



# Oracle Database on VMware vSAN 6.7

## Table of contents

Oracle Database on VMware vSAN 6.7 .....	3
Executive Summary .....	3
Business Case .....	3
Solution Overview .....	3
Key Highlights .....	3
Audience .....	3
Technology Overview .....	4
VMware vSphere .....	4
VMware vSAN .....	4
VMware Cloud on AWS .....	4
Oracle Database .....	5
Samsung NVMe SSD .....	5
Solution Configuration .....	6
Architecture Diagram .....	6
Hardware Resources .....	6
Software Resources .....	7
Network Configuration .....	7
Oracle Database VM and Database Storage Configuration .....	7
Solution Validation .....	9
Test Overview .....	9
Test and Performance Data Collection Tools .....	9
vSAN Configurations Used in this Solution .....	10
Single Oracle VM Workload Test .....	10
Two Oracle VMs' Workload Scalability Test .....	12
vSAN Resiliency and Adaptive Resync .....	14
Recommendations for Running Oracle Database on vSAN .....	16
vSAN All-Flash Configuration Guidelines .....	16
Conclusion .....	17
Reference .....	18
Appendix A SLOB Configuration .....	19
About the Author and Contributors .....	20

## Oracle Database on VMware vSAN 6.7

### Executive Summary

This section covers the business case, solution overview, key highlights, and audience of the Oracle Database 12c on VMware vSAN 6.7.

#### Business Case

Customers deploying Oracle Database have requirements such as stringent SLAs, consistent performance, and high availability. It can be a major challenge for organizations to manage data storage in these environments due to these demanding business requirements. Common issues in using traditional storage solutions for business-critical applications include inability to easily scale-up and scale-out, storage inefficiency, complex management, high deployment, and operating costs.

VMware® vSAN™ has been widely adopted as an Hyperconverged Infrastructure (HCI) solution providing a scalable, resilient, and high-performance storage using cost-effective hardware, specifically direct-attached disks in VMware ESXi™ hosts. vSAN uses storage policy-based management, which simplifies and automates complex management workflows that exist in traditional enterprise storage systems with respect to configuration and clustering. To show the continued improvement in VMware vSAN software, we have developed this reference architecture document to demonstrate the consistent application experience by improved Oracle workload performance, scalability, and resynchronization performance.

#### Solution Overview

This solution addresses the common business challenges that organizations face today in an online transaction processing (OLTP) environment that requires predictable performance. The solution helps customers design and implement optimal configurations specifically for Oracle Database 12c on all-flash vSAN 6.7.

#### Key Highlights

The following points validate that vSAN is an enterprise-class storage solution suitable for running heavy Oracle workloads:

- Predictable Oracle OLTP performance on all-flash vSAN cluster
- Storage Policy Based Management (SPBM) to administer storage resources combined with simple design methodology that eliminates operational and maintenance complexity of traditional SAN.
- Resilient platform for Tier-1 business-critical workloads.
- Validated architecture that reduces implementation and operational risks.

#### Audience

This reference architecture is intended for Oracle Database administrators, virtualization and storage architects involved in planning, architecting, and administering a virtualized Oracle environment with vSAN.

## Technology Overview

This section provides an overview of the technologies used in this solution:

- VMware vSphere®
- VMware vSAN
- VMware Cloud on AWS
- Oracle Database
- Samsung NVMe SSD

### VMware vSphere

VMware vSphere 6.7 is the next-generation infrastructure for next-generation applications. It provides a powerful, flexible, and secure foundation for the business agility that accelerates the digital transformation to cloud computing and promotes the success in the digital economy.

vSphere 6.7 supports both existing and next-generation applications through its:

- Simplified customer experience for automation and management at scale
- Comprehensive built-in security for protecting data, infrastructure, and access
- Universal application platform for running any application anywhere

With vSphere 6.7, customers can run, manage, connect, and secure their applications in a common operating environment, across clouds and devices.

See [VMware vSphere documentation](#) for more information.

### VMware vSAN

VMware's industry leading HCI software stack consists of vSphere for compute virtualization, vSAN, vSphere native storage, and vCenter for virtual infrastructure management. VMware HCI is configurable, and seamlessly integrates with VMware NSX™ to provide secure network virtualization and/or vRealize Suite™ for advanced hybrid cloud management capabilities. HCI can be extended to the public cloud, as VMware powered HCI has native services with two of the top four cloud providers, AWS and IBM.

We are now introducing vSAN 6.7 Update 1, which makes it easy to adopt HCI with simplified operations, efficient infrastructure and rapid support resolution. With vSAN 6.7 Update 1, customers can quickly build and integrate cloud infrastructure. vSAN's automation and intelligence keeps your infrastructure stable, secure and minimizes maintenance disruptions. vSAN 6.7 Update 1 lowers TCO and makes your storage more efficient through automatic capacity reclamation, and it helps avoid overspending on storage by helping users size capacity needs correctly and incrementally. Finally, vSAN Support Insight reduces time-to-resolution while lessening customer involvement in the support process, as well as expediting self-help.

See [VMware vSAN documentation](#) for more information.

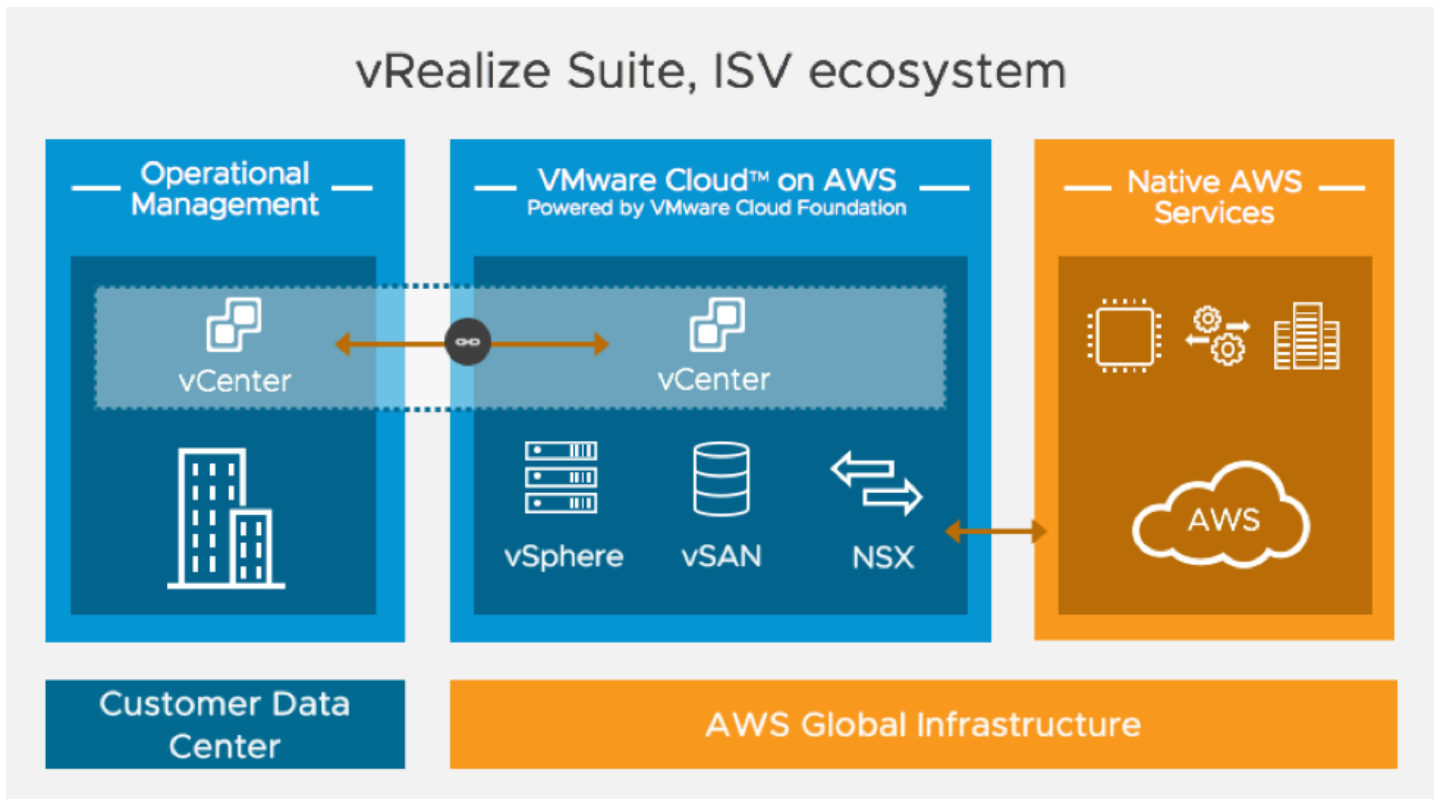
### VMware Cloud on AWS

VMware Cloud on AWS is an on-demand service that enables customers to run applications across vSphere-based cloud environments with access to a broad range of AWS services.

Powered by VMware Cloud Foundation, this service integrates vSphere, vSAN, and NSX along with VMware vCenter management, and is optimized to run on dedicated, elastic, bare-metal AWS infrastructure. ESXi hosts in VMware Cloud on AWS reside in an AWS availability zone (AZ) and are protected by VMware vSphere High Availability (vSphere HA).

With VMware Hybrid Cloud Extension, customers can easily and rapidly perform large-scale bi-directional migrations between on-premises and VMware Cloud on AWS environments.

See [VMware Cloud on AWS documentation](#) for more information.



### Oracle Database

Oracle Database is a relational database management system deployed as a single instance or as RAC (Real Application Clusters), ensuring high availability, scalability, and agility for any application.

Oracle Database provides many new features including multi-tenant architecture that simplifies the process of consolidating databases in the cloud, enabling customers to manage many databases as one without changing their application.

Oracle Database accommodates all system types, from data warehouse systems to update-intensive OLTP systems.

### Samsung NVMe SSD

Samsung is well equipped to offer enterprise environments superb solid-state drives (SSDs) that deliver exceptional performance in multi-thread applications, such as compute and virtualization, relational databases and storage. These high-performing SSDs also deliver outstanding reliability for continual operation regardless of unanticipated power loss. Using their proven expertise and wealth of experience in cutting-edge SSD technology, Samsung memory solutions helps data centers operate continually at the highest performance levels.

This solution uses the 1.6 TB Samsung PM1725 SSD as the cache tier for the vSAN cluster. It delivers consistently high performance, outstanding reliability, and high density.

See [Samsung PM1725a NVMe SSD](#) for more information.

## Solution Configuration

This section introduces the resources and configurations for the solution including:

- Architecture diagram
- Hardware resources
- Software resources
- Network configuration
- Oracle Database VM and database storage configuration

### Architecture Diagram

This solution was using a 4-Node vSAN cluster. The key designs for the vSAN Cluster solution for Oracle Database were:

- A 4-node vSAN Cluster with two vSAN disk groups on each ESXi host.
- Each disk group was created from 1 x 1.6TB NVMe (cache) and 4 x 1.75TB SSDs (capacity).
- VM with 12vCPU with 128GB RAM.
- Oracle Enterprise Linux (OEL) operating system was used for database VMs.

Note: This solution configuration was done using Oracle 12.2.0.1 and VMware vSAN 6.7. The same steps applies when using Oracle databases of higher versions. Any performance data is a result of the combination of hardware configuration, software configuration, test methodology, test tool, and workload profile used in the testing.

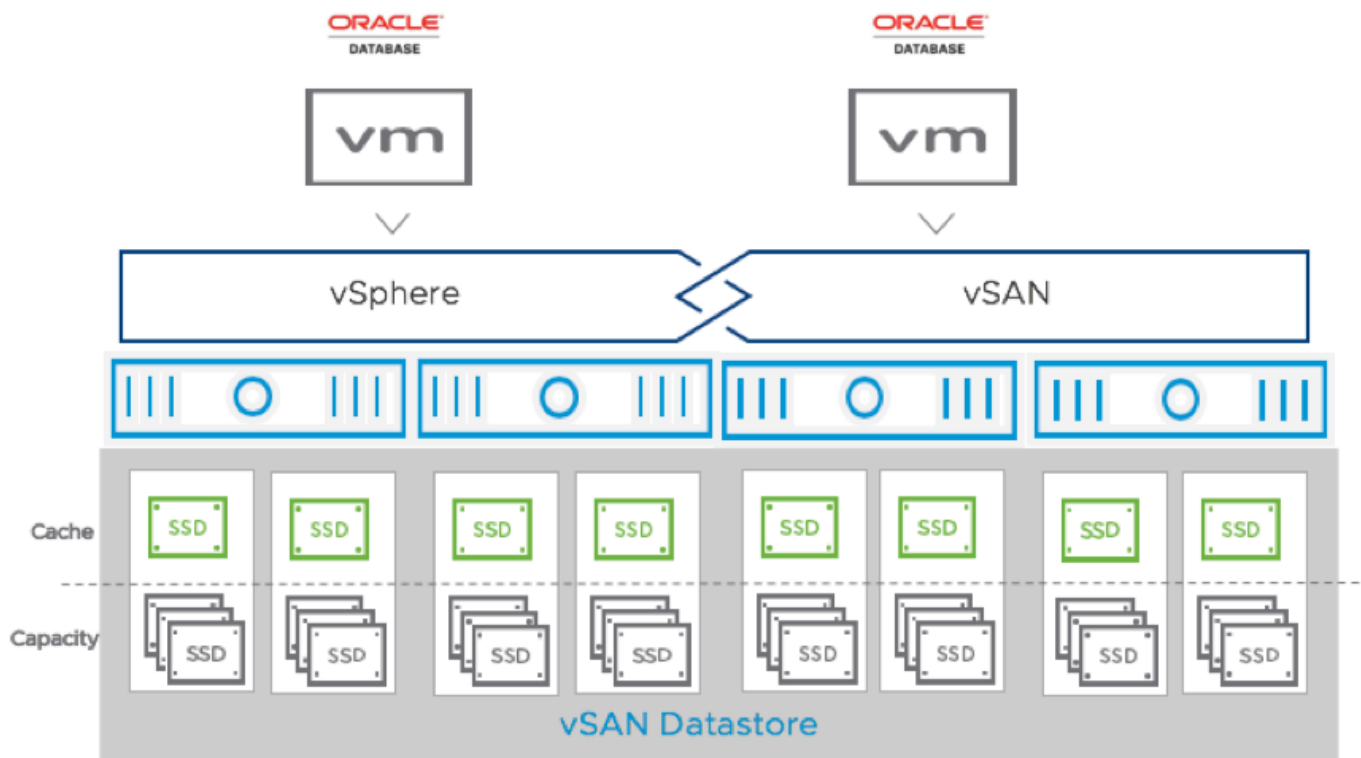


Figure 1. vSAN Cluster for Oracle Database

### Hardware Resources

Table 1 shows the hardware resources used in this solution.

Table 1. Hardware Resources

DESCRIPTION	specification
Server	4 x ESXi Server
Server Model	Dell PowerEdge R640
CPU	2 sockets with 14 core each, Intel Xeon Gold 6132 2.60GHz with hyperthreading enabled
RAM	288 GB
Storage controller	1 x DELL HBA 330 Mini
Disks	Cache—2 x 1.6TB NVMe Samsung PM1725a Capacity—8 x 1.75 TB Toshiba SAS SSD
Network	2 x 10Gb, 1 x 1Gb (Management)

The storage controller used in the reference architecture supports the pass-through mode. The pass-through mode is the preferred mode for vSAN and it gives vSAN complete control of the local SSDs attached to the storage controller.

## Software Resources

Table 2 shows the software resources used in this solution.

**Table 2. Software Resources**

Software	Version	Purpose
VMware vCenter Server® and ESXi	6.7	ESXi cluster to host virtual machines and provide vSAN Cluster. VMware vCenter Server provides a centralized platform for managing VMware vSphere environments.
VMware vSAN	6.7	Software-defined storage solution for Hyperconverged Infrastructure
Oracle Enterprise Linux (OEL)	7.4	Oracle Database server OS
Oracle Database 12c	12.2.0.1	Oracle Database
Oracle Workload Generator for OLTP	SLOB 2.4.2.1	To generate OLTP like workload

## Network Configuration

A VMware vSphere Distributed Switch™ (VDS) acts as a single virtual switch across all associated hosts in the data cluster. This setup allows virtual machines to maintain a consistent network configuration as they migrate across multiple hosts. The vSphere Distributed Switch uses two 10GbE adapters per host. Link Aggregation Control Protocol (LACP) is used to combine and aggregate multiple network connections (2 x 10GbE). When NIC teaming is configured with LACP, the load balancing of the vSAN network occurs across multiple uplinks. However, this happens at the network layer, and is not done through vSAN. The physical network switch is also configured using LACP, so the LAG (link aggregation group) is formed. For details about LAG for vSAN, refer to the [VMware vSAN Network Design](#) guide.

A port group defines properties regarding security, traffic shaping, and NIC teaming. Jumbo frames (MTU=9000 bytes) were enabled on the vSAN interface and the default port group setting was used. Two port groups were created:

- VM management port group for VMs
- vSAN port group for the kernel port used by vSAN traffic

## Oracle Database VM and Database Storage Configuration

Oracle Single Instance 12.2.0.1 Database VM was installed with Oracle Enterprise Linux 7.4 and was configured with 12 vCPU and 128 GB memory.

A large database was configured:

- Oracle ASM data disk group with external redundancy was configured with the default allocation unit size of 1M.
- All ASM Disk groups were presented on different Paravirtual SCSI controllers (PVSCSI).
- Three different ASM disk group was configured, DATA for database data file, system undo and temp files, REDO for online

redo disk and FRA to store archive logs.

Refer to Oracle on VMware best practices in the [Recommendations for Running Oracle Database on vSAN](#) chapter.

Table 3 provides Oracle VM disk layout and ASM disk group configuration.

**Table 3. Oracle Database VM Disk Layout**

Name	SCSI TYPE	SCSI ID (Controller, LUN)	Size (GB)	ASM Disk Group
Operating System (root)	Paravirtual	SCSI (0:0)	50	Not Applicable
Oracle binary disk /u01	Paravirtual	SCSI (0:1)	100	Not Applicable
FRA disk 1	Paravirtual	SCSI (0:2)	750	FRA
FRA disk 2	Paravirtual	SCSI (0:3)	750	FRA
Database data disk 1	Paravirtual	SCSI (1:0)	1024	DATA
Database data disk 2	Paravirtual	SCSI (1:1)	1024	DATA
Database data disk 3	Paravirtual	SCSI (1:2)	1024	DATA
Database data disk 4	Paravirtual	SCSI (1:3)	1024	DATA
Database data disk 5	Paravirtual	SCSI (2:0)	1024	DATA
Database data disk 6	Paravirtual	SCSI (2:1)	1024	DATA
Database data disk 7	Paravirtual	SCSI (2:2)	1024	DATA
Database data disk 8	Paravirtual	SCSI (2:3)	1024	DATA
Online redo disk 1	Paravirtual	SCSI (3:0)	20	REDO
Online redo disk 2	Paravirtual	SCSI (3:1)	20	REDO



## Solution Validation

The solution designed and deployed Oracle Single Instance Database on a vSAN Cluster focusing on ease of use, performance, resiliency, and availability. We present the test methodologies and processes used in this reference architecture.

### Test Overview

The solution validates the performance and functionality of Oracle Database running in a vSAN environment.

The solution tests include:

- Oracle OLTP like workload on a large database
- Using Storage Policy Based Management (SPBM) to provide a mix of RAID 1 mirror and RAID 5 erasure coding policy to Database VM to achieve a balance between vSAN space efficiency and performance
- Workload with Deduplication and compression enabled
- vSAN Adaptive resync in action during host failure

### Test and Performance Data Collection Tools

#### Test Tools and Configuration

##### Oracle OLTP Workload

**SLOB** is an Oracle workload generator designed to stress test storage I/O, specifically for Oracle Database using OLTP workload. We used it to validate performance of the storage subsystem without application contention.

##### SLOB and Database Configuration

- Tests were run on a single database VM and two database VMs
- Each VM was on a separate ESXi host in a 4-node cluster
- 5 TB SLOB tablespace was created and used to load SLOB schema
- TEMP files were created on DATA ASM disk group
- Online redo logs in REDO ASM disk group
- Archive logging was enabled and was located on FRA ASM disk group
- Number of users set to 32 with zero think time to hit each database with the maximum requests concurrently to generate intensive OLTP workload
- Workload is a mix of 70 percent reads and 30 percent writes to mimic a transactional database workload

Detailed SLOB configuration can be found at [Appendix A SLOB Configuration](#).

#### Performance Metrics Data Collection Tools

We measured two important workload metrics in all the tests:

- IO per second (IOPS)
- Average latency of each IO operation (ms)

IOPS and average latency metrics are important for OLTP workload.

We used the following testing and monitoring tools in this solution:

- vSAN Performance Service: [vSAN Performance Service](#) is used to monitor the performance of the vSAN environment, using the web client. The performance service collects and analyzes performance statistics and displays the data in a graphical format. You can use the performance charts to manage your workload and determine the root cause of problems.
- Oracle AWR reports with Automatic Database Diagnostic Monitor (ADDM)

Automatic Workload Repository (AWR) collects, processes, and maintains performance statistics for problem detection and self-tuning purposes for Oracle Database. This tool can generate report for analyzing Oracle performance.

The Automatic Database Diagnostic Monitor (ADDM) analyzes data in AWR to identify potential performance bottlenecks. For each of the identified issues, it locates the root cause and provides recommendations for correcting the problem.

## vSAN Configurations Used in this Solution

Several vSAN feature combinations were used during the tests. Table 4 shows the abbreviations used to represent the feature configurations.

Table 4. Feature Configurations and Abbreviations

Name	RAID Level	vSAN Deduplication and Compression
R1	1	No
[1] R15	Data disks-5 FRA disks-5 OS disks-5 Redo disks-1	No
R1+DC	1	Yes
R15+DC	Data disks-5 FRA disks-5 OS disks-5 Redo disks-1	Yes

Unless otherwise specified in the test, the vSAN Cluster was designed with the following vSAN default policy parameters:

- Failures to Tolerate of 1
- Checksum enabled
- Object Space Reservation (OSR) set to 0 percent
- Stripe width of 1

This is a common practice used in the industry. Log disks are mirrored, and data disks are configured for RAID 5. Oracle data disks, FRA (Archive log) and OS occupy the major storage capacity (usually more than 90 percent). RAID 5 policy is applied for storage efficiency. For online redo log disks, RAID 1 policy is used for performance. This practice provides substantial space savings. Since erasure coding is a storage policy, it can be independently applied to different virtual machine objects providing the simplicity and flexibility for configuring this type of workload.

## Single Oracle VM Workload Test

### Test Overview

This test focused on heavy Oracle OLTP workload on vSAN. SLOB was used to stress Oracle Databases in the vSAN Cluster.

While users can use SLOB to simulate a realistic database workload, we chose to stress the database VM with 32 users without any think time to hit each database with the most intensive database requests. The workload ran for 60 minutes.

We tested different vSAN storage policy configurations as shown in Table 4. We also scaled the number of database VM to two and ran the tests.

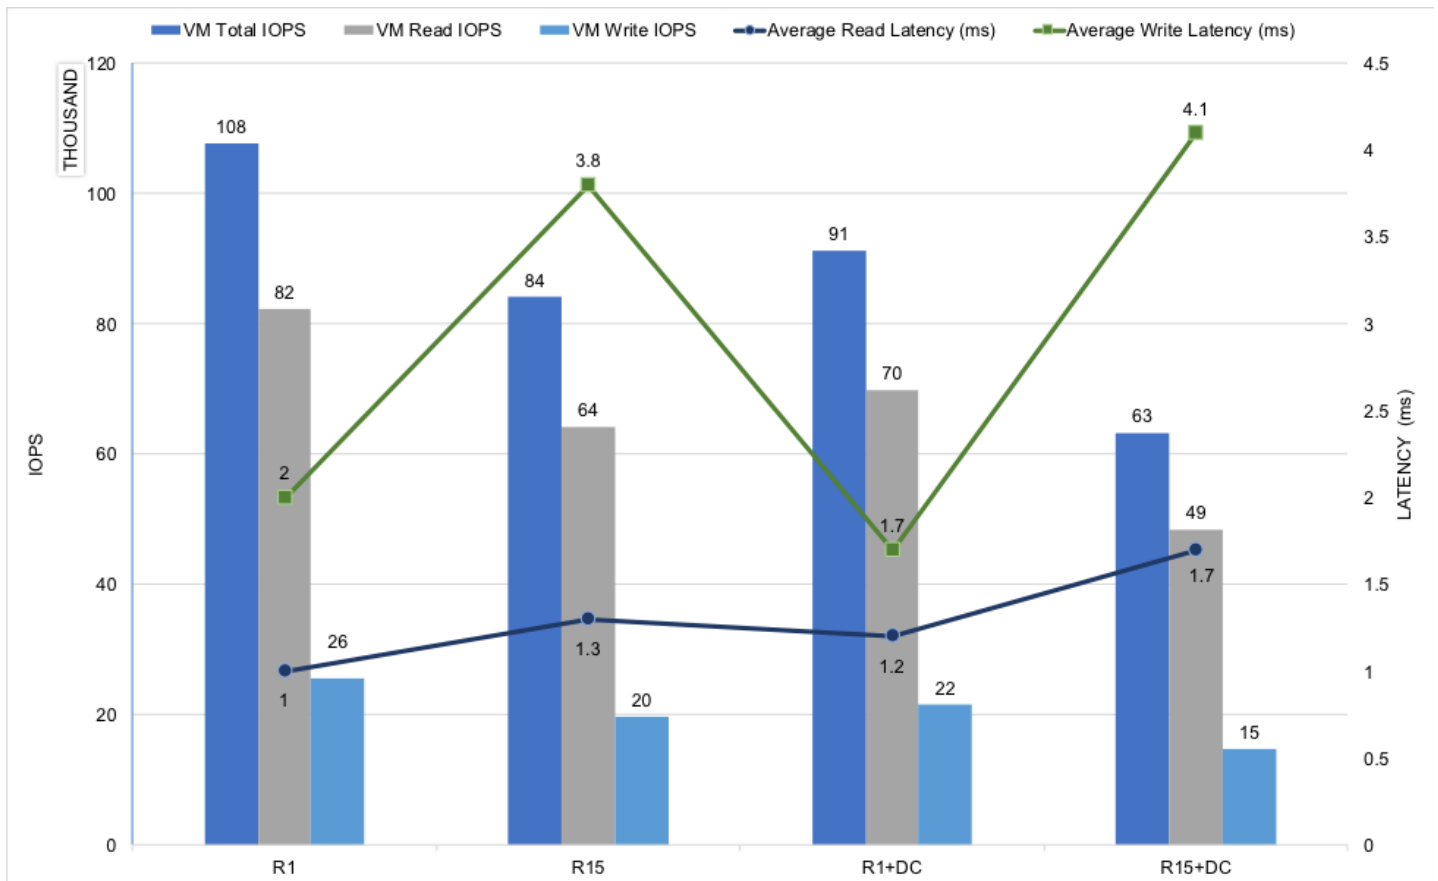
### Test Results and Observations

#### R1 Baseline

We measured the key metrics for the OLTP workload. R1 vSAN configuration was studied as baseline performance for the OLTP tests. Figure 2 shows the IOPS generated by an Oracle Database VMs during the test with different vSAN policies. For the R1 policy the average IOPS observed was of 107,800. This was observed on a single database VM. Notice the workload was a mix of 70 percent read and 30 percent write IOPS, which mimicked a transactional database workload. The read and write latency was stable across the run with 1ms and 2ms respectively.

This shows that vSAN provides reliable performance for a business-critical application such as Oracle Database despite high intensity of the workload and the size of database.

The IOPS observed at the client VM level matched with the physical read and write IO in the Oracle Database AWR reports.



**Figure 2. vSAN IOPS Single Oracle Database VM**

Latency in an OLTP test is a critical metric of how well the workload is running. Lower IO latency reduces the time CPU waits for IO completion and improves application performance.

Latency was relatively low for this solution considering the heavy IOs generated concurrently. Figure 2 shows the average read latency was 1ms and the average write latency was 2ms during this workload scenario, more realistic real-world database environments running in steady state will see much lower latencies.

The average CPU utilization on the ESXi hosting the Oracle Database was less than 45% throughout the workload.

### Compare Baseline Configuration with Other vSAN Configurations

vSAN provides built-in data reduction technologies including erasure coding, deduplication and compression.

To understand the performance impact introduced by these features, we compared the baseline configuration (R1) with the other three vSAN configurations as shown in Table 4.

- In R15 configuration, a mix of vSAN RAID 1 mirroring and erasure coding (RAID 5) policy was used. Redo disk was configured with RAID 1 mirror while the other disks were configured with RAID 5. This provided a balance between performance and cost. Average IOPS reduced from 107,000 to 84,100, which was a reduction of 21 percent. The read and write latency observed was 1.3ms and 3.8ms respectively.
- In R1+DC configuration, vSAN deduplication and compression was enabled. Average IOPS reduced from 107,000 to 91,300 with a reduction of 15 percent. The read and write latency observed was 1.2ms and 1.7ms respectively.
- In R15+DC configuration, erasure coding (RAID 5) feature was also used along with deduplication and compression. In this test, the IOPS observed was 63,300, which was a reduction of 41 percent comparing to that of baseline (R1). The read and write latency observed was 1.7ms and 4.7ms. This configuration will provide best space efficiency possible due to erasure coding and vSAN deduplication and compression.

Figure 2 also shows the IO latency under different vSAN configurations.

In case of latency-sensitive application, recommendation is to use RAID 1 (Mirror) for Data and Redo disks; otherwise, use RAID 5 (erasure coding) for Data and RAID 1 for Redo to provide space efficiency with reasonable tradeoff of performance.

Overall, while erasure coding provides a predictable amount of space savings, deduplication and compression provides a varying

amount of reduction in capacity depending upon the workload and vSAN disk group configuration.

Because the domain for deduplication is at the disk group level, smaller number of large disk groups typically yield higher overall deduplication ratios than larger number of smaller disk groups.

The disadvantage of having smaller number of large disk groups is less write-buffer capacity relative to disk group size and more data migration and resync traffic during maintenance operations (disk replacement, failure).

If database native compression is used, vSAN compression may provide reduced benefits. The space saving obtained due to deduplication and compression is highly dependent on the application workload and data set composition.

From the “[SLOB - The Simple Database I/O Testing Toolkit for Oracle Database Release 2.4.2](#)” guide:

“Due to the SLOB Method, one should not use the default SLOB schema for testing compression technology. Simply put, default SLOB data compresses too deeply to be of any use in assessing compression technology”. Hence the capacity savings reported with SLOB data set is not useful data to derive value.”

Under this workload, we observed an insignificant overhead on ESXi resources (CPU and memory) because of erasure coding, deduplication and compression.

## Summary

Any performance data is a result of the combination of hardware configuration, software configuration, test methodology, test tool, and workload profile used in the testing.

The figures above show various heavy OLTP workload tests with different vSAN configurations. Table 5 summarizes all the test results.

The IOPS and latency data in the table are from vSAN performance. Matching IOPS and latency data was observed from the Linux Operating system iostat command in each database VM.

**Table 5. Summary of OLTP Workload Tests and Key Metrics**

vSAN Configuration	Average Total IOPS	Average Read IOPS	Average Write IOPS	Average READ LATENCY (MS)	Average Write Latency (ms)
R1	107,800	82,300	25,500	1	2
R15	84,100	64,300	19,800	1.3	3.8
R1+DC	91,300	69,800	21,500	1.2	1.7
R15+DC	63,300	48,500	14,800	1.7	4.1

## Two Oracle VMs' Workload Scalability Test

### Test Overview

For an Enterprise storage system with good performance, one of the major requirements is to be able to effectively scale database workloads seamlessly with predictable IOPS and latency.

In this test, two database workloads were run concurrently on vSAN using SLOB with different SPBM settings.

One of the VM's was setup with R1 Storage Policy and the other VM used R15 storage policy. Both workloads were run concurrently for 60 minutes.

### Test Results and Observations

With both database workloads running concurrently, the average IOPS on the vSAN Cluster was 169,000 as shown in Figure 3. This IOPS was generated from two database VMs, 94,200 IOPS from the VM that used the R1 vSAN configuration and 74,800 from the VM in R15 configuration.

Another key metric for OLTP performance is having predictable latency for faster transactions. Figure 4 shows the average IO latency from OLTP VMs. Notice even with the workload from two database VMs, the latency remained was approximately the same as in the case of a single database VM workload.

On the VM using R1 storage policy, the read and write latency observed were 1.9ms and 1.1ms. On the VM using R15 storage policy, the read and write latency observed were 1.4ms and 4.3ms.

As mentioned earlier, any performance data is a result of the combination of hardware configuration, software configuration, test methodology, test tool, and workload profile used in the testing.

In conclusion, in addition to its scalability, vSAN also has the capability to use SPBM for granular control to assign R1 storage policy for production database and use R15 policy storage for other non-production databases like development and testing where space efficiency prioritizes performance.

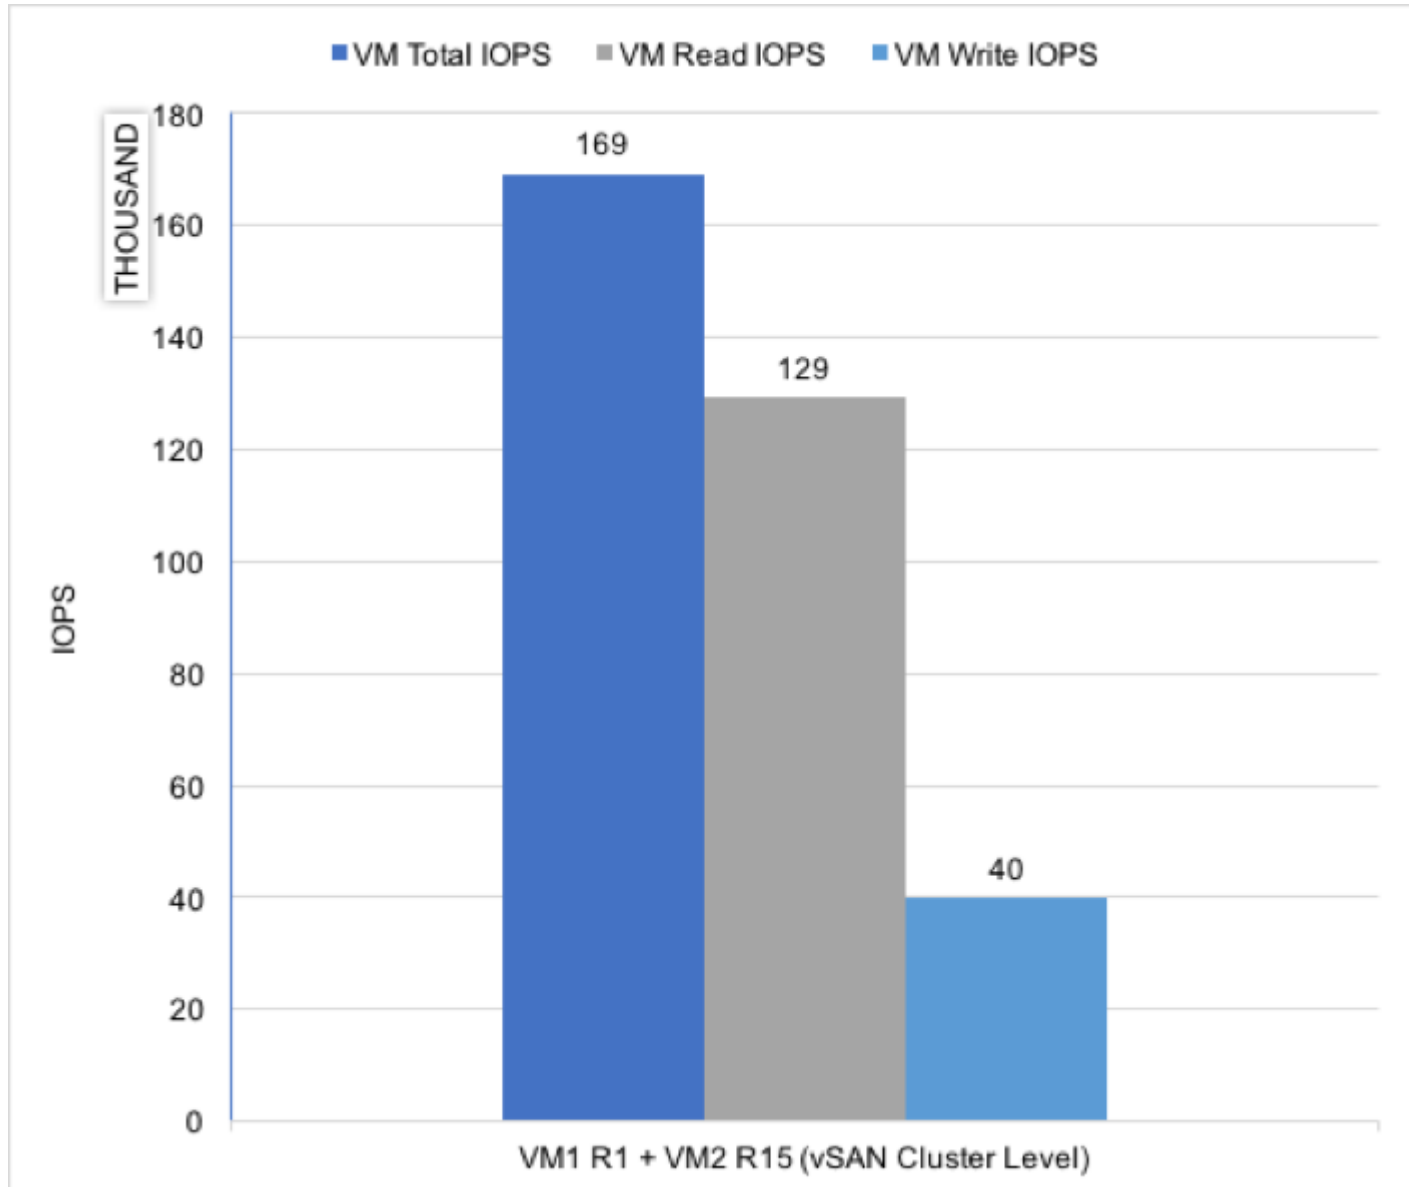


Figure 3. vSAN Cluster Level IOPS from Two Database VM Test

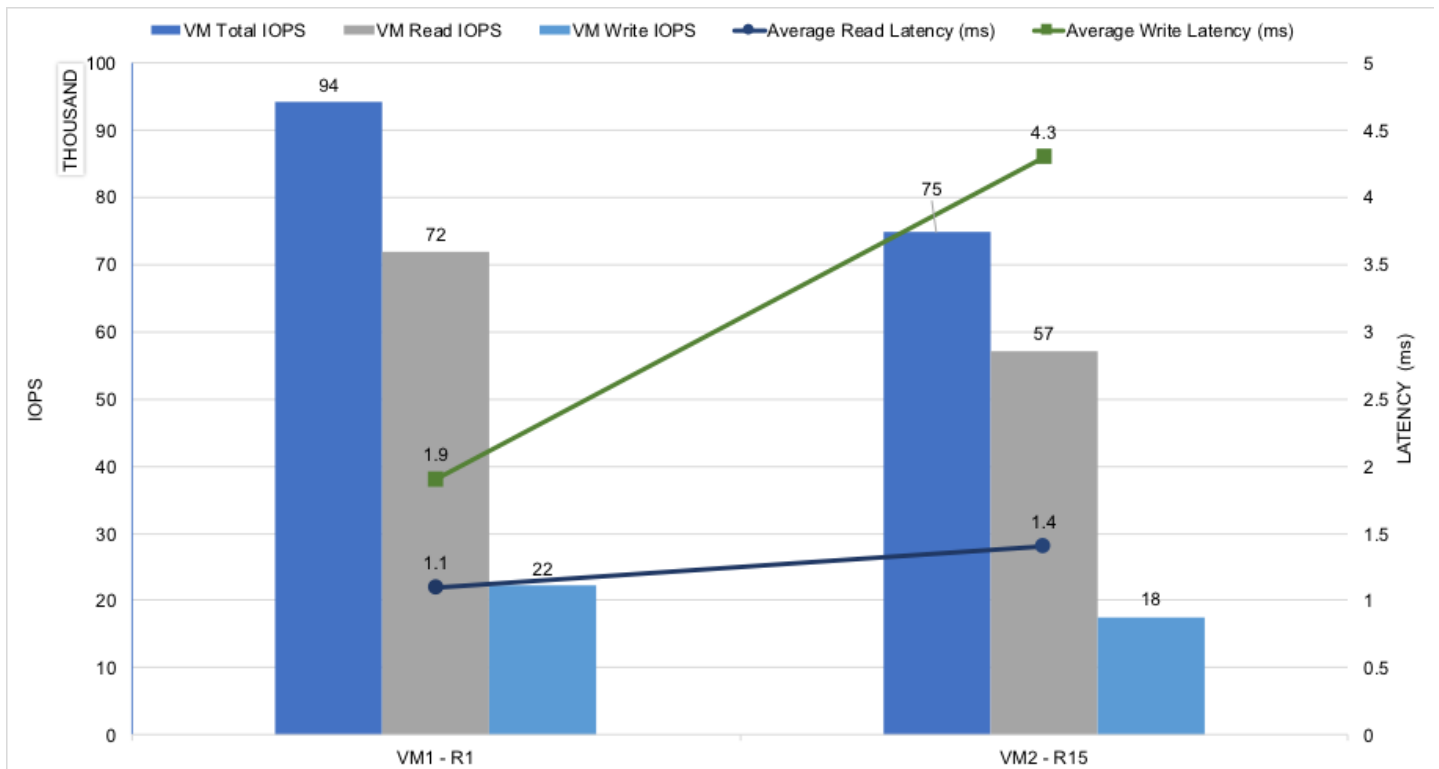


Figure 4. VM Level IOPS and Latency Two Database VM Test

## vSAN Resiliency and Adaptive Resync

### Test Overview

vSAN 6.7 introduces a new feature called “Adaptive Resync” which ensures that fair-share of resources are available to VM I/O and vSAN Resync I/O during the dynamic changes of load on the system. When I/O activity exceeds the capabilities of the bandwidth provided, the Adaptive Resync feature guarantees a level of bandwidth to ensure one type of traffic is not starved for resources.

Adaptive Resync allows more bandwidth for resync operations when there is no contention for resources. If no resync traffic exists, VM I/O may consume 100% of bandwidth, and under contention, Resync I/O is guaranteed at least 20% of the bandwidth. This provides an optimal use of resources.

See the [Adaptive Resync in vSAN 6.7](#) for more details.

This section validates that vSAN can adaptively manage resynchronization operations due to hardware failure, maintenance mode, or policy changes.

We designed the following scenarios to emulate potential real-world failures during the OLTP workload. Two Oracle Database VMs were used during this test. SLOB workload was run on both the database VMs, one using R1 storage policy and other using R15 storage policy. The same heavy SLOB workload configuration was applied.

We tested a “Host failure” scenario as part of this test. While the database workload was running, one of the ESXi server in vSAN cluster was abruptly powered down using Dell iDRAC. The server which was powered down did not host any database VM.

By default, the resync kicks in after 60 minutes. However, in this test the “Repair objects immediately” option in the vSAN health UI was used to start the rebuild immediately. The default repair delay value can be modified.

See the [VMware Knowledge Base Article 2075456](#) for steps to change the repair delay value.

### Test Results and Observations

After the host failure, the SLOB workload continued, there were no IO errors in the Linux VM or Oracle user-session disconnections. As soon as the “Repair object immediately” option is used, the resynchronization operation starts. The graph shown in Figure 5 is a vSAN backend throughput data during an Oracle SLOB workload along with the resync traffic.

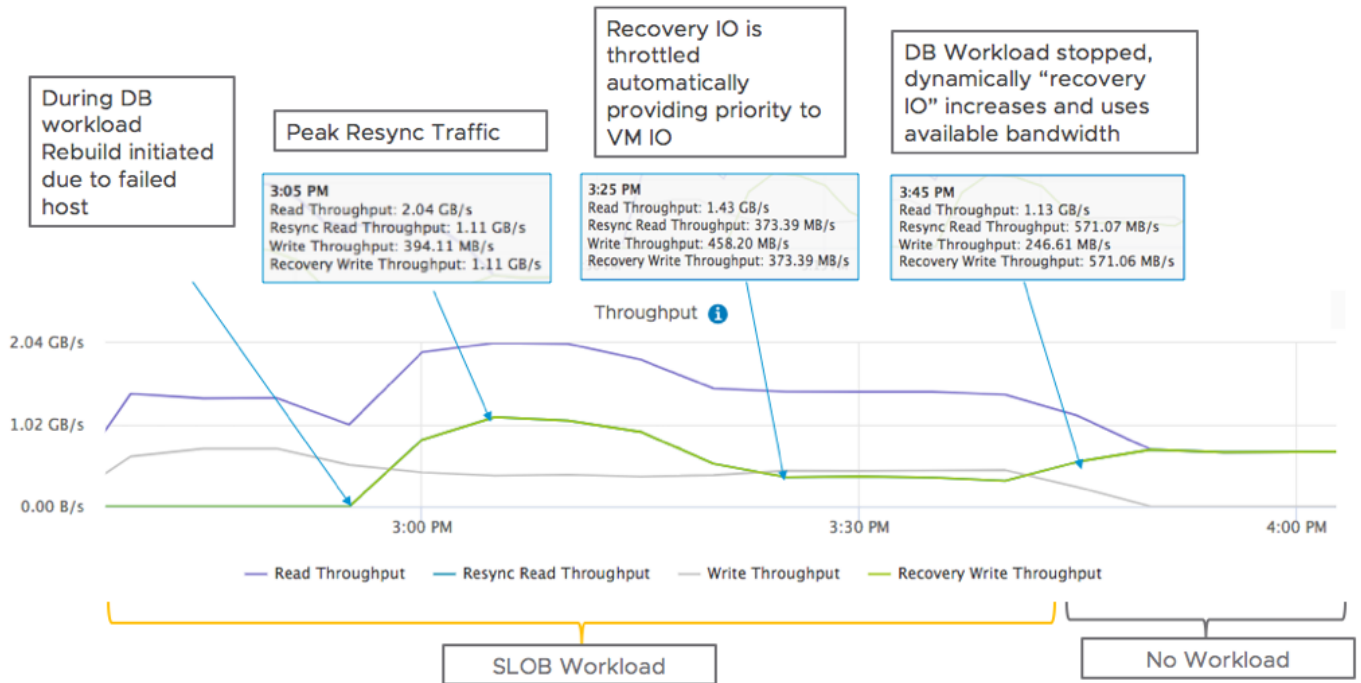
The resync traffic started at 2:55pm. The green line in the graph is the “Recovery Write Throughput” that shows the resynchronization traffic due to the failed host. It shows a gradual increase with the peak resync traffic of 1.11GB/s at 3:05pm.

vSAN bandwidth regulator sensed that the recovery IO throughput was using more than the guaranteed bandwidth and was

impacting VM IO performance. To give priority and fair share of bandwidth to VM IO, the recovery IO traffic was reduced after 3:05pm and was maintained at a steady 20% of the available bandwidth.

When the Oracle workload completed at 3:45pm, vSAN dynamically increased the recovery IO to use the available bandwidth.

This experiment demonstrates how “Adaptive Resync” feature prioritizes guest VM I/O while opportunistically allowing vSAN to use as much bandwidth as possible during any resync activities.



## Recommendations for Running Oracle Database on vSAN

This section highlights the best practices to be followed for Oracle Database 12c on vSAN 6.7.

### vSAN All-Flash Configuration Guidelines

A well designed HCI cluster powered by vSAN is key to a successful implementation of mission-critical Oracle Database. The focus of this reference architecture is vSAN best practices for Oracle Database. For information about setting up Oracle Database on VMware vSphere, refer to the [Oracle Databases on VMware Best Practices Guide](#) along with vSphere Performance Best Practices Guide for specific version of vSphere.

### vSAN All-Flash Configuration Guidelines

[vSAN Design and Sizing Guide](#) provides a comprehensive set of guidelines for designing vSAN. A few key guidelines relevant to Oracle Database are provided below:

- vSAN is a distributed object-store datastore formed from locally attached devices from the ESXi host. It uses disk groups to pool together flash devices as single management constructs. Therefore, it is recommended to use similarly configured and sized ESXi hosts for vSAN Cluster to avoid imbalance. For scale-ups, consider an initial deployment with enough cache tier to accommodate future requirements. For future capacity addition, create disk groups with similar configuration and sizing. This ensures a balance of virtual machine storage components across the cluster of disks and hosts.
- Design for availability. Depending on the failure tolerance method and setting, design with additional host and capacity that enable the cluster to be automatically recovered in the event of a failure and to be able to maintain a desired level of performance.
- vSAN SPBM provides storage policy management at virtual machine object level. Leverage it to turn on specific features like checksum, erasure coding, and QoS for required objects.
- Network: vSAN requires a correctly configured network for virtual machine IO as well as communication amount cluster nodes. With all-flash Cluster and more importantly with high speed NVMe devices, network can become a bottleneck during throughput-intensive workload and during vSAN resynchronization. For network, intensive workloads take advantage of Link Aggregation (LACP) and use larger bandwidth ports like 25Gbps if required. See the [VMware vSAN Network Design guide](#) for details.
- Workloads that are highly sensitive to latency variations should use storage policies with RAID 1 (Mirror) for both data and redo disks. If the goal is to provide the balance between space efficiency and performance, use RAID 5 (erasure coding) for data disk and RAID 1 for redo. RAID 1 (Mirror) or erasure coding can be independently applied to different virtual machine objects using SPBM, which provides simplicity and flexibility to configure database workloads.
- vSAN deduplication and compression can reduce raw storage capacity consumption, and can be used when the application-level compression is not used. The space saving obtained due to deduplication and compression is specific to the application workload and data set composition. Since the domain for deduplication is at the disk group level, smaller number of large disk groups typically yield higher overall deduplication ratios than larger number of smaller disk groups do.



## Conclusion

This section provides a solution summary of running Oracle Database 12c on vSAN 6.7.

vSAN is a cost-effective and high-performance HCI platform that is rapidly deployed, easy to manage, and fully integrated into the industry-leading VMware vSphere platform.

In this reference architecture, we ran heavy OLTP like workload against one and two Oracle Database VMs, and achieved over 108,000 and 169,000 IOPS respectively with low latency.

We also showcased how vSAN SPBM allows granular control for different Oracle Database disks to provide a balance between space efficiency and performance.

We simulated a hardware failure scenario to show how vSAN Adaptive resynchronization prioritizes Oracle Database VM IO while opportunistically allowing vSAN to use as much bandwidth as possible for resync activity.

VMware HCI architecture powered by vSAN is quite capable of running heavy OLTP database workloads for today's most demanding business-critical applications.

## Reference

This section lists the relevant references used for this document.

### White Paper

For additional information, see the following white papers:

- [Oracle Databases on VMware Best Practices Guide](#)
- [Oracle Real Application Clusters on VMware vSAN](#)
- [Oracle Database on VMware vSAN—Day 2 Operations and Management](#)

### Product Documentation

For additional information, see the following product documentation:

- [VMware vSAN Design and Sizing Guide](#)
- [VMware vSAN Network Design](#)
- [Performance Best Practices for VMware vSphere 6.7](#)
- [Oracle Database Online Documentation](#)

### Other Documentation

For additional information, see the following document:

- [SLOB Resources](#)

## Appendix A SLOB Configuration

This section provides informations on the SLOB configuration file we used in our testing.

The following file is the SLOB configuration file we used in our testing:

```

UPDATE_PCT=30
SCAN_PCT=0
RUN_TIME=3600
WORK_LOOP=0
SCALE=128G
SCAN_TABLE_SZ=1M
WORK_UNIT=64
REDO_STRESS=LITE
LOAD_PARALLEL_DEGREE=10

THREADS_PER_SCHEMA=1

DATABASE_STATISTICS_TYPE=awr # Permitted values: [statspack|awr]

#### Settings for SQL*Net connectivity:
#### Uncomment the following if needed:
ADMIN_SQLNET_SERVICE=ora12c
SQLNET_SERVICE_BASE=ora12c
#SQLNET_SERVICE_MAX="if needed, replace with a non-zero integer"
#
#### Note: Admin connections to the instance are, by default, made as SYSTEM
#         with the default password of "manager". If you wish to use another
#         privileged account (as would be the cause with most DBaaS), then
#         change DBA_PRIV_USER and SYSDBA_PASSWD accordingly.
#### Uncomment the following if needed:
DBA_PRIV_USER=sys
SYSDBA_PASSWD=password

#### The EXTERNAL_SCRIPT parameter is used by the external script calling feature of runit.sh.
#### Please see SLOB Documentation at https://kevinclosson.net/slob for more information

EXTERNAL_SCRIPT=''

#####
#### Advanced settings:
#### The following are Hot Spot related parameters.
#### By default Hot Spot functionality is disabled (DO_HOTSPOT=FALSE).

DO_HOTSPOT=FALSE
HOTSPOT_MB=8
HOTSPOT_OFFSET_MB=16
HOTSPOT_FREQUENCY=3

#### The following controls operations on Hot Schema
#### Default Value: 0. Default setting disables Hot Schema

HOT_SCHEMA_FREQUENCY=0

#### The following parameters control think time between SLOB
#### operations (SQL Executions).
#### Setting the frequency to 0 disables think time.

THINK_TM_FREQUENCY=0
THINK_TM_MIN=.1
THINK_TM_MAX=.5

```

The following is the command we used to start SLOB workload with 32 users:

```
"/home/oracle/SLOB/runit.sh 32"
```

## About the Author and Contributors

This section provides a brief background on the author and contributors of this document.

Palanivenkatesan Murugan, Solution Architect, works in the Product Enablement team of the Storage and Availability Business Unit. Palani specializes in solution design and implementation for business-critical applications on VMware vSAN. He has more than 13 years of experience in enterprise storage solution design and implementation for mission-critical workloads. Palani has worked with large system and storage product organizations where he has delivered Storage Availability and Performance Assessments, Complex Data Migrations across storage platforms, Proof of Concept, and Performance Benchmarking.

Sudhir Balasubramanian, Staff Solution Architect, works in the Cloud Platform Business Unit. Sudhir specializes in the virtualization of Oracle business-critical applications. Sudhir has more than 20 years' experience in IT infrastructure and database, working as the Principal Oracle DBA and Architect for large enterprises focusing on Oracle, EMC storage, and Unix/Linux technologies. Sudhir holds a Master Degree in Computer Science from San Diego State University. Sudhir is one of the authors of the "Virtualize Oracle Business Critical Databases" book, which is a comprehensive authority for Oracle DBAs on the subject of Oracle and Linux on vSphere. Sudhir is a VMware vExpert Ex-Member of the CTO Ambassador Program and an Oracle ACE.

Catherine Xu, Senior Technical Writer in the Product Enablement team, edited this paper to ensure that the contents conform to the VMware writing style.

