Oracle® Communications DSR Cloud Benchmarking Guide





Oracle Communications DSR Cloud Benchmarking Guide, Release 9.0.1.0.0

F87798-03

Copyright © 2014, 2024, Oracle and/or its affiliates.

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

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

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

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

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

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

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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

Contents

1.1 Reference	ces	:
Cloud Dep	loyable DSR	
Benchmar	king Cloud Deployable DSR	
3.1 Infrastru	cture Environment	:
3.1.1 Ge	neral Rules for All Infrastructures	;
3.1.1.1	Hyper-Threading and CPU Over-Subscription	;
3.1.1.2	VM Packing Rate	;
3.1.1.3	Infrastructure Tuning	;
3.1.2 KV	M (QEMU)/Oracle X9-2 – Infrastructure Environment	;
3.1.2.1	Device Drivers	;
3.1.2.2	BIOS Power Settings	;
3.1.2.3	Guest Caching Modes	;
3.1.2.4	Memory Tuning Parameters	(
3.2 Benchma	ark Testing	;
3.2.1 DA	-MP Relay Benchmark	;
3.2.1.1	Topology	;
3.2.1.2	Message Flow	3-
3.2.1.3	Indicative Alarms or Events	3-
3.2.2 RE	AR Benchmark	3
3.2.2.1	Topology	3
3.2.2.2	Message Flow	3-
3.2.2.3		3-
3.2.3 Fu	ll Address Based Resolution (FABR - SDS) Capacity	3-
3.2.3.1	. 0	3-
3.2.3.2	9	3-
3.2.3.3		3-
	ll Address Based Resolution (FABR-UDR) Capacity	3-
3.2.4.1	. 65	3-
3.2.4.2	Message Flow	3-



	3.2.4.3 Indicative Alarms or Events	3-18
	3.2.5 vSTP MP	3-19
	3.2.5.1 vSTP MP Benchmarking	3-19
	3.2.6 Policy DRA (PDRA) Benchmarking	3-20
	3.2.6.1 Topology	3-21
	3.2.6.2 Message Flow	3-22
	3.2.6.3 Indicative Alarms or Events	3-23
	3.2.7 Diameter Security Application (DSA) Benchmarking	3-24
	3.2.7.1 Topology	3-25
Part	l Appendix	
Α	DSR VM Configurations	
В	VM Networking Requirements	
С	Summary of Benchmark Data Recommendations	
D	Detailed Infrastructure Settings	
Е	Small DSR VM Configuration	
F	24 vCPU Profile Testing on Oracle X9-2	



What's New in This Guide

This section introduces the documentation updates for Release 9.0.1.0.0.

Release 9.0.1.0.0 - F87798-03, October 2024

Updated the Disk(GB) value for the iDIH Application and iDIH Mediation in the Table A-1 table.

Release 9.0.1.0.0 - F87798-02, April 2024

Updated the host hardware for the following parameters in KVM (QEMU)/Oracle X9-2 – Infrastructure Environment section:

- QEMU
- libvirt
- API QEMU

Release 9.0.1.0.0 - F87798-01, October 2023

- Updated the values in Relay Performance Benchmarking table in Message Flow.
- Added new alarms and events information in DA-MP Relay Alarms or Events table in Indicative Alarms or Events.
- Updated the CPU Peak value in RBAR Performance Benchmarking table in Message Flow.
- Updated the SDS DB Details information in Topology.
- Updated the values in SDS DP Performance Benchmarking table in Message Flow.
- Added 22004 alarm information in SDS DP Alarms or Events table in Indicative Alarms or Events.
- Updated the values in FABR with UDR Performance Benchmarking table in Message Flow.
- Added new alarms and events information in FABR with UDR Alarms or Events table in Indicative Alarms or Events.
- Updated CPU Peak values and added RAM Peak values in Feature-wise vSTP MP Benchmarking table in vSTP MP Benchmarking.
- Updated DA-MP CPU Peak values and RAM Utilization Peak values in PDRA Performance Benchmarking table in Message Flow.
- Updated CPU Peak values and RAM Utilization Peak values in DSA Performance Benchmarking table in Topology.
- Added the following sections:
 - Small DSR VM Configuration
 - 24 vCPU Profile Testing on Oracle X9-2



Acronyms

The following tables provide information about the acronyms and the terminology used in this document.

Table Acronyms

	Book and the second			
Acronym	Description			
API	Application Programming Interface			
ARR	Application Route Rule			
ART	Application Route Table			
СМ	Counter Measure			
COTS	Commercial Off the Shelf			
CPU	Central Processing Unit			
DA-MP	Diameter Agent Message Processor			
DB	Database			
DP	Database Processor			
DSA	Diameter Security Application			
DSR	Diameter Signaling Router			
EIR	Equipment Identity Register			
ETG	Egress Throttle Group			
FABR	Full Address Based Resolution			
FQDN	Fully Qualified Domain Name			
GB	Gigabyte			
HDD	Hard Disk Drive			
HP	Hewlett Packard			
HSS	Home Subscriber Server			
HTTP	Hypertext Transfer Protocol			
ID	Identification			
IDIH	Integrated Diameter Intelligence Hub			
IMI	Internal Message Interface			
IMSI	International Mobile Subscriber Identity			
I/O	Input/Output			
IOP	Interoperability			
IoT	Internet of Things			
IP	Internet Protocol			
IPCAN	Internet Protocol Connectivity Access Network			
IPFE	IP Front End			
KPI	Key Performance Indicator			
KSM	Kernel Same-page Merging			
KVM	Kernel-based virtual machine			
LSI	Large Scale Integration			
LTE	Long Term Evolution			
MME	Mobility Management Entity			
MNO	Mobile Network Operator			
MP	Message Processor			



Table (Cont.) Acronyms

Acronym	Description
MPS	Messages Per Second
MSISDN	Mobile Station International Subscriber Directory Number
мтс	Machine Type Communication
NIC	Network Interface Card
NOAM	Network Operations, Alarms, Measurements
NE	Network Element
OAM	Operations, Administration, and Maintenance
OCDSR	Oracle Communications Diameter Signaling Router
ocsg	Oracle Communications Services Gatekeeper
онс	Oracle Help Center
PCRF	Policy and Charging Rules Function
PDRA	Policy Diameter Routing Agent
PRR	Peer Route Rule
PVSCSI	Paravirtual SCSI
RAM	Random Access Memory
RBAR	Range Based Address Resolution
SAS	Serial Attached SCSC
SBR	Session Binding Repository
SBR(b)	SBR – Binding database
SBR(s)	SBR – Session database
SBR(u)	SBR – Universal database
SCS/AS	Service Centralization and Continuity Application Server
SCSI	Small Computer System Interface
SDS	Subscriber Database Server
SFF	Small Form Factor
SGSN	Serving GPRS Support Node
SMS	Short Message Service
SOAM	System (nodal) Operations, Alarms, Measurements
SS7	Signaling System #7
SSD	Solid State Drive
THP	Transparent Huge Pages
TSA	Target Set Address
TTP	Troubleshooting Trigger Point
vSTP	virtual Signaling Transfer Point
UDR-NO	User Data Repository Network OAM & Provisioning
VM	Virtual Machine
VNF	Virtual Network Function
VoLTE	Voice over LTE
WAN	Wide Area Network
XMI	External Management Interface
XSI	External Signaling Interface



Table Terminology

Term	Description
1+1 Redundancy	For every 1, an additional 1 is needed to support redundant capacity. The specific redundancy scheme is not inferred. For example, active-active and active-standby.
Geo-Diverse	Refers to DSR equipment located at geographically separated sites/datacenters.
Geo-Redundant	A node at a geo-diverse location which can assume the processing load for another DSR signaling node(s).
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 For the vSTP MP the MPS excludes the equivalent SSNM status management messages.
N+K Redundancy	For every N, an additional K is needed to support redundant capacity. The specific redundancy scheme is not inferred. For example, active-active, active-standby.
Node	A DSR node is a DSR signaling node (SOAM and subtending topology), an NOAM node or an SDS node. A node is synonymous with the network element (NE).
Site	A specific geographic location or datacenter where DSR application is installed.



1

Introduction

The Oracle Communications Diameter Signaling Router (OCDSR or DSR) is deployable in the cloud as a Virtual Network Function (VNF). With DSR's added flexibility of being cloud deployable, operators must be able to manage the capacity and performance of the DSR in the cloud.

The document provides:

- Recommendations on performance tuning the DSR.
- Benchmark data from our labs.
- Information on the key metrics used to manage DSR performance and capacity.

1.1 References

The following reference documents are available on Oracle Help Center (OHC):

- Oracle Communications Diameter Signaling Router Alarms and KPIs Guide
- Oracle Communications Diameter Signaling Router Measurement Reference Guide
- Oracle Communications Diameter Signaling Router Cloud Installation Guide
- Oracle Communications Diameter Signaling Router Policy and Charging Application User Guide

2

Cloud Deployable DSR

DSR can be deployed on the following platforms:

- OpenStack based Cloud: It provides virtualization of DSR, but does not use a cloud manager and does not co-reside with other applications. This deployment is compact and cost-effective.
- KVM based environment: It provides full virtualization. It assumes the DSR resources are managed by a Commercial Off-the-Shelf (COTS) cloud manager and that the DSR can be one of many applications in the cloud.



Benchmarking Cloud Deployable DSR

This chapter is divided into the following sections:

- Infrastructure Environment
 This section provides details of the infrastructures used for the benchmark testing, including the hardware and software. It also describes key settings and attributes, and some recommendations on configuration.
- Benchmark section for each DSR server type
 Each DSR server type is treated independently for benchmarking. Each section describes the traffic setup, and the observed results. It also provides metrics and guidelines for assessing performance on any infrastructure.

Data Usage

This data is intended to provide guidance. Recommendations may need to be adapted to the conditions in a given operator's network. Each of the following sections include metrics that provide feedback on the running performance of the application.

When planning to deploy a DSR into any cloud environment, a few steps are recommended:

- Understand the initial deployment scenario for the DSR.
 - Which features are planned?
 - How much of what type of traffic?
 This may change once deployed, and the DSR can be grown or shrunk to meet the changing needs.
- Use the DSR Cloud Dimensioning tool to get an estimate of the types of DSR virtual servers needed and an initial estimate of the quantity of the virtual machines and resources. Oracle Sales Consultant can run this tool based on DSR requirements:
 - The tool allows for a very detailed model to be built of your DSR requirements, including:
 - * Required MPS by Diameter Application ID (S6a, Sd, Gx, Rx, so on).
 - * Required DSR applications such as Full Address Based Resolution (FABR), Range Based Address Resolution(RBAR),Policy DRA (PDRA), and any required sizing information such as the number of subscribers supported for each application.
 - * Any required DSR features such as Topology Hiding, Message Copy, IPSEC, or Mediation that can affect performance.
 - * Network-level redundancy requirements, such as mated pair DSR deployments, where one DSR needs to support full traffic, when one of the DSRs is unavailable.
 - * Infrastructure information, such as OpenStack or KVM, and Server parameters.
 - The tool then generates a recommended number of VMs for each of the required VM types.

Note:

These recommendations are just guidelines. Since the actual performance of the DSR can vary significantly based on the details of the infrastructure.

- Based on the initial deployment scenario, determine if additional benchmarking is warranted:
 - For labs and trials, there is no need to benchmark performance and capacity if the goal of the lab is to test DSR functionality.
 - If the server hardware is different from the hardware used in this document then the performance differences can likely be estimated using industry standard metrics. This is done by comparing single-threaded processor performance of the CPUs used in this document with respect to the CPUs used in the customer's infrastructure. This approach is most accurate for small differences in hardware (for instance, different clock speeds for the same generation of Intel processors) and least accurate across processor generations where other architectural differences such as networking interfaces could also affect the comparison.
 - It is the operator's decision to determine if additional benchmarking in the operator's infrastructure is desired. Here is a few things to consider when deciding:
 - * Benchmark infrastructure is similar to the operator's infrastructure, and the operator is satisfied with the benchmark data provided by Oracle.
 - * Initial turn-up of the DSR is handling a relatively small amount of traffic and the operator prefers to measure and adjust once deployed.
 - * Operator is satisfied with the high-availability and geo-diversity of the DSR, and is willing to risk initial overload conditions, and adjusts once the DSR is in production.
- If required, perform benchmark testing on the target cloud infrastructure. Perform benchmark only on those types of DSR servers required for the deployment.
 For example, if full address resolution is not planned, do not waste time benchmarking the SDS, SDS SOAM, or DPs.
 - When the benchmark testing is complete, observe the data for each server type, and compare it with the baseline used for the estimate from the **DSR Cloud Dimensioning** tool.
 - If the performance estimate for a given DSR function is X and the observed performance is Y, then adjust the performance for that DSR function to Y.
 - * Re-calculate the resources needed for deployment based on the updated values.
- Deploy the DSR.
- Monitor the DSR performance and capacity as described later in the document. As the network changes additional resources may be required. If needed, increase the DSR resources as described later in this document.

3.1 Infrastructure Environment

This section describes the infrastructure that was used for benchmarking. In general, the defaults or recommendations for hypervisor settings are available from the infrastructure vendors.

Whenever possible the DSR recommendations align with vendor defaults and recommendations. Benchmarking was performed with the settings described in this section. Operators may choose different values, better or worse performance compared to the

benchmarks might be observed. When recommendations other than vendor defaults or recommendations are made, additional explanations are included in the applicable section.

There is a sub-section included for each infrastructure environment used in benchmarking.

3.1.1 General Rules for All Infrastructures

3.1.1.1 Hyper-Threading and CPU Over-Subscription

All of the tests were conducted with Hyper-Threading enabled, and with a 1:1 subscription ratio for vCPUs in the hypervisor. The hardware used for the testing were dual-processor servers with 32 physical cores each (Oracle X9-2). Thus, each server had:

```
(2 CPUs) x (32 cores per CPU) x (2 threads per core) = 128 vCPUs
```

It is not recommended to use over-subscribed vCPUs (for instance 4:1) in the hypervisor. Not only is the performance lower, but it makes the performance more dependent on the other loads running on each physical server.

Turning off Hyper-Threading is also not recommended. There is a small increase in performance of a given VM without Hyper-Threading for a given number of vCPUs. But since the number of vCPUs for each processor drops in half without Hyper-Threading, the overall throughput for each server also drops almost by half.

The vCPU sizing for each VM is provided in the DSR VM Configurations section.



The recommended configuration is: Hyper-Threading is enabled with 1:1 CPU subscription ratio.

CPU Technology

The CPUs in the servers used for the benchmarking were the Oracle X9-2. Servers with different processors does give different results. In general there are the following issues when mapping the results of the benchmarking data in this document to other CPUs:

The per-thread performance of a CPU is the main attribute that determines VM performance. The number of threads is fixed in the VM sizing as shown in DSR VM Configurations section. A good metric for comparing the per-thread performance of different CPUs is the integer performance measured by the SPECint2006 (CINT2006) defined by SPEC.ORG.

The mapping of **SPECint2006ratios** to DSR VM performance ratios isn't exact, but it's a good measure to determine whether a different CPU is likely to run the VMs faster or slower than the benchmark results in this document.

Conversely CPU clock speeds are a relatively poor indicator of relative CPU performance. Within a given Intel CPU generation (v2, v3, v4, so on) there are other factors that affect per-thread performance, such as potential turbo speeds of the CPU in comparison with the cooling solution in a given server.

Comparing between Intel CPU generations, there is a generation over generation improvement of CPU throughput in comparison with the clock speed. This means that even a newer generation chip with a slower clock speed may run a DSR VM faster.



- The processors must have enough cores that a given VM can fit entirely into a NUMA node. Splitting a VM across NUMA nodes greatly reduces the performance of that VM. The largest VM size (refer DSR VM Configurations section) is 18 vCPUs. Thus, the smallest processor that should be used is a 9-core processor. Using processors with more cores typically makes it easier to pack VMs more efficiently into NUMA nodes but should not affect individual VM CPU-related performance otherwise.
- One caveat about CPUs with very high core counts is that the user must be aware of
 potential bottlenecks caused by many VMs contending for shared resources such as
 network interfaces and ephemeral storage on the server. These tests were run on relatively
 large CPUs (32 physical cores for each chip), and no such bottlenecks were encountered
 while running strictly DSR VMs. In clouds with VMs from other applications potentially
 running on the same physical server as DSR VMs, or in future processor generations with
 much higher core counts. This potential contention for shared server resources has to be
 watched closely.



The selected VM sizes should fit within a single NUMA node, for instance 9 physical cores for the VMs that required 18 vCPUs. Check the performance of the target CPU type against the benchmarked CPU using per-thread integer performance metrics.

3.1.1.2 VM Packing Rate

The DSR doesn't require or use CPU pinning. Thus, the packing of the DSR VMs onto the physical servers is under the control of OpenStack using the affinity or anti-affinity rules given in DSR VM Configurations. Typically, the VMs do not fit exactly into the number of vCPUs available in each NUMA node, leaving some un-allocated vCPUs. The ratio of the allocated to the unallocated vCPUs is the VMPacking Ratio. For instance, on a given server if 102 out of 128 vCPUs on a server were allocated by the OpenStack, that server would have a packing ratio of ~80%. The achieved packing in a deployment depends on a lot of factors, including the mix of large VMs (DA-MPs, SBRs) with the smaller VMs, and whether the DSR is sharing the servers with other applications that have a lot or large or small VMs.

When planning the number of physical servers required for an DRS a target packing ratio of 80% is a good planning number. A packing ratio of 100% is hard to achieve and may affect the performance numbers shown in the benchmarks. Some amount of server capacity is necessary to run the Host OS for the VMs. While performing functions such as interrupt handling, a packing ratio of 95% or lower is desirable.

Note:

When planning for physical server capacity a packing ratio of 80% is a good guideline. Packing ratios of greater than 95% might affect the benchmark numbers since there aren't sufficient server resources to handle the overhead of Host OSs.

3.1.1.3 Infrastructure Tuning

The following parameters should be set in the infrastructure to improve DSR VM performance. The instructions for setting them for a given infrastructure is including the *DSR Cloud Installation Guide*.



- Txqueuelen: The default of 500 is too small. Recommendation is to set this parameter to 120000.
 - Tuned on the compute hosts.
 - Default value of 500 is too small. Our recommendation is to set to 120000. This
 increases the network throughput of a VM.
- Ring buffer increase on the physical Ethernet interfaces: The default is too small. The recommendation is to set both receive and transmit values to 4096.
- Multiqueue: Multiqueue should be enabled on any IPFE VMs to improve performance.



Refer to instructions in the DSR Cloud Installation Guide.

3.1.2 KVM (QEMU)/Oracle X9-2 - Infrastructure Environment

There are a number of settings that affect performance of the hosted virtual machines. A number of tests were performed to maximize the performance of the underlying virtual machines for the DSR application.

Host Hardware

- Oracle Server X9-2
 - CPU Model: Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz
 - 2 CPUs
 - 32 physical cores per CPU
 - RAM: 768 GB
 - HDD: 3.8 TB of NVMe storage (with Software RAID-1 configured)
 - NIC
 - Oracle Quad Port 10G Base-T Adapter

Hypervisor

QEMU-KVM Version: QEMU 6.2.0, libvirt 8.0.0, API QEMU 8.0.0

3.1.2.1 Device Drivers

VirtIO is a virtualizing standard for network and disk device drivers where just the guest's device driver knows it is running in a virtual environment and cooperates with the hypervisor. This enables guests to get high performance network and disk operations and gives most of the performance benefits of para-virtualization.

Vhost-net provides improved network performance over Virtio-net by totally by passing QEMU as a fast path for interruptions. The vhost-net runs as a kernel thread and interrupts with less overhead providing near native performance. The advantages of using the vhost-net approach are reduced copy operations, lower latency, and lower CPU usage.



The VirtIO driver was used for Test Bed setting.

3.1.2.2 BIOS Power Settings

Typical BIOS power settings (hardware vendor dependent, see relevant infrastructure hardware vendor documentation for details) provide three options for power settings:

- Power Supply Maximum: The maximum power the available PSUs can draw.
- Allocated Power: The power is allocated for installed and hot pluggable components.
- Peak Permitted: The maximum power the system is permitted to consume.

Note:

Set to **Allocated Power** or equivalent for your Hardware vendor.

Disk Image Formats

The preferred disk image file formats available when deploying a KVM virtual machine:

 QCOW2: Disk format supported by the QEMU emulator that can expand dynamically and supports Copy-on-write.

QCOW2 provides a number of benefits, such as:

- Smaller file size, even on file systems which don't support holes (such as, sparse files)
- Copy-on-write support, where the image only represents changes made to an underlying disk image
- Snapshot support, where the image can contain multiple snapshots of the images history

Test Bed Setting: QCOW2

3.1.2.3 Guest Caching Modes

The operating system maintains a page cache to improve the storage I/O performance. With the page cache, write operations to the storage system are considered completed after the data has been copied to the page cache. Read operations can be satisfied from the page cache if the data requested is in the cache. The page cache is copied to permanent storage using fsync. Direct I/O requests bypass the page cache. In the KVM environment, both the host and guest operating systems can maintain their own page caches, resulting in two copies of data in memory.

The following caching modes are supported for KVM guests:

- Writethrough: I/O from the guest is cached on the host but written through to the physical medium. This mode is slower and prone to scaling problems. Best used for a small number of guests with lower I/O requirements. Suggested for guests that do not support a writeback cache (such as, Red Hat Enterprise Linux 5.5 and earlier), where migration is not needed.
- Writeback (Selected): With caching set to writeback mode, both the host page cache and the disk write cache are enabled for the guest. Due to this, the I/O performance for applications running in the guest is good, but the data is not protected in a power failure.

As a result, this caching mode is recommended only for temporary data where potential data loss is not a concern.

- None: With caching mode set to none, the host page cache is disabled, but the disk write cache is enabled for the guest. In this mode, the write performance in the guest is optimal because write operations bypass the host page cache and go directly to the disk write cache. If the disk write cache is battery-backed, or if the applications or storage stack in the guest transfer data properly (either through fsync operations or file system barriers), then data integrity can be ensured. However, because the host page cache is disabled, the read performance in the guest would not be as good as in the modes where the host page cache is enabled, such as write through mode.
- Unsafe: The host may cache all disk I/O, and sync requests from guest are ignored.

Caching mode None is recommended for remote NFS storage, because direct I/O operations (O_DIRECT) perform better than synchronous I/O operations (with O_SYNC). Caching mode None effectively turns all guest I/O operations into direct I/O operations on the host, which is the NFS client in this environment. Moreover, it is the only option to support migration.



For Test Bed Setting, set Caching Mode to Writeback.

3.1.2.4 Memory Tuning Parameters

Swappiness

The swappiness parameter controls the tendency of the kernel to move processes out of physical memory and onto the swap disk. Since disks are much slower than RAM, this can lead to slower response times for system and applications if processes are too aggressively moved out of memory.

- vm.swappiness = 0: The kernel swaps only to avoid an out of memory condition.
- vm.swappiness = 1: Kernel version 3.5 and over, as well as kernel version 2.6.32-303 and over; Minimum amount of swapping without disabling it entirely.
- vm.swappiness = 10: This value is recommended to improve performance when sufficient memory exists in a system.
- vm.swappiness = 60: Default
- vm.swappiness = 100: The kernel swaps aggressively



For Test Bed Setting, set vm. swappiness to 10.

Kernel Same Page Merging

Kernel Same-page Merging (KSM), used by the KVM hypervisor, allows KVM guests to share identical memory pages. These shared pages are usually common libraries or other identical, high-use data. KSM allows for greater guest density of identical or similar guest operating systems by avoiding memory duplication. KSM enables the kernel to examine two or more already running programs and compare their memory. If any memory regions or pages are identical, KSM reduces multiple identical memory pages to a single page. This page is then



marked copy-on-write. If the contents of the page is modified by a guest virtual machine, a new page is created for that guest.

This is useful for virtualization with KVM. When a guest virtual machine is started, it only inherits the memory from the host qemu-kvm process. Once the guest is running, the contents of the guest operating system image can be shared when guests are running the same operating system or applications. KSM allows KVM to request that these identical guest memory regions be shared.

KSM provides enhanced memory speed and utilization. With KSM, common process data is stored in cache or in main memory. This reduces cache misses for the KVM guests, which can improve performance for some applications and operating systems. Secondly, sharing memory reduces the overall memory usage of guests, which allows for higher densities and greater utilization of resources.

The following 2 services controls KSM:

- KSM Service: When the KSM service is started, KSM shares up to half of the host system's main memory. Start the KSM service to enable KSM to share more memory.
- KSM Tuning Service: The ksmtuned service loops and adjusts KSM. The ksmtuned service
 is notified by libvirt, when a guest virtual machine is created or destroyed



For Test Bed Setting, set KSM service to active and ensure ksmtuned service running on KVM hosts.

Zone Reclaim Mode

When an operating system allocates memory to a NUMA node, but the NUMA node is full, the operating system reclaims memory for the local NUMA node rather than immediately allocating the memory to a remote NUMA node. The performance benefit of allocating memory to the local node outweighs the performance drawback of reclaiming the memory. However, in some situations reclaiming memory decreases performance to the extent that the opposite is true. In other words, in these situations, allocating memory to a remote NUMA node generates better performance than reclaiming memory for the local node.

A guest operating system causes zone to reclaim in the following situations:

- When you configure the guest operating system to use huge pages.
- When you use KSM to share memory pages between guest operating systems.

Configuring huge pages and running KSM are both best practices for KVM environments. Therefore, to optimize performance in KVM environments, it is recommended to disable zone reclaim.

Note:

For Test Bed Setting, disable zone reclaim.

Transparent Huge Pages

Transparent huge pages (THP) automatically optimize system settings for performance. By allowing all free memory to be used as cache, performance is increased.



Note:

For Test Bed Setting, enable THP.

3.2 Benchmark Testing

The way the testing was performed and the benchmark test set-up is the same for each benchmark infrastructure. Each section describes the common set-up and procedures used to benchmark, and then the specific results for the benchmarks are provided for each benchmark infrastructure.

3.2.1 DA-MP Relay Benchmark

This benchmarking case illustrates conditions for an overload of a DSR DA MP.

3.2.1.1 Topology

The below figure illustrates the logical topology used for this testing. Diameter traffic is generated by an MME simulator and sent to an HSS simulator.




Figure 3-1 DA-MP Relay Testing Topology

The dsr.cpu utilization can be further increased to higher levels by means of configuration changes with DOC/CL1/CL2 discards set to 0 and multi-queuing enabled on all hosts. With this configuration, it must be noted that all the discards are at one step CL3 for all incoming and outgoing messages.



3.2.1.2 Message Flow

The following figure illustrates the Message sequence for this benchmark case.

Figure 3-2 DA-MP Relay Message Sequence




Table 3-1 Relay Performance Benchmarking on 16 DA-MPs DSR Setup

Scenario	Call Flow Model	DSR MPS Achieved	DA-MP Flavor	DA-MP Profile	Avg Msg Size	CPU Peak	RAM Utilizatio n Peak
Relay	100% Relay	288K	12 vCPU	30K_MPS	2.0 K	19%	35%
	Relay		(Regular)				
Relay (with Multique	100%	576K	12 vCPU	40K_MPS_FAB	2.0 K	36%	32%
enable, configuration set to DOC/CL1/CL2 discards set to 0 and multi queuing enabled on all hosts)	Relay		(Regular)	R			
Relay	100%	560K	18 vCPU	35K_MPS	2.0 K	28%	26%
	Relay		(Large)				

3.2.1.3 Indicative Alarms or Events

During benchmark testing the following alarms or events were observed when it reaches congestion.

Table 3-2 DA-MP Relay Alarms or Events

Number	Severity	Server	Name	Description
5007	Minor	IPFE	Out of Balance: Low	Traffic statistics reveal that an application server is processing lower than average load.
5008	Minor	IPFE	Out of Balance: High	Traffic statistics reveal that an application server is processing higher than average load and will not receive new connections.



Table 3-2 (Cont.) DA-MP Relay Alarms or Events

Number	Severity	Server	Name	Description
22008	Info	DA-MP	Orphan Answer Response Received	An answer response was received for which no pending request transaction existed resulting in the Answer message being discarded.
22201	Minor	DA-MP	MpRxAllRate	DA-MP ingress message rate threshold crossed.
22221	Minor	DA-MP	Routing MPS Rate	Message processing rate for this DA-MP is approaching or exceeding its engineered traffic handling capacity.
22225	Minor	DA-MP	MpRxDiamAllLen	DA-MP diameter average ingress message length threshold crossed.

3.2.2 RBAR Benchmark

Range Based Address Resolution (RBAR) is a DSR-enhanced routing application that allows the routing of Diameter end-to-end transactions based on Diameter Application ID, Command Code, Routing Entity Type, and Routing Entity Addresses (range and individual) as a Diameter Proxy Agent.

A Routing Entity can be:

- · A User Identity:
 - International Mobile Subscriber Identity (IMSI)
 - Mobile Subscriber Integrated Services Digital Network (Number) (MSISDN)
 - IP Multimedia Private Identity (IMPI)
 - IP Multimedia Public Identity (IMPU)
- An IP Address associated with the User Equipment:
 - IPv4 (based upon the full 32-bit value in the range of 0x00000000 to 0xFFFFFFFF)
 - IPv6-prefix (1 to 128 bits)
- A general-purpose data type: UNSIGNED16 (16-bit unsigned value)

Routing resolves to a Destination that can be configured with any combination of a Realm and Fully Qualified Domain Name (FQDN); Realm-only, FQDN-only, or Realm and FQDN.

When a message successfully resolves to a destination, RBAR replaces the destination information (Destination-Host and/or Destination-Realm) in the ingress message with the corresponding values assigned to the resolved destination and forwards the message to the (integrated) Diameter Relay Agent for egress routing into the network.

3.2.2.1 Topology

The following figure illustrates the logical topology used for this testing. Diameter traffic is generated by an MME simulator and sent to an HSS simulator.




Figure 3-3 RBAR Testing Topology

3.2.2.2 Message Flow

The following figure illustrates the Message sequence for this benchmark case.



Figure 3-4 DA-MP RBAR Message Sequence

HSS -ULR--ULR-

Table 3-3 RBAR Performance Benchmarking on 16 DA-MPs DSR Setup

Scenario	Call Flow Model	DSR MPS Achieved	DA-MP Flavor	DA-MP Profile	Avg Msg Size	CPU Peak	RAM Utilization Peak
RBAR	100% RBAR	256K	12 vCPU (Regular)	30K_MPS	2.0 K	29%	30%



3.2.2.3 Indicative Alarms or Events

During benchmark testing the following alarms or events were observed when it reaches into congestion.

Table 3-4 DA-MP RBAR Alarms or Events

Number	Severity	Server	Name	Description
22008	Info	DA-MP	Orphan Answer Response Received	An answer response was received for which no pending request transaction existed resulting in the Answer message being discarded.
22225	Minor	DA-MP	MpRxDiamAllLen	DA-MP diameter average ingress message length threshold crossed.

3.2.3 Full Address Based Resolution (FABR - SDS) Capacity

The FABR application adds a Database Processor (DP) server to perform database lookups with a user defined key (IMSI, MSISDN, or Account ID and MSISDN or IMSI). If the key is contained in the database, the DP returns the realm and FQDN associated with that key. The returned realm and FQDN can be used by the DSR Routing layer to route the connection to the desired endpoint. Since there is additional work done on the DA-MP to query the DP, running the FABR application has an impact on the DA-MP performance. This section contains the performance of the DA-MP while running FABR as well as benchmark measurements on the DP itself.



3.2.3.1 Topology




Figure 3-5 SDS DP Testing Topology

SDS DB Details

The SDS database was first populated with subscribers. This population simulates real-world scenarios likely encountered in a production environment and ensure the database is of substantial size to be gueried against.

- SDS DB Size: 300 million routing entities (150 million MSISDNs or 150 million IMSIs)
- AVP Decoded: User-Name for IMSI

SDS profile (Large) enhances the capacity of SDS FABR database to 780 million routing entities. The Large profile is defined in DSR VM Configurations based on the below 780 million entry configuration:

- 260 million subscribers having 2 IMSI, 1 MSISDN = 780 million routing entities
- IMSI = 15 bytes
- MSISDN = 10 bytes



3.2.3.2 Message Flow

Figure 3-6 SDS DP Message Sequence




Table 3-5 SDS DP Performance Benchmarking using one SDS DP

Scenario	Call Flow Model	DP MPS Achieved	DA-MP MPS Achieved	DA-MP Flavor	DA-MP Profile	Avg Msg Size	CPU Peak	RAM Utilization Peak
FABR	100% FABR	80K	160K	12 vCPU	30K_MPS	400	23%	28%
				(Regular)				

3.2.3.3 Indicative Alarms or Events

Table 3-6 SDS DP Alarms or Events

Number	Severity	Server	Name	Description
19814	Info	DA-MP	Communication Agent Peer has not responded to heartbeat	Communication Agent Peer has not responded to heartbeat.
19825	Major, Critical	DA-MP	Communication Agent Transaction Failure Rate	The number of failed transactions during the sampling period has exceeded configured thresholds.
19832	Info	DA-MP	Communication Agent Reliable Transaction Failed	Communication Agent Reliable Transaction Failed.
22004	Info	DA-MP	Maximum pending transactions allowed exceeded	Routing attempted to select an egress Diameter connection to forward a message but the maximum number of allowed pending transactions queued on the Diameter connection is reached.
22008	Info	DA-MP	Orphan Answer Response Received	An answer response was received for which no pending request transaction existed resulting in the Answer message being discarded.
22225	Minor	DA-MP	MpRxDiamAllLen	DA-MP diameter average ingress message length threshold crossed.



Table 3-6 (Cont.) SDS DP Alarms or Events

Number	Severity	Server	Name	Description
22606	Info	DA-MP	Database or DB connection error	FABR application received service notification indicating Database (DP) or DB connection (COM Agent) Errors (DP timeout, errors or COM Agent internal errors) for the sent database query.
31000	Critical	DA-MP	S/W Fault	Program impaired by s/w fault

3.2.4 Full Address Based Resolution (FABR-UDR) Capacity

The FABR is a DSR application that provides an enhanced DSR routing capability to enable network operators to resolve the designated Diameter server (IMS HSS, LTE HSS PCRF, OCS, OFCS, and AAA) addresses based on individual user identity addresses in the incoming Diameter request messages. It offers enhanced functionalities with User Data Repository (UDR), which is used to store subscriber data. FABR routes the message as a Diameter Proxy Agent based on request message parameter content.

FABR use the services of the Diameter Plug-In for sending and receiving Diameter messages from or to the network. It uses Communication Agent to interact with off-board data repository (UDR) for address resolution. This section contains the performance of the DA-MP while running FABR.



3.2.4.1 Topology

UDR-01 UDR-02 UDR-03 UDR-04 UDR-05 UDR-06 vUDR Pairs (Active/Standby/Spare) DAMP00 DAMP01 DAMP02 DAMP03 DAMP04 DAMP05 DAMP06 DAMP07 DAMP08 DAMP09 DAMP10 DAMP11 DAMP12 DAMP13 DAMP14 DAMP15

Figure 3-7 FABR with UDR Testing Topology

UDR DB Details

The UDR database was first populated with subscribers. This population simulates real-world scenarios likely encountered in a production environment and ensure the database is of substantial size to be queried against.

- UDR DB Size: Tested with 40 Million records
- AVP Decoded: User-Name for IMSI

Following UDR profile is used for benchmarking.

Table 3-7 UDR Profile

vCPU	RAM (GB)	HDD (GB)
18	70	400

16 DA-MPs DSR Setup



3.2.4.2 Message Flow

Figure 3-8 FABR with UDR Message Sequence




Table 3-8 FABR with UDR Performance Benchmarking on 16 DA-MPs DSR Setup

Scenario	Call Flow Model	DSR MPS Achieved	DA-MP Flavor	DA-MP Profile	Avg Msg Size		RAM Utilizatio n Peak
FABR + Relay	70% FABR + 30% Relay	288K (18K/MP)	12 vCPU (Regular)	30K_MPS	2.0 K	38%	36%

3.2.4.3 Indicative Alarms or Events

Table 3-9 FABR with UDR Alarms or Events

Number	Severity	Server	Name	Description
19825	Minor/Major/ Critical	DA-MP	Communication Agent Transaction Failure Rate	The number of failed transactions during the sampling period has exceeded the configured thresholds.
19832	Info	DA-MP	Communication Agent Reliable Transaction Failed	Failed transaction between servers result in normal maintenance actions, overload conditions, software failures, or equipment failures.
22008	Info	DA-MP	Orphan Answer Response Received	An answer response was received for which no pending request transaction existed resulting in the Answer message being discarded.
22201	Minor	DA-MP	MpRxAllRate	DA-MP ingress message rate threshold crossed.
22221	Minor	DA-MP	Routing MPS Rate	Message processing rate for this DA-MP is approaching or exceeding its engineered traffic handling capacity.
22225	Minor	DA-MP	MpRxDiamAllLen	DA-MP diameter average ingress message length threshold crossed.

Table 3-9 (Cont.) FABR with UDR Alarms or Events

Number	Severity	Server	Name	Description
22606	Info	DA-MP	Database or DB connection error	FABR application received service notification indicating Database (DP) or DB connection (COM Agent) Errors (DP timeout, errors or COM Agent internal errors) for the sent database query.

3.2.5 vSTP MP

The vSTP-MP server type is a virtualized STP that supports M2PA, M3UA, and TDM. It can be deployed either with other DSR functionality as a combined DSR or vSTP, or as a standalone virtualized STP without any DSR functionality.

3.2.5.1 vSTP MP Benchmarking

The following table describes the feature wise vSTP MP benchmarking.

Table 3-10 Feature-wise vSTP MP Benchmarking

Scenario	Call Flow Model	vSTP MPS	SS7-MP Flavor	CPU Peak	RAM Peak
		Achieved			
SFAPP + MNP + GTT	2K MNP + 2K SFAPP + 6K GTT	18K/MP	8 vCPU	22	55
SFAPP + MNP + GFLEX + GTT	2K MNP + 2K SFAPP + 1K GFLEX + 4K GTT	18K/MP	8 vCPU	25	49
TIF + GTT	5K MNP + 10K GTT	20K/MP	8 vCPU	28	39
vMNP + GTT	5K MNP + 10K GTT	20K/MP	8 vCPU	28	43
GFLEX + GTT	5K MNP + 10K GTT	20K/MP	8 vCPU	27	61
INPQ + GTT	5K MNP + 10K GTT	20K/MP	8 vCPU	40	52
GTT + MTP Routing with MTP screening (M2PA & M3UA)	16K MPS	16K/MP	8 vCPU	56	50
GTT + MTP Routing	20K MPS	20K/MP	8 vCPU	56	42
vEIR	5K	5K/MP	8 vCPU	17	49
Elynx (E1/T1 Card) – GTT Relay	10K TDM + 10K GTT	20K/MP	8 vCPU	29	52
ENUM	5K	5K/MP	8 vCPU	9	26
DNS	10K	10K/MP	8 vCPU	2	25
vSTP – Home SMS	MO-FSM AllowList + BlockList Traffic (10K + 10K)	10K/SS7 (2 MPs) 20K /Proxy MP	8 vCPU	46	45



Note:

- For ENUM, new vENUM-MP is introduced. vENUM sends messages to UDR over ComAgent interface.
- Default timer values are supported when vSTP is configured to operate at 10K MPS for each MP.
- When vSTP is configured to operate at 20K MPS, then the t1Timer to t5Timer values has to be updated. For information about the updated timer values, see MMI API Specification.
- MNP processing is equivalent to two messages and SFAPP processing is equivalent to four messages.

3.2.6 Policy DRA (PDRA) Benchmarking

The Policy DRA (PDRA) application adds two additional database components, the SBR(session) (SBR-s) and the SBR (binding) (SBR-b). The DA-MP performance was also measured since the PDRA application puts a different load on the DA-MP than either running Relay or FABR traffic. There are two sizing metrics when determining how many SBR-s or SBR-g server groups (for example, horizontal scaling units) are required. The first is the MPS traffic rate seen at the DA-MPs. This is the metric that is benchmarked in this document. The second factor is the number of bindings (SBR-b) or sessions (SBR-s) that can be supported. This session or binding capacity is set primarily by the memory sizing of the VM and is fixed at a maximum of 16 million per SBR from the DSR 8.3 release. The number of bindings and sessions required for a given network are customer dependent. But a good starting place for engineering is to assume:

- The number of bindings is equal to the number of subscribers supported by the PCRFs.
- The number of sessions is equal to number of subscribers times the number of IPCAN sessions required on average for each subscriber. For instance, a subscriber might have one IPCAN session for LTE, and one for VoLTE.



The number of sessions is equal to or greater than the number of bindings.



3.2.6.1 Topology

DAMP00 SBR-**S** SG1 (SBR01) DAMP01 DAMP02 SBR-**S** SG2 (SBR02) DAMP03 DAMP04 SBR-**\$** SG3 (SBR03) DAMP06 DAMP07 DAMP08 SBR-B SG5 (SBR05) DAMP09 SBR-B SG6 (SBR06) DAMP10 DAMP11 DAMP12 SBR-B SG8 (SBR08) DAMP13 DAMP14 DAMP15 16 DA-MPs DSR Setup

Figure 3-9 SBR Testing Topology



3.2.6.2 Message Flow




Figure 3-10 PDRA Message Sequence

The following table shows the call model used for the testing. The message distribution is Oracle's baseline benchmarking may differ significantly from customer distributions based on factors such as the penetration of LTE support in comparison with VoLTE support. The Traffic Details shows the configured PDRA options. For more details on these options, see *Oracle Communications Diameter Signaling Router Policy and Charging Application User Guide*.

Table 3-11 PDRA Test Call Model

Messages			Traffic Details		
Message	Count	Distribution	Message	Distribution	
CCR-I, CCA-I	1	7.14%	Gx with MSISDN Alternative Key,	100%	
			Gx Topology Hiding		



Table 3-11 (Cont.) PDRA Test Call Model

Messages			Traffic Details	Traffic Details		
Message	Count	Distribution	Message	Distribution		
CCR-U, CCA-U	3	21.42%	Gx Topology Hiding	100%		
CCR-T, CCA-T	1	7.14%	Gx Topology Hiding	100%		
Gx RAR, RAA	3	21.42%	Gx Topology Hiding	100%		
AAR, AAA Initial	2	14.29%	Rx Topology Hiding	100%		
STR, STA	2	14.29%	Rx Topology Hiding	100%		
Rx RAR, RAA	2	14.29%	Rx Topology Hiding	100%		

Table 3-12 PDRA Performance Benchmarking on 16 DA-MPs DSR Setup

Scenario	Call Flow Model	SBR MPS Achieved	DA-MP Flavor	DA-MP Profile	Avg Msg Size	DA-MP CPU Peak	RAM Utilizatio n Peak
Single Server group (1 SBR(s), 1 SBR(b))	100% PDRA	50K	12 vCPU (Regular)	30K_MPS	600	21%	23%
Single Server group (4 SBR(s), 4 SBR(b))	100% PDRA	200K	12 vCPU (Regular)	30K_MPS	648	25%	23%

3.2.6.3 Indicative Alarms or Events

Table 3-13 PDRA Alarms or Events

Number	Severity	Server	Name	Description
19814	Info	DA-MP	Communication Agent Peer has not responded to heartbeat	Communication Agent Peer has not responded to heartbeat.
19825	Major, Critical	DA-MP	Communication Agent Transaction Failure Rate	The number of failed transactions during the sampling period has exceeded configured thresholds.
19832	Info	DA-MP	Communication Agent Reliable Transaction Failed	Communication Agent Reliable Transaction Failed
22008	Info	DA-MP	Orphan Answer Response Received	An answer response was received for which no pending request transaction existed resulting in the Answer message being discarded.

Table 3-13 (Cont.) PDRA Alarms or Events

Number	Severity	Server	Name	Description
22328	Minor	DA-MP	IcRate	Connection ingress message rate threshold crossed
22704	Info	DA-MP	Communication Agent Error	Policy and Charging server to SBR server communication failure
22705	Info	DA-MP	SBR Error Response Received	Policy and Charging server received response from SBR server indicating SBR errors
22714	Info	SBR	SBR RAR Initiation Error	SBR encountered an error while processing PCA initiated RAR requests
22716	Info	SBR	SBR Audit Statistics Report	SBR Audit Statistics Report
22718	Info	DA-MP	Binding Not Found for Binding Dependent Session Initiate Request	Binding record is not found for the configured binding keys in the binding dependent session-initiation request message
22741	Info	DA-MP	Failed to route PCA generated RAR	RAA with Unable To Deliver (3002) error is received at PCA for the locally generated RAR
31232	Critical, Minor	DA-MP	HA Late Heartbeat Warning	High availability server has not received a message on specified path within the configured interval
31236	Major	DA-MP	HA Link Down	High availability TCP link is down

3.2.7 Diameter Security Application (DSA) Benchmarking

Diameter Security application (DSA) applies counter measures for ingress messages received from external foreign network and for egress messages sent to external foreign network. Different counter measure profiles can be created for different IPX or roaming partners by enabling or disabling counter measures individually for different IPX provider or roaming partner Diameter Peers.


Enabling Diameter Security Application (DSA) Benchmarking

DSA application is enabled on DA-MP and it uses vUDR to store context information.



3.2.7.1 Topology

Figure 3-11 DSA Testing Topology



The following stateful and stateless counter measure application configuration and the modes of operations used in benchmarking tests.

Table 3-14 Stateful and Statelss Counter Measures

Application Configuration Data		General Options Settings		
Table Name	Count of Configured Entries	Options	Values	
AppCmdCst_Config	2	Opcodes Accounting	Disabled	
AppIdWL_Config	1	Max. UDR Queries per Message	5	
AVPInstChk_Config	48	Max. Size of Application State	4800	
Foreign_WL_Peers_Cfg_Sets	14	Logging of Vulnerable Messages	Enabled	
MCC_MNC_List	11			
MsgRateMon_Config	1			
Realm_List	6			
Security_Countermeasure_Config	19			
SpecAVPScr_Config	1	Application Threads	•	
System_Config_Options	1	Request	6	
TimeDistChk_Config	2000	Answer	4	
TTL_Config	5	SbrEvent	4	
VplmnORCst_Config	1	AsyncEvent	2	
TimeDistChk_Country_Config	2			
TimeDistChk_Exception_List	0			
TimeDistChk_Continent_Config	15			
VplmnORCst_Config	1			
RealmIMSICst_Config	210			
Exception_Rule_Config	0			



Table 3-14 (Cont.) Stateful and Statelss Counter Measures

Application Configuration Data		General Options Settings			
Table Name	Count of Configured Entries	Options	Values		
All Exception Types Table IMSI_Exception_Config MCC_MNC_Exception_Config Origin_Host_Exception_Config Realm_Exception_Config VPLMN_ID_Exception_Config	0				

Note:

The following error is received during performance run, if the call rate is more than 1.7k in each MP DSA:

UDR Internal Error: Create record failed. Error Code = SendError

This is caused due to comagent connection getting timeout due to ttl expired.

Communication Agent Reliable Transaction Failed} .. GN_INFO/INF Failure reason = Time to live limit exceeded

To avoid this, run the following commands from Active DSR NOAM before running performance traffic:

iset -fvalue=400 ComAgtConfigParams where "name='IntraNe Maximum Timeout Value'" iset -fvalue=3 ComAgtConfigParams where "name='Maximum Number Of Retries'"

Table 3-15 DSA Performance Benchmarking on 2 DA-MPs DSR Setup

Counter Measure (CM)	Call Flow Model	DSR MPS Achieved (Per DA- MP)	DA-MP Flavor	UDR Flavor	DA-MP Profile	Avg Msg Size	CPU Peak	RAM Utilizatio n Peak
Previous_Locatio n_Check	15% Vulnerable and 85% Non Vulnerable traffic	9.2K	12 vCPU (Regular)	18vCPU	30K_MPS	2.0 K	29%	63%
Time_Distance_ Check	15% Vulnerable and 85% Non Vulnerable traffic	9.8K	12 vCPU (Regular)	18vCPU	30K_MPS	2.0 K	34%	63%
Source_Host_Va lidation_Hss	15% Vulnerable and 85% Non Vulnerable traffic	9.6K	12 vCPU (Regular)	18vCPU	30K_MPS	2.0 K	34%	62%

Table 3-15 (Cont.) DSA Performance Benchmarking on 2 DA-MPs DSR Setup

Counter Measure (CM)	Call Flow Model	DSR MPS Achieved (Per DA- MP)	DA-MP Flavor	UDR Flavor	DA-MP Profile	Avg Msg Size	CPU Peak	RAM Utilizatio n Peak
Source_Host_Va lidation_Mme	15% Vulnerable and 85% Non Vulnerable traffic	10K	12 vCPU (Regular)	18vCPU	30K_MPS	2.0 K	33%	63%
Message_Monito ring_Rate	15% Vulnerable and 85% Non Vulnerable traffic	10K	12 vCPU (Regular)	18vCPU	30K_MPS	2.0 K	30%	63%
Session_Integrit y_Validation_Ch k	15% Vulnerable and 85% Non Vulnerable traffic	10K	12 vCPU (Regular)	18vCPU	30K_MPS	2.0 K	39%	63%
All Stateful	15% Vulnerable and 85% Non Vulnerable traffic	7.02K	12 vCPU (Regular)	18vCPU	30K_MPS	2.0 K	15%	30%
All Stateful + All stateless	15% Vulnerable and 85% Non Vulnerable traffic	5.25K	12 vCPU (Regular)	18vCPU	30K_MPS	2.0 K	14%	31%



Appendix

This section provides information about the following:

- Summary of VM configurations
- Networking characteristics of different VMs
- DSR Benchmarking values on KVM/Oracle X9-2 server



DSR VM Configurations

The following table provides a summary of the VM configurations used for the benchmarking data and the affinity rules for deploying those VMs. Using VM sizes different from these tested values may give unexpected results because the application profiles are tuned to this number of vCPUs and memory sizes.

Table A-1 VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redunda ncy Models	Affinity/Placement Rules (Per Site)	Notes
DSR NOAM (Regular)	4	14	120	1 Pair	Active/ Standby	2 VMs per DSR network in any site. VMs to be deployed on separate servers.	
DSR NOAM (Large)	8	14	120	1 Pair	Active/ Standby	2 VMs per DSR network in any site. VMs to be deployed on separate servers.	It is recommended to use large NOAM profile if the deployment is more than 32 C level servers. If SOAM profile is large, NOAM profile must be large. In large profiles, the scheduled two measurement exports run in parallel.
DSR SOAM (Regular)	4	14	120	1 Pair per DSR NF	Active/ Standby or Active/ Standby/ Spare	2 VMs per site. VMs to be deployed on separate servers.	Redundancy model Active/Standby/Spare model is used for PCA mated-pair deployments. For all other deployments Active or Standby model is used.

Table A-1 (Cont.) VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redunda ncy Models	Affinity/Placement Rules (Per Site)	Notes
DSR SOAM (Large)	8	14	120	1 Pair per DSR NF	Active/ Standby or Active/ Standby/ Spare	2 VMs per site. VMs to be deployed on separate servers.	Redundancy model Active/Standby/Spare model is used for PCA mated-pair deployments. For all other deployments Active or Standby model is used. If the diameter connections are higher than 4K, it is recommended to use SOAM large profile. If NOAM profile is large, SOAM profile must be large. In large profiles, the scheduled two measurement exports will run in parallel.
DA-MP (Regular)	12	16	120	32 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss.	The limit of 32 is the combined total of DA-MPs, DA-MPs with IWF, and DA-MP with EIR. The vSTP MPs do not count against this 32. Cannot max out all types in one DSR (for instance 32 DA-MPs AND 32 vSTPs). Any solution using more than 500 ART (Application Route Tables)/ARR (Application Route Rules)+PRR (Peer Route Rules) beyond 20k please use the below profile (DAMP w/IWF) which is with 24GB RAM.



Table A-1 (Cont.) VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redunda ncy Models	Affinity/Placement Rules (Per Site)	Notes
DA-MP (Large)	18	24	120	32 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss.	The limit of 32 is the combined total of DA-MPs, DA-MPs with IWF, and DA-MP with EIR. The vSTP MPs do not count against this 32. Cannot max out all types in one DSR (for instance 32 DA-MPs and 32 vSTPs). Any solution using more than 500 ART (Application Route Tables)/ARR (Application Route Rules) + PRR (Peer Route Rules) beyond 20k please use the below profile (DAMP w/IWF) which is with 24GB RAM.
vSTP MP	8	12	120	32 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss.	The vSTP MPs do not count against the 32 DA-MP limits in a single OCDSR node, so a DSR can have up to 32vSTP MPs. Cannot max out all types in one DSR (for instance 32 DA-MPs and 32 vSTPs).
HomeSMS C Service MP	8	12	120	32 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss.	The Service MPs do not count against the 32 DA-MP limits in a single OCDSR node, so a DSR can have up to 32service MPs. Cannot max out all types in one DSR (for instance 32 DA-MPs and 32 service).
vENUM Service MP	8	12	150	32 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss.	The vENUM MPs do not count against the 32 DA-MP limits in a single OCDSR node, so a DSR can have up to 32 ENUM MPs. Cannot max out all types in one DSR (for instance 32 DA-MPs and 32 service).



Table A-1 (Cont.) VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redunda ncy Models	Affinity/Placement Rules (Per Site)	Notes
IPFE	6	16	120	2 pairs per DSR NF	Active/ Standby	Each VM in a pair must be deployed on separate server.	Deployed in pairs. Max 2 pairs (4VMs).
SBR(s)	12	25	120	8 Server Groups per SBR(b)	Active/ Standby/ Spare	Active/Standby VMs to be deployed on separate servers,Spare is typically at another geographic location for geo-redundancy.	Can be either Active/ Standby/Spare or Active/Standby depending on customer geo- redundancy requirements.
SBR(b)	12	32	120	8 Server Groups per SBR(b)	Active/ Standby/ Spare	Active/Standby VMs to be deployed on separate servers,Spare is typically at another geographic location for geo-redundancy.	Can be either Active/ Standby/Spare or Active/Standby depending on customer geo- redundancy requirements.
SBR(u)	12	24	120	64 Server Groups per SBR(b	Active/ Standby/ Spare	Active/Standby VMs to be deployed on separate servers,Spare is typically at another geographic location for geo-redundancy.	Can be either Active/ Standby/Spare or Active/Standby depending on customer geo- redundancy requirements.
SDS NOAM (Regular)	4	32	350	1 Pair per Network	Active/ Standby	Anti-affinity between the Active/Standby VMs	Active/Standby. An optional "Disaster Recovery" SDS is supported that would typically be located at a different data center to provide georedundancy.
Query Server (Regular)	4	32	350	1 per SDS NOAM	N/A since non- redundant	Non, non-redundant	Optional 1 per site. Can have one for the primary SDS-NOAM and one for the Disaster Recovery SDS-NOAM
SDS SOAM (Regular)	4	12	175	1Pair per DSR NF	Active/ Standby	2 VMs per site. VMs to be deployed on separate servers.	
SDS DP (Regular)	6	10	175	10 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss.	To be evenly distributed across servers to minimize capacity loss.



Table A-1 (Cont.) VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redunda ncy Models	Affinity/Placement Rules (Per Site)	Notes
SDS NOAM (Large)	4	128	890	1 Pair per Network	Active/ Standby	Anti-affinity between the Active/Standby VMs	Active/Standby. An optional "Disaster Recovery" SDS is supported that would typically be located at a different data center to provide georedundancy.
Query Server (Large)	4	128	890	1 per SDS NOAM	N/A since non- redundant	Non, non-redundant	Optional 1 per site. Can have one for the primary SDS-NOAM and one for the Disaster Recovery SDS-NOAM.
SDS SOAM (Large)	4	64	450	1 Pair per DSR NF	Active/ Standby	2 VMs per site. VMs to be deployed on separate server.	Supports
SDS DP (Large)	24	64	450	10 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss.	To be evenly distributed across servers to minimize capacity loss.
UDR NO (Small)	6	16	220	n (Active,Sta ndby)	Active/ Standby/ Spare	Active/Standby/Spare VMs to be deployed on separate servers, Spare is typically at another geographic location for geo- redundancy	Redundancy model Active/Standby/Spare model is used. Active/ Standby on Site 1 and Spare on Site 2. UDR is scaled by adding UDR NOs. The Standby UDR NO also receives query traffic from STP-MP and DA- MP.
UDR NO (Medium)	12	32	400	n (Active,Sta ndby)	Active/ Standby/ Spare	Active/Standby/Spare VMs to be deployed on separate servers, Spare is typically at another geographic location for geo- redundancy	Redundancy model Active/Standby/Spare model is used. Active/ Standby on Site 1 and Spare on Site 2. UDR is scaled by adding UDR NOs. The Standby UDR NO also receives query traffic from STP-MP and DA- MP.



Table A-1 (Cont.) VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redunda ncy Models	Affinity/Placement Rules (Per Site)	Notes
UDR NO (Regular)	18	70	400	n (Active,Sta ndby)	Active/ Standby/ Spare	Active/Standby/Spare VMs to be deployed on separate servers, Spare is typically at another geographic location for geo- redundancy	Redundancy model Active/Standby/Spare model is used. Active/ Standby on Site 1 and Spare on Site 2. UDR is scaled by adding UDR NOs. The Standby UDR NO also receives query traffic from STP-MP and DA- MP.
UDR NO (Large)	32	120	800	n (Active,Sta ndby)	Active/ Standby/ Spare Active/Standby/Spare VMs to be deployed on separate servers, Spare is typically at another geographic location for georedundancy		Redundancy model Active/Standby/Spare model is used. Active/ Standby on Site 1 and Spare on Site 2. UDR is scaled by adding UDR NOs. The Standby UDR NO also receives query traffic from STP-MP and DA- MP.
UDR NO (Extra Large)	56	256	800	n (Active,Sta ndby)	Active/ Standby/ Spare	e/ Active/Standby/Spare Redundancedby/ VMs to be deployed on Active/Stand	
VNFM	8	10	80	NA	N/A since non- redundant	None, non-redundant	VNF Manager deployment
iDIH Applicatio n	4	8	70	1 per Site	N/A since non- redundant	None, non-redundant	Optional component for Diameter traffic monitoring
iDIH Mediation	4	8	70	1 per Site	N/A since non- redundant	None, non-redundant	Optional component for Diameter traffic monitoring
iDIH Database	4	8	120GB + 100GB (ephemera I)	1 per Site	N/A since non- redundant	None, non-redundant	Optional component for Diameter traffic monitoring



B

VM Networking Requirements

This section gives information on the networking characteristics of the different VMs. The traffic is broken down into signaling traffic handled on the XMI network, and OAM traffic carried on the IMI and XMI networks.

The Diameter Traffic requirements on the XSI networks can be calculated from the MPS. Treating the OCDSR as black box, this network traffic is simply the average Diameter message size (for requests and answers) times the MPS rate for the OCDSR node. The complication is that some Diameter traffic is likely to go through both an ingress DA-MP and an egress DA-MP. The most conservative consumption is that any ingress message is equally likely to go out any of the DA-MPs. Thus, if a DSR has X DA-MPs, and Y total MPS per DA-MP, the average Diameter signaling traffic through a DA-MP is:

```
((Average Diameter message size including IP overhead) * Y) * (1+ ((2X-1)/X))
```

As an example, if the average Diameter message size is 2,000 bytes including overhead, the overall DSRMPS is 10000 MPS, and the number of DA-MPs is three, the calculation would be:

```
(2,000 \text{ bytes} * 8 \text{ bits/byte} *10,000 \text{ MPS}) * (1+ (2*3 DA-MPs) -1)/(3 DA-MPs)) = (160,000 \text{ kb/s}) * (2.66) = 426,666 \text{ kb/s per DA-MP}
```

For the MP types other than the DA-MPs simply substitute the average size of signaling types, for instance SS7 messages for the vSTP MP. Since typically SS7 messages are much smaller than Diameter messages (for instance ~200 bytes for SMS), the vSTP MP bandwidth is much smaller than the DA-MP bandwidth.

The OAM traffic on the VMs can be much more variable since it's dependent to customerspecific usage patterns such as the number of reports requested and the number of periodic activities (backups and restores). The notes for each VM type give some background on the network impacts of these customer-driven activities.

Table B-1 VM Networking Utilization Characteristics

VM Name	Networks Supported	Management Networks (Gb/s)	Traffic Networks (Gb/s)	Notes
DSR NOAM DSR SOAM	XMI IMI	1	N/A	Activities such as backups can generate higher network utilization but runs at the rate of the bandwidth available since they are not real-time activities.
DA MP DA MP w/IWF vSTP MP	XMI IMI XSI	0.2	MPS Dependent	See explanation above for how to calculate the signaling network traffic.

Table B-1 (Cont.) VM Networking Utilization Characteristics

VM Name	Networks Supported	Management Networks (Gb/s)	Traffic Networks (Gb/s)	Notes
IPFE	XMI IMI XSI	0.2	MPS Dependent	The peak networking capacity supported by the IPFE is 3.2 Gb/s. Typically, the IPFE is deployed only on the ingress (towards clients such as MMEs) side of the DA-MP, so the total traffic through the IPFE is ½ the total bandwidth of the DA-MPs.
SBR(s) SBR(b) SBR(u)	XMI	1.0	N/A	The given OAM bandwidth is for routine operations. Some recovery operations such as synchronizing the database between the active and standby servers after a prolonged disconnection can consume an order of magnitude or more of network bandwidth. The required amount of bandwidth for these recovery operations is very dependent on customer-factors such as number of subscribers, the MPS rate, and the amount of networking downtime.
SDS NOAM	XMI IMI	1.0	N/A	The maximum bandwidth required by the SDS NOAM is determined primarily by the provisioning rate from external customer systems along with the size of the customer records.
DP SOAM	XMI IMI	1.0	N/A	All of the subscriber data provisioned at the SDSNOAM is passed down to each DP SOAM, which then distributes the data to any attached DPs.
DP	XMI IMI	1.0	N/A	The DP receives writes of new subscriber records from the SOAM, and database queries from the DA-MPs.
Query Server	XMI IMI	1.0	N/A	The Query Server is synchronized to the changes in the SDS NOAM. In addition, there is some network traffic due to customer search requests, but this traffic is small compared to the synchronization traffic.
UDR NO	XMI IMI XSI	1.0	N/A	UDR NO receives internal query from STP MP and DAMP.

The following table shows some guidelines for mapping the logical OCDSR networks (XMI, IMI, so on) to interfaces. There is nothing fixed about these assignments in the application, so they can be assigned as desired if the customer has other requirements driving interface assignment.

Table B-2 Typical OCDSR Network to Device Assignments

VM Name	OAM (XMI)	Local (IMI)	Signaling A (XSI1)	Signaling B (XSI2)	Signaling C (XSI3)	Signaling ()	Signaling D (XSI6)	Replicati on (SBR Rep)	DIH Internal
DSR NOAM	eth0	eth1							



Table B-2 (Cont.) Typical OCDSR Network to Device Assignments

VM Name	OAM (XMI)	Local (IMI)	Signaling A (XSI1)	Signaling B (XSI2)	Signaling C (XSI3)	Signaling ()	Signaling D (XSI6)	Replicati on (SBR Rep)	DIH Internal
DSR SOAM	eth0	eth1							
DA-MP	eth0	eth1	eth2	eth3	eth4		eth17	eth18	
IPFE	eth0	eth1	eth2	eth3	eth4		eth17		
SBRB	eth0	eth1						eth2	
SBRS	eth0	eth1						eth2	
vSTP	eth0	eth1	eth2	eth3	eth4		eth17		
UDRNO	eth0	eth1	eth2	eth3	eth4		eth17		
SDS NOAM	eth0	eth1							
DP	eth0	eth1							
Query Server	eth0	eth1							



C

Summary of Benchmark Data Recommendations

The information shown in the following table is a summary of the benchmark data described throughout the document. This data is intended to provide guidance and is based solely on the observed results from the test setups described in this document. Recommendations may need to be adapted to the conditions in a given operator's cloud, such as differences in traffic patterns, feature utilization patterns, and infrastructure differences.

Table C-1 Benchmark Data Summary

Benchmark Run	Openstack/KVM
Application Software	DSR 9.0 (running Oracle Linux 8.x)
Host VM	OpenStack Train, KVM
HW	Oracle Server X9-2
VM Profiles/Flavors	DSR VM Configurations

Table C-2 Recommended Maximum Engineering Targets

VM Name	VM Purpose	Recommended Maximum Engineering Targets	
		Unit	Quantity
DSR NOAM	Network Operation, Administration, Maintenance (and Provisioning)	VM	1+1
DSR SOAM	Site (node/Network Element) Operation,Administration,Maintenance (and Provisioning)	VM	1+1
DA MP (Relay) (Regular Profile)	Diameter Agent Message Processor	MPS	18,000
DA MP (Relay) (Regular Profile) configuration set to DOC/CL1/CL2discards set to 0 and multi queuing enabled on all hosts	Diameter Agent Message Processor	MPS	36,000
DA MP (Relay) (Large Profile)	Diameter Agent Message Processor	MPS	35,000
DA MP (RBAR)	Diameter Agent Message Processor	MPS	16,000
DA MP (FABR with UDR)	Diameter Agent Message Processor	MPS	18,000
			(70% FABR and 30% Relay)
DSA (All Stateful)	Diameter Agent Message Processor	MPS	7.02K
DSA (All Stateful + All Stateless)	Diameter Agent Message Processor	MPS	5.25K
vSTP MP	Virtual STP for M3UA and M2PA message Processing	MPS	20,000
vSTP w/ EIR	Virtual STP Message processor with EIR application	MPS	10000

Table C-2 (Cont.) Recommended Maximum Engineering Targets

VM Name	VM Purpose	Recommended Maximum Engineering Targets		
		Unit	Quantity	
EIR + GTT	Virtual STP Message processor with EIR application and GTT	MPS	5000 EIR + 8000 GTT	
SBR(s) (Single Server Group)	Subscriber Binding Repository (session) for Policy DRA	Diameter sessions	16,000,000	
		MPS	50,000	
SBR(s) (4 Server Groups)	Subscriber Binding Repository (session) for Policy DRA	Diameter sessions	64,000,000	
		MPS	200,000	
SBR(b) (Single Server Group)	Subscriber Binding Repository (binding) for Policy DRA	Diameter sessions	16,000,000	
		MPS	50,000	
SBR(b) (4 Server Groups)	SR(b) (4 Server Groups) Subscriber Binding Repository (binding) for Policy DRA		64,000,000	
		MPS	200,000	
SDS SOAM	Database Processor Site (node) Operation, Administration, Maintenance for address resolution and subscriber location functions	VM	1+1	
SDS DP	Database Processor for address resolution and subscriber location functions	MPS requiring DP lookups (usually 50% of FABR traffic)	80,000	
SDS (NOAM)	Subscriber Database Processor for address resolution and subscriber location functions	Routing Entities (SDS Large Profile)	780 million (260 million subscribers having 2 IMSI, 1 MSISDN)	
		Routing Entities (SDS Regular Profile)	300 million Routing Entities (150 M MSISDNs/150 M IMSIs)	
		Provisioning TPS	200	
SDS Query Server	Allows customers to query FABR subscriber data via a MySQL interface	N/A	N/A	



D

Detailed Infrastructure Settings

Table D-1 Detailed Infrastructure Settings

Attribute	KVM/Oracle X9-2
Model	Oracle Server X9-2
Processor Type	Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz
vCPUs	2 CPUs (32 physical cores per CPU)
RAM	768 G
CPU Cache Memory	L1d Cache: 48K
	L1i Cache: 32K
	L2 Cache: 1280K
	L3 Cache: 49152K
Number and Type of NICs	Oracle Quad Port 10G Base-T Adapter
BIOS Power Settings	Power Supply Maximum: Maximum power the available PSUs can draw Allocated Power: Power allocated for installed and hot pluggable components
	Peak Permitted: Maximum power the system is permitted to consume (set to Allocated Power)
HDD	3.8 TB of NVMe storage (with Software RAID-1 configured)



Е

Small DSR VM Configuration

Many customers do not need the capacity provided by even a single pair of standard size DSR DA-MP VMs. The fixed configuration in this section is for customers who need the following configuration:

- Relay or RBAR applications
- 6k MPS or less (assuming an infrastructure equal to or better than the X5-2 processors used for the benchmark tests)
- 5K or less RBAR entries
- No IPFE

This configuration should be run with the 6k DA-MP profile.

Table E-1 Minimal System VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redundanc y Models	Affinity/Placement Rules (Per Site)	Notes	
DSR NOAM	4	8	120	1 Pair	Active/ Standby	2 VMs for each DSR network in any site.	Two NOAMs are always required to	
DSR SOAM	4	8	120	1 Pair per DSR NF	Active/ Standby	VMs to be deployed on separate servers if possible.	support upgrades, but they can be on the same server if only one server is available.	
DA-MP	4	8	120	1 or 2 per DSR NF	1 or 1+1	The 1+1 configuration should have the DA-MPs on different servers.	There is a little value in having redundant DA-MPs (1+1) if they are on the same server, since the server is the thing that is most likely to fail.	

Topology

Figure E-1 vDSR Setup




Table E-2 Performance Benchmarking Figures (Small DSR VM Profile)

Scenario	Call Flow Model	DSR MPS Achieved	DA-MP Flavor	DA-MP Profile	Avg Msg Size	CPU Peak	RAM Utilizatio n Peak
Relay	100% Relay	12K (6K/MP)	4 vCPU	6K_MPS	2.0 K	25%	42%
RBAR	100% RBAR	12K (6K/MP)	4 vCPU	6K_MPS	2.0 K	35%	35%



F

24 vCPU Profile Testing on Oracle X9-2

This section provided the DSR Benchmarking figures on KVM/Oracle X9-2 server with 24 vCPU Profile.

Topology

Figure F-1 Topology for KVM/Oracle X9-2 server with 24 vCPU Profile




Table F-1 Performance Benchmarking on Oracle X9-2

Scenario	Call Flow Model	DSR MPS Achieved	DA-MP Flavor	SDS DP Flavor	DA-MP Profile	Avg Msg Size	CPU Peak	RAM Utilizatio n Peak
Relay	100% Relay	720K (45K/MP)	24 vCPU	24 vCPU	40K_MPS_FAB R	2.0 K	26%	24%
FABR + Relay	70% FABR + 30% Relay	664K (41.5K/MP)	24 vCPU	24 vCPU	40K_MPS_FAB R	2.0 K	36%	23%