ORACLE
COMMUNICATIONS

Oracle Communications
Diameter Signaling Router

Release 8.2 Virtual Network Function Whitepaper

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Table of Contents

Table of Figures

Table of Tables

Introduction

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. Figure 1 captures a typical LTE/IMS network and the role of DSR in the signaling path.

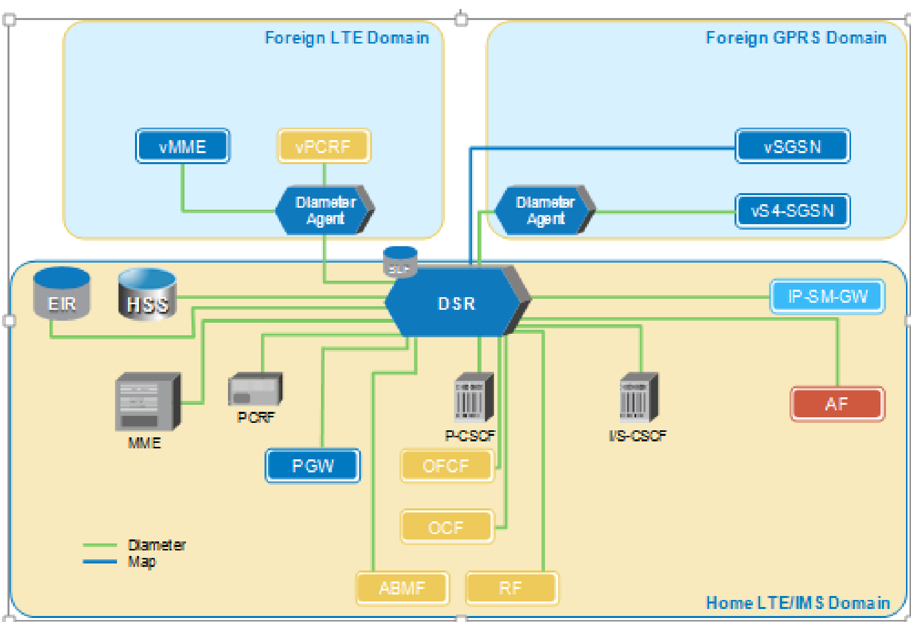

Figure 1: An Example LTE and IMS Network with DSR

Acronyms

Acronym	Description
COTS	Commercial Off The Shelf
CPU	Central Processing Unit
DA	Diameter Agent
DP	Database Processor
DSR	Diameter Signaling Routing
HDD	Hard Disk Drive
IDIH	Integrated Diameter Intelligence Hub
IPFE	IP Front End
KPI	Key Performance Indicator
MP	Message Processor
MPS	Messages Per Second
NIC	Network Interface Card
NOAM	Network Operations, Alarms, Measurements
NE	Network Element
OCDSR	Oracle Communications Diameter Signaling Router
P-DRA	Policy DIAMETER Routing Agent
RAM	Random Access Memory
SBR	Session Binding Repository
SBR(b)	SBR – subscriber binding database
SBR(s)	SBR – session database
SOAM	System (nodal) Operations, Alarms, Measurements
vSTP	virtual Signaling Transfer Point
WAN	Wide Area Network

DSR as a Virtualized Network Function

The main purpose of the OCDSR is to perform 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 point. The OCDSR system consists of the following components as defined in Figure 2.

- Operations, Administration, Maintenance and Provisioning (OAMP)
  - Network OAMP
  - System OAM per signaling node
- 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)
- Data Processor — System OAM per signaling node (DP SOAM)

- Query Server (QS)
- Integrated Diameter Intelligence Hub (IDIH)

Each component plays a key role, the OAM and DA MP Components are the mandatory components of the system. Detailed information pertaining to these OCDSR components is available in the OCDSR user Guide.

Details of the hypervisor and Infrastructure managers tested by OCDSR at a minimum is provided below

- Hypervisor/Virtual Infrastructure Managers
  - QEMU-KVM Version: QEMU 1.5.3, libvirt 1.2.8, API QEMU 1.2.8/Oracle OpenStack (based on Newton Release)
  - VMWare ESXi version 6.0 /VMware vSphere

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 defined subset of diameter peers (defined as an element to which the DSR has a direct transport connection).

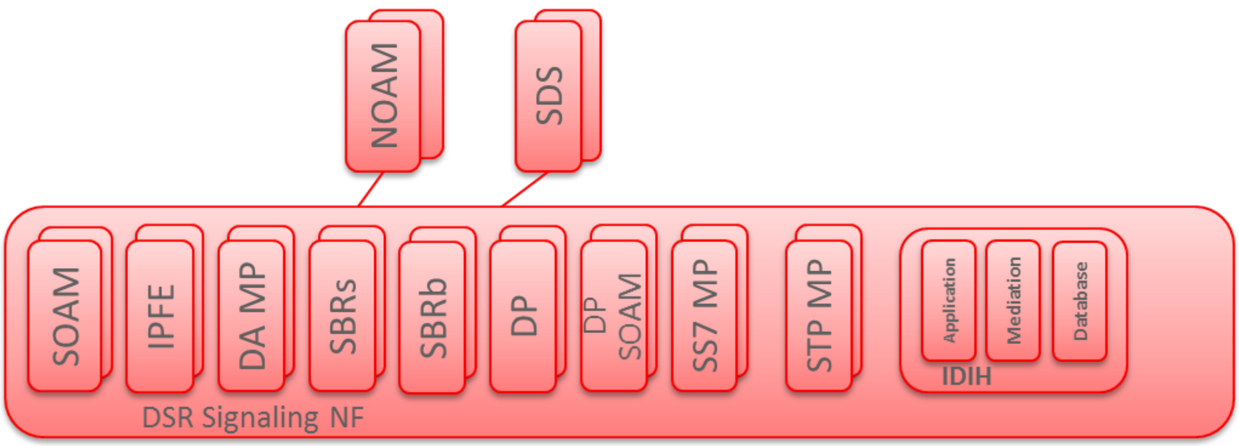

Figure 2: OCDSR Virtualized Components

OCDSR VNF Component Profiles

Each Oracle Communication DSR Virtual Network Function Components (VNFC) requires a profile definition. The VM profile is defined by vCPUs, Memory, and HDD allocated to each component as depicted in Figure 3. OCDSR VNFC profiles are selected because of following.

- Profiles provide optimal performance and capacity for typical deployments.
- VM profile are used for Benchmark purposes.

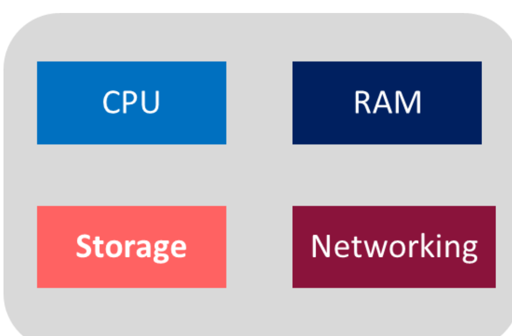

Figure 3: VNFC Profile

The profile definition is independent of hypervisor and the performance benchmarking of the profile is independent of hardware-specific optimizations. Table 1 provides information on the OCDSR VNFC profile.

DSR VNFC profile definition information is listed below:

- Profile specifies supported component VM definitions
- Profiles define CPU, Memory, Hard Disk sizing, and Networking
- Profile definitions are hypervisor independent. Profile performance benchmarking assumes no hardware-specific enhancements.
- Profile assumes host software network configuration
    - Support for OAM, Signaling and Local interfaces

Table 1: DSR VM Profiles

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Interface	Redundancy Models	Affinity/Placement Rules (Per Site)	Notes
DSR NOAM	4	6	70	1 Pair	XMI & IMI	Active/Standby	2 VMs per DSR network in any site. VMs to be deployed on separate servers	
DSR SOAM	4	6	70	1 Pair per DSR NF	XMI & IMI	Active/Standby or Active/Standby/Spare	2 VMs per site. VMs to be deployed on separate servers.	Redundancy model Active/Standby/Spare model is used for PCA mated-pair deployments. For all other deployments Active/Standby model is used.
DA MP	12	16	70	16 per DSR NF	XMI, IMI & XSI	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss	The limit of 16 is the combined total of DA-MPs and DA-MPs with IWF. The vSTP MPs and SS7 MPs don't count against this 16.
DA MP w/IWF	12	24	70	16 per DSR NF	XMI, IMI & XSI	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss	The 24 GB RAM requirement is a hard minimum if the DA-MP VM will be used with the IWF application.
vSTP MP	8	6	70	10 per DSR NF	XMI, IMI & XSI	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss	The vSTP MPs don't count against the 16 DA-MP limit in a single OCDSR node, so a DSR can have up to 16 DA-MPs and up to 10 vSTP MPs.
IPFE	6	16	70	2 pairs per DSR NF	XMI, IMI & XSI	Active/Standby	Each VM in a pair must be deployed on separate server.	Deployed in pairs. Max 2 pairs (4 VMs).
SS7 MP	12	24	70	8 per DSR NF	XMI, IMI & XSI	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss	
SBR(s)	12	25	70	8 Server Groups per SBR(s)	XMI & IMI	Active/Standby/Spare	Active/Standby VMs to be deployed on separate servers, Spare is typically at another geographic location for geo-redundancy.	Can be either Active/Standby/Spare or Active/Standby depending on customer geo-redundancy requirements.
SBR(b)	12	32	70	8 Server Groups per SBR(b)	XMI & IMI	Active/Standby/Spare	Active/Standby VMs to be deployed on separate servers, Spare is typically at another geographic location for geo-redundancy.	Can be either Active/Standby/Spare or Active/Standby depending on customer geo-redundancy requirements.
DP SOAM	4	12	70	1Pair per DSR NF	XMI & IMI	Active/Standby	2 VMs per site. VMs to be deployed on separate servers.	
DP	6	10	70	10 per DSR NF	XMI & IMI	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss	To be evenly distributed across servers to minimize capacity loss
SDS NOAM	4	32	300	1 Pair per Network	XMI & IMI	Active/Standby	Anti-affinity between the Active/Standby VMs.	Active/Standby. An optional "Disaster Recovery" SDS is supported that would typically be located at a different data center to provide geo-redundancy.
Query Server	4	32	300	1 per SDS NOAM	XMI & IMI	N/A since non-redundant	Non, non-redundant	Optional 1 per site. Can have one for the primary SDS-NOAM and one for the Disaster Recovery SDS-NOAM
iDIH Application	4	8	64	1 per Site	XMI & IMI	N/A since non-redundant	None, non-redundant	Optional component for DIAMETER traffic monitoring
iDIH Mediation	4	8	64	1 per Site	XMI & IMI	N/A since non-redundant	None, non-redundant	Optional component for DIAMETER traffic monitoring
iDIH Database	4	8	200	1 per Site	XMI & IMI	N/A since non-redundant	None, non-redundant	Optional component for DIAMETER traffic monitoring

Virtualized OCDSR can be deployed in High Availability (HA) and Geo-redundant configurations to provide operators with necessary resiliency to support signaling traffic. In a virtualized configuration, placement of VMs is an important consideration to enable higher availability. Information pertaining to VM placement (affinity rules), redundancy scheme per VNFC and benchmark information pertaining to latest DSR release 8.2 is captured in Table 1.

OCDSR Performance and Capacity

OCDSR component capacity is depends on CPU, memory and storage of a bare metal server or a virtual machine. OCDSR VNFCs performance and capacity is characterized by Messages Per Message (MPS) for non-data base nodes whereas database node performances are based on number of entries. Nodal capacity information pertaining to the OCDSR VMs are provided in the Appendix.

OCDSR and HEAT Template

Heat is the main project in the OpenStack Orchestration program. It implements an orchestration engine to launch multiple composite cloud applications based on templates. A Heat template allows a user to specify an application in more abstract terms. Heat then translates this specification into the same direct API calls, and executes them, automatically. OCDSR VNFCs can be described using the HEAT template with information pertaining to resources, flavor, initial configuration, etc.

What is a HEAT template? Following points may help:

- Heat templates are text files. They use the YAML markup language to describe an application deployment.
- Heat templates are re-usable. They can be generalized, so that the same base templates can be used to describe many deployments of an application.

- Heat templates are modular. For a complex application composed of multiple sub-components, small Heat templates can be created to describe the sub-components.
- Heat templates can invoke other Heat templates. A master template can invoke component templates to create a complex application as provided in Figure 4.

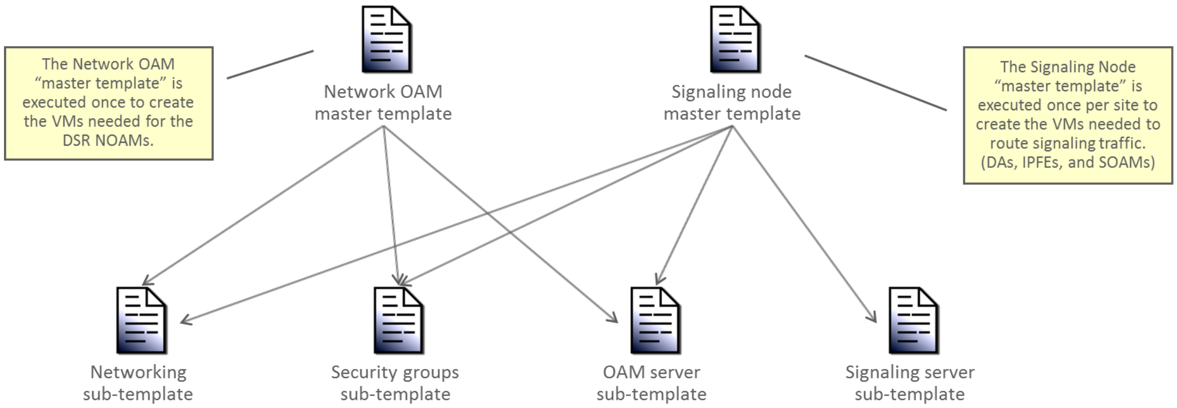

Figure 4: Bare Metal High Availability

Summary

The OCDSR solution can be deployed in virtualized environments using industry leading hypervisors such as KVM and VMWare, as well as bare metal servers.

Oracle tests and validates the performance of the defined OCDSR VNFC profiles on combinations of hypervisors and HP Gen 8, HP Gen 9, and Oracle X5-2 and X7-2 Servers. In addition to this service providers can run OCPM on other hardware platforms using the specified KVM and VMWare hypervisors as long as the underlying resources are available to support OCDSR VNF component profiles. In this case the service provider will have to validation the performance of the virtualized system in their labs and production environments or can optionally use the Oracle Consulting Services team.

The service provider can also decide to run different profile sizes (vCPUs, memory, HDD) than has been verified by Oracle with the specified KVM, VMWare, and OVM hypervisors. In this case the service provider will have to validation performance and capacity of these profiles in their lab and/or production environments or can optionally use the Oracle Consulting Services team.

Appendix A. DSR Cloud-Based Nodal Capacity

Table 2 provides the DSR cloud-based nodal capacity for KVM and VMWare hypervisors based on the latest test results.

Table 2: Redundancy Model Decision Making

	DSR Function	Optional Capacity			
		DSR 7.1.0	DSR 7.1.1	DSR 7.2.0 - 7.4	DSR 8.x[1]
Nodal	Relay MPS	50,000	50,000	50,000	250,000
	Database MPS				250,000
	Session Stateful MPS				250,000
	Diameter Connections per DSR	8,000	8,000	8,000	8,000
	Peer Routing Rules	10,000	10,000	10,000	10,000
	Range-Based Routing Rules	1M	1M	1M	1M
	Max IPFE client Connections	8,000	8,000	8,000	8,000
Per IPFE	IPFE Bandwidth (Gb/s)	1.0	1.0	1.0	2.5
Network	Address Resolution Subscriber Database Entries				300M
	Max subscriber binds per network				128M
	Max concurrent sessions per network				256M

[1] Capacities reflect configurations in the DSR Benchmarking Guide. Benchmarked OVM capacities are lower. See the DSR Benchmarking Guide

ORACLE

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

Connect with us:

- blogs.oracle.com/oracle
- facebook.com/oracle
- twitter.com/oracle
- oracle.com

Integrated Cloud Applications & Platform Services

E93245-01, Oracle Communications Diameter Signaling Router Planning Guide. Release 8.2

May 2018

Oracle is committed to developing practices and products that help protect the environment

Integrated Cloud Applications & Platform Services