

Oracle Communications Diameter Signaling Router Cloud Benchmarking Guide

Release 8.1.X

ORACLE COMMUNICATIONS | APRIL 2021



Table of Tables

Table 1: Acronyms	1
Table 2: Terminology	1
Table 3: Virtual Network Interrupt Coalescing and SplitRX Mode	12
Table 4: Power Management Profiles	13
Table 5: Virtualization Performance by Processor	13
Table 6: DA-MP Test Set-up	15
Table 7: DA-MP Alarms/Events	17
Table 8: DA-MP Utilization Metrics	18
Table 9: DA-MP Connection Metrics	19
Table 10: SDS DP Message Details	21
Table 11: SDS DP Alarms/Events	24
Table 12: SDS DP Utilization Metrics	25
Table 13: DP SOAM Metrics	26
Table 14: SS7 MP Message Detail	27
Table 15: SS7 MP Alarms/Events	28
Table 16: SS7 MP Metrics	28
Table 17: PDRA Test Call Model	32
Table 18: SBR (b) Alarms/Events	35
Table 19: Session SBR (b) Blades Metrics	35
Table 20: Binding SBR (b) Server Metrics	36
Table 21: Network OAM Metrics	37
Table 22: System OAM Metrics	38
Table 23: IPFE Throughput	39
Table 24: IPFE Metrics	41
Table 25: VM Configurations and Affinity Rules	43
Table 26: VM Disk Utilization Characteristics	46
Table 27: VM Networking Utilization Characteristics	48
Table 28: Typical ODCSR Network to Device Assignments	50
Table 29: Benchmark Data Summary	51
Table 30: Detailed Infrastructure Settings	53
Table 31: Minimal System VM Configurations and Affinity Rules	54

Table of Figures

Figure 1 – DA-MP Testing Topology	15
Figure 2 – DA-MP Message Sequence	15
Figure 3 – DA-MP CPU Occupancy	16
Figure 4 – DA-MP Round-Trip Latency vs. Message Rate	17
Figure 5 - SDS DP Testing Topology	20
Figure 6 - SDS DP Message Sequence	21
Figure 7 – DA-MP CPU Utilization (%) vs. MPS Running FABR Traffic	22
Figure 8 – DA-MP Latency vs. MPS Running FABR Traffic	23
Figure 9 - DP CPU Util vs Traffic Rate (DP Queries)	24
Figure 10 - SS7 MP Testing Topology	27
Figure 11 - SS7 MP Message Flow	27
Figure 12 - SBR Testing Topology	30
Figure 13 - PDRA Message Sequence	31
Figure 14 – DA-MP CPU Utilization (%) vs. MPS Running PDRA Traffic	33
Figure 15 – DA-MP Latency vs. MPS Running PDRA Traffic	34
Figure 16: SBR (b) CPU Occupancy vs. Message Rate	34
Figure 17 – IPFE on Ingress Side Only	39
Figure 18 – IPFE on both Ingress and Egress Sides	40

Acronyms

Table 1: Acronyms

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

Terminology

Table 2: 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 (e.g. active-active, 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	<p>A measure of the DSR Diameter message processing volume in messages per second. For this measure, a message is defined as:</p> <ol style="list-style-type: none"> 1. 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. 2. 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. <p>Messages excluded from this measure are:</p> <ul style="list-style-type: none"> • 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 <p>For the vSTP MP the MPS excludes the equivalent SSNM status management messages.</p>
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 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.

References

1. Performance Best Practices for VMware vSphere® 6.0
(<http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmware-perfbest-practices-vsphere6-0-white-paper.pdf>)
2. DSR Alarms, KPIs, and Measurements – Available at Oracle.com on the Oracle Technology Network (OTN)
3. DSR Cloud Deployable Installation Guide - Available at Oracle.com on the Oracle Technology Network (OTN)
4. Policy and Charging Application User's Guide – Available at Oracle.com on the Oracle Technology Network (OTN)

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.

This document focuses on:

- How to benchmark DSR performance and capacity in a cloud deployed DSR
- Provides recommendations on performance tuning the DSR
- Provides benchmark data from our labs
- Provides information on the key metrics used to manage DSR performance and capacity
- Provides recommendations on how to use the data obtained from the metrics

About Cloud Deployable DSR

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

- » Bare-metal and hybrid (mixture of bare metal and virtual machines) - is the original deployment configuration of the DSR. It scales to very high performance and is widely deployed.
- » Fully virtualized - was introduced shortly after bare-metal. It provides virtualization of the DSR, but does not use a cloud manager, and does not co-reside with other applications. Provides a compact, cost-effective footprint and is widely deployed.
- » Cloud deployable –It provides full virtualization, assumes the DSR resources will be managed by a COTS cloud manager, and that the DSR can be one of many applications in the cloud. Cloud deployable DSR is the focus of this document.
- » Mix and match – DSR is a network of DSR signaling sites. The deployment infrastructure at each site can vary. E.g. bare-metal at one site, and then cloud deployed at another location.

What is a cloud deployable DSR?

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 cloud 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.

Infrastructure matters

The DSR is capable of running on a huge variety of infrastructures, but not all infrastructures are the same. The performance, capacity, and latency of the DSR can vary dramatically based on the chosen infrastructure and how it is deployed. In general, the DSR works best in a high bandwidth, low-latency, high processing power environment (carrier grade cloud). Some considerations that impact DSR performance, capacity, latency:

- » Hardware – the CPUs, and NICs (network interface cards)
- » Hypervisor settings/configuration

- » Uplinks, switches, WAN latency
- » Storage configuration (local, networked)

DSR has excellent high availability and geo-diversity resiliency mechanisms that work in concert with cloud manager capabilities. Obviously the needed scale, availability, and resiliency of the deployment also impact the resource and infrastructure requirements.

Flexibility

DSR is flexibly deployed into many different clouds. It is unlikely that any two clouds are exactly the same and operators need to optimize for different reasons (e.g. power consumption may be critical for one operator, and WAN latency at another), varying sets of applications, and differing operational requirements. The performance and capacity of the DSR varies in each cloud, and the DSR application can no longer provide a guaranteed level of performance and capacity. However, the operator still needs to:

- » Plan their networks – DSR's use resources, what impact will DSR have on their datacenters?
- » Deploy DSR with predictable (if not exact) performance and capacity.
- » Manage the capacity and performance of the DSR in their datacenters.

Methodology

There is a set of DSR specific tools, methods and documentation to assist in planning, deploying, and managing the capacity and performance of a cloud deployable DSR. This toolset is expected to be used in conjunction with information and tools provided by the infrastructure (hardware, cloud manager, hypervisor) vendors.

- » Planning for cloud deployable DSR
 - » Estimating required resources for a given DSR cloud deployment
 - Please contact your Oracle Sales Consultant. They have access to the DSR Cloud Dimensioning tool which estimates DSR cloud resources. This tool takes into account many factors not covered in this benchmarking guide, such as the overhead for optional DSR features not covered in the benchmarking guide, and recommended margins for redundancy.
 - » DSR Cloud Customer Documentation
 - Can be found with the DSR customer documentation at Oracle.com on the Oracle Technology Network (OTN).
 - Look under the topic: "Cloud Planning, Installation, Upgrade, and Disaster Recovery".
- » Deploy DSR with predictable performance and capacity
 - » It is recommended that the DSR is run through a benchmark on the target cloud infrastructure to determine the likely capacity and performance in the target infrastructure. This information can then be used to adjust the initial deployment resources (if needed), and to help predict future resource requirements if and when the DSR grows.
 - » This document provides information on how to benchmark DSR performance and capacity. It also provides comprehensive benchmark results for a few select infrastructures. More benchmark data will be added to the document as it becomes available.
 - » This document also provides performance recommendations and observed differences for performance tuning decisions.
- » Manage the capacity and performance of the DSR
 - » The customer network is always changing- traffic patterns change, new applications are introduced. The infrastructure changes – new hardware, software/firmware updates. The operator needs to monitor and adjust the DSR resources for the changing conditions of the network and infrastructure.
 - » This document provides the key metrics and recommendations for monitoring the capacity and performance of a cloud deployed DSR.

Benchmarking Cloud Deployable DSR

This document is divided into several 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.
- » A benchmark section for each DSR server type - Each DSR server type is given independent treatment for its benchmark. Each section describes the traffic setup, and the observed results. It also provides metrics and guidelines for assessing performance on any infrastructure.

What to do with all this data?

This data is intended to provide guidance. Recommendations may need to be adapted to the conditions in a given operator's network. Each section below provides 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? Of course, 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. Your Oracle Sales Consultant can run this tool for you based on your 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, etc.).
 - Required DSR applications such as Full Address Based Resolution (FABR) and 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 VMware vs. KVM, and Server parameters.
 - » The tool will then generate a recommended number of VMs for each of the required VM types.
 - » As noted below, 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 comparing single-threaded processor performance of the CPUs used in this document vs. the CPUs used in the customer's infrastructure. This approach will be 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's 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 will adjust once the DSR is production.
- » If desired, execute benchmarking testing on the target cloud infrastructure. Only benchmark those types of DSR servers needed for the deployment (e.g. if full address resolution is not planned, don't waste time benchmarking the SDS, SDS SOAM, or DPs).
- » Once that benchmarking is completed, take a look at the data for each server type, and compare it to the baseline used for the estimate (from the 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.
 - Recalculate 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. Increase the DSR resources as described later in this document as needed.

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 (e.g. ESXi vendor recommendations and defaults found in [1]). 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 may be observed. When recommendations other than vendor defaults or recommendations are made, additional explanations are included in the applicable section.

There is a subsection included for each infrastructure environment used in benchmarking.

General Rules for All Infrastructures

Hyper-Threading and CPU Over-subscription

All of the tests were conducted with Hyper-Threading enabled, and a 1:1 subscription ratio for vCPUs in the hypervisor. The hardware used for the testing were dual-processor servers with 18 physical cores each (Xeon E5-2699v3). Thus each server had:

$$(2 \text{ CPUs}) \times (18 \text{ cores per CPU}) \times (2 \text{ threads per core}) = 72 \text{ vCPUs}$$

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

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

The vCPU sizing per VM is given in Appendix A.

Recommendation: *Hyper-Threading enabled and 1:1 CPU subscription ratio*

CPU Technology

The CPUs in the servers used for the benchmarking were the Intel Xeon E5-2699v3. Servers with different processors are going to give different results. In general there are two issues when mapping the results of the benchmarking data in this document to other CPUs:

- 1) The per-thread performance of a CPU is the main attribute that determines VM performance since the number of threads is fixed in the VM sizing as shown in Appendix A. 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 SPECint2006 ratios 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, etc.) there are other factors that affect per-thread performance such as potential turbo speeds of the CPU vs. the cooling solution in a given server. Comparing between Intel CPU generations there is a generation over generation improvement of CPU throughput vs. clock speed that means that even a newer generation chip with a slower clock speed may run a DSR VM faster.
- 2) The processors must have enough cores that a given VM can fit entirely into a NUMA node. Splitting a VM across NUMA nodes will greatly reduce the performance of that VM. The largest VM size (see Appendix A) is 12 vCPUs. Thus the smallest processor that should be used is a 6 core processor. Using processors with more cores typically makes it easier to "pack" VMs more efficiently into NUMA nodes, but shouldn't affect individual VM CPU-related performance otherwise (see the next note though).
- 3) 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 (18 physical cores per 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 will have to be watched closely.

Recommendation: *The selected VM sizes should fit within a single NUMA node (for instance 6 physical cores for the VMs that required 12 vCPUs. Check the performance of the target CPU type against the benchmarked CPU using per-thread integer performance metrics.*

VM Packing Rate

The OCDSR doesn't require or use CPU pinning. Thus the packing of the OCDSR VMs onto the physical servers is under the control of Openstack using the affinity/anti-affinity rules given in Appendix A: DSR VM Configurations. Typically the VMs will not fit exactly into the number of vCPUs available in each NUMA node, leaving some unallocated vCPUs. The ratio of the allocated to the unallocated vCPUs is the VM Packing Ratio. For instance, on a given server if 60 out of 72 vCPUs on a server were allocated by the Openstack, that server would have a packing ratio of ~83%. The achieved packing in a deployment will depend on a lot of factors, including the mix of large VMs (DA-MPs, SBRs) with the smaller VMs, and whether the OCDSR is sharing the servers with other applications that have a lot of large or small VMs. When planning the number of physical servers required for an OCDSR 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, performing functions such as interrupt handling. A packing ratio of 95% or lower is desirable.

Recommendation: 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.

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 [3].

- Txqueuelen: The default of 500 is too small. Recommendation is to set this parameter to 30,000.
 - Tuned on the compute hosts.
 - Default value of 500 is too small. Our recommendation is to set to 30000. 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. Already enabled by default on ESXi, needs to be set for Openstack.
- Advanced NUMA settings (ESXi only): The SwapLoadEnabled and SwapLocalityEnabled options should be disabled. This prevents the ESXi scheduler from moving VMs around from one NUMA node to another trying to optimize performance. These settings aren't appropriate for VMs that are processing real-time loads since messages might be delayed during the move.

-
- **Recommendation:** Follow instructions in the DSR Installation Guide.
-

KVM (QEMU) / Oracle X5-2 – 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.

A Host Hardware

- » Oracle Server X5-2
 - CPU Model: Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
 - 2 CPUs
 - 18 physical cores per CPU
 - RAM: 128 GB
 - HDD: 2.3 TB of solid state drive (SSD) storage
 - Nic:
 - 4 x Intel Ethernet Controller 10-Gigabit x540-AT2

B Hypervisor

- » QEMU-KVM Version: QEMU 1.5.3, libvirt 1.2.8, API QEMU 1.2.8

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 bypassing 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.

Recommendation: *Vhost-net driver is recommended.*

BIOS Power Settings

Typical BIOS power settings (hardware vendor dependent, see your 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

Recommendation: *Set to Allocated Power or equivalent for your Hardware vendor*

Disk Image Formats

The two 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.
- » **Raw Dump Representation:** Unstructured disk image format

QCOW2 provides a number of benefits over raw dump such as:

- » Smaller file size, even on filesystems which don't support holes (i.e. 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

The recommended Container format is "bare" – no container or metadata envelope for the disk image. The container format string is not currently being used by OpenStack components.

Recommendation: *QCOW2 (Since DSR does not involve processes which are disk I/O intensive.)*

Recommendation: *"Bare" Container format*

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:** With caching set to writeback mode, both the host page cache and the disk write cache are enabled for the guest. Because of 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 [Selected]:** 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.

Recommendation: *Caching Mode = None*

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. Because 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 will swap 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 will swap aggressively.

Recommendation: *vm.swappiness = 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 control KSM:

- » **KSM Service:** When the ksm service is started, KSM will share 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.

Recommendation: *'ksm' service set to active and '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 reclaim in the following situations:

- » When you configure the guest operating system to use huge pages.
- » When you use Kernel same-page merging (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.

Recommendation: *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.

Recommendation: *Enable THP*

VMware (ESXi) – Infrastructure Environment

There are a number of ESXi (VMware hypervisor) 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.

- Hypervisor versions tested
 - » ESXi 5.5

» ESXi 6.0 U2

Virtual Sockets vs. Cores per Virtual Socket

When defining a virtual machine the number of vCPUs must be assigned to a server. The user has options for setting the number of “Virtual Sockets” and the number of “Cores per Virtual Socket”. The product of these two parameters determines the number of vCPUs available to the virtual machine.

In following the *VMware* best practices, the default value of 1 core per socket was used. This configuration is referred to as “wide” and “flat.” This will enable vNUMA to select and present the best virtual NUMA topology to the guest operating system, which will be optimal on the underlying physical topology.

Recommendation: 1 core per socket, virtual socket set to the number of vCPUs required by the server role.

Network Settings

Network Adapters

There is a number of networking adapter choices when deploying a virtual machine:

- » **E1000:** This adapter is an emulated version of Intel 82545EM Gigabit Ethernet Controller. VMXNET3 adapter is the next generation of Para virtualized NIC designed for performance.
- » **VMXNET3:** This adapter has less CPU overhead compared to e1000 or e1000e. Also, VMXNET3 is more stable than e1000 or e1000e. . VMXNET3 adapter is the next generation of Para virtualized NIC designed for performance. This is the vSphere default setting.
 - VMXNET family implements an idealized network interface that passes network traffic between the virtual machine and the physical network interface cards with minimal overhead.

Recommendation: VMXNET3. No observable differences were noticed between E1000 and VMXNET3 for DSR application testing.

Virtual Network Interrupt Coalescing and SplitRx Mode

- **Virtual network Interrupt Coalescing:** This option reduces number of interrupts thus potentially decreasing CPU utilization. This may however increase network latency. By default this is enabled in ESX 5.5 and 6.0.
- **SplitRxMode:** This option uses multiple physical CPUs to process network packets received in single network queue. By default this is enabled in ESX 5.5 and 6.0 for VMXNET3 adapter type.

Table 3: Virtual Network Interrupt Coalescing and SplitRX Mode

Network Setting:	Default	Virtual network Interrupt Coalescing: Disabled	SplitRxMode: Disabled
DSR.Cpu (Avg / Max)	~40.7% / ~44.5%	~42% / ~45.5%	~38.8% / ~40.6%
System.CPU_UtilPct (Avg / Max)	~44.4% / ~53%	~44.4% / ~55.5%	~41.8% / ~53.3%
Latency	Observed as same in DSR application benchmarking		

Recommendation: Virtual network interrupt coalescing: Enabled, SplitRxMode: Enabled.

Power Settings

VMware ESXi allows assignment of power management profiles. These profiles allow the user to configure the host to save power while balancing performance. The power management profiles use the host's processor ACPI power setting. Many host manufacturer's bios overrides the ESXi settings.

Table 4: Power Management Profiles

ESXi Power Mode:	High Performance	Balanced Performance
System.CPU UtilPct (Avg / Max)	~40% / ~60%	~38% / ~55%
Dsr.Cpu (Avg / Max)	~38% / ~48%	~36% / ~44%
Used % / Run % / System %	~472 / ~388 / ~49%	~462 / ~376 / ~49%
Wait % / Idle %	~407% / ~1013	~419% / ~1023

The data in the table above is collected from a DA MP, but similar trends are observed on the other DSR virtual server types. A small but significant difference was observed between balanced and high performance power settings. However, the data did not indicate a large enough deviation to vary from the hardware vendor's guidelines. DSR benchmark testing was performed with ESXi Power Mode set to Balanced Performance.

Recommendation: Refer to host hardware vendor power management guidelines for virtualization.

Hardware Assisted Virtualization

VMware ESXi automatically determines if a virtual machine can use hardware support for virtualization based on processor type. Several settings were selected for assessing performance:

- A. Automatic
- B. Use software for instruction set and MMU virtualization.
- C. Use Intel® VT-x/AMD-V™ for instruction set virtualization and software for MMU virtualization
- D. Use Intel® VT-x/AMD-V™ for instruction set virtualization and Intel® EPT/AMD RVI for MMU virtualization
 - Also testing with "Node Interleaving" setting Enabled [i.e. NUMA disabled], with no noticeable changes in performance.

Table 5: Virtualization Performance by Processor

MMU Virtualization Setting:	A.	B.	C.	D.
System CPU UtilPct (Max / Avg)	57.5/38%	71.5/43%	71.5/43%	53/38%
Dsr.Cpu (Max / Avg)	43.5/36.3%	50/38.6%	50/38.5%	43/36.3%

The data in the table above is provided from a DA MP. Trends for other servers are similar. The automatic (default) settings provide performance better than options B and C above, and fairly equivalent to option D.

Recommendation: Refer to host hardware vendor guidelines for virtualization. Defaults recommended.

Virtual Machine Storage Configuration

Storage Type Adapter

Testing was performed with the default “LSI Logic Parallel” option. No testing was performed against recent virtual SCSI adapters (LSI Logic SAS and VMware para-virtualized (PVSCSI).) At the time of testing the default was considered as the most stable and compatible.

Recommendation: Default “LSI Logic Parallel”

Disk Provisioning

The following disk provisioning options are available when deploying a virtual machine:

- » **Thick Provision Lazy Zeroed:** All space needed for the VM is allocated during creation. Data on the host disk is zeroed out at a later time on first write from the virtual machine.
- » **Thick Provision Eager Zeroed:** All space needed for the VM is allocated during creation. The data on the host disk will be zeroed out during creation. Time to deploy the virtual machine will be increased with this option. This option will support fault tolerant features provided by the infrastructure.
- » **Thin Provision:** This option uses only the amount needed by the virtual machine disk. The image will grow as needed until the allocated capacity is reached.

With the high availability of the DSR, storage should be allocated at the time the VM is created, so thin provisioned is not recommended. When instantiating a fairly typical DSR VM with 60G of storage, the lazy zeroed disk was created almost instantaneously. Whereas the eager zeroed disk took about 7 minutes to initialize. Lazy zeroed is recommended.

Recommendation: “Thick Provisioned Lazy Zeroed”

Large Memory Pages

VMware ESXi Large-page support enables server applications to establish large-page memory regions. The use of large pages can potentially increase TLB access efficiency and thus improve program performance. By default Large page support is enabled on VMware ESXi Server although the full benefit of large pages comes only when guest OS and applications use them as well. By default large page support is not enabled in the DSR product.

The following settings were evaluated:

- Default settings [i.e. Large memory pages support enabled on host and large pages configured as 0 on guest]
- Large memory pages support enabled on host and 1024 large pages configured on guest
- Large memory pages support disabled on host

Recommendation: Default settings. No visible advantage was observed when comparing iterative memory stats as observed through /proc/meminfo. No visible advantage could be observed in using large pages.

Benchmark Testing

The way the testing was performed and the benchmark test set-up is the same for each benchmark infrastructure. Each section below describes the common set-up and procedures used to benchmark, and then the specific results for the benchmarks are provided for each benchmark infrastructure. In general the benchmarking results for VMware/ESXi vs. Openstack/KVM are close enough that only one set of numbers are shown.

DA MP Relay Benchmark

Overview

This benchmarking case illustrates conditions for an overload of a DSR DA MP. Simulator message rate is increased until the DA-MP Overload mechanisms are triggered causing messages to be discarded.

Topology

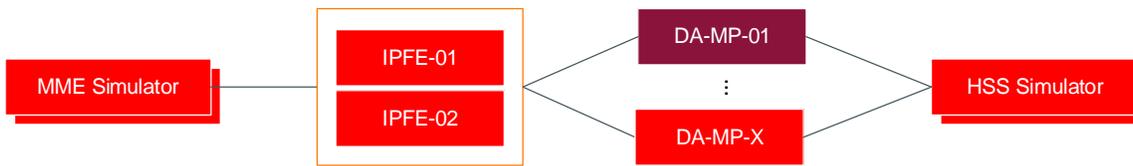


Figure 1 – DA-MP Testing Topology

Figure 1 illustrates the logical topology used for this testing. DIAMETER traffic is generated by an MME simulator and sent to an HSS simulator. The traffic is ramped up in increments until congestion is seen on the DA-MPs indicating that the maximum normal traffic handling capacity has been reached for the current message flow.

Message Flow

Figure 2 illustrates the Message sequence for this benchmark case.

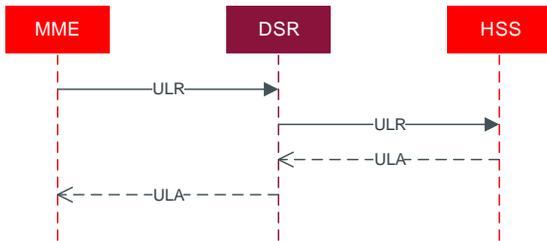


Figure 2 – DA-MP Message Sequence

Table 6: DA-MP Test Set-up

Messages		Traffic Details	
Message	Distribution	Detail	Distribution
ULR, ULA	100%	Average Message Size	2.2K
		Cross DSR Routing	75 or higher%

Observations

Figure 3 – DA-MP CPU Occupancy details the utilization of CPU of the affected DA MP. As the message rate increases the amount of CPU necessary to process these incoming messages increases. The DA MP will be marked as congested and messages will be rejected as the Danger of Congestion (DOC) threshold is reached at 56% CPU utilization for the Peak utilization values shown below. At 25k MPS the DA-MP has crossed this threshold. When this threshold is reached the DSR is shedding traffic based on the DOC settings, is generating congestion alarms, and the average latency for requests that are handled has started to increase significantly. Since some headroom should be provided between the point where the DA-MP goes into overload and the target for maximum routine traffic handling, Oracle’s recommendation is to use 18k MPS as a maximum engineering number for relay traffic.

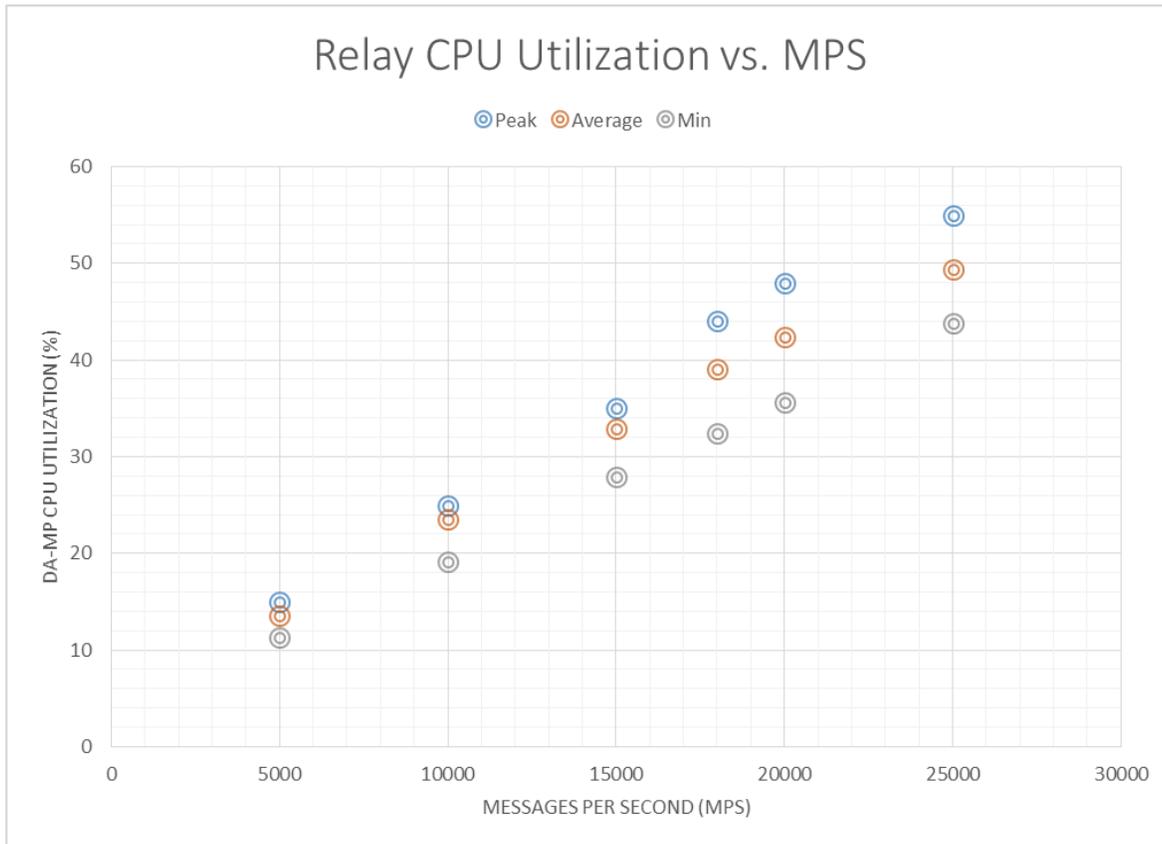


Figure 3 – DA-MP CPU Occupancy

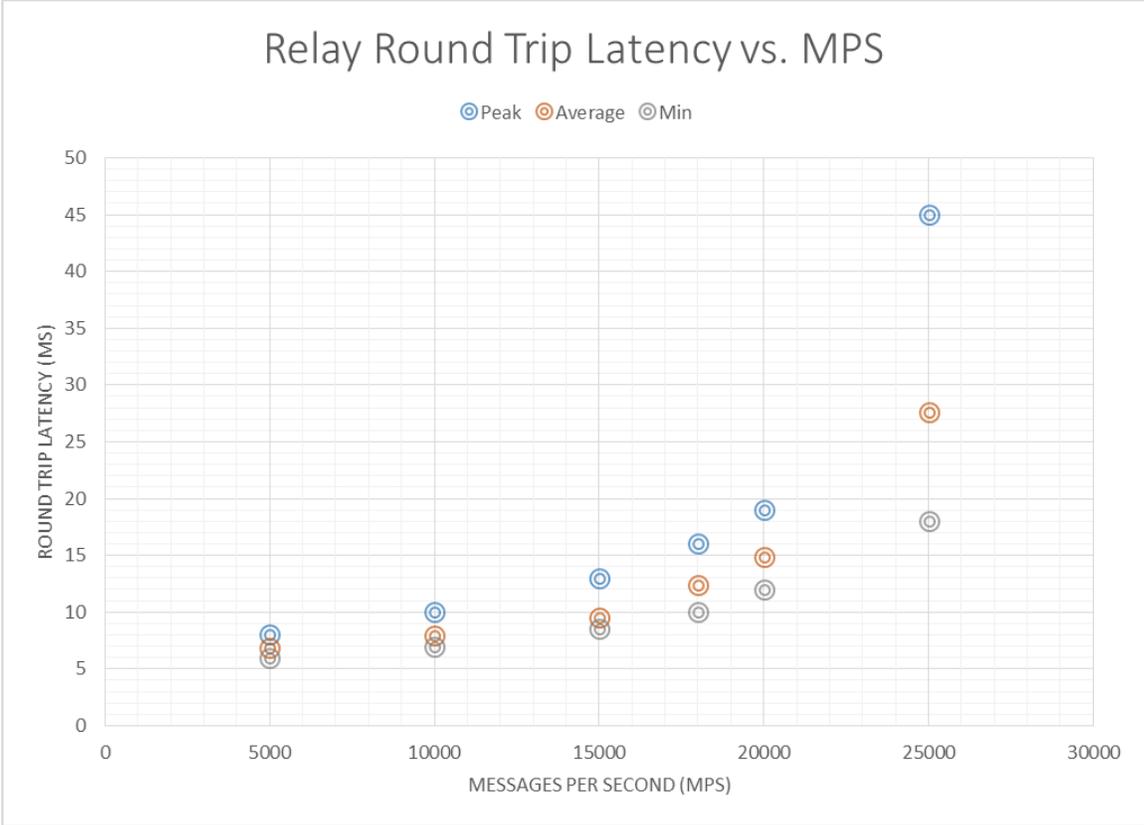


Figure 4 – DA-MP Round-Trip Latency vs. Message Rate

It was observed as the number of incoming messages increased the round trip message latency increases. Figure 4 illustrates the increase in latency gradually increases as the message rate increases. Once the system goes into overload above 20k MPS there is a significant spike in the latency.

Indicative Alarms / Events

During benchmark testing the following alarms/events were observed as it crossed into congestion.

Table 7: DA-MP Alarms/Events

Number	Severity	Server	Name	Description
5005	Minor	IPFE	IPFE Backend in Stasis	A backend server not accepting new connections but continuing to process existing connections
22200	Major	DA-MP	Communication Agent Routed Service Congested	The Diameter process is approaching or exceeding its engineered traffic handling capacity
22215	Major	DA-MP	Ingress Message Discarded: DA MP Overload Control	Ingress message discarded due to DA MP (danger of) CPU congestion.

DA-MP Indicators for Capacity Expansion

The following DSR system measurements, KPIs and events can be used to monitor the performance of the DA-MP in a field deployment, or to drive the orchestration of DA-MP expansion.

DA-MP Utilization

In this section, only the key recommended metrics for planning expansions of the DA-MP are discussed. There are many more measurements available on the DA-MP, and these can be found in [2].

The key metrics for managing the DA-MP are:

Table 8: DA-MP Utilization Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
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 as needed.
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, anomalous traffic spikes 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	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	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	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

					the rated maximum capacity, OR exceeds 60% of the rated capacity when running without an active mate.	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

In this section, only the key recommended metrics for planning expansions of the DA-MP connections are discussed. There are many more measurements available on the DA-MP connections, and these can be found in [2].

The key metrics for managing the DA-MP connections are:

Table 9: DA-MP Connection Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
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

Suggested Resolution

If congestion alarms shown in Table 7: DA-MP Alarms/Events then add additional DA-MPs to avoid CPU congestion. However, if the connection alarm shown in Table 9: DA-MP Connection Metrics is seen, than adding additional connections for that peer will help distribute the load and alleviate the connection alarm.

In general, the growth mechanism for DA MPs is via horizontal scaling. That is by adding additional DA MPs. The current maximum number of the DA MPs per DSR signaling NE is 16. If the number of DA MPs at given DSR Signaling NE is nearing its limits, the consider adding an additional DSR signaling NE. A DSR network supports up to 32 DSR signaling NEs.

Full Address Based Resolution (FABR) Capacity

Overview

The FABR application adds a Database Processor (DP) server is 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 will return 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.

Topology

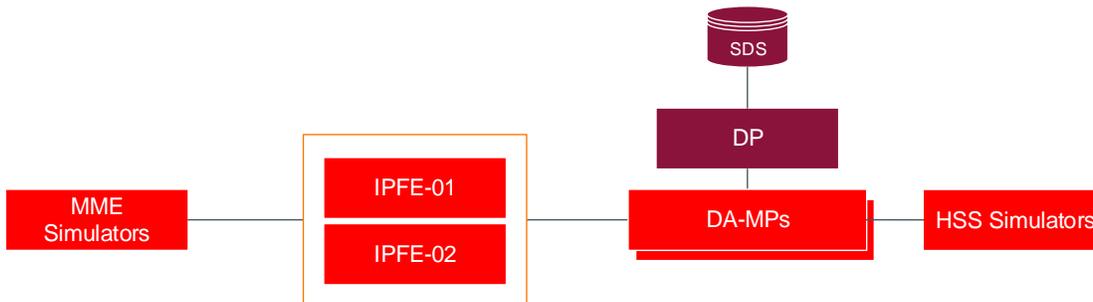


Figure 5 - SDS DP Testing Topology

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 queried against.

- » SDS DB Size: 300 Million Routing Entities (150 M MSISDNs / 150 M IMSIs)
- » AVP Decoded: User-Name for IMSI

Message Flow

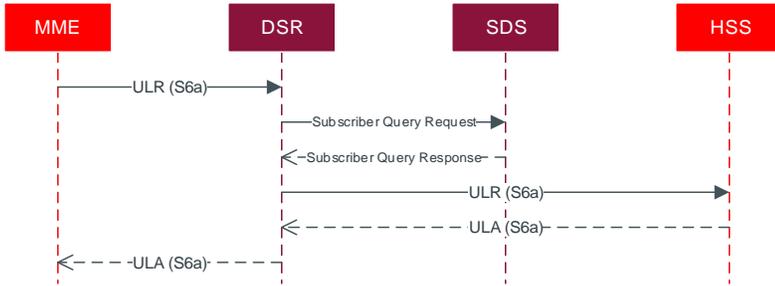


Figure 6 - SDS DP Message Sequence

Table 10: SDS DP Message Details

Messages		Traffic Details	
Message	Distribution	Detail	Distribution
ULR, ULA	100%	Average Message Size	2.2K
		Cross DSR Routing	75% or higher

DA-MP FABR Performance Observations

Figure 7 shows the CPU utilization % vs MPS for the DA-MP while running FABR traffic. Since there's more work required per request message due to the database lookup, the CPU utilization per MPS is higher than for the Relay traffic as shown in Figure 3. For FABR the Peak CPU utilization crosses the 56% DOC threshold around 17k MPS level. Figure 8 shows the end to end latency versus MPS for the FABR traffic. The delay is well-behaved increases slowly up until around the 18k MPS point after which it jumps significantly. Since some headroom should be provided between the point where the DA-MP goes into overload and the target for maximum routine traffic handling, Oracle's recommendation is to use 16k MPS as a maximum engineering number for FABR traffic. Please see the previous section DA-MP benchmarking for all of the DA-MP indicators for capacity expansion.

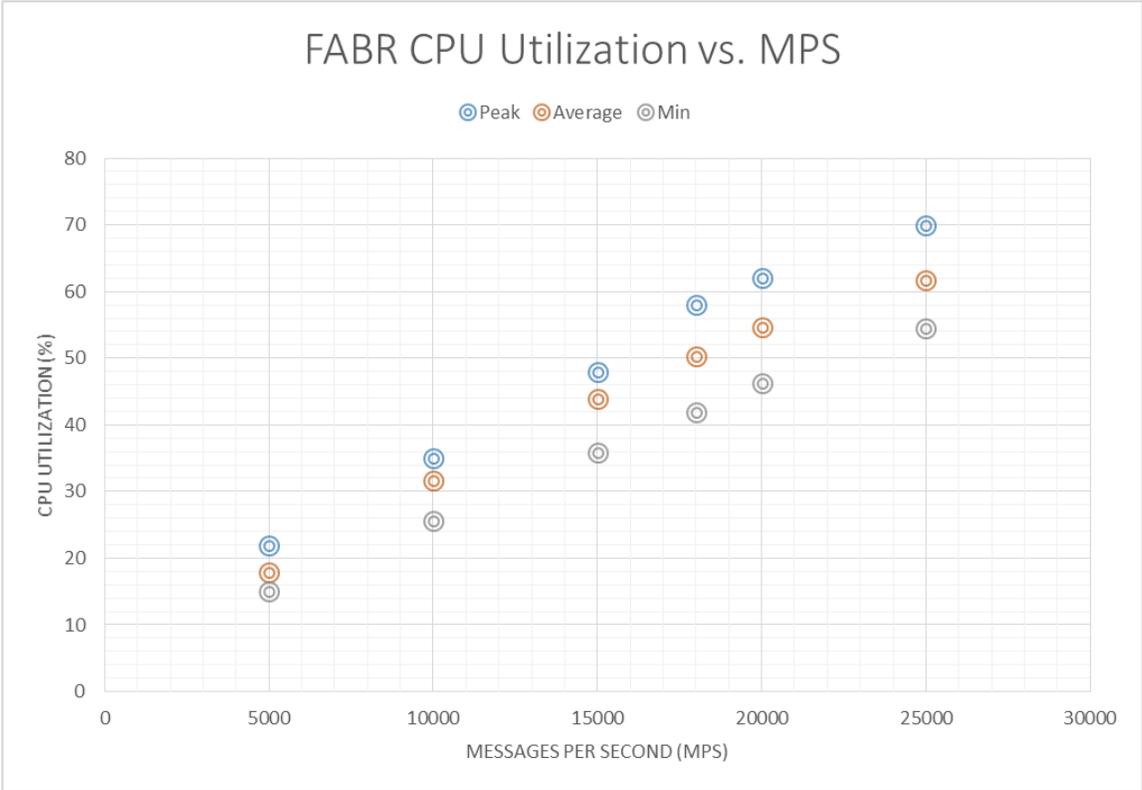


Figure 7 – DA-MP CPU Utilization (%) vs. MPS Running FABR Traffic

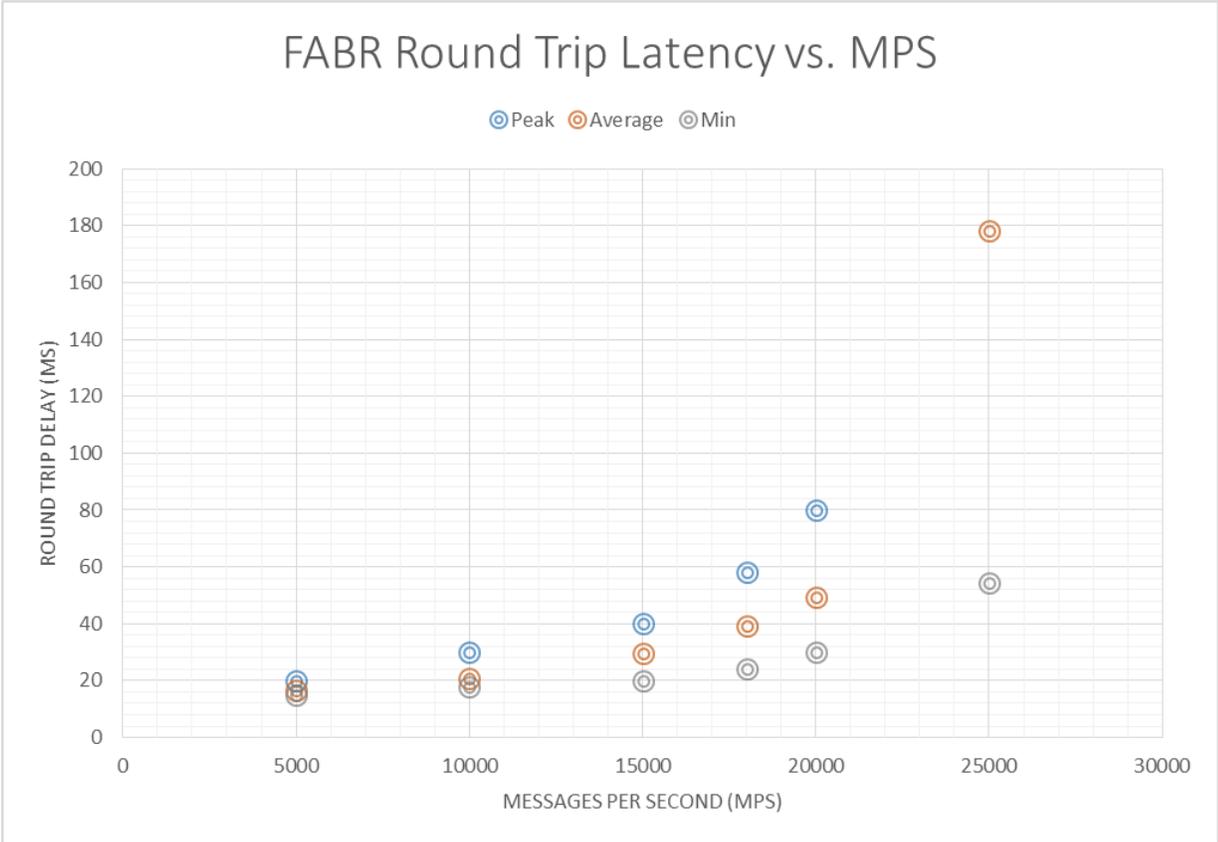


Figure 8 – DA-MP Latency vs. MPS Running FABR Traffic

DP CPU Utilization Observations

The CPU utilization vs. MPS is shown in Figure 9. As the number of DP Lookups increases so does the CPU occupancy of the Database Processor (DP). Typical only 50% of the DSR’s traffic needs to do a lookup (i.e. responses don’t need to lookup), and depending of the message flow, it could be less.

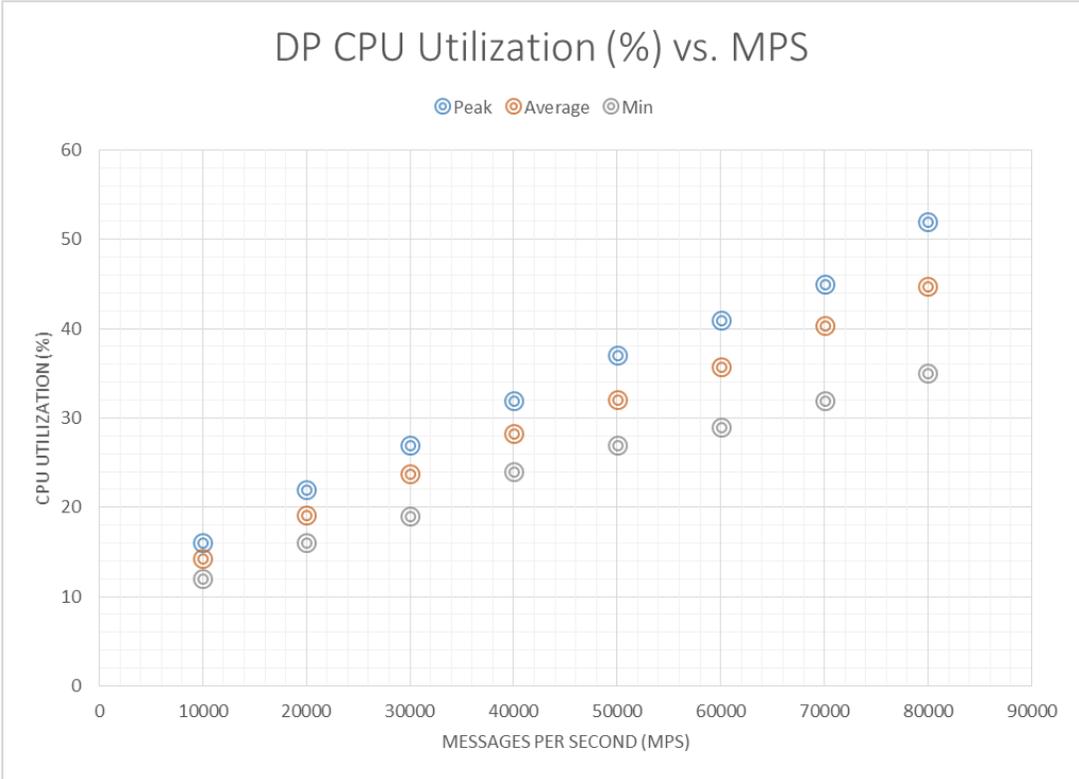


Figure 9 - DP CPU Util vs Traffic Rate (DP Queries)

DP Indicators for Capacity Expansion

Table 11: SDS DP Alarms/Events

Number	Severity	Server	Name	Description
19900	Minor	sdsDP	Process CPU Utilization	The Process, which is responsible for handling all traffic, is approaching or exceeding its engineered traffic handling capacity.
19822	Major	DA-MP	Communication Agent Routed Service Congested	Communication Agent Routed Service Congested
19825	Major / Critical	DA-MP	Communication Agent Transaction Failure Rate	The number of failed transactions during the sampling period has exceeded configured thresholds.
19826	Major	sdsDP	Communication Agent Connection Congested	Communication Agent Connection Congested
19831	Info	DA-MP	Communication Agent Service Operational State Changed	Communication Agent Service Operational State Changed, Instance DPService
19816	Info	DA-MP	Communication Agent Connection state Changed	Configuration Mode = Configured, Admin State = Enabled, Connect Mode = Server, Operational State = Degraded, Congestion Level = 1, Overload Level = 1, Transport Congestion Level = 0

Measuring DP Utilization

In this section, only the key recommended metrics for managing the performance of the DP are discussed. There are many more measurements available on the DP, and these can be found in [2].

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 DPs are:

Table 12: SDS DP Utilization Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
4170	DpQueriesReceived	DP	System (per DP)	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 benchmarked maximum capacity, OR exceeds 60% of the benchmarked 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 DPs should be planned for.
31056	RAM_UtilPct_Average	System	System (per DP)	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 (per DP)	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 benchmarked 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 for the entire DSR network.

The OAM application related to the DPs (DP SOAM) runs at each DSR Signaling NE requiring the Full Address Resolution feature. Currently these are fixed sized VMs with no horizontal or vertical scaling recommended as no need for scaling these VMs has been observed. The following two metrics should be monitored,

Table 13: DP SOAM Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
31056	RAM_UtilPct_Average	System	System (per DP SOAM)	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 (per DP SOAM)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	If the average CPU utilization exceeds 80% utilization	Contact Oracle

Suggested Resolution

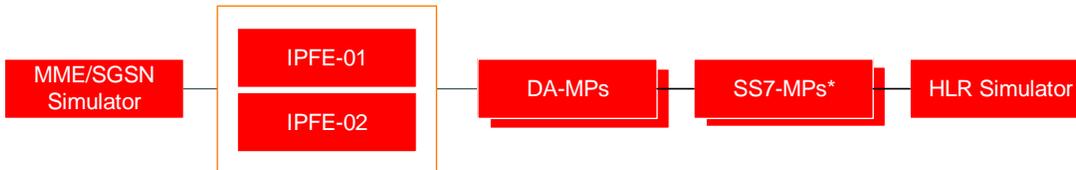
The growth mechanism for DPs is via horizontal scaling by adding additional DPs. The current maximum number of the DPs per DSR signaling NE is 10. This amount of scaling currently well exceeds capacities of the DA MPs driving queries to the DPs.

SS7 MP

Overview

The SS7-MP server type is responsible for transforming messages between SS7 and Diameter protocols. Both Diameter and MAP messages were sent from the simulator to the DSR. The SS7-MP is based on similar technology to the DA-MP benchmarked in previous sessions. The licensed capacity of the SS7-MP is currently limited to 12k MPS per SS7-MP, even though the performance is similar to the DA-MP. This section is here to explain the functions of the SS7-MP, and note the SS7-MP specific events and measurements used to monitor it.

Topology



*System under test

Figure 10 - SS7 MP Testing Topology

Message Flow

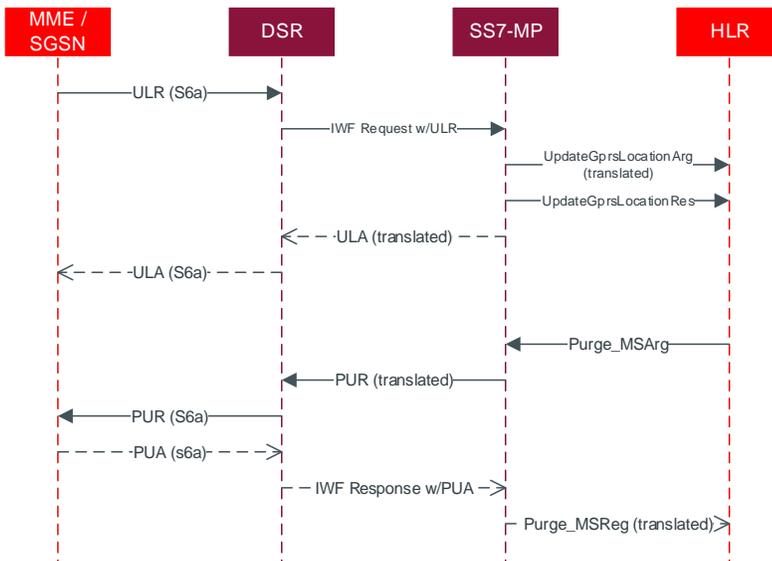


Figure 11 - SS7 MP Message Flow

Table 14: SS7 MP Message Detail

Detail	Distribution
Diameter originated	50%

MAP originated	50%
----------------	-----

Indicative Alarms / Events

Table 15: SS7 MP Alarms/Events

Number	Severity	Server	Name	Description
19250	Minor	DA-SS7MP	SS7 Process CPU Utilization	The SS7 Process, which is responsible for handling all SS7 traffic, is approaching or exceeding its engineered traffic handling capacity.

SS7 MP Indicators for Capacity Expansion

Measuring SS7 MP Utilization

In this section, only the key recommended metrics for planning expansions of the SS7 MP are discussed. There are many more measurements available on the SS7 MP, and these can be found in [2].

The key metrics for managing the SS7 MP and associated IWF DA-MP are:

Table 16: SS7 MP Metrics

Measurement ID (or KPI ID)	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
11054	MAP Ingress Message Rate	MAP Diameter Interworking function (MD-IWF)	SS7 MP	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 benchmarked maximum, OR exceeds 60% of the benchmarked 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 MPs 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 (SS7 MP)	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 (SS7 MP)	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%, OR exceeds 60% when running	If additional growth in the system is anticipated, then consider adding an additional SS7 MP.



					without an active mate.	
--	--	--	--	--	-------------------------	--

Suggested Resolution

Add additional SS7 MPs to accommodate additional MAP traffic.

The growth mechanism for SS7 MPs is via horizontal scaling by adding additional SS7 MPs.

vSTP MP

Overview

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

For the 8.1 release only a single vSTP MP is supported per DSR VNF, with a capacity of 10k MPS. Since STPs are always deployed in mated pairs, the redundancy is at the VNF (mated pair) level rather than at DSR/vSTP level.

SS7 MP Indicators for Capacity Expansion

In the 8.1 release only a single vSTP MP is supported per OCDSR node.

Policy DRA (PDRA) Benchmarking

Overview

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 (i.e. horizontal scaling units) are required. The first is the MPS traffic rate seen at the DA-MPs in Figure 12. 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/binding capacity is set primarily by the memory sizing of the VM, and is fixed at a maximum of 16 million per SBR for the DSR 8.1 release. The number of bindings and sessions required for a given network will be very 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. Note that the number of sessions will always be equal to or greater than the number of bindings.

Topology



Figure 12 - SBR Testing Topology

Message Flow

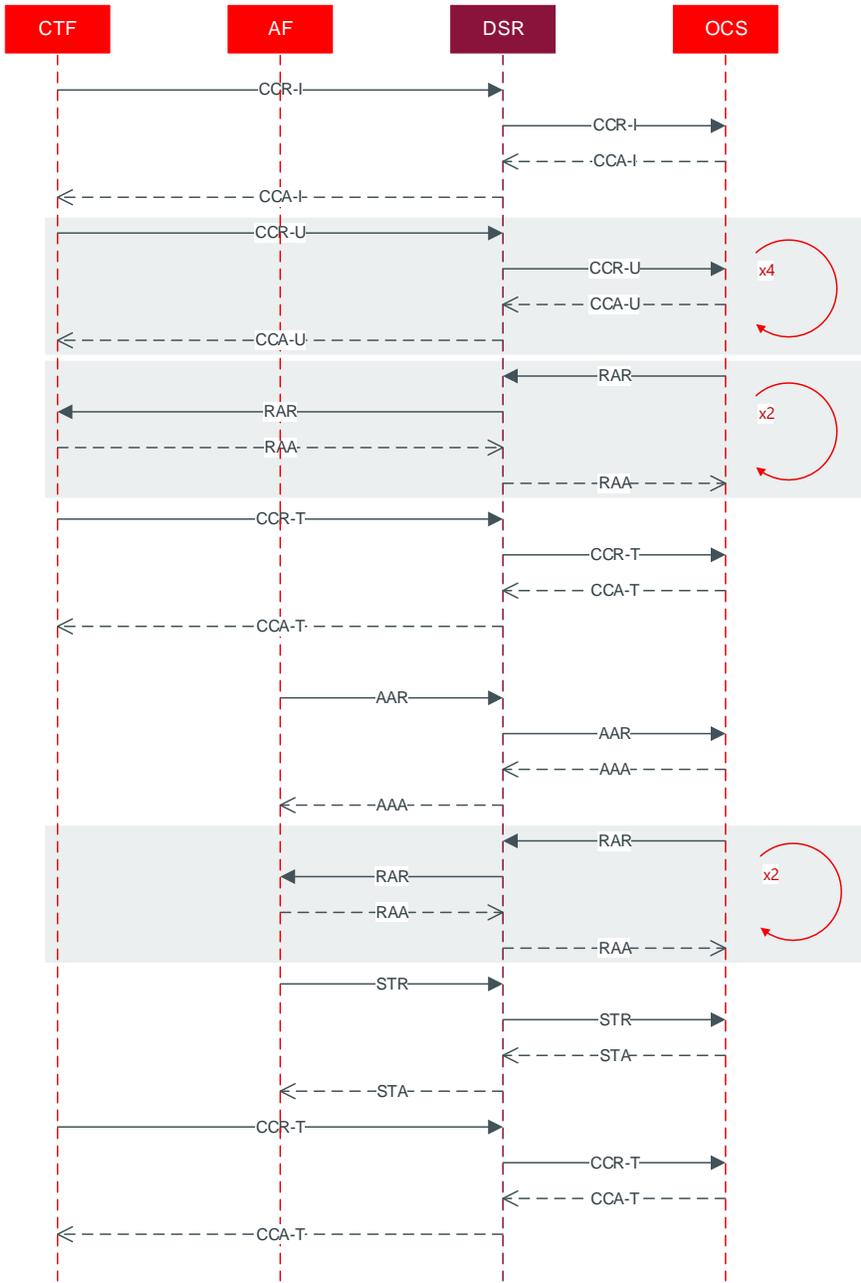


Figure 13 - PDRA Message Sequence

Table 17 shows the call model used for the testing. The message distribution is Oracle's baseline benchmarking, and may differ significantly from customer distributions based on factors such as the penetration of LTE support vs. VoLTE support. The Traffic Details shows the configured PDRA options. For more details on these options please see [4].

Table 17: PDRA Test Call Model

Messages		Traffic Details	
Message	Distribution	Detail	Distribution
CCR-I, CCA-I	8.5%	Gx w/ IPv6 Alternate Key	100%
CCR-U, CCA-U	25.5%	Gx w/ IPv4 Alternate Key	0%
CCR-T, CCA-T	8.5%	Gx with MSISDN Alternative Key	100%
Gx RAR, RAA	25.5%	Gx Topology Hiding	0%
AAR, AAA Initial	12.8%	Rx Topology Hiding	0%
STR, STA	12.8%		
Rx RAR, RAA	6.4%		

DA-MP PDRA Performance Observations

Figure 14 shows the CPU utilization % vs MPS for the DA-MP while running PDRA traffic. Since it can have to write database entries as well as read them, the PDRA application has more overhead than either Relay or FABR. For PDRA the Peak utilization crosses the 56% DOC threshold around 14k MPS. Figure 15 shows the end to end latency versus MPS for the PDRA traffic. The delay increases slowly up until around the 16k MPS point after which it jumps significantly (not shown in graph). Since some headroom should be provided between the point where the DA-MP goes into overload and the target for maximum routine traffic handling, Oracle's recommendation is to use 13k MPS as a maximum engineering number for PDRA traffic. Please see the previous section DA-MP benchmarking for all of the DA-MP indicators for capacity expansion.

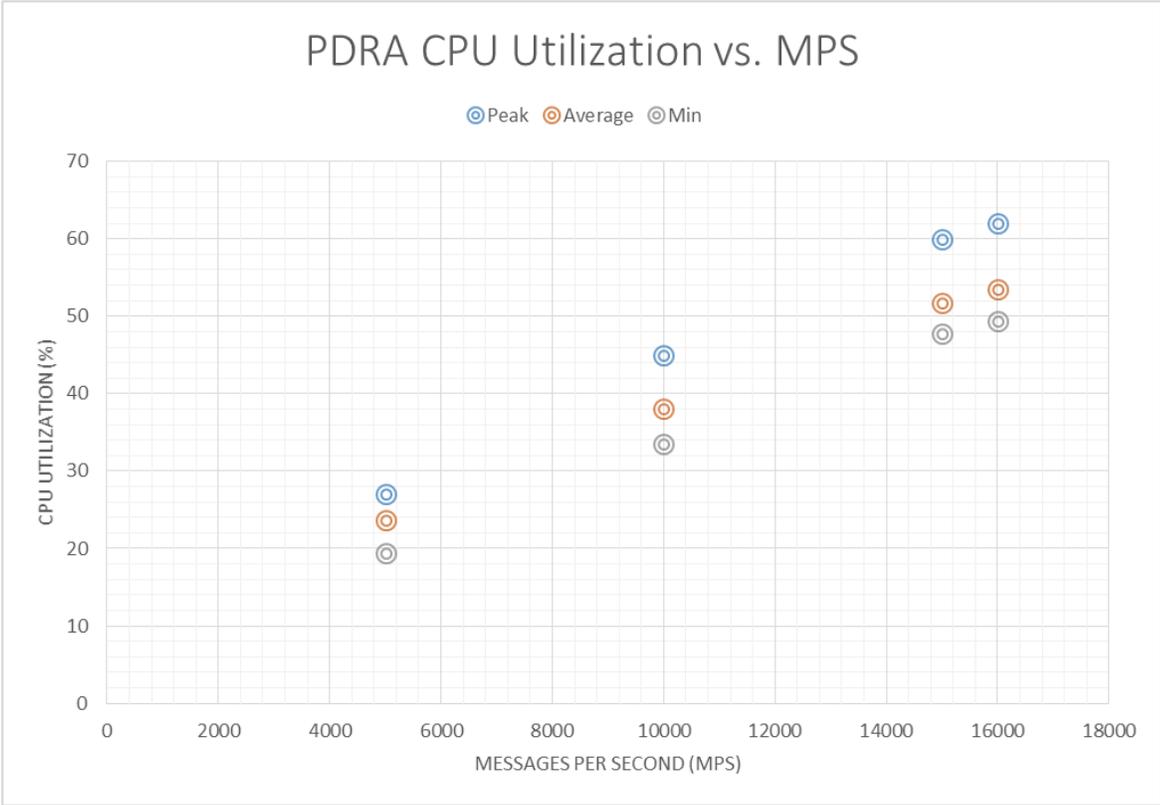


Figure 14 – DA-MP CPU Utilization (%) vs. MPS Running PDRA Traffic

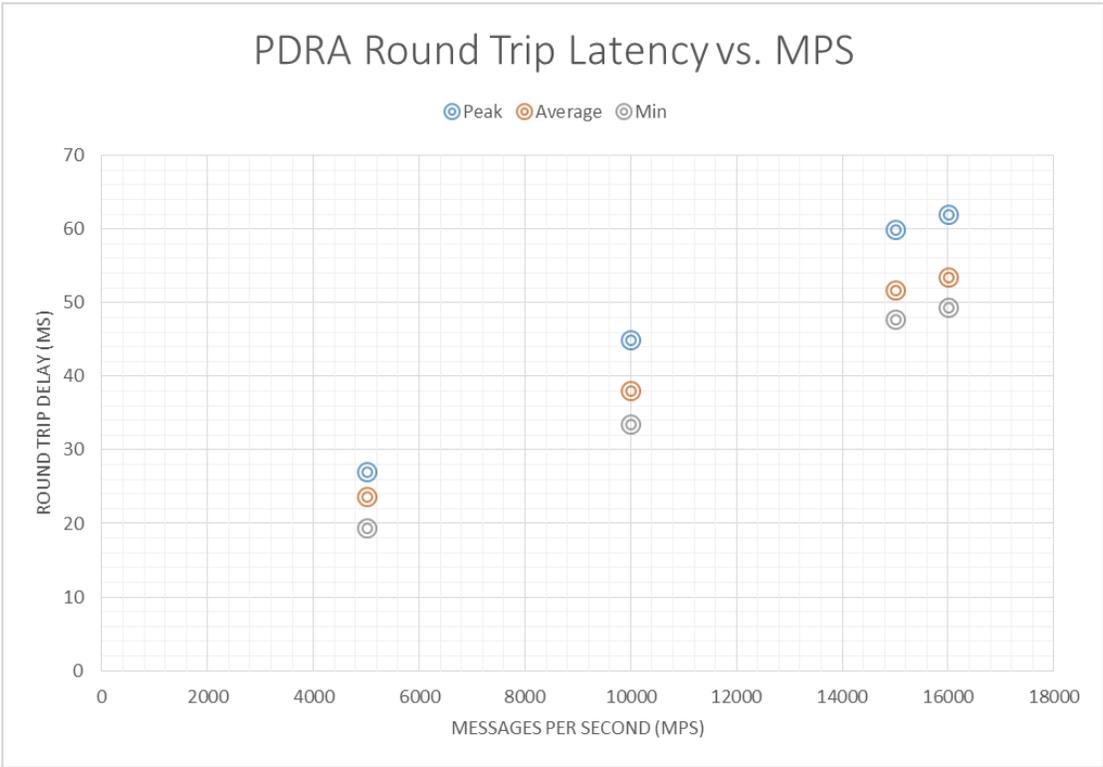


Figure 15 – DA-MP Latency vs. MPS Running PDRA Traffic

SBR Observations

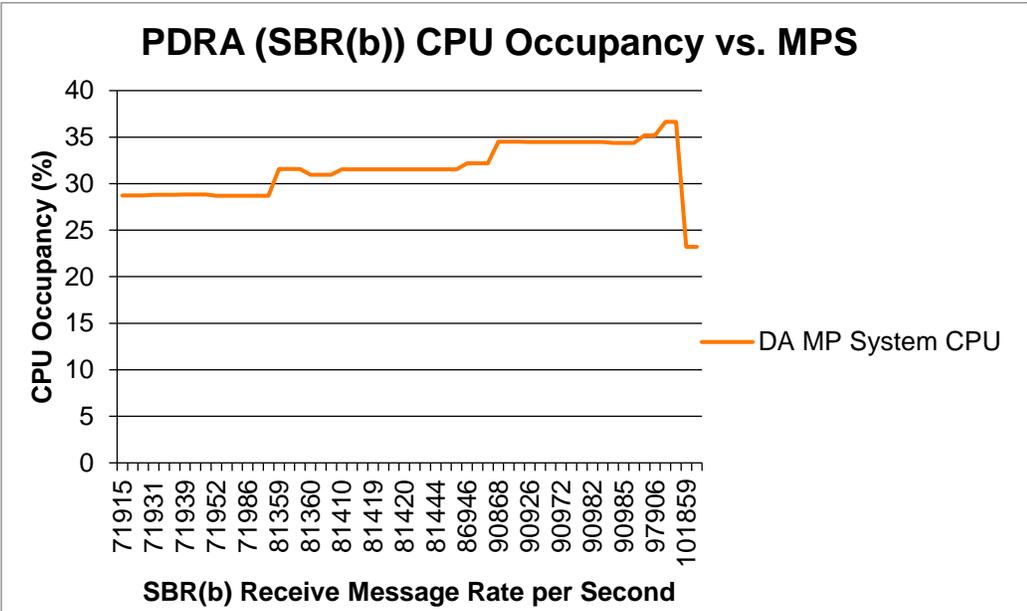


Figure 16: SBR (b) CPU Occupancy vs. Message Rate

Like with the other messaging intensive components of the DSR, this processor too is recommended to stay at or below 60% CPU occupancy for planning purposes. From Figure 16 with 5 Million sessions, the SBR (b) runs about ~100K receive messages per second at 60% CPU occupancy. Note that this number needs to be derated for other potential activities such as SBR Reconfiguration that can put additional load on the SBRs. The suggested engineering for the SBR-B is 50k MPS per Server Group running the standard call model given above.

Indicative Alarms / Events

Table 18: SBR (b) Alarms/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 configured thresholds.
19826	Major	DA-MP, SBR(s)	Communication Agent Connection Congested	Communication Agent Connection Congested
19846	Major	DA-MP, SBR(s)	Communication Agent Resource Degraded	Communication Agent Resource Degraded
22051	Critical	SOAM	Peer Unavailable	Unable to access the Diameter Peer because all of the diameter connections are Down.
22101	Major	SOAM	Connection Unavailable	Connection is unavailable for Diameter Request/Answer exchange with peer.
22715	Minor	SBR(s)	SBR Audit Suspended	SBR audit is suspended due to congestion.
22725	Minor / Major	SBR(s)	SBR Server In Congestion	SBR server operating in congestion.
22732	Minor / Major	SBR(s)	SBR Process CPU Utilization Threshold Exceeded	SBR process CPU utilization threshold has been exceeded.

Measurements

Key metrics for managing the Session SBR (b) blades are:

Table 19: Session SBR (b) Blades Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
31052	CPU_UtilPct_Average	System	System (SBR)	The average CPU usage from 0 to 100% (100% indicates that all cores are completely busy).	When this measurement exceeds 60% utilization	Contact Oracle
31056	RAM_UtilPct_Average	System	SBR	The average committed RAM usage as a percentage of the total physical RAM.	If the average Ram utilization exceeds 80% utilization	Contact Oracle
11372	SbrPolicySessionRecsAvg	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 benchmarked capacity. If both P-DRA and OC-DRA are enabled this average	Contact Oracle

					must be combined with the SbrOcSessionRecsAvg and the combined average exceeds benchmarked capacity.	
11441	SbrOcSessionRecsAvg	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 benchmarked capacity. If both P-DRA and OC-DRA are enabled this average must be combined with the SbrPolicySessionRecsAvg and the combined average exceeds benchmarked capacity.	Contact Oracle

Key metrics for managing the Binding SBR (b) servers are:

Table 20: Binding SBR (b) Server Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
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% occupancy.	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	SbrPolicyBindingRecsAvg	SBR Binding Performance	Server Group	Average number of active SBR Policy bindings.	When this average exceeds benchmarked capacity.	Contact Oracle

Suggested Resolution

If either additional Bindings or MPS capacity is required is required then additional Server Groups may be added to an existing SBR(b) using the SBR reconfiguration feature. There can be up to 8 Server Groups in the SBR(b).

NOAM

Overview

Specific benchmark data for the DSR NOAM is not provided in this release as the DSR Cloud deployable footprint is modest and system testing of the DSR indicates that NOAM growth is not currently needed.

Indicative Alarms / Events

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.

Measurements

Measuring Network OAM Utilization

In this section, only the key recommended metrics for managing the performance of the Network OAM are discussed. There are many more measurements available, and these can be found in [2].

The key metric for managing the Network OAM Servers are:

Table 21: Network OAM Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
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

Suggested Resolution

The NOAM can be vertically scaled; however this action is not anticipated to be necessary with the DSR 7.1.1 cloud deployable footprint. Please contact Oracle support for additional guidance as needed.

SOAM

Overview

Specific benchmark data for the DSR NOAM is not provided in this release as the DSR Cloud deployable footprint is modest and system testing of the DSR indicates that NOAM growth is not currently needed.

Indicative Alarms / Events

A key metric for managing the System OAM blades is:

Table 22: System OAM Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
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

Suggested Resolution

Vertical and horizontal scaling of the DSR is not supported or indicated in this release. Please contact Oracle support for additional guidance as needed.

IPFE

Overview

The IPFE was exercised in both VMware and KVM environments. Table 23 shows the measurement capacity of the IPFE. Note that there are three main factors that determine the throughput limits:

- The number of TSAs (one or more) on the IPFE
- Whether there are more than 2,000 connections
- Whether the average message size is less than the MTU size.

Under most conditions the throughput of the IPFE is 2 Gbit/sec. However under the worst case of all three of the above conditions the throughput of the IPFE drops to 1.6 Gb/sec.

When monitoring IPFE capacity both the guest and host CPU utilization should be monitored. Much of the IPFE work is done at the host kernel level so the CPU utilization numbers returned by the IPFE application level don't fully reflect all of the IPFE overhead on the system.

	Single TSA on IPFE Pair		2 or more TSAs on IPFE Pair (Total on both TSAs)	
	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	2 Gbit/sec	2 Gbit/sec	2 Gbit/sec	2 Gbit/sec
More than 2,000 Connections	2 Gbit/sec	1.6 Mbits/sec	2 Gbit/sec	2 Gbit/sec

Table 23: IPFE Throughput

Topology



Figure 17 – IPFE on Ingress Side Only

Figure 17 shows the typical IPFE configuration, with IPFEs used for “ingress” side traffic between the DIAMETER clients. This configuration is typical since there are typically many more DIAMETER clients such as MMEs than there are DIAMETER servers such as HSSs. The bandwidth numbers given in Table 23 are for the traffic flowing between the clients (MME in the figure above) and the DSR DA-MPs.



Figure 18 – IPFE on both Ingress and Egress Sides

Figure 18 shows another possible IPFE configuration where the IPFE uses connection “initiator” functionality to set up the connections between the DA-MPs and the DIAMETER servers (the HSS simulator in this figure). The IPFE bandwidth requirements need to be calculated separately for the egress side. If all of the DSR traffic goes through an IPFE on both ingress and egress sides, then the required IPFE bandwidth is:

$$(2 * MPS) * (\text{average DIAMETER message size including IP overhead})$$

Indicative Alarms / Events

In this section, only the key recommended metrics for managing the performance of the IPFE are discussed. There are many more measurements available on the IPFE, and these can be found in [2].

Measurements

The key metrics for managing the IPFE blades are:

Table 24: IPFE Metrics

Measurement ID	Name	Group	Scope	Description	Recommended Usage	
					Condition	Actions
5203	RxIpfeBytes	IPFE Performance	Server Group	Bytes received by the IPFE	If the number of (bytes * 8 bits/byte)/(time interval in s) is > benchmarked capacity (Gbps)	If the traffic is expected to grow then, consider adding an additional IPFE pair
31052	CPU_UtilPct_Average	System	System (IPFE)	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% occupancy, OR exceeds 60% occupancy when running without an active mate.	Contact Oracle
31056	RAM_UtilPct_Average	System	System (IPFE)	The average committed RAM usage as a percentage of the total physical RAM.	If the average Ram utilization exceeds 80% utilization	Contact Oracle

Suggested Resolution

Horizontal scaling by adding an additional pair of IPFEs per DSR Signaling NE as indicated.



IDIH

Overview

The IDIH (IDIH application, IDIH mediation, IDIH Database VMs) are considered a best effort trouble shooting tool for the DSR. Benchmarking data is not currently provided for the IDIH VMs.

Suggested Resolution

Contact Oracle support.

Appendix A: DSR VM Configurations

The information shown below is 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 since the application profiles are tuned to this number of vCPUs and memory sizes. In particular the 70 GB of storage for each VM is an absolute minimum required to load the DSR images.

Table 25: VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redundancy Models	Affinity/Placement Rules (Per Site)	Notes
<i>DSR NOAM</i>	4	6	70	1 Pair	Active/Standby	2 VMs per DSR network in any site. VMs to be deployed on separate servers	
<i>DSR SOAM</i>	4	6	70	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/Standby model is used.
<i>DA MP</i>	12	16	70	16 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss	All DA-MP types (This row, IWF, SS7, vSTP) count against the same 16 DA-MP limit.
<i>DA MP w/IWF</i>	12	24	70	16 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss	The 24 GB RAM requirement is a hard minimum if the DA-MP VM will be used with the IWF application.
<i>vSTP MP</i>	8	6	70	1 per DSR NF	Active	Since only a single vSTP MP is supported in this release there are no affinity rules required.	Only 1 vSTP MP is supported in the 8.1 release. The vSTP MPs don't count against the 16 MP limit in a single OCDSR node, so a DSR can have up to 16 DA-MPs and 1 vSTP MP.
<i>IPFE</i>	6	16	70	2 pairs per DSR NF	Active/Standby	Each VM in a pair must be deployed on separate server.	Deployed in pairs. Max 2 pairs (4 VMs).

<i>SS7 MP</i>	12	24	70	16 per DSR NF	Active Cluster (N+0)	Should be spread over as many servers as possible to minimize capacity loss on server loss	
<i>SBR(s)</i>	12	25	70	8 Server Groups per SBR(s)	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.
<i>SBR(b)</i>	12	32	70	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.
<i>DP SOAM</i>	4	12	70	1Pair per DSR NF	Active/Standby	2 VMs per site. VMs to be deployed on separate servers.	
<i>DP</i>	6	10	70	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
<i>SDS NOAM</i>	4	32	300	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 geo-redundancy.
<i>Query Server</i>	4	32	300	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
<i>iDIH Application</i>	4	8	64	1 per Site	N/A since non-redundant	None, non-redundant	Optional component for DIAMETER traffic monitoring
<i>iDIH Mediation</i>	4	8	64	1 per Site	N/A since non-redundant	None, non-redundant	Optional component for DIAMETER traffic monitoring



<i>iDIH Database</i>	4	8	120GB + 100GB (ephemeral)	1 per Site	N/A since non-redundant	None, non-redundant	Optional component for DIAMETER traffic monitoring
----------------------	---	---	---------------------------	------------	-------------------------	---------------------	--

Appendix B: DSR VM Disk Requirements

This appendix provides guidance on the disk requirements for the OCDSR VMs. Characterizing disk requirements can be tricky since there are many variables that can affect disk usage, such as how many reports are being run on the OAM systems, or how often backups are run. Peak disk utilization can also very different from average disk utilization, for instance during backups or restore operations. While these guidelines are provided for the disk usage of the different VM types, customers should verify their disk usage under their own conditions since they are more driven by how the customer uses their system than by easier to calculate factors such as CPU utilization per MPS.

The OCDSR has been designed as a low disk-utilization application, with all critical call processing applications performed in memory. There is also no swap disk utilization in any of the VMs. As a background for all of these numbers, the OCDSR has been run for years on “bare metal” deployments with a single pair of industry-standard 10k RPM, 2½ inch disk drives in Raid 1. So even maximum sized OCDSR configurations run successfully on the approximately 120 IOPs provided by those disks. When run on higher-performance disk subsystems such as SSDs, high disk utilization tasks such as background report generation just complete faster. The notes for each VM type give some of the factors that can drive different disk utilization levels. For instance, the primary traffic handling VMs, the IPFEs and the different types of MPs, have a fairly constant disk utilization independent of the traffic level. This is because the primary disk utilization is for saving statistics, then forwarding them to the SOAM.

Table 26: VM Disk Utilization Characteristics

VM Name	Disk (GB)	Routine Disk	Peak Disk	Disk Usage Modes	Notes
		Utilization (IOPs) ¹	Utilization (IOPs) ²		
<i>DSR NOAM</i>	70	100	800	Periodic (30 second) small writes to collect statistics Large block reads to run reports and backups	Background disk utilization is mostly statistics collection from managed DSRs. Peak disk utilization driven by customer report generation and maintenance activities such backups.
<i>DSR SOAM</i>	70	100	800	Periodic (30 second) small writes to collect statistics Large block reads to run reports and backups	Background disk utilization is mostly statistics collection from managed DSRs. Peak disk utilization driven by customer report generation and maintenance activities such backups.
<i>DA MP</i>	70	50	500	Writes statistics to disk at 30 second intervals, reads them at 5 minute intervals to send to SOAM	Disk utilization is independent to traffic levels. Is affected by the size of the DSR configuration (number of connections for instance) and the utilization of features that create measurements such as ETGs and TTPs
<i>DA MP w/IWF</i>	70	50	500	Writes statistics to disk at 30 second intervals, reads them at 5 minute intervals to send to SOAM	Disk utilization is independent to traffic levels. Is affected by the size of the DSR configuration (number of connections for instance) and the utilization of features that create measurements such as ETGs and TTPs

<i>vSTP MP</i>	70	50	500	Writes statistics to disk at 30 second intervals, read them at 5 minute intervals to send to SOAM	Disk utilization is independent to traffic levels. Is affected by the size of the vSTP configuration such as the number of local and remote peers.
<i>SS7 MP</i>	70	50	500	Writes statistics to disk at 30 second intervals, reads them at 5 minute intervals to send to SOAM	Disk utilization is independent to traffic levels. Is affected by the size of the DSR configuration.
<i>IPFE</i>	70	20	100	Writes statistics to disk at 30 second intervals, reads them at 5 minute intervals to send to SOAM	IPFE has relatively few configuration items and statistics, and very low disk utilization.
<i>SBR(s)</i>	70	50	800	Synchronizes changes to in-memory database to disk. Mostly write application.	Peak disk utilization driven by recovery activities between active/standby servers.
<i>SBR(b)</i>	70	50	800	Synchronizes changes to in-memory database to disk. Mostly write application.	Peak disk utilization driven by recovery activities between active/standby servers.
<i>SDS NOAM</i>	300	100	800	Synchronizes changes to in-memory database to disk. Mostly write application.	SDS can maintain multiple copies of large subscriber database. Peak disk utilization is mostly driven by creating new backups.
<i>DP SOAM</i>	70	50	800	Synchronizes changes to in-memory database to disk. Mostly write application.	
<i>DP</i>	70	80	500	Synchronizes changes to in-memory database to disk. Mostly write application.	
<i>Query Server</i>	300	100	800	Synchronizes changes to in-memory database to disk. Reads are driven by customer queries.	The query server is not a real-time system. The amount of disk reads is driven entirely by manual customer database queries.

¹ The “routine” disk utilization is the minimum engineered IOPs for the proper functioning of the VM. Average disk utilization is typically lower.

² The “Peak” disk utilization is number of IOPs the VM is capable of using given sufficient resources.

Appendix C: VM Networking Requirements

This appendix 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, the average DIAMETER signaling traffic through a DA-MP is:

$$((\text{Average DIAMETER message size including IP overhead}) * Y) * (1 + ((2X-1)/X))$$

As an example, if the average DIAMETER message size is 2,000 bytes including overhead, the overall DSR MPS is 10,000 MPS, and the number of DA-MPs is three, the calculation would be:

$$(2,000 \text{ bytes} * 8 \text{ bits/byte} * 10,000 \text{ MPS}) * (1 + (2*3 \text{ DA-MPs} - 1)/(3 \text{ DA-MPs})) =$$

$$(160,000 \text{ kb/s}) * (2.66) = 426,666 \text{ kb/s per DA-MP}$$

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.

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

Table 27: VM Networking Utilization Characteristics

VM Name	Networks Supported	Management	Traffic	Notes
		Networks (Gb/s)	Networks (Gb/s)	
DSR NOAM	XMI	2	N/A	Activities such as backups can generate higher network utilization, but will run at the rate of the bandwidth available since they aren't real-time activities.
	IMI			
DSR SOAM	XMI	1	N/A	Activities such as backups can generate higher network utilization, but will run at the rate of the bandwidth available since they aren't real-time activities.
	IMI			
DA MP	XMI	0.2	MPS Dependent	See explanation above for how to calculate the signaling network traffic.
	IMI			
	XSI			
DA MP w/IWF	XMI	0.2	MPS Dependent	See explanation above for how to calculate the signaling network traffic.
	IMI			

	XSI			
<i>vSTP MP</i>	XMI IMI XSI	0.2	MPS Dependent	See explanation above for how to calculate the signaling network traffic.
<i>SS7 MP</i>	XMI IMI XSI	0.2	MPS Dependent	See explanation above for how to calculate the signaling network traffic.
<i>IPFE</i>	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.
<i>SBR(s)</i>	XMI IMI	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.
<i>SBR(b)</i>	XMI IMI	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.
<i>SDS NOAM</i>	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.
<i>DP SOAM</i>	XMI IMI	1.0	N/A	All of the subscriber data provisioned at the SDS NOAM is passed down to each DP SOAM, which then distributes the data to any attached DPs.
<i>DP</i>	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.
--------------	------------	-----	-----	--

Table 28 shows some guidelines for mapping the logical OCDSR networks (XMI, IMI, etc.) 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 28: Typical OCDSR Network to Device Assignments

VM Name	OAM (XMI)	Local (IMI)	Signaling A (XSI1)	Signaling B (XSI2)	Signaling C (XSI3)	Signaling D (XSI4)	Replication (SBR Rep)	DIH Internal
DSR NOAM	eth0	eth1						eth0
DSR SOAM	eth0	eth1						eth0
DA-MP	eth0	eth1	eth2	eth3	eth4-17	eth18		eth0
IPFE	eth0	eth1	eth2	eth3	eth4-17			eth0
SS7 MP	eth0	eth1	eth2	eth3	eth4-17	eth18		eth0
SBRB	eth0	eth1				eth2		eth0
SBRS	eth0	eth1				eth2		eth0
vSTP	eth0	eth1	eth2	eth3	eth4-17			
iDIH App	xmi							int
iDIH Med	xmi	imi						int
iDIH DB	xmi							Int
SDS NOAM	eth0	eth1						
SDS NOAM	eth0	eth1						
DP	eth0	eth1						
Query Server	eth0	eth1						

Appendix D: Summary of Benchmark Data Recommendations

The information shown below 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 29: Benchmark Data Summary

Benchmark characteristics:

Benchmark Run	VMware/ESXi	Openstack/KVM
Application Software	DSR 8.1 (running Oracle Linux)	DSR 8.1 (running Oracle Linux)
Host VM	VMware 6.0	Openstack Newton/KVM
HW	Oracle Server X5-2	Oracle Server X5-2, HP Gen9v1
VM Profiles/Flavors	Appendix A: DSR VM Configurations	Appendix A: DSR VM Configurations

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)	Diameter Agent Message Processor	MPS	18,000
DA MP (Database)	Diameter Agent Message Processor	MPS	16,000
DA MP (Statefull)	Diameter Agent Message Processor	MPS	13,000
DA MP w/ IWF	Diameter Agent Message Processor with IWF	MPS	12,000
vSTP	Virtual STP Message processor	MPS	10,000
IPFE	IP Front End	Gb/s per IPFE pair	2.0 ¹
		Connections per IPFE Pair	2,000
		Target Set Address (TSA) per IPFE pair	16
SS7 MP	SS7 Message Processor for MAP Diameter interworking function	MPS	12,000
SBR(s) (Single Server Group)	Subscriber Binding Repository (session) for Policy DRA	Diameter sessions	16,000,000
		MPS	50,000
SBR(s) (Max 8 Server Groups)	Subscriber Binding Repository (session) for Policy DRA	Diameter sessions	128,000,000
		MPS	400,000
SBR(b) (Single Server Group)	Subscriber Binding Repository (binding) for Policy DRA	Subscriber Bindings	50,000
		MPS	50,000
SBR(b) (Max Network-wide with 8 Server Groups)	Subscriber Binding Repository (binding) for Policy DRA	Subscriber Bindings	128,000,000
		MPS	400,000
DP SOAM	Database Processor Site (node) Operation, Administration, Maintenance for address resolution and subscriber location functions	VM	1+1



DP	Database Processor for address resolution and subscriber location functions	MPS requiring DP lookups (usually 50% of FABR traffic)	80,000
SDS	Subscriber Database Processor for address resolution and subscriber location functions	Routing Entities (typically 2x subscriber count)	300,000,000
		NAI User Routing Entries	500,000
		Provisioning TPS	200
		XML Interface reads/second	200
Query Server	Allows customers to query FABR subscriber data via a MySQL interface	N/A	N/A

¹ See also Table 23 for factors that can lower this number.

Appendix E: Detailed Infrastructure Settings

Table 30: Detailed Infrastructure Settings

Attribute	KVM/Oracle X5-2	HP Gen9v1
Model	Oracle Server X5-2	ProLiant DL380c Gen9v1
Processor Type	Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz	Intel(R) Xeon(R) CPU E5-2680 v3 @ 2.50GHz
vCPUs	72 [2 CPU Sockets [18 x 2 Cores, each with Hyper threading Active]	48 [2 CPU Sockets [12 x 2 Cores, each with Hyper threading Active]
RAM	128 G [DDR4-2133]	256 GB (DDR4-2133)
CPU Cache Memory	45 MB (Intel® Smart Cache)	30MB (Intel® Smart Cache)
Number and Type of NICs	4 [Intel Corporation Ethernet Controller 10-Gigabit X540-AT2]	2 x HP 10Gb 2-port 560M 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)	HP Static High Performance Mode
HDD	2.3 TB of solid state drive (SSD) storage	HP Smart Array P244br Controller Disk Drive Interface 6Gb/s SAS (Serial Attached SCSI) 2 x HP 900GB 12G SAS 10K rpm SFF (2.5-inch) SC Enterprise

Appendix F: Small DSR VM Configuration

Many customers don't need the capacity provided by even a single pair of standard size DSR DA-MP VMs. The fixed configuration in this appendix is for customers that only need:

- 1) Relay or RBAR applications
- 2) 6k MPS or less (assuming an infrastructure equal to or better than the X5-2 processors used for the benchmark tests).
- 3) 5K or less RBAR entries
- 4) No IPFE

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

Table 31: Minimal System VM Configurations and Affinity Rules

VM Name	vCPU	RAM (GB)	Disk (GB)	Max Config	Redundancy Models	Affinity/Placement Rules (Per Site)	Notes
<i>DSR NOAM</i>	4	6	70	1 Pair	Active/Standby	2 VMs per DSR network in any site. VMs to be deployed on separate servers if possible.	Two NOAMs are always required to support upgrades, but they can be on the same server if only one server is available.
<i>DSR SOAM</i>	4	6	70	1 Pair per DSR NF	Active/Standby	2 VMs per site. VMs to be deployed on separate servers if possible.	Two SOAMs are always required to support upgrades, but they can be on the same server if only one server is available.
<i>DA MP</i>	4	8	70	1 or 2 per DSR NF	1 or 1+1	The 1+1 configuration should have the DA-MPs on different servers	There's little value in have redundant DA-MPs (1+1) if they are on the same server since the server is the thing most likely to fail.

Appendix G: DSR DA-MP Profiles

DA-MP profiles are used to set the various parameters associated with performance of the DA-MP such as:

- 1) Internal resource allocations such as the number of various types of threads created by that VM.
- 2) The absolute limit of traffic that a DA-MP will handle, set to 2.5 times the stated capacity of the VM.
- 3) Fixed alarms for various traffic conditions such as message rates, average message sizes and hold times.
- 4) User-configurable values for message discard percentages at the different DA-MP congestion levels.

In the cloud environment the correct profile should be selected based on the expected DA-MP throughput. This throughput will be based on three major factors:

- 1) The performance of the customer's infrastructure relative to the infrastructure used for the benchmarking described in this document.
- 2) The mix of application types (relay, database, session) being run on the DA-MP.
- 3) The VM configuration such as the number of vCPUs.

These factors are useful for estimating the VM performance in advance, but the customer should also monitor the actual performance of the different VM types using the KPIs, alarms and measurements described for each VM type in the previous sections.

Infrastructure Performance Differences

The benchmarks in this document were done on Oracle X5-2 servers configured with Intel Xeon X5-2699v3 processors. When trying to estimate the performance of a different type of server with a different processor type, the main factor is the single threaded performance of the chip, such as SPECint® 2006 from Spec.org. The overall performance of the server/chip is less important since the VM configurations will assign the same number of vCPUs/threads to the VM independent of the number of threads supported by processor (i.e. you can get more VMs on a server with more vCPUs, but the performance for a given number of vCPUs is set by the single-thread performance). Historically, within a couple of Intel processor generations the performance of the OCDSR VMs varies relatively linearly with the difference in single-thread performance, i.e. a processor with a SPECint2006 rating of 66 will have 10% better OCDSR performance than a processor with a SPECint2006 rating of 60.

Traffic Mix

The different classes of applications tested (Relay, Database, Statefull) have significantly different MPS results for a fixed infrastructure configuration (server type, VM size). The capacity for an OCDSR running a mixture of these types can be calculated by using percentage of traffic of a given type. For example, using the values from Table 29, consider an OCSR with the following traffic mix:

- 1) Relay 40% (18k MPS)
- 2) Database (30%) (16k MPS)
- 3) Statefull (30%) (13k MPS)

The effective throughput for this traffic mix would be:

$$(40\%*18k) + (30\%*16k) + (30\%*13k) = 15.9k \text{ MPS}$$



The caveat to this calculation is that the percentage for the lower performing traffic types (Database and Statefull) should reflect their worst case values, for instance to handle traffic spikes of that traffic type caused by failure conditions of other components in the network.

The performance of the vSTP application is calculated separately since it has dedicated MPs that run only vSTP traffic.

VM Configuration

The tested VM configurations are given in Appendix A and Appendix G. It's recommended that these VM configurations are used as tested for the most predictable results. If for some reason the customer desires to change them for a given installation, here are some guidelines that should be kept in mind:

- 1) The installation and upgrade procedures for the VMs requires a minimum of 70GB of disk storage. While assigning less storage than this may appear to work during installation, it will likely cause failures during the upgrade procedures.
- 2) The IWF function requires 24Gbytes of memory to run. It will not come into service with any smaller memory allocation.
- 3) Adding vCPUs to the configurations may increase performance, but only up to a point since there may not be enough threads to efficiently take advantage of the extra vCPUs.
- 4) Reducing the vCPU counts should not be done for any VM except for the DA-MPs. The issue is that configurations that appear to run fine under normal traffic conditions may not have sufficient capacity under recovery conditions or under heavy loads (for instance running reports on a SOAM while it's doing a backup). The DA-MP vCPU number should not be lower than 4 vCPUs (see appendix G for an example small DSR configuration.)

Selecting a DA-MP Traffic Profile

The following DA-MP traffic profiles for cloud deployments are supported in OCDSR 8.1:

- 1) 6k MPS
- 2) 8k MPS
- 3) 12k MPS
- 4) 14k MPS
- 5) 16k MPS
- 6) 18k MPS
- 7) 21k MPS
- 8) 24k MPS
- 9) 27k MPS
- 10) 30k MPS

It is possible to change the DA-MP profile assigned to a DA-MP VM, however the DA-MP application will have to be restarted for it take effect. Since the results of the traffic mix calculation and any adjustments due to different processor performance are unlikely to land exactly on any of these numbers, the question is whether to use the next



larger or next smaller DA-MP profile. In general, the selecting the next larger DA-MP profile (for instance selecting the 16k MPS profile for the 15.9k MPS calculated in the Traffic Mix example above) will provide the best solution.

However, some considerations:

- 1) While the larger profiles assign more system resources (message buffers, etc.) the differences between these resource allocations between any two consecutive profiles is relatively small (proportional to their MPS difference, or about typically about 20%). Since these allocated system resources are a relatively small portion of the over VM memory usage (which includes, for instance, the resources used by the Guest OS and the base OCDSR application), assigning the higher rated profile has ensures sufficient resource allocation. The standard VM profiles given in Appendix A have sufficient memory to support up to the 24k MPS profiles. Above that level additional memory may be required depending on the traffic mix.
- 2) Over a small range of DA-MP profiles (plus or minus one at least) the DA-MP profile will likely have no impact on the actual throughput of the VM. The DA-MP overload controls are driven by CPU percentage (i.e. actual utilization), not by any of these calculations. The only hard limit on the throughput of a given profile is 250% of the nominal value. For instance the 10k MPS DA-MP profile will start shedding traffic (using the same algorithms as DA-MP overload) at 25k MPS, even if the CPU percentage hasn't hit the CPU utilization level required to trigger the CL1 congestion level.
- 3) Selecting the DA-MP profile just under the calculated MPS capacity (for instance selecting the 14k DA-MP profile for the 15.9k MPS calculated in the Traffic Mix example above) will cause the various system alarms to happen earlier. For instance, the minor traffic alarm will trigger at 8400 MPS (60% of 14k) instead of at 9200 MPS (60% of 16k) if the next larger profile was selected. Thus selecting the smaller DA-MP profile (14k in tis example) will give more conservative alarming, and the higher profile (16K MPS) will give less conservative alarming. But as noted in the previous item, it won't affect the actual throughput of the VM since that's driven by CPU usage.
- 4) Selecting a DA-MP profile that's a lot larger than the nominal capacity calculated in the previous section won't increase the throughput of the VM, and will cause the capacity alarms to trigger later than expected. Selecting a DA-MP profile that's a lot smaller than the nominal capacity will cause the alarms to trigger sooner than expected, and may limit the capacity of the VM due to insufficient system resources being allocated.

