ORACLE®

COMMUNICATIONS

Oracle Communications
Diameter Signaling Router

Release 8.2 Planning Guide

ORACLE®

Table of Contents

Table of Figures

Table of Tables

Introduction

The Diameter Signaling Router (DSR) creates a Diameter signaling core that relieves LTE and IMS endpoints of routing, traffic management and load balancing tasks and provides a single interconnect point to other networks. The resulting architecture is below. Refer to the DSR Feature Guide for currently supported functionality.
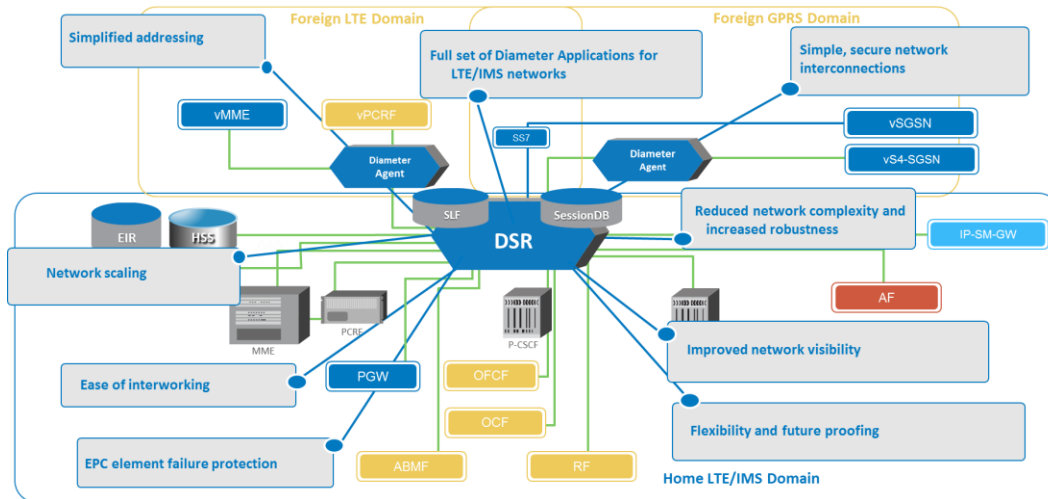
Figure 1: DSR in the Diameter Network

Acronyms

Acronym	Description
DA	Diameter Agent
DIH	Diameter Intelligence Hub
DP	Database Processor
DSR	Diameter Signaling Routing
FV RMS	Fully Virtualized Rack Mount Server
IDIH	Integrated Diameter Intelligence Hub
IMR	Ingress Message Rate
IPFE	IP Front End
MP	Message Processor
MPS	Messages Per Second
N-IDIH	Network Integrated Diameter Intelligence Hub
NOAM	Network Operations, Alarms, Measurements
NE	Network Element
NFV	Network Function Virtualization
OTN	Oracle Technology Network
PMAC	Platform, Management, and Control
RMS	Rack Mount Server
SBR	Session Binding Repository
SBR(b)	SBR – subscriber binding database

Acronym	Description
SBR(s)	SBR – session database
SBR(u)	SBR-Universal application database
SOAM	System (nodal) Operations, Alarms, Measurements
SS7	Signaling System 7
TOR	Top-Of-Rack (Switches)
VM	Virtual Machine

Terminology

Term	Description
Co-Located	Two or more products in the same cabinet.
Co-Mingled	Two or more products in the same enclosure.
Site	A specific geographic location where DSR equipment is installed.
Geo-Diverse	Refers to DSR equipment located at geographically separated sites
Geo-Redundant	A node at a geo-diverse location which can assume the processing load for another node(s)
1+1 Redundancy	For every 1, an additional 1 is needed to support redundant capacity. The specific redundancy scheme is not inferred (e.g. active-active, active-standby).
N+K Redundancy	For every N, an additional K is needed to support redundant capacity. The specific redundancy scheme is not inferred (e.g. active-active, active-standby).
Node	A DSR node is either a DSR signaling node, an NOAM node or an SDS node.
Ingress Message Rate	A measure of the total Diameter messages per second ingressing the DSR. For this measure, a message is defined as any Diameter message that DSR reads from a Diameter peer connection independent of how the message is processed by the DSR.
Messages Per Second	A measure of the DSR Diameter message processing volume in messages per second. For this measure, a message is defined as: • DSR processing of an ingress Diameter message and either transmitting a single outgoing Diameter message or discarding the ingress message. The outgoing message may be a variant of, or a response to, the ingress message. • DSR transmission of any Diameter message, as required by DSR configuration, that is associated with incremental actions/events associated with #1 above. For example, the re-routing of a Request upon connection failure or the copying of a Request. Messages excluded from this measure are: • Diameter peer-to-peer messages: CER/CEA, DWR/DWA, and DPR/DPA • Ingress Diameter messages discarded by the DSR due to Overload controls • Answers received in response to Message Copy
X5-2	Refers to any type of X5-2 server: AC/DC and NEBS compliant Netra X5-2 or the AC ORACLE X5-2

References

[1] DSR Alarms and KPIs Reference – Available at Oracle.com on the Oracle Help Center (OHC)

[2] Diameter Measurements Reference – Available at Oracle.com on the Oracle Help Center (OHC)

[3] Platform Feature Guide – Available upon request

[4]   DSR Feature Guide – Available at Oracle.com on the Oracle Help Center (OHC)

[5]   DSR Cloud Benchmarking Guide – Available at Oracle.com on the Oracle Help Center (OHC)

[6]   VM Placement and CPU Socket Pinning Tool – Available at Oracle.com on the Oracle Help Center (OHC)

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

DSR Platform

Historically, the DSR is sold and supported as an engineered system. That is, the software, exact hardware, precise configuration of both the hardware and software are specified and tuned for optimal performance, capacity and reliability. The telecommunications market is moving rapidly and many operators are deploying telecom applications on a common and shared infrastructure with many other telecom applications. This effort, along with the automation of application lifecycle is referred to as Network Functions Virtualization (NFV).

Oracle Communications Diameter Signaling Router (OCDSR or DSR) is deployed on a number of platforms. The DSR has a number of deployment scenarios:

- Bare-metal is the original deployment configuration of the DSR. It scales to very high performance and is widely deployed.

- Fully virtualized RMS (FV RMS) provides virtualization of the DSR, but does not use a cloud manager, and does not co-reside with other applications. Provides compact, cost-effective footprint and is widely deployed.

- Cloud deployable provides full virtualization, assumes the DSR resources are managed by a COTS infrastructure manager, and that the DSR can be one of many applications in the cloud.

The bare-metal deployments are supported over HP c-Class while the FV RMS options are available on rack-mount servers (HP or Oracle). The cloud deployable DSR is a DSR that is ready and able to be deployed into a number of different cloud environments, including but not limited to:

- A customer provided shared resource infrastructure. The DSR is simply one of many applications.

- A dedicated private cloud. The DSR may be the only application, or one of a small set of applications. Services and infrastructure may also be provided by Oracle and deployed at customer's sites. Often (but not necessarily) this is a deployment tuned specifically for the DSR.

- A hosted cloud. The DSR is deployed in an Oracle or operator hosting cloud, and end-customers rent or lease the DSR application from the hosting provider.

The DSR is capable of running on a huge variety of infrastructures, but not all infrastructures are the same and performance, capacity, and latency can vary dramatically based on the chosen infrastructure and how it is deployed. In general, the DSR assumes a high bandwidth, low-latency, high processing power environment (carrier grade cloud). A DSR mated pair must be deployed using the same platform type – for example both sites deployed as fully virtualized RMS. Across a network, pairs can be mixed. As an example one pair may be bare-metal and another pair in the same network may be deployed as fully virtualized RMS. In this scenario, every pair runs at its own engineered capacity.

The DSR Cloud Benchmarking Guide – Available on the OTN, includes additional details about the cloud deployment model and provides details of the infrastructure used for benchmark testing, including the hardware and software. It also describes key settings and attributes, and some recommendations on configuration.

The remainder of this document focuses on bare metal and fully virtualized RMS but does not reflect requirements for a cloud deployable DSR. This document is focused on Release 8.2. Earlier versions of this document applicable to earlier releases can be found on the OTN. Please refer to the DSR Cloud Benchmarking guide for cloud deployments which is also available on the OTN.

Key Components

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). DSR requires three tiered OAM architecture with a centralized Network OAM as shown below. The Network OAM can be virtualized onto the System OAM blades at one of the signaling sites (for small systems) or deployed as standalone rack mount servers either co-located or

separate from the signaling sites. Please note that the server icons shown in the diagram below represent either a physical blade or a virtual machine depending on the deployment scenario.
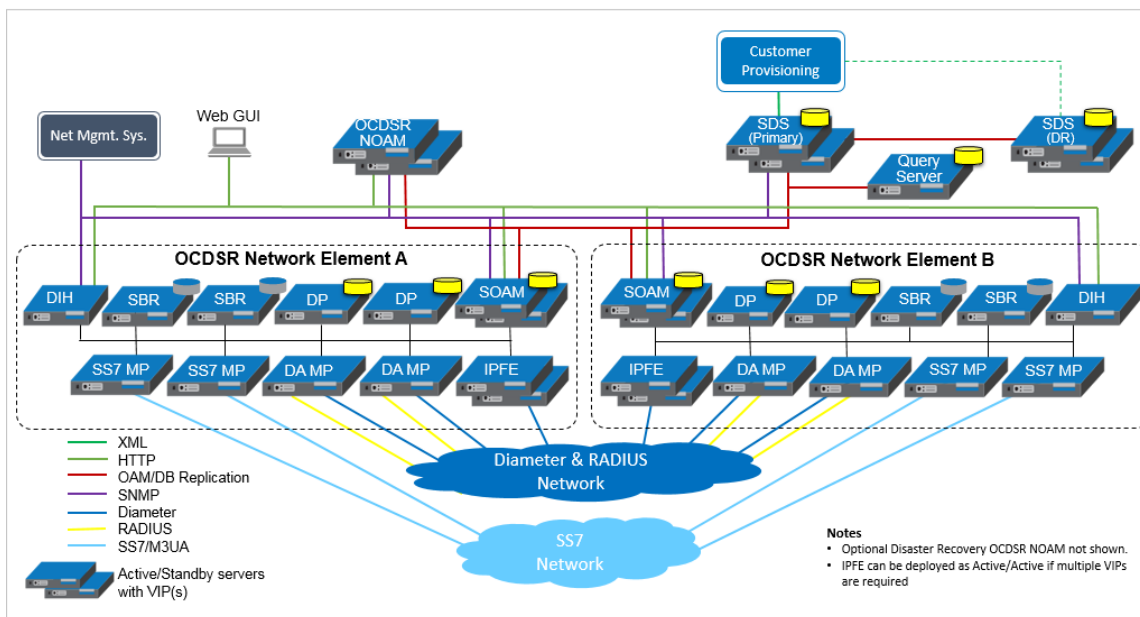
Figure 2: DSR System Architecture

A DSR network consists of:

- A Network OAM network element that can be co-located with the database provisioning node or signaling node, and may optionally include a disaster recovery node

- At least one signaling node, and may optionally include up to 31 additional nodes

- A single network may contain up to 32 DSR signaling nodes

- An optional database provisioning node (SDS), and may optionally a include disaster recovery node

- One database provisioning system (SDS nodes) normally serves all the DSRs in a network.

The key components of the solution are:

- Operations, Alarms and Maintenance (OAM) - Required

  - System OAM per signaling node
  - Network OAM
- Diameter Agent Message Processor (DA-MP) – Required

  - Also includes support for RADIUS signaling and RADIUS IWF
- Management Server (PMAC) - Required

- IP Front End (IPFE) - Optional

- Integrated Diameter Intelligence Hub (IDIH) – Optional

- Diameter Intelligence Hub (DIH) – Existing deployments only

- Subscriber Database Server (SDS) - Optional

- Query Server (QS) - Optional

- Database Processor (DP) - Optional

- Session Binding Repository (SBR) – Optional

- SS7 Message Processor (SS7 MP) – Optional

- Virtual Signal Transfer Point (vSTP MP) - Optional

Operations, Alarms and Maintenance

The Operations, Administration and Maintenance components of the DSR include the System OAM located at each signaling node and the Network OAM (NOAM). DSR requires three tiered OAM topology.

Network OAM 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, SBR are all "MPs" managed by NOAM
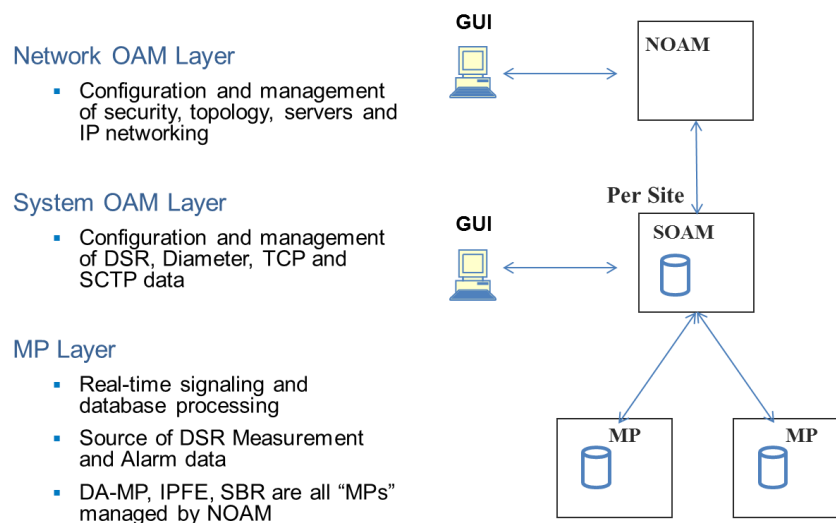
Figure 3: Three Tier OAM Topology

System OAM

The System OAM supports the following:

- System OAM which is the primary DSR OAM application supporting the System OAM, Diameter Agent, IP Front End, SS7 MP and the Session Binding Repositories

- DP SOAM provides the nodal management and distribution of the DPs and subscriber database. This is an optional functional component and is only present when DPs are used.

System OAM applications have certain times when bursts of data are transferred to or from the DSR. For example: transferring measurements data off the System OAM, or transferring the subscriber database to the DP OAM from the SDS during initialization of the system. The applications are designed to handle this type of bursty transfer.

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

- centralized OAM interface for the node

- maintains local copy of the configuration database

- supports SNMP northbound interface to operations support systems for fault management

- maintains event and security logs

- centralizes collection and access to measurements and reports

Network OAM

Key characteristics of the Network OAM are as follows:

- centralized OAM for the DSR network

- supports SNMP northbound interface to operations support systems for fault management

- optionally supports Disaster Recovery site for geographic redundancy

- provides configuration and management of topology data

- provides mechanism to create user groups with various access levels

- maintains event and security logs

- centralizes collection and access to measurements and reports

Platform, Management and Control (PMAC)

Provides the management of the overall infrastructure.

Key characteristics of the management server are as follows:

- not a signaling critical component of the DSR

- deployed as a single server per site (may be redundant) or as a VM on one of the RMS server configurations

Diameter Agent Message Processor (DA-MP)

The Diameter Agent message processor (DA-MP) provides the primary signaling processing and routing capabilities of the DSR. DA-MPs are a critical component of any DSR signaling node. DA-MP processor capabilities are primarily bound by the processing speed of the blade, read/write speeds on I/O ports of the blade and with increased message size as described below, RAM utilization.

DSR supports Diameter message sizes up to 60K bytes. Normal message size is defined as messages up to 16K bytes and large messages are defined as messages between 16K and 60K bytes. Typically, less than 1% of the messages traversing the DSR are large messages. The percentage of large message sizes is likely to grow in some networks and varies from network to network. For these reasons, engineering the system to at least 5% mix of large messages provides support for growth of message sizes in the networks for the foreseeable future.

As the average message size grows in the network, so do the bandwidth requirements. Engineering/planning assumes an average 2K message size, the bandwidth available to the Diameter Agent may limit the actual average message size that can be supported in a given configuration.

Key characteristics of a DA-MP are as follows:

- provides application specific handling of real-time Diameter and RADIUS messages and Diameter-RADIUS interworking

- accesses DPs for real-time version of the subscriber DB, as needed

- accesses SBRs for real-time session and binding data, as needed

- interfaces with System OAM

- interfaces with SS7 MP for MAP-Diameter IWF

IP Front End (IPFE)

The DSR IP Front End provides TCP/SCTP connection based load balancing that makes a large DSR accessible to incoming connections through a minimal number of IP addresses.

Key characteristics of an IPFE are as follows:

- optional but highly recommended component of the DSR

- supports up to two active / standby pairs each supporting 3.2 Gbps bandwidth or 1Gbps for fully virtualized deployments

- supports TCP, uni-homed SCTP and multi-homed SCTP

Integrated Diameter Intelligence Hub (IDIH)

The IDIH is an optional component of the DSR and provides advanced troubleshooting for Diameter traffic handled by the DSR.

Key characteristics of the IDIH are as follows:

- optional but highly recommended component of the DSR

- integrates the configuration of trace filters directly into the DSR

- trace filters can be applied against all messages processed by DSR

- transaction data records containing both Diameter Requests and Answers are captured for additional analysis and filtering

- ladder diagrams facilitate troubleshooting concerning message sequence and peer connection issues

- supports detailed decoding of AVPs associated with Diameter protocols

Diameter Intelligence Hub (DIH)

The DIH is supported for existing deployments only. The DIH supports advanced troubleshooting for Diameter traffic handled by the DSR.

Subscriber Database 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 OAMs 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

- provides web based GUI for provisioning, configuration and administration of the data

- supports SNMP 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 or deployed as VMs in fully virtualized deployments

- optionally supports Disaster Recovery site for geographic redundancy

Query Server (QS)

The Query Server contains a replicated copy of the local SDS database and supports a northbound SQL interface for free-form verification queries of the SDS Provisioning Database. The Query Server's northbound SQL 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 SQL access for offline queries to the SDS Provisioning Database

Database Processor (DP)

The DP is the repository of provisioned address resolution data on the individual DSR node elements. The DP hosts the full address resolution database and scales by adding blades/servers/VMs. The database processor provides an in-memory database for extremely fast and reliable access to provisioned information. DPs are an optional component of a DSR signaling node, and are needed when an application using a subscriber database is utilized. Following are the primary resources being utilized on the DP blade:

- Memory – bounds the size of the in-memory database on each DP

- I/O & processor – bounds the number of database accesses that can occur

Key characteristics of a DP are as follows:

- provides high capacity real-time database query capability to DA-MPs

- interfaces with DP SOAM for provisioning of subscriber data and for measurements reporting across all DPs

Session 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)). Other stateful applications such as the Security Application and the Zero Balance offload applications use another variation of the SBR called the SBR(u). 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.

- SBR(b) - A subscriber binding database – which maps subscribers to specific PCRFs (or other element)
- SBR(s) - A session database – which maintains the state of policy and/or charging related Diameter sessions
- SBR(u) – A database that enables Stateful applications and stores application specific data.

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

SS7 Message Processor (SS7 MP)

The SS7 MP provides the MAP 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

vSTP Message Processor (vSTP MP)

The vSTP MP brings the SS7 signal transfer point functionality to the DSR.

Key characteristics of a vSTP MP are as follows:

- Service providers can optimize the use of network resources, manage subscribers and also interoperate between networks with disparate technologies
- vSTP functionality supported for Cloud deployment models only and not available on Bare-metal or FV RMS
- 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

DSR Deployment and Scalability Options

This document focuses on requirements for bare metal, hybrid, and hardened/engineered virtualization. Please refer to DSR Cloud Benchmarking Guide – for all information regarding cloud deployable DSR.

General Node Topics

Messages Per Second (MPS)

The DSR's main traffic processing attribute is messages per second (MPS). Please refer to DSR Licensing Information User Manual – Available on Oracle Help Center (OHC) for details on the definitions and examples for calculating the MPS correctly.

Network Architecture

Customers have differing requirements for how to implement DSR signaling nodes into their networks. These requirements include support for both layer 2 and layer 3 demarcations and support for physical separation of the different types of traffic. The following diagram shows the network architecture options for DSR signaling nodes. Green indicates the default shipping baseline.
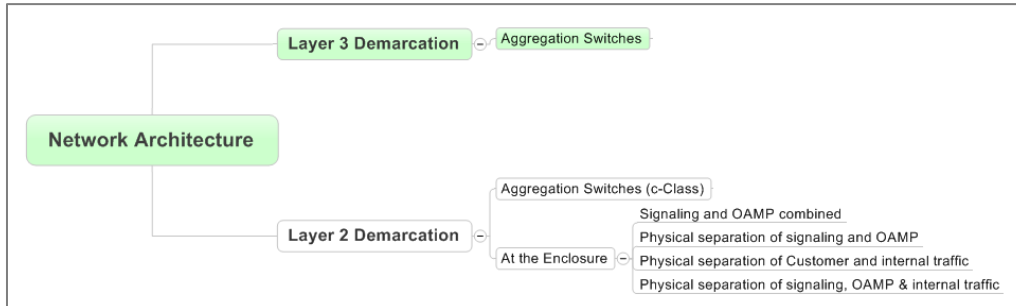
Figure 4: Bare Metal Network Architecture

Figure 5: RMS Network Architecture

Note: For systems with up to 50K MPS, 1G aggregation switch is supported. For systems over 50K MPS, the 1G aggregation switch is not supported and 10G aggregation switch is required and must be sourced by the customer.

These options result in multiple network topologies which are supported based on customer needs. Topologies differ from each other in number of physical networks supported and whether aggregation switches are present. Each Topology can have multiple variations based on whether 1G or 10G connectivity is required and whether a backup network is needed.

Table 1: DSR Bare Metal Signaling Node Network Topologies

Topology	Equipment Needed	Purpose
1, Layer 3	One pair of enclosure switches per enclosure. One pair of aggregation switches	Provide aggregated uplinks to the customer network with Layer-3 demarcation.
1, Layer 2	One pair of enclosure switches per enclosure. One pair of aggregation switches	Provide aggregated uplinks to the customer network with Layer-2 demarcation.
3	One pair of enclosure switches per enclosure.	Provide integration with Customer L2 domain where aggregation switches are not used.
4	Two pairs of enclosure switches per enclosure.	Provide physical separation of traffic between OAM and Signaling.
5	Three pairs of enclosure switches per enclosure.	Provide physical separation of traffic between OAM and two separate Signaling networks.
6	Three pairs of enclosure switches per enclosure.	Provide physical separation of inter-enclosure traffic from customer's OAM and Signaling networks.

Table 2: DSR Fully Virtualized Rack Mount Server (FV RMS) Signaling Node Network Topologies

Equipment Needed	Purpose	FV RMS Server Platform
One pair of aggregation switches	Provide aggregated uplinks to the customer network with Layer-3 demarcation.	HP Gen8, HP Gen8 v2
One pair of aggregation switches	Provide aggregated uplinks to the customer network with Layer-2 demarcation.	HP Gen8, HP Gen8 v2
N/A	Provide physical separation of traffic between OAM and Signaling. Provide direct connection from RMS to customer network(s).	HP Gen8, HP Gen9, Oracle X5-2, ORACLE X6-2, Netra X5-2, ORACLE X7-2

Bandwidth

Bare Metal Signaling Node

Enclosure bandwidth is driven by the traffic sent to and from the blades within a given enclosure. The amount of bandwidth needed in certain situations drives the size and number of switches required.

- DSR supports 1G enclosure switches. This implies a 1G connection from the blade to the enclosure switch.

- DSR supports 10 G enclosure switches. This implies a 10G connection from the blade to the enclosure switch.

- DSR supports 10 G fiber connection from the enclosure to the customer's switches, to Cisco 4948E-F, or to another enclosure switch when 10G enclosure switches are used.

- DSR supports 4 x 1 G connections from the enclosure to the aggregation switches, customer's switches or to another enclosure switch when 10G or 1G enclosure switches are used.

When multiple physical networks are used, and therefore multiple sets of enclosure switches are used, then each physical network may use 1G or 10G enclosure switches depending on the bandwidth requirements of the specific deployment.

FV RMS Signaling Node

The DSR rack mount servers support eight 1G connections from the server to the aggregation switches and/or customer switches for FV RMS up to 50K RMS (HP Gen8 and HP Gen8 v2 only).

The Fully Virtualized Rack Mount servers support up to eight 10G connections from the server to the aggregation switches and/or customer switches for FV RMS greater than 250K MPS (HP Gen9, X5-2,X6-2 and X7-2 servers).

High Availability

High availability is an expected attribute of any core signaling product. DSR is no exception. A typical DSR node is engineered with high levels of redundancy to ensure a very high availability of the system at all times. A DSR site supports redundant power. Below is the general configuration for Bare Metal high availability networks.
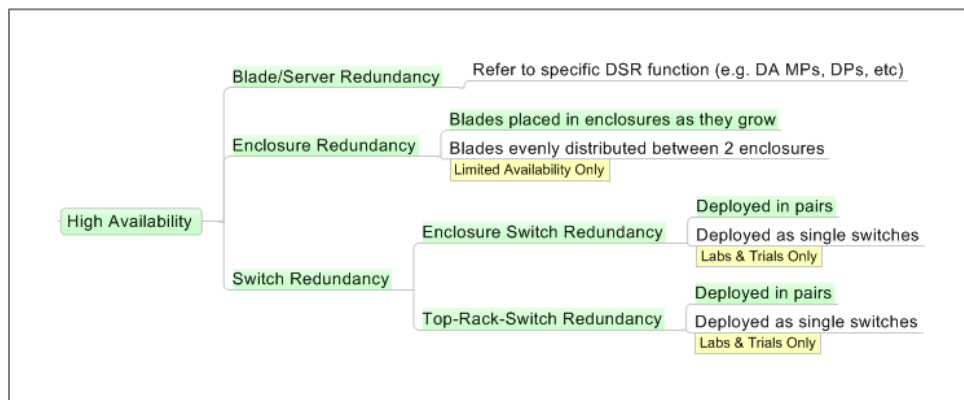
Figure 6: Bare Metal High Availability

When redundant enclosures are used, SOAM, DA-MP, IPFE, SBR, SS7 MP, and DP blades are distributed across enclosures to maximize availability in the case of an enclosure failure.

The management server is not a message processing critical component of the DSR, and is deployed as a single server per site or as a VM on one of the existing servers in the fully virtualized deployment model. In cases where DSR is co-mingled with another product that uses a redundant management server, the DSR supports the redundant management server.

The following diagram shows the configuration for FV RMS high availability networks.

Figure 7: FV RMS High Availability

The Diameter Intelligent Hub and Integrated Diameter Intelligence Hub are not message processing critical component of the DSR, and are deployed as a single server or blade (existing configurations only) per node.

Site Redundancy

The DSR can be deployed in mission critical core routing of networks, or in less critical areas of the network such as off-line billing. Since there are a variety of needs for high system availability, the DSR can be deployed with or without geo-redundancy.

Figure 8: Site Redundancy

The DSR can be deployed as a single node at a single site with and without local high-availability or with geo-redundancy. DSR also supports a geo-redundant model that allows for a mated-triple. In this model, any one of the three nodes forming the mated triple could take all the traffic of the three nodes. Although mated triplets are supported, they are not normally recommended due to the cost, added complexity of configuration and operational maintenance of an extra signaling node.

When HA is not deployed in both DSRs in a mated pair, the target availability may not be achieved. When a single DSR is used to provide geo-redundancy for multiple other DSR nodes that have diverse configurations (i.e. DSR is not deployed in a mated pair or mated-triplet configuration), the following restrictions may apply:

- Deployment of DSR address resolution applications (RBAR, FABR) requires that either 1) subscriber identities/ranges are unique to each DSR -OR- 2) all address resolution table entries in all DSRs be identical (i.e. all entries can resolve to common destinations).

- Deployment of DSR stateful applications (P-DRA, OC-DRA) may require common state databases shared by all DSRs.

Power

The DSR signaling node supports DC and AC International and AC North America power supplies.

Figure 9: Power

Spares

To maintain a high availability system, spare parts for key components are required.

Figure 10: Spares

Co-Location

The DSR may be co-located with one or more products. Co-location is the ability for the DSR to be physically located in the same site footprint (cabinets) as other products.

Bare Metal Signaling Node
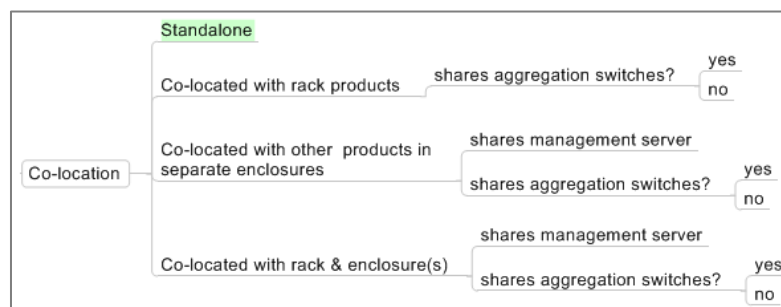
Figure 11: DSR Co-location

Co-location allows for different products at the same site to share certain resources:

- Aggregation switches – when in use
- Management server
- Physical cabinet and power

Here are some examples of co-located products with a DSR:

- DSR in an enclosure + rack mount SDM
- DSR in an enclosure + rack mount SDS (DSR subscriber database provisioning system)
- DSR in an enclosure + 2nd DSR in a second enclosure
- DSR in an enclosure + Policy Management in a second enclosure
- DSR in an enclosure + rack mount PIC PMF

Provided there is physical space, the DSR and one or more additional products are allowed to co-locate within the same physical cabinet. The customer needs to ensure the aggregation switch can support the desired combination of products.

Co-locating products:

- share the same power supplies provided the power supplies are appropriately sized for the full co-located solution
- share the same aggregation switches provided the co-locating solutions are using aggregation switches
- share the same management server provided the co-locating solutions are using a management server

Co-Mingling

The Bare Metal DSR may be co-mingled with one or more products. Co-mingling is the ability for the DSR to be physically located in the same enclosure(s) as other products.

Figure 12: Bare Metal DSR Co-mingling

Provided there is physical space, the DSR and one or more additional products are allowed to co-mingle within the same physical enclosure.

- Co-mingling and co-location are allowed at the same site deployment.

- Co-mingling products are allowed to share the same power supplies provided the power supplies are appropriately sized for the full co-mingled solution.

- Co-mingling products are allowed to share the same aggregation switches provided the co-mingling solutions are using aggregation switches.

- Co-mingling products are allowed to share the same enclosure switches.

- Co-mingling products can use the same management server.

FV RMS Signaling Node

Figure 13: FV RMS Co-location

Co-location allows for different products at the same site to share certain resources:

- Physical cabinet and power

Some examples of co-located products with an FV RMS DSR are:

- FV RMS DSR signaling node + SDS rack mount

- FV RMS DSR signaling node + Policy rack mount

System OAM

The DSR requires that System OAM is present in each DSR node. The platform requires a PMAC server at each site. These two components form the basis of the overall management of the DSR signaling node.

In the rack mount server form factor, the System OAM and PMAC server functions are virtualized onto the appropriate servers with the DA-MP, NOAM (if applicable) and IPFE (when used) functions. DSR also supports simplex configurations for labs and trials.

Redundancy Model

System OAM blades are deployed in active-standby configuration with two blades per node.

DSR uses a non-redundant PMAC rack-mount server. This PMAC server may be shared with other products using Bare Metal at the same site.

Network OAM Network Element

Each DSR network requires at least one Network OAM and optionally a disaster recovery Network OAM.

Network Architecture

The network architecture shown below applies to the cases where the NOAM NE is deployed on rack mount servers. In the cases where the deployment of the Network OAM is co-resident on Bare Metal, the networking for the DSR Signaling node is applicable.
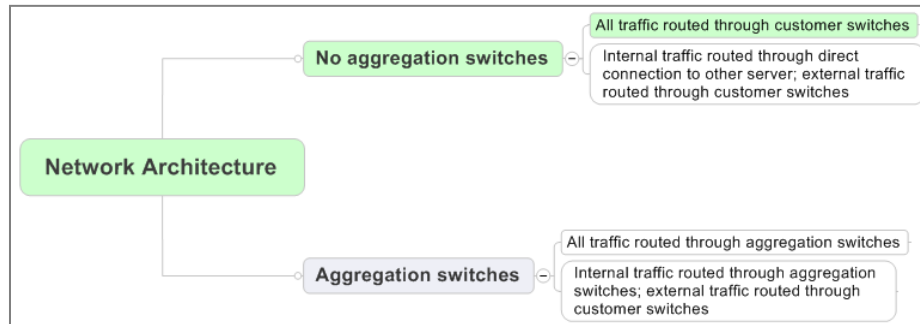
Figure 14: Network OAM Network Architecture

High Availability

The NOAM provides a network OAM and is not considered to be a real-time signaling critical component of a DSR network deployment. However, customers generally require that the NOAM is highly available.

Figure 15: NOAM High Availability

Redundancy Model

The NOAM is not considered to be a real-time critical component of a DSR network deployment. However a disaster recovery node is optionally available to provide site-redundancy in cases of a complete loss of a site. The Network OAM deployment alternatives are as follows:

Bare Metal Configurations

- The NOAM can be deployed on dedicated rack-mount servers either at a standalone site or co-located at one of the DSR signaling sites. If co-located at one of the DSR signaling sites, the active or standby NOAM can be co-resident on the management server located at the DSR signaling site, as long as the IDIH is not hosted by the management server as well.

- The NOAM can also be deployed on the System OAM blades at one of the DSR signaling sites. This option is supported for small systems with no more than two signaling nodes.

- If the system is beyond one pair, then the system may be required to update the NOAM profile to the medium or large size.

FV RMS Configurations

- For details on VM redundancy models please see Planning for Fully Virtualized Rack Mount Servers.

Power

The NOAM can be deployed with either AC or DC power. DC power is supported on HP hardware only.

Co-Location

The rack mount server form factor of the NOAM may be co-located with one or more products. Co-location is the ability for the NOAM to be physically located in the same site footprint (cabinets) as other products.


Figure 16: NOAM Co-location

Co-located products share resources such as cabinet and power. Examples of co-located products with NOAM are:

- NOAM rack mount + DSR signaling node in an enclosure

- NOAM rack mount + SDS rack mount

- NOAM rack mount + Policy rack mount

PMAC

For Bare Metal configurations the PMAC is located on a dedicated server called the management server.

For FV RMS configurations, the PMAC is configured and assigned as a VM.

See the Platform, Management and Control (PMAC) section for information on the PMAC.

Diameter Agent Message Processors

The DSR Diameter Agent message processors (DA-MPs) provide the core routing and processing functionality for the DSR.

The DSR bare metal form factor supports DA-MPs in the following configurations per DSR signaling node:

- Active-Active configuration for scaling up to sixteen total DA-MPs per DSR signaling node.

- An Active-Standby configuration for two DA-MPs per DSR signaling node. (This configuration is supported but not recommended.)

- An Active Only simplex configuration for labs and trials.

SS7 Message Processor (SS7 MP)

The SS7 Message Processor is an optional component of the DSR which is required to support the MAP Diameter interworking function (IWF). MAP Diameter IWF is supported for both Bare Metal and Fully Virtualized rack mount server signaling node configurations.

Below are the redundancy models supported for the SS7 MP.

- Active-Active configuration for scaling up to eight total SS7 MPs per DSR signaling node.

- An Active Only simplex configuration for labs and trials.

Subscriber Database Processor

The DSR subscriber database processor (DP) provides real-time access to a subscriber database. This is an optional component in a DSR signaling node. DPs are deployed in conjunction with the Subscriber Database Servers (SDS) to provide full subscriber address resolution capability.

Below are the redundancy models supported for DPs:

- Active-Active configuration for scaling up to ten total SS7 DPs per DSR signaling node.

IP Front End

The DSR IP Front End provides TCP/SCTP connection based load balancing to hide the internal hardware architecture and IP addresses from the customer network. This is an optional component in a DSR signaling node.

IPFE blades can be deployed with an active-active configuration. In either configuration, the maximum number of supported IPFE blades is four per DSR signaling node.

The DSR Bare Metal form factor supports IPFE in the following configurations per DSR signaling node:

- Active-Active configuration for scaling up to four IPFEs per DSR signaling node.
- An Active Only simplex configuration for labs and trials.

A FV RMS DSR always includes IPFE and it is a deployment decision whether or not to use it. In the rack mount server form factor, the IPFE function is virtualized onto the servers along with the DA-MP, NOAM, System OAM and PMAC functions.

Session Binding Repository

The OC-DRA uses the session database but does not currently use the subscriber binding database. Policy DRA uses both the session database and the subscriber binding database.

DSR supports an optional Gateway Location Application (GLA) on nodes that have access to the SBR databases. This application provides a Diameter signaling mechanism to access a subset of the state in the SBR databases. If GLA is implemented, one additional blade is equipped per network to allow for the query capability.

The SBRs are pooled and distributed resources for maintaining state information. The state information is scoped differently based on the database function. For example, in the Policy DRA application:

- Subscriber binding state is scoped for the entire Policy DRA network. The pool of SBR(b)'s maintaining subscriber binding are used by the entire set of DSR signaling nodes running the Policy DRA application in a given network.

  Example: in an eight node network, all eight nodes communicate and use the subscriber binding database.

- Session state is scoped for a DSR mated pair or mated triplet. The pool of SBR(s)'s maintaining the session state is shared between a signaling node and its mate. Note: a mate is not always used by all operators in which case the SBR(s) cover only a single DSR.

  Example: continuing the eight node example above – this network is comprised of four mated pairs. Each mated pair has its own session database. So there are four session databases.

The example shown above assumes that the Diameter Peers (PCRFs and Policy Clients) have geo-redundant Diameter connections to a mated pair of DSRs. However, some DSR customers may have Peers that don't support geo-redundant connections, or may not want to configure their networks with geo-redundant connections. In this case the DSRs will typically not be configured as mated pairs, and there will be a separate SBR(s) for each DSR.

In summary, the session database is active in every DSR signaling node running the Policy DRA application, but the subscriber binding database is normally only active in a subset of the DSRs. The following figure shows the SBR relationship with the two databases and the deployment architecture (form factor).

To describe the configuration of the SBRs, the concept of a server group is introduced. A server group is a collection of servers that work together to provide a specified functional behavior. In this case, a server group provides SBR functionality. Each SBR Server Group consists of one, two, three or four servers, depending on the redundancy model desired.
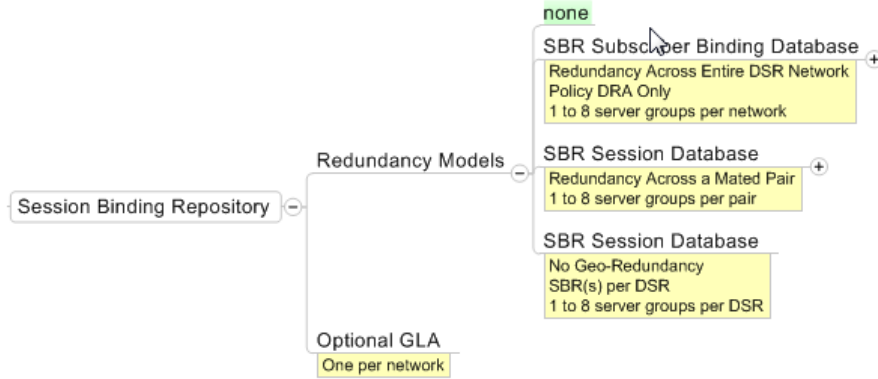
Figure 17: DSR SBR

The following information provides guidance on the decision making to determine which redundancy model is appropriate for a particular configuration.

Table 3: Redundancy Model Decision Making

Redundancy	Highly Available	Geo-redundant	Min # of Rqd Nodes	Pros	Cons	Typically Used…
Active Only	No	No	1	Small footprint/power. Lease Cost Option. Functionality with no replication.	Low availability/reliability; Can't upgrade without scheduled downtime	Not allowed in production. Labs/trials only. Used for demonstrating behavior when HA/geo/capacity testing is not needed.
Active-Standby	Yes	No	1	High availability. Small footprint/power. Lower Cost Option. Single site high availability.	Not geo-redundant.	Labs/trials where HA testing is needed. Singleton production DSRs where lower availability is acceptable.
Active-Standby-Spare	Yes	Yes	2	Geo-redundant, highly available. Fast switchover to standby on loss of active.	Higher cost, footprint, power. Asymmetrical deployment model for mated pairs (i.e. one node can end up with more servers than the other).	Typical production model when HA & geo are needed.
Active-Spare	Yes	Yes	2	Geo-redundant, highly available. Symmetrical deployment model for mated pairs. Most cost effective for HA/Geo.	Slow switchover to spare on loss of active. Moderate availability due to longer switchover time.	Acceptable production model when HA & geo are needed. Consider for lower cost solutions where moderate availability is acceptable

Subscriber Database Server (SDS) Node

The SDS consists of a primary and optional backup Subscriber Database Server and an optional Query Server. The SDS is only used in conjunction with DSR Signaling sites when full address resolution features are needed.

Network Architecture

SDS supports layer 2 demarcation to the customer network at the server level.

Figure 18: SDS Network Architecture

High Availability

The SDS provides a provisioning feed to the subscriber database and is not considered to be a real-time critical component of a DSR network deployment. However, customers requiring a subscriber database are generally hosting large network solutions and still require that the SDS is highly available.

Figure 19: SDS High Availability

- SDS is deployable in a highly available configuration where primary and backup servers are deployed as a pair.

- In configurations where a highly available SDS is not required (e.g. lab), then the SDS is deployed using a primary server only.

Site Redundancy

DSR optionally supports an SDS Disaster Recovery site for geographic redundancy.

Figure 20: SDS Site Redundancy

- The SDS can be deployed as a single site.

- The SDS can be deployed as a pair: the primary site, and a disaster recovery site.

- The disaster recovery site, when used, has a functionally duplicated configuration to the primary site.

Power

The SDS supports both AC and DC power options.

Query Server

The query server allows non-real-time applications to query the provisioned SDS database.

Figure 21: SDS Query Server

SDS allows for an optional query server.

Co-location

The SDS may be co-located with one or more products. Co-location is the ability for the SDS to be physically located in the same site footprint (cabinets) as other products.

Figure 22: SDS Co-Location

Co-location allows for different products at the same site to share certain resources such as the physical cabinet and power. Here are some examples of co-located products with an SDS:

- SDS rack mount + DSR in an enclosure

- SDS rack mount + policy in an enclosure

- SDS rack mount + NOAM rack mount

Integrated Diameter Intelligence Hub (IDIH)

The Integrated Diameter Intelligence Hub is an optional but normal component of most DSR deployments. IDIH provides troubleshooting capabilities for the DSR and supersedes the Diameter Intelligence Hub described below. All new deployments requiring DSR troubleshooting use IDIH.

Diameter Intelligence Hub (DIH)

The Diameter Intelligence Hub is an optional component of DSR deployments and is superseded by the Integrated Diameter Intelligence Hub described above. The DIH is supported for existing deployments only. The DIH provides troubleshooting capabilities for the DSR.

For Bare Metal deployments, the DIH function is provided on a single server per signaling node. DIH deployment requires use of external switches to provide port mirroring for traffic capture. For FV RMS deployments, each DIH VM (Oracle DB Application and Mediation) can be located on separate physical RMS – based on configuration.

Bare Metal

Enclosures

The bare metal version of DSR is supported on c-Class hardware. DSR is supported in up to three HP C-Class enclosures per node. Additional enclosures are supported for future expansion and for co-mingling applications.

Figure 23: Enclosures

The DSR blades may be placed in any position within an enclosure except as noted below. There are no restrictions (beyond physical) on where blades can be placed in an enclosure. That is, DA-MPs, DPs, DIH, and System OAMs, can go in any position within an enclosure. The following are the current known physical limitations:

- The DIH storage sidecar must go to the immediate right of the DIH blade (existing deployments only).

- Dividers to support ½ height blades span 2 slots (odd + even slot to the right)

In a multi-enclosure DSR configuration, the DSR blades may be placed in any enclosure. There are no restrictions (beyond physical) on which types of blades can be placed in which enclosures. In a multi-enclosure DSR configuration, blades with the same function may be split into different enclosures. For example: one System OAM, DA-MP, and DP in one enclosure with the redundant mates in a second enclosure.

Fully Virtualized Rack Mount Server (FV RMS) Signaling node

DSR supports a fully virtualized rack mount server (FV RMS) form factor for DSR signaling nodes. The rack mount server configuration makes use of virtualization to host multiple functions per server.

Engineering configurations for the servers running virtual machines are set up differently to provide server optimization and scalability. For more details on the rules and configuration please see VM Placement and CPU Socket Pinning Tool – Available at Oracle.com on the Oracle Help Center (OHC). These servers can be 'packed' with the VMs to optimize usage. The figure below is a screenshot of the tool.
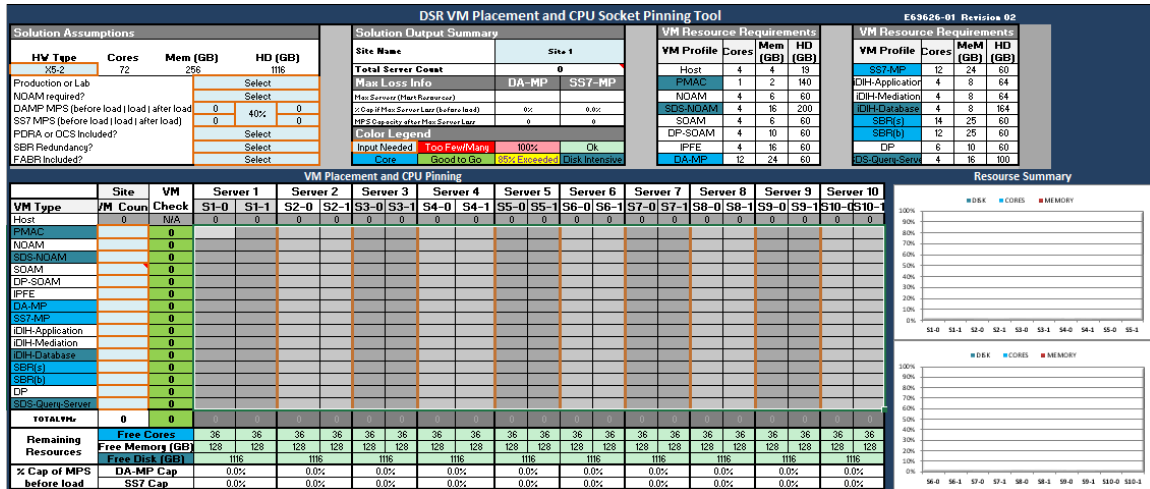
Figure 24: Tool for Fully Virtualized RMS with Packable VMs

DSR Planning

For bare metal and fully virtualized RMS the following form factors apply:

Table 4: DSR Form Factors

System	Hardware Supported
Bare Metal System	HP Gen8, HP Gen8 v2, HP Gen9, HP Gen9 v2 – all half height blades
FV RMS System	HP Gen8*, HP Gen8 v2*, HP Gen9, ORACLE X5-2, ORACLE X6-2, ORACLE X7-2, Netra X5-2, servers

* HP Gen8 and HP Gen8 v2 FV RMS deployments have the following limitations

- Up to 50K MPS
- No support for the following applications: Stateful (PDRA and OCS), Database (FABR)
- Must have dedicated IDIH/DIH servers

Any references to X5-2 imply the support for all of the following servers:

- Netra X5-2: AC or DC and NEBS compliant
- ORACLE X5-2: AC

General Planning

System OAM

The DSR nodal OAM is handled in DSR by a pair of active-standby blades / VMs referred to as System OAM. The System OAM manages the OAM for entire DSR node. The System OAM is sized to accommodate the OAM requirements for DSR and for co-resident OAM components as described below. There is no expansion or scaling of the System OAM blades / VMs.

Table 5: DSR System OAM System Capacity

	System OAM
Number of System OAMs	Primary/standby (1+1)
Bare-Metal System: Number of blades	48
FV RMS System	Please see Table 27 for VM information
System OAM supported components	DSR SOAM, DP OAM, DSR NOAM, Spare SOAM

Network OAM

The NOAM is appropriately sized to handle the DSR network. Only some types of servers are included in the NOAM managed server count:

- SOAM, DA MP, IPFE, SBR and SS7 MP servers count towards the NOAM managed server total in a network. If any of these functions is deployed as a virtual machine (VM), each VM counts as a server.

- DP, DIH and IDIH do not count towards to the total number of NOAM managed servers supported in a network.

Examples:

1. Node 1 has 2 SOAMs, 2 IPFEs, 4 DA MPs and 2 DPs. The server count for purposes of NOAM managed servers is 8 (DPs are not included)
2. Node 2 has 2 SOAMs, 2 IPFEs, 16 DS MPs and 12 SBRs. The server count for purposes of NOAM managed servers is 32.

To calculate the number of NOAM subtended servers in a network, take the sum of the server counts in each node.

The DSR Network OAM is potentially a RAM intensive function. The Network OAM is designed not to exceed the available memory; however RAM is the most likely resource constraint.

Table 6: DSR Network OAM Capacity

	Capacity
Number of NOAMs	primary/backup, plus optional disaster recovery site
Number of DSR signaling nodes supported	32
	2 (when co-resident on System OAM blades)

*The NOAM NE is validated to support up to 768 subtended servers per network.

Platform, Management and Control (PMAC)

The PMAC provides the management of the overall DSR platform. For bare metal environments, the PMAC function runs on a server called the management server. For FV RMS environments, the PMAC runs as another virtual machine.

Table 7: DSR System OAM System Capacity

	PMAC
Number of PMAC	Primary/backup, plus optional disaster recovery site
c-Class System: Number of servers	1
RMS System	Please see Table 27 for VM information.

Diameter Agent MP

MPS Calculations for DA-MP

The different DA-MPs have differing capacity limits, however, the methodology for dimensioning and determining how many DA-MPs are needed is exactly the same. The dimensioning of the DA-MPs needs to consider two primary factors:

- The set of capabilities and applications running (e.g. policy proxy, core routing, address resolution), and

- The desired processing throughput and its impact on critical resources (i.e. MPS, number of connections, IPSEC MPS).

To keep the management of the DA-MP performance simple, performance is divided into three groups:

- Core – this group consists of all routing, mediation, screening features for both diameter and radius as well as anything that doesn't fit into the next two groups.

- Database – this group consists of all database features including HSS/HLR address resolution features both range based and subscriber based address resolution (RBAR and FABR), and the inter-working function (IWF) required for MAP Diameter interworking.

- Stateful – this group consists of all subscriber/session stateful features including policy DRA, online charging DRA and includes Gateway Location Application (GLA).

If multiple applications are running on the same DA-MP, then the application with the lowest rating applies to the entire DA-MP.

In addition to the above, the use of IPSEC has a significant impact on the performance of the DA-MP and has its own capacity rating. The IPSEC de-rating only applies to the portion of the traffic running IPSEC.

Another factor that impacts performance is the support of application chaining. Application chaining refers to the case where a single DSR message encounters more than one application. The DSR supports application chaining for IWF with RBAR or FABR, Policy DRA with RBAR or FABR and Online Charging with RBAR or FABR. The application chaining de-rating applies when more than 10% of the traffic on a given DA-MP is subject to application chaining.

Message size and number of connections also have an impact on the overall performance of the blade; however Oracle has factored these items into the capacity ratings of the blade. The following formulas can be used to determine the required number of DA-MPs:

Number of DA-MPs = greater of (MPS required/MPS per blade or VM, connections required/connections per blade or VM)*

Note: If the customer desires DA-MP redundancy at the DSR node level, they should configure sufficient servers and DA-MPs such that the loss of a physical server leaves a sufficient number of active DA-MPs.

If IPSEC is in use, then adjust the MPS per blade value by 40% de-rating for the % of traffic using IPSEC.

Examples:

- Customer A uses HSS address resolution and is running 20K MPS S6a traffic. In this case, a pair of bare metal DA-MPs easily supports the traffic.

- Customer B is running the policy proxy and needs to run 100K MPS policy proxy traffic using bare metal. Then 4+1= 5 blades are required.

Table 8: DSR DIAMETER Agent MP Capacity per Hardware Type

		Performance Grouping		
		Core Routing	Database IWF IWF with RBAR Chaining	Stateful Applications (including GLA) IWF with FABR Chaining
Bare Metal	MPS	45,000	40,000	30,000
	Connections	2000	2000	2000
	IPSEC	40%1 degradation	40%1 degradation	40%1 degradation
FV RMS	MPS	10,000	10,000	10,000
	Connections	1000	1000	1000
	IPSEC	40%1 degradation	40%1 degradation	40%1 degradation

Degradation applies only to the portion of traffic running IPSEC.

Capacities advertised are commercial capacities: 1+1 non-scaled DA-MP performs significantly better.

Entire blade / server is rated at the capacity of the slowest application running on the blade / server. For example: a blade / server running core routing and policy proxy (a stateful application) is rated at 30K MPS on the Gen8 blade.

When application chaining is used and makes up more than 10% of the overall traffic, then a 15% derating of the blade is applied.

Use of diameter mediation has an impact on the diameter agent performance. Please refer to the Mediation user's guide to learn more about the performance impact of using mediation.

Table 9: DSR DIAMETER Agent Site Capacity

		Performance Grouping		
		Core Routing	Database	Stateful Applications
Universal	Peer Route Tables (PRTs)	500	500	500
	Routing Rules per PRT	1,000	1,000	1,000
	Routing Rules per System	10,000	10,000	10,000
	Peer Route Lists	2,000	2,000	2,000
	Peer Routing Groups	6,000	6,000	6,000
	Range based routing rules	1,000,000	1,000,000	1,000,000
Bare Metal	Max Number of DA-MPs (HP Gen8, Gen8v2,Gen9, Gen9v2)	16*	16*	16*
	Total MPS	675,000	640,000 for all HW except Gen9v2 675,000 for Gen9v2	policy and online charging DRA 480,000 (for all HW except Gen9v2) 675,000 for Gen9v2
	Total Connections	16,000	16,000	16,000

		Performance Grouping		
		Core Routing	Database	Stateful Applications
FV RMS	Max Number of Servers	8 X5-2 and Netra X5-2 8 X6-2 8 X7-2 10 HP Gen9/HP Gen9 v2		
	Total MPS	(Relay) 250,000 for X5-2, Netra X5-2, X6-2, X7-2 and HP Gen9 (Database) 200,000 for X5-2, Netra X5-2, X6-2, X7-2 and HP Gen9 (Stateful) 160,000 for X5-2, Netra X5-2, X6-2, X7-2 and HP Gen9 (Relay)HP Gen8 / HP Gen8 v2 limited to 50K		
	Total Connections	8,000	8,000	8,000

*All DPs are active. A customer should plan capacity of the site based on their desired redundancy model. For example, if the customer desires to operate with no impact from the loss of a single blade, then the site should be engineered with one extra DA-MP blade, making the capacity of a DSR with N blades being (N-1) times the capacity of a blade. For FV RMS, the VM Placement and CPU Socket Pinning Tool – Available at Oracle.com on the Oracle Help Center (OHC) calculates the percentage of capacity lost due to a single server failure.

Database Processor

Subscriber and Processing Calculations for DP

Entities may be either text based strings (text based IMPU/IMPIs) or numeric strings (e.g. IMSI, MSISDN). Each subscriber may have one or more entities associated with it. Differing applications and specific needs of a specific application implementation dictates how many entities are needed for each subscriber. The maximum number of subscribers varies based on the number of entities required for each subscriber.

The following formula can be used to determine the required number of available subscribers:

Number of subscribers = Max number of entities / number of entities per subscriber

The max number of entities can be supported by a single DP blade or VM. An additional blade / VM is used for redundancy.

The number of database accesses varies on the amount of traffic, type of traffic, and application implementations needing subscriber database access. For example:

- Customer B has a DSR node which is sized to process 200K MPS. The customer is running HSS address resolution on S6a traffic. Approximately 30% (60K MPS) of the traffic is S6a, of which 30K MPS is S6a Request messages and 30K MPS is S6a Answer Messages. Since only S6a Request messages need to access the database the required rate is 30K database accesses per second.

The following formulas can be used to determine the number of database accesses, and the number of DPs needed:

Number of database accesses/s = MPS per node x %traffic on applicable interfaces x %msgs needing DB access

Number of DPs = round up (database accesses/s / DB lookups /s) + 1

Table 10: DSR Database Processor Capacity

	Performance Grouping	
	Bare Metal Database	**FV RMS Database**
Number of Numeric Entities	1,000,000,000	60,000,000
Number of Text Entities	10,000,000	500,000
DB lookups per second	140,000	140,000

Table 11: DSR Database Processor System Capacity

	Performance Grouping	
	Bare Metal Database	**FV RMS Database**
Max Number of DPs	10*	10
Max Number of Numeric Entries	1,000,000,000	60,000,000
DB Lookups per second	140,000	140,000

*All DPs are active. A customer should plan capacity of the site based on their desired redundancy model. For example, if the customer desires to operate with no impact from the loss of a single blade, then the site should be engineered with one extra DA-MP blade, making the capacity of a DSR with N blades being (N-1) times the capacity of a blade. For FV RMS, the VM Placement and CPU Socket Pinning Tool – Available at Oracle.com on the Oracle Help Center (OHC) calculates the percentage of capacity lost due to a single server failure.

IP Front End

Bandwidth Calculations for IPFE

The IPFE is managed by bandwidth. The bandwidth of an IPFE pair is 3.2 Gbps. For the IPFE bandwidth and capacity planning, customers should consider the impact of IPFE blade failure resiliency and DSR node redundancy in their network. Examples:

- For an IPFE blade pair at a given DSR site, when the aggregate traffic across the IPFE blade pair exceeds 3.2 Gbps then a single IPFE blade failure at the site would cause the traffic on the remaining IPFE blade to exceed its rated capacity.

- If you ARE engineering for a 'local site IPFE blade failure AND a mate site DSR node failure' simultaneously, then, the aggregate traffic across the redundant IPFE blade pairs (i.e. all 4 IPFEs across the 2 sites) must not exceed 3.2 Gbps, since under such a scenario, all of the Client traffic would traverse the single remaining IPFE blade in this failure mode.

- If you are NOT engineering for a 'local site IPFE failure AND a mate site DSR node failure' simultaneously, then, the normal volume of traffic across each IPFE blade and the TSA redundancy scheme across the mated DSR pair dictates the capacity planning threshold.

- If you have a normal distribution of half the Client traffic to each IPFE blade pair, in the worst case all traffic ends up being carried by a single IPFE blade in the remaining DSR, the capacity planning threshold for IPFE bandwidth could be a percentage below 1.6 Gbps for the aggregate traffic across each IPFE blade pair.

In summary, customers should consider following network and deployment attributes for IPFE bandwidth and capacity planning:

- The traffic on a given IPFE blade

- The aggregate traffic across an IPFE blade pair at a given DSR site

- The aggregate traffic across redundant IPFE blade pairs in a mated DSR pair (i.e. all 4 IPFEs across the 2 sites)

To determine how much bandwidth is needed for an IPFE deployment, use the average Diameter message size (from measurement RxMsgSizeAvg or predicted message size) and the MPS to lookup the required bandwidth in the table below. The table indicates the resultant DSR MPS for various average Diameter Request message sizes assuming the following:

- Under normal operating load, clients will perform standard aggregation of Diameter PDUs into TCP packets, up to the path MTU. The path MTU between clients and the IPFE is assumed to be at least 1500 bytes.

Table 12: IPFE Bandwidth and Resultant DSR MPS

IPFE Bandwidth	Average Diameter Message Size (bytes)	Resultant DSR MPS
0.5 Gbps	500	125,000
	1,000	62,500
	1,500	41,666
	2,000	31,250
1.0 Gbps	500	250,000
	1,000	125,000
	1,500	83,333
	2,000	62,500
1.5 Gbps	500	375,000
	1,000	187,500
	1,500	125,000
	2,000	93,750
2.0 Gbps	500	500,000
	1,000	250,000
	1,500	166,667
	2,000	125,000
2.5 Gbps	500	625,000
	1,000	312,500
	1,500	208,333
	2,000	156,250
3.2 Gbps	500	800,000
	1,000	400,000
	1,500	266,666
	2,000	200,000

IPFE System Engineering

Table 13: DSR IPFE MP Capacity

	Bare Metal	FV RMS
Bandwidth Supported	3.2 Gbps	Please see Table 15: IPFE Capacity on Virtualized DSR
Connections	16000	8000
Number of active DA-MPs load-balanced per IPFE pair	16	
Supported Protocols	TCP, SCTP (uni-homed or multi-homed)	

Table 14: DSR IPFE MP System Capacity

	Bare Metal	FV RMS
Max Number of IPFE blades/VMs	4 = 2+ 2	4 = 2+ 2
Total Bandwidth	6.4 Gbps	Please see Table 15: IPFE Capacity on Virtualized DSR
Total Connections	16000	8000

Table 15: IPFE Capacity on Virtualized DSR

	Single TSA on IPFE Pair		2 or more TSAs on IPFE Pair	
	Avg Msg Size < 1 MTU	Avg Msg Size >= 1 MTU	Avg Msg Size < 1 MTU	Avg Msg Size >= 1 MTU
2,000 Connections or less	1 Gbits/sec	1 Gbits/sec	1 Gbits/sec	1 Gbits/sec
More than 2,000 Connections	1 Gbits/sec	800 Mbits/sec	1 Gbits/sec	1 Gbits/sec

SS7 Message Processor

The SS7 message processor supports MAP Diameter Interworking Function (IWF) in conjunction with the DA-MPs. The performance ratings for the SS7 MPs are based specifically on MAP messages with assumptions of a 170 byte average message size with a 4K byte maximum message size and 20 second average transaction hold time.

SS7 MP System Engineering

Table 16: SS7 MP Capacity

	Performance MAP Diameter IWF	
	Bare Metal	FV RMS
MAP MPS	30,000	10,000

Table 17: SS7 MP System Capacity

	Performance MAP Diameter IWF	
	Bare Metal	FV RMS
Max Number of SS7 MPs	8	4
Max Nodal MAP MPS	210,000	40,000

vSTP Message Processor

The vSTP MP is only available in the cloud deployment model and not .available on the bare-metal or the FV RMS deployment models. Please refer to the DSR cloud benchmarking guide for vSTP performance and capacity information.

Session Binding Repository

Sessions, Subscribers and MPS Calculations for SBR

The SBRs are managed by the number of concurrent sessions, the number of active subscribers (Policy DRA only) and the amount of Policy and/or Online Charging DRA MPS. Other factors in determining the SBR architecture include the DSR deployment architecture (small, large, lab, trial, number of nodes etc.) and the SBR server group and redundancy model desired.

To deploy each of the redundancy models, varying numbers of servers are needed:

Table 18: Required Server by Redundancy Model

Redundancy	Number of Servers per Server Group			Redundancy Model
	Servers in Node A	Servers in Node B	Total Servers in Group	
Active Only	1	0	1	non-HA, non-Geo
Active-Standby	2	0	2	HA only, non-Geo
Active-Standby-Spare	2	1	3	HA, Geo
Active -Spare	1	1	2	HA, Geo
Active - Standby - Spare - Spare	See note	See note	4	HA, Geo

Note: This redundancy model is for use with mated triples and should only be used with prior approval.

Here are two examples:

- One server group is needed to support a small lab using Active Only. In this case, 1 server group = 1 server, so only a single SBR is needed.

- Four server groups are needed to support a network consisting of a DSR mated pair using Active-Standby-Spare. In this case 4X3=12 servers are needed. In the table above, the Node A and Node B designations are arbitrary. That is, for any server group, the node A and node B designations can be reversed. This allows for relatively uniform distribution of the servers among the DSRs. In this example, there are four Active-Standby servers or 4X2= 8 servers in node A, plus four standby servers 4X1=4 servers in node B. This is a valid deployment scenario, but to balance the server deployments, the first node would be assigned the node A role for two of the server groups and then would be assigned the node B role for the remaining 2 servers groups. The mate would have the node roles reversed. The result would be that each node would have 6 SBR servers deployed as shown below.

Figure 25: SBR Distribution Example (SG=Server Group)

The subscriber binding and session databases have dedicated server groups. The system scales by adding server groups - though it is good engineering practice to configure enough server groups to handle at least 18 months of traffic growth.

Examples:

- Very small DSR network with one mated pair of DSR nodes and Active-Spare server group. The number of SBRs is dependent upon the redundancy model selected. In this case, there are two SBR servers that provide subscriber binding and two SBR servers that provide session binding.

- Small DSR network with one mated pair of DSR nodes, Active-Standby-Spare server groups with dedicated session and subscriber binding databases. The number of SBRs required can be determined using the formulas below.

Number of session SBR(s) servers = Greater of:

(Required Concurrent Sessions / maximum number of concurrent sessions per SBR server group) * 3 (for Active-Standby-Spare redundancy model)

OR

(P-DRA MPS / maximum P-DRA MPS per SBR server group) * 3 (for Active-Standby-Spare redundancy model)

Number of binding SBR(b) servers = Greater of:

(Active subscribers / maximum number of subscribers per combined SBR server group) * 3 (for Active-Standby-Spare redundancy model)

OR

(P-DRA MPS / maximum P-DRA MPS per SBR server group) * 3 (for Active-Standby-Spare redundancy model)

Table 19: SBR Nodal / Mated Pair Capacity

	Performance Grouping Stateful Applications Policy DRA and/or Online Charging DRA	
	Bare Metal	**FV RMS**
Max Session Server groups per node/mated pair	8*	3*
Max Concurrent Sessions per node/mated pair	240M	48M
Max DRA MPS per node/mated pair	450K	160K
Max Subscriber Binding Server groups per node/mated pair (Policy DRA only)	8*	8*

* Actual blade count depends on the selected redundancy model

Table 20: SBR Network Capacity

	Performance Grouping Stateful Applications			
	Bare Metal Policy DRA	**Bare Metal Online Charging DRA**	**FV RMS Policy DRA**	**FV RMS Online Charging DRA**
Max P-DRA nodes per network	32 (16 pairs)		16 (8 pairs)	N/A
Max Subscriber Binding Server groups per network	8		8	N/A
Max subscriber binds per network	200M		64M	N/A
Max concurrent session binds per network	480M		384M*	N/A

* If the PCA nodes are deployed as singletons and not pairs (i.e. SBR-S scope is a node, not a pair) the max could be twice this amount.

Table 21: Subscriber SBR(b) Server Group Capacity

	Performance Grouping Stateful Applications			
	Bare Metal Policy DRA	**Bare Metal Online Charging DRA**	**FV RMS Policy DRA**	**FV RMS Online Charging DRA**
Subscriber Binds	50M	N/A	8M	N/A
MPS (of Policy DRA traffic)	250,000*	N/A	50,000	N/A

* assumes up to 17% CCR-I + CCR-T + CCA-I + CCA-T traffic mix with up to 2 alternate keys, assumes average session size of 100 bytes. If GLA application is running, each GLA application MPS counts as 3 MPS for the purposes of dimensioning the SBRs.

Table 22: Session SBR(s) Server Group Capacity

	Performance Grouping Stateful Applications Policy DRA and/or Online Charging DRA	
	Bare Metal	**FV RMS**
Concurrent Sessions	60M	16M
MPS (of Policy DRA and/or Online Charging DRA traffic)	200,000*	50,000
Topology hiding (Policy DRA only)	125K with 100% Rx TH** 135K with 100% Gx TH 100K with 100% Rx & Gx TH	50,000

* assumes up to 17% CCR-I + CCR-T + CCA-I + CCA-T traffic mix with up to 2 alternate keys, assumes average session size of 100 bytes. If GLA application is running, each GLA application MPS counts as 3 MPS for the purposes of dimensioning the pSBRs.

** assumes up to 32% Rx traffic and up to 25% Gx traffic

Planning for SDS

The SDS consists of a primary and optional backup Subscriber Database Server and an optional Query Server. The SDS is only used in conjunction with DSR Signaling sites when full address resolution features are needed.

Subscriber Database Servers

There are two scaling factors that should be considered for the SDS:

The provisioning rate. This is the rate at which new subscribers and associated information can be written to the SDS database.

Number of DSR Signaling nodes synchronized.

Table 23: DSR Subscriber Database Server (SDS) System Capacity

	Database	
	Bare Metal	**FV RMS**
Number of servers or VMs	1+1	1+1
Number of query servers (optional)	1	1
Max Number of Subscribers	500,000,000	30,000,000
Number of Numeric Entities	1,000,000,000	60,000,000
Number of Text Entities	10,000,000	500,000
Number of writes/updates per second	200	100

Table 24: DSR Subscriber Database Server (SDS) Network Capacity

	Database	
	Bare Metal	**FV RMS**
Number of SDSs (active + optional disaster recovery SDS)	2 = 1+1	2=1+1
Number of DSR signaling nodes synchronized	32	24

Query Server

The Query Server has been designed to be a best effort server for free-form variable length and complexity queries. DSR supports 1 query server, however the system has been designed to add additional query servers if needed.

A new query server may be needed if the latency of the average query-response grows to a point that is not acceptable by the querying application.

Integrated Diameter Intelligence Hub (IDIH)

 The IDIH function is provided on a single server per signaling node. IDIH is supported on the rack mount server form factor for all new deployments. IDIH can be implemented on a c-Class blade for existing DIH deployments only by repurposing the existing DIH blade. DIH rack mount server hardware at existing deployments can also be repurposed for IDIH. DSR supports one IDIH per signaling node. DIH and IDIH can run simultaneously on the same node.

New deployments of IDIH servers allow for co-residence with the PMAC function. In other words, one server can be used for IDIH plus the PMAC function.

Table 25: Integrated Diameter Intelligence Hub Capacity

	Performance Rating
IDIHs supported	up to 1 per DSR signaling node
Supported Capacity	1K TTRs/s with 2.5KB average TTR size 700 TTRs/s with 5.0KB average TTR size
Supported Interfaces	Diameter - Sh, Cx, Gq, Cc, S6a, S6d, Gx, Rx, Gy, SLg, SLh, Gxa, SWm, SWx, STa, S6b, S9, Sd, Sy, S13
FV RMS Form Factor	HP Gen8, HP Gen8 v2, HP Gen9 server plus 6 900 GB hard drives
Bare Metal Form Factor*	HP Gen8, HP Gen8 v2 ½ height blade plus attached sidecar w/12 600 GB hard drives

*For existing deployment only.

Please Note: iDIH install activities do not impede, impact or affect DSR. The DSR and iDIH are independent systems and do not know about each other before or during installation. There is no automated discovery process where the two systems discover each other. It is only after the IDIH installation is completed, that the DSR is made aware of iDIH's presence via manual configuration. There are no impacts on the DSR with respect to re-installing DIH 1.2 in the event of a unsuccessful attempt to install IDIH.

Diameter Intelligence Hub (DIH)

The DIH is an optional component of the DSR and is supported for existing deployments only. DIH provides trouble shooting capabilities for acquiring up to 20K MPS of port mirrored Diameter Agent signaling traffic and analyzing and creating xDRs for up to 10K MPS assuming average message size of 500 bytes.

Table 26: Diameter Intelligence Hub Capacity

	Performance Rating
DIHs supported	up to 1 per DSR signaling node
Supported Capacity	1K TTRs/s with 2.5KB average TTR size 700 TTRs/s with 5.0KB average TTR size
Supported Interfaces	Diameter - Sh, Cx, Gq, Cc, S6a, S6d, Gx, Rx, Gy, SLg, SLh, Gxa, SWm, SWx, STa, S6b, S9, Sd, Sy, S13
FV RMS Form Factor	HP Gen8, HP Gen8 v2 server plus 2 900 GB hard drives and 12 600 GB hard drives
Bare Metal Form Factor	HP Gen8, HP Gen8 v2 ½ height blade plus attached sidecar w/12 600 GB hard drives

Planning for Fully Virtualized Rack Mount Servers

A rack mount server based DSR signaling node consists of a number of virtual machines (VM) which provide specific capabilities and the servers they are hosted on.

The following table describes the VM to server affinity for X5-2, X6-2, HP Gen9, HP Gen8, HP Gen8 v2 servers. For HP Gen8 and HP Gen8 v2 there is no support for the following applications: Stateful (PDRA and OCS), Database (FABR) and IDIH/DIH VMs.

Table 27: VM Placement to Server Mappings

Virtual Machine	Redundancy Models	DSR Applications	Optional/ Mandatory	VMs Per Server	Max VMs Per Node	When Used?
HOST	N/A	Core	Mandatory	1	1	1 per Server. Not a VM. Reserves resources on the TVOE Host server
NOAM	Active/Standby	Core	Mandatory	1	1-2	2 VMs per DSR network in any site. VMs to be deployed on separate servers. If managing more than 2 DSR nodes, the NOAM may be configured as a medium or large VM profile. (see above for standalone NOAM planning) An optional Disaster Recovery NOAM can be located at a separate site.
SOAM	a) Active/Standby b) Active/Standby/ Spare	Core	Mandatory	1	1-3	Active/Standby VMs to be deployed on separate servers, Spare must be placed on the mated DSR node. Redundancy model (b) is used for PDRA or OCS mated-pair deployments which utilize Spare SBRs (a Spare SOAM is required in any DSR containing a Spare SBR). For all other deployments redundancy model (a) is used.
PMAC	Active/(Optional) Cold Standby	Core	Mandatory	1	1	Mandatory 1 per node, optional 2nd PMAC. 1 VM per DSR site 2 if Cold Standby. Cold Standby should be placed on a different server than active server.
DA MPs	Active Cluster	Core	Mandatory	0-1	1 to 16	A minimum of 2 servers required for any production configuration. VM's must be evenly distributed across servers to minimize capacity loss.
IPFE	Active, Active	Core	Highly Recommended but Optional	1	1-4	Deployed in pairs. Max 2 pairs (4 VMs). Each VM must be deployed on a separate server.

Virtual Machine	Redundancy Models	DSR Applications	Optional/ Mandatory	VMs Per Server	Max VMs Per Node	When Used?
SS7 MPs	Active Cluster	MAP IWF	Optional	1	1-4	SS7 MPs for MAP Diameter IWF are optionally deployed. Active cluster per site. To be evenly distributed across servers to minimize capacity loss.
SBRs*	Per Server Group a) Active/Standby b) Active/Spare c) Active/Standby, Spare	PCA	Optional	1	1-9	Active/Standby redundancy model (a) VMs to be deployed on separate servers. Spare must be placed on the mated DSR node. Redundancy model (b) or (c) model is used for PDRA or OCS mated-pair deployments.
SBRb* (1-8 server groups per network)	Per Server Group a) Active/Standby b) Active/Spare c) Active/Standby, Spare	PDRA	Optional	1	1-9	Active/Standby redundancy model (a) VMs to be deployed on separate servers. Spare must be placed on the mated DSR node. Redundancy model (b) or (c) model is used for PDRA mated-pair deployments.
DP*	Active Cluster	FABR	Optional	1	1-10	Active cluster per site. To be evenly distributed across servers to minimize capacity loss
SDS NOAM*	Active/Standby	FABR	Mandatory if FABR is deployed	1	1-2 (please see text in next column)	2 VMs per DSR Network. VMs to be deployed on separate servers. Note: An optional Disaster Recovery SDS-NOAM can be located at a separate site.
SDS SOAM*	Active/Standby	FABR	Mandatory if FABR is deployed	1	1-2	2 VMs per DSR Node when FABR is enabled at the node
Query Server*	Active Only	FABR	Optional	1	1	1 per SDS-NOAM site where query server is desired.
IDIH – Application*	N/A	Core	Optional	1	1	1 per site All 3 IDIH VMs per DSR site, no redundancy. Must have layer 2 connectivity to other IDIH VM types per DSR site.
IDIH – Mediation*	N/A	Core	Optional	1	1	
IDIH – Database*	N/A	Core	Optional	1	1	

* HP Gen8 and HP Gen8 v2 do not support these applications.

Single Server Lab

A DSR single server lab environment is provided on the X5-2, X6-2 or HP Gen9 hardware platform. This provides a fully functional non-HA lab on a single server. This cost effective option for customers to test feature functionality does the following:

- Provides the ability to test all application functionality on a single X5-2, X6-2 or HP Gen9 server

- OAM servers run in a pair (active/standby) while all other VMs run single in order to support upgrade.

- Does NOT test performance and capacity.

The single server lab capacity is as follows:

- 500 MPS

- 5M SDS (Subscribers Database Entries)

- 5M SBR (Sessions/Bindings)

Baseline Reference/Assumptions/Limitations

This document is based on Platform 7.0.2 and the Platform Feature Guide – Available upon request should be referenced for further details on the platform.

Metrics/KPIs

In this section, the key metrics for managing the performance of the various system components are shown. There are many more measurements available, and these can be found in Diameter Measurements Reference – Available at Oracle.com on the Oracle Help Center (OHC).

Measuring System OAM Utilization

Table 28: Key Metric for System OAM Blade / VM

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Action
31056	RAM_UtilPct_Average	System	Blade/VM	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle

Measuring DA-MP Utilization

Table 29: Key Metrics for DA-MP

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	**Action**
10202	RxMsgRateAvgMp	MP Performance	Server Group	Average Ingress Message Rate (messages per second) utilization on a MP server. The Ingress Message Rate is the number of ingress Diameter messages that are targeted for Relay Agent routing (non-zero Application ID).	When running in normal operation with a mate in normal operation, and this measurement exceeds 30% of the rated maximum OR exceeds 60% of the rated capacity when running without an active mate.	The operator should determine if Additional growth in diameter traffic is continuing. If so, an estimate of the additional rate of MPS should be calculated and additional DA-MP blades should be planned for. This condition could also be an anomalous spike in traffic, and the operator may choose to ignore the occurrence. Above 40% in normal operation indicates an immediate need for additional DA-MPs.
10204	EvDiameterProcessAvg	MP Performance	Server Group	Average percent Diameter Process CPU utilization (0-100%) on a MP server.	When running in normal operation with a mate in normal operation, and this measurement exceeds 30% of the rated maximum capacity OR exceeds 60% of the rated capacity when running without an active mate.	If additional growth in the system is anticipated, then consider adding an additional DA-MP. It's possible that the traffic mix is different than originally dimensioned (e.g. 40% IPSEC instead of the originally dimensioning 5%). In these cases, re-assess the dimensioning with the actual traffic/application mix and add additional DA-MPs blades as needed.

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Action
10205	TmMpCongestion	MP Performance	Server Group	Total time (in milliseconds) spent in local MP congestion state.	Any number greater than 0.	After eliminating any configuration, or major failure conditions, then is a late indication that additional DA MPs are needed. It is highly desirable that planning for additional DA-MPs happens before this condition occurs.
10133	RxMsgSizeAvg	Diameter Performance	Server Group	The average ingress message size in Diameter payload octets.	Average message size > 2000 bytes.	DA-MP dimensioning assumes 2K average message size. This information is used to dimension IPFEs and DIH/IDIH. No action required if there are no alarms associated with the PDU message pool (available memory for messages). If PDU message pool is exhausting, contact Oracle.
31056	RAM_UtilPct_Average	System (blade)	System	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle
31052	CPU_UtilPct_Average	System (blade)	System	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When running in normal operation with a mate in normal operation, and this measurements exceeds 30% of the rated maximum capacity OR exceeds 60% of the rated capacity when running without an active mate.	If additional growth in the system is anticipated, then consider adding an additional DA MP. It's possible that the traffic mix is different than originally dimensioned (e.g. 40% IPSEC instead of the originally dimensioning 5%). In these cases, re-assess the dimensioning with the actual traffic and application mix and add additional DA-MPs blades as needed.

Measuring DA-MP Connection Utilization

Table 30: Key Metrics for DA-MP Connections

Meas ID	Name	Group	Scope	Description	Recommended Usage Condition	Action
10500	RxConnAvgMPS	Connection Performance	connection	Average Ingress Message Rate (messages per second) utilization on a connection.	Minor alarm is set by default at 50%, major at 80%. Ingress message rate per connection is customer configurable with a max per connection of 10,000.	Configure additional connections

Measuring DP Utilization

There are two key components of the subscriber database within a DSR Signaling node: the database processors (DPs), and OAM component which runs on the System OAM blades. The key metrics for managing the DP blades are:

Table 31: DP Performance Metrics

Meas ID	Name	Group	Scope	Description	Recommended Usage Condition	Action
4170	DpQueriesReceived	DP	System (blade)	The total number of queries received per second.	When running in normal operation with a mate in normal operation, and this measurement exceeds 30% of the rated maximum capacity, OR exceeds 60% of the rated capacity when running without an active mate.	The operator should determine if additional growth in the number traffic requiring subscriber database look-ups is continuing to grow. If so, an estimate of the additional rate of database lookups should be calculated and additional DP blades should be planned for.
31056	RAM_UtilPct_Average	System	System (blade)	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle
31052	CPU_UtilPct_Average	System	System (blade)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When running in normal operation with a mate in normal operation, and this measurements exceeds 30% of the rated maximum capacity OR exceeds 60% of the rated capacity when running without an active mate.	Oracle considers this measurement of lesser importance to the DpQueriesReceived. However, this measurement in conjunction with DpQueriesReceived can be used to indicate the need to add additional DPs.

While memory is a consideration for the DPs, the SDS provides the centralized provisioning of the entire DSR network. It is recommended that the metrics from the SDS are used for managing the number of available subscribers instead of the metrics of individual DPs.

Measuring IPFE Utilization

Table 32: IPFE Blade Performance Metrics

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	**Action**
5203	RxIpfeBytes	IPFE Performance	Server Group	Bytes received by the IPFE.	If the number of (bytes * 8 bits/byte)/(time interval in s) is > 3.0 Gbps.	If the traffic is expected to grow then, consider adding an additional IPFE pair.
31052	CPU_UtilPct_Average	System	System (blade)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When running in normal operation with a mate in normal operation, and this measurements exceeds 30% of the rated maximum capacity, OR exceeds 60% of the rated capacity when running without an active mate.	Contact Oracle
31056	RAM_UtilPct_Average	System	System (blade)	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle

Measuring SBR Utilization

Table 33: Key Metrics for Session SBR Blades

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	**Action**
31052	CPU_UtilPct_Average	System	System (blade)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When this measurement exceeds 60% of the rated capacity.	Contact Oracle
31056	RAM_UtilPct_Average	System	Blade	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Action
11372	SbrPolicySessionRecs Avg	SBR Session Performance	Server Group	The number of policy sessions in progress.	If P-DRA function is enabled and OC-DRA is not enabled and average exceeds 48M. If both P-DRA and OC-DRA are enabled this average must be combined with the SbrOcSessionRecsAvg and the combined average exceeds 48M.	Contact Oracle
11441	SbrOcSessionRecsAv g	SBR Session Performance	Server Group	The number of online Charging sessions in progress	If OC-DRA function is enabled and P-DRA is not enabled and average exceeds 48M. If both P-DRA and OC-DRA are enabled this average must be combined with the SbrPolicySessionRecsAvg and the combined average exceeds 48M.	Contact Oracle

Table 34: Key Metrics for Managing the Binding SBR(b) Blades

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Action
31052	CPU_UtilPct_Average	System	System (blade)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When this measurement exceeds 60% of the rated capacity.	Contact Oracle
31056	RAM_UtilPct_Average	System	Blade	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization	Contact Oracle
11374	SbrPolicyBindingRecs Avg	SBR Binding Performance	Server Group	Average number of active SBR Policy bindings.	When this average exceeds 40M.	Contact Oracle

Measuring SS7 MP Utilization

Table 35: Key Metrics for SS7 MP and Associated IWF DA-MP

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	**Action**
11054	MAP Ingress Message Rate	MAP Diameter Interworking function (MD-IWF)		Average number of MAP messages (both requests and responses) received per second from SS7 network.	When running in normal operation with a mate in normal operation, and this measurement exceeds 30% of rated maximum, OR exceeds 60% of the rated capacity when running without an active mate.	If additional growth in MAP traffic is expected, an estimate of the traffic should be calculated and additional SS7 MP blades should be planned for. This condition could also be an anomalous spike in traffic, and the operator may choose to ignore the occurrence. Above 40% in normal operation indicates an immediate need for additional SS7 MPs.
31056	RAM_UtilPct_Average	System (blade)	System	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle
31052	CPU_UtilPct_Average	System (blade)	System	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When running in normal operation with a mate in normal operation, and this measurements exceeds 30% of the rated maximum capacity OR exceeds 60% of the rated capacity when running without an active mate.	If additional growth in the system is anticipated, then consider adding an additional SS7 MP.

Measuring Network OAM Utilization

Table 36: Key Metrics for Managing the Network OAM Servers

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Action
31056	RAM_UtilPct_Average	System	System (server)	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle

Measuring SDS Utilization

Table 37: Key Metrics for SDS Servers

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Action
4110	ProvMsgsReceived	Prov	Server Group	The total number of provisioning messages that have been received.	If the rate exceeds 150 message/s.	Contact Oracle
31052	CPU_UtilPct_Average	System	System (server)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When this measurements exceeds 65% CPU utilization.	Contact Oracle
31056	RAM_UtilPct_Average	System	System (server)	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle

Measuring Query Server Utilization

Table 38: Key Metrics for the SDS Query Server

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Action
31052	CPU_UtilPct_Average	System	System (server)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When this measurements exceeds 65% CPU utilization.	Contact Oracle
31056	RAM_UtilPct_Average	System	System (server)	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle

Measuring FV Rack Mount Servers

Table 39: Key Metrics for FV RMS C

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	**Action**
10202	RxMsgRateAvgMp	MP Performance	Server Group (VM)	Average Ingress Message Rate (messages per second) utilization on a MP server. The Ingress Message Rate is the number of ingress Diameter messages that are targeted for Relay Agent routing (non-zero Application ID).	When running in normal operation with a mate in normal operation, and this measurement exceeds 30% of the rated maximum, OR exceeds 60% of the rated capacity when running without an active mate.	If additional Diameter MPS growth is expected, an estimate of the expected traffic should be calculated and additional DA-MP blades should be planned for. This condition could also be an anomalous spike in traffic, and the operator may choose to ignore the occurrence. Above 40% in normal operation indicates an immediate need for additional DA-MPs.
10133	RxMsgSizeAvg	Diameter Performance	Server Group	The average ingress message size in Diameter payload octets.	Average message size > 2000 bytes.	DA-MP dimensioning assumes 2K average message size. This information is used to dimension the IPFEs and DIH/IDIH. No action required provided there are no alarms associated with the PDU message pool (available memory for messages). If PDU message pool is exhausting, contact Oracle.

Meas ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Action
11054	MAP Ingress Message Rate	MAP Diameter Interworking function (MD-IWF)		Average number of MAP messages (both requests and responses) received per second from SS7 network	When running in normal operation with a mate in normal operation, and this measurement exceeds 30% of the rated maximum, OR exceeds 60% of the rated capacity when running without an active mate.	The operator should determine if Additional growth in MAP traffic is continuing. If so, an estimate of the additional rate of MPS should be calculated and additional SS7 MP blades should be planned for. This condition could also be an anomalous spike in traffic, and the operator may choose to ignore the occurrence. Above 40% in normal operation indicates an immediate need for additional SS7 MPs.
31052	CPU_UtilPct_ Average	System	System (server)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When this measurements exceeds 65% CPU utilization.	Contact Oracle
31056	RAM_UtilPct_ Average	System	System (server)	The average committed RAM usage as a percentage of the total physical RAM.	If the average RAM utilization exceeds 80% utilization.	Contact Oracle

Server Configuration Options

The following tables provide the configuration options for the X5-2, X6-2 and HP Gen9 servers. For additional server and blade configuration options please refer to the Platform Feature Guide – Available upon request.

Table 40: Netra X5-2 Configuration Options

Component	Supported Options	Configurator Options	Quantity
Model	NX5-2 (Carrier Grade)	Fixed	Fixed
Form Factor	2U	Fixed	Fixed
CPU	Two 18-core E5-2699v3-series Intel ® Xeon ® Processor; processor shall be the highest CPU speed targeted by Oracle	Fixed	Fixed
Memory	16GB DIMM	Quantity configurable only	Variable: Eight(8), twelve(12), sixteen(16), or twenty (20)

Component	Supported Options	Configurator Options	Quantity
On-board Hard Disks	SAS3 SFF Solid State Drives with 400GB capacity each. SAS2 1.2 TB HDDs	Type and quantity configurable Type not mutually exclusive (Both drive types may be selected so long as quantity rules are followed).	Variable: Two(2), four(4), six(6) or eight(8) of either type. Maximum total quantity allowed: eight
NIC(s)	Dual port 10Gbe SFP+ NIC Dual port 10Gbe RJ45 NIC	Type & Quantity configurable Type mutually exclusive (only one NIC type may be selected).	Variable: Zero(0), one(1), or two(2) NICs of either type.
SFP(s)	Dual rate short reach 10Gbe SFP+ transceiver compatible with the 10GbE SFP+ NIC Dual rate long reach 10GbE SFP+ transceiver compatible with the 10GbE SFP+ NIC	Type configurable Quantity determined by SFP+ NIC selection[zero(0) if RJ45 NIC selected]	Automatic: Two SFPs per SFP+ NIC selected above. Valid values: Zero(0), two(2), four(4)
DVD-RW	Not supported	Fixed	Fixed
USB Ports	Minimum one internal USB port Minimum two external USB ports	Fixed	Fixed
LOM	Four 10GBase-T* ports on motherboard	Fixed	Fixed
Storage Controller	[Aspen] Raid HBA with minimum 512MB cache***	Fixed	Fixed
Integrated Lights-out management port	One 1000BASE-T port	Fixed	Fixed
Power Supply	AC or DC (and accompanying power cord, where appropriate)	Fixed	Fixed
OSA support (internal USB drive)	**Required (included in above USB ports)	Fixed	Fixed

Table 41: ORACLE X5-2 Configuration Options

Component	Supported Options	Configurator Options	Quantity
Model	X5-2 (Non-carrier Grade)	Fixed	Fixed
Form Factor	1U	Fixed	Fixed
CPU	Two 18-core E5-2600v3-series Intel ® Xeon ® Processor; processor shall be the highest CPU speed targeted by Oracle	Fixed	Fixed
Memory	16GB DIMM	Quantity configurable only	Variable: Eight(8), twelve(12) or sixteen(16)

Component	Supported Options	Configurator Options	Quantity
On-board Hard Disks	SAS3 SFF Solid State Drives with 400GB capacity each. SAS2 1.2 TB HDDs	Type and quantity configurable Type not mutually exclusive (Both drive types may be selected so long as quantity rules are followed).	Variable: Two(2), four(4), six(6) or eight(8) of either type. Maximum total quantity allowed: eight
NIC(s)	Dual port 10Gbe SFP+ NIC Dual port 10Gbe RJ45 NIC	Type & Quantity configurable Type mutually exclusive (only one NIC type may be selected).	Variable: Zero(0), one(1), or two(2) NICs of either type.
SFP(s)	Dual rate short reach 10Gbe SFP+ transceiver compatible with the 10GbE SFP+ NIC Dual rate long reach 10GbE SFP+ transceiver compatible with the 10GbE SFP+ NIC	Type configurable Quantity determined by SFP+ NIC selection[zero(0) if RJ45 NIC selected]	Automatic: Two SFPs per SFP+ NIC selected above. Valid values: Zero(0), two(2), four(4)
DVD-RW	Not supported	Fixed	Fixed
USB Ports	Minimum one internal USB port Minimum two external USB ports	Fixed	Fixed
LOM	Four 10GBase-T* ports on motherboard	Fixed	Fixed
Storage Controller	[Aspen] Raid HBA with minimum 512MB cache***	Fixed	Fixed
Integrated Lights-out management port	One 1000BASE-T port	Fixed	Fixed
Power Supply	AC only (and accompanying power cord, where appropriate)	Fixed	Fixed
OSA support (internal USB drive)	**Required (included in above USB ports)	Fixed	Fixed

Table 42: ORACLE X6-2 Configuration Options

Component	Supported Options	Configurator Options	Quantity
Model	X6-2 (Non-carrier Grade)	Fixed	Fixed
Form Factor	1U	Fixed	Fixed
CPU	Two 22-core E5-2600v4-series Intel ® Xeon ® Processor; processor shall be the highest CPU speed targeted by Oracle	Fixed	Fixed
Memory	16GB DIMM	Quantity configurable only	Variable: Eight(8), twelve(12) or sixteen(16)
On-board Hard Disks	SAS2 1.2 TB HDDs	Quantity configurable	Variable: Two(2), four(4), six(6) or eight(8) Maximum total quantity allowed: eight

Component	Supported Options	Configurator Options	Quantity
NIC(s)	Dual port 10Gbe SFP+ NIC Dual port 10Gbe RJ45 NIC	Type & Quantity configurable Type mutually exclusive (only one NIC type may be selected).	Variable: Zero(0), one(1), or two(2) NICs of either type.
SFP(s)	Dual rate short reach 10Gbe SFP+ transceiver compatible with the 10GbE SFP+ NIC Dual rate long reach 10GbE SFP+ transceiver compatible with the 10GbE SFP+ NIC	Type configurable Quantity determined by SFP+ NIC selection[zero(0) if RJ45 NIC selected]	Automatic: Two SFPs per SFP+ NIC selected above. Valid values: Zero(0), two(2), four(4)
DVD-RW	Not supported	Fixed	Fixed
USB Ports	Minimum one internal USB port Minimum two external USB ports	Fixed	Fixed
LOM	Four 10GBase-T* ports on motherboard	Fixed	Fixed
Storage Controller	[Aspen] Raid HBA with minimum 1 GB cache***	Fixed	Fixed
Integrated Lights-out management port	One 1000BASE-T port	Fixed	Fixed
Power Supply	AC only (and accompanying power cord, where appropriate)	Fixed	Fixed
OSA support (internal USB drive)	**Required (included in above USB ports)	Fixed	Fixed

Table 43: HP Gen9 Configuration Options

Component	Supported Options	Configurator Options	Quantity
Model	HP Proliant DL380 Gen9	Fixed	Fixed
Form Factor	2U	Fixed	Fixed
CPU	Two 12-core E5-2680v3 Xeon processors	Fixed	Fixed
Memory	16GB DIMM	Quantity configurable only	Variable: Eight(8), twelve(12) or sixteen(16)
On-board Hard Disks	SAS SFF 10K RPM drives with 900GB capacity SAS SFF 15K RPM drives with 600GB capacity	Type and quantity configurable Type not mutually exclusive (Both drive types may be selected so long as quantity rules are followed).	Minimum: Two (2) 900GB drives Maximum total quantity allowed: Eight (8)
NIC(s)	HP Ethernet 10Gb Dual port 560SFP+ NIC	Quantity Configurable	Variable: zero(0), one(1) or two(2)

Component	Supported Options	Configurator Options	Quantity
SFP(s)	Dual rate short reach 10Gbe SFP+ transceiver compatible with the 10GbE SFP+ NIC Dual rate long reach 10GbE SFP+ transceiver compatible with the 10GbE SFP+ NIC	Type configurable Quantity determined by SFP+ NIC selection	Automatic: Two SFPs per SFP+ NIC selected Valid Values: zero(0), two(2), or four(4)
DVD-RW	Not supported	Fixed	Fixed
USB Ports	5 total USB 3.0 ports: 1 front, 2 rear, 2 internal (secure)	Fixed	Fixed
LOM	4 x 1Gb ports embedded	Fixed	Fixed
Storage Controller	HPE Dynamic Smart Array B140i Controller	Fixed	Fixed
Integrated Lights-out management port	1 Gb dedicated	Fixed	Fixed
Power Supply	HP 800W Flex Slot-48VDC Hot Plug Power Supply HP 800W Flex Slot Platinum Hot Plug Power Supply (AC power supply)	Type configurable	Quantity fixed: Two (2)
OSA support (internal USB drive)	Included in above USB ports	Fixed	Fixed

Signaling Node Cabinet Layouts

The below diagrams are example layouts of different power options supporting various servers. Please see DSR Planning for exact configuration rules.

Sample DSR Layout (DC)		
Power Requirements (-48VDC)		
Recommended Power Feed		8 x 60A
Maximum Power (W)		4,110
Maximum BTU/hr		14,024
Average Power (W)		3,576
Average BTU/hr		12,202

	DC c-Class Cabinet							
44	Telect 4/4 Fuse Panel (PDP A)							
43	OPEN - FILLER PANEL							
42	OPEN - FILLER PANEL							
41	Telect HC Demarc Panel (PDP C)							
40	Telect HC Demarc Panel (PDP D)							
39	OPEN - FILLER PANEL							
38	OPEN - FILLER PANEL							
37	OPEN - FILLER PANEL							
36	OPEN - FILLER PANEL							
35	Plenum / Bracket							
34								
33	Plenum /Bracket							
32								
31	Management Server							
30								
29	OPEN - FILLER PANEL							
28	OPEN - FILLER PANEL							
27	OPEN - FILLER PANEL							
26	OPEN - FILLER PANEL							
25	OPEN - FILLER PANEL							
24	OPEN - FILLER PANEL							
23	OPEN - FILLER PANEL							
22	OPEN - FILLER PANEL							
21	OPEN - FILLER PANEL							
20	OPEN - FILLER PANEL							
19	OPEN - FILLER PANEL							
18	OPEN - FILLER PANEL							
17	OPEN - FILLER PANEL							
16	OPEN - FILLER PANEL							
15	OPEN - FILLER PANEL							
14	SEISMIC BRACE							
13								
12						SBR	IPFE	DA-MP
11								
10								
9	1	2	3	4	5	6	7	8
8								
7						SBR	IPFE	DA-MP
6								
5								
4	9	10	11	12	13	14	15	16
3	OPEN - FILLER PANEL							
2	OPEN - FILLER PANEL							
1	OPEN - FILLER PANEL							

Figure 26: DC c-Class Example Configuration

	Power Requirements (208VAC)	
	Recommended Power Feed	2 x 50A
	Maximum Power (W)	4,110
	Maximum BTU/hr	14,024
	Average Power (W)	3,576
	Average BTU/hr	12,202

	AC c-Class Cabinet							
42	Plenum / Bracket							
41								
40	Plenum / Bracket							
39								
38	Management Server							
37								
36	OPEN - FILLER PANEL							
35	OPEN - FILLER PANEL							
34	OPEN - FILLER PANEL							
33	OPEN - FILLER PANEL							
32	OPEN - FILLER PANEL							
31	OPEN - FILLER PANEL							
30	OPEN - FILLER PANEL							
29	OPEN - FILLER PANEL							
28	OPEN - FILLER PANEL							
27	OPEN - FILLER PANEL							
26	OPEN - FILLER PANEL							
25	OPEN - FILLER PANEL							
24	OPEN - FILLER PANEL							
23	OPEN - FILLER PANEL							
22	OPEN - FILLER PANEL							
21	OPEN - FILLER PANEL							
20	OPEN - FILLER PANEL							
19	OPEN - FILLER PANEL							
18	OPEN - FILLER PANEL							
17	OPEN - FILLER PANEL							
16						SBR	IPFE	DA-MP
15								
14								
13								
12	1	2	3	4	5	6	7	8
11						SBR	IPFE	DA-MP
10								
9								
8								
7	9	10	11	12	13	14	15	16
6	OPEN - FILLER PANEL							
5	OPEN - FILLER PANEL							
4	OPEN - FILLER PANEL							
3	OPEN - FILLER PANEL							
2	AC PDU A2							
1	AC PDU A1							

Figure 27: AC c-Class Example Configuration

AC-Enterprise Cabinet-870-3107-01 (SEE NOTE 1 BELOW)	
FILLER PANEL PN 804-2085-G01	
FILLER PANEL PN 804-2085-G01	
	NETWORK ZONE
FILLER PANEL PN 804-2085-G01	
FILLER PANEL PN 804-2085-G01	
FILLER PANEL PN 804-2085-G01	
FILLER PANEL PN 804-2085-G01	
FILLER PANEL PN 804-2085-G01	
FILLER PANEL PN 804-2085-G01	
SERVER X	
SERVER W	
SERVER V	
SERVER U	
SERVER T	
SERVER S	
SERVER R	
SERVER Q	
SERVER P	SERVERS (ORACLE SERVER X5-2/X6-2) MAXIMUM TWENTY-FOUR (24)
SERVER O	
SERVER N	
SERVER M	
SERVER L	
SERVER K	
SERVER J	
SERVER I	
SERVER H	
SERVER G	
SERVER F	
SERVER E	
SERVER D	
SERVER C	
SERVER B	
SERVER A	
SHIPPING BRACKET	
FILLER PANEL PN 804-2085-G01	
PDU A3/B3 (804-2062-GXX)	AC PWR
PDU A2/B2 (804-2062-GXX)	
PDU A1/B1 (804-2062-GXX)	
CABLE MANAGEMENT AREA	
FRONT VIEW	

Figure 28: AC RMS Example Configuration

U	DC-Enterprise Cabinet	
44	PDP A - TELECT 4/4 PANEL	
43	PDP B - TELECT 4/4 PANEL	
42	FILLER PANEL	
41	FILLER PANEL	
40	FILLER PANEL	NETWORK ZONE SUPPORTS UP TO 6 SWITCHES
39	FILLER PANEL	
38	FILLER PANEL	
37	FILLER PANEL	
36	FILLER PANEL	
35	FILLER PANEL	
34	FILLER PANEL	
33	FILLER PANEL	
32		
31		
30		
29		
28	FILLER PANEL	
27	FILLER PANEL	
26	FILLER PANEL	
25	FILLER PANEL	
24	FILLER PANEL	
23	FILLER PANEL	
22	FILLER PANEL	
21	FILLER PANEL	
20	FILLER PANEL	
19	NETRA SERVER NX5-2	Server Zone Supports up to 8 servers
18		
17	NETRA SERVER NX5-2	
16		
15	NETRA SERVER NX5-2	
14		
13	NETRA SERVER NX5-2	
12		
11	NETRA SERVER NX5-2	
10		
9	NETRA SERVER NX5-2	
8		
7	NETRA SERVER NX5-2	
6		
5	NETRA SERVER NX5-2	
4		
3	FILLER PANEL	
2	FILLER PANEL	
1	FILLER PANEL	
	FRONT VIEW	

Figure 29: DC RMS Example Configuration

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:

B   blogs.oracle.com/oracle

f   facebook.com/oracle

Y   twitter.com/oracle

O   oracle.com

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

Integrated Cloud Applications & Platform Services