2 and 3
- Messages
 - MAP_CHECK_IMEI
 - MAP_CHECK_IMEI Response
- Point codes to be supported for SS7
 - ANSI
 - ITU-I
 - ITU-N
- IMEI SV and IMSI are optional Parameters
- Messages with EIR SSN (default 9, but configurable) sent to EIR application for processing.
- Final GT routing supported

Figure 2-73 S13 EIR Call Flow



- Messages
 - ECR
 - ECA
- IMEI is Mandatory AVP
- IMSI is optional AVP
- IMEISV is optional AVP
- Diameter routing using ART

EIR leverages DSR architecture for Measurements and Alarms, Backup and restore, OAM and Congestion Control.

2.18.5 FLOBR (Flexible Linkset Optional Based Routing)

FLOBR (Flexible Linkset Optional Based Routing) features:

- When GTT mode is “FLOBR CDPA”, CDPA fields in the MSU shall be used for GTT selector search and GTT set shall be taken from “CDPA GTT SET Name” configured in the selector entry.
- When GTT mode is “FLOBR CGPA”, CGPA fields in the MSU shall be used for GTT selector search and GTT set shall be taken from “CGPA GTT SET Name” configured in the selector entry.
- When GTT hierarchy is “FLOBR CDPA and FLOBR CGPA”, GTT selectors shall be searched as defined in 1. If no selector match is found or CDPA GTTSET is not provisioned, GTT selectors shall be searched as defined in 2.

- When GTT hierarchy is “FLOBR CGPA and FLOBR CDPA”, GTT selectors shall be searched as defined in 2. If no selector match is found or CGPA GTTSET is not provisioned, GTT selectors shall be searched as defined in 1.
- If GTT selectors are not found as specified in 1, 2, 3 or 4 then VSTP will consider this as translation failure.
- With FLOBR, the user can provision a fallback option for each translation that tells us how to route an MSU under the following conditions:
 - Routing when subsequent search failed in FLOBR.
 - Routing when same GTT set name is referred more than once.
 - Limiting the number of database searches to 7 for FLOBR.
- Under the above conditions:
 - When fallback option in last matched translation is set to “No”, the GTT will fail and MSU shall be discarded.
 - When fallback option in last matched translation is set to “Yes”, the GTT will be performed based on that matched entry.

2.18.6 MBR (MAP Based Routing)

MBR provides VSTP with the ability to route messages based on their 'MAP Components'. This can be done by adding 5 new GTT set types. These new GTT set types will be linked by OPCODE set type or any of them.

- IMSI
- MSISDN
- VLRNB
- SMRPOA
- SMRPDA

The GTT Sets of the types mentioned above are allowed to be provisioned ONLY in GTA entries from a GTT Set of the type OPCODE or any of the other GTT Set types supported by this feature.

Only TCAP Package Types BEGIN, CONTINUE & END are going to be supported for MAP based routing, so “optsn” with one of the MAP GTT Set types (listed above) are allowed to be provisioned only for TOBR GTA entries that have “pkgtype” as BGN or CNT or END.

2.18.7 TOBR (TCAP Opcode Based Routing)

TOBR (TCAP Opcode Based Routing) features:

- If the message/package type is NOT one of those mentioned in the list above slide, VSTP will treat it as an unknown message type and will not proceed with the decoding.
- As part of TOBR feature, VSTP will attempt to decode TCAP portion of all UDT/UDTS/ Unsegmented XUDT/Unsegmented XUDTS queries coming to SCCP layer for GTT.
- If decoding fails, the message will still undergo GTT using some default values for the TCAP data to denote their absence in the message.
- ACN will be used for all supported ITU TCAP messages except ABORT. No attempt to retrieve ACN will be made for Abort messages. All other supported messages may have a Dialog portion containing Dialogue Request / Unidirectional Dialogue/ Dialogue Response

PDU, from which the ACN will be retrieved. If no Dialog portion is detected, then ACN is assumed to be NONE.

- TOBR will attempt to find Operation Code (Opcode) in all supported ITU TCAP messages except ABORT. These messages must contain Invoke or Return Result (Last or Not Last) as the first component. If not, Opcode is assumed to be NONE.
- TOBR will attempt to find Operation Family and Specifier in all supported ANSI TCAP messages (except ABORT) containing an INVOKE component. For all other messages, Family and Opcode are assumed to be NONE.

2.18.8 IDP Query

IDP query features:

- Message flow for INPQ Solution on vSTP.
- MSC will send INPQ request to vSTP-MP over SS7 links.
 - vSTP-MP will decode and verify the INPQ Message.
 - vSTP-MP will decode and verify the INPQ Message.
 - Check whether INPQ message has valid request (the requestedInfo parameter must be MNP Requested Info and/or Location Information).
 - Decode the MSISDN parameter from the Subscriber Identity parameter.
 - Condition the MSISDN to the international format.
- vSTP-MP will query the UDR NOAM for conditional MSISDN DB.
- UDR NOAM will look up MSISDN DB and will send response to the vSTP-MP.
- Determine whether the lookup is considered to be successful based on provisioned options. If yes, use entity information to encode INPQ ACK response and route the response to the originator. If no, send INPQ NACK response with appropriate error code.

2.18.9 Loop Detection

Loop detection features:

- The Loop Detection feature design is both simple and flexible by using a Loopset table consisting of up to 12 pointcodes (OPCs) and comparing only the OPC of an MSU on a per GTT Translation basis requires the GTT translation to have a loopset provisioned.
- No assumptions/rules are imposed, everything is provisional whether concurring or opposite direction defines a "Loop", whether Customer using Capability Point Codes or True or Secondary Point Codes, which node is Adjacent and which is not, which node is mate and which is not.
- Detects looping on "multi-hop" translations and "single-hop" ones.
- Allows provisioning "Loops" for each Global Title Translation (GTA record).
- Allows provisioning/re-using same Point Codes in different LoopSets.
- Allows re-using same LoopSets for different Global Title Translations (GTA records).
- Allows provisioning "Loops" for each GTT Action record.
- Allows re-using same LoopSets for different GTT Action records.

2.18.10 Multiple PC

Multiple PC features:

- This feature proposes to start support for Secondary Point Codes and increase the number of supported Capability point codes.
- Multiple Point Code (MPC) feature shall allow the vSTP to add the support of Secondary Point Codes (SPCs) in addition to the true point codes used by the vSTP in any of the three domains ANSI, ITUN/ITUN-Spare (14-bit or 24-bit) and ITUI/ITUI-Spare.
- This is a different concept from capability point codes. The provisioning and routing will use secondary point codes as if they were the actual point code of the vSTP. SPCs are supported for any type of link (A, B, C, D, so on).
- In addition to the one True Point Code (TPC) already supported for each of the ANSI, ITU-N (14-bit or 24-bit) and ITU-I domains, the vSTP support a pool of Secondary Point Codes (SPC), each of which may be assigned as either ANSI, ITUI/ITUI-Spare, 14-bit ITUN/ITUN-Spare, or 24-bit ITUN.
- In addition to the SPCs this feature also recommends increasing the number of CPCs supported by the vSTP.
- Currently only 1 CPC is supported. This number shall be increased to 100. The increase CPC would allow the vSTP to support more applications on the same node and allow route-on-gt functionality for routing to various application.
- MPC feature creates the framework to allow provisioning and usage of SPCs in vSTP. This change shall primary affect the provisioning and routing algorithms. In addition to the SPC's the CPC shall also be increased.

TUI-Spare/ITUN-Spare

- The vSTP to fully support ITU National Spare and ITU International Spare by providing a new PC sub type named Spare.
- ITUN_S and ITUI_S indicates a Spare point code. Spare point codes only apply to ITU-I and ITU-N point code types.

Figure 2-74 ITUI-Spare/ITUN-Spare

NI	NAME
00(03)	International
01(43)	International-Spare
10 (83)	National
11(c3)	National-Spare

- The subservice field contains the network indicator and two spare bits. The network indicator is used by signaling message handling function.
- The vSTP currently provides full support for four types of point codes (PC) – ANSI, ITU-National (NI=10 binary), ITU-National 24-bit, and ITU-International (NI=00 binary). ITU National Spare PCs (NI=11 binary) and ITU International Spare (NI= 01 binary) PCs can be primarily supported.

- If the Spare PC feature is enabled then vSTP will validate the NI value of incoming message on designated linkset . If NI value matches then link will be mark available otherwise message will be discarded.

2.18.11 MTP Screening

MTP screening features:

- MTP Screening provides a first level of security check for VSTP.
- The MTP Screening feature examines the contents of a Message Signaling Unit (MSU) attempting to enter the VSTP against predefined criteria in the VSTP database to determine if the MSU should be allowed to enter.
- When a message comes on a linkset, the associated screen set will be looked up. Based on the NSFI and rule group name associated with screen set, the rule in the corresponding rule group will be looked up.
- If rule lookup is successful, then message will go for further screening level based on NSFI and the next screen rule group name associated with it.
- If rule lookup does not find a match, then:
 - In case of BLKOPC/BLKDPC rule type, default rule for that rule group will be looked up and based on NSFI and next screen rule group name associated with default rule further screening performed.
 - In case of OPC/SIO/DPC/AFTDSTN rule type, FAIL NSFI will be performed.
- The final result of MTP screening should always be either FAIL or STOP.
 - In case of FAIL, the message will be discarded.
 - In case of STOP, the message will go for further processing in VSTP.

2.18.12 ANSI ITU Conversion

Cross domain (ANSI/ITU) features

- Cross domain (ANSI/ITU) signaling can be MTP routed or GT routed.
- Cross domain conversion combination supported across domains are:
 - ANSI, ITU-I(S) ,ITU-N(S)
 - ITU-N24, ITU-I(S)
- Based on message type, ANSI/ITU Conversion supports UDT, XUDT, UDTS and XUDTS messages.
- The feature provides SCCP management (SCMG) across network type boundaries.
- This feature also provide support of China Point Code SCCP conversions to ITU-International.
- Without the GT Modification feature, VSTP performs the following translations on a per GTA basis: DPC and DPCSSN.
- Below table describes the fields modified within the Called party portion of the SCCP data of a MSU.
- The GTT Modification feature allows modification of all the fields of SCCP layer.
- The GTT Modification feature is activated by default that is no special activation process is required.

- The GTT Modification allows users to configure the Called Party / Calling Party modification data corresponding to each Global Title Address entry.

Table 2-19 Global Title Address entry

XLAT				
Fields Modified				
DPC	OPC	DPC	RI	
DPCSSN	OPC	DPC	RI	SSN

A

Appendix A: Supported Diameter Interfaces

Appendix A: Supported Diameter Interfaces.

The following list of Diameter Interfaces are all supported via the relay function on the DSR.

Table A-1 Appendix A

Count	ID Value	Name	Reference
1	0	Diameter common message	RFC6733
2	1	NASREQ	RFC7155
3	2	Mobile IPv4	RFC4004
4	3	Diameter base accounting	RFC6733
5	4	Diameter Credit Control	RFC4006
6	5	Diameter EAP	RFC4072
7	6	Diameter Session Initiation Protocol(SIP) Application	RFC4740
8	7	Diameter Mobile IPv6 IKE (MIP6I)	RFC5778
9	8	Diameter Mobile IPv6 Auth (MIP6A)	RFC5778
10	9	Diameter QoS application	
11	10	Diameter Capabilities Update	RFC6737
12	11	Diameter IKE SK (IKESK)	RFC6738
13	12	Diameter NAT Control Application	RFC6736
14	13	Diameter ERP	RFC6942
15	14-167772 15	Unassigned	
16	16777216	3GPP Cx	3GPP TS 29.28/3GPP TS 29.229
17	16777217	3GPP Sh	3GPP TS 29.328/SGPP TS 29.329
18	16777218	3GPP Re	3GPP TS 32.296
19	16777219	3GPP Wx	3GPP TS 29.234
20	16777220	3GPP Zn	3GPP TS 29.109
21	16777221	3GPP Zh	3GPP TS 29.109
22	16777222	3GPP Gq	3GPP TS 29.209
23	16777223	3GPP Gmb	3GPP TS 29.061
24	16777224	3GPP Gx	3GPPnTS 29.210
25	16777225	3GPP Gx over Gy	3GPP TS 29.210
26	16777226	3GPP MM10	3GPP TS 29.140
27	16777227	Ericsson MSI	
28	16777228	Ericsson Zx	
29	16777229	3GPP Rx	3Gpp TS 29.211
30	16777230	3GPP Pr	3GPP TS 29.234
31	16777231	ETSI e4	[ETSI ES 283034]
32	16777232	Ericsson Charging-CIP	
33	16777233	Ericsson Mm	

Table A-1 (Cont.) Appendix A

Count	ID Value	Name	Reference
34	16777234	Vodafone Gx+	
35	16777235	ITU-T Rs	[ITU-T Recommendation Q.3301.1]
36	16777236	3GPP Rx	3GPP TS 29.214
37	16777237	3GPP2 Ty	
38	16777238	3GPP Gx	3GPP TS 29.212
39	16777239	Juniper Cluster	
40	16777240	Juniper Policy-Control-AAA	
41	16777241	iptego USPI	
42	16777242	Convergence-specific SIP routing	
43	16777243	Policy Processing	OMA PEEM V1.0
44	16777244	Juniper Policy-Control-JSRC	
45	16777245	ITU-T S-TC1	ITU-T Recommendation Q.3221
46	16777246	NSN Unified Charging Trigger Function(UCTF)	
47	16777247	3GPP2 CAN Access Authentication and Authorization	
48	16777248	3GPP2 WLAN Interworking Access Authentication and Authorization	
49	16777249	3GPP2 WLAN Interworking Accounting	
50	16777250	3GPP Sta	3GPP TS 29.273
51	16777251	3GPP S6a	3GPP TS 29.272
52	16777252	3GPP S13	3GPP TS 29.272
53	16777253	ETSI Re	ETSI TS 183 060
54	16777254	ETSI GOCAP	ETSI ES 283 039
55	16777255	SLg	3GPP TS 29.172
56	16777256	ITU-T Rw	[ITU T Rec. Q.3303.3][RFC5431]
57	16777257	ETSI a4	ITU-T Rec. Q.3305.1
58	16777258	ITU-T Rt	
59	16777259	CARA	
60	16777260	CAMA	
61	16777261	Femtocell extension to Diameter EAP Application	
62	16777262	ITU-T Ru	ITU-T Rec. Q.nacp.RuQ.nacp.Ru]
63	16777263	ITU-T Ng	[ITU-T Rec. Q.nacp.RuQ.nacp.Ru]
64	16777264	3GPP SWm	3GPP TS 29.273
65	16777265	3GPP SWx	3GPP TS 29.273
66	16777266	3GPP Gxx	3GPP TS 29.212
67	16777267	3GPP S9	3GPP TS 29.215
68	16777268	3GPP Zpn	3GPP TS 29.109
69	16777269	Ericsson HSI	
70	16777270	Juniper-Example	
71	16777271	ITU-T Ri	ITU-T Rec. Q.3307.1
72	16777272	3GPP S6b	3GPP TS 29.273
73	16777273	Juniper JGx	
74	16777274	ITU-T Rd	ITU-T Rec. Q.3306.1

Table A-1 (Cont.) Appendix A

Count	ID Value	Name	Reference
75	16777275	ADMI Notification Application	
76	16777276	ADMI Messaging Interface Application	
77	16777277	Peter-Service VSI	
78	16777278	ETSI Rr request mode	[ETSI TS 183 071]
79	16777279	ETSI Rr delegated mode	ETSI TS 183 071
80	16777280	WiMAX HRPD Interworking	3GPP2X.S0058-0 v1.0
81	16777281	WiMAX Network Accounting Diameter Application (WNADA)	WiMAX Release 1.5
82	16777282	WiMAX Network Accounting DiameterApplication (WNADA)	WiMAX Release 1.5
83	16777283	WiMAX MIP4 Diameter Application(WM4DA)	WiMAX Release 1.5
84	16777284	WiMAX MIP6 Diameter Application(WM6DA)	WiMAX Release 1.5
85	16777285	WiMAX DHCP Diameter Application(WDDA)	WiMAX Release 1.5
86	16777286	WiMAX-Location-Authentication-Authorization Diameter Application(WLAADA)	WiMAX Release 1.5
87	16777287	WiMAX-Policy-and-Charging-Control-R3-Policies Diameter Application(WiMAX PCC-R3-P)	WiMAX Release 1.5
88	16777288	WiMAX-Policy-and-Charging-Control-R3-Offline-Charging DiameterApplication (WiMAX PCC-R3-OFC)	WiMAX Release 1.5
89	16777289	WiMAX-Policy-and-Charging-Control-R3-Offline-Charging-Prime Diameter Application (WiMAX PCC-R3-OFC-PRIME)	WiMAX Release 1.5
90	16777290	WiMAX-Policy-and-Charging-Control-R3-Online-Charging DiameterApplication (WiMAX PCC-R3-OC)	WiMAX Release 1.5
91	16777291	3GPP SLh	3GPP TS 29.173
92	16777292	3GPP SGmb	3GPP TS 29.061
93	16777293	CMDI - Cloudmark Diameter Interface	
94	16777294	Camiant DRMA	
95	16777295	PiLTE Interworking DiameterApplication	3GPP2 publication X.S0057
96	16777296	Juniper-Sessions-Recovery (JSR)	
97	16777297	Vedicis LiveProxy	
98	16777298	Pi*3GPP2 Diameter Application	3GPP2 publication X.S0057A E UTRAN eHRPD
99	16777299	Sandvine Rf+	
100	16777300	Subscription Information Application	
101	16777301	Ericsson Charging-DCIP	
102	16777302	3GPP Sy	3GPP TS 29.219
103	16777303	3GPP Sd	3GPP TS 29.212
104	16777304	Ericsson Sy	
105	16777305	HP DTD	
106	16777306	M9 interface between MLM-PE(P) andMLM-PE(C)	ITU-T Q5/Sg11
107	16777307	ITU-T M13	ITU-T Q.3230
108	16777308	3GPP S7a	3GPP TS 29.272

Table A-1 (Cont.) Appendix A

Count	ID Value	Name	Reference
109	16777309	3GPP Tsp	3GPP TS 29.368
110	16777310	3GPP S6m	3GPP TS 29.336
111	16777311	3GPP T4	3GPP TS 29.337
112	16777312	3GPP S6c	3GPP TS 29.338
113	16777313	3GPP SGd	3GPP TS 29.338
114	16777314	Intrado-SLg	
115	16777315	Ericsson Diameter Signalling Controller Application (DSC)	
116	16777316	Verizon-Femto-Loc	
117	16777317	Nokia Siemens Networks (NSN) Hd Application	
118	16777318	3GPP S15	3GPP TS 29.212
119	16777319	3GPP S9a	3GPP TS 29.215
120	16777320	3GPP S9a*	3GPP TS 29.215
121	16777321	Gateway Location Application	
122	16777322	Verizon Session Recovery	
123	16777323	3GPP2 M1 Interface	3GPP2 X.S0068
124	16777324	MAGIC Client Interface Protocol (CIP)	ARINC 839
125	16777325	ITU-T Nc	ITU-T Rec. Q.nacp.Nc
126	16777326	ITU-T Ne	ITU-T Rec. Q.nacp.Nc
127	16777327	Ericsson Sx	
128	16777328	Nokia Service Extension, NS	
129	16777329	Rivada X	
130	16777330	Rivada Xm	
131	16777331	Rivada Xh	
132	16777332	Rivada Xf	
133	16777333	Rivada Xp	
134	16777334	Rivada Xa	
135	16777335	3GPP MB2-C	3GPP TS 29.468
136	16777336	3GPP PC4a	3GPP TS 29.344
137	16777337	3GPP PC2	3GPP TS 29.343
138	16777338	Juniper Domain Policy	
139	16777339	Host Observer	
140	16777340	3GPP PC6/PC7	3GPP TS 29.345
141	16777341	Nokia Sdr Applicatio	
142	16777342	3GPP Np	3GPP TS 29.217
143	16777343	Sandvine Location Relay Service	
144	16777344	Sandvine Fairshare TrafficManagement Service	
145	16777345	3GPP S6t	3GPP TS 29.336
146	16777346	3GPP T6a/T6b	3GPP TS 29.128
147	16777347	3GPP Ns	3GPP TS 29.153
148	16777348	3GPP Nt	3GPP TS 29.154
149	16777349	3GPP St	3GPP TS 29.212
150	16777350	3GPP PC2	3GPP TS 29.343

Table A-1 (Cont.) Appendix A

Count	ID Value	Name	Reference
151	16777351	3GPP Diameter Data Management	3GPP TS 29.283
152	16777352	ITU-T M1	ITU-T Recommendation Q.nacp.M1
153	16777353	ITU-T M2	ITU-T Recommendation Q.nacp.M2