



**Hewlett Packard**  
Enterprise

Reference Architecture

# **HPE Reference Architecture for Oracle Database on HPE Superdome Flex with HPE Serviceguard for Linux and Oracle Data Guard**

# Contents

Executive summary.....	3
Introduction.....	3
Solution overview.....	4
HPE Superdome Flex.....	4
HPE 3PAR StoreServ 9450 Storage.....	6
HPE Serviceguard for Linux.....	7
Oracle Data Guard.....	8
HPE Application Tuner Express (HPE-ATX).....	9
Solution components.....	10
Hardware.....	11
Software.....	11
Application software.....	11
Best practices and configuration guidance for the solution.....	12
Install HPE Foundation Software.....	12
Configure kernel boot options.....	12
RHEL OS settings.....	12
Oracle parameters.....	12
HPE-ATX configuration.....	12
The SGLX database package was configured with an entry to start the Oracle listener processes with the round robin flat policy.....	12
Network configuration.....	13
Serviceguard cluster configuration.....	13
Capacity and sizing.....	15
Workload description.....	15
Analysis and recommendations.....	15
Summary.....	18
Implementing a proof-of-concept.....	18
Appendix A: Bill of materials.....	19
Appendix B: RHEL kernel tuning parameters.....	20
Appendix C: Network teaming configuration.....	21
Appendix D: Oracle user account limits.....	21
Appendix E: Oracle initialization parameters.....	21
Appendix F: Oracle listener configuration.....	22
Resources and additional links.....	25



## Executive summary

Demands for scalable and highly available database implementations continue to escalate in today's IT organizations. Faster transaction processing speeds, capacity-based scaling, increased flexibility, high availability (HA), and business continuity are required to meet the needs of a 24/7 business.

The HPE Superdome Flex Server delivers an unmatched combination of flexibility, performance, and reliability for critical environments of any size. HPE Superdome Flex protects vital workloads with superior reliability, availability, serviceability (RAS) and security features. HPE Superdome Flex coupled with HPE 3PAR StoreServ Storage arrays is an ideal scale-up configuration. With the addition of HPE Serviceguard for Linux (SGLX) and Oracle® Data Guard (ODG), a highly available solution can be built to satisfy the needs of most Oracle database environments.

While an Oracle RAC solution is highly available and scalable, it is expensive to deploy, requires Oracle RAC skills, and is complex to manage. The solution described in this Reference Architecture provides an alternative for those customers who desire a simpler and more cost-effective environment. Deploying Oracle 18c with an HPE Superdome Flex nPartition (nPar) configuration, combined with HPE SGLX and ODG, provides a high-performing and cost-effective alternative when application availability is acceptable as per the defined service level agreements (SLA). Testing conducted for this Reference Architecture (RA) demonstrates that an Oracle Database deployed with ODG for replication can accomplish a role transition from a primary server to a standby in under one minute by utilizing the automation packages provided by HPE SGLX. Best practices for configuring this solution are also provided, along with performance comparisons of various options that can be utilized in an ODG environment.

**Target audience:** This RA is designed for IT decision makers and IT professionals who use, program, manage, or administer Oracle databases that require scalable performance and HA. Specifically, this information is intended for those who evaluate, recommend, or design new and existing IT high-performance architectures. Additionally, CIOs may be interested in this document as an aid in implementing highly available Oracle databases combined with HPE SGLX and the performance characteristics associated with that implementation.

**Document purpose:** The purpose of this document is to describe an RA that organizations can utilize to plan for their deployment of Oracle 18c on HPE Superdome Flex and HPE 3PAR StoreServ all-flash storage and also plan for highly available Oracle databases using HPE SGLX and ODG. The performance results presented are specific to this configuration. Note that variations in system configuration, HPE storage and third-party storage may also be utilized to implement this solution. HPE can provide guidance on alternative configurations.

This RA describes solution testing completed in August 2019.

## Introduction

As enterprises deploy more and more mission-critical workloads in their data centers, the need for highly available, high-performance solutions increases. As transactions processed using mission-critical data increase, data centers need to be able to quickly and easily scale to accommodate the growing need while maintaining availability, performance, and reliability. Enterprises need infrastructure solutions that maximize application uptime, standardize or consolidate hardware, and work with different availability designs.

Traditional hardware requires extended downtime for maintenance, and can often be difficult to configure with the HA required for mission-critical workloads. In addition, many data centers have hardware from multiple vendors, which can create compatibility issues. Due, in part, to this piecemeal arrangement, planning for resource additions to traditional solutions can be difficult. Adding new resources can be time consuming, difficult, and can necessitate downtime. Previous IT models also make it difficult for administrators to tailor existing resources to best fit the needs of various workloads.

For mission-critical workloads, the HPE Superdome Flex provides the ease of scale-up combined with the capacity required to also run mixed workloads. Since it's easily scaled by simply adding more chassis, there is no migration to a new server, you simply add the new resources to the existing partition.

---

### Note

HPE Superdome Flex nPars are recognized as Oracle-approved hard partitions.

---



Mission-critical resiliency is provided through end-to-end implementation of processor RAS features, redundancy of key system components, and advanced system software, to help ensure the server is up and running 24 x 7.<sup>1</sup> Additionally, in rare cases, where a problem happens that would typically bring down the entire server, the HPE Superdome Flex's modular design enables administrators to localize an issue at a chassis level, in order to continue operations.

In order to scale up a system such as the HPE Superdome Flex, a storage system with similar, scale-up capability is required. The HPE 3PAR portfolio can scale to four controller nodes for the midrange products, and eight nodes for the high-end products. In addition, capacity can be scaled from a few terabytes to over 80 PB in a four-system federation with a common OS, feature set, and management.

With HPE InfoSight predictive analysis technology, and the ability to group storage arrays together for management and aggregation, these arrays are perfectly suited for mission critical environments.

HPE InfoSight not only monitors the storage environment for problems and potential hazards, it also proactively predicts problems before they occur, and in some cases can resolve these problems without intervention. HPE InfoSight can see across your entire infrastructure, giving you a view that you may have never had before, transforming your whole support model.

In addition to hardware design considerations for a highly available Oracle solution, one must consider the software design options. At one end of the availability continuum is Oracle RAC, a relatively expensive option that adds complexities, but provides near-instant recovery options. With other less expensive and less complex options, recovery times are longer but are often still acceptable depending on an organization's SLA. One example is HPE SGLX, a software solution that provides HA and disaster recovery (DR) capabilities. Pairing SGLX with ODG is an approach that provides further resiliency from system failures and shortens the recovery time.

Every level of a data center - server, storage, network, OS, virtualization layer, virtual machine guest - must be monitored at all times. When SGLX detects a failure, it follows shutdown procedures, and moves applications to an available, appropriate resource - whether physical or virtual - to bring them back online. SGLX provides monitoring, failover, data loss prevention, and limited-to-no planned downtime. This RA shows how SGLX works in the test environment during an unplanned failure.

Testing highlighted in this RA details the HA capabilities of two HPE Superdome Flex partitions coupled with HPE 3PAR Storage, HPE SGLX, and ODG. The hardware and software components of this solution are detailed below.

## Solution overview

This solution included HPE Superdome Flex nPartitions with HPE 3PAR StoreServ 9450 All Flash storage, running Oracle Database 18c. HPE SGLX and ODG were utilized to provide clustering software and automated replication and failover of the Oracle Database in the event of a component failure. The HPE Application Tuner Express software was utilized to achieve maximum performance in a non-uniform memory access (NUMA) environment. Note that variations in system configuration, HPE storage, and third-party storage may be utilized to implement this solution. This RA is a companion to the previously published paper, "[HPE Reference Architecture for Oracle 18c OLTP and OLAP workloads on HPE Superdome Flex and HPE 3PAR Storage](#)".

## HPE Superdome Flex

The mission-critical HPE Superdome Flex platform is ideal for managing the exponentially growing data coming in and out of a business. The HPE Superdome Flex was designed with a modular approach to adapt to ever-increasing data management needs. With the HPE Superdome Flex, your Oracle database infrastructure can grow along with your data. The platform provides the following noteworthy characteristics:

- Modular architecture that scales seamlessly from 4 to 32 sockets in 4-socket increments in a single system
- Shared memory capacity from 768GB up to 48TB
- Intel® Xeon® Scalable processors 1st or 2nd generation
- Extreme RAS capabilities not available on other standard platforms such as:
  - Best-in-class predictive fault-handling Analysis Engine, predicts hardware faults and initiates self-repair without operator assistance
  - Firmware First approach to log analysis ensures error containment at the firmware level, including memory errors, before any interruption can occur at the OS layer

<sup>1</sup> For further details, see [HPE Superdome Flex server architecture and RAS](#)



- Mission-critical resiliency from end-to-end implementation of processor RAS features, to redundancy of key system components to advanced system software
- O/S and virtualization spanning SUSE Linux Enterprise Server, Red Hat® Enterprise Linux®, Oracle Linux, Microsoft Windows, Oracle VM, VMware® and KVM
- Unbounded I/O

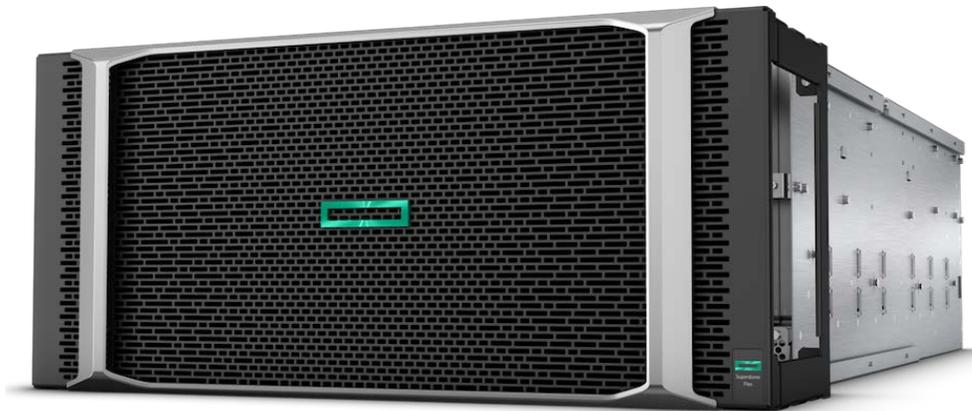
The HPE Superdome Flex can be configured with an embedded management controller (for use with up to two chassis only) or an external rack management controller (RMC).

The RMC is the administrative node for the HPE Superdome Flex. The RMC includes one network port for administrative access to the system console. All administrative functions of the HPE Superdome Flex are performed through the management controller:

- Configuring system partitions
- Network configuration
- Booting, rebooting, and shutting down the system
- Viewing hardware resources
- HPE OneView communicates with RMC (via Redfish) to monitor the system, whether on site or remotely.

Computing resources on the HPE Superdome Flex are assigned using partitions. Partitions are currently aligned along the chassis boundaries within the HPE Superdome Flex environment. This greatly enhances the ease in which CPUs can be added and memory can be expanded to the installed operating system partition.

The solution highlighted for this RA features two single-chassis partitions. Additional chassis and partitions can be added at a later time by simply adding them to the existing server. Figure 1 depicts a single Superdome Flex chassis.



**Figure 1.** Front and side view of HPE Superdome Flex chassis



## HPE 3PAR StoreServ 9450 Storage

HPE 3PAR StoreServ 9450 Storage is an enterprise-class, all-flash array that helps consolidate primary storage workloads without compromising performance, scalability, data services, or resiliency. This HPE 3PAR model is built for all-flash consolidation, delivering the performance, simplicity, and agility needed to support a hybrid IT environment. It can scale up to 6000 TiB of raw capacity, and is capable of over 2 million IOPS at sub-millisecond latency. These capabilities are complemented by enterprise-class, Tier-1 features and functionality. HPE 3PAR StoreServ All Flash is designed for 99.9999% availability with full hardware redundancy, supporting availability objectives for the most demanding environments. Enhanced storage capabilities provide continuous data access and the HPE 3PAR Priority Optimization software offers fine-grained QoS controls to ensure predictable service levels for all applications without physical partitioning of resources. Figure 2 shows the HPE 3PAR StoreServ 9450 Storage.



Figure 2. Front view of HPE 3PAR StoreServ 9450 Storage

## HPE InfoSight for HPE 3PAR

In addition, HPE 3PAR customers can benefit from HPE InfoSight for HPE 3PAR. HPE InfoSight is an industry-leading predictive analytics platform that brings software-defined intelligence to the data center with the ability to predict and prevent infrastructure problems before they happen. HPE InfoSight for HPE 3PAR provides the following capabilities:

- **Cross-stack analytics** – HPE InfoSight Cross Stack Analytics looks across the stack, enabling visibility into storage, network and host CPU and memory resources. It allows customers to quickly resolve performance problems and pinpoint the root cause of issues between the storage and host virtual machines (VMs). It also provides visibility to locate “noisy neighbor” VMs that are consuming a large amount of storage resources.
- **Global visibility** – All HPE 3PAR customers with systems that are remotely connected will see detailed performance trending, capacity predictions, health checks and best practice information across all of their HPE 3PAR arrays.
- **Predictive support** – HPE InfoSight constantly watches over 3PAR customer’s environments in order to predict problems before they occur. If it uncovers an issue, it can resolve the issue and prevent other systems from experiencing the same problem. HPE InfoSight automatically creates a support case when it detects a potential issue and will then close the case when the issue is resolved.



### HPE 3PAR benefits for Oracle

Oracle customers can use HPE 3PAR StoreServ Storage to address their most significant challenges including:

- **Performance** – HPE 3PAR delivers high throughput and low latencies in multi-tenant, mixed-workload Oracle environments. With industry leading performance and sub-millisecond latencies, HPE 3PAR provides high transactions-per-second (TPS) and minimal wait times for Oracle online transaction processing (OLTP) workloads, using features such as Priority Optimization and Adaptive Flash Cache.<sup>2</sup>
- **Efficiency and Data Reduction** – In many Oracle environments, overprovisioning the primary database has become a matter of survival. As Oracle databases grow, every added core creates additional license and support fees. HPE 3PAR helps reduce Oracle sprawl and simplifies instance consolidation, driving higher TPS and providing a reduced storage footprint. What's more, HPE 3PAR Adaptive Data Reduction technologies, such as compression, can boost storage efficiency while helping enterprises bypass costly Oracle compression license fees.
- **High Availability and Data Protection** - Many Oracle environments face challenges with a growing primary database with an increasing number of applications writing to that database. As databases get larger, backup windows, recovery point objectives (RPO), and recovery time objectives (RTO) become harder to meet. HPE Recovery Manager Central for Oracle accelerates Oracle database cloning with fewer steps and significantly less storage capacity compared to traditional Oracle Recovery Manager (RMAN) based cloning. HPE Recovery Manager Central combines these copy data management capabilities with the ability to take direct snapshot copies of the Oracle database volumes for fast data protection purposes. Also, it can enable significantly faster backup to on-premises HPE StoreOnce Systems as the target backup appliance than Oracle RMAN. HPE Storage offers Recovery Manager Central for Oracle (HPE 3PAR RMC-O) software at no additional cost. For more information about this solution, see [HPE Reference Configuration for Copy Data Management and Data Protection of Oracle environments with HPE Recovery Manager Central for Oracle](#). Also, the HPE StoreOnce Catalyst Plug-in for Oracle is bundled with HPE StoreOnce and is tightly integrated with Oracle RMAN, giving the Oracle database administrator complete visibility and control for backup and recovery tasks. For more information about this solution, see [HPE Reference Architecture for Comprehensive Oracle Backup, Restore and Disaster Recovery using HPE RMC and HPE StoreOnce](#). In addition, HPE 3PAR Peer Persistence can be deployed with Oracle RAC to provide customers with a highly available stretched cluster. Peer Persistence can also be utilized in conjunction with HPE SGLX in a single instance Oracle environment to provide automated fail-over protection with no data loss (RPO of zero).

Taken together, these features help Oracle database and storage administrators manage even the most demanding Oracle environments, delivering the performance, data protection, efficiency, and HA needed to keep critical applications and business processes up and running.

For more details about the HPE 3PAR features that benefit Oracle environments, see [Best Practices for Oracle Database on HPE 3PAR StoreServ Storage](#).

### HPE Serviceguard for Linux

HPE Serviceguard for Linux increases the availability and uptime of your mission-critical applications, allows business continuity during maintenance windows, and minimizes the impact of unplanned outages. Providing both HA and DR features, HPE SGLX packages applications and other services with their associated resources and monitors the entire package for any failure. Each package is monitored for faults related to hardware, software, OS, virtualization layer, virtual machine guests, network, and storage. When any failure is detected, SGLX shuts down the application to quickly and smartly relocate the application or service to another system with the necessary resources to bring it into production again. An HPE SGLX cluster allows for multiple node failures using redundant hardware to remove any single points of failure. This RA focuses on showing how quickly the database on the failed node returns to production after the takeover event, as well as the performance of the databases before and after the failover event.

In addition, HPE SGLX uses lock LUN or quorum server arbitration mechanisms to prevent data corruption and loss. A lock LUN can be used for clusters of up to four nodes, while a quorum server can be used for clusters of any size and must be run on a system outside of the cluster. The solutions in this RA use the lock LUN configuration because it provides faster failure detection. SGLX also minimizes planned downtimes using its Live Application Detach (LAD) and Rolling Upgrades features to perform maintenance on clusters and install upgrades for any OS and application without downtime. Take advantage of the Cluster Verification technology to find and fix cluster configuration issues before they advance and cause unplanned downtime.

### SGLX-based HA and DR solution for Oracle Database 18c

HPE SGLX provides automatic HA and DR protection for Oracle databases. This solution manages and completely automates the transition of Oracle databases from one cluster node to another in the event of a failure. Oracle is installed on each of the cluster nodes and the databases are configured such that the database data is located on shared storage which can be accessible from all the cluster nodes. HPE Serviceguard deep

<sup>2</sup> Adaptive Flash Cache utilizes SSDs as a cache for slower storage devices, and therefore is not needed in the all-flash storage arrays such as the HPE 3PAR StoreServ 9450.



integration with Oracle allows SGLX to start, stop, monitor, and recover the databases on the adoptive cluster nodes in case of failures. In this solution, SGLX recovers the database by restarting the Oracle Database on the adoptive cluster node.

---

**Note**

In this RA, the configuration where SGLX is used to recover the database by automatically restarting the database on another node is referred to as the “SGLX only” configuration. This is to distinguish it from the configuration where both SGLX and ODG are used.

---

**Oracle Data Guard**

Oracle Data Guard enables Oracle databases to survive disasters and data corruption using services that create, maintain, manage, and monitor the standby databases. ODG maintains these standby databases as consistent copies of the production database. If the production database is unavailable, in the case of a planned or unplanned outage, ODG can allow a standby database to take over the production role, thus reducing the downtime associated with the outage. ODG can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability. An ODG configuration consists of one production database, known as the primary database, and up to thirty standby databases. Each standby database is transactionally consistent with the primary (production) database, maintained by transmitting redo data from the primary database to the standby database.

There are three main types of standby databases: physical, logical, and snapshot. Physical databases are physically identical copies of the primary database on a block-for-block basis. The entire database, including schema and indexes, are identical. ODG uses Redo Apply technology to keep the standby database synchronized with the primary database. In addition, a physical standby database can be used for read-only access while maintaining synchronization, allowing applications to use standby databases for data protection and reporting.

Logical databases have the same logical information as the primary database, but the physical structure of the data can differ. ODG uses a SQL Apply feature to keep logical standby databases in sync with primary databases using SQL statements. Logical standby databases allow for upgrading Oracle Database software in a rolling fashion with next to no downtime.

ODG also has three protection modes: maximum availability, maximum performance, and maximum protection. The maximum availability protection mode is the highest level of data protection possible without compromising the availability of a primary database. In this mode, all redo data needed to recover the primary database is collected by a synchronized standby database before any transactions are committed. However, if the primary database is unable to write the redo data to a standby database, it will revert to maximum performance mode until it is able to start writing again. This will ensure that data will remain after a primary database failure so long as the standby database is intact.

The maximum performance protection mode offers the maximum amount of protection possible without affecting the performance of the primary database. This mode, which is the default mode, allows transactions to commit once the redo data is written to an online log, but before the data has been transferred to a standby database. Thus, the primary database isn't affected by slow redo writing to standby databases.

The maximum protection mode is similar to the maximum availability mode, but any time the primary database cannot write to the standby database, the primary database will be shut down. Due to the impact this mode could have on availability of the primary database, Oracle recommends any primary database set in this mode have at least two standby databases to help minimize downtime. This mode ensures no data loss on a primary database failure.

The database used in this RA used a physical standby database, and the protection mode was set to maximum availability, except for some performance tests that were run with the maximum performance protection mode to demonstrate the difference in performance between the two protection modes. Note that the maximum availability mode requires configuring synchronous redo transport, while the maximum performance mode requires configuring asynchronous redo transport.

**SGLX-based HA/DR solution for Oracle Database using ODG**

HPE SGLX along with ODG provides HA and DR protection for Oracle databases with much faster recovery times as compared to the solution with just SGLX (without ODG). HPE SGLX manages and completely automates the role transition process of ODG databases in the event of a failure. In this solution, primary and standby databases are configured with ODG. SGLX is utilized to start, stop, and monitor the databases and monitor the replication between primary and standby databases. In case of failures, the solution automatically recovers the Oracle Database by promotion of the standby database instance to primary. This mode of recovery is much faster compared to the restart-based solution as described above. The databases can be located on the same premises or in geographically dispersed data centers. SGLX performs RPO sensitive automatic role management to recover from failures.



Integrating ODG with SGLX has the following advantages:

- Provides HA and DR with fully automatic and fast recovery for Oracle databases
- Monitoring for primary and secondary DB instances and the ODG processes
- RPO sensitive recovery
- Built-in monitoring capabilities to check system resources, such as network, volume groups, files systems, etc. and failover in the case of failure of any of these components
- Protect and recover the application stacks in addition to databases
- Failure and status notifications
- Advanced features to minimize downtime for planned maintenance

---

**Note**

While ODG includes a fast-start feature to automate the role transition, HPE SGLX provides the following advantages over the fast-start feature. Note that Data Guard Broker is a tool for managing ODG configurations.

---

- Includes the following features not available with ODG: critical process monitoring, DB hang detection, failure notification (through email), and backup mode support.
- Data Guard Broker is required for fast-start failover and automatic role management. The broker by itself can be a single point of failure.
- If the Data Guard Broker is down when there is a need to make a decision in terms of promoting a standby, the fast-start failover will not happen, leading to downtime.
- Support for RPO sensitive failover is not available with Data Guard Broker.

**HPE Application Tuner Express (HPE-ATX)**

HPE-ATX is a utility for Linux® users to achieve maximum performance when running on multi-socket servers. Using this tool, you can align application execution with the data in memory resulting in increased performance. HPE-ATX is designed to improve performance of NUMA unaware applications without requiring any changes to the application itself. It helps multi-process and multi-threaded applications running on multi socket machines to achieve better NUMA placement of processes and threads that are related and share memory segments.

HPE-ATX offers the following launch policies to control the distribution of an application's processes and threads in a NUMA environment:

- Round Robin: Each time a process (or thread) is created, it will be launched on the next NUMA node in the list of available nodes. This ensures even distribution across all of the nodes.
- Fill First: Each time a process (or thread) is created, it will be launched on the same NUMA node until the number of processes (or threads) matches the number of CPUs in that node. Once that node is filled, future processes will be launched on the next NUMA node.
- Pack: All processes (or threads) will be launched on the same NUMA node.
- None: No launch policy is defined. Any child process or sibling thread that is created will inherit any NUMA affinity constraints from its creator.

For Oracle workloads, we strongly recommend launching the database listeners using HPE-ATX with the "ff\_flat" or "rr\_flat" process launch policy while affinizing to all the NUMA nodes where the database instance is running. Note that the HPE Serviceguard integration for Oracle provides a mechanism to start the listeners under the control of ATX within the Serviceguard database package.

HPE-ATX is fully supported by Hewlett Packard Enterprise and can be downloaded from the HPE [My License Portal](#).



### Solution components

This solution features the HPE Superdome Flex with two nPartitions and HPE 3PAR StoreServ 9450 all-flash storage array. A diagram of the hardware components is shown in Figure 3.

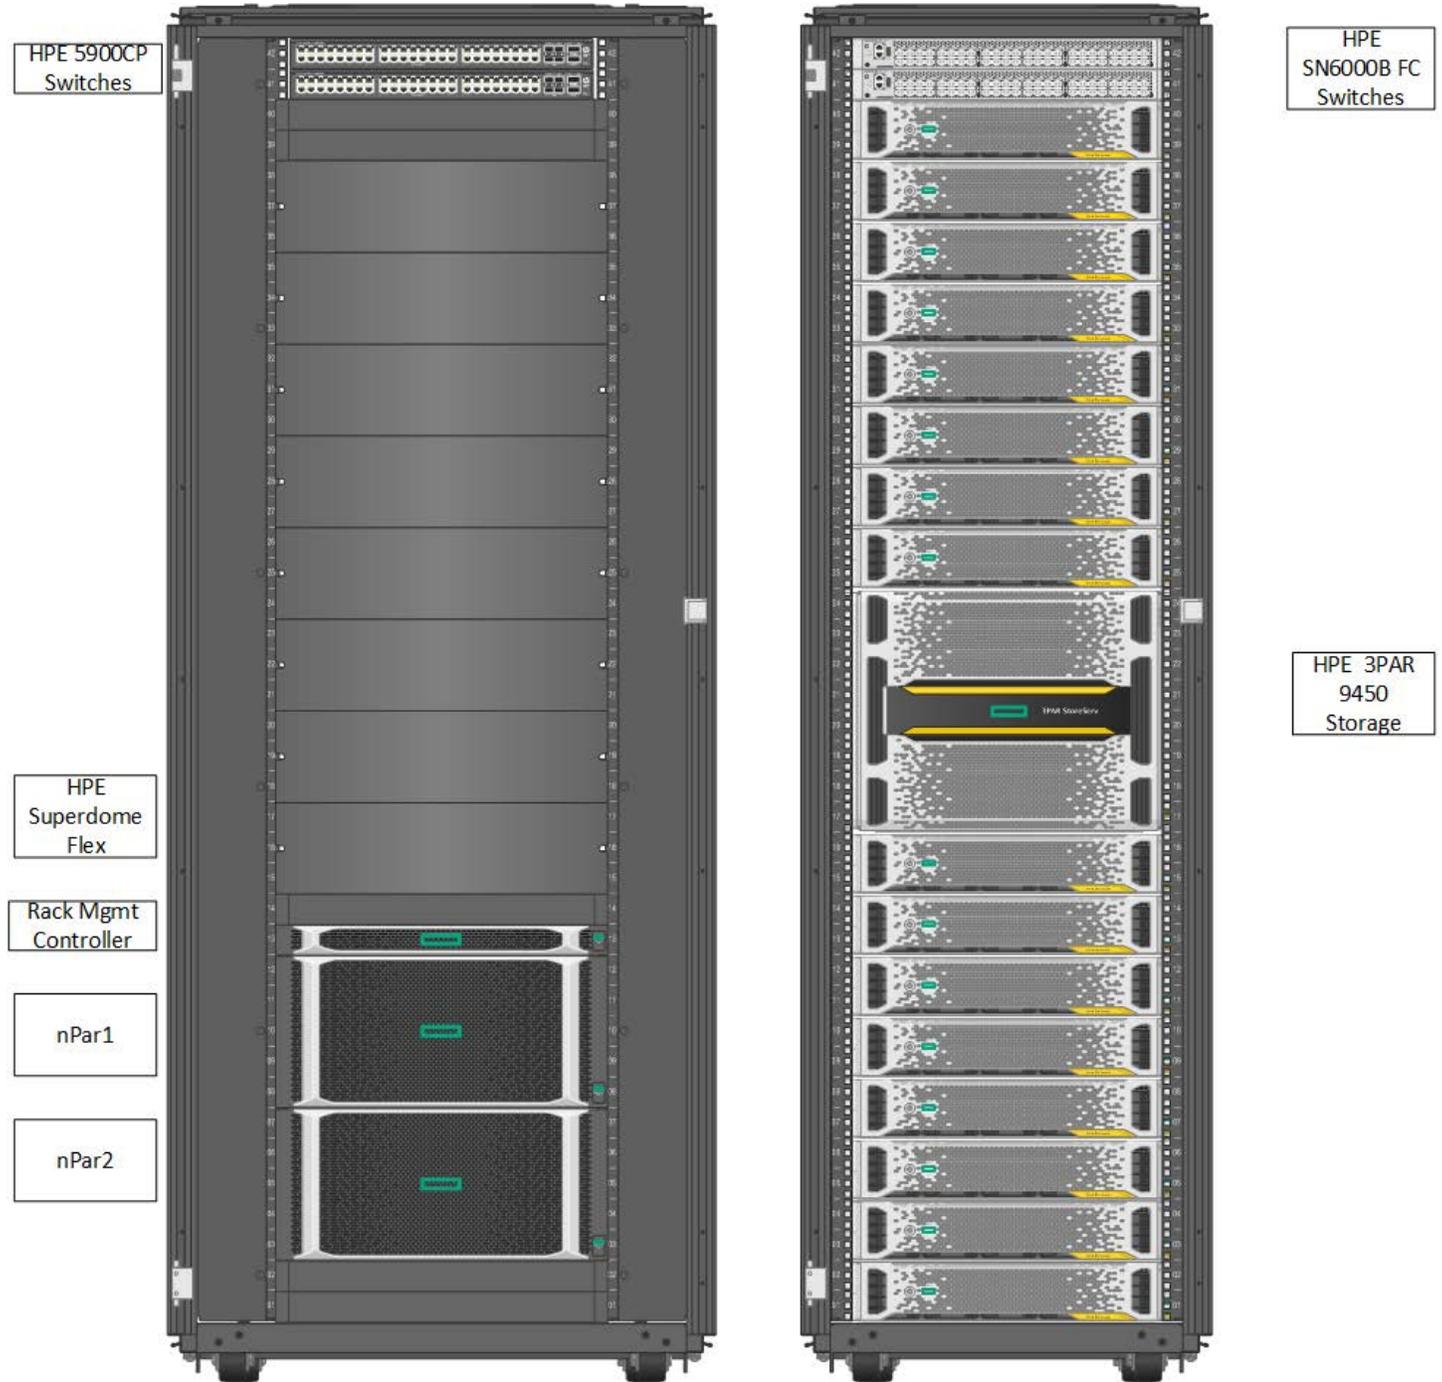


Figure 3. Solution hardware diagram



## Hardware

A four-socket chassis is the core building block for HPE Superdome Flex. Expanding beyond four sockets requires the use of expansion chassis. Each chassis supports four processors, 48 memory DIMM slots, up to 16 I/O slots, eight fans, and four power supplies. Up to seven additional chassis can be added to a single HPE Superdome Flex. For this RA, the server had one base chassis and one partition chassis, with each one configured as an nPar. Each chassis included the hardware listed in table 1.

**Table 1.** Server configuration

### HPE Superdome Flex nPar Configuration

<b>Processor</b>	Four 28-core Intel Xeon 8180 processors at 2.50GHz
<b>Memory</b>	3TB memory (48 x 64GB HPE DDR4 SmartMemory LRDIMMs)
<b>FC HBAs</b>	4 X 16Gb HPE SN1200E 16Gb DP FC HBA
<b>Built-in Network Adapters</b>	1 x 10Gb dual port Ethernet adapter, 1 x 1GbE dual port Ethernet adapter
<b>Additional Network Adapters</b>	2 x 10/25Gb dual port 631SFP28 Ethernet adapter
<b>Boot disk</b>	2 x 480GB SATA 6G MU SSD

Table 2 shows the configuration of the HPE 3PAR StoreServ 9450 all-flash storage array.

**Table 2.** Storage configuration

### HPE 3PAR StoreServ 9450 all-flash storage array

<b>Nodes</b>	4 nodes
<b>Cache</b>	896GiB
<b>Drive enclosures</b>	16
<b>Drives</b>	160 x 400GB SAS SFF SSDs
<b>FC ports</b>	8 x 4-port 16Gb Fibre Channel HBAs

## Software

The following software was configured on the system:

- Red Hat Enterprise Linux version 7.6
- HPE Foundation Software 1.5
- HPE Serviceguard for Linux A.12.40.00
- HPE Application Tuner Express 1.0.1-103.15

## Application software

Oracle Database 18c Enterprise Edition (18.3.0.0.0) was used for this RA. Note that Oracle Data Guard is included with Oracle Database Enterprise Edition, but Active Data Guard is a separately licensed product. Oracle Data Guard was used for this RA.



## Best practices and configuration guidance for the solution

### Install HPE Foundation Software

The HPE Foundation Software must be installed on the server after installing the OS. This software consists of packages designed to ensure the smooth operation of the server. It includes Data Collection Daemon (DCD) for Linux, an agentless service that proactively monitors the health of hardware components in the server. For instructions on installing this bundle, see the [HPE Superdome Flex Server OS Installation Guide](#).

### Configure kernel boot options

The HPE Foundation Software sets some kernel boot options that are key to optimal performance. Some additional parameters were also set manually. The boot options listed in table 3 were used for optimal performance for this RA. Note that `numa_balancing=0` and `transparent_hugepage=never` are both recommended by Oracle. When NUMA balancing is enabled, the kernel migrates a task's pages to the same NUMA node where the task is running. Due to the size of the Oracle SGA and frequency of task migrations, the migration of pages can be expensive, and optimal performance can be achieved by disabling NUMA balancing. Transparent huge pages are enabled by default in RHEL, and Oracle recommends disabling this setting to avoid memory allocation delays at runtime. Note that while dynamically-allocated transparent huge pages were disabled, statically-allocated huge pages were configured via the `vm.nr_hugepages` kernel parameter as described in the RHEL OS settings section below. The option `intel_idle.max_cstate=1` allows Cstate 1 transitions and encourages TurboBoost functionality.

**Table 3.** Kernel boot options for optimal performance

Kernel boot option	Description	How set
<code>numa_balancing=0</code>	Disable automatic NUMA balancing	HPE Foundation Software
<code>intel_idle.max_cstate=1</code>	Allow only C-state 1 transitions to encourage TurboBoost clock speed increases	HPE Foundation Software
<code>transparent_hugepage=never</code>	Disable transparent huge pages	HPE Foundation Software
<code>cgroup_disable=cpu</code>	Disable cgroups for the CPU	Manually added to <code>/etc/default/grub</code>

### RHEL OS settings

A complete list of the RHEL tuning parameters is shown in Appendix B. The `vm.nr_hugepages` kernel parameter was set large enough to include the Oracle SGA.

### Oracle parameters

A complete list of the Oracle parameters set for both environments (SGLX-only configuration and SGLX plus ODG configuration) is shown in Appendix E.

### HPE-ATX configuration

The SGLX database package was configured with an entry to start the Oracle listener processes with the round robin flat policy. This ensured that the listener processes were evenly distributed across all of the nodes (sockets) of the server. The following entry was included in the SGLX database package:

```
pev_atx_oracle_lsnr -p rr_flat -l /home/oracle/atx_oracle_lsnr.log
```

---

### Note

If the Oracle listener has already been started without HPE-ATX, the listener must be stopped prior to starting the SGLX database package.

---



## Network configuration

Separate networks are recommended for the Serviceguard heartbeat and ODG redo transport. Two networks were configured for the Serviceguard heartbeat to allow setting the Serviceguard member\_timeout setting to three seconds. Two (or more) networks are required to prevent an unnecessary failover due to the failure of a network device used for the heartbeat (note that one of the networks can be the data network, but there must be at least one dedicated network for the heartbeat traffic). In addition, network teaming was used to provide fault tolerance, protecting the solution from the failure of a single network device.

Network teaming provides a choice of policies to specify how traffic is distributed across the network ports. The default mode is round robin, but active backup and load balance provided better performance when using synchronous transport for ODG. For details on configuring network teaming, see the configuration file in Appendix C and the [RHEL Networking Guide](#).

Table 4 shows the network configuration for the solution with HPE SGLX and ODG.

**Table 4.** Network configuration

Function	Number of network ports	Teaming policy	Cluster.conf entry
Public network	1	N/A	STATIONARY_IP
SGLX heartbeat #1	2	Active backup	HEARTBEAT_IP
SGLX heartbeat #2	2	Active backup	HEARTBEAT_IP
ODG redo transport	2	Load balance or active backup <sup>3</sup>	STATIONARY_IP

## Serviceguard cluster configuration

With two networks configured for the Serviceguard heartbeat, the cluster setting for MEMBER\_TIMEOUT can be reduced from the default of 14 seconds down to three seconds. This provides a faster failover time and decreases downtime when the primary node becomes unavailable. A MEMBER\_TIMEOUT setting of three seconds was used for this RA.

A two-node Serviceguard cluster requires that a cluster lock be configured to prevent a split-brain scenario. The cluster lock can be configured via a lock LUN or a quorum server. For this RA, a lock LUN was used to provide faster failover. The lock LUN was located on the HPE 3PAR storage, and exported to both nodes in the cluster.

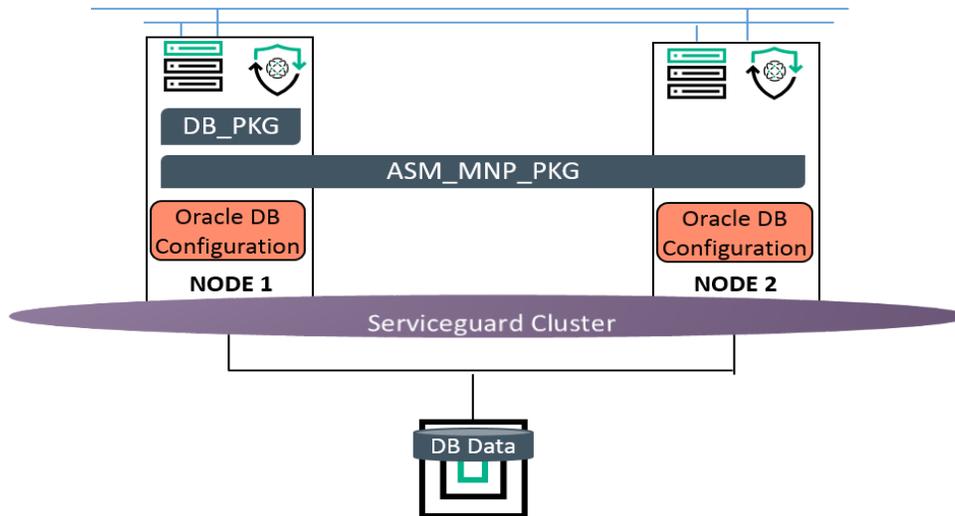
### HA solution for Oracle Database (without ODG)

- Install Oracle Database 18c Grid Infrastructure and database binaries local to each cluster node. When installing, ensure that the ASM SIDs are the same for both nodes.
- Create ASM disk groups that are accessible from all cluster nodes (shared storage) to be used as database stores.
- Create SGLX packages that manage the ASM processes and database instance(s).
  - Oracle ASM processes are running on both nodes and the multi-node package named “ASM\_MNP\_PKG” manages them on the two nodes.
  - The Oracle DB is configured to run on both the nodes and will be active only on one node at any point of time. This database is managed by the failover package named “DB\_PKG” and is active on the nodes where this package is running.
- In this configuration, the shared volume configuration used for the Oracle baseline configuration was:
  - Eight 500GB volumes were created and used for the data area disk group, and eight 500GB volumes were assigned for the fast recovery area disk group. The disk groups were configured as shared storage between the two nPars.
  - Two 5GB volumes were configured to hold the ASM configuration data, with one volume available to each of the two nPars.

<sup>3</sup> Testing was done with both load balance and active backup policies. Another option is LACP, which may be preferable for environments with non-HPE switches.



Figure 4 depicts the SGLX cluster and packages for the SGLX solution without ODG.



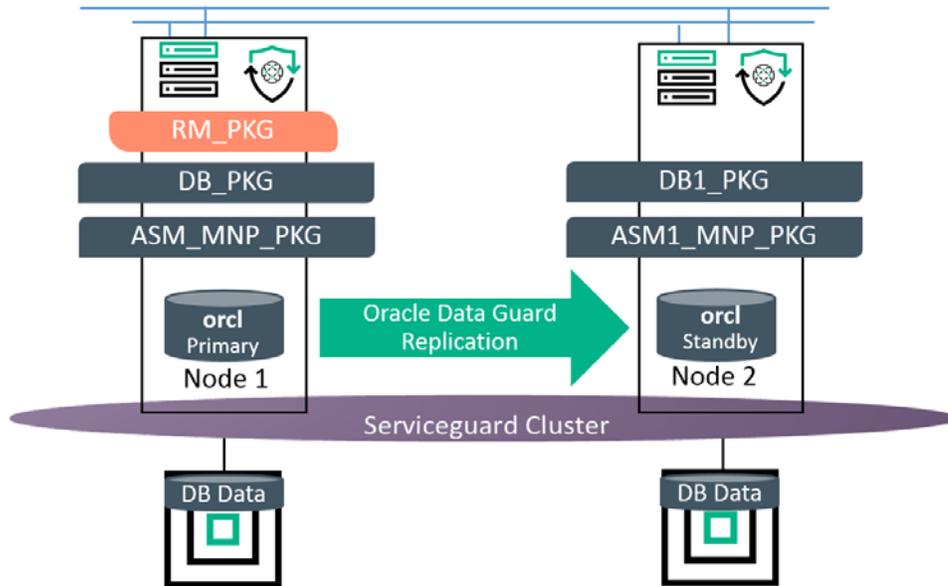
**Figure 4.** Schematic representation of cluster and packages for SGLX solution without ODG for Oracle DB

#### HA/DR solution configuration for Oracle Database using ODG

- Install Oracle Database 18c Grid Infrastructure (ASM) and database binaries local to each cluster node.
- Create ASM disk groups for application data and fast recovery area that are local to each node.
- Deploy the solution using “cmoradgworkload” utility:
  - The Oracle ASM processes are running on both nodes and are managed by the packages named “ASM\_MNP\_PKG” and “ASM1\_MNP\_PKG” on the respective nodes.
  - The Oracle DB is configured to run on both nodes with one node being the primary and the other the standby. This database is managed by the packages named “DB\_PKG” and “DB1\_PKG” on the respective nodes.
  - The DB roles are managed by the package named “RM\_PKG”. This package will always run on the node which is the primary.
  - Sixteen 500GB volumes were created and used for the two data area disk groups, and sixteen 500GB volumes were created and used for the two fast recovery areas. The ASM disk groups were not set up as shared storage. Instead, each nPar had its own ASM data and fast recovery area (8 volumes for each).



Figure 5 shows the SGLX cluster and packages for the solution with ODG.



**Figure 5.** Schematic representation of cluster and packages for SGLX solution with ODG for Oracle DB

## Capacity and sizing

### Workload description

Oracle performance tests were conducted using HammerDB, an open-source tool. For this RA, HammerDB 3.0 was used to implement an OLTP-type workload with a schema size of 5,000 warehouses and 500GB in size. HammerDB provides a real-world type scenario that consumes both CPU for the application logic and I/O. The tool implements an OLTP-type workload with small I/O sizes of a random nature. The transaction results were normalized and used to compare test configurations. Other metrics collected during the testing came from the operating system and/or standard Oracle Automatic Workload Repository (AWR) statistics reports.

### Analysis and recommendations

Two types of tests were conducted:

- Performance tests to compare the impact of adding ODG replication versus the configuration with HPE SGLX alone.
- Failover tests to measure the time to recover from a failure on the primary server and bring up the database on the standby server.

There is a tradeoff between performance and failover time, and customers need to determine the importance of each criteria for their environment. The following sections describe the results and recommendations for both types of tests.

---

### Note

In this RA, the configuration where SGLX is used to recover the database by automatically restarting the database on another node is referred to as the "SGLX only" configuration. This is to distinguish it from the configuration where both SGLX and ODG are used.

---

### Performance results

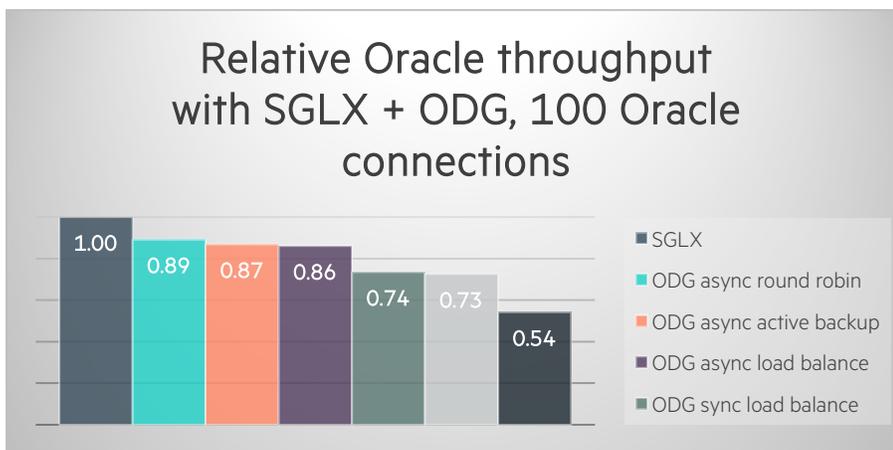
For the performance tests, multiple factors were considered, including the redo transport mode, and the network teaming mode for the redo transport network. ODG provides the option to transport the redo logs to the standby server in either synchronous or asynchronous mode. When using asynchronous mode, the protection mode is set to maximum performance. As expected this mode provided better performance than the synchronous mode with the maximum availability protection level. When configuring network teaming, options for the teaming mode include



active backup, load balance, and round robin. As shown below, the active balance and load balance modes provided better performance than round robin when using synchronous mode for the redo log transport.

Figure 6 shows the relative Oracle throughput with a workload of 100 Oracle connections (virtual users) for a baseline configuration with SGLX only, as compared to various configurations with SGLX plus ODG. When using asynchronous redo transport, Oracle throughput for the ODG configurations was 86% to 89% of the throughput for the baseline configuration. When using synchronous redo transport, Oracle throughput for the ODG configurations was 73% to 74% of the baseline configuration when the network team was configured in active backup or load balance mode for the redo transport network. The ODG configuration with the round robin network team mode only provided 54% of the performance of the baseline configuration. Based upon these results:

- Either active backup or load balance is recommended for the network team mode for the redo transport network.<sup>4</sup>
- Asynchronous redo transport can be used with ODG for better performance, and synchronous transport mode can be used when there is no tolerance for data loss due to a failure event.



**Figure 6.** Relative Oracle throughput for SGLX baseline versus various ODG configurations

### Failover results

For the failover event, the power was cycled for the primary nPar using the Remote Management Console command “power cycle npar”. All system times were synced using NTP, and a script running on the primary nPar recorded the time every second to a log file. The time to recovery for both configurations (with and without ODG) is the time elapsed between (1) the last time recorded in the script log file, which represents the beginning of the failover event, and (2) the time reported by the SGLX package when the database was open (ready to receive transactions) on the standby server. The recovery includes the time to detect the failure and reform the cluster, plus the time to start the database on the standby server.

The key finding is that with a 100 Oracle connection workload, the role transition of the Oracle Database to the standby node was completed in under 30 seconds when using HPE SGLX and ODG to automate the recovery when the primary node has a failure.

<sup>4</sup> Another network teaming option is LACP. LCAP may be preferable for environments with non-HPE switches.



Figure 7 shows the recovery time broken down into the cluster reformation time and the Oracle database startup time for a failover with no Oracle workload running on the primary server. Figure 8 shows similar information with 100 Oracle connections running at the time of the failure of the primary server. Note that the detection of the failure and cluster reformation happens within 4 to 5 seconds (when using a 3 second MEMBER\_TIMEOUT setting), and the bulk of the time is consumed in recovering and starting up the database. When there was no load on the database, it took 10 seconds for ODG to convert the standby database to a primary and open it, as compared to 62 seconds to start the database on the standby when performing a failover using HPE SGLX. With 100 Oracle connections at the time of the failure, it took 19 seconds for ODG to bring up the standby database as the primary, as compared to 287 seconds with HPE SGLX. There is substantially more time required to recover the database when ODG is not in use, demonstrating the value of combining SGLX and Data Guard for automating the role transition process.

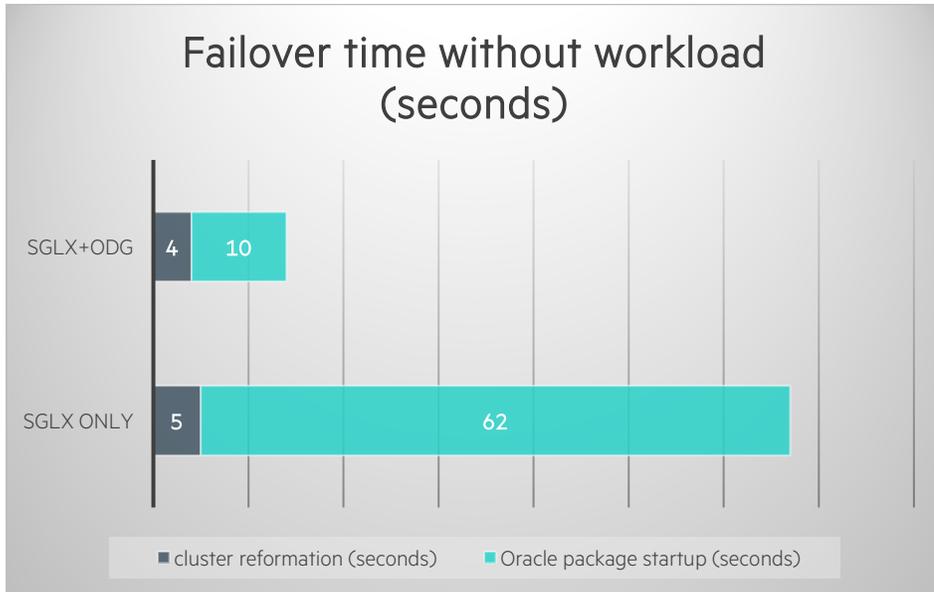
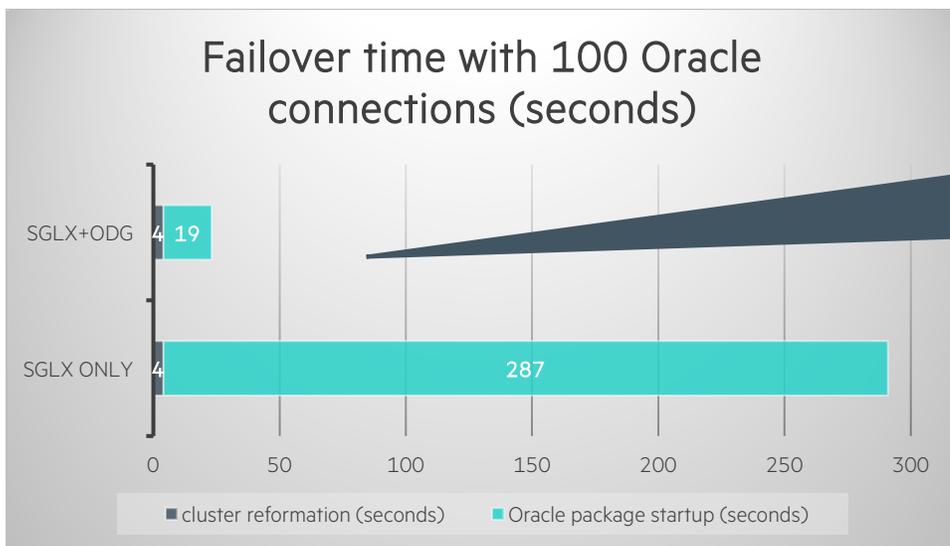


Figure 7. Failover time with no system load for configuration with Serviceguard versus configuration with Serviceguard plus ODG



Failover time under 30 seconds with HPE Serviceguard for Linux plus Oracle Data Guard automation

Figure 8. Failover time with 100 Oracle connections for configuration with Serviceguard versus configuration with Serviceguard plus ODG



In addition to ODG significantly lowering the role transition time, the automation of the role transition by the HPE Serviceguard deep integration with ODG packages should not be underestimated. If ODG had been deployed without the use of HPE SGLX, the time to recognize and initiate a failover would be much larger and would include the following manual steps:

1. In case of failure, ensure graceful shutdown of the primary when possible
2. Identify a problem
3. Look up the commands required to make the standby database the new primary database
4. Enter and execute the commands
5. Check the consistency of the new primary database
6. Manually bring up the rest of the resources like IPs, listeners, and the application stack

When the server that was originally running the primary database is brought back into service after a failure, the HPE Serviceguard solution for ODG automates bringing up the former primary database as the standby database. This saves time and frees up the database administrator from the need to follow multiple manual steps to complete the conversion. With 100 Oracle connections running on the primary node at the time of the initial role transition, it took 6 minutes and 1 second to convert the former primary database to a standby database. The bulk of the time (5 minutes and 24 seconds) was spent in flashing back the database to the point in time prior to the role transition (when the secondary became the primary).

Additionally, you have the assurance HPE SGLX will execute the same recovery commands in the same order, without error, every time. This is a level of consistency and rapidity that cannot be replicated by a human.

## Summary

Getting the most out of any server environment requires understanding the unique solution requirements so that an appropriate balance can be struck between performance and protection. A configuration that delivers the highest possible performance may not be the best option when SLAs require minimal downtime in the event of a failure. HPE Superdome Flex, HPE 3PAR Storage, and HPE SGLX combined with ODG is an ideal solution for providing HA for Oracle databases. Additionally, combining the scale-up Superdome Flex platform with a fast-failover solution minimizes the need for Oracle RAC, removing complexity and substantially reducing costs (Oracle RAC can increase Oracle licensing fees by nearly 50%). This RA demonstrates that:

- HPE SGLX combined with ODG automate the role transition process from a primary to a standby node, ensuring minimal downtime for mission critical environments.
- An Oracle Database under load can be failed over from a primary to a standby node in under 30 seconds.
- If minimal failover time with no data loss is required, implement synchronous redo transport for ODG.
- If performance is more important, use asynchronous redo transport for ODG to obtain Oracle database throughput close to performance without replication to a standby database.

## Implementing a proof-of-concept

As a matter of best practice for all deployments, Hewlett Packard Enterprise recommends implementing a proof-of-concept using a test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HPE Services representative ([hpe.com/us/en/services/consulting.html](http://hpe.com/us/en/services/consulting.html)) or your HPE partner.



## Appendix A: Bill of materials

### Note

Part numbers are at time of publication/testing and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult with your HPE Reseller or HPE Sales Representative for more details at [hpe.com/us/en/services/consulting.html](http://hpe.com/us/en/services/consulting.html).

**Table 5.** Bill of materials

Qty	Part number	Description
<b>Rack Infrastructure</b>		
1	P9K16A	HPE 42U 800x1200mm Advanced Shock Rack <sup>5</sup>
1	P9K16A-001	HPE Factory Express Base Racking Service
2	JG838A	HPE FlexFabric 5900CP 48XG 4QSFP+ Switch
<b>HPE Superdome Flex</b>		
1	Q2N05A	HPE Superdome Flex Base Chassis
4	Q2N43A <sup>6</sup>	HPE Superdome Flex 480GB SATA SSD
2	Q2N42A	HPE Superdome Flex DVD-R Drive
4	817718-B21	10/25Gb dual port 631SFP28 Ethernet adapter
1	Q6L89A	HPE Superdome Flex 4s Partition Chassis
8	Q6L90A	HPE Superdome Flex Intel Xeon-Platinum 8180 (2.5GHz/28-core/205W) Processor Kit
24	Q2N39A	HPE Superdome Flex DDR4 256 GB (4x64GB) Memory Kit
2	Q2N08A	HPE Superdome Flex PCIe FH 12-slot 3 Riser Kit
8	Q0L14A	HPE SN1200E 16Gb 2p FC HBA
1	Q9Z03A	HPE Superdome Flex 8-socket Partition Activation Kit
1	Q2N07A	HPE Superdome Flex Mgmt Controller
4	P9Q61A	HPE G2 Basic 3Ph 17.3kVA/C13 NA/JP PDU
1	H6J85A	HPE Rack Hardware Kit
2	Q7N11A	HPE Foundation SW 2 RHEL Lic RTU
<b>HPE 3PAR 9450 Storage</b>		
1	P9K10A	HPE 42U 600x1200mm Adv G2 Kit Shock Rack
1	Q0E92A	HPE 3PAR 9450 2N+SW Storage Base
2	Q7F41A	HPE 3PAR 9450+SW Storage Node
4	Q0E96A	HPE 3PAR 9000 4pt 12Gb SAS HBA
8	Q0E97A	HPE 3PAR 9000 4pt 16Gb FC HBA
16	Q0E95A	HPE 3PAR 9000 2U SFF Drive Enclosure
160	Q0F40A	HPE 3PAR 9000 400GB+SW SFF SSD
32	716197-B21	HPE Ext 2.0m MiniSAS HD to MiniSAS HD Cbl

<sup>5</sup> This rack is no longer available. See the QuickSpecs for the HPE Superdome Flex for current racks.

<sup>6</sup> This disk drive is superseded by part number R2A72A.



1	Q0F86A	HPE 3PAR StoreServ RPS Service Processor
1	Q1H95A	HPE 3PAR 1U Rack Accessories Kit
4	P9Q41A	HPE G2 Basic 4.9kVA/(20) C13 NA/JP PDU
1	BW932A	HPE 600mm Rack Stabilizer Kit
1	L7F19A	HPE 3PAR All-in Sngl-sys SW Latest Media
2	QK753C	HPE SN6000B 16Gb 48/24 FC Switch
48	QK724A	HPE B-series 16Gb SFP+ Short Wave Transceiver

There are multiple options for SGLX licenses, and table 6 lists these options, including ones for an SGLX solution with ODG.

**Table 6.** Bill of materials (sample for a solution consisting of individual HPE BU Products. NOTE: remove this comment before publication)

Part number	Description
BB095AC	HPE SGLX Advanced x86 1y24x7 PSL Fix LTU
BB095ACE	HPE SGLX Advanced x86 1y24x7 PSL E-LTU
R1T31AAE	HPE SGLX x86 Adv 1yr Sub 24x7 PSL E-LTU
R1T34AAE	HPE SGLX x86 Adv 3yr Sub 24x7 PSL E-LTU
<b>SGLX with ODG solution</b>	
BB097AC	HPE SGLX Ent x86 1y24x7 PSL Fix LTU
BB097ACE	HPE SGLX Enterprise x86 1y24x7 PSL E-LTU
R1T32AAE	HPE SGLX x86 Ent 1yr Sub 24x7 PSL E-LTU
R1T35AAE	HPE SGLX x86 Ent 3yr Sub 24x7 PSL E-LTU

## Appendix B: RHEL kernel tuning parameters

The RHEL tuning parameters that were set in /etc/sysctl.conf are listed below. The shmmax parameter and nr\_hugepages were set large enough to accommodate the memory requirements of the Oracle SGA.

```
kernel.sem = 250 32000 100 128
kernel.shmall = 4294967295
kernel.shmmax = 3298534883328
fs.file-max = 6815744
kernel.shmmni = 4096
fs.aio-max-nr = 1048576
net.ipv4.ip_local_port_range = 9000 65500
net.core.rmem_default = 1048576
net.core.wmem_default = 1048576
net.core.rmem_max=26214400
net.core.wmem_max=26214400
net.ipv4.tcp_rmem = 1048576 1048576 4194304
net.ipv4.tcp_wmem = 1048576 1048576 1048576
vm.nr_hugepages = 237850
vm.hugetlb_shm_group = 507
kernel.numa_balancing = 0
```



## Appendix C: Network teaming configuration

The network configuration file for the network team for the ODG redo transport network is shown below. Three different modes were tested: round robin, active backup, and load balance. A similar configuration file was used for the SGLX heartbeat network, which was configured with the active backup mode.

```
# cat ifcfg-team1
DEVICE=team1
DEVICETYPE=Team
ONBOOT=yes
BOOTPROTO=none
IPADDR=192.168.3.107
PREFIX=24
#TEAM_CONFIG='{"runner": {"name": "roundrobin"}}'
#TEAM_CONFIG='{"runner": {"name": "activebackup"}, "link_watch": {"name": "ethtool"}}'
TEAM_CONFIG='{"runner": {"name": "loadbalance", "tx_hash": ["eth", "ipv4", "l4"], "tx_balancer": {"name": "basic"}}}'
MTU=9000
```

## Appendix D: Oracle user account limits

The following settings were included in the file /etc/security/limits.d/oracle-limits.conf:

```
oracle soft nofile 1024
oracle hard nofile 65536
oracle soft nproc 16384
oracle hard nproc 16384
oracle soft stack 10240
oracle hard stack 32768
oracle hard memlock 2765214764
oracle soft memlock 2765214764
```

## Appendix E: Oracle initialization parameters

The following Oracle parameters were set in the init.ora initialization file on the primary server. The parameters listed below the comment "Params for ODG" were used only for the ODG tests. The parameters listed below the comment "ODG params for standby role" came into play when the database was brought up in standby mode after a failover of the primary database to the other server. The second option to the parameter `log_archive_dest_2` was changed from `SYNC` to `ASYNC` when testing ODG with the asynchronous redo transport mode. The listener named ODG\_O was added for the ODG configuration, and was used for the redo log transport.

```
*.audit_file_dest='/u01/app/oracle/admin/orcl/adump'
*.audit_trail='db'
*.compatible='18.0.0'
*.control_files='+DATA/ORCL/CONTROLFILE/current.260.1008435665','+FRA/ORCL/CONTROLFILE/current.256.1008435665'
*.db_block_size=8192
*.db_create_file_dest='+DATA'
*.db_name='orcl'
*.db_recovery_file_dest='+FRA'
*.diagnostic_dest='/u01/app/oracle'
*.dispatchers='{(PROTOCOL=TCP) (SERVICE=orclXDB)}'
*.local_listener='LISTENER_ORCL','ODG_O'
*.open_cursors=3000
*.pga_aggregate_target=158380m
*.processes=3000
```



```
*.remote_login_passwordfile='EXCLUSIVE'
*.sga_target=475648m
*.undo_tablespace='UNDOTBS1'
_high_priority_processes='VKTM*|LG*'
lock_sga=TRUE
use_large_pages='ONLY'
_fast_cursor_reexecute=true
_enable_NUMA_support=true
_ash_size=512M
# Params for ODG
db_recovery_file_dest_size = 3T
db_unique_name='orcl'
log_archive_config='DG_CONFIG={orcl,orcl_stby}'
log_archive_dest_1='location=USE_DB_RECOVERY_FILE_DEST valid_for=(ALL_LOGFILES,ALL_ROLES)
db_unique_name=orcl'
log_archive_dest_2='service=orcl_stby SYNC NOAFFIRM valid_for=(ALL_LOGFILES,PRIMARY_ROLE)
db_unique_name=orcl_stby'
log_archive_format='%t_%s_%r.arc
# ODG params for standby role
fal_server='orcl_stby'
db_file_name_convert='+data/orcl_stby/','+data/orcl/','+fra/orcl_stby/','+fra/orcl/'
log_file_name_convert='+data/orcl_stby/','+data/orcl/','+fra/orcl_stby/','+fra/orcl/'
standby_file_management='auto'
```

For the Serviceguard-only configuration, the init.ora file was identical on both servers. For the ODG configuration, the following settings were specific to the standby server, with all other settings remaining the same on both servers:

```
*.control_files='+DATA/ORCL_STBY/CONTROLFILE/control01.ct1','+FRA/ORCL_STBY/CONTROLFILE/control02.ct1'
*.db_file_name_convert='+data/orcl/','+data/orcl_stby/','+fra/orcl/','+fra/orcl_stby/'
*.db_unique_name='orcl_stby'
*.fal_server='orcl'
*.log_archive_dest_1='location=USE_DB_RECOVERY_FILE_DEST valid_for=(ALL_LOGFILES,ALL_ROLES)
db_unique_name=orcl_stby'
*.log_archive_dest_2='service=orcl SYNC NOAFFIRM valid_for=(ALL_LOGFILES,PRIMARY_ROLE)
db_unique_name=orcl'
*.log_file_name_convert='+data/orcl/','+data/orcl_stby/','+fra/orcl/','+fra/orcl_stby/'
```

## Appendix F: Oracle listener configuration

The Oracle listener and TNS names files were configured as follows.

The listener.ora file for the SGLX-only configuration:

```
ORCL =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP)(HOST = sdflexp1)(PORT = 1521))
      (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC1521))
    )
  )
```

The tnsnames.ora file for the SGLX-only configuration:

```
LISTENER_ORCL =
```



```

[ADDRESS = {PROTOCOL = TCP}{HOST = sdflexp1}{PORT = 1521}]
ORCL =
  [DESCRIPTION =
    [ADDRESS = {PROTOCOL = TCP}{HOST = sdflexp1}{PORT = 1521}]
    [CONNECT_DATA =
      [SERVER = DEDICATED]
      [SERVICE_NAME = orcl]
    ]
  ]
]

```

The listener.ora file for the ODG primary server (the file for the standby server was identical, other than the HOST parameter for both listeners):

```

ORCL =
  [DESCRIPTION_LIST =
    [DESCRIPTION =
      [ADDRESS = {PROTOCOL = TCP}{HOST = sdflexp1}{PORT = 1521}]
      [ADDRESS = {PROTOCOL = IPC}{KEY = EXTPROC1521}]
    ]
  ]
# added for ODG
ODG_0 =
  [DESCRIPTION_LIST =
    [DESCRIPTION =
      [ADDRESS = {PROTOCOL = TCP}{HOST = sdflexp1-odg}{PORT = 1600}]
      [SEND_BUF_SIZE=104857600]
      [RECV_BUF_SIZE=104857600]
    ]
    [UR=A]
  ]
]
# this is for static registration for ODG
SID_LIST_ODG_0 =
  [SID_LIST =
    [SID_DESC =
      [SDU=65535]
      [GLOBAL_DBNAME = orcl]
      [ORACLE_HOME = /u01/app/oracle/product/18.0.0/dbhome_1]
      [SID_NAME = orcl]
      [UR=A]
    ]
  ]
]
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER=ON # line added by Agent
VALID_NODE_CHECKING_REGISTRATION_LISTENER=ON # line added by Agent

```

The tnsnames.ora file for the primary server:

```

LISTENER_ORCL =
  [ADDRESS = {PROTOCOL = TCP}{HOST = sdflexp1}{PORT = 1521}]

ORCL,ODG_0 =
  [DESCRIPTION =
    [ADDRESS = {PROTOCOL = TCP}{HOST = sdflexp1-odg}{PORT = 1600}]
  ]

```



```

[CONNECT_DATA =
  [SERVER = DEDICATED]
  [SERVICE_NAME = orcl]
  [SID = orcl]
]
[SDU=65535]
[SEND_BUF_SIZE=10485760]
[RECV_BUF_SIZE=10485760]
]

ORCL_STBY =
[DESCRIPTION =
  [ADDRESS = [PROTOCOL = TCP][HOST = sdflexp2-odg][PORT = 1600]]
  [CONNECT_DATA =
    [SERVER = DEDICATED]
    [SERVICE_NAME = orcl_stby]
    [SID = orcl]
  ]
  [SDU=65535]
  [SEND_BUF_SIZE=10485760]
  [RECV_BUF_SIZE=10485760]
]

```

The tnsnames.ora file for the standby server:

```

LISTENER_ORCL =
  [ADDRESS = [PROTOCOL = TCP][HOST = sdflexp2][PORT = 1521]]

ORCL =
  [DESCRIPTION =
    [ADDRESS = [PROTOCOL = TCP][HOST = sdflexp1-odg][PORT = 1600]]
    [CONNECT_DATA =
      [SERVER = DEDICATED]
      [SERVICE_NAME = orcl]
      [SID = orcl]
    ]
    [SDU=65535]
    [SEND_BUF_SIZE=10485760]
    [RECV_BUF_SIZE=10485760]
  ]

ORCL_STBY,ODG_0 =
  [DESCRIPTION =
    [ADDRESS = [PROTOCOL = TCP][HOST = sdflexp2-odg][PORT = 1600]]
    [CONNECT_DATA =
      [SERVER = DEDICATED]
      [SERVICE_NAME = orcl_stby]
      [SID = orcl]
    ]
    [SDU=65535]
    [SEND_BUF_SIZE=10485760]
    [RECV_BUF_SIZE=10485760]
  ]

```



### Resources and additional links

HPE Mission Critical Servers Reference Architectures, [hpe.com/info/missioncritical-ra](http://hpe.com/info/missioncritical-ra)

HPE Reference Architectures, [hpe.com/info/ra](http://hpe.com/info/ra)

Oracle Solutions, [hpe.com/info/oracle](http://hpe.com/info/oracle)

HPE Superdome High-End Servers, [hpe.com/superdome](http://hpe.com/superdome)

HPE Superdome Flex Server architecture and RAS, <https://h20195.www2.hpe.com/V2/getpdf.aspx/A00036491ENW.pdf>

HPE Superdome Flex Server Partitioning Guide, [https://support.hpe.com/hpsc/doc/public/display?docId=a00050296en\\_us](https://support.hpe.com/hpsc/doc/public/display?docId=a00050296en_us)

HPE Superdome Flex Server OS Installation Guide, [https://support.hpe.com/hpsc/doc/public/display?docId=a00038168en\\_us](https://support.hpe.com/hpsc/doc/public/display?docId=a00038168en_us)

Running Linux on HPE Superdome Flex Server, [https://support.hpe.com/hpsc/doc/public/display?docId=a00058577en\\_us](https://support.hpe.com/hpsc/doc/public/display?docId=a00058577en_us)

Managing HPE Serviceguard for Linux A.12.40.00

[https://support.hpe.com/hpsc/doc/public/display?docLocale=en\\_US&docId=emr\\_na-a00069643en\\_us&withFrame](https://support.hpe.com/hpsc/doc/public/display?docLocale=en_US&docId=emr_na-a00069643en_us&withFrame)

HPE Serviceguard Toolkit for Oracle on Linux User Guide

[https://support.hpe.com/hpsc/doc/public/display?docLocale=en\\_US&docId=emr\\_na-a00052274en\\_us&withFrame](https://support.hpe.com/hpsc/doc/public/display?docLocale=en_US&docId=emr_na-a00052274en_us&withFrame)

HPE Serviceguard Toolkit for Oracle Data Guard User Guide

[https://support.hpe.com/hpsc/doc/public/display?docLocale=en\\_US&docId=emr\\_na-a00052284en\\_us&withFrame](https://support.hpe.com/hpsc/doc/public/display?docLocale=en_US&docId=emr_na-a00052284en_us&withFrame)

Oracle Data Guard Concepts and Administration, <https://docs.oracle.com/en/database/oracle/oracle-database/18/sbydb/data-guard-concepts-and-administration.pdf>

HPE Application Tuner Express, <https://myenterpriselicense.hpe.com/cwp-ui/evaluation/HPE-ATX>

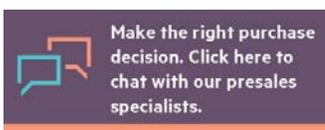
HPE 3PAR StoreServ Storage, [hpe.com/3par](http://hpe.com/3par)

HPE 3PAR Red Hat Enterprise Linux Implementation Guide, <https://support.hpe.com/hpsc/doc/public/display?docId=c04448818>

Best Practices for Oracle Database on HPE 3PAR StoreServ Storage, [https://support.hpe.com/hpsc/doc/public/display?docId=emr\\_na-a00038978en\\_us&docLocale=en\\_US](https://support.hpe.com/hpsc/doc/public/display?docId=emr_na-a00038978en_us&docLocale=en_US)

Red Hat Enterprise Linux 7 Networking Guide, [https://access.redhat.com/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/7/pdf/Networking\\_Guide/Red\\_Hat\\_Enterprise\\_Linux-7-Networking\\_Guide-en-US.pdf](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/pdf/Networking_Guide/Red_Hat_Enterprise_Linux-7-Networking_Guide-en-US.pdf)

To help us improve our documents, please provide feedback at [hpe.com/contact/feedback](http://hpe.com/contact/feedback)



Share now 

Sign up for updates



---

© Copyright 2019 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Red Hat and Red Hat Enterprise Linux are trademarks of Red Hat, Inc. in the United States and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Oracle is a registered trademark of Oracle and/or its affiliates. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware® is a registered trademark of VMware, Inc. in the United States and/or other jurisdictions.