



Mixed Workloads on VMware vSAN All-Flash

Table of contents

Mixed Workloads on VMware vSAN All-Flash	3
Executive Summary	3
Business Case	3
Solution Overview	3
Key Results	4
Audience	4
Technology Overview	5
VMware vSphere 6.7 U1	5
VMware vSAN 6.7 U1	5
Microsoft SQL Server 2016	5
Microsoft Exchange Server 2016	5
Samsung NVMe SSD	5
Solution Configuration	7
Overview	7
Architecture Diagram	7
Hardware Resources	8
Software Resources	8
Network Configuration	9
vSAN Storage Policy Configuration	9
Mixed-Workload Virtual Machine Configuration	9
Mixed Workload Database Configuration	10
Solution Validation	12
Test Overview	12
Testing Tools	12
Mixed Workload Performance Test	13
Mixed Workload Failure Test	17
Mixed Workload Backup Test	21
Best Practices	23
Conclusion	24
Summary	24
Reference	25
About the Author	26

Mixed Workloads on VMware vSAN All-Flash

Executive Summary

This section covers the business case, solution overview, key results, and audience of mixed workloads on VMware vSAN.

Business Case

The traditional approach to building out compute and storage infrastructure for multiple enterprise workloads has been dedicated resources based on the applications that run in a separated environment, with a hesitance on business-critical applications sharing resources with others. Infrastructure administrators are more likely to plan new workload deployments on dedicated hardware so that the original workloads are not impacted, which ensures that both performance and availability service level agreements (SLAs) are met. While this is appropriate in some cases, it is more complex and often unnecessary. Dedicated infrastructure hardware for multiple environments and applications can lead to a higher cost and lower resource utilization.

Hyperconverged infrastructure (HCI) makes it easier to plan for multiple enterprise workloads consolidation within several clusters. It helps combine hardware silos, reduce deployment complexity and management difficulty, and save capital and operational expenditures. VMware vSAN™ is VMware’s premier storage solution for HCI, which provides the broadest set of HCI deployment choices for enterprise mixedworkload deployment and provisioning. VMware vSAN is also an ideal platform to migrate your current mixed-workload environment from traditional infrastructure with the least performance and operational compromise for mission-critical applications.

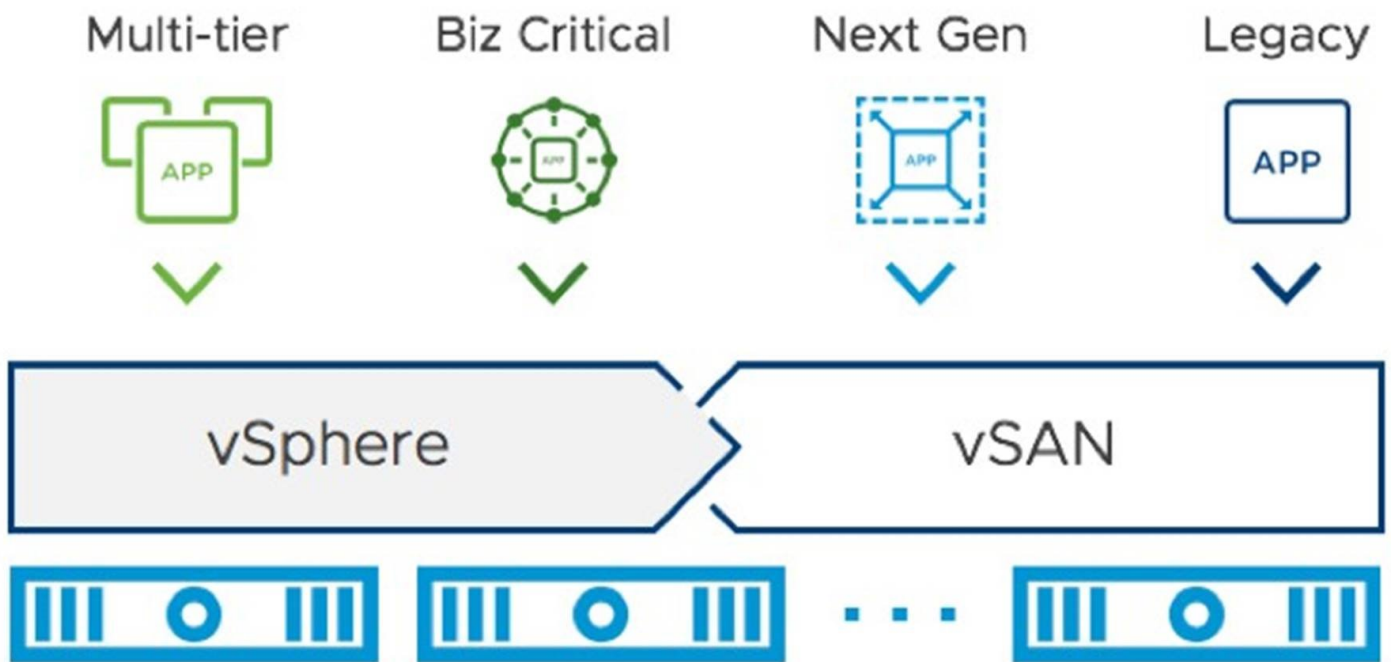


Figure 1. Enterprise Mixed Workloads Running on VMware vSAN

In this solution, we provide design guidance and best practices for enterprise infrastructure administrators and application owners to run mixed workloads on VMware vSAN platform. We combined the two workloads together as an example: Microsoft SQL Server—the most popular relational database platform and Exchange Server—the common mailbox system. We explore the performance capability and consistency, workload resiliency and how the two application workloads perform in daily operations.

Solution Overview

This reference architecture is a showcase of using VMware vSAN all-flash as HCI for operating and managing mixed workloads in a VMware vSphere® environment:

- We measure the performance impact and consistency of running both SQL Server and Exchange together within a single vSAN cluster.
- We demonstrate the vSAN resiliency against different kind of failures, which maintains business continuity and prevents data loss for the mixed-workload environment.

- We perform daily backup operations for SQL Server and Exchange.

Key Results

This reference architecture:

- Designs an architecture of deploying SQL Server and Exchange in a vSAN cluster.
- Showcases predictable performance of running SQL Server and Exchange together within a vSAN cluster.
- Demonstrates high availability features of running SQL Server AlwaysOn Availability Group (AAG) and Exchange Database Availability Group (DAG) on vSAN with minimal performance impact.
- Performs backup and restore guidance for mixed workloads on vSAN.
- Demonstrates various vSAN fault tolerance offering for business continuity of mixed workloads.
- Provides best practices and guidance.

Audience

This solution is intended for IT administrator, SQL Server DBAs, Exchange mailbox administrators, virtualization and storage architects involved in planning, architecting, and administering virtualized enterprise mixed workloads on VMware vSAN 6.7.

Technology Overview

This chapter introduces the technology components as bulleted list:

- VMware vSphere 6.7 U1
- VMware vSAN 6.7 U1
- Microsoft SQL Server 2016
- Microsoft Exchange Server 2016
- Samsung NVMe SSD

VMware vSphere 6.7 U1

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

VMware vSAN 6.7 U1

VMware vSAN is the industry-leading software powering VMware's software defined storage and HCI solution. vSAN helps customers evolve their data center without risk, control IT costs and scale to tomorrow's business needs. vSAN, native to the market-leading hypervisor, delivers flash-optimized, secure storage for all of your critical vSphere workloads. vSAN is built on industry-standard x86 servers and components that help lower TCO in comparison to traditional storage. It delivers the agility to easily scale IT and offers the industry's first native HCI encryption.

vSAN 6.7 U1 simplifies day-1 and day-2 operations, and customers can quickly deploy and extend cloud infrastructure and minimize maintenance disruptions.

Secondly, vSAN 6.7 U1 lowers the total cost of ownership with more efficient infrastructures. vSAN 6.7 U1 automatically reclaims capacity, using less storage at the capacity tier for popular workloads.

In addition, [vSAN ReadyCare](#) rapidly resolves support requests. vSAN ReadyCare is a marketing name introduced to capture the significant investments VMware has made to support vSAN customers. VMware continues to invest in ReadyCare support, and new ReadyCare simplifies support request resolution and expedites diagnosis of issues.

Microsoft SQL Server 2016

Microsoft SQL Server is the database management system with innovative security and compliance features, industry-leading performance, mission-critical availability and advanced analytics to data workloads, and support for big data built-in. Gartner has already rated SQL Server as having the most complete vision of any operational database management system. With SQL Server 2016, performance is enhanced with a few new technologies, including new features and enhancements regarding enterprise-grade performance, security, availability, and scalability.

Microsoft Exchange Server 2016

Microsoft Exchange Server is an enterprise-class email, calendar, collaboration and data management platform in Windows Server environment. Exchange Server helps organizations to efficiently manage large amounts of email content, files, and documents while improving collaboration, inbox management, scheduling, security, and compliance. With Exchange Server 2016, IT administrators can work more efficiently due to a simplified architecture, enhanced security and compliance tools, data loss and recovery features, and much more.

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. Samsung has the added advantage of being the sole manufacturer of all its SSD components, ensuring end-to-end integration, quality assurance, and the utmost compatibility.

The Samsung PM1725a NVMe SSD delivers:

- Extreme performance: The highest levels with unsurpassed random read speeds and an ultralow latency rate using Samsung's highly innovative 3D vertical-NAND (V-NAND) flash memory and an optimized controller.
- Outstanding reliability: Features five DWPDs (drive writes per day) for five years, which translates to writing a total of 32 TB per day during that time. This means users can write 6,400 files of 5 GB-equivalent data or video every day, which represents a level of reliability that is more than sufficient for enterprise storage systems that have to perform ultrafast transmission of large amounts of data.
- High capacities: Depending on your storage requirements and applications, 800 GB, 1.6 TB, 3.2 TB and 6.4 TB capacities are available.

This solution chooses the 1.6 TB Samsung PM1725a SSD as the cache tier for the vSAN cluster.

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

Solution Configuration

This section introduces the resources and configurations:

- Architecture diagram
- Hardware resources
- Software resources
- Network configuration
- Mixed workload virtual machine
- Mixed workload database

Overview

In this solution, we created a 4-node vSAN cluster for mixed workload validation. We deployed Microsoft SQL Server 2016 and Exchange Server 2016 together in the vSAN cluster, and generated a mixed environment with a brokerage system and mailbox service system and two domain controllers, as described in Figure 2.

The workload of the brokerage system is powered by TPC-E like benchmark, simulating a traditional OLTP database profile on vSAN. We deployed two 600 GB databases for the standalone configuration, and configured two groups of AAG replicas in high availability validation.

The workload of the mailbox service system is simulated by [Microsoft Exchange Server Jetstress 2013 Tool](#), with a mailbox profile for insert, delete, replace, and read operations. We configured total 16,000 1GB size mailboxes of average 1.0 IOPS per mailbox. A background Exchange database maintenance job is also simulated throughout the validations.

The mission-critical database system and enterprise mailbox service system consist of a mixed-workload environment in terms of different user profiles and block size deployed on vSAN.

Architecture Diagram

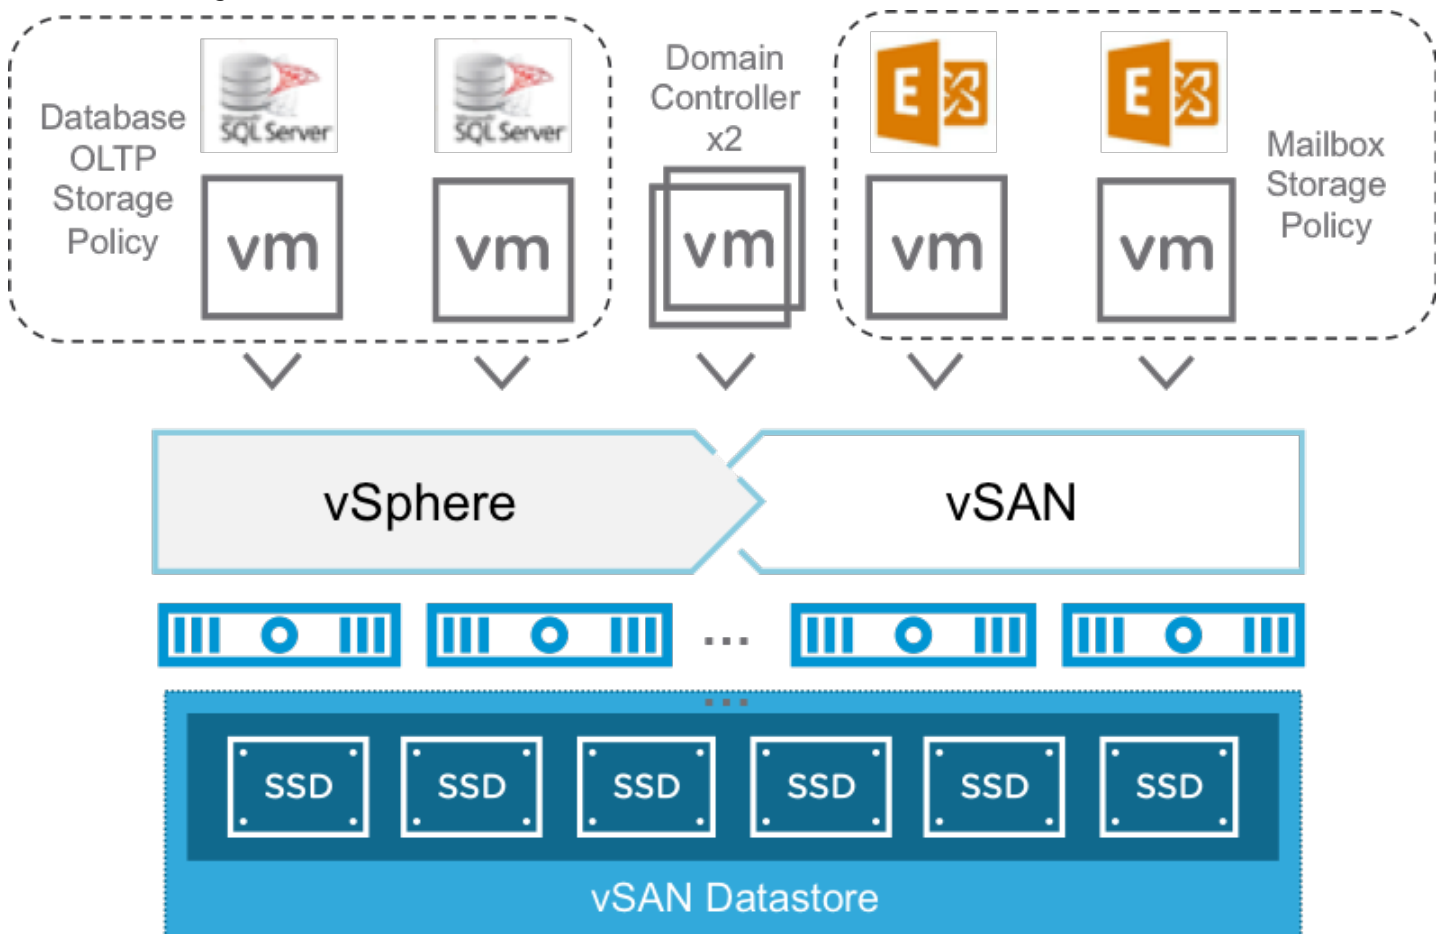


Figure 2. Mixed Workloads on vSAN All-Flash Architecture**Hardware Resources**

In this solution, we used four DELL PowerEdge R640 servers, a 1U platform for density, performance and scalability, and with optimized application performance. Each VMware ESXi™ host contains two disk groups, and each disk group consists of one cache-tier NVMe SSD and four capacity-tier SAS SSDs. We configured the pass-through mode for the capacity-tier storage controller, which is a preferred mode for vSAN with complete control of the local SSDs attached to the storage controller.

Each ESXi server in the vSAN cluster has the following configuration:

Table 1. Hardware Configuration for vSAN Cluster

PROPERTY	SPECIFICATION
Server model name	4 x DELL PowerEdge R640
CPU	2 x Intel(R) Xeon(R) Gold 6132 CPU @ 2.60GHz, 14 core each
RAM	288GB
Network adapter	2 x Intel(R) Ethernet Controller 10G X550T port 2 x Intel Corporation I350 Gigabit Network Connection port
Storage adapter	1 x Dell HBA330 Mini 2 x NVMe SSD Controller 172Xa/172Xb
Disks	Cache—2 x 1.6TB NVMe Samsung PM1725a Capacity—8 x 1.92TB Samsung PM1633a SAS SSD

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 U1 (vSAN 6.7 U1 is included)	vSphere 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 U1	Software-defined storage solution for hyperconverged infrastructure
Windows Server 2016	Datacenter	Operating system
Microsoft SQL Server 2016	SP2, 13.0.5026.0	Database server platform
Benchmark Factory for Databases	8.0.1	SQL Server OLTP workload generate tool
Microsoft Exchange Server 2016	15.01.1591.008	Mailbox server platform
Microsoft Exchange Jetstress 2013	15.01.1019.000	Exchange workload generate tool

Network Configuration

We created a VMware vSphere Distributed Switch™ to act as a single virtual switch across all associated hosts in the data cluster. The vSphere Distributed Switch uses two 10GbE adapters for the teaming and failover. A port group defines properties regarding security, traffic shaping, and NIC teaming. To isolate vSAN, virtual machine, and VMware vSphere vMotion® traffic, we used the default port group settings except for the uplink failover order. We assigned one dedicated NIC as the active link and assigned another NIC as the standby link. vSAN and vSphere vMotion are on separate VLANs and the uplink order is reversed. See Table 3.

Table 3. Network Configuration

DISTRIBUTED PORT GROUP	ACTIVE UPLINK	STANDBY UPLINK
VMware vSAN	Uplink2	Uplink1
Virtual machine and vSphere vMotion	Uplink1	Uplink2

vSAN Storage Policy Configuration

In this solution, we created separate storage policy for SQL Server and Exchange virtual machines. The detailed configuration is defined in table 4.

Table 4. vSAN Storage Policy Configuration for Mixed Workloads

Settings	Value	Description
Failure to Tolerate	1	Defines the number of disks, host or fault domain failures a storage object can tolerate.
Erasur Coding	RAID 1 (Mirroring)	Defines the method used to tolerate failures. By default, Exchange database will preserve two copies on vSAN as storage level protection.
Number of disk stripes per object	1	The number of capacity disks across which each replica of a storage object is striped.
Checksum	Enabled	Checksum is calculated by default to prevent from Exchange data corruption.

Mixed-Workload Virtual Machine Configuration

We configured two SQL Server virtual machines and two Exchange virtual machines in this solution. Table 5 described the virtual machine configuration details. We populated the SQL Server TPC-E like database using the [Benchmark Factory for Databases](#) tool, and Exchange database using the Jetstress tool. We set VMware PVSCSI controller for each of the database virtual disks. All the virtual disks were configured with thin provisioning. We provisioned one primary domain controller and one backup domain controller for active directory services for both applications.

Table 5. Virtual Machine Configuration

VM Role	vCPU	memory (GB)	vm Count	Virtual Disks	SCSI ID (Controller, LUN)	SCSI Type
SQL Server VM	32	128	2	OS disk: 40 GB x 1	SCSI (0, 0)	LSI Logic
				SQL Server data disk 1 ~ 4: 250 GB x 4	SCSI (1, 0) ~ SCSI (1, 3)	VMware Paravirtual
				SQL Server tempdb and log disk: 250 GB x 1	SCSI (2, 0)	VMware Paravirtual
				SQL Server AAG data disk 5 ~ 8: 250 GB x 4	SCSI (3, 0) ~ SCSI (3, 3)	VMware Paravirtual
				SQL Server AAG log disk: 250 GB x 1	SCSI (2, 1)	VMware Paravirtual
				Backup disk : 1TB x 1	SCSI (2, 2)	VMware Paravirtual
Exchange VM	8	64	2	OS disk: 40 GB x 1	SCSI (0, 0)	LSI Logic
				Exchange data disk 1~8: 1 TB x 8	SCSI (1, 0) ~ SCSI (1, 7)	VMware Paravirtual
Domain Controller VM	2	8	2	OS disk: 40 GB x 1	SCSI (0, 0)	LSI Logic

Mixed Workload Database Configuration

We provisioned two 600 GB OLTP databases with the Benchmark Factory tool on each SQL Server instance, and eight 800 GB Exchange databases on each Exchange server VM. The total vSAN datastore capacity utilization was around 65 percent including all the mixed-workload virtual machines and domain controllers’ virtual machines.

Table 6 shows the SQL Server and Exchange database configuration.

Table 6. SQL Server and Exchange Database Configuration

Application	Item	Test configuration
SQL Server	SQL Server memory allocation	120 GB
	Database size	1 x 600 GB per instance
	Workload profile	Industry brokerage system, TPC-E like OLTP workload, 90/10 read/write ratio, 8KB majority block size. Simulated 50 virtual users/connections.
	AAG configuration	1 primary AAG database, 1 secondary AAG database per instance
Exchange Server	Exchange ESE.dll file version	15.01.1591.010
	Database Size	8 x 800 GB databases
	Workload profile	Jetstress workload. 70/30 read/write ratio, 32 KB typical block size. Simulated 16 thread count. <ul style="list-style-type: none"> • 16,000 user mailboxes • 1GB mailbox size • 1.0 IOPS per user mailbox (1,200 messages per day per mailbox) • 24x7 background database maintenance job enabled
	DAG configuration	Mailbox resiliency with 2 database copies, active/passive configuration

Solution Validation

In this section, we present the test methodologies and results to validate this solution.

Test Overview

The solution validates both performance and resiliency for running SQL Server and Exchange mixed workloads on vSAN. The backup testing is also included as supplement for administrators managing day-2 operations for mixed workloads on vSAN.

The test scenarios are as follows:

- Mixed-workload performance test
 - Baseline test 1: SQL Server standalone only test
 - Baseline test 2: Exchange standalone only test
 - Mixed-workload test 1: SQL Server standalone + Exchange standalone test
 - Baseline test 3: SQL Server AAG only test
 - Baseline test 4: Exchange DAG only test
 - Mixed-workload test 2: SQL Server AAG + Exchange DAG test
- Mixed-workload failure test
 - Disk failure test (capacity disk)
 - Disk group failure test (cache disk)
 - Host failure test
- Backup and restoring test
 - Backup test with no workload
 - Backup test with mixed workloads

Testing Tools

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

- Monitoring tools

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.

vSAN Health Check

[vSAN Health Check](#) delivers a simplified troubleshooting and monitoring experience of all things related to vSAN. Through the web client, it offers multiple health checks specifically for vSAN including cluster, hardware compatibility, data, limits, physical disks. It is used to check the vSAN health before the mixed-workload environment deployment.

VMware vRealize® Operations™

vSphere and vSAN 6.7 and later releases include [vRealize Operations within vCenter](#). This new feature allows vSphere customers to see a subset of intelligence offered up by vRealize Operations through a single vCenter user interface. Light-weight purpose-built dashboards are included for both vSphere and vSAN. It is easy to deploy, provides multi-cluster visibility, and does not require any additional licensing. Figure 3 shows the vRealize Operations Manager console.

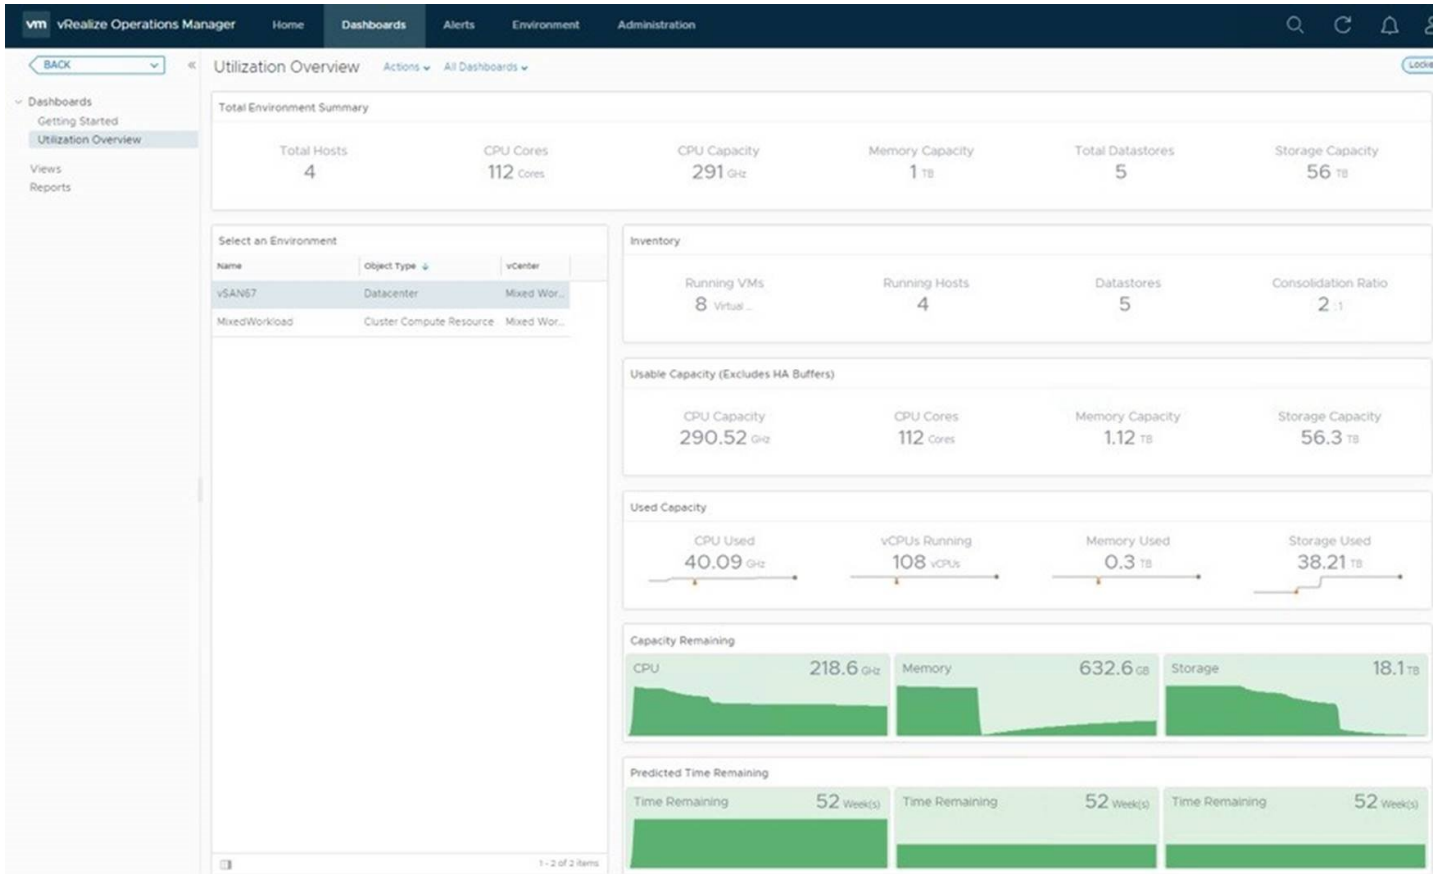


Figure 4. vRealize Operations Manager Console

ESXTOP

ESXTOP is a command line tool that can be used to collect data and provide real-time information about the resource usage of a vSphere environment such as CPU, disk, memory, and network usage. We measure the ESXi Server performance by this tool.

Windows Performance Monitor

Windows Performance Monitor is a Windows tool that enables users to capture statistics about CPU, memory, and disk utilization from operating system levels. It also provides counters for monitoring SQL Server and Exchange performance and status.

- Mixed-workload generation tool

Benchmark Factory for Databases

Benchmark Factory for Databases is a database workload generation tool that can conduct industry-standard benchmark testing and scalability testing. Benchmark Factory for Databases is a database workload generation tool that can conduct industry-standard benchmark testing and scalability testing. With Benchmark Factory for databases, you can make changes to your database environment, and mitigate the risks of planned database changes. We used this tool to generate SQL Server database OLTP workload.

Jetstress 2013

Jetstress is a tool for simulating Exchange database I/O load without requiring Exchange to be installed. It is primarily used to validate physical deployments against the theoretical design targets that were derived during the design phase.

To simulate the complex Exchange database I/O pattern effectively, Jetstress makes use of the same ESE.DLL that Exchange uses in production. Therefore, it is vital that Jetstress uses the same version of the Extensible Storage Engine (ESE) files that your Exchange infrastructure is built with in production.

Mixed Workload Performance Test

Test Objective

This test was designed to demonstrate the minimal impact and consistent performance of running SQL Server and Exchange mixed workloads within a single all-flash vSAN cluster.

Test Scenario

We designed two test scenarios as follows:

- Standalone performance test: Mixed workload with SQL Server and Exchange standalone configuration
- HA performance test: Mixed workload with SQL Server AAG Group and Exchange DAG configuration.

Test Procedures

The testing procedures were described as follows:

1. Perform SQL Server standalone TPC-E like test with two virtual machines, and get SQL Server standalone baseline result. User workload set for each SQL instance is 50.
2. Perform Exchange standalone Jetstress test with two virtual machines, and get Exchange standalone baseline result. The thread count set for each Exchange server is 16.
3. Mixed-workload test: Perform the test in step 1 and 2 together, measure the performance impact of tests conducted in step 1, 2, and 3.
4. Reconfigure AAG on two SQL Server VMs (each VM hosts one primary AAG and one secondary AAG, respectively). Perform the same test as described in step 1 and get SQL Server AAG baseline result.
5. Perform Exchange DAG (2 copies) Jetstress test with two virtual machines, and get Exchange DAG baseline result. The thread count set for Exchange server is 16.
6. Mixed-workload test: Perform the test in step 4 and 5 together, measure the performance impact of tests conducted in step 4, 5, and 6.

Test Results

Figure 4 shows the mixed workload test result of SQL Server and Exchange standalone performance combined on single vSAN cluster. The baseline performance for two SQL Server VMs was 2,152.37 and 2,141.92 TPS respectively; The baseline performance for two Exchange Server VMs was 16,519.88 and 14,226.59 transactional IOPS respectively.

After we mixed the two workloads on vSAN with the same amount of virtual users and threads, SQL Server only got less than 4 percent impact, while keeping the transaction response time at the same level. On the Exchange side, though there was 12 percent transactional IOPS decrease observed, it still supported the Exchange profile defined in the baseline test (8,000 mailboxes per VM, 1.0 IOPS per mailbox).



Figure 4. Mixed Workload Test Result of SQL Server and Exchange Standalone Performance

Table 7 and Table 8 show the detailed test result metrics.

Table 7. Performance Test Result Metrics—SQL Server

SQL Server Metrics	Baseline Test	Mixed Workload Test	Impact
Transaction per second (TPS)	4294.29	4130.7	-3.81%
Transaction response time (ms)	23	24	+1
VM achieved IOPS	12801.8	11821.83	-7.65%
Disk read average latency (ms)	0.404	0.557	+0.153
Disk write average latency (ms)	1.275	1.842	+0.576
CPU utilization	87.59%	86.87%	-0.72%

Table 8. Performance Test Result Metrics—Exchange

Exchange Metrics	Baseline Test	Mixed Workload Test	Impact
Transactional I/O per second	30746.47	27034.52	-12.07%
Database read average latency	0.557	0.636	+0.079
Database write average latency	2.766	2.891	+0.125
Log write average latency	1.304	1.369	+0.065

Figure 5 shows the mixed workload test result of SQL Server AAG and Exchange DAG performance combined on single vSAN cluster. The baseline performance for two SQL Server AAG VMs was 2,084.66 and 2,031.03 TPS respectively; The baseline performance for two Exchange Server VMs was 16,450.05 and 14,042.68 transactional IOPS respectively.

After we mixed the two workloads on vSAN with the same amount of virtual users and threads, the performance impact was negligible. The TPS drop was only 1.54 percent, and the transaction response time was kept at the same level of 24 milliseconds. On the Exchange side, the result was about the same as standalone tests. Though 14 percent transactional IOPS decrease was observed, it still supported the Exchange profile defined in the baseline test (8,000 mailboxes per VM, 1.0 IOPS per mailbox).

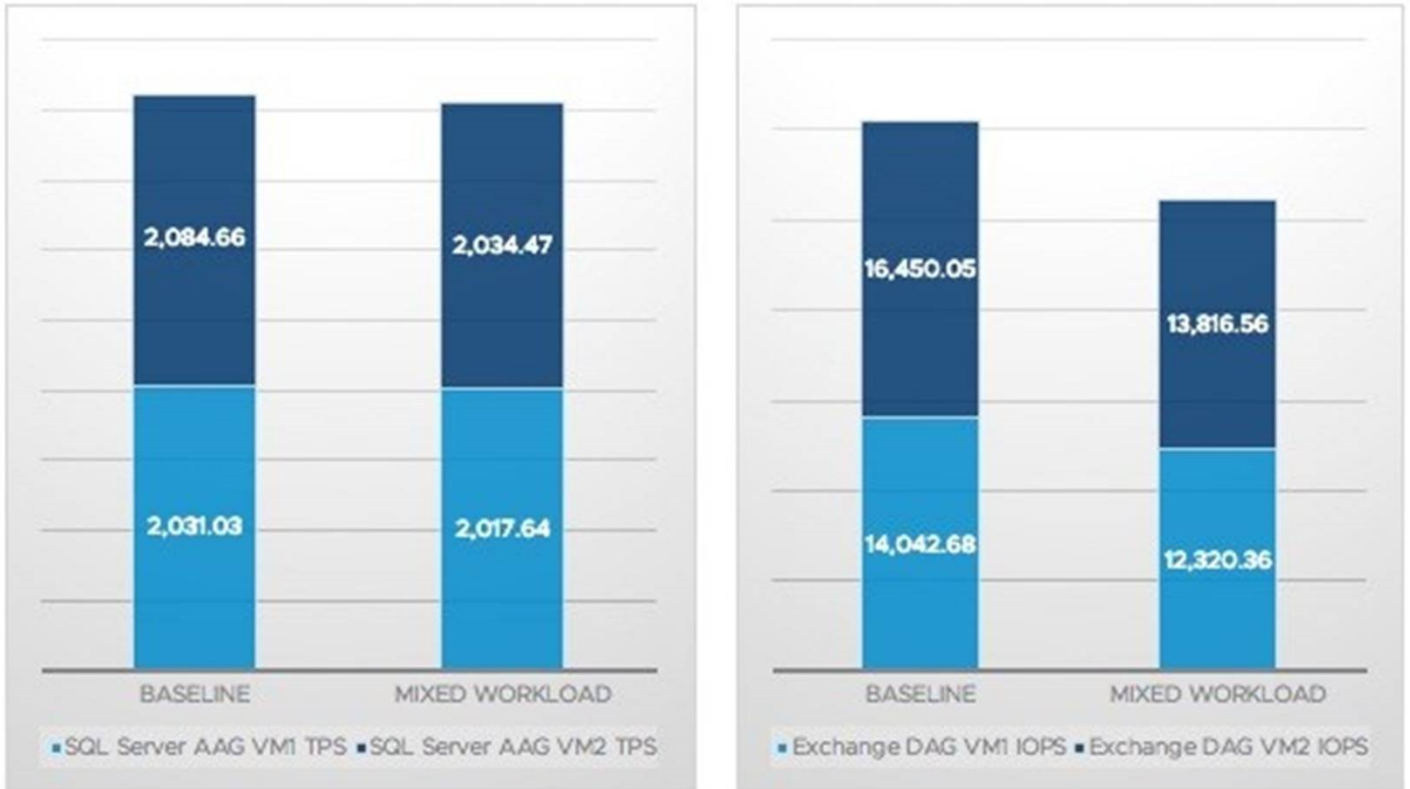


Figure 5. Mixed Workload Test Result of SQL Server AAG and Exchange DAG Performance

Table 9 and Table 10 show the detailed test result metrics.

Table 9. Performance Test Result Metrics—SQL Server AAG

SQL Server AAG Metrics	Baseline Test	Mixed Workload Test	Impact
Transaction per second (TPS)	4115.69	4052.11	-1.54%
VM achieved IOPS (primary + secondary)	15538.86	14459.6	-6.95%
Transaction response time (ms)	24	24	0
Disk read average latency (ms)	0.487	0.531	+0.044
Disk write average latency (ms)	1.765	2.818	+1.053
CPU utilization	88.93%	88.05%	-0.88%

Table 10. Performance Test Result Metrics—Exchange DAG

Exchange DAG Metrics	Baseline Test	Mixed Workload Test	Impact
Transactional I/O per second	30492.74	26136.91	-14.28%
AAG log replication I/O	74.66	64.75	-13.27%
Database read average latency	0.606	0.647	+0.041
Database write average latency	2.846	2.937	+0.091
Log write average latency	1.335	1.424	+0.089

Figure 6 shows the virtual machine performance after running mixed workload combining SQL Server and Exchange. The result was similar for both standalone and AAG/DAG configuration. The first 15 minutes was the pre-sampling period. The SQL Server memory was filled up quickly as new transaction executed, so there was a read performance spike at the very beginning of the test.

The following 45 minutes was the sampling period. The mixed workload performance observed for both SQL Server and Exchange was stabilized. The NVMe all-flash vSAN was capable of handling the two workloads and providing consistent performance expectations.

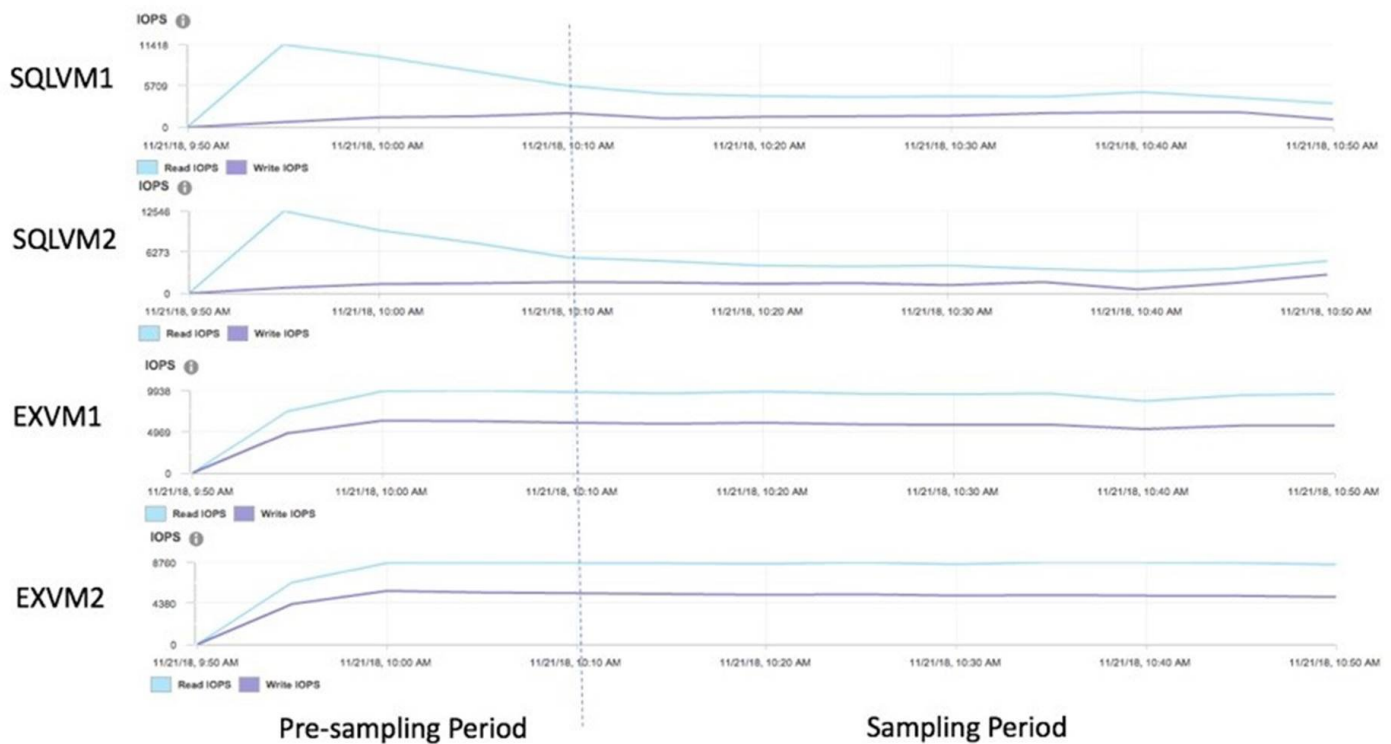


Figure 6. Virtual Machine Performance after Running Mixed Workloads Combining SQL Server and Exchange

In conclusion, the mixed workload performance test demonstrated:

- vSAN could support sustained performance for SQL Server and Exchange mixed workloads both with standalone and with application high availability protection configuration. The performance impact for combined workload of SQL Server and Exchange was minimized, and the mixed workload performance was consistent and stabilized.
- Through all the tests, vSAN all-flash configuration with Samsung NVMe cache-tier could maintain sub-millisecond read latency and less than 3 ms write latency for heavy mixed workloads.

Mixed Workload Failure Test

Test Objective

This test was designed to demonstrate vSAN resiliency features that ensure business continuity of mixed workloads running on vSAN.

Test Scenario

We designed the following test scenarios:

- Disk failure: Evaluate how vSAN deals with disk failure to ensure sustainability of mixed workloads
- Disk group failure: Evaluate vSAN deals with disk group failure to ensure sustainability of mixed workloads.
- Host failure: Evaluate how vSAN deals with host failure to ensure sustainability of mixed workloads.

Test Procedures

The test procedures were described as follows:

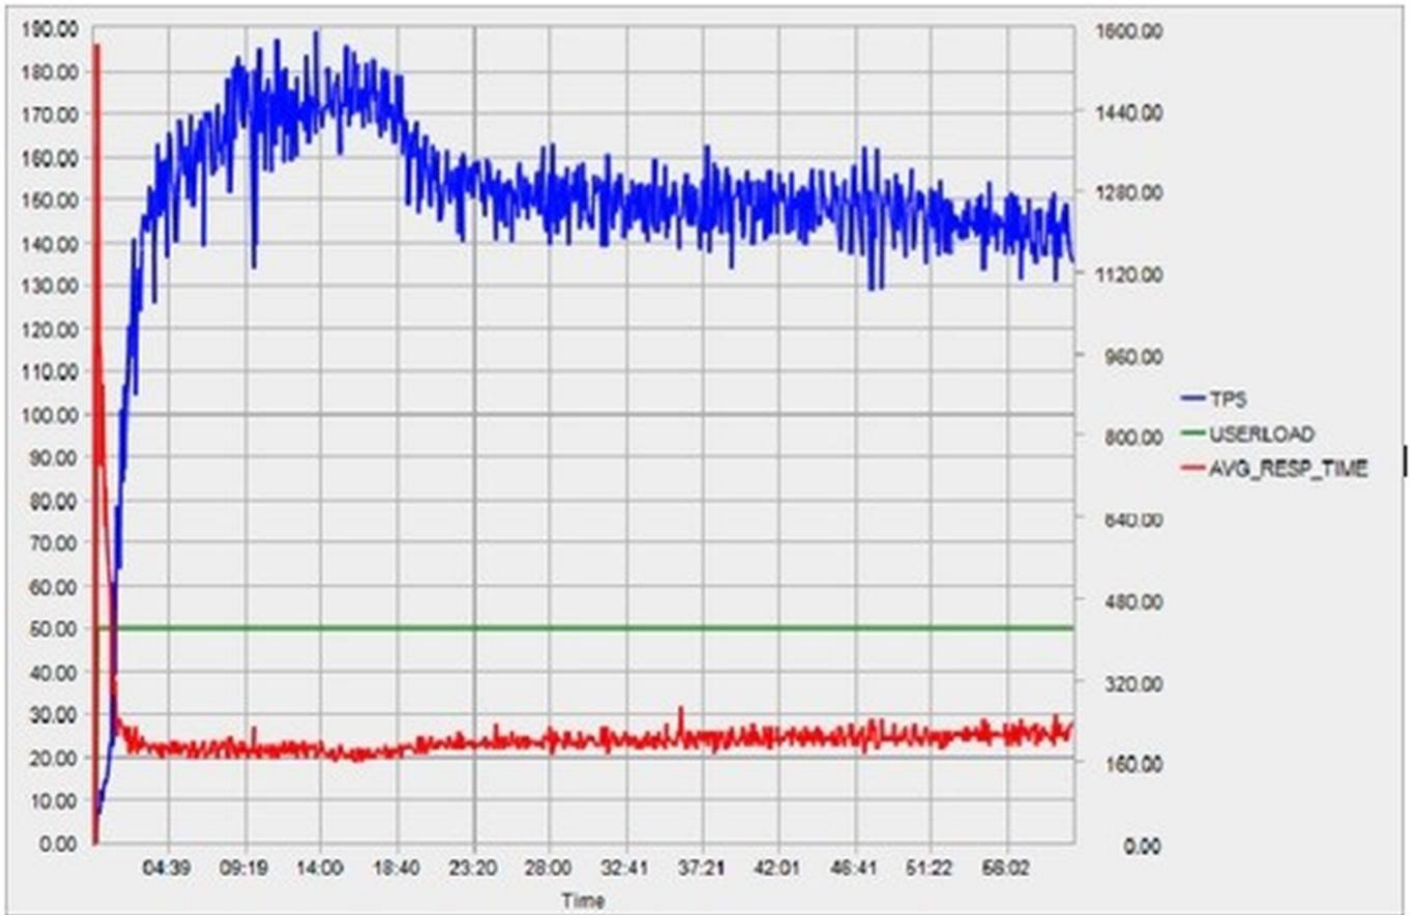
1. Perform SQL Server and Exchange mixed workload test as conducted before. Manually inject a disk error on a capacity SSD. Collect and measure the performance before and after the disk failure.
2. Repeat the test in step 1, instead inject a disk error on a cache SSD which will cause an entire disk group failure. Collect and measure the performance before and after the disk group failure.
3. Repeat the test in step 1, instead force shut down a host in the vSAN cluster. Collect and measure the performance before and after the host failure.

Test Results

The result overview was summarized as follows:

- In all of the three failure scenarios, vSAN could maintain the mixed workload continuity and consistency. The performance impact in the failure tests was subject to certain data spread across the vSAN capacity disks and hosts.
- In the host failure scenario, as vSAN maintained a 60 minutes interval to trigger the resync operation. The “temporary” missing object was marked as absent, so no obvious performance impact was observed during this testing period. Notice the rebuild process will be triggered if the host failure is not recovered in 60 minutes by default. If the host comes back shortly, vSAN also automatically triggers the rebuild process if there is data inconsistency between these objects.
- In the disk and disk group failure, vSAN immediately resynced the components hosting on the failure disks. We showcased the disk group failure testing result as a reference.

Figure 7 shows the SQL Server behavior in the disk group failure test. After the cache disk error injected at the time around 18:40, the TPS result from Benchmark Factory dropped immediately by 10 percent, and the transaction response time increased from about 22 ms to 24 ms. A resync operation was triggered at the same time to handle the failure, with a total amount of 5.4 TB resync data and estimated time 107 minutes as shown in Figure 7. vSAN adaptive resync feature also helps ensure a smooth SQL Server TPS result after the failure, at the cost of minimum 20 percent of available bandwidth.



Object repair timer: 60 minutes ⓘ

Resyncing Objects	30
Bytes left to resync	5.30 TB
ETA to compliance	107 minutes
Scheduled resyncing ⓘ	None

Figure 7. SQL
Figure 7. SQL Server Behavior in the Disk Group Failure Test

Figure 8 shows the Exchange result after the cache disk failure. The result of Exchange was similar to that of SQL Server: decrease of the transactional IOPS and increase of the database disk response time. Overall no obvious fluctuation of Exchange performance was observed after the failure.




Figure 8. Exchange Result after the Cache Disk Failure

Mixed Workload Backup Test

Test Objective

This test was designed to demonstrate the mixed workload sustainability and stability during daily backup operations.

Test Scenario

We designed the following test scenarios:

- Backup with no performance workload on vSAN: Simulate a situation where administrators perform backup operations during no mixed workload running in the environment.
- Backup with performance workload on vSAN: Simulate a situation where administrators perform backup operations during mixed workloads running in the environment.

Test Procedures

The test procedures were described as follows:

1. Perform SQL Server native backup and Exchange backup with Jetstress for no mixed workload enabled. Collect and measure the backup bandwidth.
2. Start SQL Server AAG and Exchange DAG mixed workload test as conducted in step 1. Perform SQL Server native backup from AAG secondary replicas. Collect and measure the backup bandwidth and performance impact.

Test Results

The backup performance with no mixed workload enabled was summarized in Table 11. The mixed workload full backup job was completed with 52 minutes and 39 seconds for a total backup dataset of 2.84 TB. The average peak backup bandwidth was close to 1.2 GB/s. And the vSAN backend peak bandwidth was up to 3.41 GB/s including read write operations on vSAN, as shown in Figure 9.

Table 11. Backup Performance Result with No Mixed Workload Enabled

Virtual Machine	Actual Backup Dataset (GB)	Throughput (MB/s)	Elapse Time
SQL VM1	600GB	297.68	28min29sec
SQL VM2	600GB	357.59	23min41sec
Exchange VM1	820GB	265.5	52min39sec
Exchange VM2	820GB	279.17	51min17sec
Total	2840GB	1199.94	52min39sec




Figure 9. vSAN Backend Performance

Table 12 shows the backup performance with mixed workload enabled. With SQL Server AAG configuration, we chose the backup setting to “Prefer Secondary”, which means the backup operation was redirected to the AAG secondary replica while the primary database was servicing the production OLTP workload.

The overall SQL Server backup bandwidth result was about 579 MB/s, and 3841.62 TPS for SQL Server OLTP performance and 22,375.34 transactional IOPS for Exchange mailbox. This result was still consistent and comparable to the mixed workload performance result described in Table 7 and Table 8, as the backup bandwidth traded off to certain performance impacts for mixed workloads.

Table 12. Backup Performance Result with Mixed Workload Enabled

SQL Server Virtual Machine	Transaction per second (TPS), Primary database	Transaction Response Time (ms)	Backup Performance (MB/s) Secondary	Elapse Time
SQL VM1	1927.34	25	290.49	29min01sec
SQL VM2	1914.28	26	288.17	29min05sec
Exchange Virtual Machine	Transactional IOPS	Read Response Time (ms)	Write Response Time (ms)	Log Latency (ms)
Exchange VM1	11545.59	0.758	2.841	1.461
Exchange VM2	10829.75	0.767	2.804	1.525

Best Practices

This section provides the recommended best practices for this solution.

The following best practices and guidance are recommended for running Microsoft SQL Server and Exchange Server mixed workloads on VMware vSAN:

- Choose to consolidate compute/storage intensive workloads and capacity-oriented workloads on vSAN to improve resource utilization and overall TCO. For example, SQL Server and Exchange demonstrated in this solution. The generic best practices for SQL Server and Exchange are also applied to this solution:
 - For SQL Server best practices on vSAN, visit [here](#).
 - For Exchange Server best practices on vSAN, visit [here](#).
- SPBM considerations:
 - For each type of mixed workload, create a dedicated vSAN storage policy for management and isolation purpose.
 - Use “IOPS limit for object” to limit certain mixed-workload impacts against another.
 - Consider different “Failures to tolerate” options with RAID 5 or RAID 6 to improve space efficiency and protection SLAs.
 - Consider increasing “Number of disk stripes per object” for better object spread if it is a bandwidth demanding workload.
- Refer to [Troubleshooting vSAN Performance](#) as additional resources for performance optimization on vSAN.

Conclusion

Summary

VMware vSAN is an ideal HCI platform to service the majority of user workloads. By running mixed workloads on vSAN, IT administrators can easily achieve better scalability for performance, reduced impact for workload consistency, improved resiliency for data protection, and optimized expenditure for total TCO.

In this solution, we validated mixed workloads running on vSAN with Microsoft SQL Server 2016 and Exchange Server 2016, which demonstrated the feasibility and flexibility of consolidating enterprise mission-critical application with minimized deployment and operational overhead.

We validated a mixture of SQL Server OLTP database performance with TPC-E like workload, and Exchange mailbox database performance with Jetstress workload. The test result showed a minimized performance impact for both workloads running on vSAN, and the application workloads were consistent and stabilized.

We also demonstrated that vSAN offers resiliency and business continuity for mixed workloads against disk failure and host failure situations.

Furthermore, we verified a maximum bandwidth and minimal production impact for mixed workloads in day-2 operations such as backup scenarios.

Reference

For additional information, see the following deliverables:

- [vRealize Operations within vCenter](#)
- [vSAN ReadyCare](#)
- [Troubleshooting vSAN Performance](#)
- [Benchmark Factory for Database](#)
- Microsoft Exchange Server Jetstress 2013 Tool
- [Samsung PM1725a NVMe SSD](#)

About the Author

Mark Xu, Senior Solutions Engineer in the Product Enablement team of the Storage and Availability Business Unit, wrote the original version of this paper.

Pete Koehler, Senior Technical Marketing Architect in the Storage Product Marketing team of the Storage and Availability Business Unit, contributed to the paper contents.

Catherine Xu, Senior Technical Writer in the Product Enablement team of the Storage and Availability Business Unit, edited this paper to ensure that the contents conform to the VMware writing style.

