

# Cisco HyperFlex All-Flash Systems for Oracle Real Application Clusters Reference Architecture



# Contents

## Executive summary

Cisco HyperFlex HX Data Platform all-flash storage

Why use Cisco HyperFlex all-flash systems for Oracle RAC deployments

Oracle RAC Database on Cisco HyperFlex systems

Oracle RAC with VMware virtualized environment

Audience

Oracle Database scalable architecture solution overview

## Hardware components

- Cisco HyperFlex system

- Cisco HyperFlex HX220c M4 Node and HX220c M4 All Flash Node

- Cisco UCS 6200 Series Fabric Interconnects

## Software components

- VMware vSphere

- Oracle Database 12c

## Storage architecture

- Storage configuration

Oracle Database configuration

Network configuration

Storage configuration

- Paravirtual SCSI queue depth settings

## Engineering validation: Performance testing

- Oracle RAC node scale test

- Oracle user scale test

## Reliability and disaster recovery

- Disaster recovery

- Oracle Data Guard

- Cisco HyperFlex HX Data Platform disaster recovery

- Oracle RAC Database backup and recovery: Oracle Recovery Manager

Conclusion

For more information

## Executive summary

Oracle Database is the choice for many enterprise database applications. Its power and scalability make it attractive for implementing business-critical applications. However, making those applications highly available can be extremely complicated and expensive.

Oracle Real Application Clusters (RAC) is the solution of choice for customers needing to provide high availability and scalability to the Oracle Database. Originally focused on providing best-in-class database services, Oracle RAC has evolved over the years and now provides a comprehensive high-availability stack that also provides scalability, flexibility, and agility for applications.

With the Cisco HyperFlex™ solution for Oracle RAC databases, organizations can implement RAC databases using a highly integrated solution that scales as business demand increases. RAC uses a shared-disk architecture that requires all instances of RAC to have access to the same storage elements. Cisco HyperFlex systems use the [Multi-writer option](#) to enable virtual disks to be shared between different RAC virtual machines.

This reference architecture provides a configuration that is fully validated to help ensure that the entire hardware and software stack is suitable for a high-performance clustered workload. This configuration follows the industry best practices for Oracle Database in a VMware virtualized environment. Additional details about deploying Oracle RAC on VMware can be found [here](#).

## Cisco HyperFlex HX Data Platform all-flash storage

Cisco HyperFlex systems are designed with an end-to-end software-defined infrastructure that eliminates the compromises found in first-generation products. With all-flash memory storage configurations and a choice of management tools, Cisco HyperFlex systems deliver a tightly integrated cluster that is up and running in less than an hour and that scales resources independently to closely match your Oracle Database requirements. For an in-depth look at the Cisco HyperFlex architecture, see the Cisco white paper [Deliver Hyperconvergence with a Next-Generation Platform](#).

## Why use Cisco HyperFlex all-flash systems for Oracle RAC deployments

Oracle Database acts as the back end for many critical and performance-intensive applications. Organizations must be sure that it delivers consistent performance with predictable latency throughout the system. Cisco HyperFlex all-flash hyperconverged systems offer the following advantages:

- Low latency with consistent performance: Cisco HyperFlex all-flash systems, when used to host virtual database instances, deliver low latency and consistent database performance.
- Data protection (fast clones, snapshots, and replication factor): Cisco HyperFlex systems are engineered with robust data protection techniques that enable quick backup and recovery of applications to protect against failures.
- Storage optimization (always-active inline deduplication and compression): All data that comes into Cisco HyperFlex systems is by default optimized using inline deduplication and data compression techniques.
- Dynamic online scaling of performance and capacity: The flexible and independent scalability of the capacity and computing tiers of Cisco HyperFlex systems allow you to adapt to growing performance demands without any application disruption.
- No performance hotspots: The distributed architecture of the Cisco HyperFlex HX Data Platform helps ensure that every virtual machine can achieve the maximum number of storage I/O operations per second (IOPS) and use the capacity of the entire cluster, regardless of the physical node on which it resides. This feature is especially important for Oracle Database virtual machines because they frequently need higher performance to handle bursts of application and user activity.
- Nondisruptive system maintenance: Cisco HyperFlex systems support a distributed computing and storage environment that helps enable you to perform system maintenance tasks without disruption.

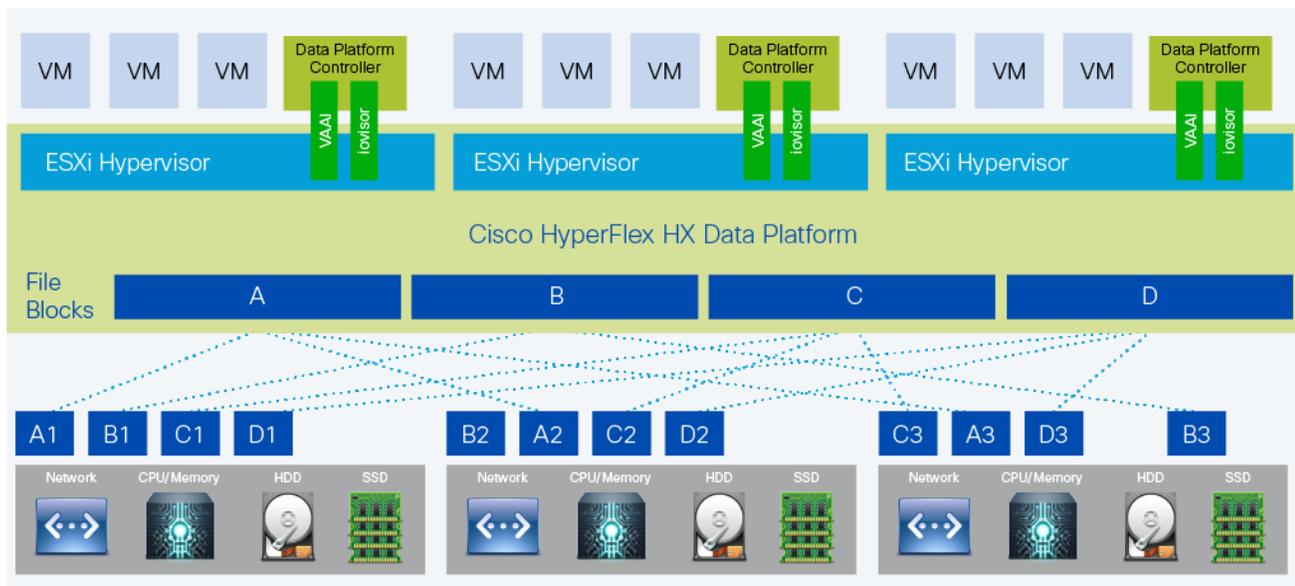
Several of these features and attributes are particularly applicable to Oracle RAC implementations, including consistent low-latency performance, storage optimization using always-on inline compression, dynamic and seamless performance and capacity scaling, and nondisruptive system maintenance.

### Oracle RAC Database on Cisco HyperFlex systems

This reference architecture guide describes how Cisco HyperFlex systems can provide intelligent end-to-end automation with network-integrated hyperconvergence for an Oracle RAC Database deployment. Cisco HyperFlex systems provide a high-performance, easy-to-use, integrated solution for an Oracle Database environment.

The Cisco HyperFlex data distribution architecture allows concurrent access to data by reading and writing to all nodes at the same time. This approach provides data reliability and fast database performance. Figure 1 shows the data distribution architecture.

**Figure 1.** Data distribution architecture



This hyperconverged solution integrates servers, storage systems, network resources, and storage software to provide an enterprise-scale environment for Oracle Database deployments. This highly integrated environment provides reliability, high availability, scalability, and performance for Oracle virtual machines to handle large-scale transactional workloads. The solution uses four virtual machines to create a single 4-node Oracle RAC database for performance, scalability, and reliability. The RAC node uses the Oracle Enterprise Linux operating system for the best interoperability with Oracle databases.

This reference architecture uses a cluster of four Cisco HyperFlex HX220c M4 All Flash Nodes to provide fast data access. Use this document to design an Oracle RAC database solution that meets your organization’s requirements and budget.

Cisco HyperFlex systems also support other enterprise Linux platforms such as SUSE and Red Hat Enterprise Linux (RHEL). For a complete list of virtual machine guest operating systems supported for VMware virtualized environments, see the [VMware Compatibility Guide](#).

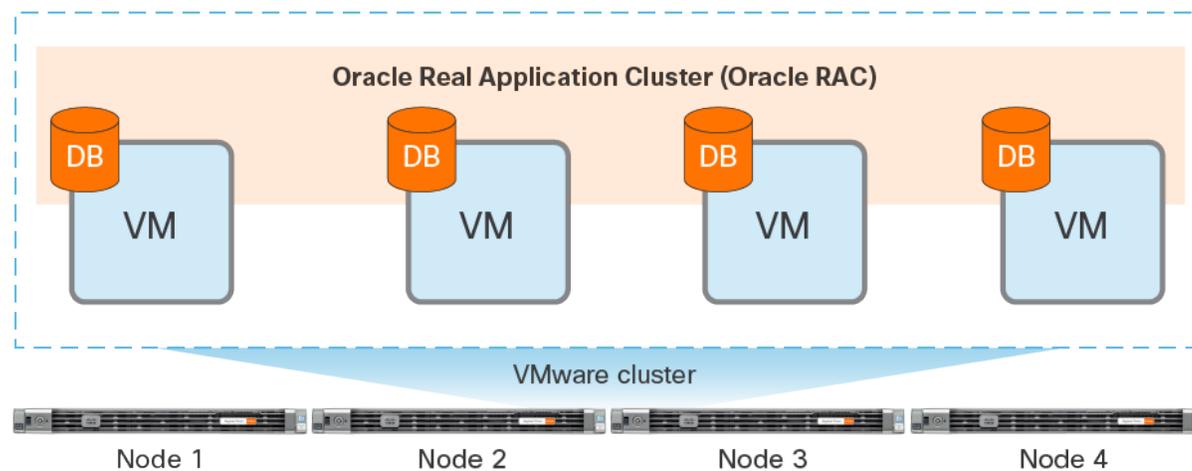
## Oracle RAC with VMware virtualized environment

This reference architecture uses VMware virtual machines to create Oracle RAC with four nodes. Although this solution guide describes a 4-node configuration, this architecture can support other all-flash Cisco HyperFlex configurations and RAC node and virtual machine counts and sizes as needed to meet your deployment requirements.

**Note:** For best availability, Oracle RAC virtual machines should be hosted on different VMware ESX servers. With this setup, the failure of any single ESX server will not take down more than a single RAC virtual machine and node with it.

Figure 2 shows the Oracle RAC configuration used in the solution described in this document.

**Figure 2.** Oracle Real Application Cluster configuration



Oracle RAC allows multiple virtual machines to access a single database to provide database redundancy while providing more processing resources for application access. The distributed architecture of the Cisco HyperFlex system allows a single RAC node to consume and properly use resources across the Cisco HyperFlex cluster.

The Cisco HyperFlex shared infrastructure enables the Oracle RAC environment to evenly distribute the workload among all RAC nodes running concurrently. These characteristics are critical for any multitenant database environment in which resource allocation may fluctuate.

The Cisco HyperFlex all-flash cluster supports large cluster sizes, with the capability to add computing-only nodes to independently scale the computing capacity of the cluster. This approach allows any deployment to start with a small environment and grow as needed, using a pay-as-you-grow model.

## Audience

This reference architecture document is written for the following audience:

- Database administrators
- Storage administrators
- IT professionals responsible for planning and deploying an Oracle Database solution

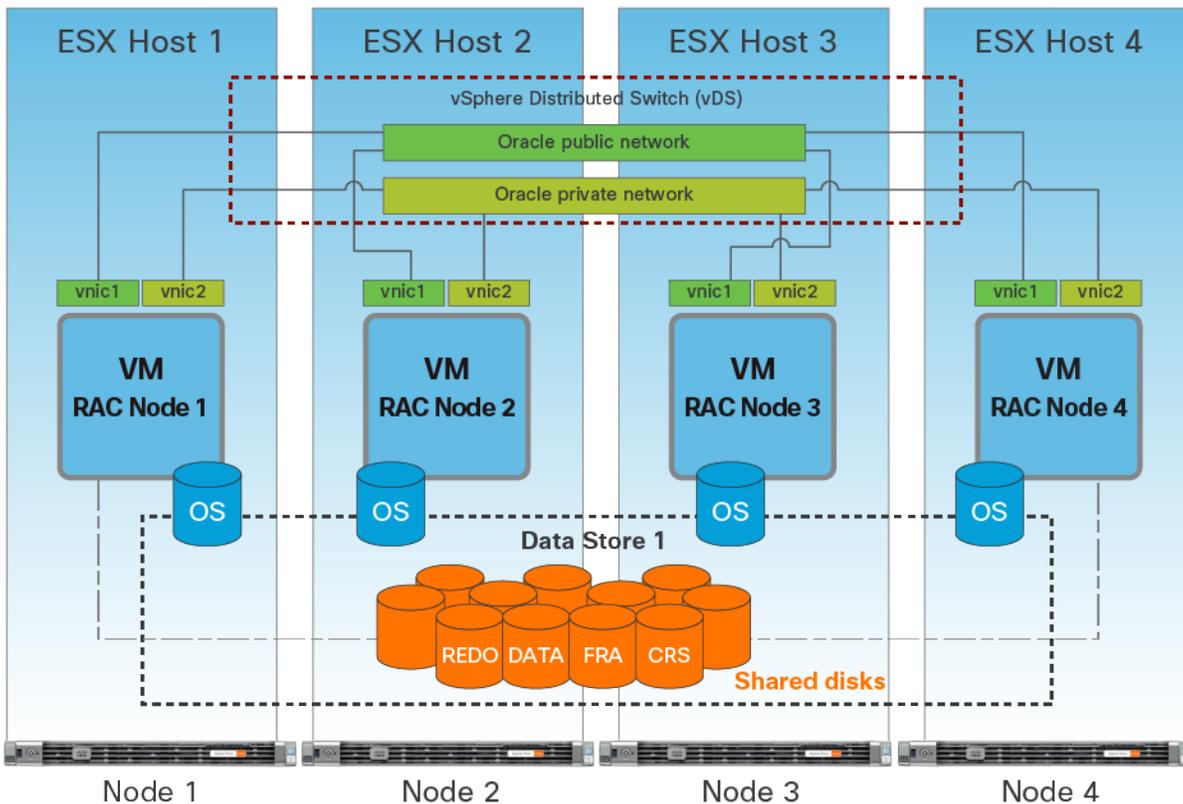
To benefit from this reference architecture guide, familiarity with the following is required:

- Hyperconvergence technology
- Virtualized environments
- Solid-state disk (SSD) and flash storage
- Oracle Database 12c
- Oracle Automatic Storage Management (ASM)
- Oracle Enterprise Linux

### Oracle Database scalable architecture solution overview

This section describes how to implement Oracle RAC Database on a Cisco HyperFlex system using a 4-node cluster. This reference configuration helps ensure proper sizing and configuration when you deploy a RAC database on a Cisco HyperFlex system. This solution enables customers to rapidly deploy Oracle databases by eliminating engineering and validation processes that are usually associated with deployment of enterprise solutions. Figure 3 provides a high-level view of the environment.

**Figure 3.** High-level solution design



This solution uses virtual machines for Oracle RAC nodes. Table 1 summarizes the configuration of the virtual machines with VMware.

**Table 1.** Oracle virtual machine configuration

Resource	Details for Oracle VM
<b>Virtual machine specifications</b>	24 virtual CPUs (vCPUs) 128 GB of vRAM <b>Note:</b> Several different virtual machine configurations were tested. Additional details are provided in the sections that follow.
<b>Virtual machine controllers</b>	4 Paravirtual SCSI (PVSCSI) controllers
<b>Virtual machine disks</b>	1 × 500-GB VMDK for virtual machine OS 4 × 500-GB VMDKs for Oracle Data 3 × 70-GB VMDKs for Oracle Redo log 2 × 80-GB VMDKs for Oracle Fast Recovery Area (FRA) 3 × 40-GB VMDKs for Oracle Cluster Ready Services (CRS) and voting disk Commvault Data Platform

Table 2 summarizes the hardware components of the solution.

**Table 2.** Hardware components

Description	Specification	Notes
<b>CPU</b>	2 Intel Xeon processor E5-2690 v4 CPUs at 2.60 GHz	
<b>Memory</b>	24 × 16-GB DIMMs	
<b>Cisco FlexFlash SD card</b>	2 × 64-GB SD cards	Boot drives
<b>SSD</b>	120-GB SSD	Configured for housekeeping tasks
	800-GB SSD	Configured as write log
	6 × 3.8-TB SSD (960-GB SSD option also available)	Capacity disks for each node
<b>Hypervisor</b>	VMware vSphere 6.0 U1B	Virtual platform for HX Data Platform software
<b>Cisco HyperFlex HX Data Platform software</b>	Cisco HyperFlex Data Platform Release 2.1(1b)	

Table 3 summarizes the Cisco HyperFlex HX220c M4 node configuration for the cluster.

**Table 3.** Cisco HX220c M4 node configuration

Hardware	Description	Quantity
<b>Cisco HyperFlex HX220c M4 and M4 All Flash Node servers</b>	Cisco 1RU hyperconverged nodes that allow cluster scaling in a small footprint	4
<b>Cisco UCS 6200 Series Fabric Interconnects</b>	Fabric interconnects	2

Table 4 summarizes the software components used for this solution.

**Table 4.** Software components

Software	Version	Function
Cisco HyperFlex HX Data Platform	Release 2.5(1a)	Data platform
Oracle Enterprise Linux	Version 7	OS for Oracle RAC
Oracle Grid and Automatic Storage Management (ASM)	Version 12c Release 2	Automatic storage management
Oracle Database	Version 12c Release 2	Oracle Database system
Oracle Swingbench, Order Entry workload	Version 2.5	Workload suite
Oracle Recovery Manager (RMAN)		Backup and recovery manager for Oracle Database

## Hardware components

This section describes the hardware components used for this solution.

### Cisco HyperFlex system

The Cisco HyperFlex system provides next-generation hyperconvergence with intelligent end-to-end automation and network integration by unifying computing, storage, and networking resources. The Cisco HyperFlex HX Data Platform is a high-performance, flash-optimized distributed file system that delivers a wide range of enterprise-class data management and optimization services. HX Data Platform is optimized for flash memory, reducing SSD wear while delivering high performance and low latency without compromising data management or storage efficiency.

The main features of the Cisco HyperFlex system include:

- Simplified data management
- Continuous data optimization
- Optimization for flash memory
- Independent scaling
- Dynamic data distribution

Visit [Cisco's website](#) for more details about the Cisco HyperFlex HX-Series.

### Cisco HyperFlex HX220c M4 Node and HX220c M4 All Flash Node

Nodes with flash storage are integrated into a single system by a pair of Cisco UCS® 6200 or 6300 Series Fabric Interconnects. Each node includes two Cisco® Flexible Flash (FlexFlash) Secure Digital (SD) cards, a single 120-GB SSD data-logging drive, a single SSD write-log drive, and up to six 1.2-TB SAS hard-disk drives (HDDs) or up to six 3.8-terabyte (TB) or six 960-GB SATA SSD drives, for a contribution of up to 22.8 TB of raw storage capacity to the cluster. The nodes use Intel® Xeon® processor E5-2600 v4 family CPUs and next-generation DDR4 memory and offer 12-Gbps SAS throughput. They deliver significant performance and efficiency gains and outstanding levels of adaptability in a 1-rack-unit (1RU) form factor.

This solution uses four Cisco HyperFlex HX220c M4 Nodes and HX220c M4 All Flash Nodes for a 4-node server cluster to provide 1-node failure reliability.

See the [Cisco HX220c M4 All Flash Nodes data sheet](#) for more information.

### Cisco UCS 6200 Series Fabric Interconnects

The Cisco UCS 6200 Series Fabric Interconnects are a core part of the Cisco Unified Computing System™ (Cisco UCS), providing both network connectivity and management capabilities for the system. The 6200 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions.

The Cisco UCS 6200 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the 6200 Series Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the 6200 Series provides both LAN and SAN connectivity for all blades within the domain.

The Cisco UCS 6200 Series uses a cut-through networking architecture, supporting deterministic, low-latency, line rate 10 Gigabit Ethernet on all ports, switching capacity of 2 terabits (Tb), and bandwidth of 320 Gbps per chassis, independent of packet size and enabled services. The product family supports Cisco low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the interconnect. Significant savings in total cost of ownership (TCO) come from an FCoE optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

**Note:** Although the testing described here was performed using Cisco UCS 6200 Series Fabric Interconnects, the Cisco HyperFlex HX Data Platform does include support for Cisco UCS 6300 Series Fabric Interconnects, which provide higher performance with 40 Gigabit Ethernet.

### Software components

This section describes the software components used for this solution.

#### VMware vSphere

VMware vSphere helps you get performance, availability, and efficiency from your infrastructure while reducing the hardware footprint and your capital expenditures (CapEx) through server consolidation. Using VMware products and features such as VMware ESX, vCenter Server, High Availability (HA), Distributed Resource Scheduler (DRS), and Fault Tolerance (FT), vSphere provides a robust environment with centralized management and gives administrators control over critical capabilities.

VMware provides product features that can help manage the entire infrastructure:

- **VMware vMotion and Storage vMotion:** vMotion allows nondisruptive migration of both virtual machines and storage. Its performance graphs allow you to monitor resources, virtual machines, resource pools, and server utilization.
- **VMware Distributed Resource Scheduler:** DRS monitors resource utilization and intelligently allocates system resources as needed.
- **VMware High Availability:** HA monitors hardware and OS failures and automatically restarts the virtual machine, providing cost-effective failover.
- **VMware Fault Tolerance:** FT provides continuous availability for applications by creating a live shadow instance of the virtual machine that stays synchronized with the primary instance. If a hardware failure occurs, the shadow instance instantly takes over and eliminates even the smallest data loss.

For more information, visit the [VMware website](#).

## Oracle Database 12c

Oracle Database 12.2 provides massive cloud scalability and real-time analytics, offering you greater agility, faster time to business insights, and real cost advantages. Designed for the cloud, Oracle Database enables you to lower IT costs and become more agile in provisioning database services, and it gives you the flexibility to elastically scale up, scale out, and scale down IT resources as required. New innovations in Oracle Database 12.2 include:

- Massive savings for consolidated and software-as-a-service (SaaS) environments with up to 4096 pluggable databases
- Increased agility with online cloning and relocation of pluggable databases
- Greatly accelerated in-memory database performance
- Offloading of real-time in-memory analytics to active standby databases
- Native database shading
- Massive scaling with Oracle RAC
- JavaScript Object Notation (JSON) document store enhancements

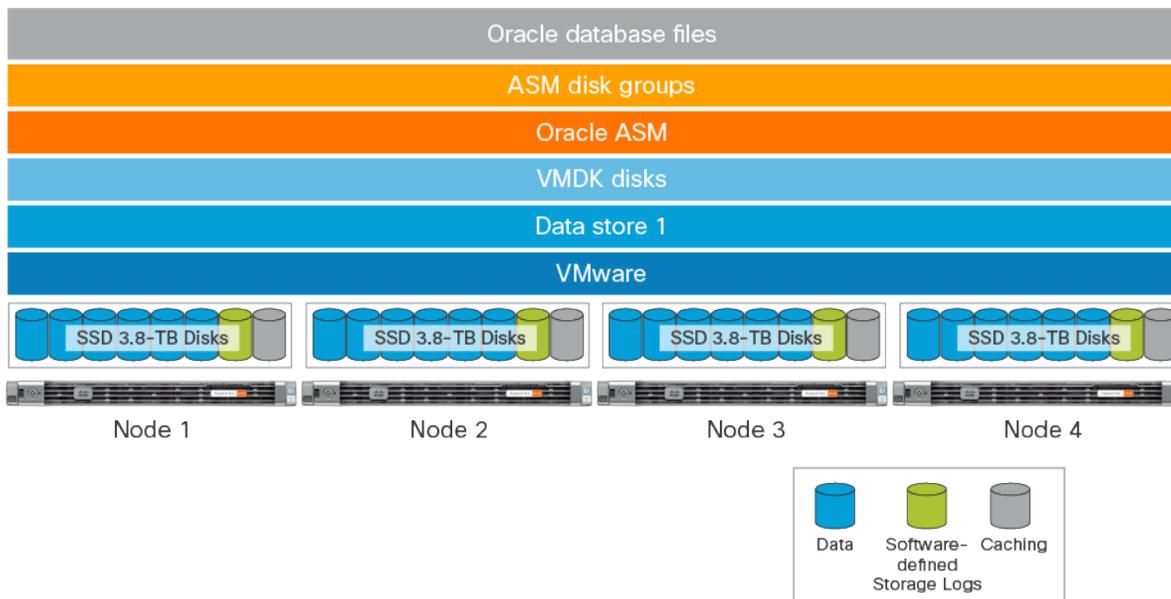
For more information, visit the [Oracle website](#).

**Note:** The validated solution discussed here uses Oracle Database RAC 12c Release 2. Limited testing shows no issues when Oracle Database 12c Release 1 and 11g Release 2 are used for this solution.

## Storage architecture

This reference architecture uses an all-flash configuration. The Cisco HyperFlex HX220c M4 All Flash Nodes allow eight SDDs. However, two are reserved for cluster use. SDDs from all four nodes in the cluster are striped to form a single physical disk pool. (For an in-depth look at the Cisco HyperFlex architecture, see the Cisco white paper Deliver Hyperconvergence with a Next-Generation Platform.) A logical data store is then created for placement of Virtual Machine Disk (VMDK) disks. Figure 4 shows the storage architecture for this environment. This reference architecture uses 3.8-TB SSDs. However, a 960-GB SSD option is also available.

**Figure 4.** Storage architecture



### Storage configuration

This solution uses VMDK disks to create shared storage that is configured as an Oracle Automatic Storage Management, or ASM, disk group. Because all Oracle RAC nodes must be able to access the VMDK disks concurrently, you should configure the Multi-writer option for sharing in the virtual machine disk configuration. For optimal performance, distribute the VMDK disks to the virtual controller using Table 5 for guidance.

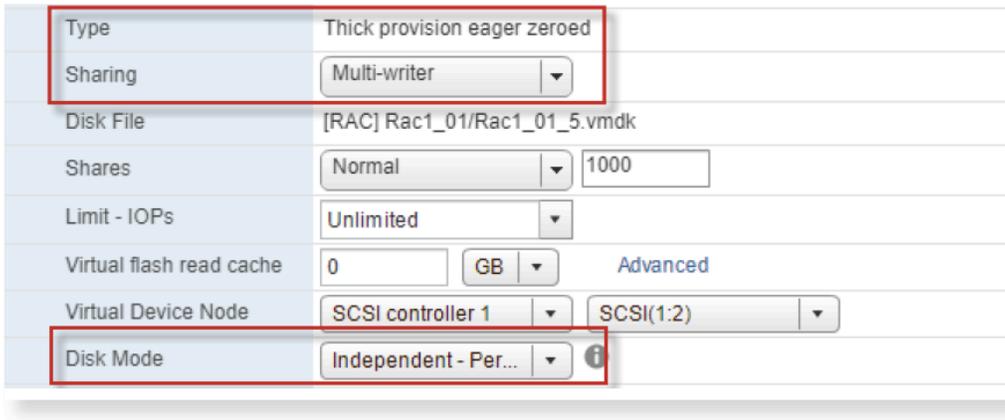
**Table 5.** Assignment of VMDK disks to SCSI controllers;  
Storage layout for each virtual machine (all disks are shared with all 4 Oracle RAC nodes)

SCSI 0 (Paravirtual)	SCSI 1 (Paravirtual)	SCSI 2 (Paravirtual)	SCSI 3 (Paravirtual)
<b>200 GB, OS disk</b>	250 GB, Data 1	250GB, Data 2	70 GB, Log 1
	250 GB, Data 3	250GB, Data 4	70 GB, Log 2
	80 GB, FRA 1	80 GB, FRA 2	70 GB, Log 3
	40 GB, CRS 1	40 GB, CRS 2	
	40 GB, CRS 3	2000 GB, RMAN	

Configure the following settings on all VMDK disks shared by Oracle RAC nodes (Figure 5):

- **Type:** Thick provision eager zeroed
- **Sharing:** Multi-writer
- **Disk Mode:** Independent-Persistent

**Figure 5.** Settings for VMDK disks shared by Oracle RAC nodes



For additional information about the Multi-writer option and the configuration of shared storage, see this [VMware knowledgebase article](#).

Table 6 summarizes the Cisco ASM disk groups for this solution that are shared by all Oracle RAC nodes.

**Table 6.** Oracle ASM disk groups

Oracle ASM disk group	Purpose	Stripe size	Capacity
<b>DATA-DG</b>	Oracle database disk group	4 MB	1000 GB
<b>REDO-DG</b>	Oracle database redo group	4 MB	210 GB
<b>CRS-DG</b>	Oracle RAC Cluster Ready Services disk group	4 MB	120 GB
<b>FRA-DG</b>	Oracle Fast Recovery Area disk group	4 MB	160 GB

## Oracle Database configuration

This section describes the Oracle Database configuration for this solution. Table 7 summarizes the configuration details.

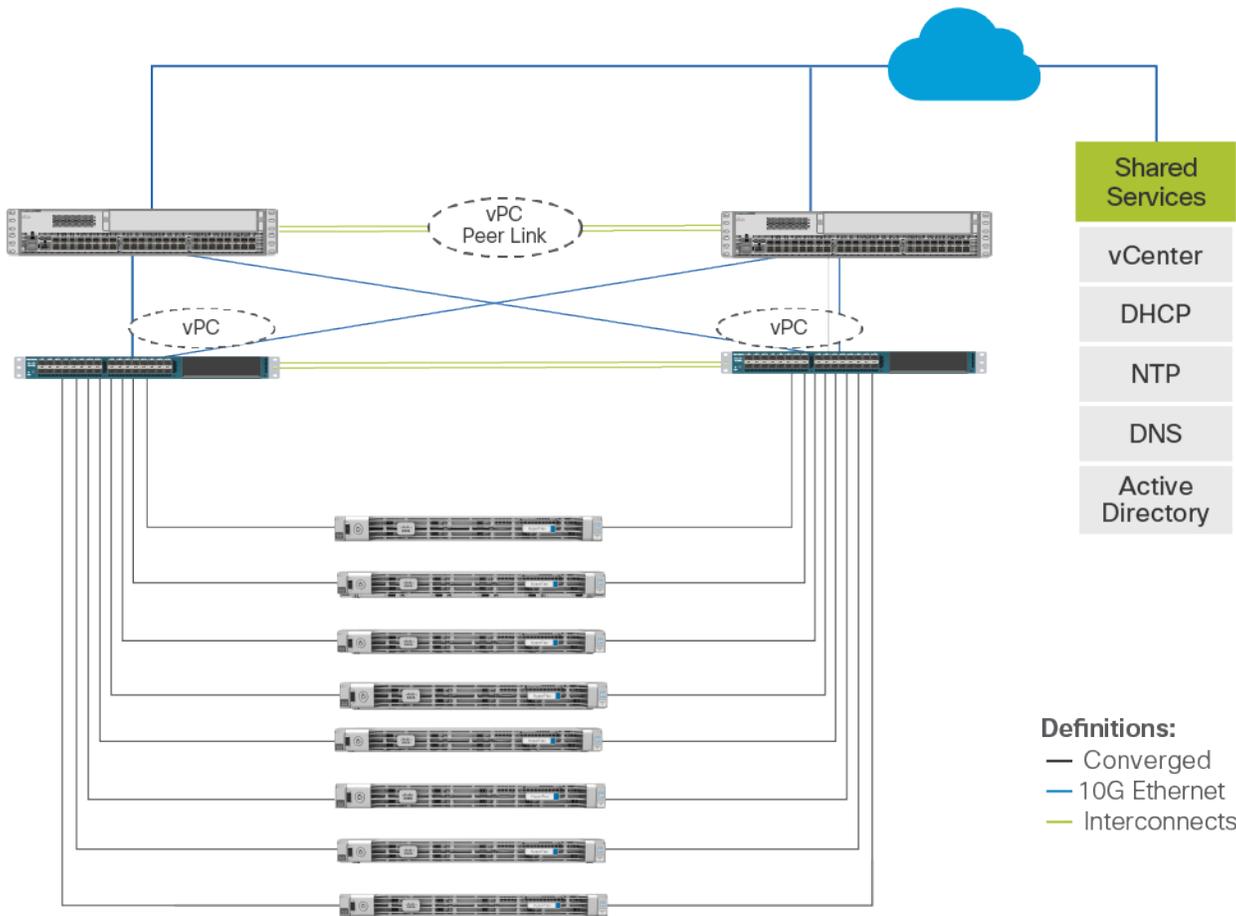
**Table 7.** Oracle Database configuration

Settings	Configuration
<b>SGA_TARGET</b>	96 GB
<b>PGA_AGGREGATE_TARGET</b>	30 GB
<b>Data files placement</b>	ASM and DATA DG
<b>Log files placement</b>	ASM and REDO DG
<b>Redo log size</b>	30 GB
<b>Redo log block size</b>	4 KB
<b>Database block</b>	8 KB

## Network configuration

The Cisco HyperFlex network topology consists of redundant Ethernet links for all components to provide the highly available network infrastructure that is required for an Oracle Database environment. No single point of failure exists at the network layer. The converged network interfaces provide high data throughput while reducing the number of network switch ports. Figure 6 shows the network topology for this environment.

**Figure 6.** Network topology



## Storage configuration

For most deployments, a single data store for the cluster is sufficient, resulting in fewer objects that need to be managed. The Cisco HyperFlex HX Data Platform is a distributed file system that is not vulnerable to many of the problems that face traditional systems that require data locality. A VMDK disk does not have to fit within the available storage of the physical node that hosts it. If the cluster has enough space to hold the configured number of copies of the data, the VMDK disk will fit because the HX Data Platform presents a single pool of capacity that spans all the hyperconverged nodes in the cluster. Similarly, moving a virtual machine to a different node in the cluster is a host migration; the data itself is not moved.

In some cases, however, additional data stores may be beneficial. For example, an administrator may want to create an additional HX Data Platform data store for logical separation. Because performance metrics can be filtered to the data-store level, isolation of workloads or virtual machines may be desired. The data store is thinly provisioned on the cluster. However, the maximum data store size is set during data-store creation and can be used to keep a workload, a set of virtual machines, or end users from running out of disk space on the entire cluster and thus affecting other virtual machines. In such scenarios, the recommended approach is to provision the entire virtual machine, including all its virtual disks, in the same data store and to use multiple data stores to separate virtual machines instead of provisioning virtual machines with virtual disks spanning multiple data stores.

Another good use for additional data stores is to assist in throughput and latency in high-performance Oracle deployments. If the cumulative IOPS of all the virtual machines on a VMware ESX host surpasses 10,000 IOPS, the system may begin to reach that queue depth. In ESXTOP, you should monitor the Active Commands and Commands counters, under Physical Disk NFS Volume. Dividing the virtual machines into multiple data stores can relieve the bottleneck. The default value for ESX queue depth per data store on a Cisco HyperFlex system is 1024.

Another place at which insufficient queue depth may result in higher latency is the SCSI controller. Often the queue depth settings of virtual disks are overlooked, resulting in performance degradation, particularly in high-I/O workloads. Applications such as Oracle Database tend to perform a lot of simultaneous I/O operations, resulting in virtual machine driver queue depth settings insufficient to sustain the heavy I/O processing (the default setting is 64 for PVSCSI). Hence, the recommended approach is to change the default queue depth setting to a higher value (up to 254) as suggested in [VMware knowledgebase article here](#).

For large-scale and high-I/O databases, you always should use multiple virtual disks and distribute those virtual disks across multiple SCSI controller adapters rather than assigning all of them to a single SCSI controller. This approach helps ensure that the guest virtual machine accesses multiple virtual SCSI controllers (four SCSI controllers maximum per guest virtual machine), thus enabling greater concurrency using the multiple queues available for the SCSI controllers.

### Paravirtual SCSI queue depth settings

Large-scale workloads with intensive I/O patterns require adapter queue depths greater than the PVSCSI default values. Current PVSCSI queue depth default values are 64 (for device) and 254 (for adapter). You can increase the PVSCSI queue depths to 254 (for device) and 1024 (for adapter) in a Microsoft Windows or Linux virtual machine.

The following parameters were configured in the design discussed in this document:

- `vmw_pvscsi.cmd_per_lun=254`
- `vmw_pvscsi.ring_pages=32`

For additional information about PSCSI queue depth settings, see the [VMware knowledgebase article here](#).

### Engineering validation: Performance testing

The performance, functions, and reliability of this solution were validated while running Oracle Database in a Cisco HyperFlex environment. The Oracle Swingbench test kit was used to create and test the Order Entry workload, an online transaction processing (OLTP)-equivalent database workload.

This section describes the results that were observed during the testing of this solution. To better understand the performance of each area and component of this architecture, each component was evaluated separately to help ensure that optimal performance was achieved when the solution was under stress. The transactions-per-minute (TPM) metric in the Swingbench benchmark kit was used to measure the performance of Oracle Database.

Note that this section describes the results from some of the tests that were performed as part of the qualification process. These results are presented to provide some data points for the performance observed during the testing. They are not meant to provide comprehensive sizing guidance. For proper sizing of Oracle or other workloads, please use the Cisco HyperFlex Sizer available at <https://hyperflexsizer.cloudapps.cisco.com/>.

**Oracle RAC node scale test**

The node scale test validates the scalability of the Oracle RAC cluster when running the Swingbench test with 200 users. The scale testing consists of four tests using four RAC nodes. The test was run for 60 minutes. The TPM average for the test is reported in Table 8.

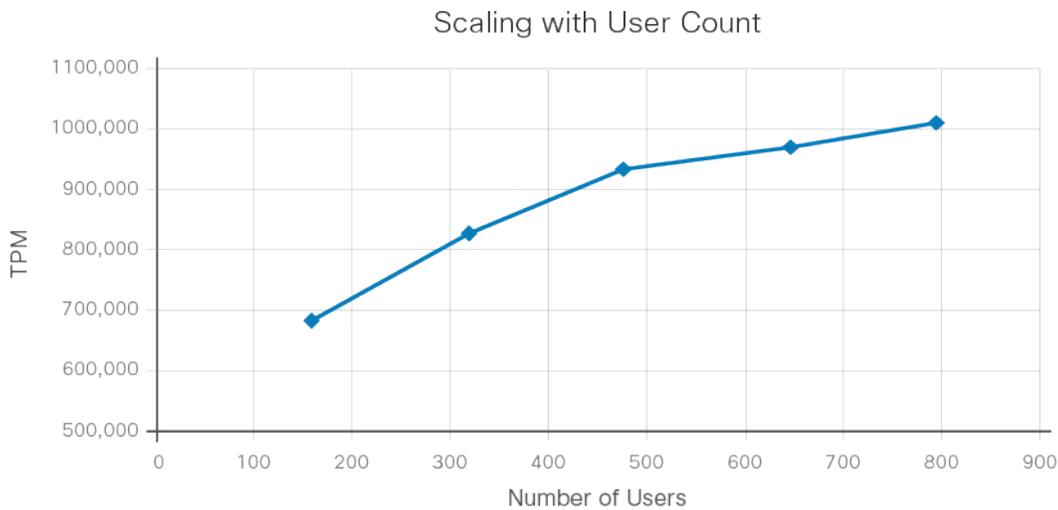
**Table 8.** Oracle node scale test

	Average TPM
4 nodes	552,864

**Oracle user scale test**

The user scale test validates the capability of Oracle RAC to scale with user count. The test starts with 160 users and increases to 800 users. This test shows TPM gains as the user count increases, with the 4-node Cisco HyperFlex HX220c All Flash cluster starting to saturate toward the higher end of the user count (Figure 7).

**Figure 7.** Oracle user scale test



## Reliability and disaster recovery

This section describes some additional reliability and disaster-recovery options available for use with Oracle RAC databases.

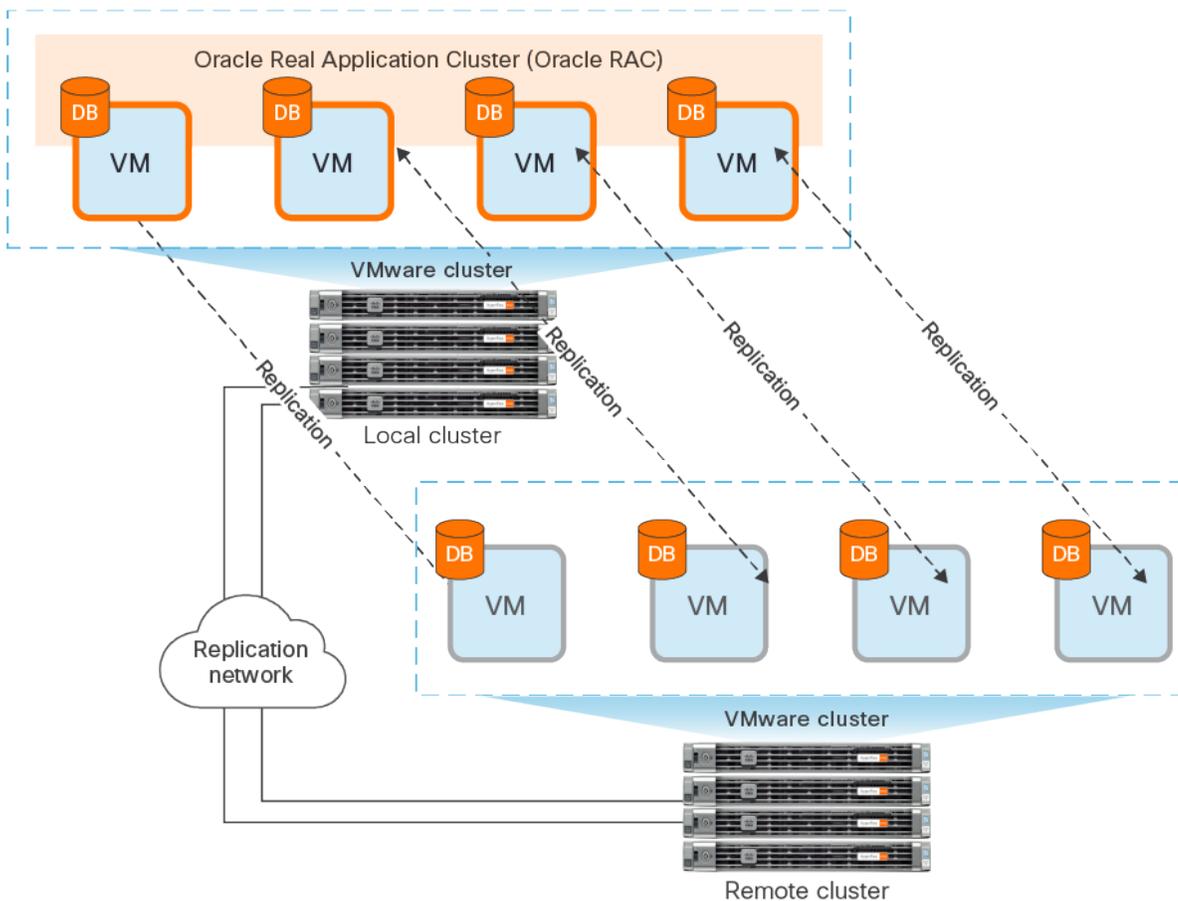
### Disaster recovery

Oracle RAC is usually used to host mission-critical solutions that require continuous data availability to prevent planned or unplanned outages in the data center. Oracle Data Guard is an application-level disaster-recovery feature that can be used with Oracle RAC.

The Cisco HyperFlex HX Data Platform also provides disaster recovery. Note, though, that as of HX Data Platform Release 2.5, you cannot use this feature to protect Oracle RAC virtual machines.

Figure 8 shows a typical disaster-recovery setup. It uses two HX Data Platform clusters: one for the local site and one for the remote site over a WAN connection.

**Figure 8.** Disaster-recovery configuration



## Oracle Data Guard

Oracle Data Guard helps ensure high availability, data protection, and disaster recovery for enterprise data. It provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruptions. Data Guard maintains these standby databases as copies of the production database. Then, if the production database becomes unavailable because of a planned or an unplanned outage, Data Guard can switch any standby database to the production role, reducing the downtime associated with the outage.

Data Guard can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability. Data Guard transport services are also used by other Oracle features such as Oracle Streams and Oracle GoldenGate for efficient and reliable transmission of redo logs from a source database to one or more remote destinations.

## Cisco HyperFlex HX Data Platform disaster recovery

The Cisco HyperFlex HX Data Platform disaster recovery feature allows you to protect virtual machines from a disaster by setting up replication by running virtual machines between a pair of network-connected clusters. Protected virtual machines running on one cluster replicate to the other cluster in the pair, and vice versa. The two paired clusters typically are located at a distance from each other, with each cluster serving as the disaster-recovery site for virtual machines running on the other cluster.

**Note:** The HX Data Platform disaster-recovery feature, as of HX Data Platform Release 2.5, does not support Multi-writer mode virtual machines and so cannot be used protect Oracle RAC virtual machines. Therefore, at this time, you should use Oracle Data Guard for disaster recovery for RAC on Cisco HyperFlex systems.

## Oracle RAC Database backup and recovery: Oracle Recovery Manager

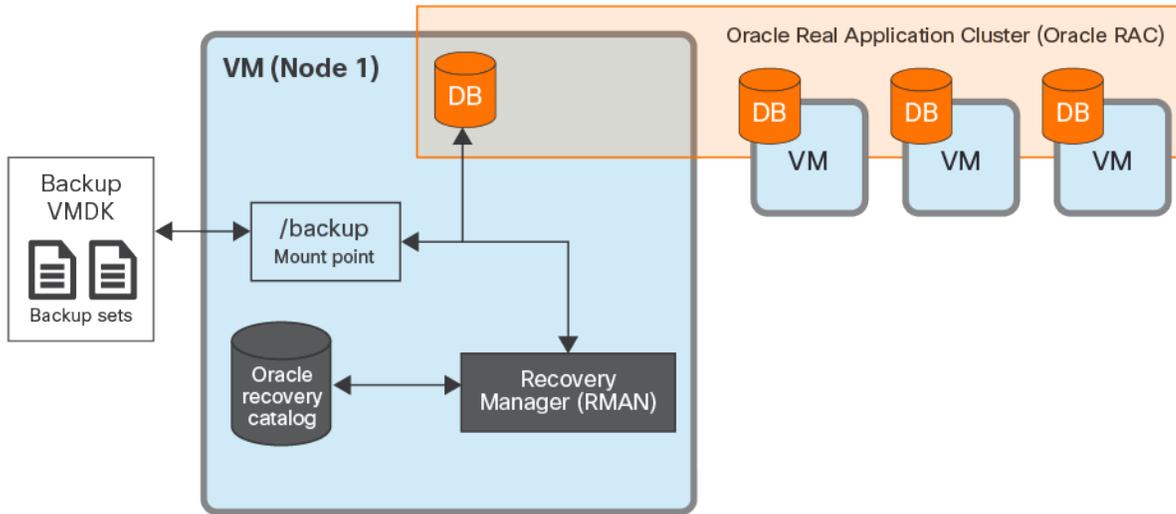
Oracle Recovery Manager (RMAN) is an application-level backup and recovery feature that can be used with Oracle RAC. It is a built-in utility of Oracle Database. It automates the backup and recovery processes. Database administrators (DBAs) can use RMAN to protect data in Oracle databases. RMAN includes block-level corruption detection during both backup and restore operations to help ensure data integrity.

In the lab environment, RMAN was configured on one RAC node for testing. A dedicated 2000-GB VMDK disk was configured as the backup repository under the /backup mount point. By default, RMAN creates backups on disk and generates backup sets rather than image copies. Backup sets can be written to disk or tape.

The RMAN test was performed running a 200-user Swingbench test workload on the Oracle RAC database. The backup operation did not affect the TPM count in the test. However, CPU use on Node 1 increased as a result of the RMAN processes.

Figure 9 shows the RMAN environment for the backup testing. The figure shows a typical setup. In this case, the backup target is the mount point /backup.

**Figure 9.** RMAN test environment



Two separate RMAN backup tests were run: one while the database was idle with no user activity, and one while the database was under active testing with 200 user sessions. Table 9 shows the test results.

**Table 9.** RMAN backup results

	During idle database	During active Swingbench testing
<b>Backup type</b>	Full backup	Full backup
<b>Backup size</b>	204 GB	204 GB
<b>Backup elapsed time</b>	00:14:15	00:15:45
<b>Backup size</b>	204 GB	204 GB
<b>Backup throughput</b>	249 MBps	225 MBps

Table 10 shows the average TPM for two separate Swingbench tests: one when RMAN was used to perform a hot backup of the database, and one for a Swingbench test with no RMAN activity. Test results show little impact on the Oracle RAC database during the backup operation.

**Table 10.** Impact of RMAN backup operation

	Average TPM
<b>RMAN hot backup during Swingbench test</b>	330,493
<b>Swingbench only</b>	395,000

**Note:** These tests and several others were performed to validate the functions. The results presented here are not intended to provide detailed sizing guidance or guidance for performance headroom calculations, which are beyond the scope of this document.

## Conclusion

Cisco is a leader in the data center industry. Our extensive experience with enterprise solutions and data center technologies enables us to design and build an Oracle RAC Database reference architecture on hyperconverged solutions that is fully tested, protecting our customer's investment and offering a high-level ROI. The Cisco HyperFlex architecture helps enable databases to achieve optimal performance with very low latency—features that are critical for enterprise-scale applications.

Cisco HyperFlex systems provide the design flexibility needed to engineer a highly integrated Oracle Database system to run enterprise applications that use industry best practices for a virtualized database environment. The balanced, distributed data-access architecture of Cisco HyperFlex systems supports Oracle scale-out and scale-up models, reducing the hardware footprint and increasing data center efficiency.

This solution delivers many business benefits, including the following:

- Increased scalability
- High availability
- Reduced deployment time with a validated reference configuration

Cisco HyperFlex systems provide the following benefits for this reference architecture:

- Optimized performance for transactional applications with little latency
- Balanced and distributed architecture that increases performance and IT efficiency with automated resource management
- Capability to start with a small investment and grow as business demand increases

## For more information

For additional information, consult the following resources:

- [Cisco HyperFlex systems white paper, Deliver Hyperconvergence with a Next-Generation Data Platform](#)
- [Cisco HyperFlex systems solution overview](#)
- [Cisco HyperFlex all-flash systems for Oracle Database deployments](#)
- [Oracle Databases on VMware Best Practices Guide, Version 1.0, May 2016](#)

Americas Headquarters  
Cisco Systems, Inc.  
San Jose, CA

Asia Pacific Headquarters  
Cisco Systems (USA) Pte. Ltd.  
Singapore

Europe Headquarters  
Cisco Systems International BV Amsterdam,  
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Printed in USA

C11-739679-00 09/17