

Cisco HyperFlex All-NVMe Systems for Oracle Database: Reference Architecture

Contents

Executive summary	3
Cisco HyperFlex HX Data Platform all-NVMe storage	4
Why use Cisco HyperFlex All-NVMe systems for Oracle Database deployments	4
Oracle Database 19c on Cisco HyperFlex systems	6
Solution overview	10
Solution components	11
Solution design	15
Engineering validation	19
Performance comparison between all-NVMe and all-flash clusters	23
Conclusion	24
For more information	26

Executive summary

Organizations of all sizes require an infrastructure for their Oracle Database deployments that provides high performance and enterprise-level reliability for transactional online transaction processing (OLTP) workloads. Oracle Database is a leading relational database management system (RDBMS). It has a large installed base that is now running on outdated hardware architecture that often no longer provides the levels of performance and scale that enterprises require. With the Cisco HyperFlex™ solution for Oracle Database, organizations can implement Oracle databases using a highly integrated solution that scales as business demand increases. 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 OLTP workload, enabling rapid deployment of Oracle databases. This configuration meets the industry best practices for Oracle Database in a VMware virtualized environment.

Cisco HyperFlex HX Data Platform all-NVMe 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-Non-Volatile Memory Express (NVMe) 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](#)

An all-NVMe storage solution delivers more of what you need to propel mission-critical workloads. For a simulated Oracle OLTP workload, it provides 71 percent more I/O operations per second (IOPS) and 37 percent lower latency than our previous-generation all-flash node. The behavior discussed here was tested on a Cisco HyperFlex system with NVMe configurations, and the results are provided in the "[Engineering validation](#)" section of this document. A holistic system approach is used to integrate Cisco HyperFlex HX Data Platform software with Cisco HyperFlex HX220c M5 All NVMe Nodes. The result is the first fully engineered hyperconverged appliance based on NVMe storage.

- **Capacity storage:** The data platform's capacity layer is supported by Intel 3D NAND NVMe solid-state disks (SSDs). These drives currently provide up to 32 TB of raw capacity per node. Integrated directly into the CPU through the PCI Express (PCIe) bus, they eliminate the latency of disk controllers and the CPU cycles needed to process SAS and SATA protocols. Without a disk controller to insulate the CPU from the drives, we have implemented reliability, availability, and serviceability (RAS) features by integrating the Intel Volume Management Device (VMD) into the data platform software. This engineered solution handles surprise drive removal, hot pluggability, locator LEDs, and status lights.
- **Cache:** A cache must be even faster than the capacity storage. For the cache and the write log, we use Intel® Optane™ DC P4800X SSDs for greater IOPS and more consistency than standard NAND SSDs, even in the event of high-write bursts.
- **Compression:** The optional Cisco HyperFlex Acceleration Engine offloads compression operations from the Intel® Xeon® Scalable CPUs, freeing more cores to improve virtual machine density, lowering latency, and reducing storage needs. This helps you get even more value from your investment in an all-NVMe platform.
- **High-performance networking:** Most hyperconverged solutions consider networking as an afterthought. We consider it essential for achieving consistent workload performance. That's why we fully integrate a 40-Gbps unified fabric into each cluster using Cisco Unified Computing System™ (Cisco UCS®) fabric interconnects for high-bandwidth, low-latency, and consistent-latency connectivity between nodes.
- **Automated deployment and management:** Automation is provided through Cisco Intersight™, a software-as-a-service (SaaS) management platform that can support all your clusters from the cloud to wherever they reside in the data center to the edge. If you prefer local management, you can host the Cisco Intersight Virtual Appliance, or you can use Cisco HyperFlex Connect management software.

All-NVMe solutions support most latency-sensitive applications with the simplicity of hyperconvergence. Our solutions provide the first fully integrated platform designed to support NVMe technology with increased performance and RAS.

Why use Cisco HyperFlex All-NVMe systems for Oracle Database deployments

Oracle Databases act 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-NVMe hyperconverged systems offer the following advantages, making them well suited for Oracle Database implementations:

-
- **High performance:** NVMe nodes deliver the highest performance for mission-critical data center workloads. They provide architectural performance to the edge with NVMe drives connected directly to the CPU rather than through a latency-inducing PCIe switch.
 - **Ultra-low latency with consistent performance:** Cisco HyperFlex all-NVMe systems, when used to host the virtual database instances, deliver extremely 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 storage IOPS capability and make use of 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.

Oracle Database 19c on Cisco HyperFlex systems

This reference architecture guide shows how Cisco HyperFlex systems can provide intelligent end-to-end automation with network-integrated hyperconvergence for an Oracle Database deployment. The Cisco HyperFlex system provides a high-performance, integrated solution for Oracle Database environments.

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.

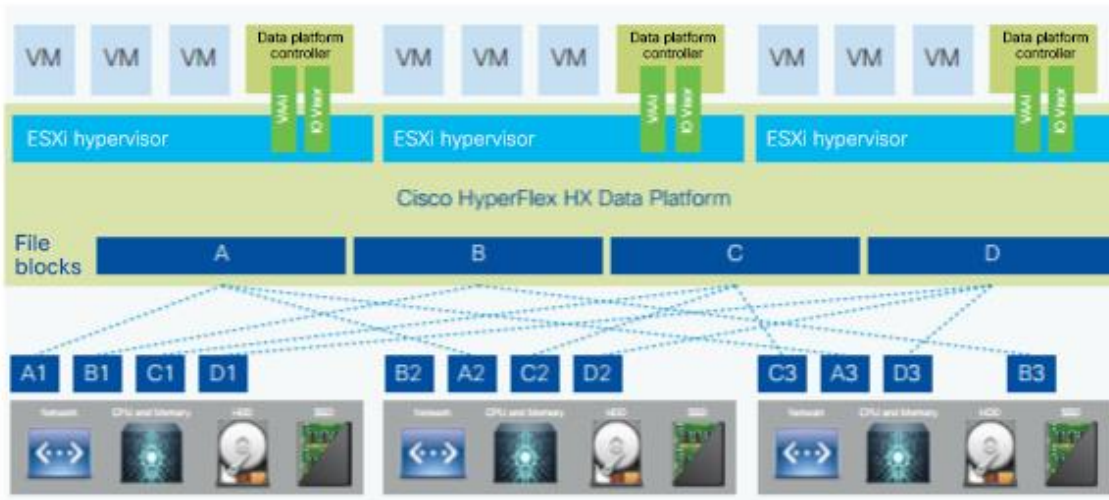


Figure 1.
Data distribution architecture

This reference architecture uses a cluster of four Cisco HyperFlex HX220c M5 All NVMe nodes to provide fast data access. Use this document to design an Oracle Database 19c solution that meets your organization's requirements and budget.

This hyperconverged solution integrates servers, storage systems, network resources, and storage software to provide an enterprise-scale environment for an Oracle Database deployment. This highly integrated environment provides reliability, high availability, scalability, and performance to handle large-scale transactional workloads. This is an Oracle on VMware solution utilizing Oracle Enterprise Linux 7 as the OS across two VM profiles to validate performance, scalability, and reliability which achieves the best interoperability for Oracle databases.

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](#).

This reference architecture document is written for the following audience:

- Database administrators
- Storage administrators
- IT professionals with the responsibility of planning and deploying an Oracle Database solution

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

- Hyperconvergence technology
- Virtualized environments
- SSD and all-NVMe storage
- Oracle Database 19c
- Oracle Automatic Storage Management (ASM)
- Oracle Enterprise Linux

Oracle Database scalable architecture overview

This reference architecture uses two different Oracle virtual server configuration profiles to validate the scale-up and scale-out architecture of this environment. Table 1 summarizes the two profiles. The use of two different configuration profiles allows additional results and performance characteristics to be observed. The Oracle Medium (MD) and Large (LG) profiles are base configurations commonly used in many production environments. However, this architecture can support different database configurations and sizes to fit your deployment requirements.

Table 1. Oracle server configuration profiles

Specification	Oracle Medium (MD)	Oracle Large (LG)
Virtual CPU (vCPU)	8	16
Virtual RAM (vRAM)	64 GB	96 GB
Database size	300 GB	300 GB
Schema	32	128

Customers have two options for scaling their Oracle databases: scale up and scale out. Either of these is well suited for deployment on Cisco HyperFlex systems.

Oracle scale-up architecture

The scale-up architecture shows the elasticity of the environment, increasing virtual resources to virtual machines on demand as the application workload increases. This model enables the Oracle database to scale up or down for better resource management. In this model, customers can grow the Oracle MD virtual machine into the Oracle LG configuration with no negative impact on Oracle Database data and performance. The scale-up architecture is shown in Figure 2.

The distributed architecture of the Cisco HyperFlex system allows a single Oracle database running on a single virtual machine to consume and properly use resources across all cluster nodes, thus allowing a single database to achieve the highest cluster performance at peak operation. These characteristics are critical for any multitenant database environment in which resource allocation may fluctuate.

A Cisco HyperFlex all-NVMe cluster supports up to 16 nodes, with the capability to extend to 32 nodes by adding 16 external computing-only nodes. This approach allows any deployment to start with a smaller environment and grow as needed, for a "pay as you grow" model.

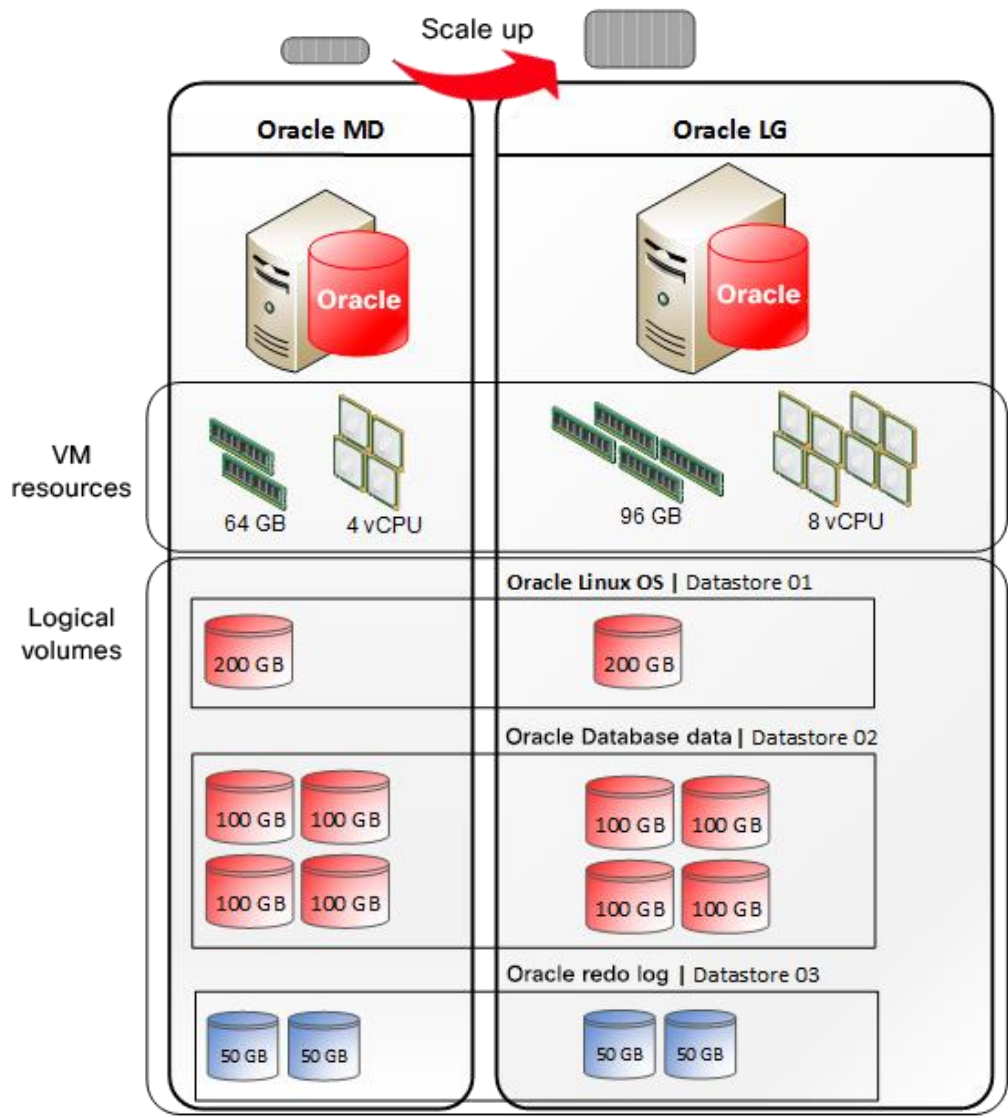


Figure 2.
Oracle scale-up architecture

Oracle scale-out architecture

The scale-out architecture validates the capability to add Oracle Database instances within the same cluster to support scale-out needs such as database consolidation to increase the ROI for infrastructure investment. This model also enables rapid deployment of Oracle Real Application Clusters (RAC) nodes for a clustered Oracle Database environment. Figure 3 shows the scale-out architecture.

Oracle RAC is beyond the scope for this reference architecture. This document focuses on single-instance Oracle Database deployments and validates different sizes of Oracle VM and the use of several virtual machines running simultaneously

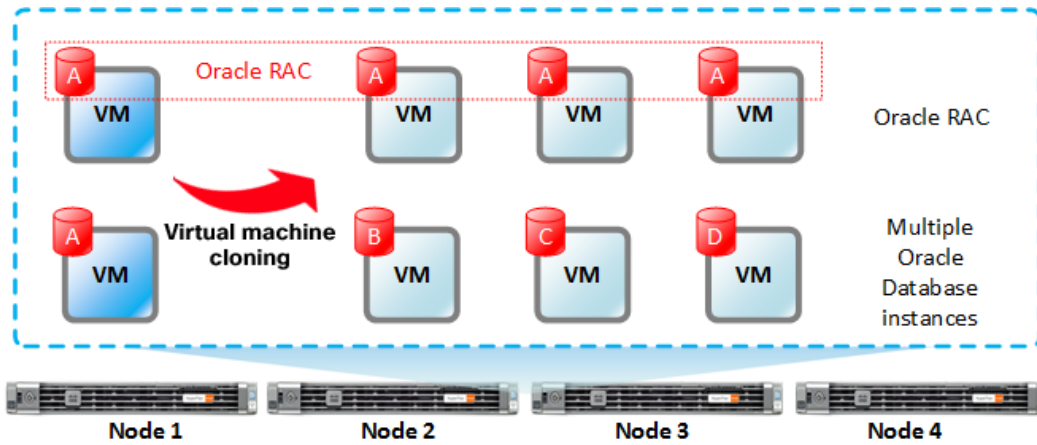


Figure 3.
Oracle scale-out architecture

Solution overview

This section describes how to implement Oracle Database 19c on a Cisco HyperFlex system using a 4-node cluster to provide 2-node failure redundancy. This reference configuration helps ensure proper sizing and configuration when you deploy Oracle Database 19c 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. During the validation and testing phases, only virtual machines that are under test are powered on. Figure 4 presents a high-level view of the environment.

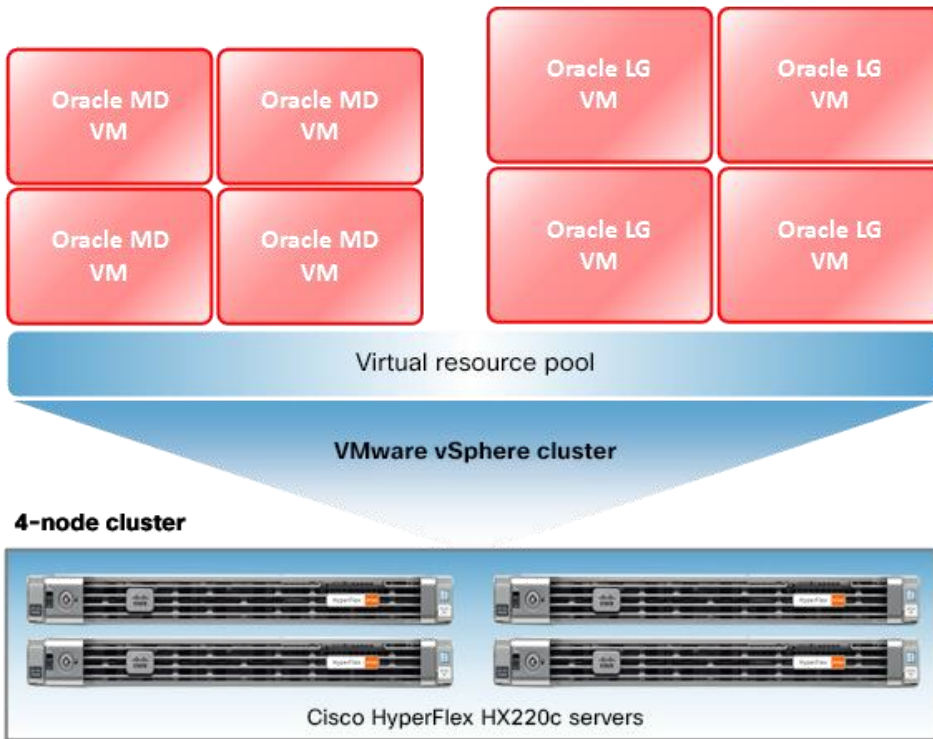


Figure 4.
High-level view of solution design

Solution components

This section describes the components of this solution. Table 2 summarizes the main components of the solution. Table 3 summarizes the HX220c M5 Node configuration for the cluster.

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 M5 All NVMe Nodes

Nodes with all-NVMe 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 (Flex Flash) Secure Digital (SD) cards, a single 120-GB SSD data-logging drive, a single SSD write-log drive and up to six 4-TB NVMe SSD drives, for a contribution of up to 24-TB of raw storage capacity to the cluster. The nodes use Intel Xeon processor 6140 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 M5 All NVMe Nodes for a four-node server cluster to provide two-node failure reliability.

See the [Cisco HX220c M5 All NVMe Node data sheet](#) for more information.

Cisco UCS 6200 Series Fabric Interconnects

The Cisco UCS 6200 Series Fabric Interconnects are a core part of 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 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 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.

Table 2. Reference architecture components

Hardware	Description	Quantity
Cisco HyperFlex HX220c M5 All NVMe Node servers	Cisco 1RU hyperconverged nodes that allow cluster scaling in a small footprint	4
Cisco UCS 6200 Series Fabric Interconnects	Fabric interconnects	2

Table 3. Cisco HyperFlex HX220c M5 Node configuration

Description	Specifications	Notes
CPU	2 Intel Xeon Gold 6140 CPUs at 2.30 GHz	
Memory	24 x 32-GB DIMMs	
Cisco FlexFlash SD card	240-GB SD cards	Boot drives
SSD	500-GB SSD	Configured for housekeeping tasks
	375-GB SSD	Configured as cache
	6 x 4-TB NVMe SSD	Capacity disks for each node
Hypervisor	VMware vSphere, 6.5.0	Virtual Platform for HX Data Platform software
Cisco HyperFlex HX Data Platform software	Cisco HyperFlex HX Data Platform Release 4.0(1b)	

Software components

This section describes the software components used for this solution. Table 4 lists the software and versions for the solution.

VMware vSphere

VMware vSphere helps you get the best performance, availability, and efficiency from your infrastructure while dramatically reducing the hardware footprint and your capital expenses 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:

- **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.
- **Distributed Resource Scheduler:** DRS monitors resource utilization and intelligently allocates system resources as needed.
- **High Availability:** HA monitors hardware and OS failures and automatically restarts the virtual machine, providing cost-effective failover.
- **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 19c

Oracle Database 19c now provides customers with a high-performance, reliable, and secure platform to easily and cost-effectively modernize their transactional and analytical workloads on-premises. It offers the same familiar database software running on-premises that enables customers to use the Oracle applications they have developed in-house. Customers can therefore continue to use all their existing IT skills and resources and get the same support for their Oracle databases on their premises.

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

Note: The validated solution discussed here uses Oracle Database 19c Release 3. Limited testing shows no issues with Oracle Database 19c Release 3 or 12c Release 2 for this solution.

Table 4. Reference architecture software components

Software	Version	Function
Cisco HyperFlex HX Data Platform	Release 4.0(1b)	Data platform
Oracle Enterprise Linux Oracle UEK Kernel	Version 7.6 4.14.35-1902.3.1.el7uek.x86_64 x86_64	OS for Oracle RAC Kernel version in Oracle Linux
Oracle Grid and Automatic Storage Management (ASM)	Version 19c Release 3	Automatic storage management
Oracle Database	Version 19c Release 3	Oracle Database system
Oracle Silly Little Oracle Benchmark (SLOB)	Version 2.4	Workload suite

Solution design

This section describes the architecture for this reference environment. It describes how to implement scale-out and scale-up Oracle Database environments using Cisco HyperFlex HX220c M5 nodes on All-NVMe SSDs for Oracle transactional workloads.

Storage architecture

This reference architecture uses an all-NVMe configuration. The HX220c M5 All NVMe Nodes allow eight NVMe SSDs. However, two per node are reserved for cluster use. NVMe SSDs 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 datastore is then created for placement of Virtual Machine Disk (VMDK) disks. The storage architecture for this environment is shown in Figure 5. This reference architecture uses 4-TB NVMe SSDs.

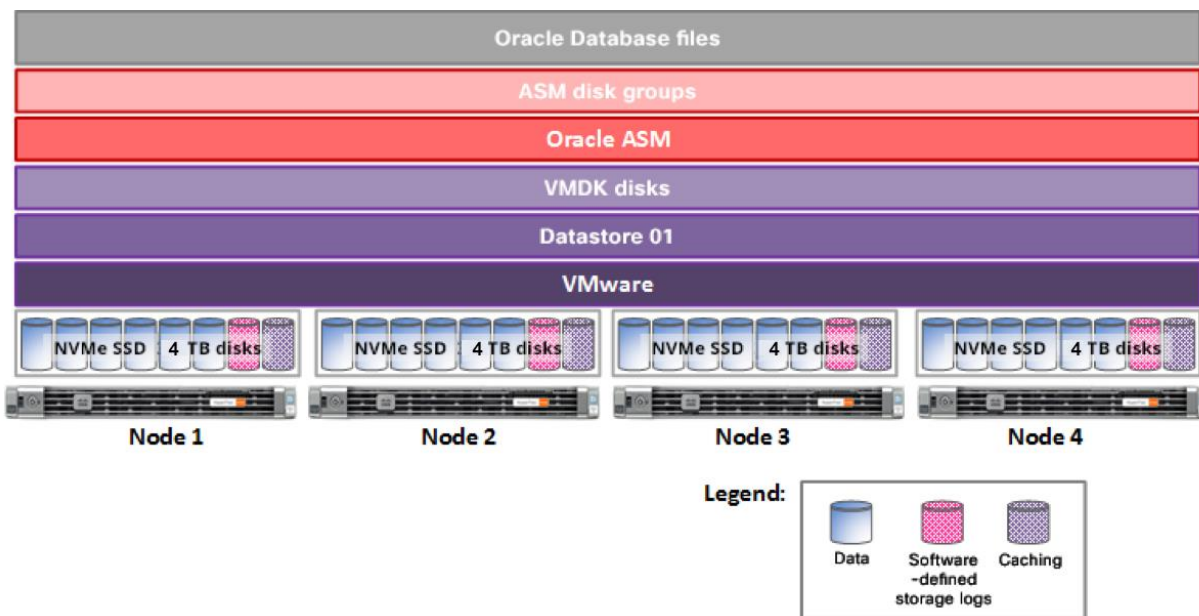


Figure 5.
Storage architecture

Storage configuration

This solution uses VMDK disks to create Linux disk devices. The disk devices are used as ASM disks to create the ASM disk groups (Table 5).

Note: In general both HX and Oracle ASM provide RF factor. But in our test environment, we are only using the RF factor provided by HX not by the Oracle ASM. The capacities vary depending on the RF factor being set (If RF is set to 2, actual capacities are one-half of raw capacity and If RF is set to 3, actual capacities are one-third of raw capacity).

Table 5. Oracle ASM disk groups

Oracle ASM disk group	Purpose	Stripe size	Capacity
DATA-DG	Oracle Database disk group	4 MB	400 GB
REDO-DG	Oracle Database redo group	4 MB	100 GB

Virtual machine configuration

This section describes the Oracle Medium and Oracle Large virtual machine profiles. These profiles are properly sized for a 300-GB Oracle database (Table 6).

Table 6. Oracle VM configuration

Resource	Details for Oracle MD	Details for Oracle LG
VM specifications	8 × vCPU 64-GB vRAM	16 × vCPU 96 GB vRAM
VM controllers	4 × Paravirtual Small Computer System Interface (PVSCSI) controller	4 × Paravirtual Small Computer System Interface (PVSCSI) controller
VM disks	1 × 200-GB VMDK for VM OS 4 × 100-GB VMDK for Oracle data 2 × 50-GB VMDK for Oracle redo log	1 × 200-GB VMDK for VM OS 4 × 100-GB VMDK for Oracle data 2 × 50-GB VMDK for Oracle redo log

Oracle Database configuration

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

Table 7. Oracle Database configuration

	Oracle MD	Oracle LG
SGA_TARGET	53 GB	77 GB
PGA_AGGREGATE_TARGET	10 GB	10 GB
Data files placement	ASM and DATA DG	
Log files placement	ASM and REDO DG	
Redo log size	30 GB	30 GB
Redo log block size	4 KB	4 KB
Database block	8 KB	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.

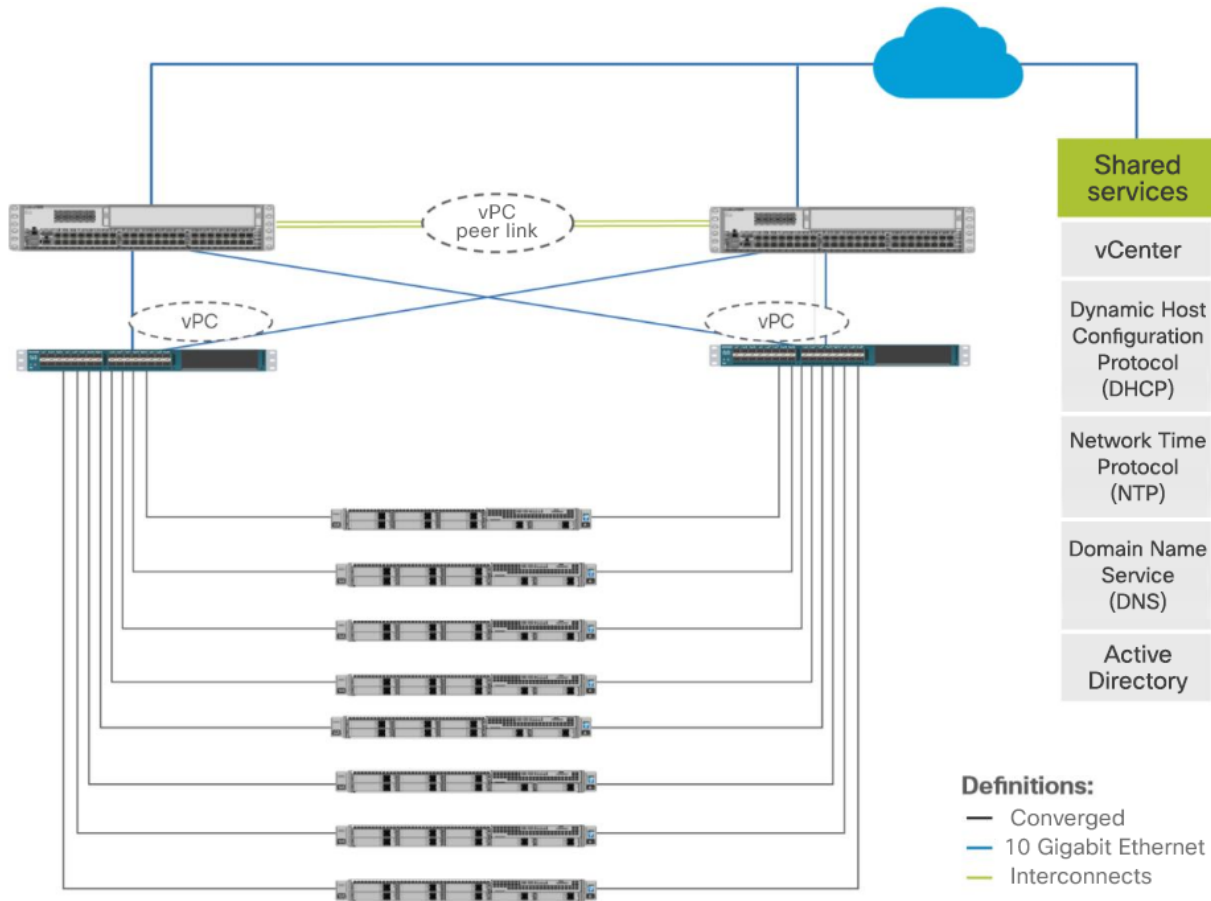


Figure 6.
Network topology

Storage configuration

For most deployments, a single datastore for the cluster is enough, resulting in fewer objects that need to be managed. The 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 datastores may be beneficial. For example, an administrator may want to create an additional HX Data Platform datastore for logical separation. Because performance metrics can be filtered to the data-store level, isolation of workloads or virtual machines may be desired. The datastore is thinly provisioned on the cluster. However, the maximum datastore 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 datastore and to

use multiple datastores to separate virtual machines instead of provisioning virtual machines with virtual disks spanning multiple datastores.

Another good use for additional datastores is to assist in throughput and latency in high-performance Oracle deployments. If the cumulative IOPS of all the virtual machines on an 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 datastores or increasing the ESX queue limit (the default value is 256) to up to 1024 can relieve the bottleneck.

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 this [VMware knowledgebase article](#).

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.

Another operation that helps reduce application-level latency is to change the amount of time that the virtual machine NIC (vmnic) spends trying to coalesce interrupts. (More details about the parameter can be found [here](#).)

```
esxcli network nic coalesce set -n vmnic3 -r 15 -t 15
```

```
esxcli network nic coalesce set -n vmnic2 -r 15 -t 15
```

Engineering validation

This section describes the methodology used to test and validate this solution. The performance, functions, and reliability of this solution were validated while running Oracle Database in a Cisco HyperFlex environment. The SLOB test kit was used to create and test an OLTP equivalent database workload.

The test includes:

- SLOB profile of 75 percent read and 25 percent write operations
- Testing of scale-up options using different virtual machine sizes and multiple virtual machines running simultaneously

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/>.

Test methodology

The test methodology validates the computing, storage, and database performance advantages of Cisco HyperFlex systems for Oracle Database. These scenarios also provide data to help you understand the overall capabilities when scaling Oracle databases.

This test methodology uses the SLOB test suite to simulate an OLTP-like workload. It consists of 75 percent read and 25 percent write I/O distribution to mimic an online transactional application.

Test results

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.

Single Oracle Large virtual machine test

This test uses a single Oracle Large virtual machine configuration. Disk-level performance was captured at both at the virtual machine and Cisco HyperFlex levels. The test duration is 60 minutes with 128 users. Figure 7 and Table 8 show the application-level performance for this test.

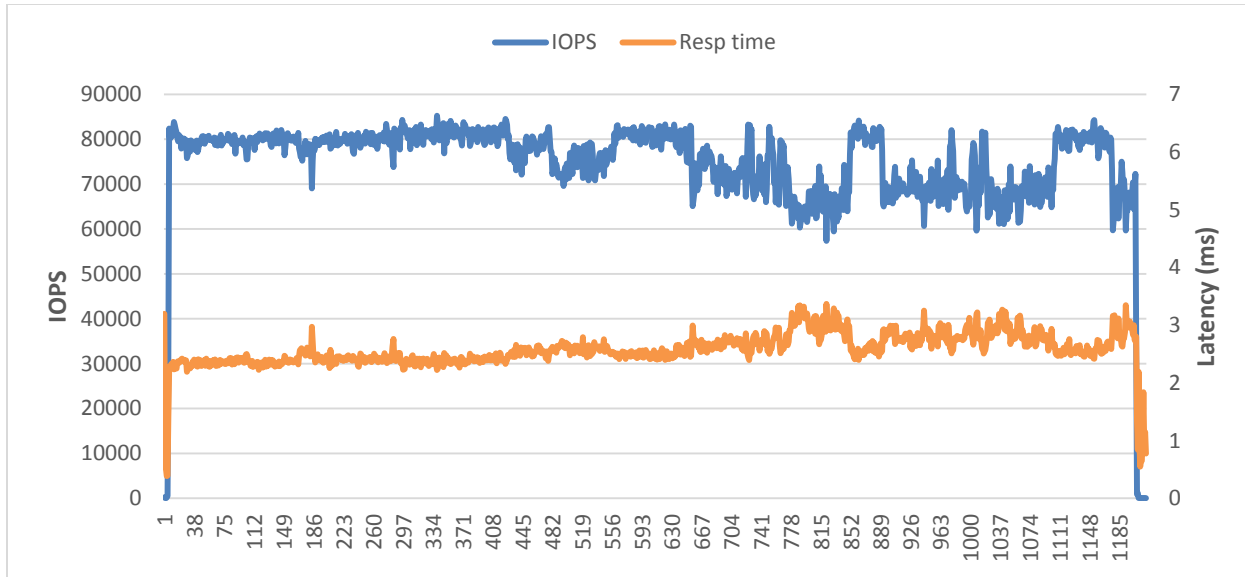


Figure 7.
Single Oracle Large virtual machine test application-level performance

The results show that the IOPS values are consistent between the virtual machine and Cisco HyperFlex underlying storage.

Table 8. Table 8 Application-level performance

Average IOPS	Average latency (ms)
74,692.8	2.5

Scaling test using Oracle Medium virtual machines

This test uses four Oracle Medium virtual machines to show the capability of the environment to scale as additional virtual machines were used to run the SLOB workload. Scaling is very close to linear. Figure 8 shows the scale testing results.

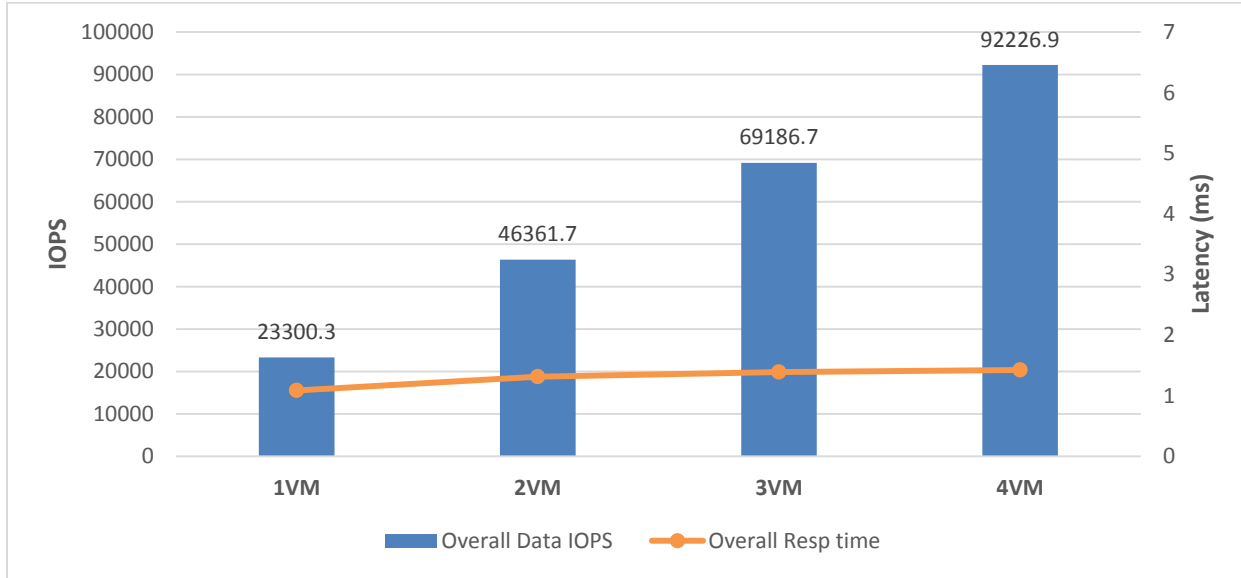


Figure 8.
Virtual machine scale test results

This feature gives you the flexibility to partition the workload according to the needs of the environment.

Concurrent workloads test

The concurrent workloads test uses four Oracle Medium virtual machine instances to test the solution's capability to run many Oracle instances concurrently. The test ran for one hour and revealed that this solution stack has no bottleneck and that the underlying storage system can sustain high performance throughout the test. Figure 9 shows the performance results for this test. Cisco HyperFlex maximum IOPS performance for this specific all-NVMe configuration was achieved while maintaining the ultra-low latency. Note that the additional IOPS come at the cost of a moderate increase in latency.

Additional workloads can be added to this test, but at the cost of slightly higher latency. Figure 9 and Table 9 show the results of the concurrent workload test.




Figure 9. Concurrent workload application-level performance

Table 9. Concurrent workload application latency test summary

Average IOPS	Average latency (ms)
200,265.3	3.9

This test demonstrates the unique capability of the Cisco HyperFlex system to provide a single pool of storage capacity and storage performance resources across the entire cluster. You can scale your Oracle databases using the following models:

- Scale the computing resources dedicated to the virtual machine running Oracle Database, thus removing a computing bottleneck that may be limiting the total overall performance
- Scale by adding virtual machines running Oracle Database or use the additional performance headroom for other applications

Performance comparison between all-NVMe and all-flash clusters

Increasing performance at the end-user level requires a holistic approach in designing your solutions. Simply adding low-latency storage is not enough. Therefore, Cisco HyperFlex systems have been designed to provide balanced high performance with very low latency. Configurations using all-flash and all-NVMe nodes, together with our standard high-throughput network and fast computing, supports consistently high performance even for large databases. The distributed architecture provides every virtual machine with access to high IOPS regardless of the physical location of the virtual machine. This capability is important for virtual machines running Oracle Database because they frequently need higher performance to handle bursts of application or user activity. All-NVMe storage performs even faster (as shown by the testing reported here) and is excellent for databases that require ultra-low latency.

Table 10 compares the performance of the two cluster types.

Table 10. Comparison matrix

This section describes the key differences between All NVMe and All Flash test setups:	All-NVMe test setup	All-flash test setup
Cisco HyperFlex HX Data Platform	Release 4.0(1b)	Release 2.1(1b)
Cisco HyperFlex servers	HX220c M5 All NVMe Node servers	HX220c M4 All Flash Node servers
Oracle Database	Version 19c Release 3	Version 12c Release 2
Oracle SLOB	Version 2.4	Version 2.3
Oracle Grid and ASM	Version 19c Release 3	Version 12c Release 2
Replication factor	3	2
Hypervisor	Version 6.5.0	Version 6.0

See the [Cisco HyperFlex and Oracle white paper](#) about all-flash setup for detailed information.

Test results comparing all-NVMe and all-flash clusters

This section describes the results that were observed during the testing of all-NVMe and all-flash solutions. To better understand the performance of each area and component of this architecture (both all-NVMe and all-flash solutions), each component was evaluated separately to help ensure that optimal performance was achieved when the solution was under stress.

Concurrent workload test

The test results presented here show a difference in the test setups of the all-NVMe and all-flash clusters. The concurrent workload scenario for the all-NVMe cluster uses four Medium test virtual machines, and the scenario for the all-flash cluster uses only three Medium test virtual machines. However, the graphs displayed in Figure 10 shows that application-level performance for the all-NVMe solution is much higher at a cost of very low latency than the all-flash solution. One of the major advantages of all-NVMe storage is higher performance at ultra-low latency, and similar behavior is observed in the test results here. Although the load generated is higher on the all-NVMe cluster than on the all-flash cluster, the all-NVMe cluster shows a 30 percent gain in IOPS at very low latency. The test ran for a duration of one hour. The solution had no bottleneck, and the underlining storage system sustained high performance throughout the test. Figure 10 shows the performance results for this test. Cluster- and application-level latencies are significantly lower for the all-NVMe system.

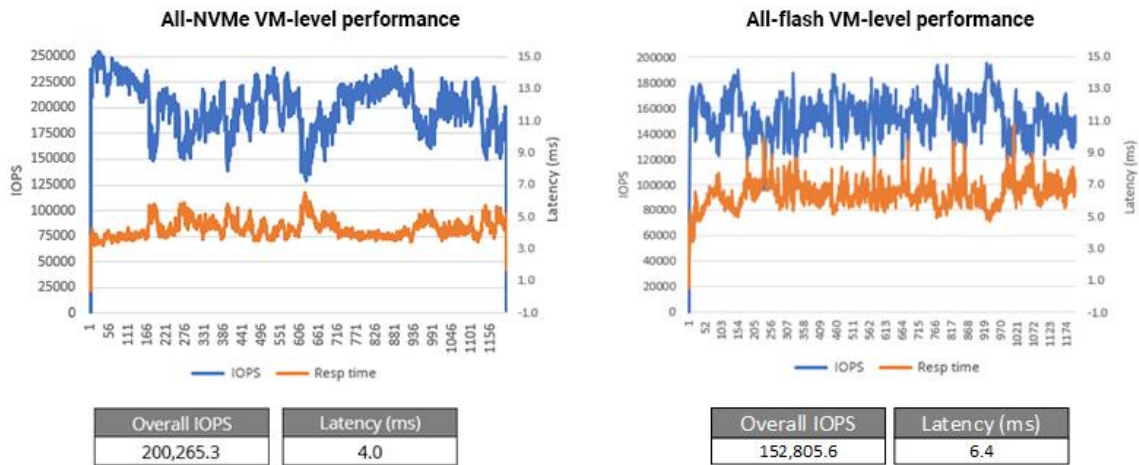


Figure 10. Concurrent workload test

Scaling test using Oracle Medium virtual machines

This test used four Oracle Medium virtual machines to show the capability of the environment to scale as additional virtual machines were used to run the SLOB workload. The performance was observed at the virtual machine level on both the all-NVMe and all-flash clusters. The graphs in Figure 11 show that performance on the all-NVMe cluster is higher than that on the all-flash cluster. In this scale test, as the number of virtual machines increases, the all-NVMe cluster shows a 10 to 15 percent gain in IOPS at each load point at very low latency compared to the all-flash cluster.

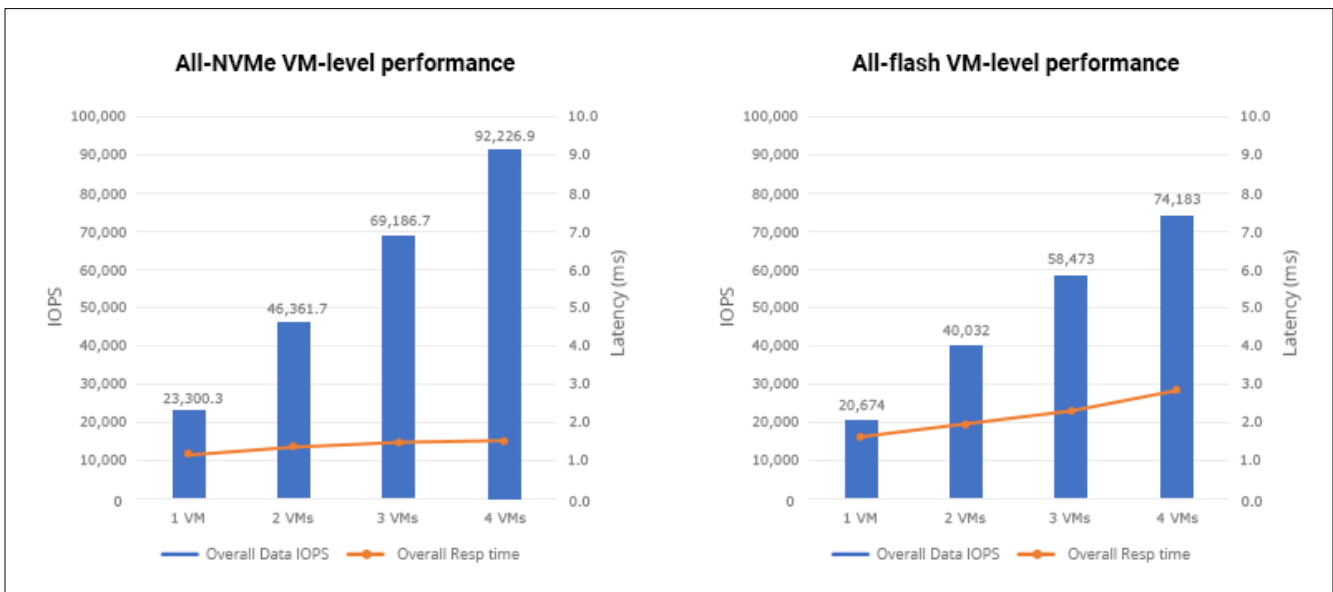


Figure 11. Virtual machine scale test results

Conclusion

Cisco's extensive experience with enterprise solutions and data center technologies enables us to design and build Oracle Database reference architecture on hyperconverged solutions that is fully tested to protect our customers' investments and

offer high ROI. Cisco HyperFlex systems enable databases to perform at optimal performance with very low latency: capabilities 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 meet industry best practices for a virtualized database environment. The balanced and distributed data access architecture of Cisco HyperFlex systems support Oracle scale-out and scale-up environments that reduce the hardware footprint while increasing data center efficiency.

By deploying Cisco HyperFlex systems with all-NVMe configurations, you can run your database deployments on an agile platform that delivers insight in less time and at less cost.

The comparison graphs provided in the "Engineering validation" section of this document show that the all-NVMe storage configuration has higher performance capabilities at ultra-low latency compared to the all-flash cluster configuration. Thus, to handle Oracle OLTP workloads at very low latency, the all-NVMe solution is preferred over the all-flash configuration.

This solution delivers many business benefits, including the following:

- Increased scalability
- Highly available
- Reduced deployment time with a validated reference configuration
- Cost-based optimization
- Data optimization
- Cloud-scale operation

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

- Optimized performance for transactional applications with very low latency
- Balanced and distributed architecture that increases performance and enhances IT efficiency with automated resource management
- Capability to start with a smaller investment and grow as business demands increase
- Enterprise-application-ready solution
- Efficient data storage infrastructure
- Scalability

For more information

For additional information, consult the following resources:

- Cisco HyperFlex white paper: Deliver Hyperconvergence with a Next-Generation Data Platform:
<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/white-paper-c11-736814.pdf>
- Cisco HyperFlex systems solution overview:
<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/solution-overview-c22-736815.pdf>
- Oracle Databases on VMware Best Practices Guide, Version 1.0, May 2016:
<http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutions/vmware-oracle-databases-on-vmware-best-practices-guide.pdf>
- HyperFlex All NVMe at-a-glance: <https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/le-69503-aag-all-nvme.pdf>
- Hyperconvergence for Oracle: Oracle Database and Real Application Clusters:
<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/le-60303-hxsql-aag.pdf>

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 <https://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: <https://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)