Oracle Communications Diameter Signaling Router

Release 8.4.0.5 Feature Guide

F12323 Revision 04

ORACLE
COMMUNICATIONS

ORACLE

Table of Contents

elop applications and business models of the future. To support this next gen tech. the intelligence hub of the 5G network is architected on the principles of Service Based Architecture (SBA) where different Network Functions (NFs) are deployed and seamlessly work together to provide various mobile network functionalities. To fully exploit the true capabilities of a 5G network the 3GPP defines a secure and reliable way to expose the network

.

List of Terms

Acronym	Meaning
ACL	Access Control List
API	Application programming Interface
APDE	Automatic Performance Data Export
ART	Application Routing Table
AVP	Attribute Value Pair
CLI	Command Line Interface
CPA	Charging Proxy Application
DA	Diameter Agent
DA-MP	Diameter Agent Message Processor
DAS	Diameter Application Server
DEA	Diameter Edge Agent
DIH	Diameter Intelligence Hub
DNS	Domain Name Server
DOIC	Diameter Overload Indication Conveyance
DP	Database Processor
DR	Disaster Recovery
DTLS	Datagram Transport Layer Security
EIR	Equipment Identity Register
ECC	Extended Command Code
EMS	Element Management System
EPC	Evolved Packet Core
FABR	Full Address Based Resolution
FQDN	Fully Qualified Domain Name
GUI	Graphical User Interface
HSS	Home Subscriber Server
IDIH	Integrated Diameter Intelligence Hub
ILO	Integrated Lights Out
IMI	Internal Management Interface
IMS	IP Multi-media System
IOT	Interoperability Tests
IWF	Interworking Function
KPI	Key Performance Indicator
LTE	Long Term Evolution

MEAL	Measurements, Events, Alarms, and Logging
MME	Mobility Management Entity
MP	Message Processor
MPS	Messages per Second
MAP-IWF	Map to Diameter Interworking Framework
NAI	Network Access Identifier
NE	Network Element
NGN-PS	Next Generation Network Priority Service
NMS	Network Management System
NRF	Network Repository Function
NFV	Network Fuction Virtualization
OAM	Operations, Administration, Maintenance
OAM&P	Operations, Administration, Maintenance and Provisioning
OC-DRA	Online Charging Diameter Routing Agent
OCF	On-line Charging Function
OCSG	Oracle Communications Services Gatekeeper
PAT	Pending Answer Timer
PCA	Policy and Charging Application
PCRF	Policy Control and Charging Rules Function
P-CSCF	Proxy-Call Session Control Function
P-DRA	Policy Diameter Routing Agent
PDU	Protocol Data Unit
PRT	Peer Routing Table
QS	Query Server
RBAR	Range Based Address Resolution
RD-IWF	Radius to Diameter Interworking
RSR	RADIUS Signaling Router
ROS	Routing Option Set
SCEF	Service Capability Exposure Function
SBR	Session Binding Repository
SDS	Subscriber Data Server
SLF	Subscriber Location Function
SS7 MP	Signaling System 7 Message Processor
TLS	Transport Layer Security
TTP	Traffic Throttling Point

UDP	User Datagram Protocol
UDR NO	Universal Data Repository Network OAM & Provisioning
VIP	Virtual IP Address
XMI	External Management Interface
XSI	External Signaling Interface
vSTP	Virtual Signaling Transfer Point
VNFM	Virtualized Network Function Manager

References

[1]   DSR Operation, Administration, and Maintenance (OAM) Guide – Available at Oracle.com on the Oracle Help Center (OHC)

[2]   DSR Alarms, KPIs, and Measurements – Available at Oracle.com on the Oracle Help Center (OHC)

[3]   Platform Feature Guide – Available upon request

[4]   Diameter Signaling Router (DSR) 7.0 Security Guide (E61125)– Available at Oracle.com on the Oracle Help Center (OHC)

[5]   DSR Licensing Information User Manual – Available on Oracle Help Center (OHC)

INTRODUCTION TO DIAMETER SIGNALING ROUTER

Mobile data traffic is skyrocketing, fueled by the introduction of smartphones, laptop dongles, flat-rate plans, social networking and applications like mobile video. Operators are looking to all-Internet protocol (IP) networks such as long term evolution (LTE) and IP multimedia subsystem (IMS) to provide the bandwidth required to support data-hungry devices and applications and to cost effectively address the growing gap between traffic and revenue growth.

The 3GPP Evolved Packet core (EPC) and IP Multimedia Subsytem (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 these interfaces are illustrated in the figure below. Diameter is also defined by 3GPP and ETSI standard bodies as the foundation for Authentication, Authorization and Accounting (AAA) functions in the Next Generation Network (NGN).
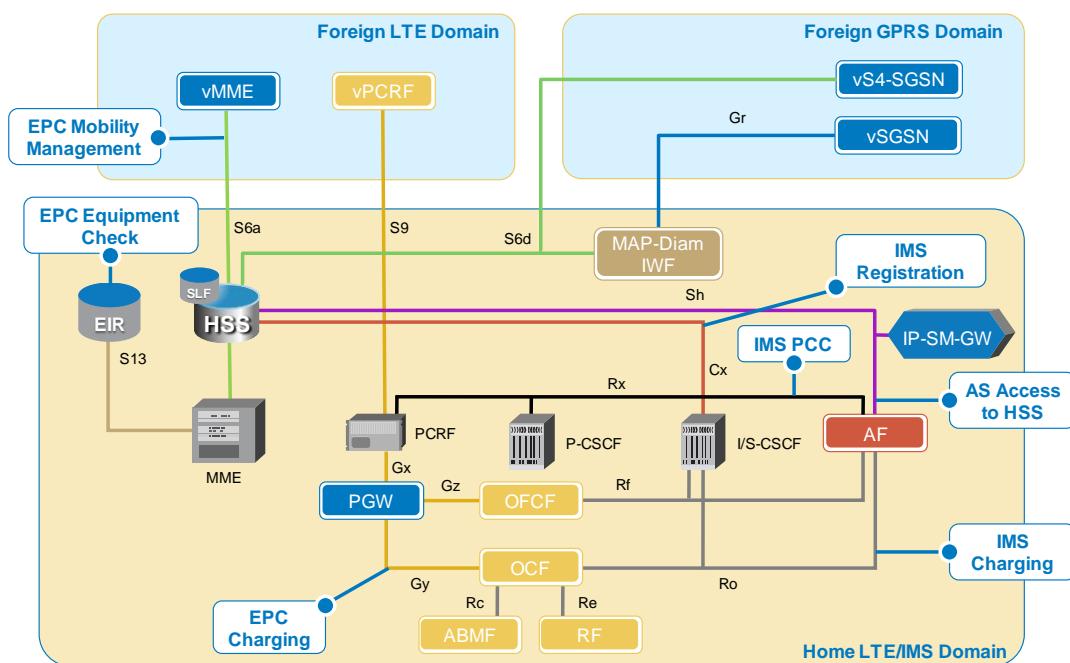
Figure 1 - Selected Diameter Interfaces in LTE and IMS

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 will expect the same network behavior and robustness as they enjoy with SS7 networks today.

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 as well as the status of each, placing a heavy burden on the endpoints as the number of nodes grows. This burden is made more complex with the responsibility of load balancing placed on each end point.
- » **Congestion control:** Diameter lacks the 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 interworking 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 interworking is required as transitions are made and LTE is 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 HSSs in the network, subscribers may be homed on different HSSs. 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.
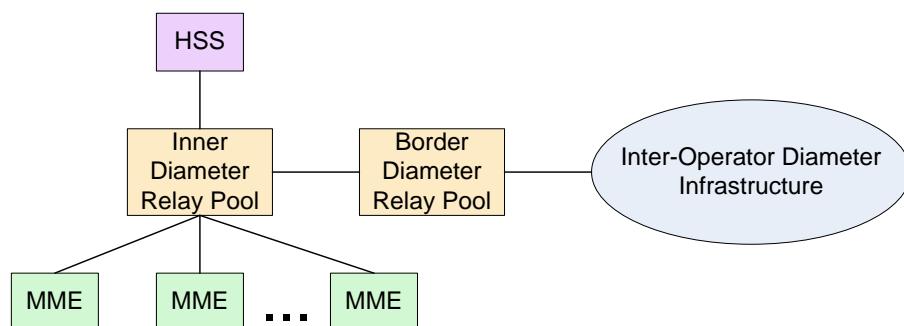
Figure 2 - 3GPP Inter/Intra-operator Diameter infrastructure
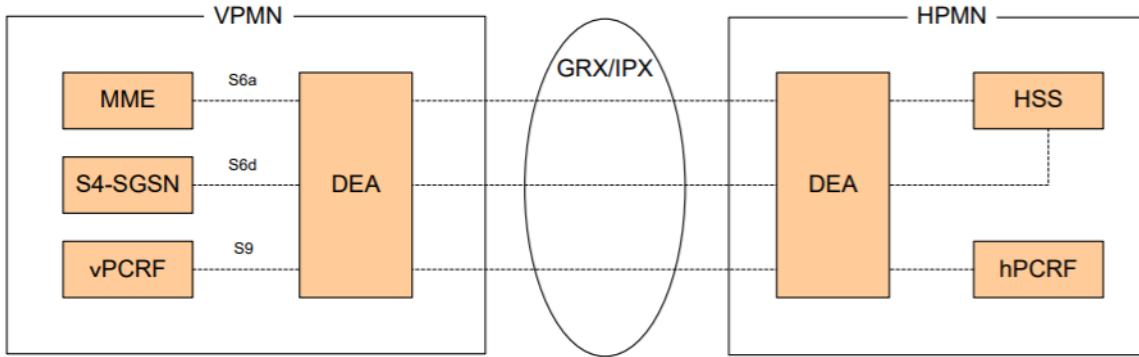
Figure 3 - GSMA roaming implementation architecture

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/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.

A centralized signaling architecture:

» 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); and

» 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/IMS core network with DSR.
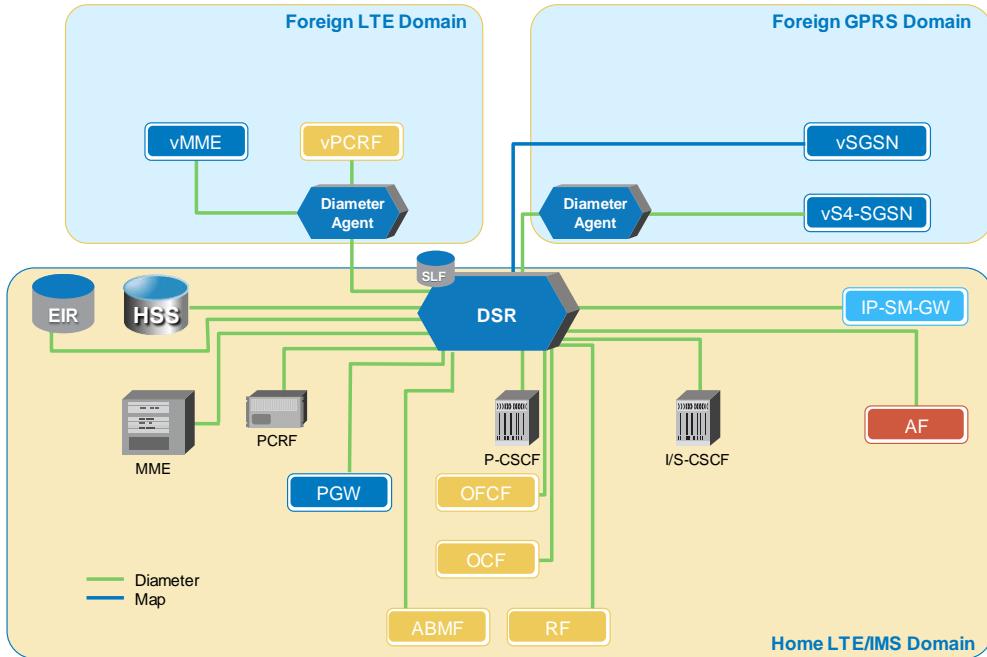
Figure 4 – Example of Operator's EPC/IMS Core network with DSR

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.

DSR FEATURES AND FUNCTIONS

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).

The figure below shows an overview of a DSR system architecture. Only single elements are shown for simplicity. 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
» Service Capability Exposure Function (SCEF) API Gateway (OCSG) – Optional
  » Gateway App Server
  » Gateway Admin Server
  » Gateway DB Server

List of non-compatible features in DSR 8.3: *
» Diameter Intelligence Hub (DIH).
» SS7-Message Processor (MAP-IWF).
» Gateway Location Application (GLA).
» Radius.

In rest of this document any reference to these non-complaint features should not be considered for feature enabling.

These components are described at a high 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 please see DSR Licensing Information User Manual – Available on Oracle Help Center (OHC)
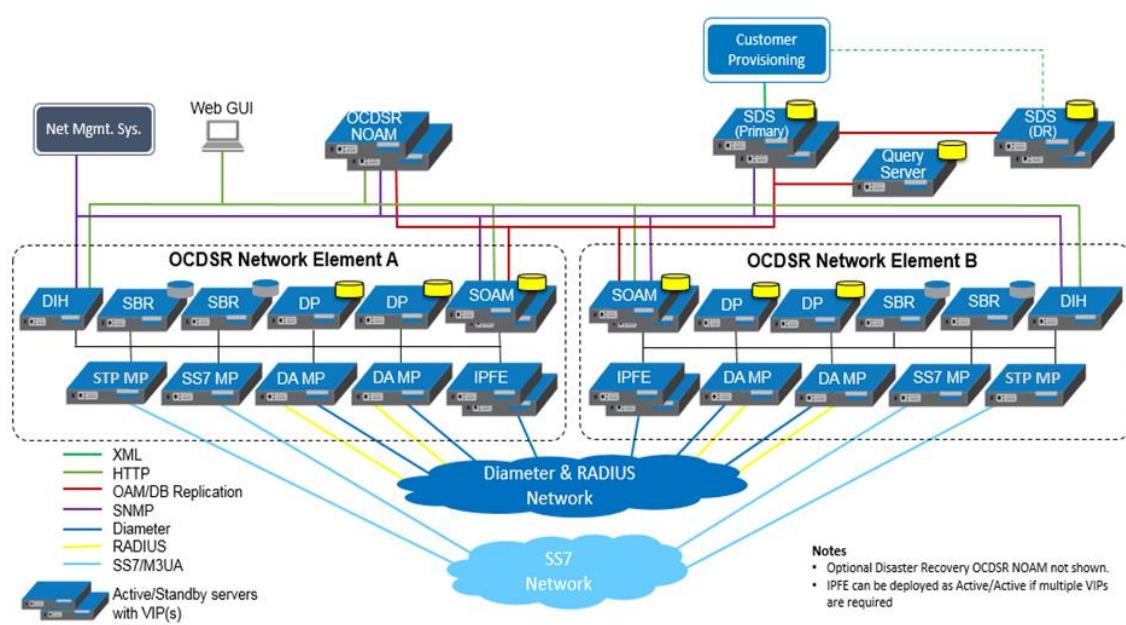
Figure 5 - DSR 7.x Architecture

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 and 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, etc.)

» 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 and access to measurements and reports

» centralized view of key operational metrics which identifies potential operational issues

Diameter Agent Message Processor (DA MP)

The DA MP hosts Proxy applications such as Address Resolution, Policy and Charging Application, Charging Proxy etc. 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 / IDIH

SS7 Message Processor (SS7 MP)

The SS7 MP provides the MAP to Diameter interworking function.

Key characteristics of an SS7 MP are as follows:

» performs message content conversion between MAP and Diameter.

» performs address mapping between SS7 (SCCP/MTP) and Diameter.

» supports 3G<->LTE authentication interworking as needed.

» interfaces with System OAM

STP Message Processor (STP MP)

The STP MP 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

IP Front End (IPFE)

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 / standby pairs with 3.2 Gbps bandwidth per active/standby pair

» Supported with SCTP Multi-homing

Session / Subscriber Binding Repository (SBR)

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 db to another: 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.

Subscriber Data Server (SDS)

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

Database Processor (DP)

The DP is the repository of subscriber data on the individual DSR node elements. The DP 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 DPs
» can also host other Oracle SDS based applications

Query Server (QS)

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

Integrated Diameter Intelligence Hub (IDIH)

The IDIH supports advanced troubleshooting for Diameter traffic handled by the DSR.   The IDIH is an optional feature of the DSR that enable the selective collection and storage of diameter traffic and provides nodal diameter troubleshooting.

DSR OAM&P

The DSR has a 3-tiered topology as described in the diagram below.

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 and process Diameter and RADIUS messages.

Network OAMP Layer

- Configuration and management of security, topology, servers and IP networking

System OAM Layer

- Configuration and management of DSR, Diameter, TCP and SCTP data

MP Layer

- Real-time signaling and database processing
- Source of DSR Measurement and Alarm data
- DA-MP, IPFE, SS7-MP class, SBR are all "MPs" managed by NOAMP
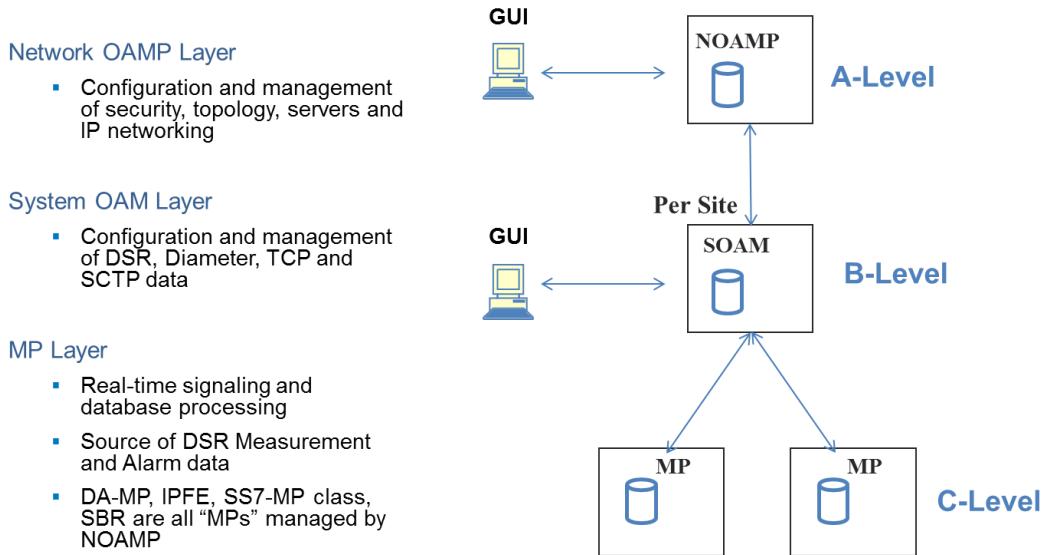
Figure 6 – DSR 3-tiered Topology Architecture

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: Interface to the DSR's internal management network. 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.

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.

Operations and Provisioning

» Operations and Provisioning of the DSR can be accomplished via one of the 10 GUI sessions that are made available to the User through an 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

- » HA (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 and data collection, and define HA 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).

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 and critical alarms and events for a wide variety of conditions. These are immediately sent up to the OAM system and can also be sent to the operator's network management system using SNMP. Alarm/event logs at the OAM are stored for up to seven days. The OAM provides a dashboard view of all alarms on the downstream MPs. This information is maintained locally for up to three days.
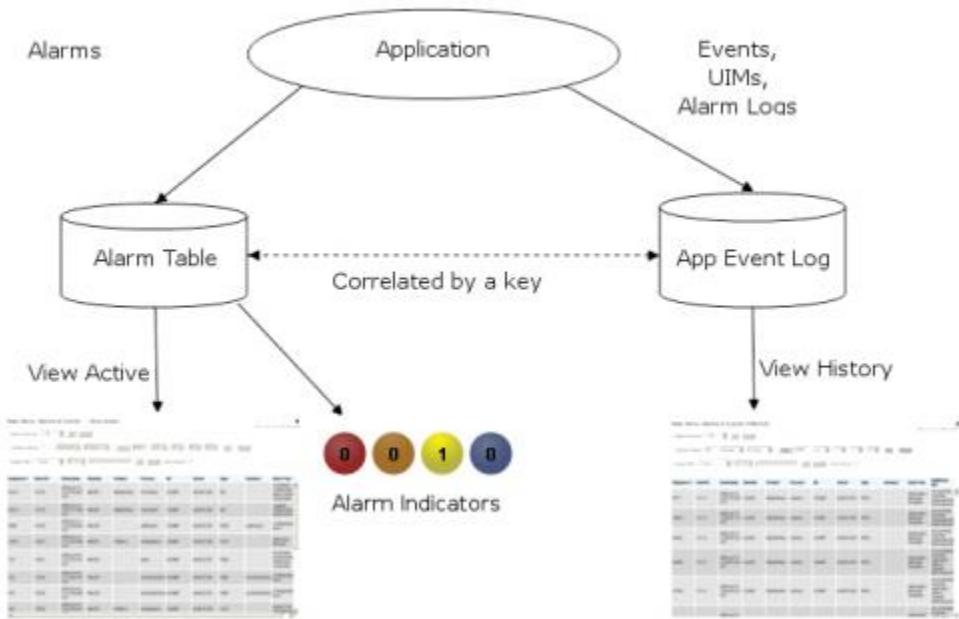
Figure 7 - Flow of Alarms

Below 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 .

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 1: DSR KPI SUMMARY

KPI Category	KPI Examples
Server Element KPIs	A group of KPIs that appear regardless of server role such as CPU, Network Element, etc
CAPM KPIs	Counters related to Computer-Aided Policy Making such as Active Templates, Test Templates, etc
Charging Proxy Application KPIs	KPIs related to the CPA feature such as CPA Answer Message Rate, CPA Ingress Message Rate, cSBR Query Error Rate, etc

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 Avg 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: 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 DSR Alarms, KPIs, and Measurements document found on the Oracle Help Center (OHC) area of http://docs.oracle.com.

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 and 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 measurements supported in DSR can be found in the DSR Alarms, KPIs, and Measurements document found on the Oracle Help Center (OHC) area of http://docs.oracle.com.

TABLE 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.
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
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 frm 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 IPFFE 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
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

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 & deletion of provisioned data. Both the import & export operations will generate log files.

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 via 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 8 - 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 (e.g., 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.

Automatic Performance Data Export (APDE)

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.

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 please see the Platform Feature Guide.

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 (i.e. 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.

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.

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 & 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 (e.g. configure a new OAM, initiate a backup, view alarm log)

Please see the Diameter Signaling Router (DSR) 8.3 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

Machine/Machine Interface (MMI)

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.

DSR Nodes (Identity)

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 Figure 9 for a sample configuration.
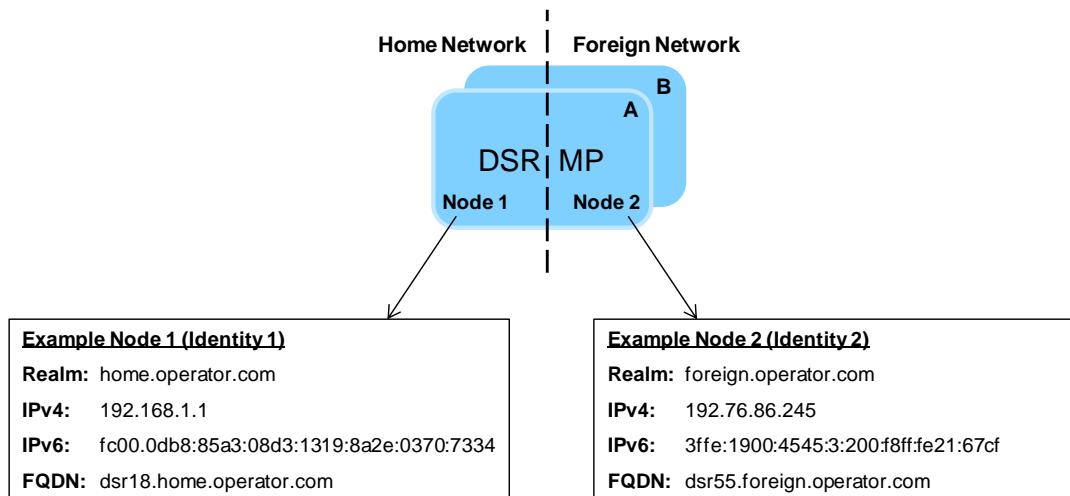
Figure 9 - Multiple Nodes per Message Processor

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.
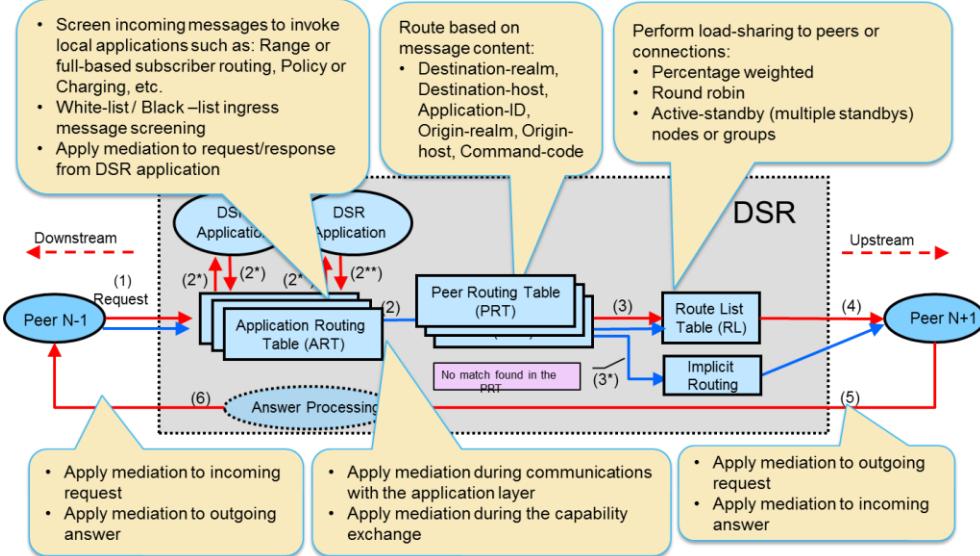
Figure 10 - High Level Message Processing and Routing in DSR

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 3 Route Groups with associated Route Group Priority level (1-3)
» 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/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 Rout 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
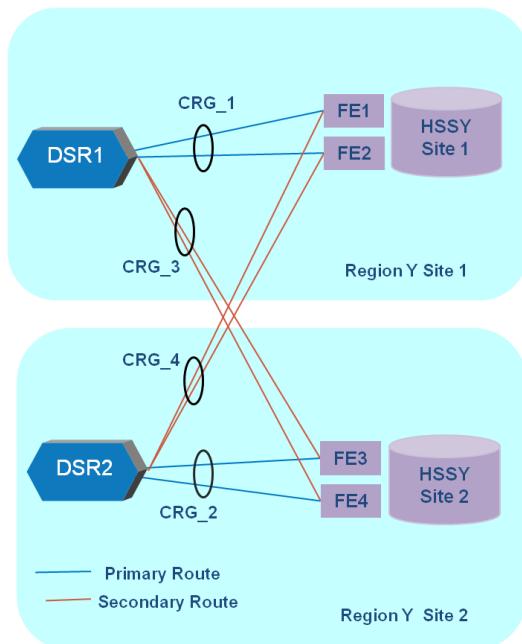
Figure 11 - Connection Route Group

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.

Route List Table

Route Group−1,Pri=1
Peer1, Wt=50
Peer2, Wt=50

Route Group−2,Pri=2
Peer3, Wt=60
Peer4, Wt=40

Route Group−3,Pri=1
Peer5, Wt=100

Route Group−4,Pri=2
Peer6, Wt=25
Peer7, Wt=25
Peer8, Wt=30
Peer9, Wt=20

Route Group−3,Pri=1
Peer5, Wt=100

Route List-1

Route List-2
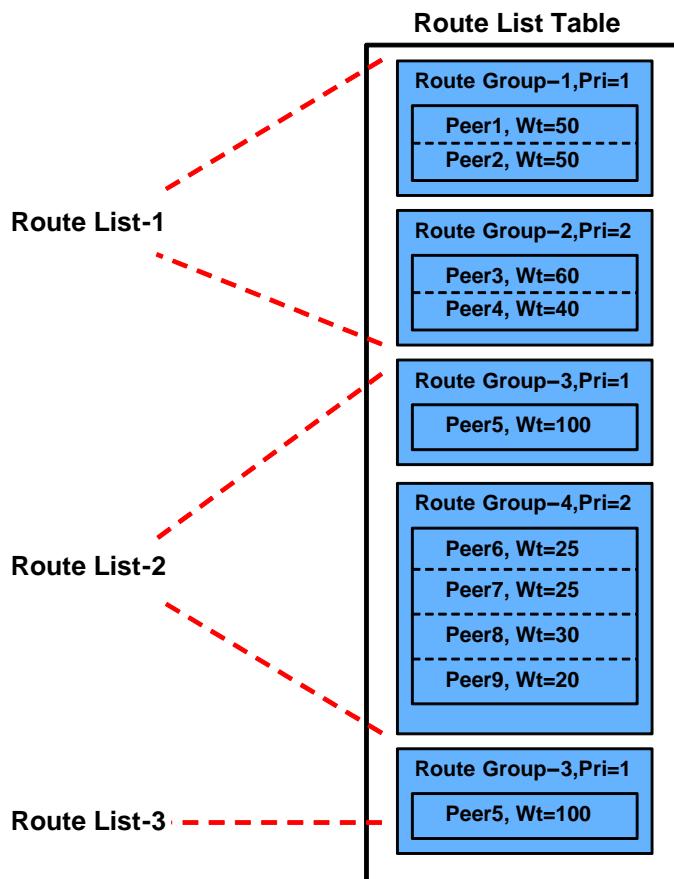
Route List-3

Figure 12 - 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. DSR supports provisioning up to 160 routes in a route group (same priority) and allows for provisioning of 3 route groups per route list.

DSR

W=40 → Peer1
W=30 → Peer2
W=30 → Peer3

Route Group-1
(Routes with Pri=1)

W=60 → Peer4
W=40 → Peer5

Route Group-2
(Routes with Pri=2)

W=50 → Peer6
W=50 → Peer7

Route Group-3
(Routes with Pri=3)

Route List-1

Route Group-1, Pri=1
Route-1, Pri=1, Wt=40, Peer=1
Route-2, Pri=1, Wt=30, Peer=2
Route-3, Pri=1, Wt=30, Peer=3
Route Group-2, Pri=2
Route-1, Pri=2, Wt=60, Peer=4
Route-2, Pri=2, Wt=40, Peer=5
Route Group-3, Pri=3
Route-1, Pri=3, Wt=50, Peer=6
Route-2, Pri=3, Wt=50, Peer=7

Figure 13 - Load Balancing Based on Route Groups and Peer Weights

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 of 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.

Extended Command Codes (ECC)

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).

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.
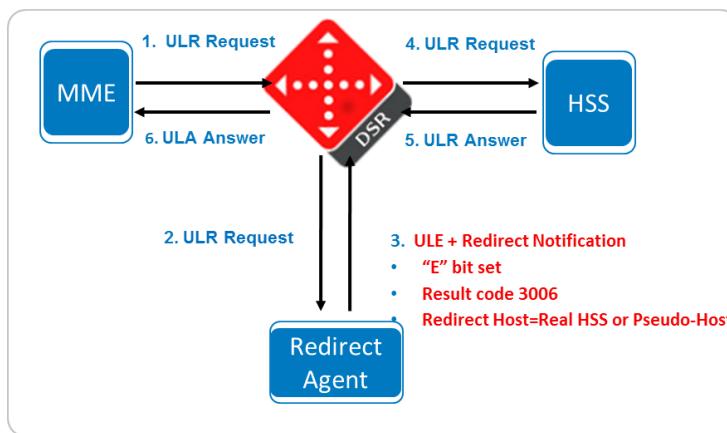
Figure 14 - Redirect Agent

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 4: Modified Routing and Transaction Parameter Selection Precedence Order

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

Note 1: For multiple DRA Application invocation on the same message, the applications can select a different ART and override the core routing ART precedence.

Note 2: Local DSR applications can select a different PRT and override this core routing PRT precedence

Note 3: 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.

Peer Routing Table (PRT)

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

Routing rules can be prioritized (1 – 1000) for cases where an inbound Diameter request may match multiple user-defined routing rules. 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 1000 rules and the total number of rules across all PRTs cannot exceed 50,000. 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:

- o   Route to Peer - use Route List Table

- o   Send Answer Response - an Answer response is sent with a configurable Result-Code and no further message processing occurs

- o   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 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	No	Yes	Don't Care	Don't Care	Yes

Selected Transaction Configuration Group					
Local App PRT	Yes	Don't Care	Don't Care	Yes	Yes

Application Routing Table (ART)

An application route table contains one or more application routing rules that can be used for routing request messages to DSR applications. Up to 400 application routing rules can be configured per application route table. Up to 1,500 application route tables can be configured per DSR network element; a total of upto 50,000 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:

- o Route to Application - route the message to the local Application associated with this Rule

- o Forward to Egress Routing - ART search stops and moves on to PRT

- o 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)

- o Forward to Application Route Table - ART forwards the request message to the specified Application Route Table.

- o 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 NodeApplication 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.

Routing Option Sets (ROS)

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): (Refer to Table 4.)

» 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. Otherwise,
» The Routing Option Set is assigned to the ingress peer node. Otherwise,
» The Routing Option Set is assigned to the default TCG. Otherwise,
» 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

Pending Answer Timer (PAT)

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 to16 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): (Refer to Table 4.)

» If the Transaction Configuration Group is select 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. Otherwise,
» The PAT from the ROS assigned to the ingress peer node is used. Otherwise,
» The PAT assigned to the egress peer node is used. Otherwise,
» The PAT assigned to the default TCG is used. Otherwise,

» The System default PAT is used.

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 & IPv6 networks but IPv4 and IPv6 cannot co-exist on the same connection.
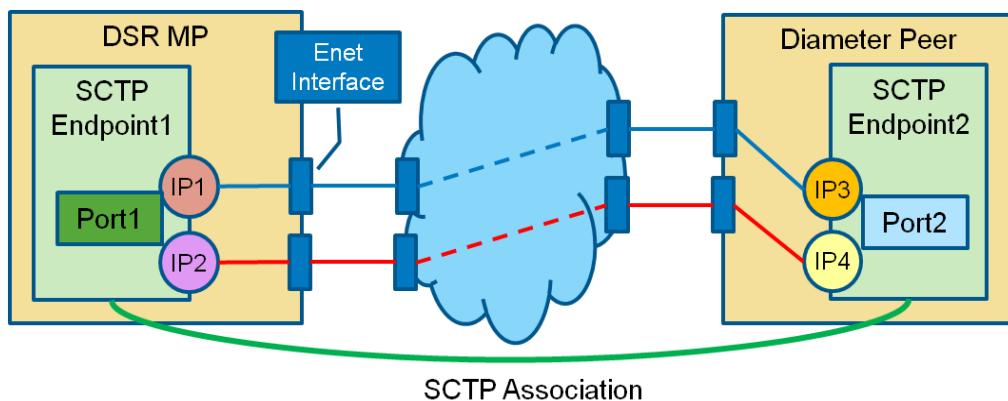
Figure 15 - SCTP Multi-Homing
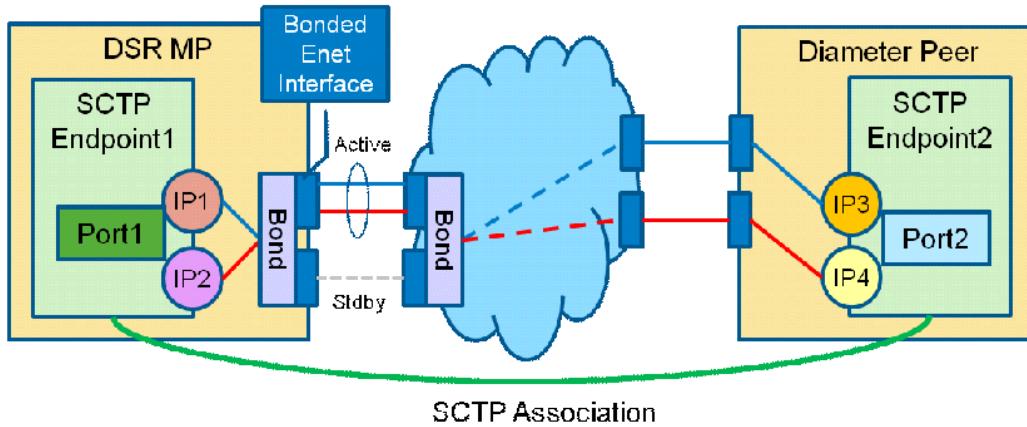
Figure 16 - SCTP Multi-Homing via Port Bonding

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.

Diameter Routing Message Priority (DRMP)

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 i.e. 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.

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 10 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

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.

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.

1. **initiator mode**: DSR discovering the last-hop Diameter peers
1. **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.
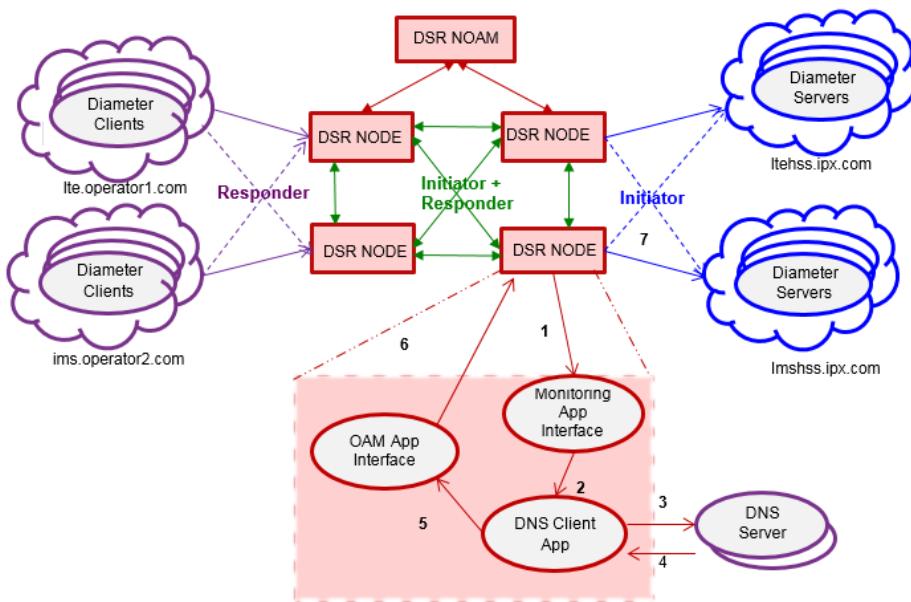
Figure 17 – Dynamic Diameter Peer Discovery: Example

In the above example for 'initiator' mode:

Each DSR node does the following:

1. monitors configuration changes

2. creates tags required for Diameter extended NAPTR (S-NAPTR) query

3. invokes DNS Client Application Interface for query resolution towards configured DNS Servers

4. provides DNS Client Application Interface, processes the DNS responses and resolves NAPTR, SRV, A, AAAA lookups

5. performs target server resolution mapping to Diameter peer attributes for specified realm

6. invokes OAM interface to update discovered Diameter peer attributes in DSR configuration managed objects

7. 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.

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.
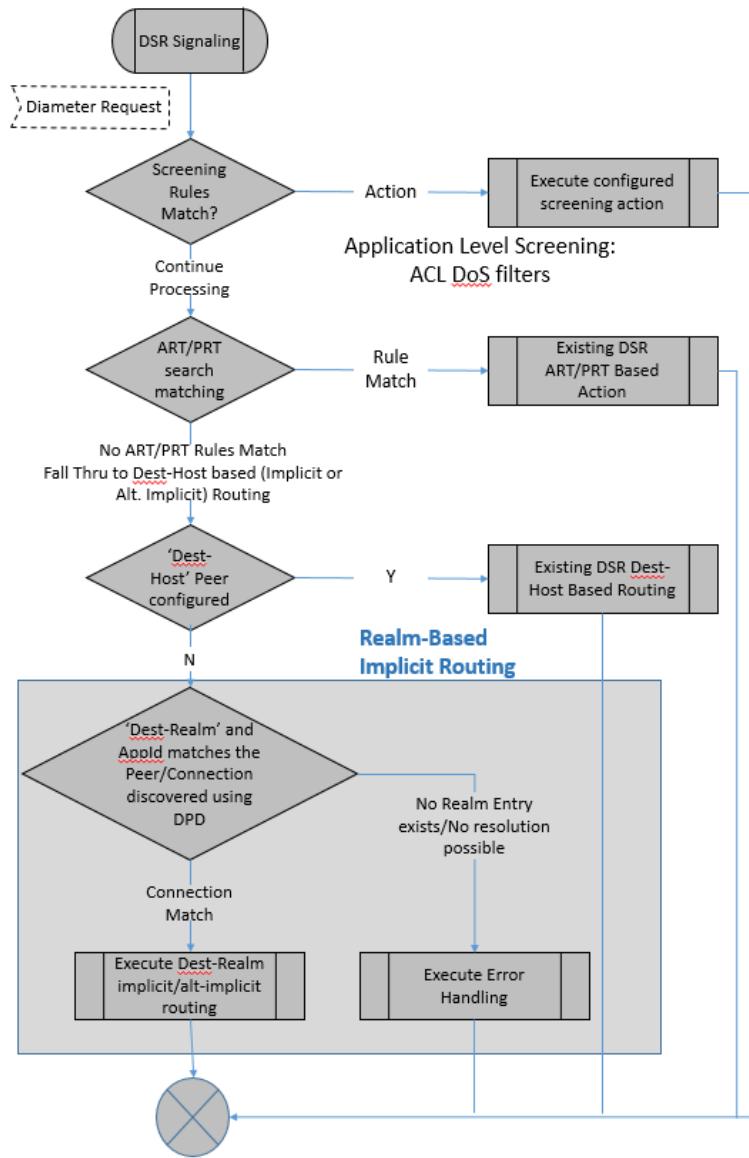
Figure 18 – High Level DSR Routing Flow – Fall through to Dest-Realm Based Implicit Routing

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.

Signaling Firewall

Signaling Firewall feature is the network security feature of DSR which configures native Linux 'iptable' 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.
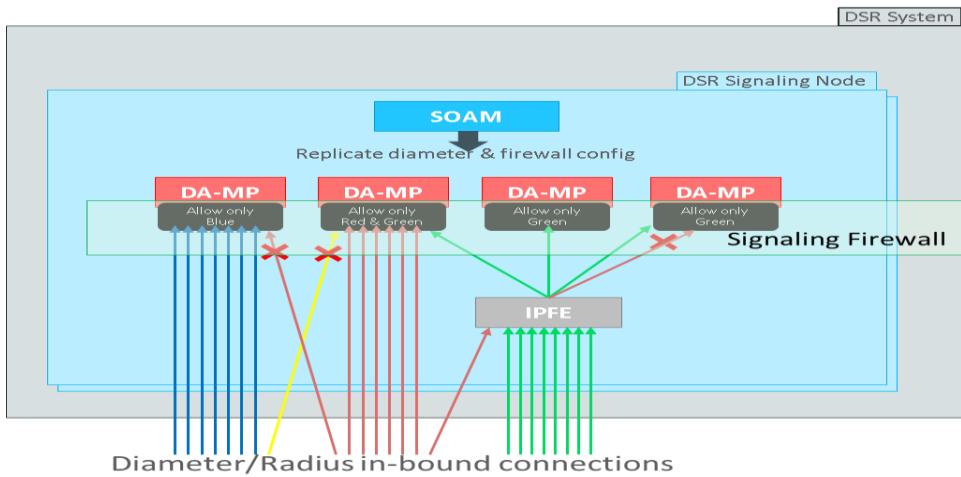
Figure 19 - DSR Signaling Firewall

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

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.
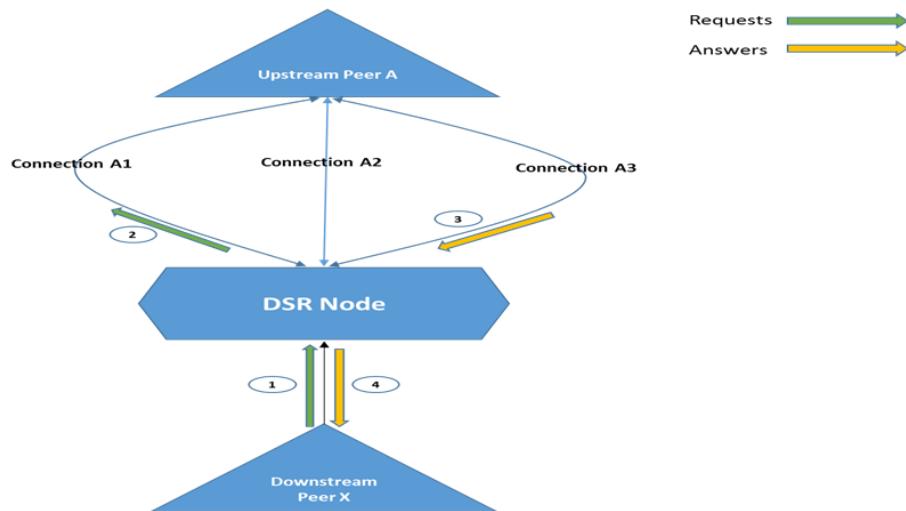
Figure 20 - Answer processing on connection different from the connection used to send Request message

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.   Please see   for more information
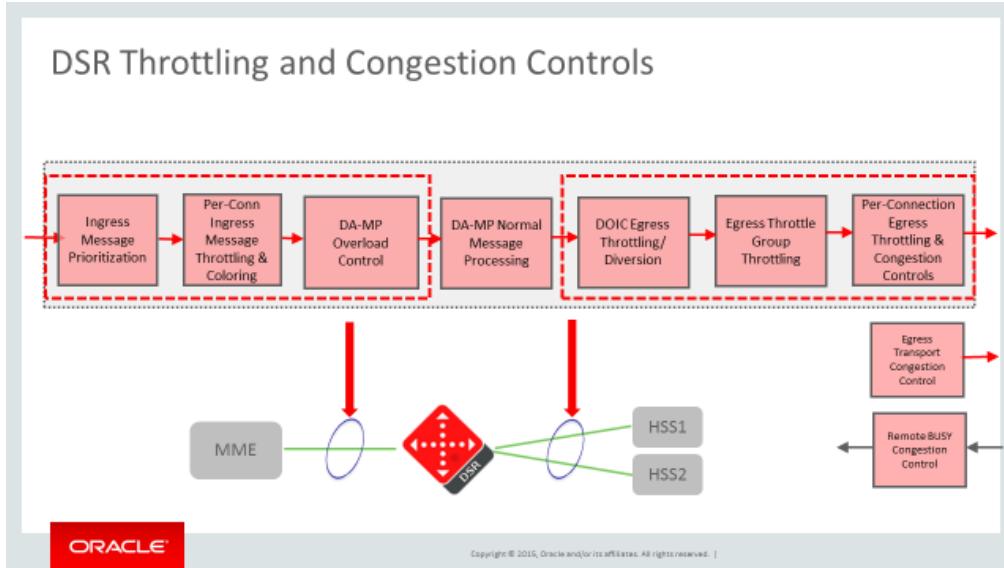
Figure 21 - Congestion Control

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 (i.e. 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 (ex: 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 (i.e. Request/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 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
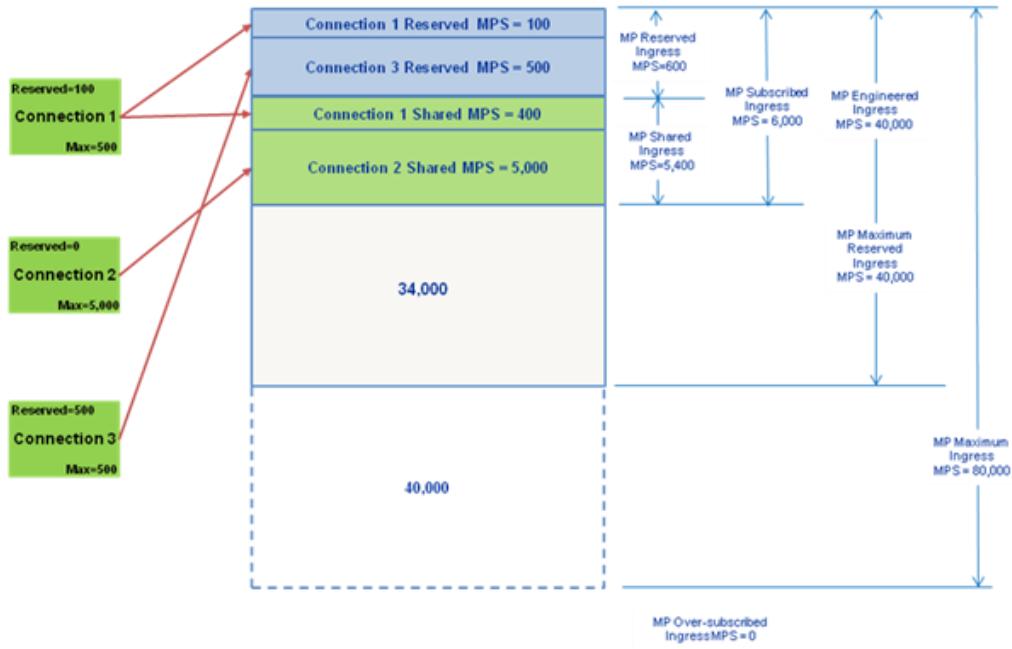
Figure 22 – DSR Ingress MPS Configuration Example 1 – Normal Case

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.
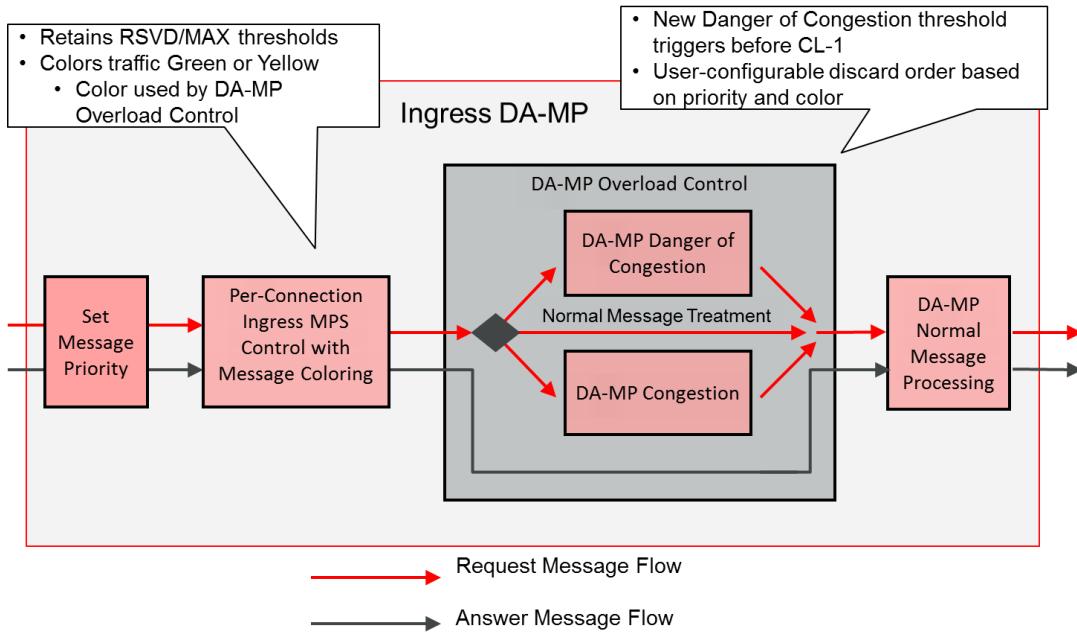
Figure 23 – Message Coloring and Priority/Color-based DA-MP Overload Control

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 (i.e. message & event handling), allowing a sufficient CPU budget for other non-critical (i.e. 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.

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

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.
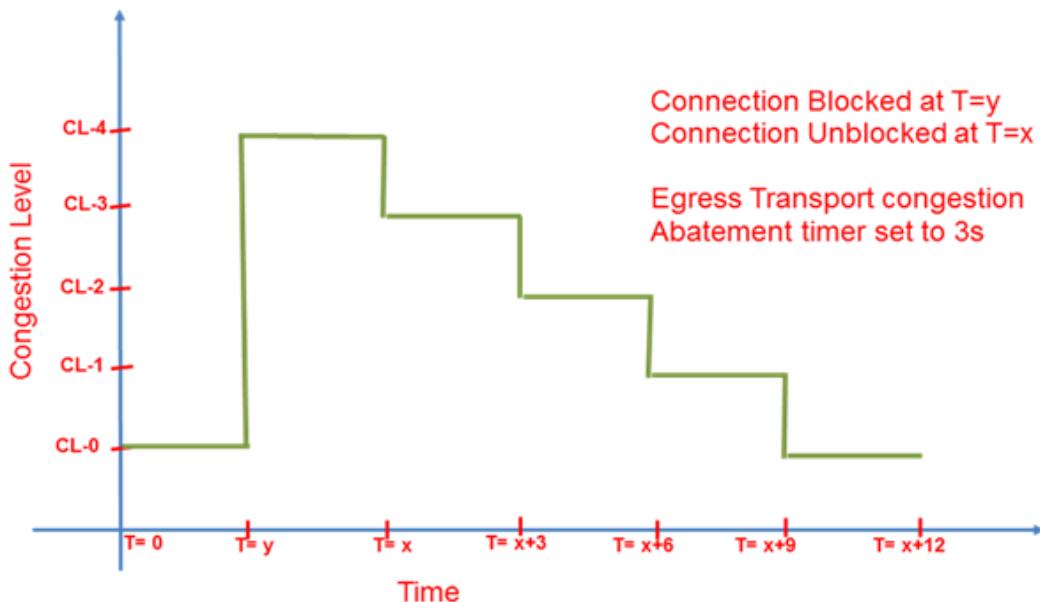
Figure 24 – Example Congestion level abatement

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.

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

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)
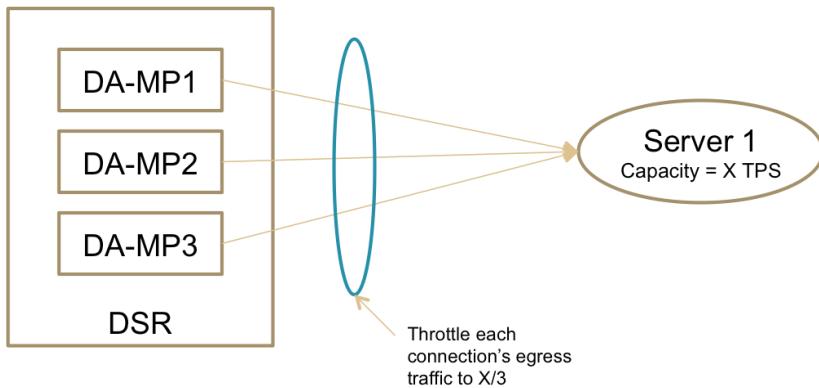
Figure 25 – DSR Per-Connection Egress Throttling

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 be 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)
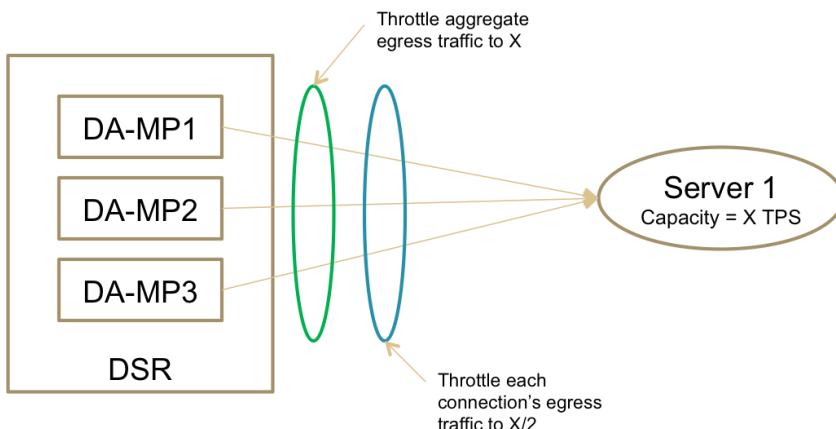
Figure 26 – DSR Aggregate and Per-Connection Egress Throttling

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.

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.

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.
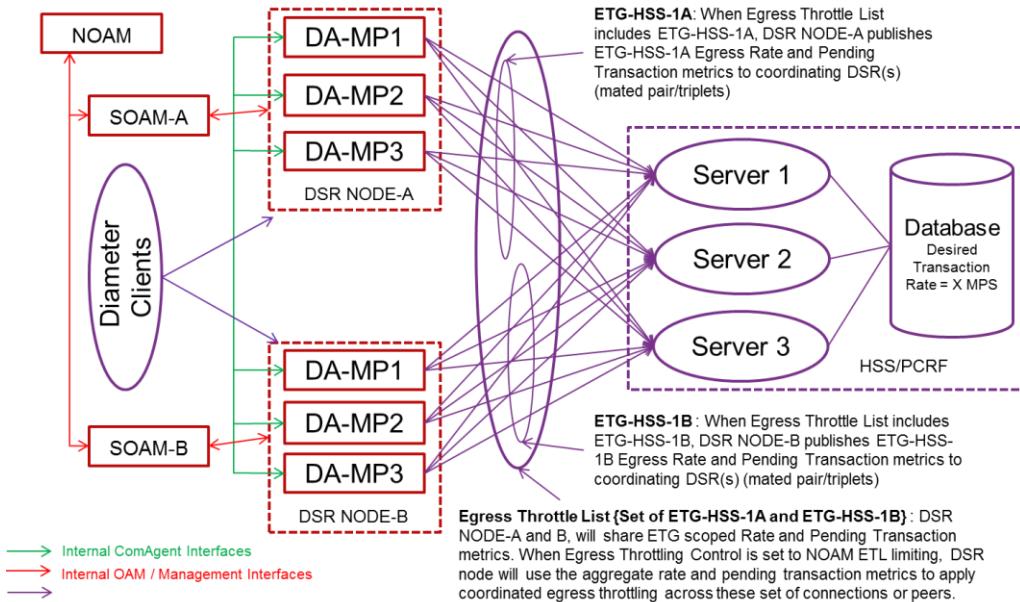
Figure 27 – 2 DSR Sites: Coordinated Egress Throttling Example

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' **and**
» 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.
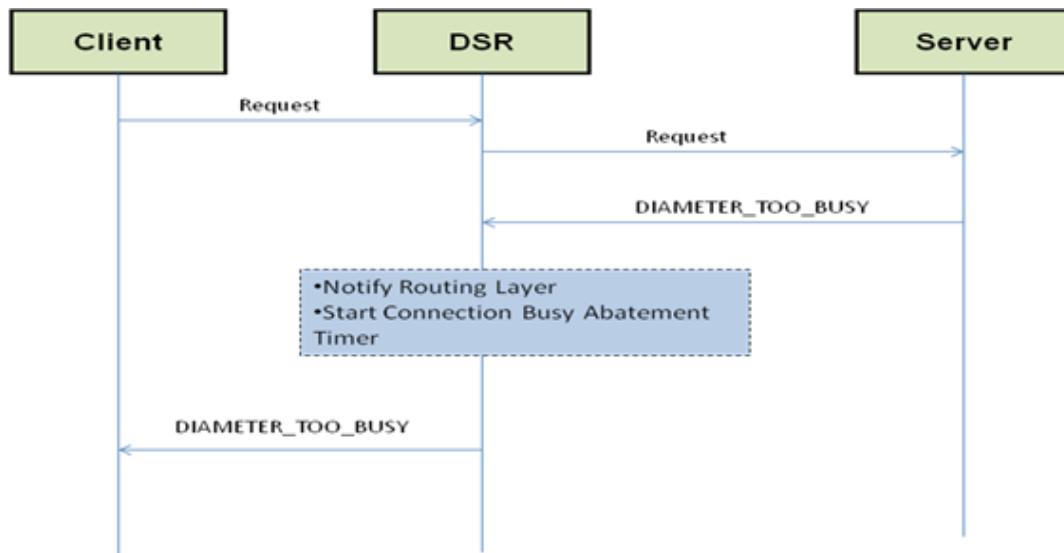
Figure 28 – Connection Busy

TABLE 7: CONGESTION LEVELS BASED ON REMOTE 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
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.
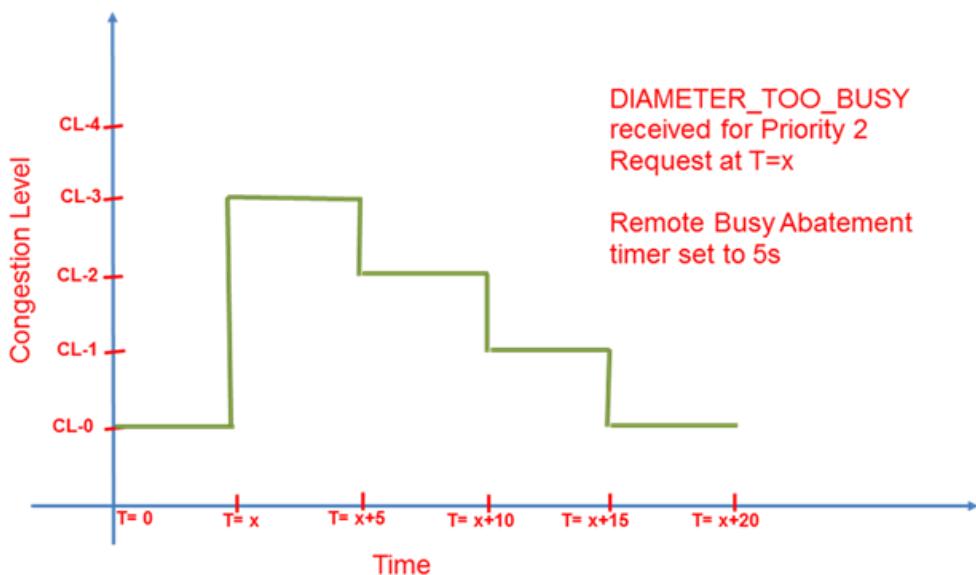
Figure 29 – Congestion level abatement over time for Remote Busy

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 / transmit buffer sizes are user-configurable for the system.

Diameter Overload Indication Conveyance (DOIC)

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:

1. 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.

2. 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.
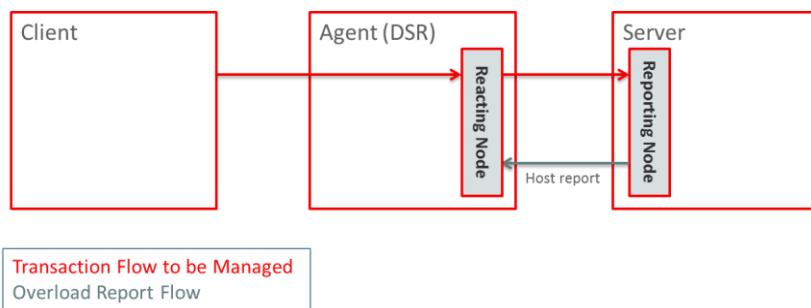
Figure 30 - Reacting Node Role for DOIC

DOIC Capabilities Announcement (DCA)

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.

1. 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).

8. 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.

9. 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.

10. 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.
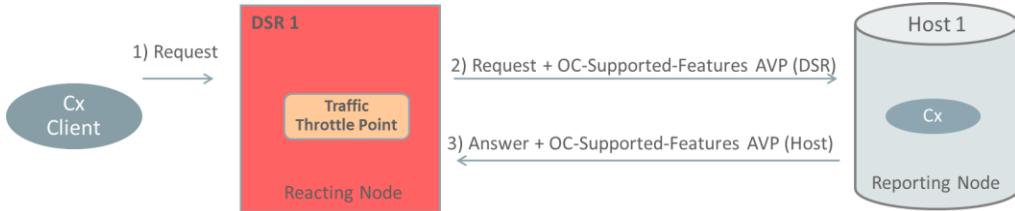
Figure 31 - DOIC Capabilities Announcement

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.
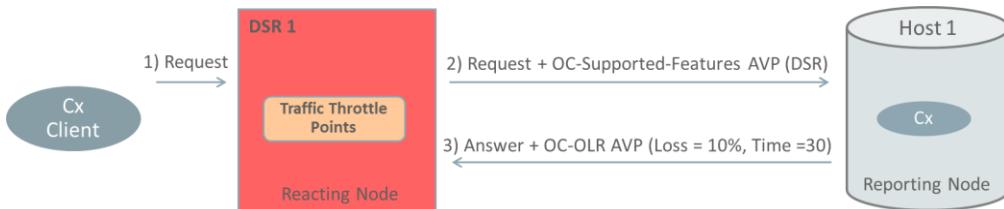
Figure 32 – DOIC Overload Reports Example

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.

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.
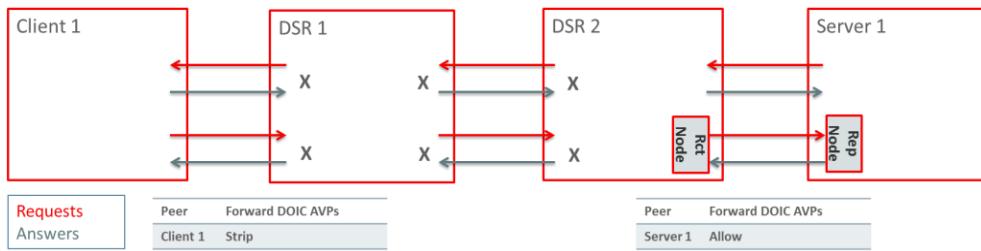
Figure 33 - DOIC Security Setting Example

The following three scenarios are supported:

1. All DOIC AVPs are stripped on all requests and answers sent and received on a connection to a given peer.

2. 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.

3. 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.

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.

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.
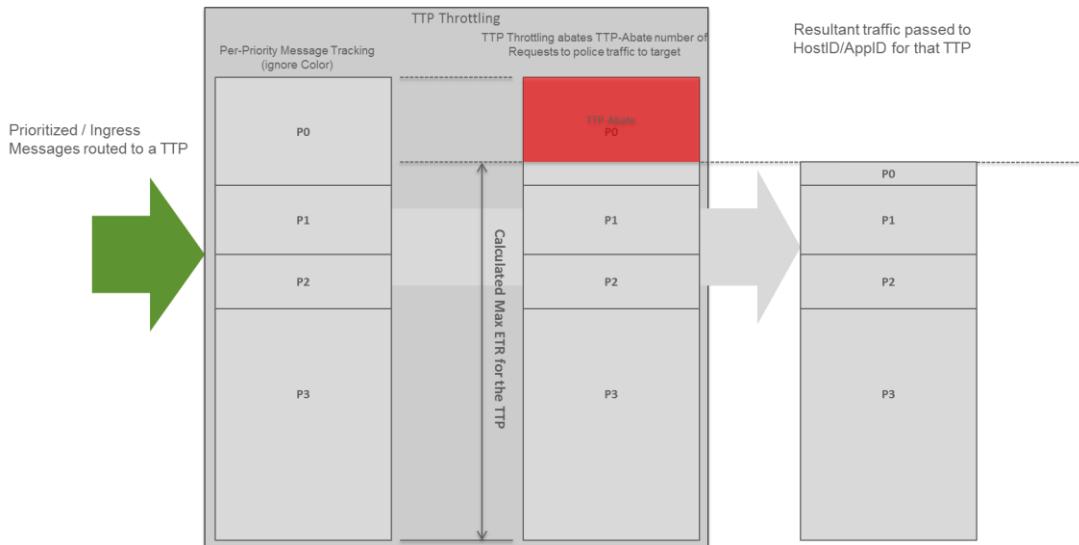
Figure 34 - Throttle by Message Priority Only when support for 5 message priorities is enabled

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.
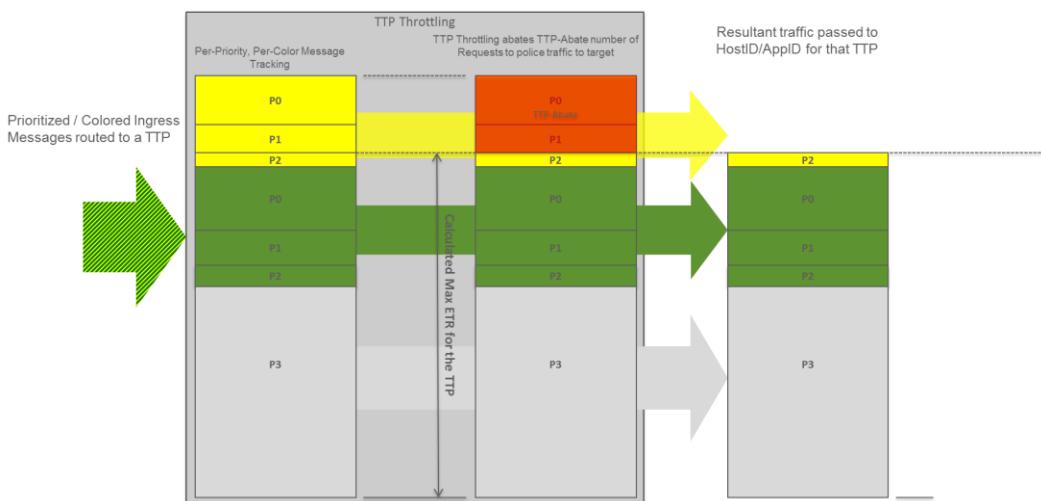
Figure 35 - Throttle by Color then Message Priority when support for 5 message priorities is enabled

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.

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.

Traffic Throttle Point (TTP)

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, etc.)

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 8: 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	
	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 (i.e. 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 (i.e. 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 (i.e. ignores DOIC loss data).
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 % (i.e. is the TTP degraded), and the operational status of the underlying peer.
	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 > 0 » Recovering means that the DSR is ramping up the traffic after an overload state has ended.	Current DOIC Status
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/Priorities needs to be displayed.	Total OTR
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 9: 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.
	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.
	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.

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, don't 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 that the DOIC capabilities layer on top of the existing routing functions. Unless specifically noted, all of the current DSR routing functions continue to operate.

1) Traffic Throttling Group (TTG)

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.

TTP FE1	
Max EMR	2000
Current EMR	900
Loss %	-20%
.....	

TTP FE2	
Max EMR	1000
Current EMR	500
Loss %	-10%
....	

TTP FE3	
Max EMR	2000
Current EMR	1000
Loss %	-30%
...	

TTG 1	
Loss %	22%

Loss % = (sum across TTPs of (Loss %*Max EMR))/(sum across TTPs of Max EMR)
-22% = ((2000*-20%)+(1000*-10%)+(2000*-30%))/(2000+1000+2000)
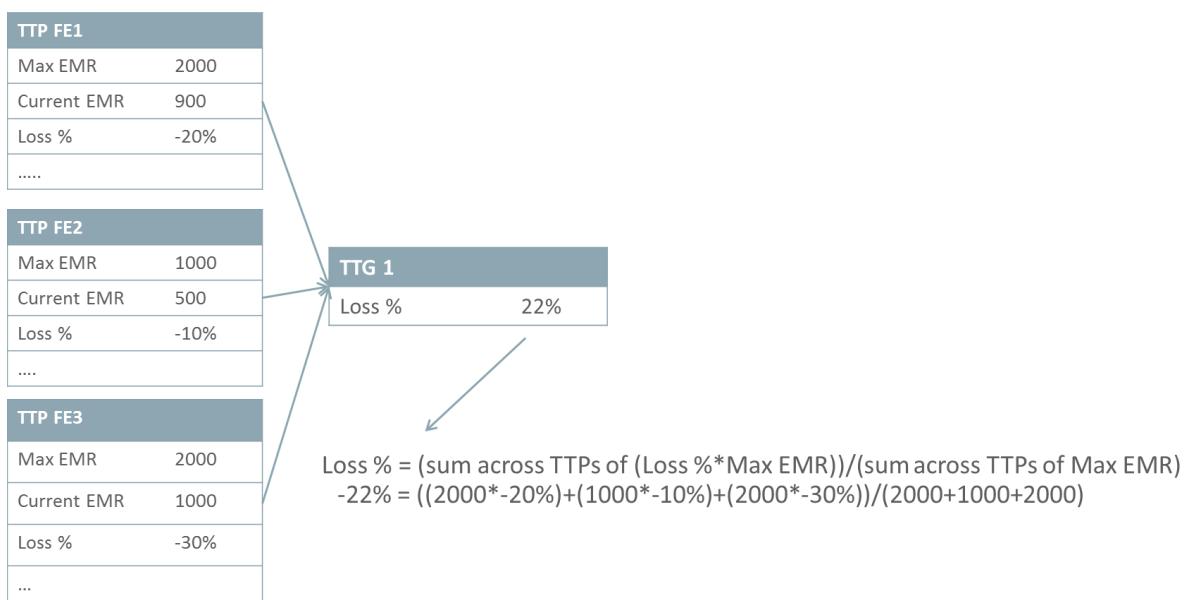
Figure 36 - TTP to TTG Rollup Logic

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.

TTG CONFIGURATION AND STATUS DATA

The table below lists the data required in the TTG.

Table 10: 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.
	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) Congestion-Aware Route Lists

The DSR can use the congestion information in the TTGs to skip Route Groups in the Route List that don't 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.

Route List 1 Definition

RG	RG Priority	Max % Loss	TTG
RG1	1	100%	DSR1_S6b
RG2	2	50%	DSR1_2_S6b
RG3	3	30%	DSR2_S6b

Current TTG Status

TTG	Current Loss Rate
DSR1_S6b	20%
DSR1_2_S6b	0%
DSR2_S6b	0%
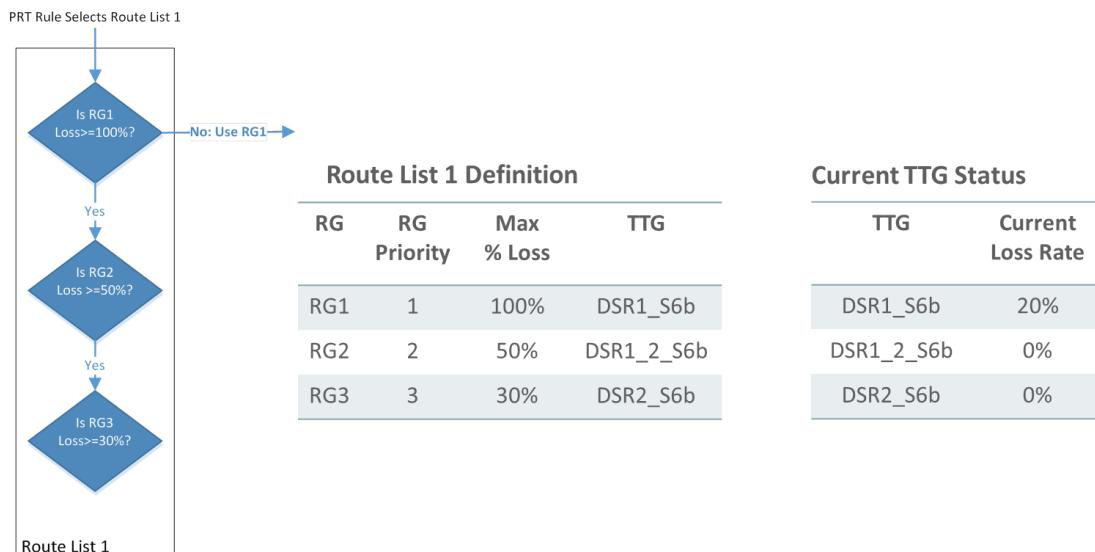
Figure 37 - Congestion-Aware Route List Logic

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.
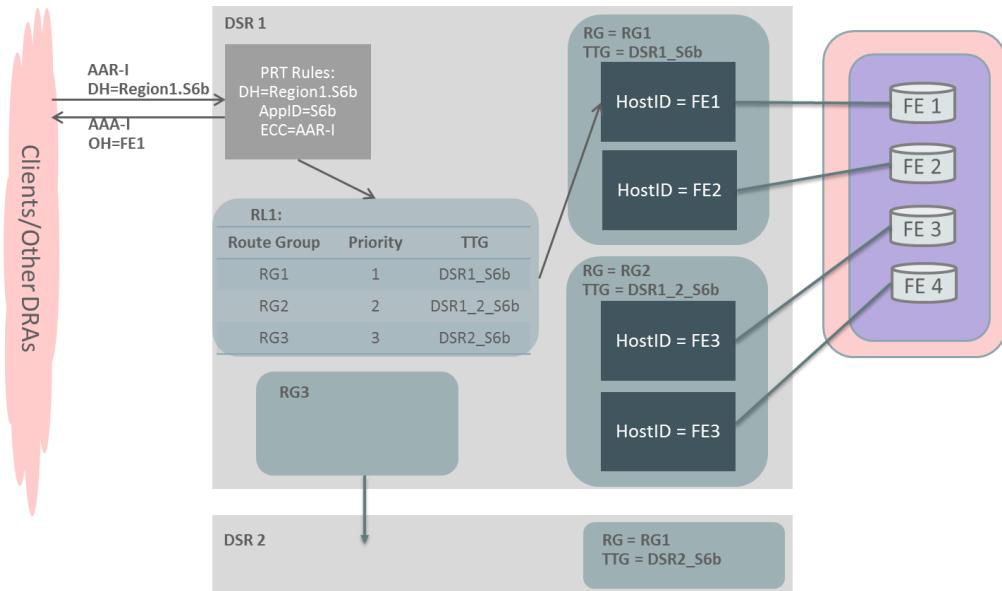
Figure 38 - Congestion-Aware Route List Example 1

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.
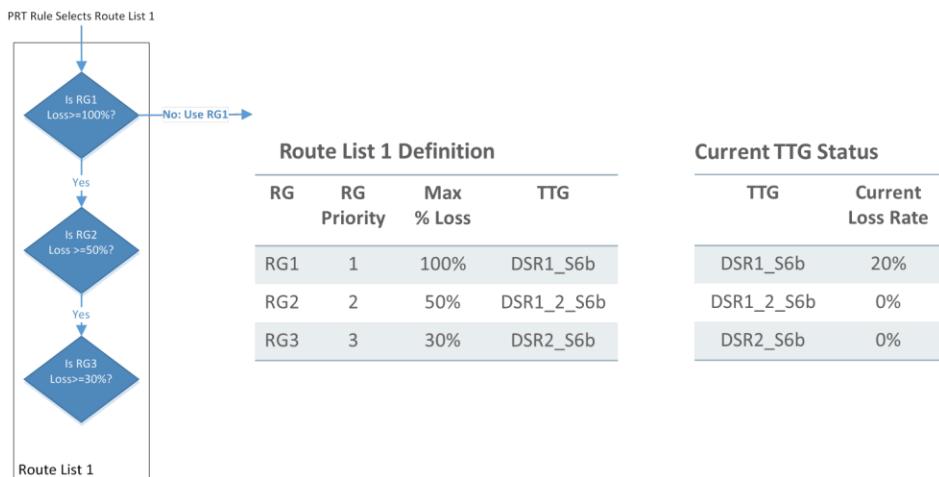
Figure 39 - Congestion-Aware Route List Logic Example 1a

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.

Route List 1 Definition

RG	RG Priority	Max % Loss	TTG
RG1	1	80%	DSR1_S6b
RG2	2	50%	DSR1_2_S6b
RG3	3	30%	DSR2_S6b

Current TTG Status

TTG	Current Loss Rate
DSR1_S6b	90%
DSR1_2_S6b	60%
DSR2_S6b	60%

Figure 40 - Congestion-Aware Route List Logic Example 1b

3) 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.

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.

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:

1. Yellow/priority 0
11. Yellow/priority 1
12. Green/priority 0
13. Green/priority 1
14. Yellow/priority 2/DOIC Override
15. Green/priority 2/DOIC Override

Next Generation Network Priority Service (NGN-PS)

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.
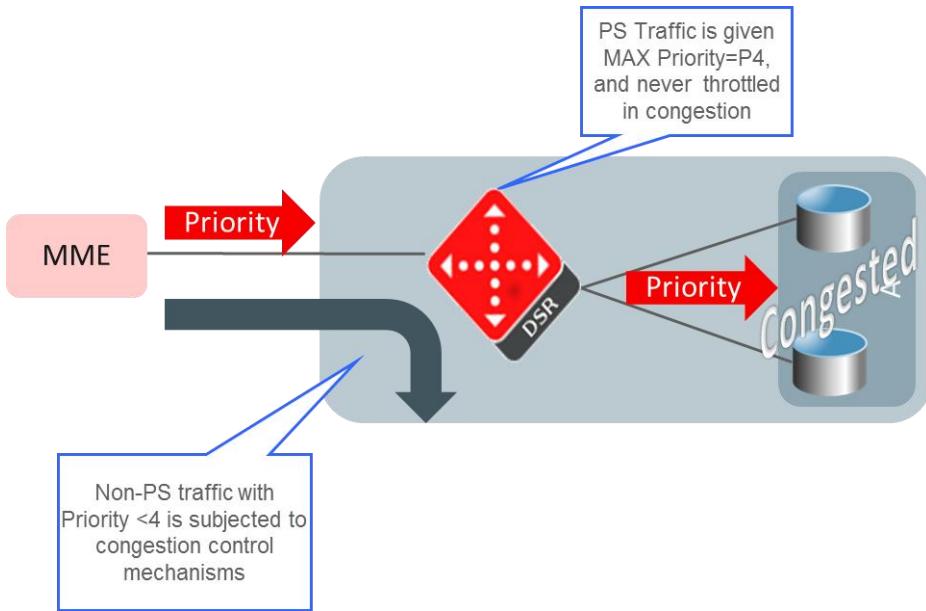
Figure 41 - NGN-PS Support

IP Front End (IPFE)

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.

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, etc.) 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 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 + 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

Each IPX interconnect uses Initiator + Responder Connections. No need to to negotiate connection initiator / responder role with each interconnect.
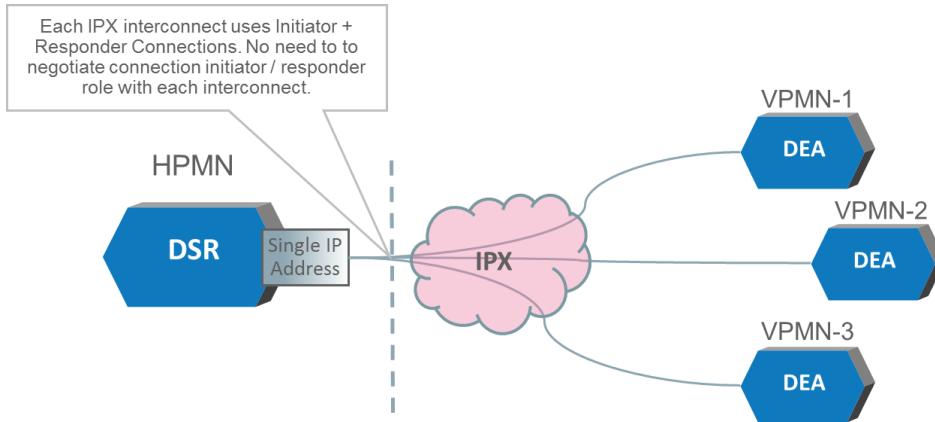
Figure 42 - IPFE Initiator + Responder Support

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

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.

RADIUS Signaling Router (RSR)

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.

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.
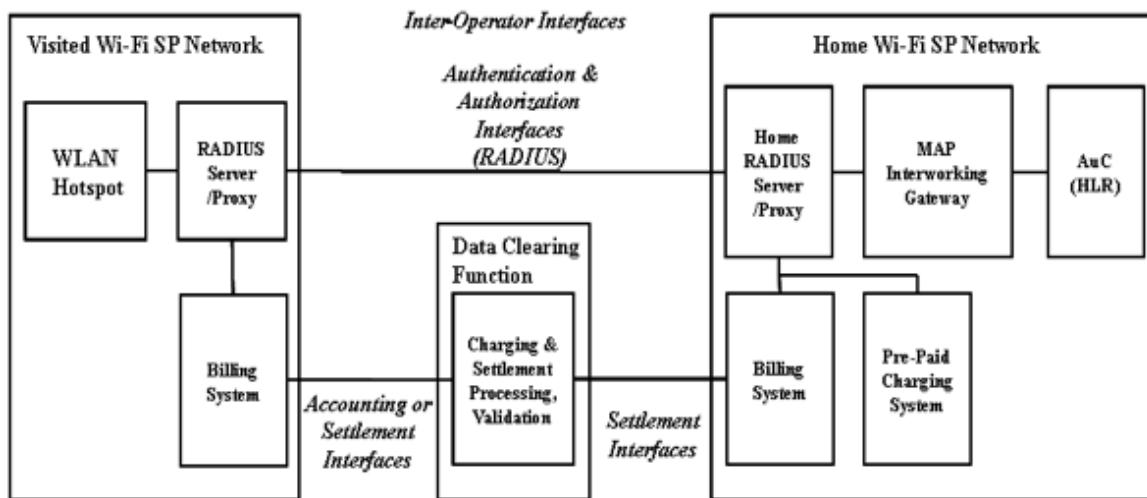
Figure 43 – RADIUS Interfaces in WLAN Roaming Architecture[1]

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/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.

RADIUS Message Format

[1]   GSMA IR.61(v5.0) – WLAN Roaming Guidelines (Inter-Operator Handbook)

The figure below shows the basic RADIUS message data format.
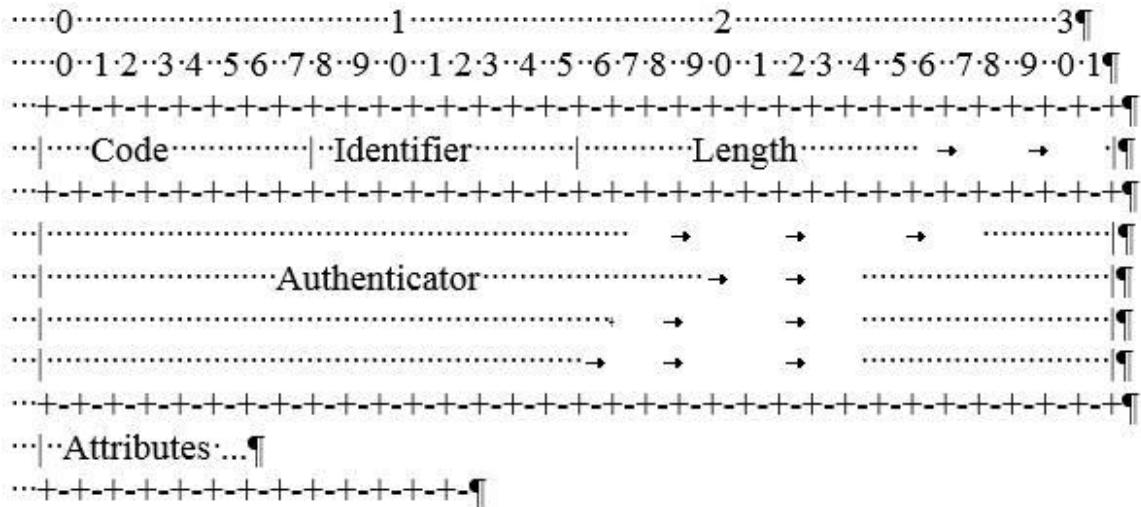
Figure 44 – RADIUS Data Format

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.

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.

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.

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 11: 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, etc. 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.

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.

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)

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.
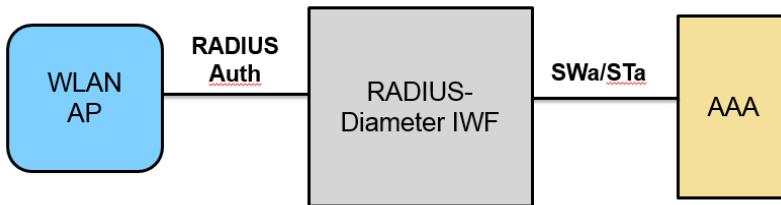
Figure 45 - RADIUS-Diameter IWF for WLAN Authentication

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 red lines 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
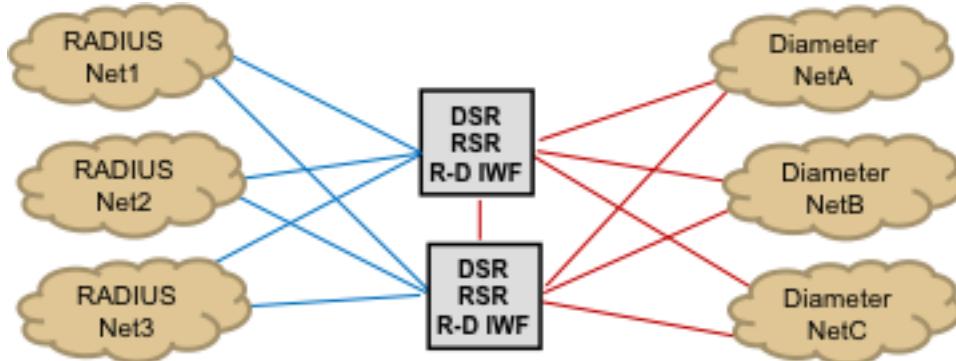
Figure 46 - RSR and R-D IWF Deployment

As shown, 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.

DSR

BNG — RCL — RD-IWF / DRL — AAA

1) Access-Request
- User-name
- EAP-message
- Calling-Station-ID
- NAS-Port-Type

2) DER
- Appl-ID=STa
- Origin-Host=BNG
- Origin-Realm=BNG_Realm
- Dest-Realm=derived from User-Name
- Session-ID=0
- Proprietary-AVP=RADIUS content

3) DER
- Session-ID=YYY (RD-IWF generated)
- Auth-Request-Type=Authorize_Authenticate
- User-name
- EAP-payload
- Calling-Station-ID
- RAT-Type=derived from NAS-Port-Type
- ANID=derived from NAS-Port-Type

4) DEA
- Origin-Host=AAA
- Origin-Realm=AAA_Realm
- Session-ID=YYY
- User-name
- EAP-payload
- Result-Code=1xxx
- Other

6) Access-Challenge
- User-name
- EAP-message
- State=YYY+Origin-Host from DEA

5) DEA
- Message #4
- Proprietary-AVP=RADIUS content

7) Access-Request
- State=same as msg #6
- Other attributes same as msg #1

8) DER
- Proprietary-AVP=RADIUS content

9) DER
- Destination-Host=Origin-Host from State
- Session-ID=YYY
- EAP-payload
- Other AVPs same as msg #3

10) DEA
- EAP-payload
- EAP-Master-Session-Key
- Result-Code=2xxx
- Other

12) Access-Accept
- User-Name
- EAP-message
- MS-MPPE-Recv-Key
- MS-MPPE-Send-Key

11) DEA
- Message #10
- Proprietary-AVP=RADIUS content

Figure 47 - 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

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 S9 interface (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 inter operator 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.

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 flags

» checking for any component of the Diameter message (flags, appl-id, cmd-code, etc)

» 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.

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 deployed without adversely impacting the performance of the DSR.

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/proxy/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
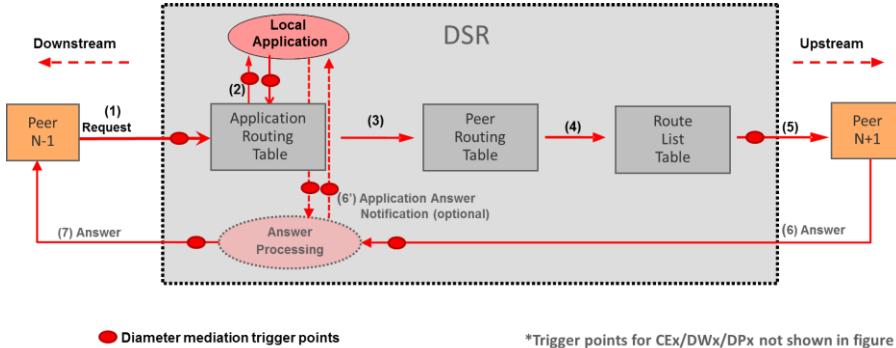
Figure 48 - Mediation Trigger Points

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

Measurements Associated with Rules

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 (i.e. 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.

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.

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, PCRFs,

AFs, and pCSCFs are deployed. Nor can the peer service providers derive any deployment architecture information through inspection of host names.

S6a/S6d Topology Hiding

S6a/S6d MME/SGSN 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.
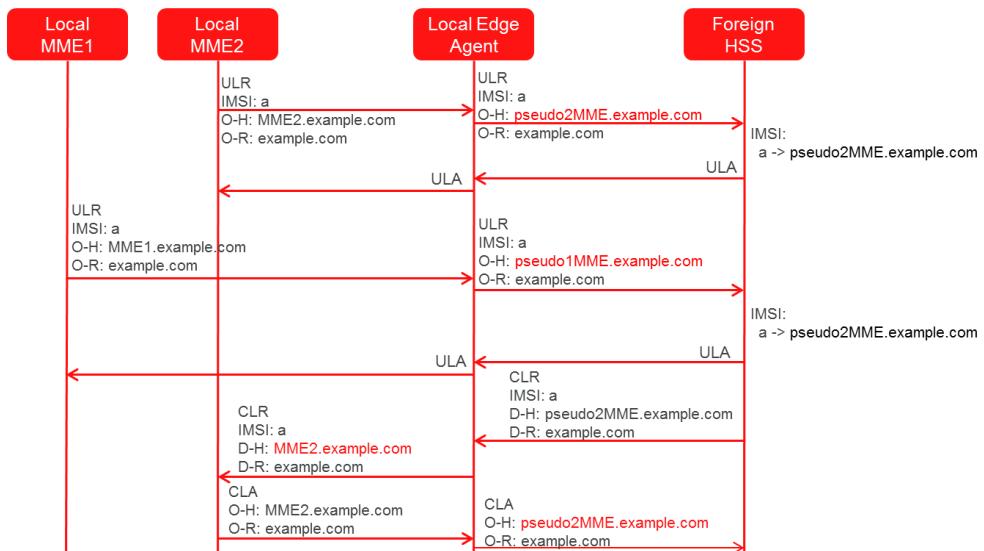
Figure 49 - MME/SGSN Topology Hiding

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. Restoral 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, or,

   » 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 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.

» 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 50 - S6a/S6d HSS Topology Hiding - ULR Message Flow

mme.foreign.com   lea.example.com   hss.example.com

```
   ┌──────┐                        ┌──────────┐                        ┌──────┐
   │ MME  │                        │  Local   │                        │ HSS  │
   │      │                        │  Edge    │                        │      │
   │      │                        │  Agent   │                        │      │
   └──┬───┘                        └────┬─────┘                        └──┬───┘
```

                                                                        CLR
                                               Destination -Host: mme.foriegn.com
                                                      Origin -Realm: example.com
                                                      Origin-Host: hss.example.com
                                             Session -Id: hss.example.com;3307

                                  CLR
                       Origin-Realm: example.com
                   Origin-Host: pseudohost.example.com
                 Session -Id: pseudohost.example.com ;3307

                    Where pseudohost is a configured value

   CLA
   Session-Id: pseudohost.example.com;3307
                                                   CLA
                                                   Session -Id: hss.example.com ;3307

Figure 51 - 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.com

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 12: MME/SGSN PSEUDO-HOST NAME MAPPING EXAMPLE

MME/SGSN Real Host Name	MME/SGSN Pseudo-Host Name(s)
mme1.westregion.example.com	mme042.example.com
	mme123.example.com
mme2.westregion.example.com	mme533.example.com
mme1.eastregion.example.com	mme922.example.com
mme2.eastregion.example.com	mme411.example.com
	mme218.example.com
	mme331.example.com

mme1.texasregion.example.com	mme776.example.com
	mme295.example.com
	mme333.example.com

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.

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:

3. 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.

4. 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/Proxy Agent 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 52 - 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.

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 S9 Reference Point. This S9 PCRF Topology Hiding feature encompasses:

» PCRF Topology Hiding in inbound and outbound roaming use cases – Hiding of PCRF host names in S9 messages 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 S9 Reference 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 S9 Reference 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 S9 Reference 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.

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:

» S9 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.

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)

» 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 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-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

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.

Full Address Based Resolution (FABR)

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 Relay Agent 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.
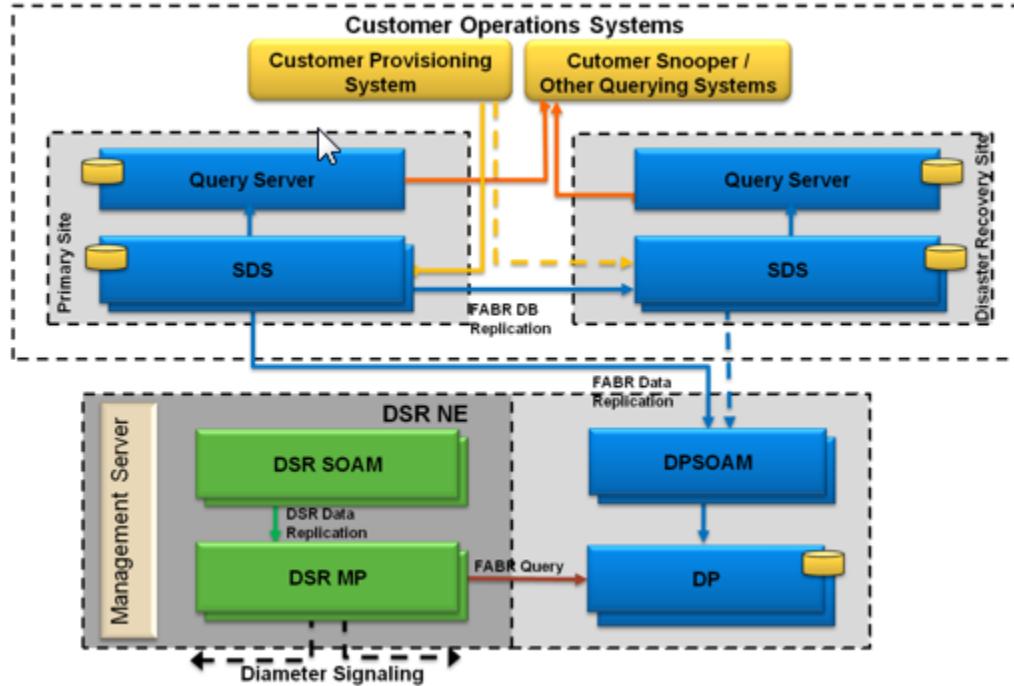
Figure 53 - Subscriber Data Server Architecture

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.

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 without the + sign). The blacklisted IMSIs and MSISDNs are provisioned via the SDS 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.

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.

Policy and Charging Application (PCA)

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

1. Online Charging Proxy (also known as Online Charging Diameter Routing Agent (OC-DRA))
2. 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.
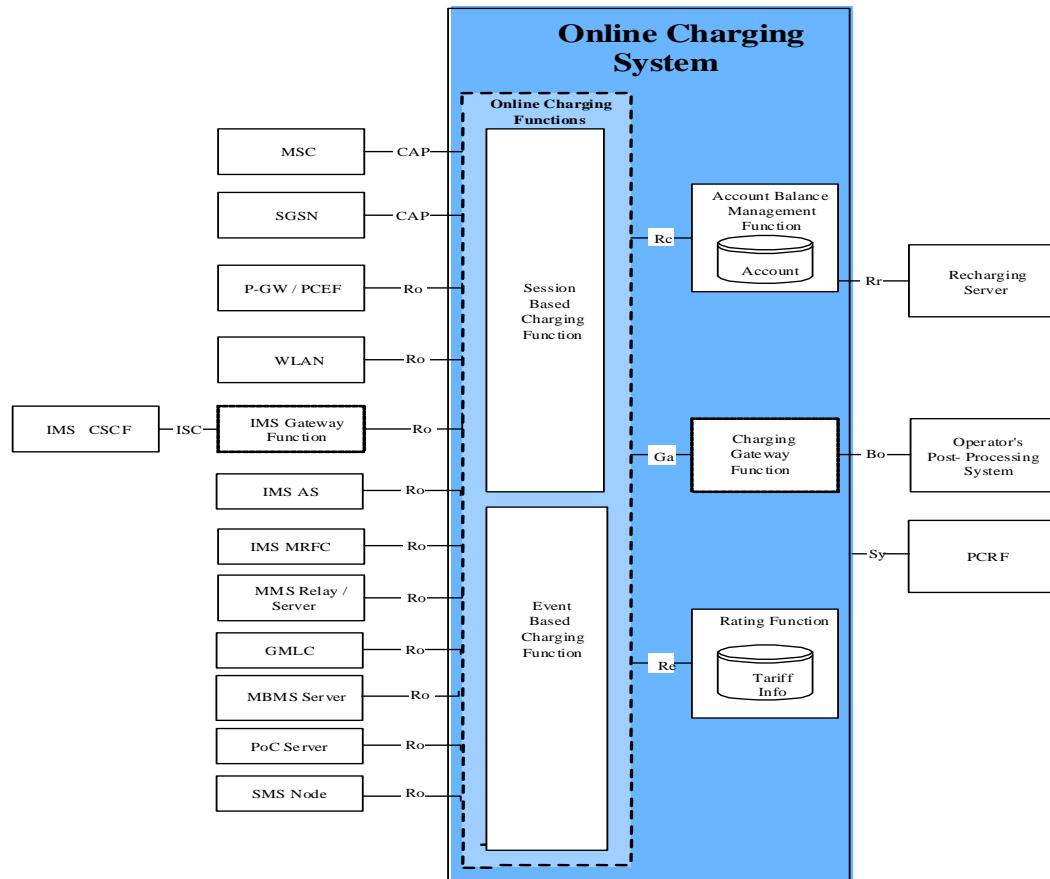
Figure 54 - Online Charging System and Architecture

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 so,

» 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.
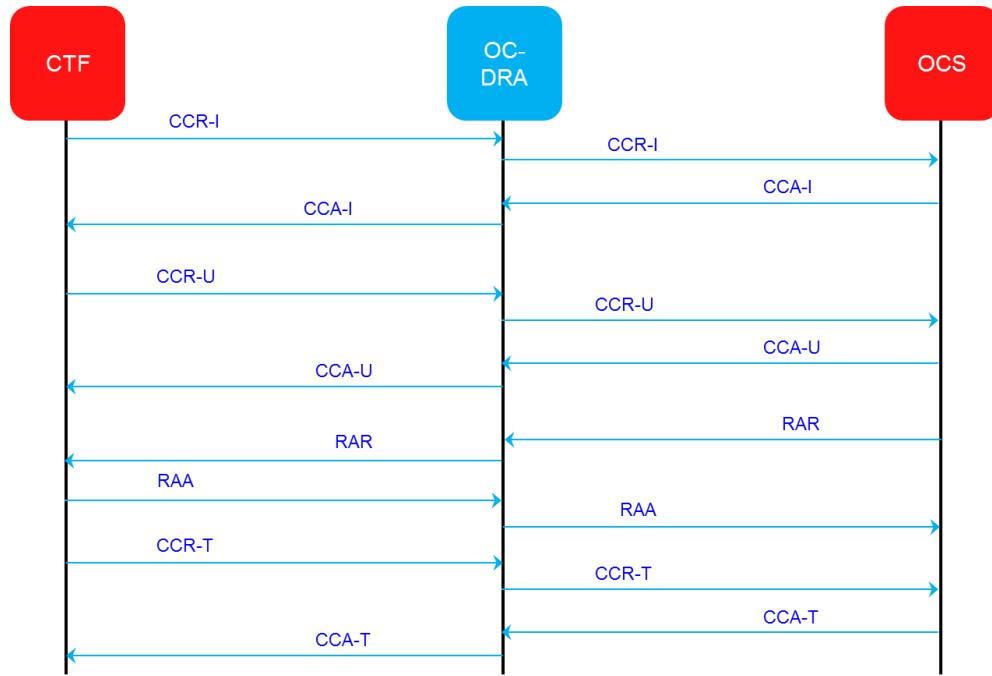
Figure 55 - A typical Online Charging Session

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 (i.e. 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.
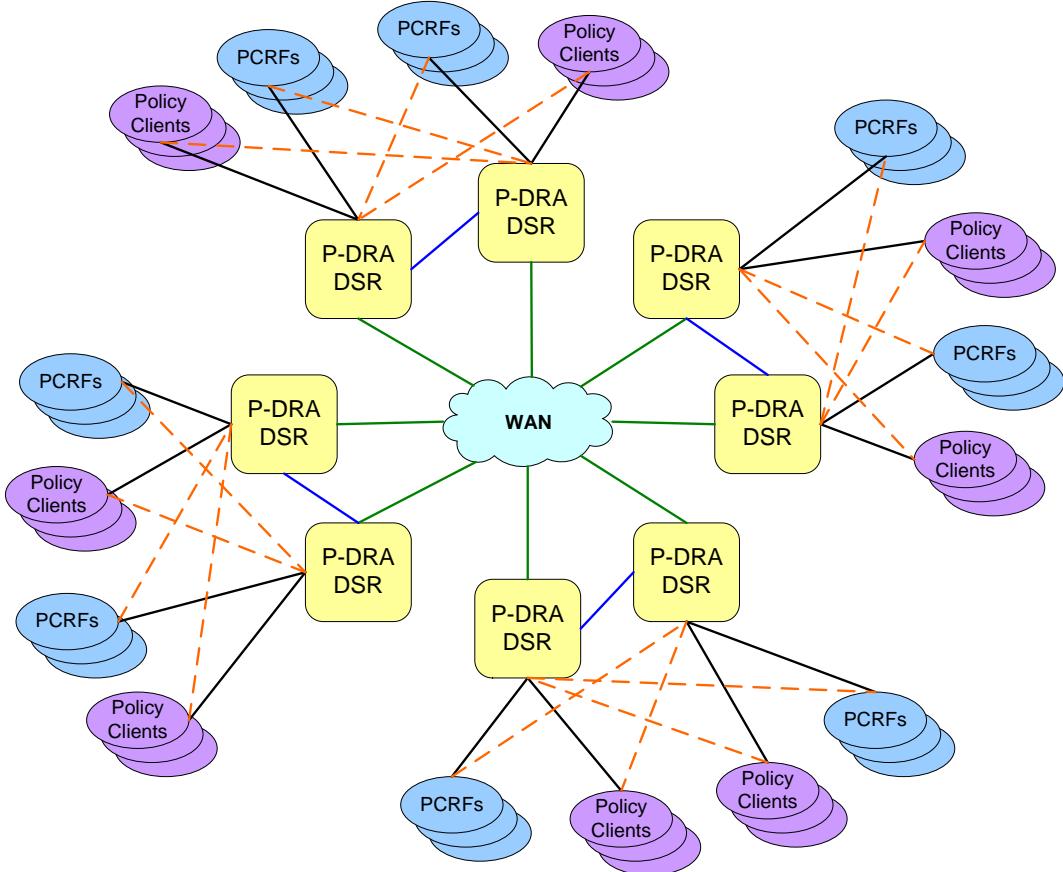
Figure 56 – Network View of P-DRA Mated Pairs

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.
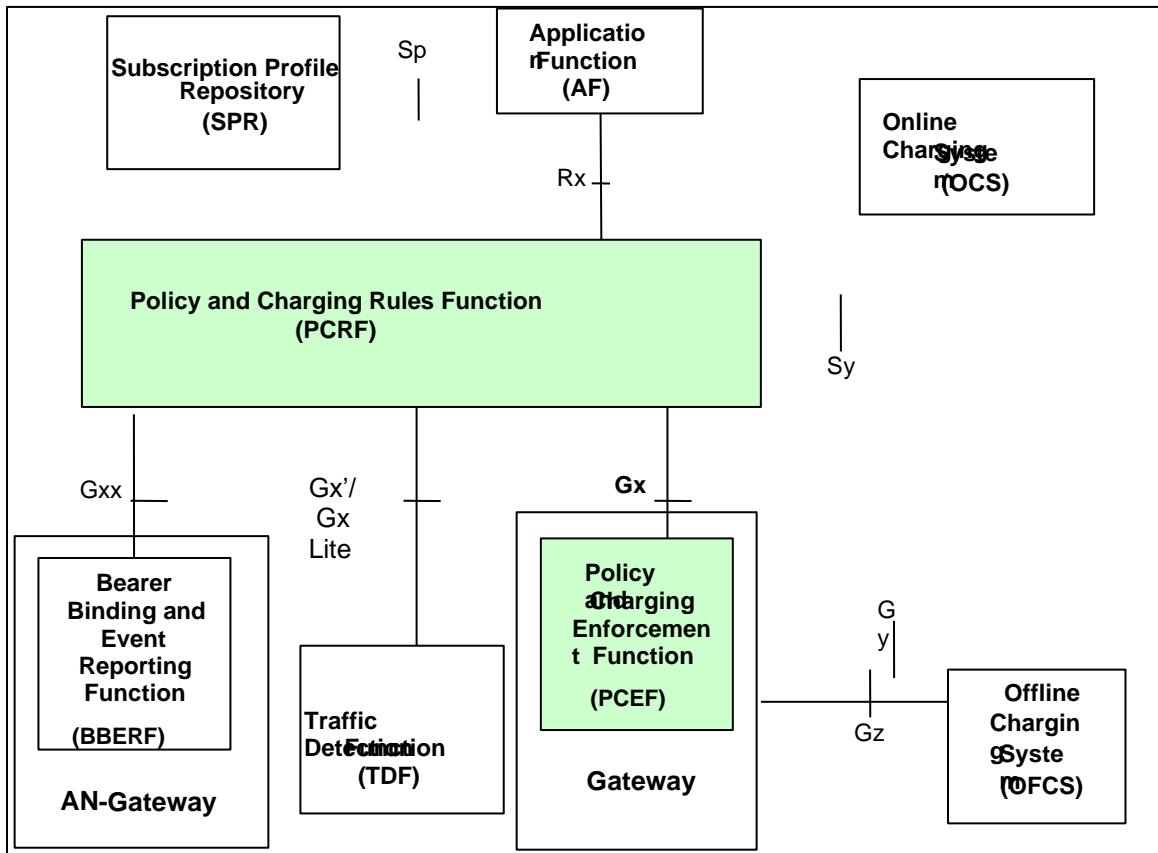
Figure 57 - Overall PCC logical architecture (non-roaming)

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-DRAs 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.

4) 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 the requests, 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 database and a lookup is required for subsequent requests in the session.

5) 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.
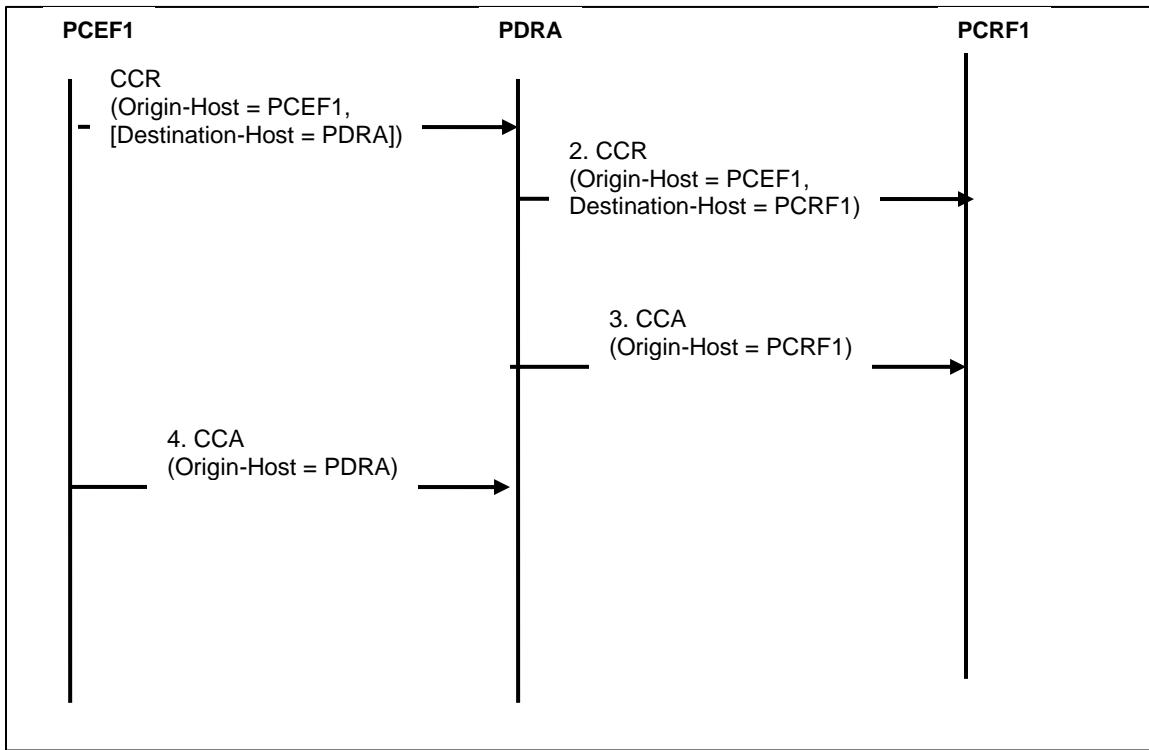
Figure 58 - PCRF topology hiding

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.

6) 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 IPcan 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 PDRAs in different 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.

As shown below, there is a many to one relationship between APNs and PCRF pools. New sessions for the same IMSI can come from multiple APNs and map to the same PCRF Pool.
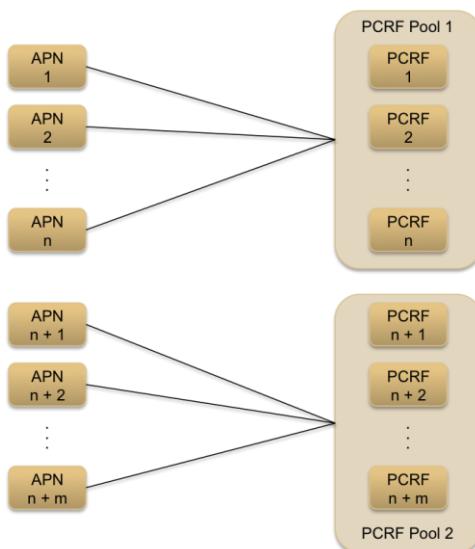

Figure 59 - Relationship between APNs and PCRF Pools

The figure below illustrates the relationship between IMSI and PCRF pool. The same IMSI is able to have active bindings to multiple PCRF pools.
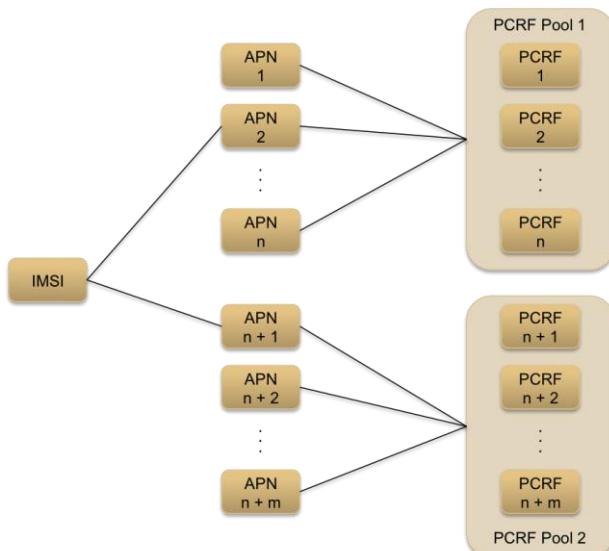
Figure 60 - Relationship between IMSIs and PCRF pools

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 function always 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/OCSs are cross-connected such that all three PCA DSRs have connections to all policy clients and all PCRFs/OCSs associated with the mated triplet.

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

Figure 61 - PCA Example Deployment
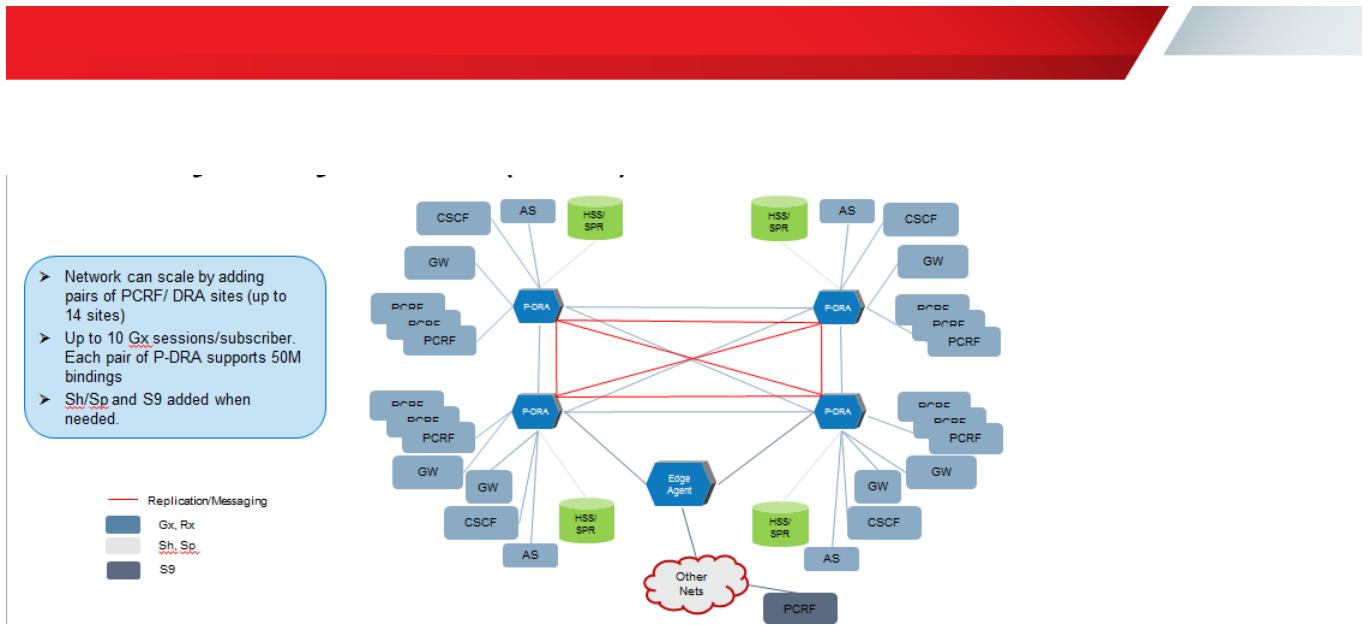
Diameter Security Application (DSA)

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 62 - 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 counter measures. 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.
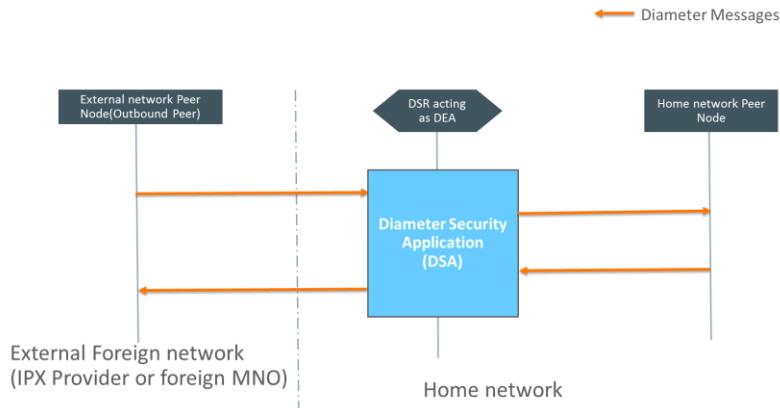
Figure 63 – Diameter Security Application (DSA)

DSA security countermeasures are applied to two types of subscribers:

- Inbound roaming subscribers: Countermeasures are applied for visited network subscribers roaming in home network

- Outbound roaming subscribers: Countermeasures are applied for home network subscribers roaming in visited 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 countermeasures.

DSA security countermeasures can be classified as:

- Stateless Security Countermeasures: Countermeasures which can be applied without maintaining the state information at DSA. Some of the important countermeasures are:

✓ Application-Id whitelist screening	✓ Visited-PLMN-ID and Origin-Realm consistency check
✓ Application-Id and Command Code consistency Check	✓ Realm and IMSI consistency check
✓ Origin Realm and Destination Realm whitelist screening	✓ Destination-Realm and Origin-Realm match Check
✓ Subscriber Identity validation	✓ AVP Multiple Instance Check
✓ Origin host and Origin Realm consistency check	✓ Origin host and Origin Realm consistency check

- Stateful countermeasures: Countermeasures which 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

DSA – Cross Protocol Security

Signaling Security Solution to cater to needs of SS7 and Diameter security using UDR as a common database.Multiple protocol technologies are available in the network today i.e. 2G, 3G and 4G.Each Protocol class is solving security issues independently.There is class of threat that may affect the networks supporting 2G/3G and 4G, protocols together.The hacker could potential mimic an MME, sending messages for an IMSI for a Subscriber who is the 2G/3G Coverage Area.Similarly a Hacker can mimic VLR, for an IMSI of a Subscriber in the 4G Coverage Area.The intention of this solution is to provide a way to provide security across protocol classes.This Solution covers Velocity Check/Time Distance Check.

Service Capability Exposure Function (SCEF)

SCEF provides a means to securely expose the services and capabilities provided by 3GPP network interfaces The SCEF provides access to network capabilities through homogenous network application programming interfaces (e.g. Network APIs) defined over T8 interface. The SCEF abstracts the services from the underlying 3GPP network interfaces and protocols. The MTC-IWF provides a means for initiating the MT-SMS based device trigger from IoT Application Server (AS)/Service Capability Server (SCS) and also provides the capability for IoT devices to communicate with AS/SCS using SMS.

SCEF/MTC-IWF functionality at DSR can be split into two functional components as depicted in figure 67:

SCEF API Gateway provides the northbound interface between SCEF and AS/SCP based on T8 interface as specified TS 29.122. The T8 APIs are a set of RESTful APIs defining the related procedures and resources for the interaction between the SCEF and the SCS/AS. It comprises of Gateway App Server, Gateway Admin Server, and Gateway DB Server. DSR implements the following functionalities at API Gateway:

- Implement the northbound T8 RESTful interface as specified by TS 29.122.

- Provide the authentication and authorization functionality for API requests from AS/SCS:

  - Identification of the API consumer (i.e. AS/SCS)

  - Profile management of AS/SCS and management of service level agreements per SCS/AS

  - Support ACL (access control list) management for individual AS/SCS

- API Firewall Functionality to protect from security attacks through T8 interface:

  - Protection against malformed messages received from SCS/AS

  - Protection against oversized messages received from SCS/AS

  - Whitelisting of IP Addresses of SCS/AS

Core SCEF/MTC-IWF application on the DA-MP provides southbound interface towards core network elements like HSS, MME/SGSN and PCRF. Following MTC call procedures are implemented at Core SCEF/MTC-IWF component with DSR acting as SCEF and MTC-IWF network elements:

  - Device Triggering Function (MTC-IWF functionality)

  - Non IP Data Delivery

  - Monitoring event

  - Enhanced Coverage Restriction Control
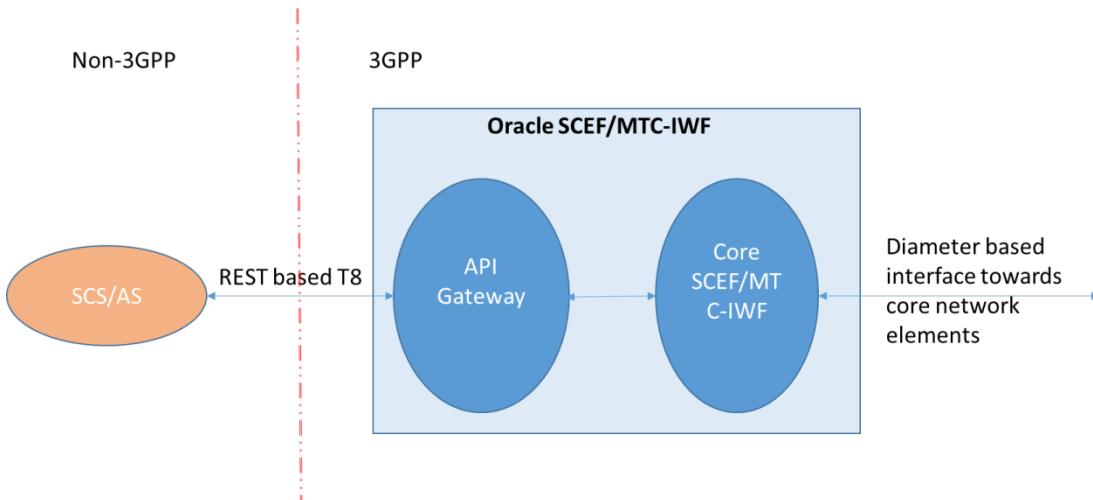
Figure 64: Functional view of SCEF/MTC-IWF functionality at DSR
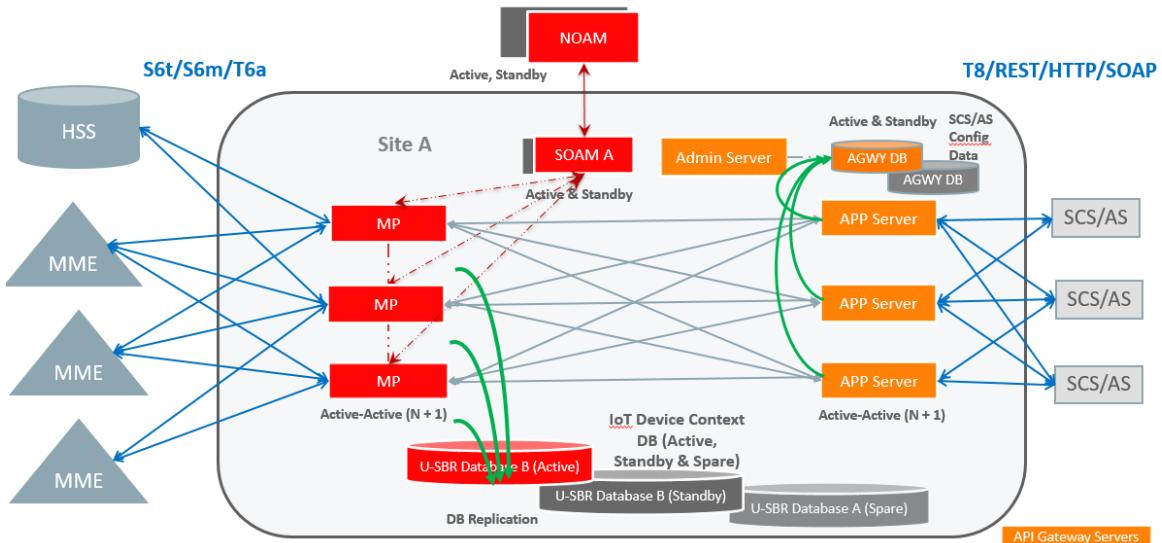
Figure 65: SCEF site architecture example

SBR(u) is used by SCEF to store IoT Device context information related to different MTC procedures. Each of the MTC procedures will add a separate record into U-SBR database.

QoS Control and Sponsored Data connectivity

For IOT connectivity attach with PDN scenario, The 3rd party SCS/AS may request that a data session to a UE that is served by the 3rd party service provider is set up with a specific QoS (e.g. low latency or jitter) and priority handling. This feature provides the functionality where Oracle SCEF can enable PCRF setting up of the QoS via Rx Interface as requested from SCS/AS.

SCEF API based charging for MTC procedures

OCDSR SCEF provides a means to securely expose the services and capabilities provided by 3GPP network interfaces and MTC-IWF provides a means for initiating the MT-SMS based device trigger from IoT Application Server (AS)/Service Capability Server (SCS) and also provides the capability for IoT devices to communicate with AS/SCS using SMS.

With the introduction of API based charging at SCEF, it enables Operators to charge Enterprises (SCS/AS) for their usage of the APIs over the T8 interface.

SCEF as IoT Connectivity Service Gateway for both NB-IoT, LTE-M and other IoT networks

SCEF acts as a LWM2M Server acts as cross proxy converting the CoAP messages to HTTP messages and vice versa to enable IP device communication between IoT devices and SCS/AS.

CoAP Gateway enables LWM2M protocol based IoT device communication for following use cases:
- Access Control
- Device Management
- Connectivity
- Firmware Update
- Location
- Connectivity Statistics

SCEF acts as a MQTT Broker (MB) shall act as cross proxy converting the MQTT messages to HTTP messages and vice versa to enable IP device communication between IoT devices and SCS/AS

High level functional procedures of SCEF CoAP Gateway

- Device Id database storing Device identifiers and Data Flow Type
- CoAP IP Device Registration (MO Data Flow)
- CoAP IP Device MT Data Call Flow
- CoAP IP Device MO Data Call Flow
- CoAP IP Device Monitoring event procedure
- Unified T8 interface for IP and non IP devices
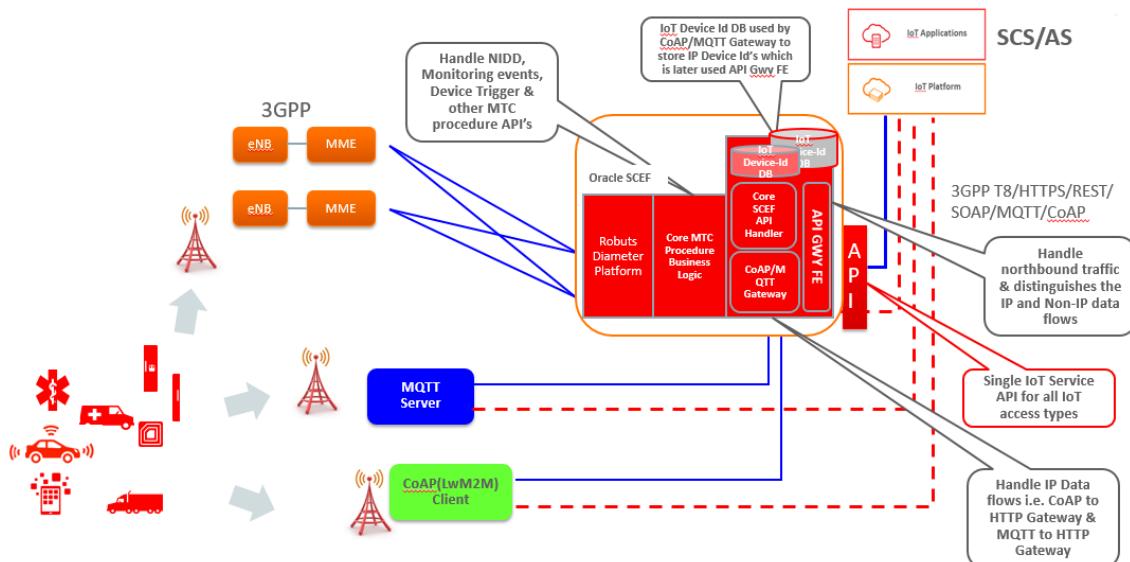
Figure 66: SCEF LWM2M Server, MQTT architecture example

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 RFC6733 compliant 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.
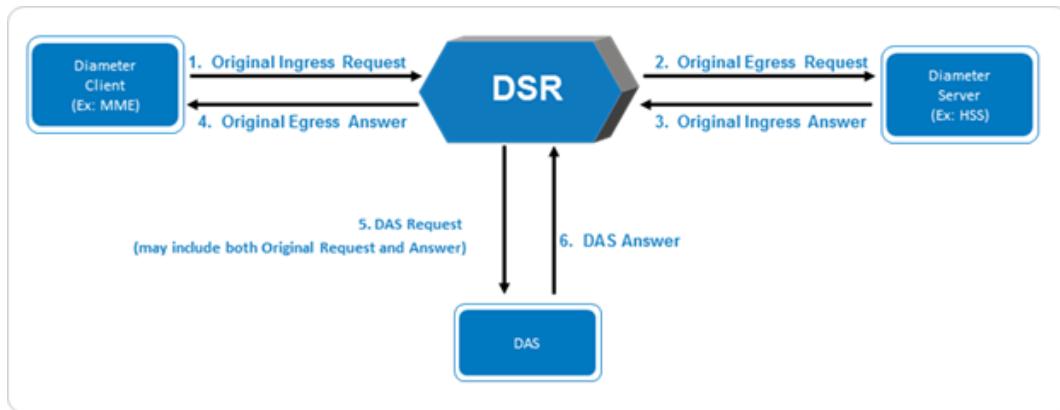
Figure 67 - Message Copy Overview

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

» PRT based triggering
» Using DSR's mediation rules
» DSR application triggering (e.g. FABR)

Virtualized network functions manager VNFM

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 8.3 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 wherein subset of VNF specific function provided in ETSI NFV SOL 003 can be supported. DSR VNFM provides supports ETSI direct model to interface with OpenStack and providing HEAT templates to deploy DSR/vSTP/5G applications.
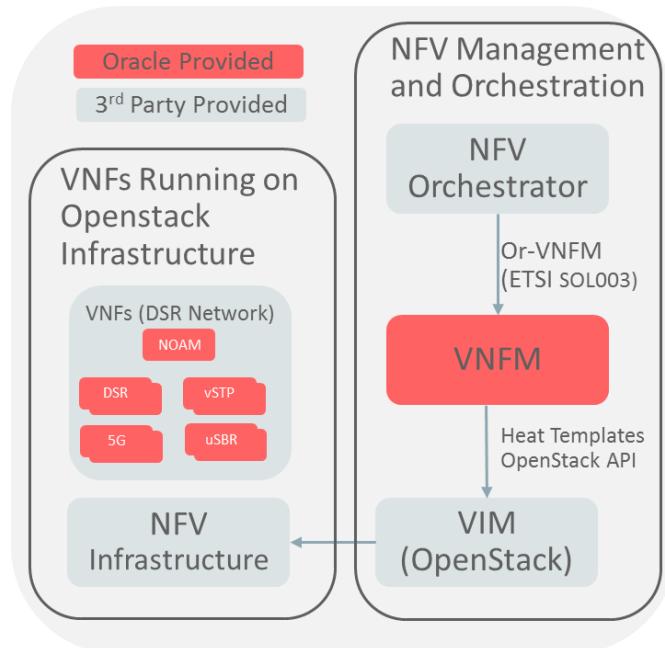
Figure 68 – DSR VNFM

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

1. Create VNF Identifier
2. Instantiate VNF
3. VNFC Scale Out capabilities
4. Query Life Cycle Management
5. Auto Discover capabilities
6. Terminate VNF

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.
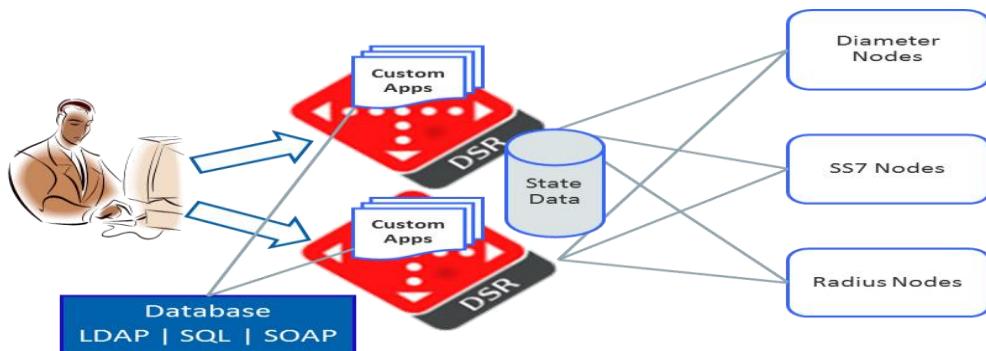
Figure 69 – Custom Application Framework (CAF)

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, etc.

CAF Application Life Cycle Management

CAF application lifecycle starts with activation of CAF application at NOAM and takes effect after the CAF has been administratively enabled. CAF provides the following application lifecycle 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's

  ➢ 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
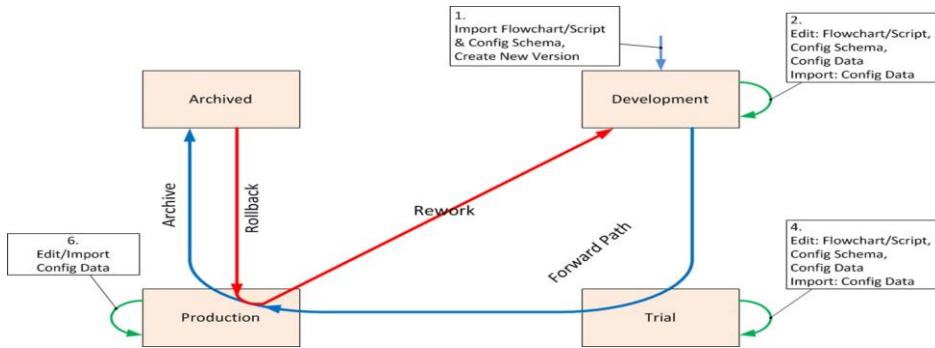
Figure 70 – Custom Application Framework (CAF) Application State Transitions

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.

Enchanced DCA Framework to creat and send Diameter Message.

- DCA Framework has been enhanced to create and send a dimeter 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.

Integrated Diameter Intelligence Hub (IDIH)

Integrated DIH is an integrated troubleshooting capability for the DSR that provides detailed information on how specific messages are processed within the DSR. Integrated DIH allows the user to create trace filters on DSR to capture messages needed for troubleshooting service issues, and presenting those traces to the user via the graphical visualization capabilities provided by IDIH. This feature provides the ability to configure and manage traces from the DSR, as well as filtering, viewing, and storing their results with IDIH.
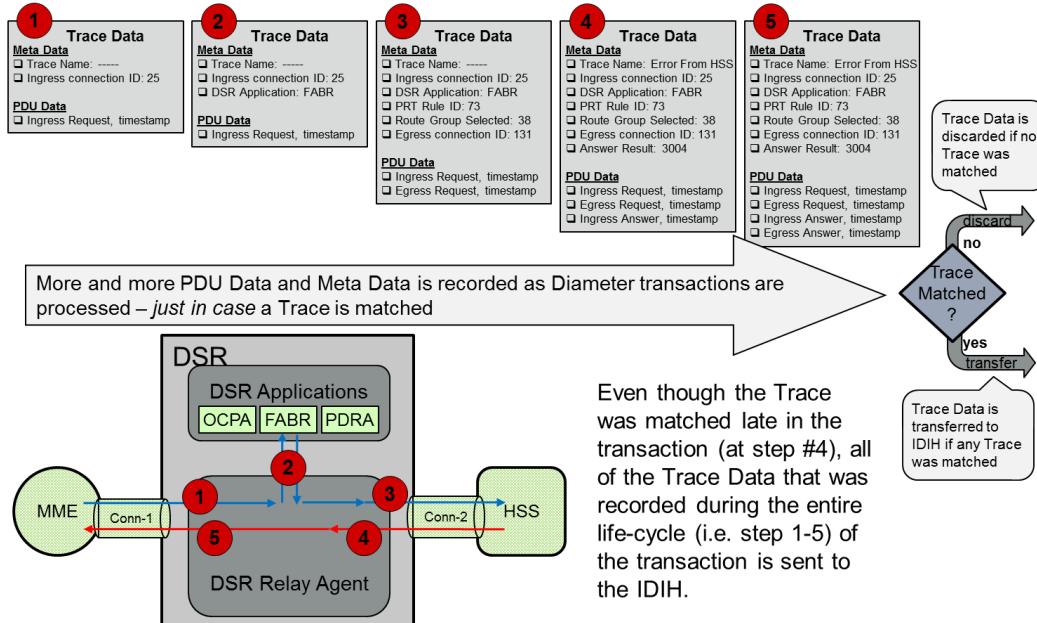
Figure 71 - IDIH Trace Data

The integration of trouble shooting capabilities into the DSR product provides a high value proposition for customers to be able to troubleshoot issues that might be identified with the Diameter traffic that transits the DSR. These troubleshooting capabilities can supplement other network monitoring functions provided by the customer's OSS and network support centers to help to quickly pinpoint the root cause of signaling issues associated with connections, peer signaling nodes, or individual subscribers.

The capabilities provided by this feature are distributed between the DA-MP(s) and an instance of Integrated DIH. The DSR plays the role of determining which messages should be captured, based on trace criteria that are created and activated by the user. The trace criteria identifies the "scope" as well as the "content".   "Scope" refers to the non-protocol-related elements (such as connections or peers) that are used to select messages for trace content evaluation. "Content" refers to the protocol-related elements (such as command codes, AVPs, etc.) that are used to refine the trace criteria. Any trace filter, regardless of scope and content, can be defined as either a "site trace" or a "network trace". A site trace is the default behavior. A network trace results in capturing TTRs that meet the trace filter criteria on any DA-MP within the network. As request and answer messages are processed by the DSR, they are analyzed for matching any of the active trace definitions, and if so, transfer message components along with supplemental information to the IDIH called trace data. A network trace also captures the path that both the Diameter request and answer take as they traverse through multiple DA-MPs within the network. The IDIH can assemble the trace data, and present it to the user leveraging graphical visualization interfaces for additional filtering and analysis. There are three options for then exporting the trace: export the TTR in HTML, export the TTR in PCAP, or export the trace in PCAP.

This feature provides the ability to manage the processing resources associated with capturing trace information as well as the bandwidth for communicating trace data between the DSR and IDIH so that it does not impact the rated signaling capacity of the DSR.
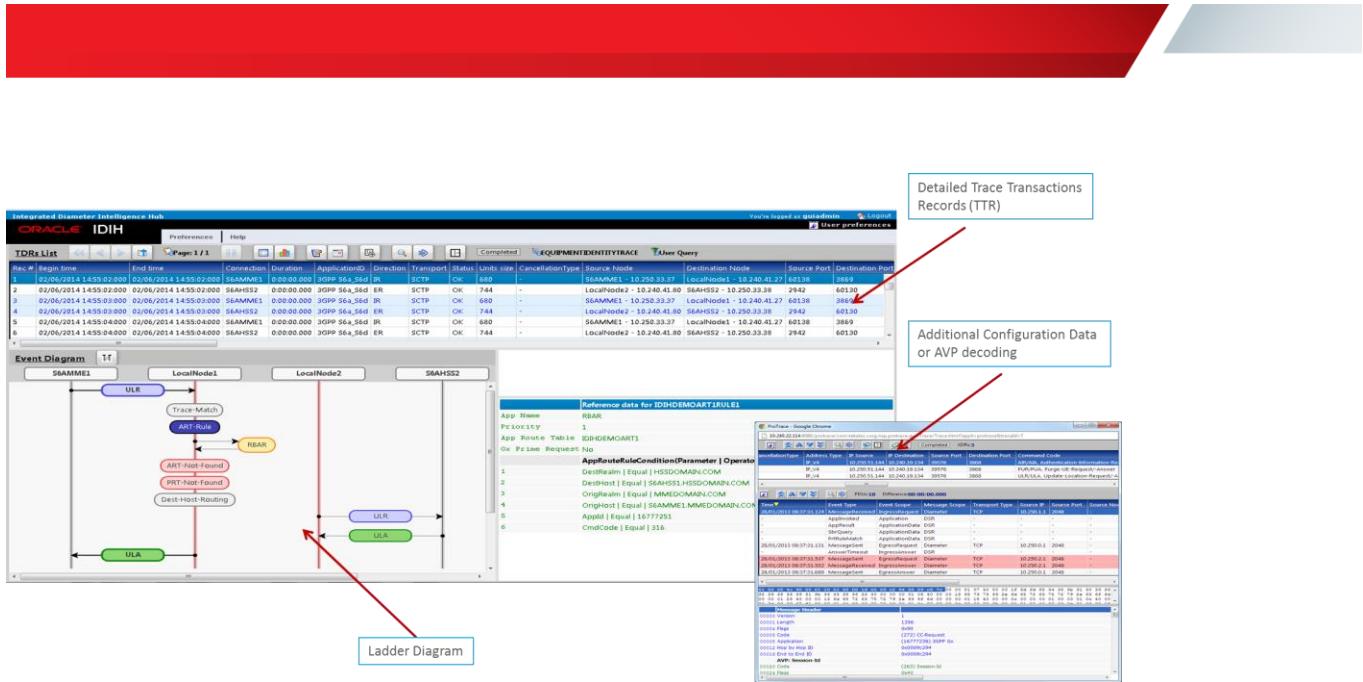
Figure 72 – IDIH Visualization GUI

Network IDIH (N-IDIH)

Operators with multiple DSRs have a need to diagnose and troubleshoot problems in their Diameter network with end-to-end visibility. N-IDIH provides support for network-wide IDIH trigger installation and trace analysis allowing centralized, end-to-end troubleshooting of transactions traversing any DSR in the network.

Whenever a Diameter message matches the trace criteria at a given site, the network trace also captures the path that the message took as it traversed through multiple DA-MPs within the network. Whenever a network trace is created, the trace criteria associated with the trace becomes active at each DA-MP within the network. Whenever a DA-MP determines that a particular diameter request or answer matches the trace criteria for an active network trace, the DA-MP captures the TTR associated with the Diameter transaction and forwards the TTR to the IDIH. In addition, the DA-MP compels any subsequent DSR node through which the Diameter message traverses to also capture TTR data associated with the Diameter message. Each DA-MP that was compelled forwards the captured TTRs to the IDIH associated with its site. The craftsperson can then use the DSR maintenance GUI from any DSR site to visualize the captured trace data, which includes TTRs captured at every site within the network.
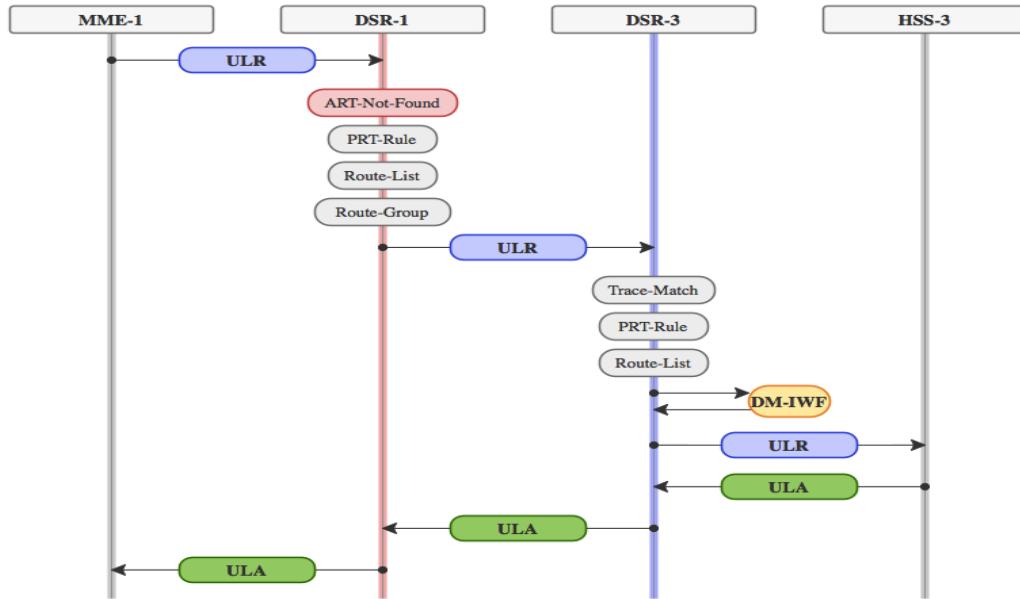
Figure 73 – Message Flow for Network Wide Trace

Supported Interfaces

IDIH supports a variety of Diameter Interfaces as a part of the rendering and visualizing messages within captured traces. In addition, DSR allows trace filters to be created for user identity, which is integrated with each of the supported interfaces. IDIH can render and visualize messages for other diameter interfaces beyond those that are officially supported, but any AVPs specific to those interfaces will not be available in the summary record of the TTR. IDIH cannot provide a full decode of AVPs specific to interfaces that are not specifically supported.

IDIH currently supports the following interfaces:

» Diameter (Base Protocol) – (can be used on all interfaces, but provides minimal information)
» Diameter Sh
» Diameter Cx
» Diameter Gq'
» Diameter S6a/d
» Diameter Gx
» Diameter Rx
» Diameter Gy
» Diameter SLg
» Diameter SLh
» Diameter Gxa
» Diameter SWm
» Diameter SWx
» Diameter Sta
» Diameter S6b
» Diameter S9
» Diameter Sd

» Diameter Sy

» Diameter S13

» Diameter Zh

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

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:
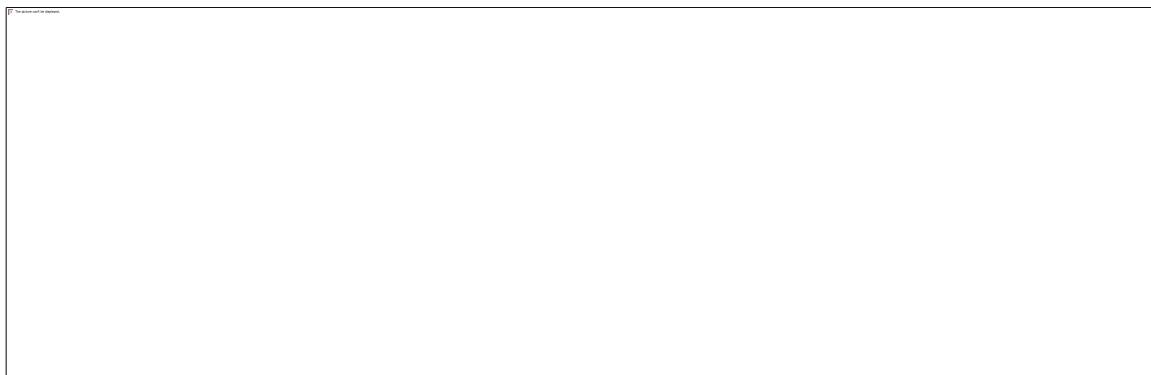

Figure 74 – DSR Architecture with STP VNF Message Processor

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.

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 protocols:

- 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.

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 inbound 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 13: 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:

1. 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.

2. 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.

3. DPC and SSN, route on GT –This result indicates that the final destination SSN should be determined by translation and be placed in the called party address. Further translation is required at the new node to determine the final DPC. The address indicator routing flag remains set to "route on GT."

4. 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."

5. 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:

1. Solitary –the GTT has a single node subsystem.

2. 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).

3. Load sharing –the load is shared equally between replicated subsystems.

4. 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, etc. node/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 Such a 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 SSN,

  - 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.

TCAP Opcode based routing enables new GTTSET Types

○ Opcode

○ IMSI

○ MSISDN

- 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.

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 handing for certain subscribers.
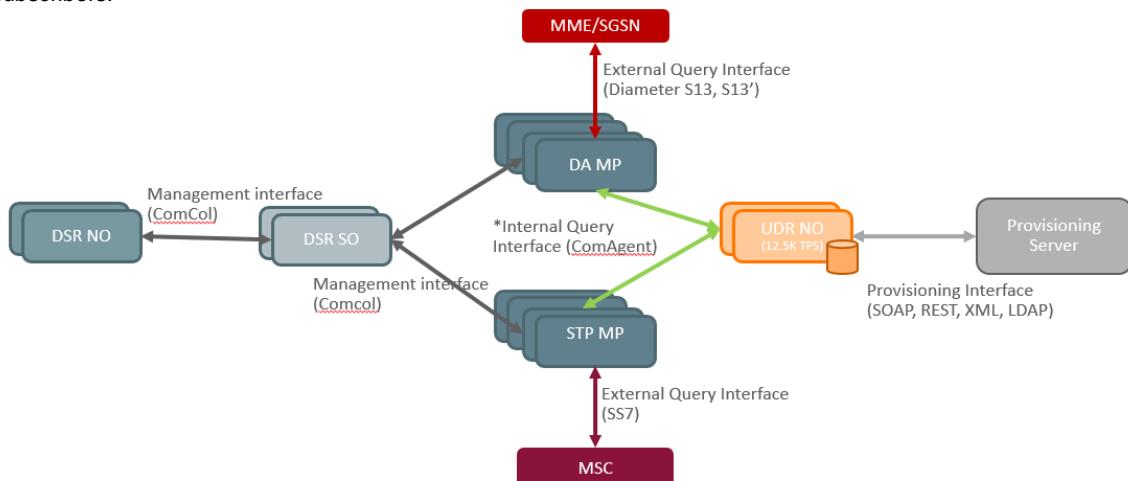

Figure 75 – EIR solution architecture.

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
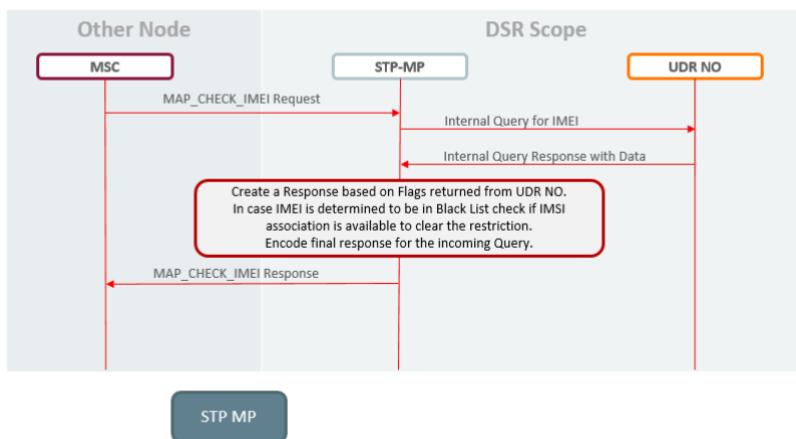
Figure 76 – SS7 EIR call flow
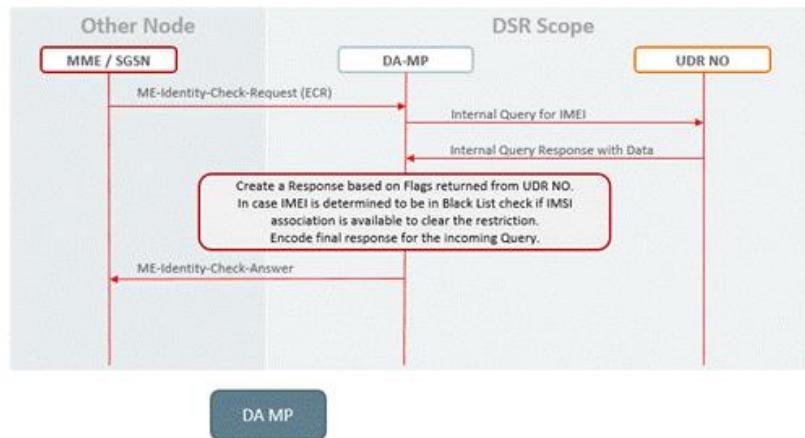
Figure 77 – S13 EIR Call Flow

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

ANSI ITU Conversion

- 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.

GTT Modification

- 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 i.e. 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.

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

MTP Screening

- MTP Screening feature 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.

Multiple PC

- 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, etc.).

- 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 applications.

- 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.

ITUI-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.

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 functions

- The vSTP currently provides full support for four types of point codes (PC) – ANSI, ITU-National (NI=10binary), 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.

Loop Detection

- The Loop Detection Feature design is both simple and flexible by using a Loopset table consisting of up to 12 point codes (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.

- Will detect 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.

IDP Query

- 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

    - Check whether INPQ message has valid request (the requestedInfo parameter must be MNPRequestedInfo 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.

Map based Routing

TOBR (TCAP Opcode Based Routing)

- 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 that 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.

FLOBR (Flexible Linkset Optional Based Routing)

- 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.

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 settypes will be linked by OPCODE settype or any of them.

  - IMSI
  - MSISDN
  - VLRNB
  - SMRPOA
  - SMRPDA

- The GTTSets of the types mentioned above are allowed to be provisioned ONLY in GTA entries from a GTTSet 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.

Cloud Native Diameter Routing Agent

- Overview
- Product Highlights.
- Deployed Components.
  - Overview of Micro services.
  - Brief of Helm Chart based Deployments.
  - Initiator and Responder Connections overview.

Overview

- Oracle **CNDRA** is a Cloud-Native Diameter Routing Agent based on micro-services architecture. It enables 3G/4G Diameter Routing functionalities. It leverages Oracle expertise from scalable ,reliable and robust DRA solution.

- CNDRA brings in the Synergies between Best-in-Breed Diameter Routing Solution and Cloud Native Delivery Model.

- Supports Container-based Cloud infrastructure to exploit the advantages of Cloud Native Model.

- It Leverages a common Oracle Communications cloud native framework.

- It is unprecedented advantage in Infrastructure reusability as per traffic pattern of the MNO.

CNDRA Content

- Introduction of the Service Based Architecture.
- DRAWorker and Overseer Service .
- With this release there is not support for GUI
- Support for TCP Protocol.
- Support for Fixed Initiator Mode and Responder Mode Connection.
- Partially Integrated with the Common CNE Services.
- Logging Display on EFK.
- Helm Chart Based Installation.
- RBAR Application.
- Core Routing Function.
- MMI.

Product Highlights.

Monolithic to Micro services.

- Change in the Architecture from NO/SO/DAMP to Micro services.
- Currently CNDRA encompasses micro services namely, "**overseer**" and "**dra"** service.
- Configuration functionalities are taken care by "**Overseer"** Service.
- Signaling Traffic processing is taken care by "**dra"** Service

One Step Deployment

- Helm Chart is used to deploy micro services "**overseer**" and "**dra**" service.
- Helm felicitates deployment of CNDRA on Bare Metal CNE environment.
- Deployment when compared to Legacy system is much quicker.


Cloud Native Load Balancer
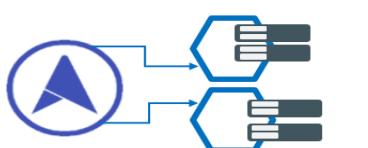
- MetalLB Load balancer utilized for traffic distribution. The Legacy IPFE will no longer be needed.
- Connection Load balancing among "Worker-Nodes" is provided by MetalLB.
- Load balancing among the PODs at given Node is provided by Kube-Proxy.


Enhanced Dashboards

- With the power EFK framework it is possible to create customizable dashboards.
- Fluentd collects data and pass to elastic search. It provide storage for collected data.
- Kibana gets collected data and display on the Kibana GUI.
-

- **CNDRA Deployment**
- CNDRA uses Helm for deployment, it helps to define, install and manage Kubernetes applications. Helm uses packaging format called charts.
- A chart is a collection of files that describe a related set of Kubernetes resources.
- Helm will deploy "**overseer**","**draworker**" service and Job to configure PODs. Helm install, will be blocking until "**JOB**" POD completes successfully. i.e. OAM PODs are configured in COMCOL topology.
- Once done, user will be able to configure Diameter configuration using MMI/RESTful commands.
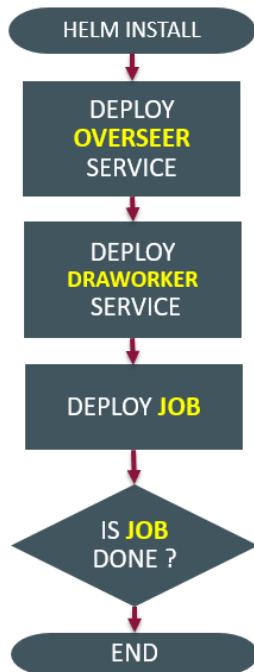
CNDRA 1.0 : MICRO SERVICES

DRAWorker Service
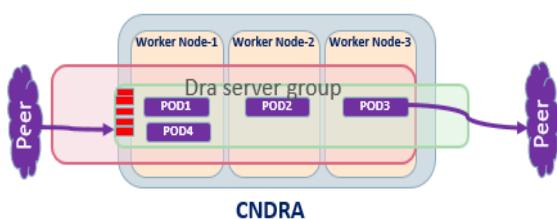
It is similar to DA-MP Node in the Legacy Architecture.

•This service is responsible to process all the signaling traffic.

•It is consisted of only those PODs which are solely responsible for Routing, Traffic and Connection Management .
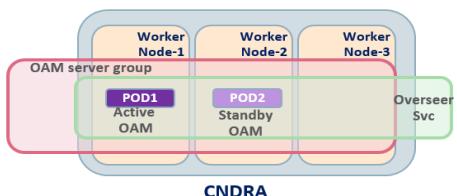
- This service provides the OAM Functionality.

- Using MMI Commands, all the diameter configuration is done.

CNDRA 1.0 : CONNECTIONS

- Diameter peer will use mp-service Load Balancer's IP when initiating a connection toward CNDRA.

- Any mpPod in topology can accept the incoming connection request from Peer.
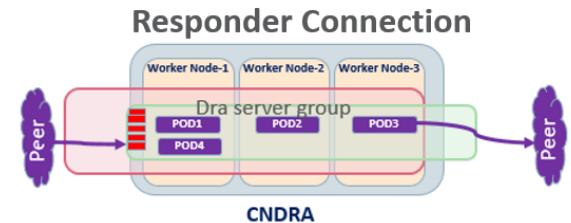
- The CNDRA has pre-defined 5 service ports, exposed for Diameter service. (Note - pre-defined ports are 3868, 3869, 3870, 3871, 3872)

- However user can specify/overwrite the set of port during deployment using HELM's custom value.yaml.
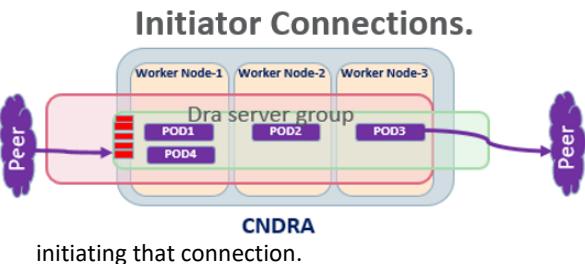
- Initiator connections are distributed based on number of message processors configured in topology.

- During the initiator connection configuration, the connection will be assigned an mpPod. The mpPod will be responsible for initiating that connection.

5GC NF

Service Proxy function (SPF) and Network Repository function (NRF)

The Service Proxy Function (SPF) aims at addressing the challenges introduced by the 5G Service Based Architecture and the usage of HTTP/2 for signaling. They key benefits are in providing an effective congestion control mechanism, simplifying the network by reducing the connection mesh to a manageable number of connections etc. The Network Repository Function (NRF) aims to provide service Registration & Discovery functions of 5G NFs. Discovery services allow for NFs to discover each other, again without manual configuration.

SPF + NRF functionality supported in 8.4 release are as follows:

SPF services supported:

- Oracle's Service Communication Proxy(SCP) resolves the challenges introduced by the 5G Service Based Architecture(SBA).
  – Routing/Selection – consumers need to define routing rules, refresh cache, and handle application failures/redirects
  – Load Balancing – based on static capacity, no feedback loop, NF specific
  – NF Degradation and Failures – timeouts can result in cascaded failures, need for circuit breaker
  – Traffic Prioritization - No mechanism in SBA for prioritizing important service requests, 'stream priority' is inadequate
  – Congestion and Overload – no congestion control, primitive overload mechanisms, DOIC like mechanisms under study
  – HTTP/2 for Signaling – New to telecom, built for large # of short-lived connection, protocol limit requires re-establishment of TCP connections
  – Visibility into Signaling Core – distributed architecture makes it hard for common visibility
  – Authorization and Authentication – SA3 has defined NRF as OAuth 2.0, certification mgmt. is not defined
  – Interoperability EPS/5GS & Diameter/HTTP – not defined in 3GPP

- SCP provide additional value in service framework in terms of signaling optimizations
  - Efficient use/Load Balancing across NF instances
  - Offloading NFs with alternate Routing and Retries
  - Limiting the effect of Cascaded failures
  - Provides visibility into the 5G network
  - Resolve interoperability issues
- SCP extends Istio/Envoy for use as a 5G-aware service mesh.
- SCP deploys into cloud native environments using Kubernetes (k8s).

SCP features

- Relay Routing for all interfaces
- Load Balancing and Alternate routing for PCF, UDM interfaces
  - Within site
  - Visibility via Metrics, KPIs and traces
- NF Egress congestion control
- Remote Busy
- SCP overload control
- Memory based
- Support for setting priority header
- Support for NF Outlier Detection
- Support for NF Circuit Breaking
- Support for NF Canary Testing
- NRF client services
- Local redundancy
- Static Routing to NFs
- Performance/Capacity –
  - Max TPS supported per worker pod
  - Max TPS supported per site

NRF Overview

- Oracle Communication Network Repository Function (NRF) serves as the repository for all 5G network functions and allow them to register and discover each other via a standards based API.
- Supports the following standard NRF Services:
  - NF Management,
  - NF Discovery
- Follow micro-service architecture principles.
- Support containerized deployment in the Cloud Native Environment (CNE)
- Supports Service Communication Proxy (SCP)
- Can be deployed without SCP but provides additional synergies and benefits when deployed with Oracle SCP

NRF features

- NF Management functions
  - NFRegister

- o NFUpdate
- o NFDeregister
- o NFStatusSubscribe
- o NFStatusUnsubscribe
- o NFStatusNotify
- Discovery functions
  - o NFDiscover
- Enhanced discovery
  - o Configuration provided to limit the instances returned in discovery result.
- Metrics, KPIs, Alerts, tracing and Logging
- Auto-Scaling up/down
- Support for Custom NF
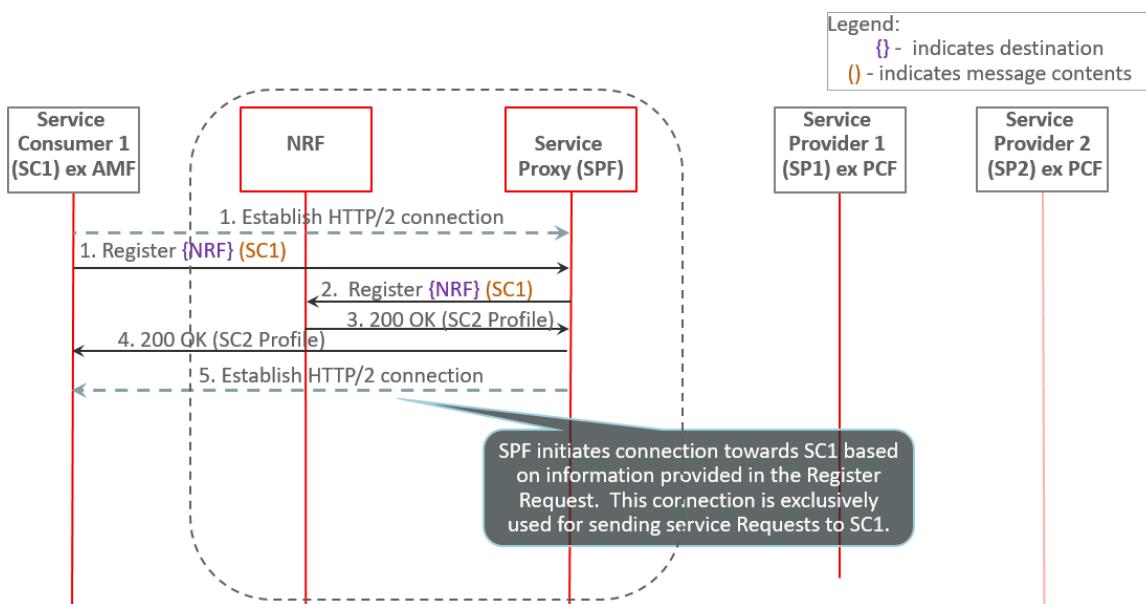


Figure 78 – NF registration call flow - Consumer
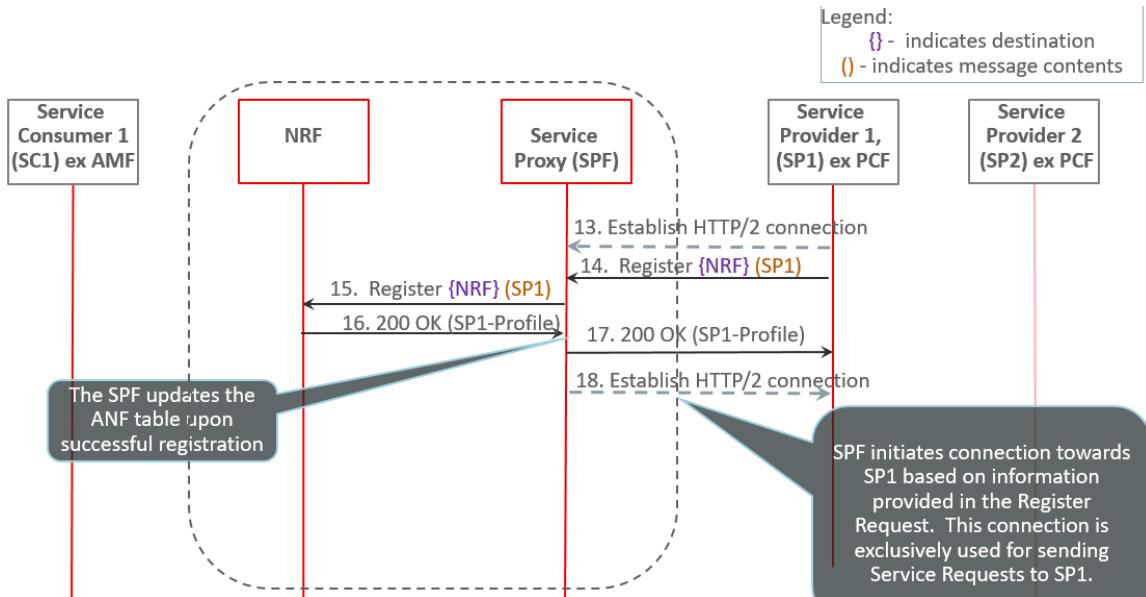
Figure 79 – NF registration call flow - Provider
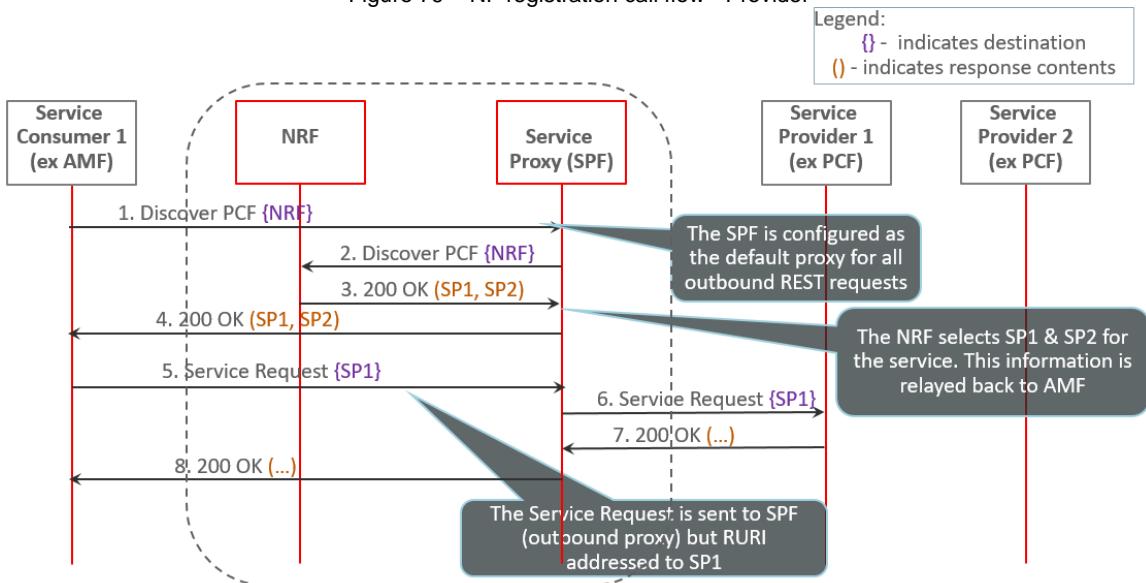

Figure 80 – Discover & Service Request Routing

CNE

Overview

Nothing has been more of a stimulus for agility than the Cloud, especially for the webscale companies whose services pose the greatest threat to the telco industry.

Why is a Cloud native platform significant for CSPs?

- Provides businesses the ability to rapidly develop, test and launch software applications that drive business growth at a rate and scale not seen before.
- Through virtualization, Cloud enables the abstraction of the IT stack so that one physical resource is used logically like many, creating more flexibility and scalability, and opening the door to new service models.
- The addition of automation further allows the focus on other issues such as core business logic, security, and analysis, as opposed to just provisioning and maintaining the resources.

All of this together supports the "anything-as-a-service" paradigm of Cloud, driving a mostly static and transactional world into a more dynamic and subscription-based economy.

In addition Oracle Communications Common Signaling NF CNE also provides:

- A common platform that removes the need to have different operational and maintenance processes for different 4G and 5G Network functions.
- The power and innovations of industry leading 3rd party open source tools that are in use by Oracle and other major IT businesses.
- A 5G Service Based Architecture leveraging technologies such as microservices, 3 tier product architectures, lifecycle management, and CI/CD.

| DevOps | Continuous Delivery | Microservices | Containers |

Figure 1. Component of a Cloud native environment

A Cloud native environment enables CSPs to move between their private Cloud and public Cloud. DevOps and CI/CD are key concepts for a successful Cloud deployment, these practices significantly shortens the software delivery times, and enable full test automation, and operations. The microservices architecture enables flexibility, agility, and reuse (e.g. a micro-service can be reused across various platforms such as 5G signaling, IP Multimedia Subsystem (IMS), and Session Border Controller (SBC)). It is also possible to run multiple versions of same microservice, enabling blue/green and canary deployments. Lastly container management solves network orchestration problem and allow CSPs to focus on service orchestration.

Product description

Oracle Communications Common Signaling Network Function Cloud Native Environment (NF CNE) is based on leading industry open source services and tools, it enable operation and management of 5G Network Functions (NFs) in a production and carrier grade environment. With the NF CNE, and leveraging Oracle's vast experience in Core network implementation and Cloud services implementation, Oracle Communications provides robust and reliable 5G implementation that is secure and operationally viable.
Oracle Communications Common Signaling NF CNE provides a holistic set of services to efficiently manage the 5G NFs. The flagship services of NF CNE includes:

- Runtime services (virtualization and networking)
- Orchestration and management services
- Storage services
- Security services
- Observability services (logging, tracing, and metrics)

Oracle Communications provides a secure, operationally viable, and robust and reliable 5G implementation with its NF CNE.

NSSF Overview
Overview

Networks are evolving quickly to meet the demands of a modern digital world. The latest evolution of 5G offers more efficiency, security, and flexibility than any "G" that has come before. It will allow for new, highly customized, innovative services across many industries. This is significantly enabled by network slicing which puts an end to the one-size-fits all network approach, providing the opportunity for communication service providers to personalize connectivity services to the precise requirements of particular applications, users, or devices. This will empower CSPs to launch and evolve custom-fit network slices-as-a-service rapidly with lower capital and operating costs that can provide opportunities for new revenue-generating digital business services.
The service based architecture of 5G along with cloud native core allows service providers to operate multiple dedicated networks on a single platform, referred to as 'network slicing'. A network slice is an independent end-to-end logical network and may include: core network control and user plane network functions, 5G Radio Access Network and N3IWF (Non-3GPP Inter Working Function) to the non-3GPP Access Network. This provides CSPs more flexibility to configure network slices based on application, transport, network, radio, and connectivity requirements.

Network slicing not only optimizes a network for a business requirement, but it also reduces the cost by logically partitioning the same network to serve multiple use cases that may require the same network characteristics. The customizable network capabilities in network slicing includes data speed, quality, latency, reliability, security, and services.
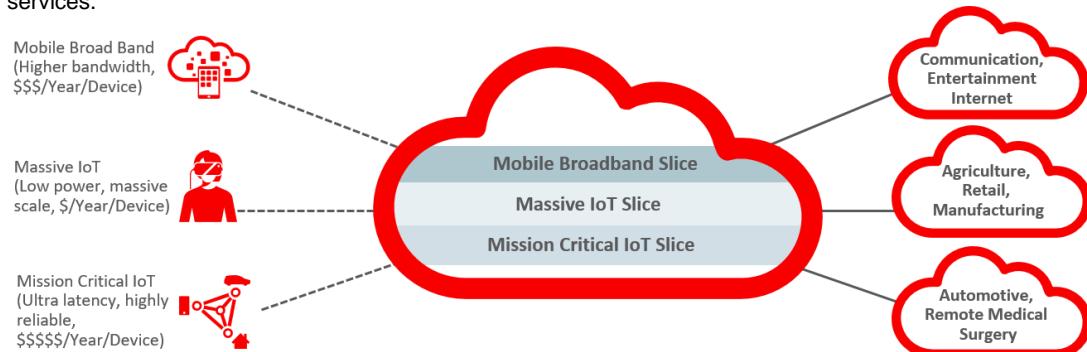
Figure 1. 5G Network Slicing

Product description

Oracle Communication's Network Slice Selection Function (NSSF) enables network slicing by determining the individual network function pertaining to a slice. Designed and built as a microservices on Cloud Native principles, Oracle Communication's NSSF runs on top of a Docker container managed by Kubernetes, this allows easier upgradation and scaling of the solution. NSSF allows dynamic slice selection for building use cases as per 3GPP release 15 standards.
Slice selection is performed based on various parameters. The prominent parameters that Oracle Communication's NSSF uses to configure a network slice are as follows:

Table 1: Parameters used by Oracle Communications NSSF for network slicing

PARAMETER	DECSRIPTION
NSSAI	Network Slice Selection Assistance Information

PARAMETER	DECSRIPTION
TA	Tracking Area
PLMN	Public Land Mobile Network
RAT	Radio Access Technology
NRF	Network Repository Function
NSI ID	Network Slice Instance Identifier

Oracle Communication's NSSF also deploys significant parameters like roaming, slice requirement info, load and analytics data to perform slice selection function.

NSSF Core Features

Oracle Communication's NSSF is enriched with best of industry standard features and agile and microservices engineering practices. Its core features include:

- Slice Selection capabilities by determining authorized slice information based on Network Slice Selection Assistance Information (NSSAIs) provided by the Access and Mobility Management Function (AMF) as part of registration procedure.

- Enables User Equipment (UE) initial registration procedure assisting Access and Mobility Management Function (AMF) to determine a specific network function or contact the slice specific Network Repository Function (NRF).

- NSSF supports Protocol Data Unit (PDU) session establishment procedure assisting AMF to select the appropriate NRF or the Network Slice Instance Identifier (NSI ID) thereby enabling AMF determine the Session Management Function (SMF).

- NSSF provides network and PDU session establishment processes in non- roaming scenario.

- NSSF supports operator configuration to update Single Network Slice Selection Assistance Information (S-NSSAI) support per Tracking Area (TA) per Public Land Mobile Network (PLMN) to determine NSI/Target AMF, NRF, etc.

Oracle Communication's NSSF supports 3GPP Technical Specification 29.510 specified NF registration procedure with the NRF by sending NF HEARTBEAT messages towards the NRF. NSSF also supports NF discovery procedures enabling NRF to provide the best possible AMF which can be used during an initial registration. Oracle Communication's NSSF extends basic KPI support for registration and PDN (Packet Data Network) session establishment use cases. It stores operator policies and network slice information in its embedded database, and supports logging and tracing features through CNE (Cloud Native Environment)
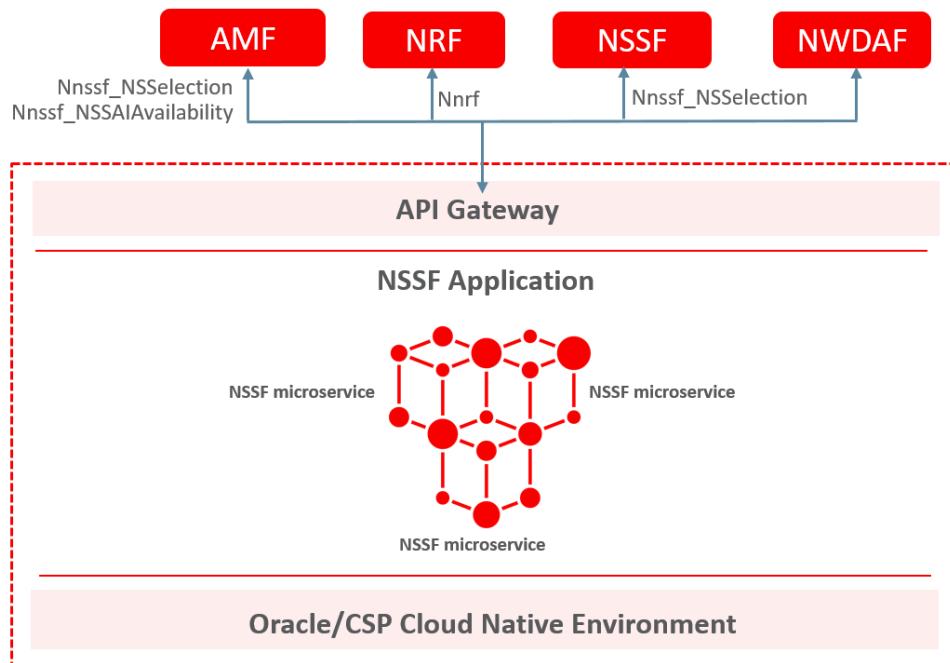
Figure 2. Oracle Communication's NSSF system Architecture

UDR

- Oracle 5G UDR, implemented as cloud native function, offers a unified database for storing application, subscription, authentication, service authorization, policy data, session binding and Application state information. UDR is one of the main key component of the 5G Service Based Architecture.

- Oracle's 5G UDR :

    – leverages a common Oracle Communications Cloud Native Framework

    – compliant to 3GPP Release 15 specification

    – has tiered architecture providing separation between the connectivity, business logic and data layers

    – uses MySQL NDB Cluster as the backend database in the Data Tier

    – registers with NRF in the 5G network, so the other NFs in the network can discover UDR through NRF.

    – In the initial release, only PCF data is supported by UDR.

- UDSF(Unstructured Data Storage Function) supports storage and retrieval of unstructured data by any NF. This functionality is part of Oracle's 5G UDR solution.

BSF Overview

Overview

Binding Support Function (BSF) is providing Binding Support Management Service.
The Nbsf_Management service is used for the BSF to provide a PDU session binding functionality, which ensures that an AF request for a certain PDU Session reaches the relevant PCF holding the PDU Session information.
This service:
- allows NF service consumers to register, update and remove the binding information; and

- allows NF service consumers to retrieve the binding information.

The Binding Support Management Service (Nbsf_Management) is exhibited by the Binding Support Function (BSF).
Known consumers of the Nbsf_management service are:
- Policy Control Function (PCF)

- Network Exposure Function (NEF)

- Application Function (AF)

The BSF is a function that can be deployed standalone or can be the functionality provided by other network functions, such as PCF, UDR, NRF, SMF.

Product description

Oracle BSF allows PCF to register/deregister itself with subscriber session binding information to be discoverable by any NF service consumers. It also support Rx to enable operators to continue using existing IMS deployment until IMS transition to SBA is made
Solution highlights -
• Cloud Native based on Micro-services architecture
• Elastic backend database that enables redundancy, scaling
• Binding Support Management Service support
• 5G-LTE Interworking support
• Integration with DevOps workflows

BSF Core Features

Oracle Communication's BSF is enriched with best of industry standard features and agile and microservices engineering practices. Its core features include:

• Suppprt 3GPP Technical Speciification 29.521 for BSF Management Service

- Register Service - registers the binding information in the BSF for a UE when an IPv4 address and/or IPv6 prefix is allocated, or a MAC address is used for the PDU session

- Deregister Service – removes the binding information in the BSF when an IPv4 address and/or IPv6 prefix is released, or a MAC address is not used for the PDU Session.

- Discovery Service – discover the selected PCF by using the Nbsf_Management_Discovery service operation when it is allowed to interact directly with the policy framework for policy control.

• Diameter Rx Proxy

- Diameter Rx proxy handling (creation, modification and termination)

- Supporting for seamless interworking with P-CSCF/DRA

- Suspect/Stale binding and session cleanup

- Support 3GPP Technical Speciification 29.510 for NRF Management Service

  - Register Service - It allows BSF to register its profile in the NRF; it includes the registration of the general parameters of the BSF, together with the list of services exposed by the BSF

  - Update Service - It allows BSF Instance to replace, or update partially, the parameters of BSF profile (including the parameters of the associated services) in the NRF; it also allows to add or delete individual services offered by the BSF Instance.

  - Deregister Service - It allows BSF Instance to deregister BSF profile in the NRF, including the services offered by the BSF Instance.

  - Subscribe Service - It allows BSF Instance to subscribe to changes on the status of BSF Instances registered in NRF.

  - Notify Service - It allows the NRF to notify subscribed BSF Instances of changes on the status of BSF Instances.

  - Heat-Beat Service - Each BSF that has previously registered in NRF shall contact the NRF periodically (heart-beat), by invoking the BSF Update service operation, in order to show that the BSF is still operative.

- Platform features:

  - Configuration Management

  - Local and Geo-redundancy Support (DB layer)

  - Basic KPI, Stats report through Cloud Native Environment

  - HTTPS and HTTP/2 support
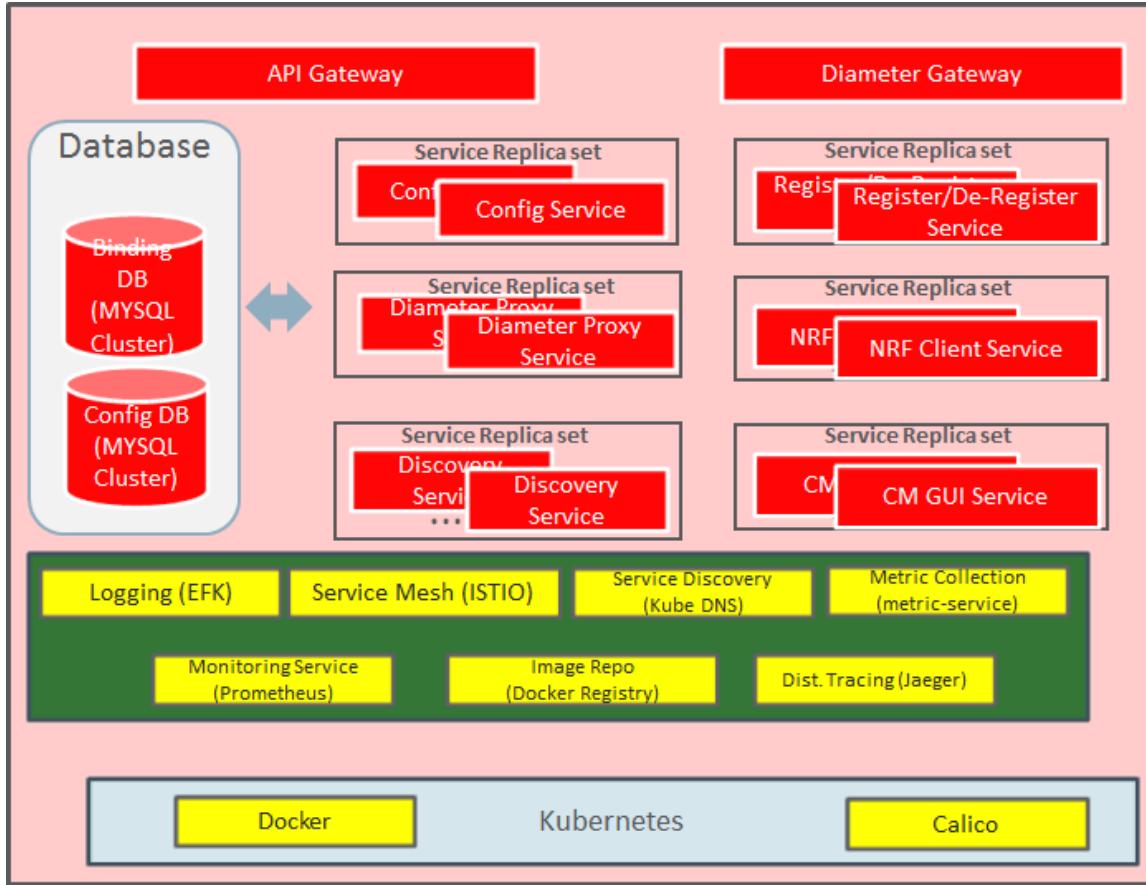
SEPP

Oracle Communication Common Signaling Security Edge Protection Proxy (SEPP) acts as an non transparent

proxy sitting at the perimeter of the Public Land Mobile Network (PLMN) network enabling secured inter Network

Function (NF) communication across PLMN networks. The SEPP makes 5G core secure and less susceptible to

any kind of attack or malpractice by maintaining confidentiality and integrity of inter NF messages.

Product description

Oracle Communication Common Signaling, Security Edge Protection Proxy (SEPP) provide authentication, confidentiality and integration protection for inter-PLMN Service Based Interfaces (SBI) signaling traffic between 5GC NFs. The SEPP is an evolution of 4G Diameter Edge Agent (DEA) and acts as a reverse proxy providing a single point of access and control to internal NFs from interconnect.
Oracle Communication SEPP comes from the same secure and reliable Oracle engineering that puts security as its number 1 priority. The prominent features of the SEPP are listed below:

- Implements N32 interface to support inter-PLMN communication between NFs.

- Provides robust congestion control to protect against the roaming signaling storms. SEPP performs threshold based roaming traffic rate limiting to protect core network NF from overloading. SEPP supports traffic (re)prioritization which allows intelligent throttling of signaling traffic during overload and congestion conditions.

- Provides topology hiding functionality to hides the internal information of a PLMN which is not required outside the PLMN by changing/removing from all egress messages.

- Provides flexible routing control and resiliency allowing routing of ingress roaming signaling to core network NF's and routing of egress roaming signaling to interconnect. SEPP also supports alternate routing in the case of an error message from the vSEPP or retry the request with the same provider in the case of a timeout.

- Provides the built-in signaling firewall where SEPP applies different security countermeasures to roaming traffic. SEPP determines if the ingress roaming signaling message is a vulnerable message. These security countermeasures will allow operators to mitigate different security attacks like DoS, Location tracking, fraud attacks, etc. initiated from interconnect roaming traffic.

- Provides the mediation capability which allows operators to pre-configure the policy rules. These rules are used to modify the ingress and egress roaming signaling messages at SEPP. SEPP supports modification of messages using pre-configured protection policies negotiated with MNO and supports integrity protection of modified messages in IP exchange (IPX) proxy mode.

- Supports traffic shadowing, it allows operators to copy the required roaming signaling and perform offline analysis. SEPP also supports analytic feed of signaling firewall statistics and roaming signaling pattern to analytic engine.

- Supports traffic steering of ingress roaming registration procedure signaling messages based on pre-configured roaming agreement rules.
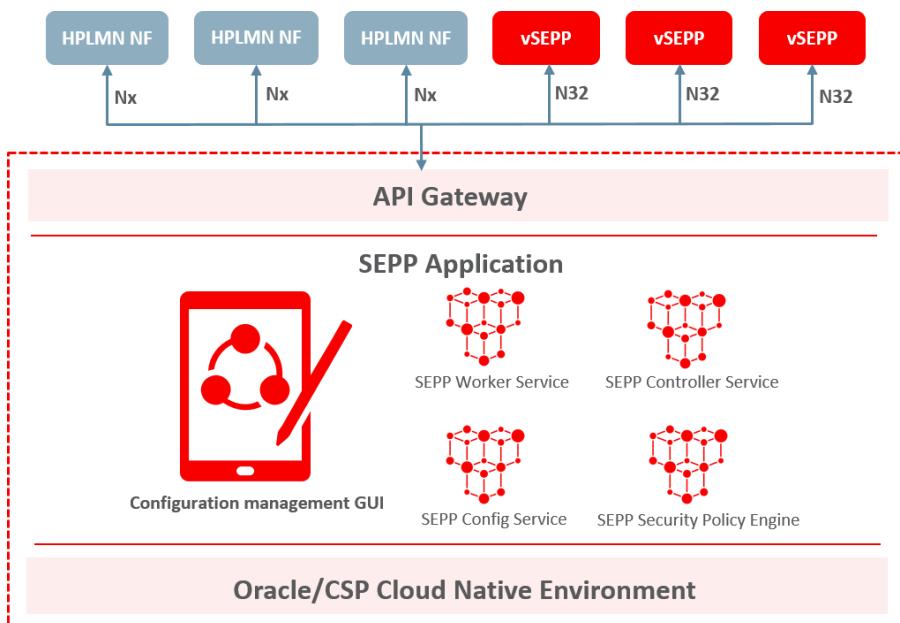
IWF

Oracle Communication Common Signaling InterWorking and Mediation Function (IWF) enables 5G core Network Functions (NFs) to communicate effectively with Evolved Packet Core (EPC) network elements. The IWF serves as a bridge between a 4G EPC and a 5G core and eases the process of transitioning to 5G for CSPs.

Overview

The accelerated evolution of telecom industry to 5G is being driven by the growing demand for Enhanced Mobile Broadband, Ultra-Reliable and Low Latency Communication, and Massive IoT. CSPs have already invested and created an ecosystem around a 4G EPC, transitioning to a 5G NGC would not happen in one go. CSPs across the world are looking forward to incremental deployment of 5G elements, for this they need to place an intelligent and reliable InterWorking technology that can enable effective communication between a 5G NGC and a 4G EPC.

Product description

Oracle Communication Common Signaling InterWorking and Mediation Function (IWF) connects a 5G NGC to a 4G EPC by enabling seamless communication between 5G NFs and EPC network elements. Oracle Communications IWF can be deployed as independent network function in the 5GC network or as the part of Oracle 5GC NFs like NRF, SEPP, SCP, etc. as an independent microservice within the 5GC NF.

Oracle Communication IWF is designed with world class Oracle Engineering and industry best practices. Its prominent features include:

- **Protocol Translation**

Protocol translation feature of IWF enables interworking between EPC network elements and 5GC NFs by performing mapping of HTTP/2 to Diameter messages and vice versa. It supports following functionalities:

- Provides Diameter-HTTP/2 interworking based on generic framework using Information Elements (IE) mapping configuration.

- Supports HTTP/2 to HTTP/1.1 protocol translation.

- Supports payload transformation from JSON to XML.

- Supports customization of inter-working IEs of HTTP/2 and Diameter messages.

- Provides load balancing of 5G Core NF traffic across EPC Network Elements.

- Supports congestion control and overload control to protect 5G Core NFs and EPC Network Elements.

Figure 1. Oracle Communications IWF Protocol Translation Service

- **Message Mediation**

Mediation framework is a rules based engine that alleviates interoperability issues through easy creation of mediation rules via IWF Mediation Wizard and allows for HTTP/2 extensions. It supports following functionalities:

- Allows to extend HTTP/2 behavior and introduce new complex use cases rapidly.

- Provides customer with a framework to configure their networks to adapt to changes.

- Near zero performance impact.

- Supports versioning of mediation rules using GIT repository.

- Supports execution of the mediation rules in test mode before deployment.

Figure 2. Oracle Communications IWF Mediation Framework

Oracle Communication Common Signaling IWF is built as a microservice on Cloud native architecture.
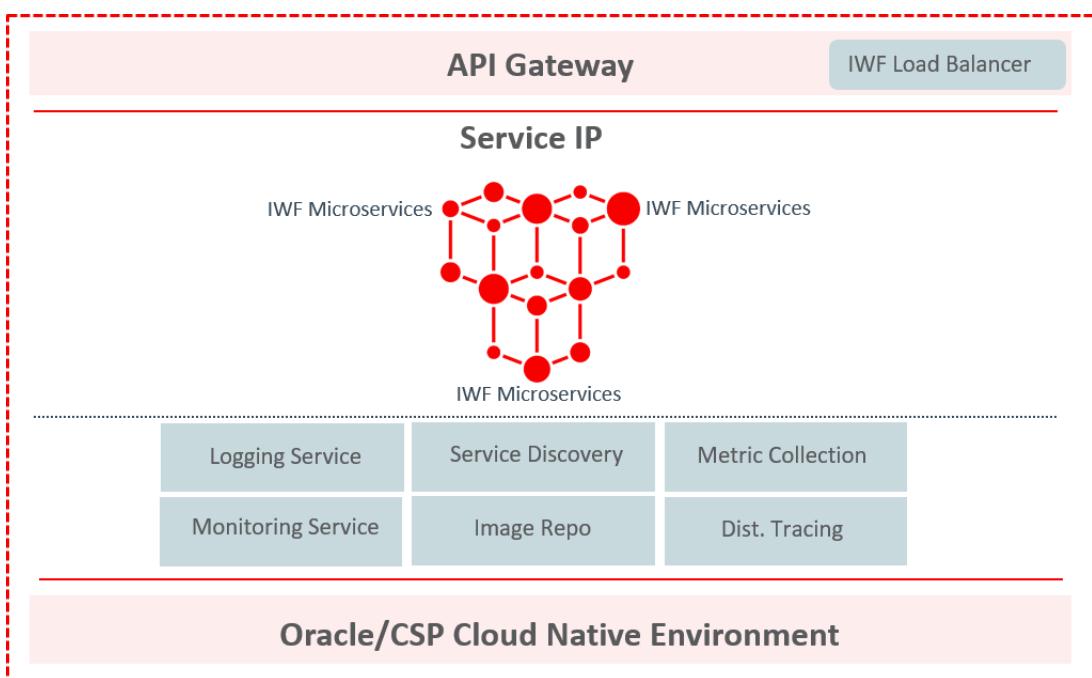

Figure 2. Oracle Communications IWF System Architecture

NEF

Oracle Communication Common Signaling Network Exposure Function (NEF) securely expose the 3GPP network capabilities of 5G to Application Functions (AFs). The NEF act as a centralized point for service exposure and play a key role in authorizing all access requests originating from outside 3GPP network to enable cellular IoT, non-IoT, edge computing and API gateway use cases for operators.

Overview

5G opens a sea of opportunities for businesses to dev

elop applications and business models of the future. To support this next gen tech. the intelligence hub of the 5G network is architected on the principles of Service Based Architecture (SBA) where different Network Functions (NFs) are deployed and seamlessly work together to provide various mobile network functionalities. To fully exploit the true capabilities of a 5G network the 3GPP defines a secure and reliable way to expose the network capabilities of the 5G network to application developers in the enterprises, enabling them to easily develop Over-The-Top (OTT) applications. CSPs, therefore are required to deploy a robust, secure and reliable exposure function in their 5G network to boost their revenue and provide best features and applications to their customers.

Product description

Oracle Communication Common Signaling NEF securely exposes the 3GPP network capabilities of 5G to Application Functions of different services thereby enabling cellular IoT, non-IoT, edge computing and API gateway use cases for operators. The NEF is designed based on the specification defined by the 3GPP Release 15 standard.

Oracle Communications NEF architecture consists of two main components:

1): Common API Framework (CAPIF) based API Gateway: Provides the secured northbound interface towards both trusted and non-trusted AFs. Oracle Communications NEF's CAPIF API Gateway provides the following API management and governance functionalities:

o Make APIs available through publishing and packaging and provides security to ensure that only the authorized AFs can have access through API authentication.

o Provides built-in API Firewall to protect against different security attacks.

o Provide support for on-boarding and managing partners.

o Supports metering and charging features to monetize the API usage.

o Enables API creation by curating the APIs with GUI based tools.

o Supports Policy and Service Level Agreements by providing specific policies toward the consumer, provider and network.

o Supports monitoring and analytics to track API usage that can be used to improve features and services.

o Supports versioning and retirement by enabling the use of multiple versions of the API and retirement of these versions as needed.

Oracle NEF CAPIF API Gateway supports both centralized and distributed deployment models. Distributed deployment model provides Mobile Network Operator (MNO) the flexibility to deploy NEF catering to multiple API provider domains, allowing NEF API gateway to be used as generic API gateway even for non 3GPP service exposure services.

2): Network Exposure Services: Exposes following 3GPP capabilities to AFs through northbound interfaces:

o Enables monitoring event procedure by allowing IoT AF to monitor the device state location information, roaming status, IoT device reachability, connectivity status, etc.

o Provides application trigger procedure by allowing IoT AF to deliver the device trigger to IoT devices in the form of SMS.

o Provides network parameter provisioning by allowing external party to provision information which can be used for the IoT Device in 5GC. This service allows provisioning of Predictable Communication Patterns (CP) of an IoT device that may be provided by the Application Server to the NEF in order to enable network resource optimizations for such IoT Device(s). This capability can also be used by IoT AF to provide the 5G System (5GS) with suggested parameter values for maximum latency, maximum response time and suggested number of downlink packets for the IoT devices using Extended Discontinuous Reception (eDRX) and power saving modes.

o Supports Policy and Charging procedures.

o Provides QoS Control & Sponsored data connectivity.

o Provides background data transfer scheduling and AF Traffic influence.

o Supports small data communication by allowing transfer of small infrequent IP and non-IP data along with reliable data service support.

o Provides enhanced coverage restriction control by allowing the IoT use case deployment with poor network coverage area.

Oracle Communications NEF also supports EPC-5GC interworking to enable 3rd party service provider get access to authorize the available 3GPP network exposure services when the User Equipment (UE) is either on EPC or on 5GC. This is achieved using common northbound interface from NEF and Service Capabilities Exposure (SCEF) towards AF and using shared context information with SCEF.
Oracle Communications Common Signaling NEF is engineered as a microservices based Cloud native architecture. It is architected into three logical grouping:

o **Connectivity**: Composed of components interfacing with external entities

o **Business Logic**: Application layer that runs the CAPIF Gateway, Network Exposure Services and other services.

o **Data Mgmt:** Data layer responsible for storing various types of persistent data
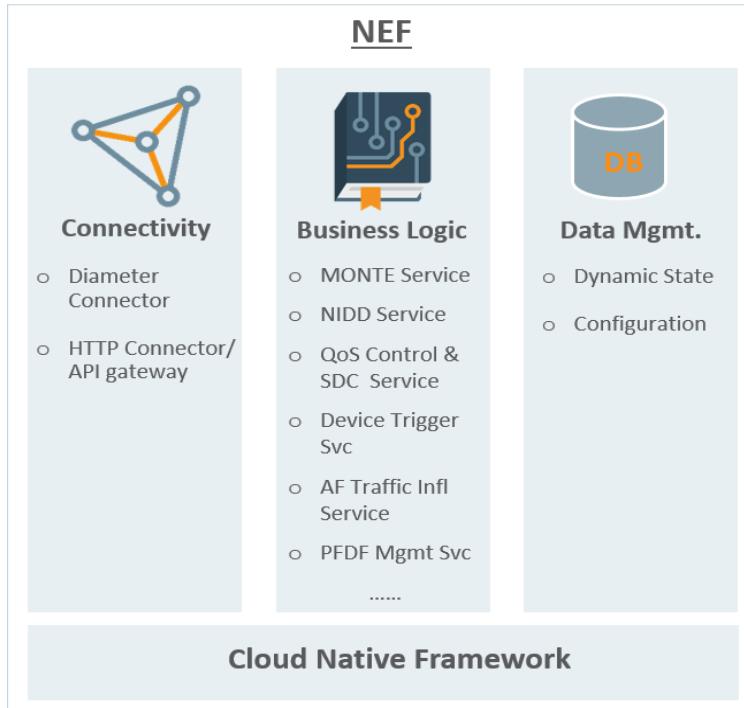
Figure 1. Oracle Communications NEF System Architecture

Flexible IP Addressing

The DSR supports Ipv4 and Ipv6 simultaneously for local DSR node addressing. Optionally, either an Ipv4 or Ipv6 address can be defined for each Diameter connection. The DSR supports both Layer 2 and Layer 3 connectivity at the customer demarcation using 1GB and optionally 10 GB (signaling only) uplinks.

The Oracle DSR supports establishing Diameter connections with Ipv4 and Ipv6 peers as follows:

» Multiple Ipv4 and Ipv6 IP addresses can be hosted simultaneously on a DSR MP utilizing dual-stack capability in the DSR operating system.
» Each Diameter connection (SCTP or TCP) configured in the DSR will specify a local DSR node and an associated local Ipv4 or Ipv6 address set for use when establishing the connection with the peer.
» Each Diameter connection (SCTP or TCP) configured in the DSR will specify a Peer Node and optionally the Peer Node's Ipv4 or Ipv6 address set.
» If the Peer Node's IP address set is specified, it must be of the same type (Ipv4 or Ipv6) as the local DSR IP address set specified for the connection.
» If the Peer Node's IP address set is not specified, DSR will resolve the Peer Node's FQDN to an Ipv4 or Ipv6 address set by performing a DNS A or AAAA record lookup as appropriate based on the type (Ipv4 or Ipv6) of the local DSR IP address set specified for the connection.

The DSR supports Ipv4/Ipv6 adaptation by allowing connections to be established with Ipv4 and Ipv6 Diameter peers simultaneously and allowing Diameter Requests and Answers to be routed between the Ipv4 and Ipv6 peers.

Full IPV6 Support

As the global public Ipv4 address pool is getting exhausted, more and more customers are requiring Ipv6 support to either field new deployments or to grow existing deployments. This feature provides Ipv6 support on all internal and external management interfaces. This along with DSR's support for Ipv6 on signaling networks, allows DSR to be deployed without the need for Ipv4.

This feature provides support for Ipv6 for all Communication Agent functionality. This includes all solution components such as DP, all variations of SBRs, and the internal bus communication agent. Also, Ipv6 is supported across all server group functions (including SS7 blades), SDS, and IDIH. There is also Ipv6 support for the SDS provisioning interface. And finally, for any OAM interface (internal/external devices) the user can either configure an Ipv4 or Ipv6 or both Ipv4 and Ipv6 address. The external servers currently supported by the OAM are LDAP servers, export servers, DNS servers and SNMP servers.

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 & deletion of provisioned data. Both the import & export operations will generate log files.

High-Availability

The DSR is built on a field proven platform and supports 99.999% availability when deployed in geographically redundant pairs. DSR signaling network elements are configured for geographic redundancy with either site able to support the total required signaling traffic in the event of a loss of the mated site. Geographic redundancy requires the originating network element to support alternate routing in the event the primary route becomes unavailable.

The platform supports fully redundant and isolated power architecture. Refer to the Platform Feature Guidefor more information.

Multiple DA MPs are supported in an active-active configuration up to a maximum of 32 DA MPs per DSR signaling node.

The IP layer from the MP to the customer network interface is fully redundant. Enclosure switches and aggregation switches are deployed in redundant pairs. Refer to the Platform Feature Guide for more information on the networking components of the platform.

The DSR factors in the availability of Diameter peers when routing. It maintains the status of each peer. If a peer is not available, the traffic destined to that peer is redistributed to other peers, if available, that provide the same application. The DSR also supports the unique ability to choose alternate routes based on Answer responses. Refer to the Routing and Load Balancing section of this document for more information.

The DSR maintains the status of the connection (SCTP association or TCP socket) and application of each peer. Transport status considers connection status and congestion level. Application status is determined via standard Diameter heartbeat mechanisms.

Capacity and Performance

Capacity and performance values are specific to the platform hardware on which DSR is deployed. Please refer to the DSR Planning Guide for details on capacity and performance.

Appendix A: Supported Diameter Interfaces

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

Table 14: Supported Diameter Interfaces by DSR

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	[RFC5866]
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	14-16777215	Unassigned	
16	16777216	3GPP Cx	3GPP TS 29.228 / 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	3GPP TS 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 283 034]
32	16777232	Ericsson Charging-CIP	_
33	16777233	Ericsson Mm	_
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)	http://www.3gpp.org/ftp/Specs/
47	16777247	3GPP2 CAN Access Authentication and Authorization	ftp://ftp.3gpp2.org/TSGX/Projects
48	16777248	3GPP2 WLAN Interworking Access Authentication and Authorization	ftp://ftp.3gpp2.org/TSGX/Projects
49	16777249	3GPP2 WLAN Interworking Accounting	ftp://ftp.3gpp2.org/TSGX/Projects
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	_
58	16777258	ITU-T Rt	ITU-T Rec. Q.3305.1
59	16777259	CARA	_
60	16777260	CAMA	_
61	16777261	Femtocell extension to Diameter EAP Application	_
62	16777262	ITU-T Ru	[ITU-T Rec. Q.nacp.Ru Q.nacp.Ru]
63	16777263	ITU-T Ng	[ITU-T Rec. Q.nacp.Ru Q.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
75	16777275	ADMI Notification Application	_
76	16777276	ADMI Messaging Interface Application	_
77	16777277	Peter-Service VSI	_
78	16777278	ETSI Rr request model	[ETSI TS 183 071]
79	16777279	ETSI Rr delegated model	[ETSI TS 183 071]
80	16777280	WIMAX HRPD Interworking	3GPP2 X.S0058-0 v1.0
81	16777281	WiMAX Network Access Authentication and Authorization Diameter Application (WNAAADA)	WiMAX Release 1.5
82	16777282	WiMAX Network Accounting Diameter Application (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

	16777286	WiMAX-Location-Authentication-Authorization Diameter Application (WLAADA)	
86			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 Diameter Application (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 Diameter Application (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 Diameter Application	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) and MLM-PE(C)	ITU-T Q5/Sg11
107	16777307	ITU-T M13	ITU-T Q.3230
108	16777308	3GPP S7a	3GPP TS 29.272
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, NSE	_
129	16777329	Rivada Xd	_
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 Application	_
142	16777342	3GPP Np	3GPP TS 29.217
143	16777343	Sandvine Location Relay Service	_
144	16777344	Sandvine Fairshare Traffic Management 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

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

Integrated Cloud Applications & Platform Services