Microsoft SQL Server 2014 on VMware VSAN 6 Hybrid

First Published On: 07-20-2016
Last Updated On: 07-20-2016

Table of Contents

1. Executive Summary

Microsoft SQL Server 2014 offers an array of new and improved capabilities with emphasis on reliability, availability, serviceability, and performance, all of which are flexible and scalable.

1.1 Business Case

Microsoft SQL Server 2014 offers an array of new and improved capabilities with emphasis on reliability, availability, serviceability, and performance, all of which are flexible and scalable.

It is critical for administrators to effectively deploy storage solutions with Microsoft SQL Server database. With more and more production servers being virtualized, the demand for highly converged server-based storage is surging. VMware® vSAN™ aims at providing highly scalable, available, reliable, and high performance storage using cost-effective hardware, specifically direct-attached disks in VMware ESXi™ hosts. vSAN adheres to a new policy-based storage management paradigm, which simplifies and automates complex management workflows that exist in enterprise storage systems with respect to configuration and clustering.

Solution Overview

This solution addresses the common business challenges that CIOs face today in an online transaction processing (OLTP) environment that requires availability, reliability, scalability, predictability and cost-effective storage, which helps customers design and implement optimal configura tions specifically for Microsoft SQL Server on vSAN.

Key Results

The following highlights validate that vSAN is an enterprise-class storage solution suitable for Microsoft SQL Server:

- Predictable and highly available SQL Server OLTP performance on vSAN.
- Simple design methodology that eliminate soperational and maintenance complexity of traditional SAN.
- Susta inable solution for enterprise Tier-1 Database Management System (DBMS) application platform.
- Validated architecture that reduce simplementation and operational risks.
- Integrated technologies to provide unparalleled availability, business continuity (BC), and disaster recovery (DR).

2. VSAN SQL Server AlwaysOn Reference Architecture

This reference architecture validates vSAN's ability to support industry-standard TPC-E like workloads in a SQL Server AlwaysOn Availability Groups environment.

2.1 Purpose

This reference architecture validates vSAN's ability to support industry-standard TPC-E like workloads in a SQL Server AlwaysOn Availability Groups environment.

AlwaysOn Availability Groups on vSAN ensures a desired level of storage performance, while providing high availability for the underlying data protection mechanism.

This paper also demonstrates that by integrating with proven VMware technologies including VMware vSphere® Data Protection™ and VMware Site Recovery Manager™ combining with VMware vSphere Replication™, vSAN offers a streamlined business continuity and disaster recovery solution with the vSAN Stretched Cluster deployment.

2.2 Scope

This reference architecture:

- Demonstrates storage performance scalability and resiliency of enterprise-class SQL Server AlwaysOn Availability Groups in a vSAN environment.
- Shows vSAN resiliency in various failure scenarios including drive, disk group, and storage host failures.
- Illustrates the benefits of uninterrupted business-critical application production environment through the combination of synchronous protection across geographical stretched clustering, database backup, and 15-minute recovery point objective (RPO) remote-site disaster recovery.

2.3 Audience

This reference architecture is intended for SQL Server 2014 database administrators and storage architects involved in planning, architecting, or administering a SQL Server environment with vSAN.

2.4 Terminology

This paper in cludes the following terminologies as shown in Table 1.

Table 1. Terminology

TERM	DEFINITION
Failure to tolerate (FTT)	FTT, one of the storage policies, defines the replica for VMware Virtual Machine Disk Format (VMDK). The maximum value is 3.

Storage Policy Base Management (SPBM)	The foundation of the Software Defined Storage Control Panel and enables vSphere administrators to overcome upfront storage provisioning challenges, such as capacity planning, differentiated service levels, and managing capacity headroom.
Stripe width	This is one policy setting for vSAN SPBM. The disk object can be striped across multiple devices in the vSAN capacity layer. The maximum stripe width per object is 12.
AlwaysOn Availability Groups	A high availability(HA) and disaster recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, Always On vailability Groups maximize the availability of a set of user databases.
Availability replica	An instantiation of an availability group that is hosted by a specific instance of SQL Server and maintains a local copy of each availability database that belongs to the availability group. Two types of availability replicas exist: one primary replica and one to eight secondary replicas.
Data synchronization	Process by which changes to a primary database are reproduced on a secondary database.

Primary replica	The availability replica that makes the primary databases available for read and write connections from clients and sends transaction log records for each primary database to every secondary replica.
Readable secondary replica	Secondary replica databases that are configured to allow read-only client connections.
Secondary replica	An availability replica that maintains a secondary copy of each availability database, and serves as a potential failover targets for the availability group. Optionally, a secondary replica can support read-only access to secondary databases and can support creating backups on secondary databases.

3. Technology Overview

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

3.1 Technology Overview

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

- VMware vSphere
- VMware vSAN
- VMware vSAN Stretched Cluster
- VMware vSphere Data Protection™
- VMware Site Recovery Manager
- Microsoft SQL Server 2014
- Microsoft SQL Server AlwaysOn Availability Groups

3.2 VMware vSphere

VMware vSphere is the industry-leading virtualization platform for building cloud infrastructures. It enables users to run business-critical applications with confidence and respond quickly to business needs. vSphere accelerates the shift to cloud computing for existing data centers and underpins compatible public cloud offerings, forming the foundation for the industry's best hybrid-cloud model.

3.3 VMware vSAN

VMware vSAN is VMware's software-defined storage solution for hyper-converged infrastructure, a software-driven architecture that delivers tightly integrated computing, networking, and shared storage from x86 servers. vSAN delivers high performance, highly resilient shared storage by clustering server-attached flash devices and hard disks (HDDs).

vSAN delivers enterprise-class storage services for virtualized production environments along with predictable scalability and all-flash performance, all at a fraction of the price of traditional, purpose-built storage arrays. Just like vSphere, vSAN provides users the flexibility and control to choose from a wide range of hardware options and easily deploy and manage it for a variety of IT workloads and use cases.

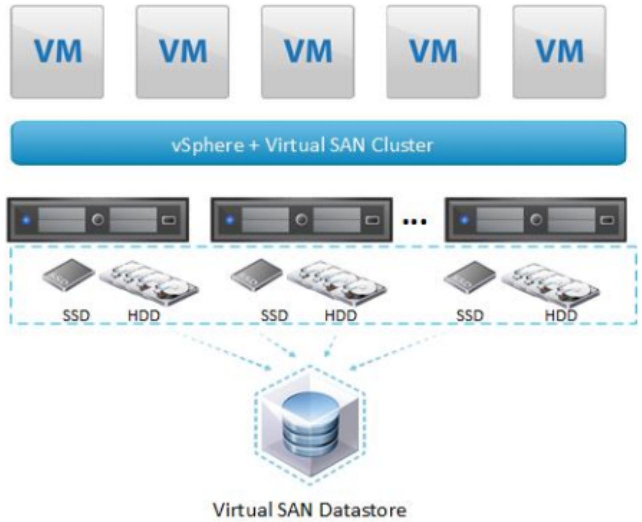

Figure 1. Virtual SAN Cluster Datastore

vSAN can be configured as a hybrid or an all-flash storage. In a hybrid disk architecture, vSAN leverages flash-based devices for performance and magnetic disks for capacity. In an all-flash disk architecture, vSAN can use flash-based devices(PCIe SSD or SAS/SATA SSD) for both caching and persistent storage. It is a distributed object storagesystem that leverages the vSAN SPBM feature to deliver centrally managed, application-centric storage services and capabilities. Administrators can specify storage attributes, such as capacity, performance, and availability, as a policy on a per VMDK level. The policies dynamically self-tune and load balance the system so that each virtual machine has the right level of resources.

This solution explores the hybrid disk architecture for Microsoft SQL Server.

3.4 VMware vSAN Stretched Cluster

vSAN 6.1 introduced the Stretched Cluster feature. vSAN Stretched Clusters provides customers with the ability to deploy a single vSAN Cluster across multiple data centers. vSAN Stretched Cluster is a specific configuration implemented in environments where disaster or downtime avoidance is a key requirement.

vSAN Stretched Cluster builds on the foundation of Fault Domains. The Fault Domain feature introduced rack awareness in vSAN 6.0. The feature allows customers to group multiple hosts into failure zones across multiple server racks in order to ensure that replicas of virtual machine objects are not provisioned onto the same logical failure zones or server racks. vSAN Stretched Cluster requires three failure domainsbased on three sites (two active/active sites and one witness site). The witness site is only utilized to host witness virtual appliances that stores witness objects and cluster metadata information and also provide cluster quorum services during failure events.

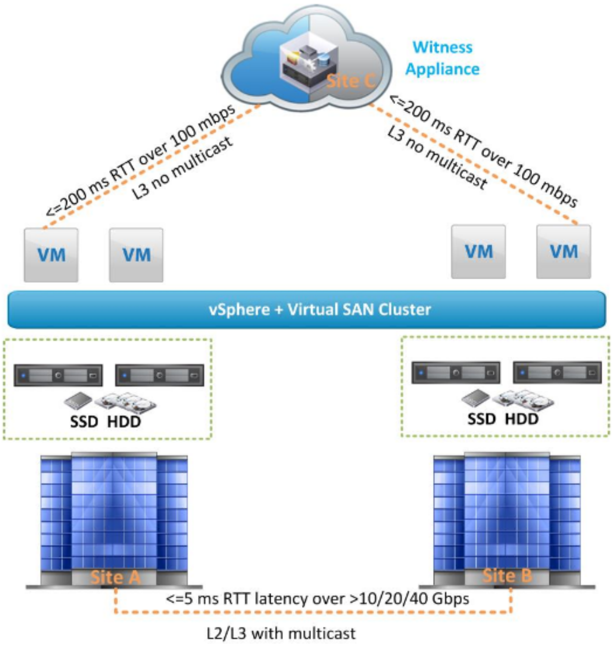

Figure 2. Virtual SAN Stretched Cluster

As shown in Figure 2, vSAN Stretched Cluster can support 5 milliseconds (ms) round trip network latency from site A to site B. The witness of the vSAN is on the site C with maximum 200ms round trip network latency to site A or site B.

3.5 VMware vSphere Data Protection

VMware vSphere Data Protection is a robust, simple to deploy, disk-based backup and recovery solution. vSphere Data Protection is fully integrated with the VMware vCenter™ and enables centralized and efficient management of backup jobs while storing backups in deduplicated destination storage locations.

VMware vSphere Data Protection 6.0 is a software-based solution that is designed to create image-level backups of virtual machines, virtual servers, databases. vSphere Data Protection can utilize application plugins to back up Microsoft SQL Server, Exchange Server, and SharePoint Server.

vSphere Data Protection 6.0 supports backup and recovery of databases on SQL Server database, SQL Server AlwaysOn, and SQL Server Failover Clusters. The vSphere Data Protection agent for SQL Server has the ability to select a replica backup for SQL Server AlwaysOn. You can deploy external proxies in remote locations where vSphere Data Protection does not have direct access to the storage on which the protected VMs are running.

3.6 VMware Site Recovery Manager

VMware Site Recovery Manager 6.1 is an extension to VMware vCenter that provides disaster recovery, site migration, and non-disruptive testing capabilities to VMware customers. It is fully integrated with VMware vCenter Server and VMware vSphere Web Client.

Site Recovery Manager works in conjunction with various replication solutions including vSphere Replication to automate the process of migrating, recovering, testing, reprotecting, and failing-back virtual machine workloads.

Site Recovery Manager Servers coordinate the operations of the VMware vCenter Server at two sites. When virtual machines at the protected site are shutdown, copies of these virtual machines at the recovery site start up. By using the data replicated from the protected site, these virtual machines assume responsibility for providing the same services.

3.7 VMware vSphere Replication

vSphere Replication is an extension to VMware vCenter Server that provides hypervisor-based virtual machine replication and recovery.

vSphere Replication is an alternative to a storage-based replication. With vSphere Replication, you can replicate servers to meet your load balancing needs. After you set up the replication infrastructure, you can choose the virtual machines to be replicated at a different recovery point objective (RPO). You can enable multi-point in time retention policy to store more than one instance of the replicated virtual machine. After recovery, the retained instances are available as snapshots of the recovered virtual machines.

3.8 Microsoft SQL Server 2014

Microsoft SQL Server 2014 is the latest version of Microsoft database management and analysis system for e-commerce, line-of-business, and data warehousing solutions. This solution depicts one of the latest replication features of SQL Server 2014 and Always On (AlwaysOn Availability Groups in particular).

AlwaysOn

AlwaysOn refers to the new comprehensive HA and DR solution for SQL Server. AlwaysOn presents new and enhanced capabilities for both specific databases and entire instances, providing flexibility to support various high availability configurations through:

- AlwaysOn Failover Cluster Instances
- AlwaysOn Availability Groups

The highlighted features of AlwaysOn of SQL Server 2014 include:

- The maximum number of secondary replicas is eight.

- When disconnected from the primary replica or during cluster quorum loss, readable secondary replicas remain available for read workloads.

These features bring much resilience to the production SQL Server to sustain single point of failure with flexible service level.

This solution explores AlwaysOn Availability Groups, with a focus on the transaction-level replication feature to support industry-standard TPC-E like workloads.

AlwaysOn Availability Groups

AlwaysOn Availability Groups is an HA and DR solution introduced in SQL Server 2012, enabling administrators to maximize availability for one or more user databases. SQL Server instances are configured so that a single or a group of primary databases can have up to eight secondary database copies residing on Windows Server Failover Cluster (WSFC) nodes.

Availability Replicas and Roles

Availability groups consist of a set of two or more failover partners referred to as an availability replica. Each availability replica is hosted on a separate instance of SQL Server that in turn resides on separate nodes of a WSFC cluster. Each of the SQL Server instances is a SQL Server Failover Cluster Instance, or a standalone instance with AlwaysOn Availability Groups enabled, as shown inFigure 3.

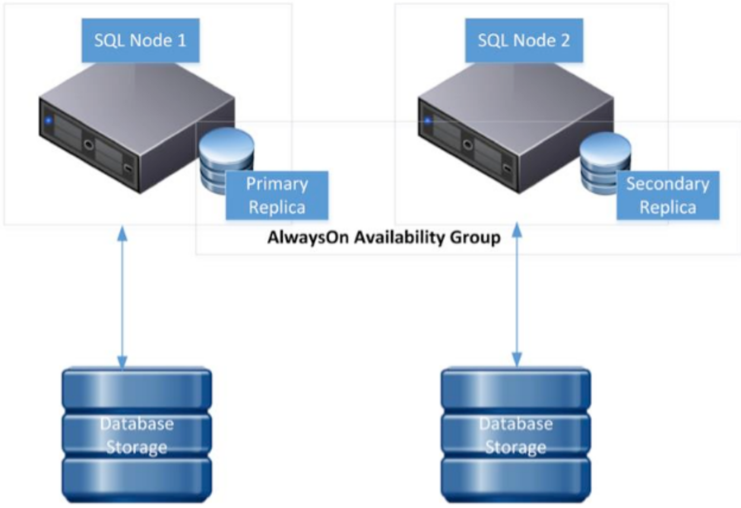

Figure 3. SQL Server AlwaysOn Availability Groups

Each availability replica hosts a copy of the availability databases in the availability group, with each availability replica assigned an initial role (primary or secondary):

- Primary replica
  Holds the primary role and there can be only one. A primary replica hosts the read-write databases that are known as the primary databases.
- Secondary replica
  There can be up to eight replicas, each holding the secondary role and hosting the read-only databases.
  Each secondary replica has the ability to be transitioned to the primary replica during a failover.

AlwaysOn Availability Groups supports two modes:

- The asynchronous-commit mode
  In the asynchronous-commitmode, the primary replica commits a transaction without acknowledgement that an asynchronous-commit replica has hardened the log. The asynchronous-commit mode minimizes the transaction latency, allowing the secondary database to lag behind the primary database. This might cause data loss.

- The synchronous-commit mode
  In the synchronous-commit mode, the primary replica waits for acknowledgment that a synchronous-commit secondary has hardened the log before committing a transaction. Synchronous-commit mode increases the transaction latency but protects against data loss. If the secondary database is in a synchronized state with the primary database, the committed transactions are fully protected.

  SQL Server AlwaysOn Availability Group with the synchronous commit mode emphasizes availability over performance. This synchronous commit mode ensures a transaction on the mirroring principal server or primary replica does not commit until it receives a message back from the mirror or secondary replica that the transaction has been hardened to the transaction log. This increases the transactional latency, particularly when the servers are in different physical locations.

4. Solution Configuration

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

4.1 Solution Configuration

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

- Architecture diagrams
- Hardware resources
- Software resources
- Network configuration
- VMware ESXi Server: storage controller mode
- vSAN configuration
- SQL Server virtual machine configuration
- SQL Server availability groups configuration

4.2 Architecture Diagrams

The key designs for the vSAN Cluster (Hybrid) solution for SQL Server are:

- vSAN Cluster: no fault domain is defined. By default, each node is one fault domain, which means every component of a virtual machine can be placed on each node in the vSAN Cluster. The virtual disk components of the SQL Server virtual machine are distributed across the four nodes. There are multiple ways that components could be placed on the hosts and drives. Per SPBM definition, different vSAN FTT settings define different replica number and each replica can have multiple stripes.
  Figure 4 is an example to show that the virtual machine has two virtual disks with the default storage policy (FTT=1, stripe width=1), which means the virtual machine has four replica components (DATA1 and DATA1', DATA2 and DATA2') and two witnesses.

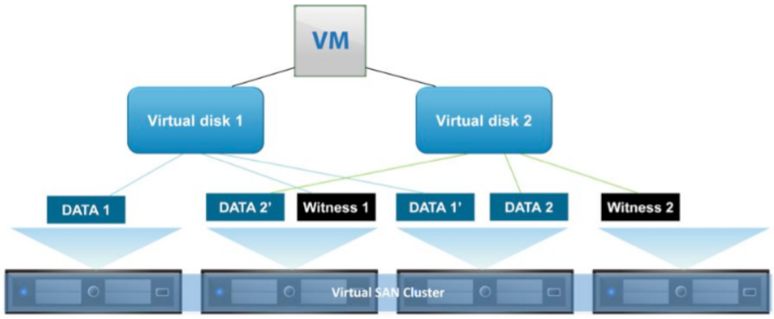

Figure 4. Component Distribution of the Virtual Machine in Virtual SAN Cluster

vSAN Stretched Cluster:4-node vSAN Cluster, with two nodes on one site and two nodes on the other site. The witness is a virtual machine as shown in Figure 5.

We defined two fault domains and deployed an ESXi Server Appliance on the witness site. The virtual disk components of the SQL Server virtual machine are distributed on the two fault domains and the witness of the components is on the witness host. There are multiple ways that components could be placed on the hosts and drives. Per SPBM definition, vSAN Stretched Cluster FTT setting should be one but each replica can have multiple stripes. Figure 5 is an example to show that a virtual machine has two virtual disks, with each VMDK having two replicas and one witness. In the vSAN Stretched Cluster, two replica components (DATA1 and DATA1', DATA2 and DATA2') should be on the fault domain A (site A) and the fault domain B (site B), and all witnesses should be on fault domain C (site C).

The vSAN replica provides the data availability from the storage perspective, which ensures each object of the virtual machine can provide the data-store-level service continuity when hardware failure happens. And the database availability group provides the database availability(including the logical database damage) when the SQL Server virtual machine hosting the database replica fails.

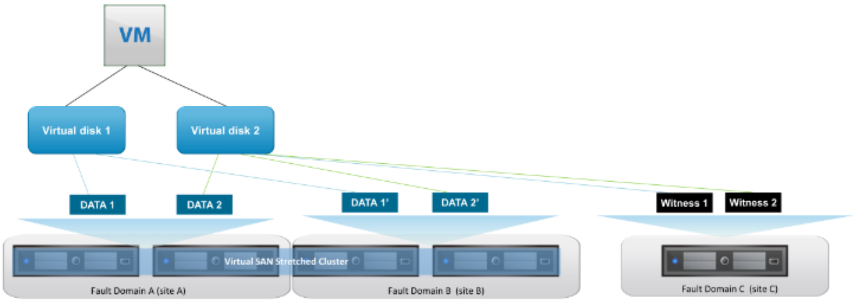

Figure 5. Component Distribution of the Virtual Machine in Virtual SAN Stretched Cluster

Figure 6 is the logical archit ecture of vSAN Cluster. In this architecture, the four SQL Server virtual machines are deployed on a 4-node vSAN Cluster. Each virtual machine can be on any of the ESXi servers. To balance computing resource utilization and reduce contention, the recommended virtual machine placement on the vSAN Cluster is to place the virtual machines on different ESXi hosts. Each SQL Server virtual machine hosts a primary replica of an availability group as well as a secondary replica of theavail ability group. We chose the Node and File Share Majority as the Windows cluster quorum mode to align with the recommendation by Microsoft for the even number of nodes.

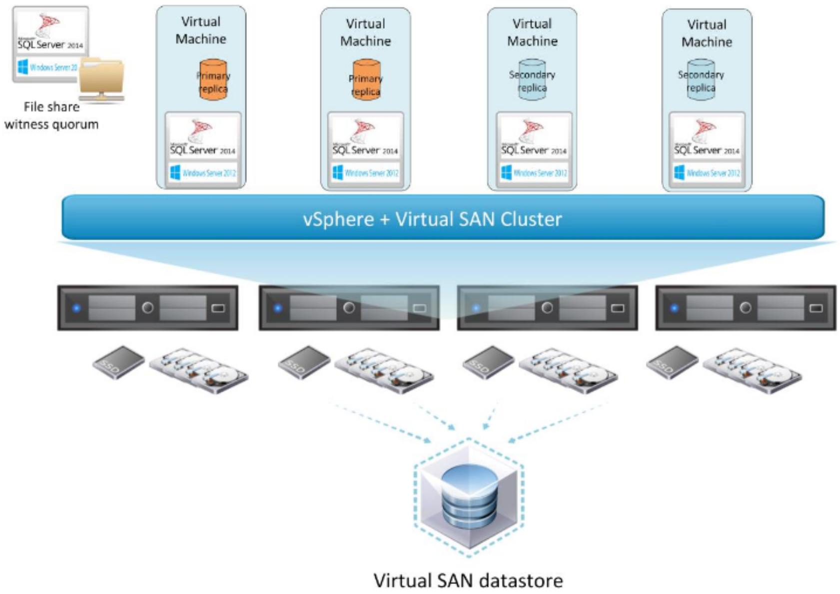

Figure 6. Logical Components of Virtual SAN Cluster

Figure 7 is the logical architecture of the vSAN Stretched Cluster. In this architecture, the four SQL Server virtual machines are on the 4-node vSAN Stretched Cluster. The four nodes are two and two geographically dispersed with the maximum 4ms round-trip network latency from site A to site B. The witness of the vSAN is on the site C with the maximum 200ms round-trip network latency to site A orto site B. Similarly, for the virtual machine placement, we deployed the SQL Server virtual machines on different ESXi hosts regardless of the site. Each SQL Server virtual machine hosts a primary replica of an availability group as well as a secondary replica of the vailability group. vSAN ensures that redundant virtual disk component copies of every AlwaysOn replica are distributed to both sites.

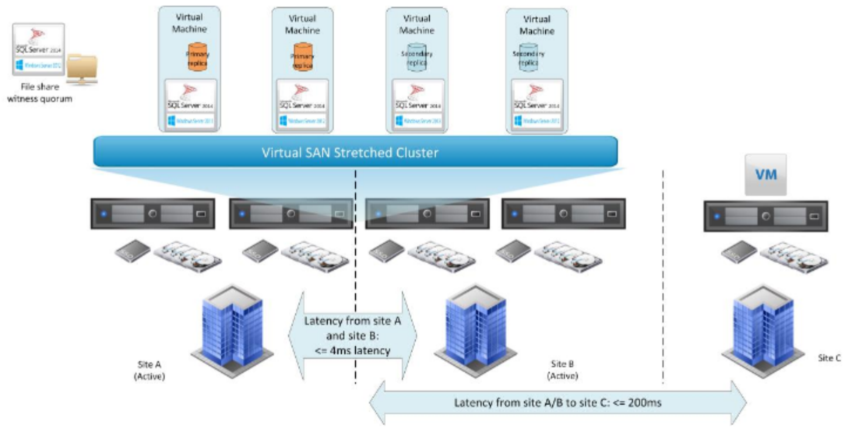

Figure 7. **Logical Components of Virtual SAN Stretched Cluster**

4.3 Hardware Resources

We used direct-attached SSDs and HDDs on ESXi host to provide vSAN solution, which is a scalable and enterprise-class storage solution. Each ESXi host has two disk groups each consisting of one SSD and six HDDs. We combined the disk groups to form a vSAN datastore. This next-generation storage platform combines powerful and flexible hardware with advanced efficiency, management, and software-defined storage.

Each ESXi server in the vSAN Cluster has the following configuration as shown in Figure 8.

Table 2. ESXi Host Configuration

PROPERTY	SPECIFICATION
CPU	2 sockets, 10 cores each
RAM	256GB
Network adapter	2 x 10Gb and 1 x 1Gb
Storage adapter	2x 12Gbps SAS 3008 PCI-Express
Disks	12 x 900GB 10k RPM SAS 2 x 400GB SATA SSD

4.4 Software Resources

Table 3 shows the software resources used in this solution

Table 3 Software Resources

SOFTWARE	VERSION	PURPOSE
VMware vCenter and ESXi	6.0 U1	ESXi cluster to host virtual machines and provide vSAN Cluster. VMware vCenter Server provides a centralized platform for managing VMware vSphere environments.
VMware vSAN	6.1	Software-defined storage solution for hyper-converged infrastructure
vSphere Data Protection	6.0	VMware Application backup and restore solution
Site Recovery Manager	6.1	VMware disaster and recovery solution
vSphere Replication	6.0	VMware hypervisor-based virtual machine replication solution
Microsoft SQL Server 2014	Enterprise Edition, RTM	Database software
Windows Server 2012	2012 R2 x64 SP1, Enterprise Edition	SQL Server database virtual machines Load generation virtual machines Domain controller VMware vCenter server
Benchmark Factory	7.2	TPC-E like data generator and workload test client

4.5 Network Configuration

A VMware vSphere Distributed Switch™ acted 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. We used two 10GbE adapters per host for the vSphere Distributed Switch.

Table 4 shows the security property settings used in this solution

Table 4. Port Group Properties—vSphere Distributed Switch

PROPERTY	SETTING	VALUE
	Promiscuous mode	Accept
	MAC address changes	Accept
	Forged transmits	Accept

In this solution, because the witness of this Stretched Cluster is a nested ESXi Server, changing the promiscuous mode from Reject to Accept is required for the underlying nested ESXi VM vmnic to provide connectivity for the underlying nested VMs to connect to the virtual switch.

Network I/O control was enabled for the distributed switch. Table 5 shows the settings and share values applied on the resource allocation.

Table 5. Network Resource Allocation

NETWORK RESOURCE POOL	HOST LIMIT (MB/S)	PNICSHARES	SHARES
VMware vSphere vMotion®	7,500	Low	25
Management	Unlimited	Normal	50
Virtual machines	Unlimited	High	100
vSAN traffic	Unlimited	Normal	50

The vSphere Distributed Switch uses two 10GbE adapters per host as shown in Figure 9.

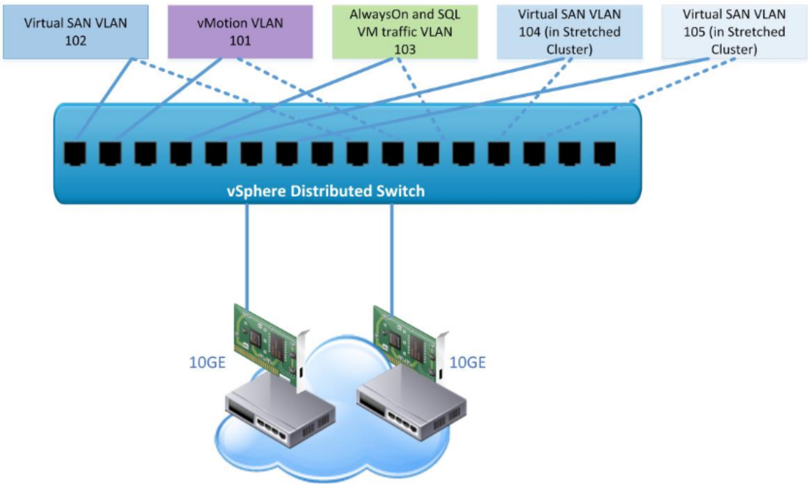

Figure 9. vSphere Distributed Switch Configuration

A port group defines properties regarding security, traffic shaping, and NIC teaming. Default port group setting was used except the uplink failover order as shown in Table 6. It also shows the distributed switch port groups created for different functions and the respective active and standby uplink to balance traffic across the available uplinks. VLAN 103 and VL AN 104 were configured in the vSAN Stretched Cluster.

Table 6. Distributed Switch Port Groups

DISTRIBUTED SWITCH PORT GROUP NAME	VLAN	ACTIVE UPLINK	STANDBY UPLINK
AlwaysOn and SQL traffic	103	Uplink1	Uplink2
vSphere vMotion	101	Uplink1	Uplink2
vSAN (site A)	102	Uplink2	Uplink1
vSAN (site B)	104	Uplink2	Uplink1
vSAN (site C)	105	Uplink2	Uplink1

4.6 VMware ESXi Server: Storage Controller Mode

The storage controller supports both pass-through and RAID mode. We used pass-through mode in the testing. Pass-through the preferred storage controller mode giving vSAN complete control of the local SSDs and HDDs attached to the storage controller.

The storage controller supports both pass-through and RAID mode. We used pass-through mode in the testing. Pass-through the preferred storage controller mode giving vSAN complete control of the local SSDs and HDDs attached to the storage controller.

4.7 vSAN Configuration

vSAN configur ation includes the vSAN Cluster design for SQL Server 2014, and the setup and configuration for the vSAN Stretched Cluster. The storage design for the SQL Server virtual machines, such as the storage policy and virtual machine configuration are applied to the vSAN Stretched Cluster as well.

vSAN Storage Policy for SQL Server 2014

vSAN can set availability, capacity, and performance policies per virtual machine. We designed and implemented the following storage policiesas shown in Table 7.

Table 7. vSAN Storage Setting for SQL Server

STORAGE CAPABILITY	SETTING
Number of FTT	1
Number of disk stripes per object	1
Flash read cache reservation	0%
Object space reservation	100%

Number of FTT-The FTT policy defines how many concurrent host, network, or disk failures can occur in the cluster and still ensure the availability of the object. The onfiguration contains FTT+1 copies of the object to ensure that the object's data is available even when the number of tolerated failures occurs.

Object Space Reservation—By default, a virtual machine created on vSAN is thin-provisioned. It does not consume any capacity until data is written. You can change this setting between 0 to 100 percent of the virtual disk size.

The combination of the object space reservation percentage and the FTT settings applied to the virtual machines on the vSAN datastore determines the usable capacity of datastore.

Number of Disk Stripes per Object—This policy defines how many physical disks across each copy of a storage object are striped. The default value (recommended) of one was sufficient for our tested workloads.

4.8 SQL Server Virtual Machine Configuration

We configured four SQL Server virtual machines for the solution tests. Table 8 lists the SQL Server 2014 VMconfiguration. The virtual machines shared the same vSAN datastore, providing the following benefits:

- vSphere vMotion between hosts is possible to provide application mobility and availability.
- The SPBM-based storage makes the storage management and configuration simpler compared to traditional SAN-based solution.

We designed four SQL Server virtual machines for the performance test for the two availability groups. The data bases were created by the Benchm ark Factory with scale 18 and 35 respectively. The data and index consumed around 180GB and 350GB space. We assigned 12 and 16 vCPUs to each VM hosting the 180GB and 350GB databases. The VM hosting the secondary replica database has the identical configuration with the primary replica. We set the maximum server memory to 120GB for each SQL Server instance with **Lock Pages in Memory** privilege granted to the SQL Server instance startup account, and left 8GB memory to the operating system and test client running on the primary SQL Server instance.

Table 8. SQL Server 2014 VM Configuration

SQLVMROLE	VCPU	RAM(GB)	OPERATINGSYSTEM
VM1—180GB primary replica DB	12	128	Windows Server 2012 Datacenter 64-bit
VM2—180 GB secondary replica DB	12	128	Windows Server 2012 Datacenter 64-bit
VM3—350 GB primary replica DB	16	128	Windows Server 2012 Datacenter 64-bit
VM4—350 GB secondary replica DB	16	128	Windows Server 2012 Datacenter 64-bit

Database VM and Disk Layout

For the TPC-Elike workload, the database size is based on the actual disk space requirement and additional space for the database growth. According to the design (see SQL Server Availability Group Configuration for details), one VM hosts one primary database and the secondary replica is on another VM. To reduce the file allocation contention and the management complexity, we created multiple virtual disks for data files to spread the load equally across the files and virtual disks. For the 180GB database virtual machine, each virtual machine has one OS disk, four data file and log disks, one tempdb data and one tempdb log disk. For the 350GB database
virtual machine, six data file disks are configured and the other disk layout is the same as the 180GB database VM. The VM hosting the secondary replica (VM2 and VM4) has the exact same disk layout as the one hosting the primary replica. Table 9 details the virtual disk layout for the four virtual machines. Data, log, and tempdb disks are separated on to different VMware Paravirtual SCSI (PVSCSI) controllers.

Table 9. VM Virtual Disk Layout

VM1/VM2 DISK LAYOUT			SCSI CONTROLLER
40GB	Drive C:	Windows OS disk	LSI Logic SAS
100GB	Drive E:	Database disk	PVSCSI Controller 1
100GB	Drive F:	Database disk	PVSCSI Controller 1
100GB	Drive G:	Database disk	PVSCSI Controller 1
100GB	Drive H:	Database disk	PVSCSI Controller 1
125GB	Drive L:	SQL log disk	PVSCSI Controller 2
160GB	Drive P:	Page file disk	LSI Logic SAS
80GB	Drive S:	tempdb data file disk	PVSCSI Controller 3
250GB	Drive T:	tempdb log file disk	PVSCSI Controller 1
VM3/VM4 DISK LAYOUT			SCSI CONTROLLER
40GB	Drive C:	Windows OS disk	LSI Logic SAS
100GB	Drive E:	Database disk	PVSCSI Controller 1
100GB	Drive F:	Database disk	PVSCSI Controller 1
100GB	Drive G:	Database disk	PVSCSI Controller 1
100GB	Drive H:	Database disk	PVSCSI Controller 2
100GB	Drive I:	Database disk	PVSCSI Controller 2

100GB	Drive J:	Database disk	PVSCSI Controller 2
125GB	Drive L:	SQL log disk	PVSCSI Controller 2
160GB	Drive P:	Page file disk	LSI Logic SAS
80GB	Drive S:	tempdb data file disk	PVSCSI Controller 3
250GB	Drive T:	tempdb log file disk	PVSCSI Controller 2

5. SQL Server Availability Group Configuration

This section outlines the process for implementing AlwaysOn Availability Groups across three SQL Server instances.

5.1 Overview

This section outlines the process for implementing AlwaysOn Availability Groups across three SQL Server instances. The solution places all three SQL Server instances in the same vSAN datastore.

5.2 Prerequisites

See Windows System Requirements and Recommendations for the prerequisites for SQL Server 2014 AlwaysOn.

5.3 Windows Server Failover Clustering (WSFC) Setup

As part of the SQLServer 2014 availability group creation, we configured 2 x 2-node clusters on Windows Server 2012 with file share quorum.

5.4 WSFC Quorum and vSAN Witness

WSFC uses a quorum-based approach to monitor Windows cluster node status. In case of a host failure, Windows cluster can vote to determine whether Windows Cluster service can continue to work or not, while the witness in vSAN provides cluster quorum services during failure events of the disk objects.

WSFC uses a quorum-based approach to monitor Windows cluster node status. In case of a host failure, Windows cluster can vote to determine whether Windows Cluster service can continue to work or not, while the witness in vSAN provides cluster quorum services during failure events of the disk objects.

5.5 WSFC Quorum Modes and Voting Configuration

SQL Server AlwaysOn Availability Groups take advantage of WSFC technology. WSFC uses a quorum-based approach to monitor overall cluster health and provide the maximum node-level fault tolerance. Because two 2-node SQL Server Clusters are designed in this solution, the node and file share majority that is recommended for clusters with an even number of nodes is used as the quorum configuration for each cluster. For more information, see the Understanding Quorum Configurations in a Failover Cluster topic.

5.6 Availability Group Configuration

We designed two availability groups in this solution:

- Four virtual machines supported the two database groups to fully leverage the virtual machine computing resource:
  - We placed the 180GB primary database replica on the virtual machine that has 12 vCPUs and 128GB memory.
  - We placed the 350GB database replica on the virtual machine that has 16 vCPUs and 128GB memory.

This solution deployed the availability groups with sync hronous mode for the two database groups to ensure high availability of the user database, while achieving excellent application performance.

Each database group has one listener configured. The listener client connectivity can be provided to a database of the availability group. This unique DNS name serves as a virtual network name (VNN) and provides a set of resources that are attached to an availability group to direct client connections to the appropriate availability replica. Also this listener is used for the vSphere Data Protection backup for the AlwaysOn Availability Groups. For a detailed description about availability group listener, see Availability Group Listeners, Client Connectivity, and Application Failover (SQL Server). For a detailed description about the vSphere Data Protection backup, see the SQL Server Database Backup and Restore on vSAN section.

6. Solution Validation

SQL Server is an I/O-intensive activity process and the database is normally larger than the amount of the memory configured, regardless of whether it is virtualized or not.

6.1 Overview

SQL Server is an I/O-intensive activity process and the database is normally larger than the amount of the memory configured, regardless of whether it is virtualized or not. Therefore, SQL Server has to pull data from disk to satisfy queries. In addition, because data in databases is constantly changing for typical OLTP workload, these changes need to be written to disk, asynchronously to data files and synchronously to log files.

Performance is an importantpart of SQL Server deployment so we need to make sure the I/O subsystem is not the bottleneck. This solution built production SQL Server with AlwaysOn feature enabled to provide sustainable and scalable performance on the vSAN platform.

This solution validated the storage resiliency capability of vSAN against hardware failures with running OLTP workload.

For the production SQL Server OLTP database, backup and restore, disaster and recovery are highly demanded. This solution leveraged vSphere Data Protection and Site

Recovery Manager with vSphere Replication to address these requirements.

SQL Server is an I/O-intensive activity process and the database is normally larger than the amount of the memory configured, regardless of whether it is virtualized or not. Therefore, SQL Server has to pull data from disk to satisfy queries. In addition, because data in databases is constantly changing for typical OLTP workload, these changes need to be written to disk, asynchronously to data files and synchronously to log files.

Performance is an importantpart of SQL Server deployment so we need to make sure the I/O subsystem is not the bottleneck. This solution built production SQL Server with AlwaysOn feature enabled to provide sustainable and scalable performance on the vSAN platform.

This solution validated the storage resiliency capability of vSAN against hardware failures with running OLTP workload.

For the production SQL Server OLTP database, backup and restore, disaster and recovery are highly demanded. This solution leveraged vSphere Data Protection and Site Recovery Manager with vSphere Replication to address these requirements.

SQL Server is an I/O-intensive activity process and the database is normally larger than the amount of the memory configured, regardless of whether it is virtualized or not. Therefore, SQL Server has to pull data from disk to satisfy queries. In addition, because data in databases is constantly changing for typical OLTP workload, these changes need to be written to disk, asynchronously to data files and synchronously to log files.

Performance is an importantpart of SQL Server deployment so we need to make sure the I/O subsystem is not the bottleneck. This solution built production SQL Server with AlwaysOn feature enabled to provide sustainable and scalable performance on the vSAN platform.

This solution validated the storage resiliency capability of vSAN against hardware failures with running OLTP workload.

For the production SQL Server OLTP database, backup and restore, disaster and recovery are highly demanded. This solution leveraged vSphere Data Protection and Site

Recovery Manager with vSphere Replication to address these requirements.

6.2 Test Overview

The solution validated the performance and functionality of enterprise-class SQL Server instances in a virtualized VMware environment running SQL Server 2014 with the AlwaysOn feature deployed on the vSAN storage platform. Test scenarios include:

- Various database size performances with OLTP workload on vSAN
- vSAN resiliency in hardware failure scenarios (single HDD failure, disk group failure, and storage host failure)
- Highly available application deployed on vSAN Stretched Cluster
- Disaster Recovery with vSAN Stretched Cluster and Site Recovery Manager
- SQL Server backup and restore by vSphere Data Protection on vSAN

6.3 TPC-E like OLTP Workload

TPC-E like workload is a type of OLTP workload. It is a mixture of read-only and update-intensive transactions that simulate the activities found in complex OLTP application environments. We designed the database schema, data population, transactions, and implementation rules to present modern OLTP systems.

The I/O access pattern of TPC-E is read-intensive. TPC-E workload is 8k dominant random I/O. This kind of workload can benefit from using SSD.

Read Cache in vSAN

In vSAN, read cache is used only in the case of hybrid vSAN Clusters, where each disk group consists of one SSD and one or more HDDs. vSAN uses the SSD device as the performance tier of each disk group and uses the HDD devices as the capacity tier of the disk group. Part of the SSD is used as the corresponding disk group for the read operations. The purpose is to serve the highest possible ratio of read operations from data staged and minimize the portion of read operations that are served by the HDDs. Generally, if the required blocks can mostly hit from a vSAN cache, hybrid vSAN will be the best choice as a cost-effective solution. However, if the read pattern is too random to hit from the vSAN cache, consider using an all-flash vSAN. Also, the performance result on the vSAN storage depends on various design factors that are addressed in the Best Practices of Virtualized SQL Server on vSAN chapter in this documentation.

6.4 TPC-E like OLTP Validation Tools

We used the OLTP workload simulation tools to run tests with the desired parameters and start vSAN Observer as well as Windows Performance Monitor in each VM.

- Benchmark Factory

  Benchmark Factory for Databases is a database performance testing tool that enables you to conduct database industry-standard benchmark testing and scalability testing.
  See Benchmark Factory for Databases for more information.

MONITORCOUNTER	DESCRIPTION

Average response time	The time frame from when SQL is sent to a response is received.
Transaction per second (TPS)	Measures the transactions in the user database. The ideal value is as large as possible for a designed SQL Server.

- Windows Performance Monitor

Performance Monitor is a Windows tool that enables users to capture statistics about SQL Server, memory usage, I/O throughput from SQL Server instances and operating system levels. There are several counters related to I/O and they are located under physical and logical disks. Table 10 lists there corded performance monitor counters to measure the database performance.

Table 10. Performance Monitor Counter

MONITOR COUNTER	DESCRIPTION
Average disk sec/read and average disk sec/write	Measures disk latency of read and write. Microsoft's recommended optimal value for average over time for OLTP is equal or less than 20ms. If the realized throughput is reasonable for the particular configuration, higher latencies may be acceptable for heavy workload periods.
Disk transfers per second	Measures the I/O numbers per second. The ideal value is as large as possible for a given I/O subsystem.

- vSAN Observer

vSAN Observer is designed to capture performance statistics and bandwidth for a VMware vSAN Cluster. It provides an in-depth snapshot of IOPS, bandwidth, and latencies at different layers of vSAN, read cache hits and missesratio, outstanding I/Os, and congestion. This information is provided at different layers in the vSAN stack to help troubleshoot storage performance. For more information about the VMware vSAN Observer, see the Monitoring VMware vSAN with vSAN Observer documentation.

7. Database Performance - OLTP Workload on VSAN

We created two TPC-E like databases for the data performance validation on vSAN.

7.1 Overview

We created two TPC-E like databases for the data performance validation on vSAN. We used Benchmark Factory with scale 18 and 35 respectively and it resulted in approximately a 180GB and a 350GB database. We created two AlwaysOn database groups with the synchronous mode. For every database on two virtual machines, we allocated same size for data, log, and tempdb virtual disks.

7.2 Test Scenario

The solution validated the performance and functionality of enterprise-class SQL Server instances in a virtualized VMware environment running SQL Server 2014 with the AlwaysOn feature on the vSAN storage platform.

We validated the AlwaysOn enabled SQL Server 2014 in a virtualized environment backed by vSAN to understand the performance characteristics of SQL Server databases. We ran the test by 100 users for 180GB database and 350GB database to generate an acceptable workload on the system. Every test run was two hours. The first one hour was the ramp-up time to achieve a steady state and we gathered the statistics based on the remaining one hour.

Note: This test validation is to achieve the optimal transaction result with acceptable storage subsystem performance. And the transaction result is related with the vCPU and memory of the virtual machines that are bound to the physical hosts. Refer to Table 8 for the virtual machine configurations: better transaction result needs more powerful physical hosts.

7.3 Test Results

To establish a performance baseline, we ran the OLTP workload on a 4-node vSAN Cluster with AlwaysOn enabled. Each VM contained asingle instance of SQL Server and a single user database. We adjusted the number of users within Benchmark Factory to produce a realistic production workload in terms of CPU and I/O, and measured the performance in terms of TPS and average response time in millisecond. We ran multiple tests with similar workload and observed the similar test result. As shown in Figure 10, the 4-node vSAN Cluster with the designed configuration can support up to 2,654 TPS with equaling or less than 25ms average response time.

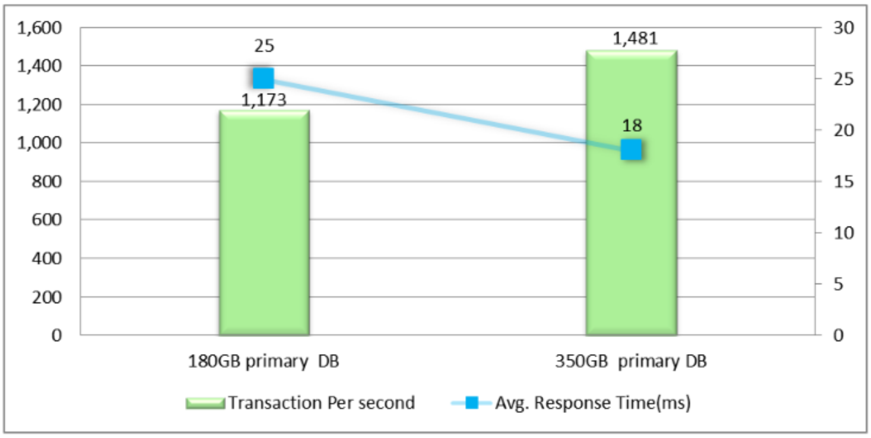

Figure 10. **TPS and Average Response Time of AlwaysOn Performance on Virtual SAN**

In the AlwaysOn database group with synchronous-comm it, the secondary SQL Server needs completing log hardening before acknowledging the transaction; thus the committed transactions are fully protected. This protection comes at the cost of increased transaction latency and additional I/Os on the secondary SQL Server. The results showed that the 4-node vSAN Cluster supported primary workload with the synchronous-comm it modeat around 6, 475 IOPS.

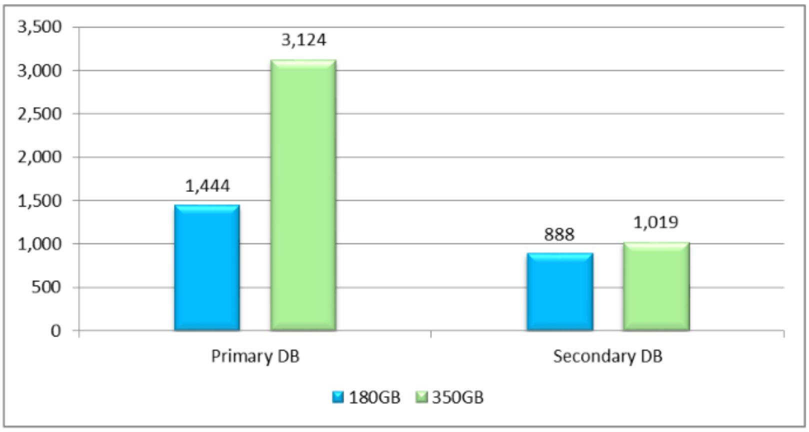

Figure 11. **IOPS of AlwaysOn Performance on Virtual SAN**

The average virtual disk latency per second measured disk latency for the underlying storage sub system of SQL Server. Microsoft's recommended optimal value for average transactional log is less than 5ms, and the optimal value for database data file is less than 20ms. We measured every data virtual disk and log virtual disk. The maximum data disk latency is the maximum value of the 4 data disks for 180GB databases (primary and secondary replica) and 6 data disks for 350GB databases. The test result showed the maximum data disk latency was 15ms. And the average transactionallog disk write latency was 2ms. These values are all optimal to run such an OLTP workload.

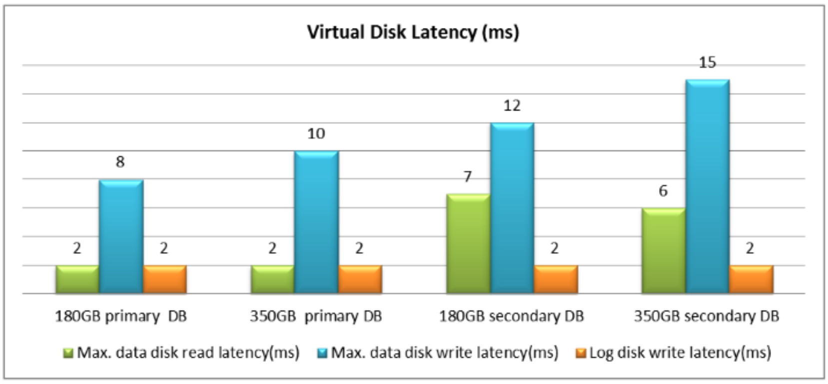

From multiple experimental tests, we concluded the 4-node vSAN Cluster can support 2,654 TPC-E like transactions per second as shown in Figure 11. The average application response time was less than or equal to 25ms as shown in Figure 10. The disk latency was acceptable for AlwaysOn Availability Groups as shown in Figure 12. The results demonstrated the vSAN as an ideal hyper-converged solution for Always On multiple database groups.

8. vSAN Resiliency

This section validated vSAN resiliency in handling disk, disk group, and host failures.

8.1 Overview

This section validated vSAN resiliency in handling disk, disk group, and host failures. We also studied the performance impact on the running SQL Server OLTP workload during failure and recovery. A Benchmark Factory with scale 18 (with approximate 180GB database with 100 configured user number) was initialized for every failure test. We designed the following scenarios to simulate real-world component failures:

- Disk failure
  This test evaluated the impact on the virtualized SQL Server when encountering one HDD failure. The HDD stored the VMDK component of the user database. Inject a permanent disk error to the HDD to simulate a disk failure on one of the nodes to observe whether it has functional or performance impact on the production SQL Server database.
- Disk group failure
  This test evaluated the impact on the virtualized SQL Server when encountering a disk group failure. Inject a permanent disk error to the SSD of the disk group to simulate a disk group failure to observe whether it has functional or performance impact on the production SQL Server database.
- Storage host failure
  This test evaluated the impact on the virtualized SQL Server when encountering one vSAN host failure. Power off one storage host in the vSAN Cluster to simulate host failure to observe whether it has functional or performance impact on the production SQL Server databases.

8.2 Test Scenarios

Single Disk Failure

This test validated the vSAN resiliency in surviving single disk failure and exploited the performance impact of one disk failure on the running OLTP workload. An availability group was created by Benchmark Factory with scale 18. The availability group status was in the synchronized mode before the failure test, which means there is no data loss if manual or automatic failover occurs.

vSAN Resiliency with Single Disk Failure

We simulated an HDD disk failure by injecting a permanent disk error to the HDD when Benchmark Factory was generating workload on a 180GB primary replica of the availability group on the 4-node vSAN. Table 11 shows the failed disk and it had 10 components that stored the VMDK for SQL Server databases. The status of the disk changed from Normal to Degraded. After vSAN dete cted the degraded disk, it started recreating the lost component on the damaged disk. This operation caused component resynchronization.

Table 11. HDD Disk Group

FAILED DISK NAA ID	ESXIHOST	NO OF COMPONENTS	TOTAL CAPACITY (GB)	USED CAPACITY (%)
naa.5XXXXXXXXXXXc05f	1x.xxx.177.96	10	829.97	829.97

Disk Group Failure

We simulated a disk group failure by injecting a permanent disk error to the SSD of the disk group while Benchmark Factory was generating workload on the 180GB database availability group. The failure of SSD rendered an entire disk group inaccessible.

The failed disk group had the SSD and HDD backing disks as shown in Table 12. The table also shows the total number of the affected components that stored the VMDK for SQL Server databases.

Table 12. Failed vSAN Disk Group—Physical Disks and Components

DISKDISPLAY NAME	ESXIHOST	DISKTYPE	NO OF COMPONENTS	TOTAL CAPACITY (GB)	USED CAPACITY (%)
naa.5XXXXXXXXXXX0db6	1xx.xx.177.98	SSD	NA	372.61	37.93
naa.5XXXXXXXXXXX0a73	1xx.xx.177.98	HDD	13	829.97	72.14
naa.5XXXXXXXXXXXeb3f	1xx.xx.177.98	HDD	10	829.97	75.25
naa.5XXXXXXXXXXXd5d7	1xx.xx.177.98	HDD	7	829.97	47.95
naa.5XXXXXXXXXXX2ccb	1xx.xx.177.98	HDD	11	829.97	61.76
naa.5XXXXXXXXXXX0a73	1xx.xx.177.98	HDD	11	829.97	65.69
naa.5XXXXXXXXXXXabdb	1xx.xx.177.98	HDD	8	829.97	59.28

Storage Host Failure

We shut down the ESXi Server that did not host an availability group VM to simulate a storage host failure, when Benchmark Factory was generating workload on a 180GB availability group on the 4-node vSAN. This was intentionally set up to understand the impact of losing a storage host alone on vSAN and avoid SQL VM nterruption.

The failed storage host had two disk group swith the following SSD and HDD backing disks as shown in Table 13.

Table 13. Failed ESXi Host vSAN Disk Group—Physical Disks and Components

DISK GROUP	DISK DISPLAY NAME	ESXI HOST	DISK TYPE	NO OF COMPONENTS	TOTAL CAPACITY (GB)	USED CAPACITY (%)
Diskgroup -1	naa.5XXXXXXXXXXX0bae	1xx.xx.177.98	SSD	NA	372.61	20.47
Diskgroup -1	naa.5XXXXXXXXXXX12af	1xx.xx.177.98	HDD	9	829.97	45.81
Diskgroup -1	naa.5XXXXXXXXXXXcad7	1xx.xx.177.98	HDD	15	829.97	67.78

Diskgroup -1	naa.5XXXXXXXXXXX09cf	1xx.xx.177.98	HDD	13	829.97	67.35
Diskgroup -1	naa.5XXXXXXXXXXXc917	1xx.xx.177.98	HDD	15	829.97	79.93
Diskgroup -1	naa.5XXXXXXXXXXX66f7	1xx.xx.177.98	HDD	19	829.97	79.06
Diskgroup -1	naa.5XXXXXXXXXXXa503	1xx.xx.177.98	HDD	11	829.97	49.99
Diskgroup -2	naa.5XXXXXXXXXXX0db6	1xx.xx.177.98	SSD	NA	372.61	37.93
Diskgroup -2	naa.5XXXXXXXXXXX0a73	1xx.xx.177.98	HDD	13	829.97	72.14
Diskgroup -2	naa.5XXXXXXXXXXXeb3f	1xx.xx.177.98	HDD	10	829.97	75.25
Diskgroup -2	naa.5XXXXXXXXXXXd5d7	1xx.xx.177.98	HDD	7	829.97	47.95
Diskgroup -2	naa.5XXXXXXXXXXX2ccb	1xx.xx.177.98	HDD	11	829.97	61.76
Diskgroup -2	naa.5XXXXXXXXXXX0a73	1xx.xx.28.3	HDD	11	829.97	65.69
Diskgroup -2	naa.5XXXXXXXXXXXabdb	1xx.xx.28.3	HDD	8	829.97	59.28

8.3 Test Results

As shown in Table 14, there were impacts on the performance of the running application during all types of the failures. None of the test reported IO error on Windows VM or Benchmark Factory ODBC disconnections, which demonstrated the resiliency of vSAN during the component failure.

Table 14. vSAN Resiliency Test Result

FAILURE TYPE	TPS BEFORE FAILURE	LOWEST TPS AFTER FAILURE	TIME TAKEN FOR RECOVERY TO STEADY STATEPS AFTER FAILURE(SEC)

Disk failure (HDD)	1,650	900	69
Disk group failure (SSD +HDD)	1,550	850	115
Storage host failure	1,400	520	350

Component Recreation Duration

vSAN tolerates the disk failure by rebuilding objects stored in the failed disk to other disks in the cluster in a failed state. There are two failure statuses: **Absent** and **Degraded**. For the **Absent** status, vSAN waits for the repair delay, which is 60 minutes by default before starting to recreate missed components on the disk. You can change the parameter per the VMware Knowledge Base Article 2075456. For the **Degraded** status, vSAN recreates the components instantly. In the recreating process, a new component on another host appeared in the **Reconfiguring** status. The repair of the affected components on the disk was triggered after the vSAN detected the error.

In the single HDD failure test scenario, we measured the component recreation duration for the degraded components. These components contain the virtual disks for the affected SQL VM (affected capacity of the running SQL VM) and virtual disks for other VMs (total affected capacity). The repair was going on when the workload was running and the total duration was around 122 minutes for approximate 285GB data:

- Failed component:single HDD
- Affected capacity of the running SQL VM: 193GB
- Total affected capacity: 285GB
- Repairing time for the total affected capacity: 122 minutes

The repairing rate was around 2.3GB per minute. Multiple test results show the average repairing rate was 2-3GB per minute.

Note: vSAN actively throttles the storage and network throughput used for reconfiguration to minimize the impact of reconfiguration of objects to normal workload execution.

9. VSAN Stretched Cluster Performance and Resiliency

vSAN Stretched Cluster enables customers to extend the existing vSAN storage to geographically dispersed data centers through a high bandwidth and low latency link.

9.1 Overview

vSAN Stretched Cluster enables customers to extend the existing vSAN storage to geographically dispersed data centers through a high bandwidth and low latency link. Customers can run virtual machines with VMware fully fledged features such as VMware vSphere Distributed Resource Scheduler™ and HA on the active/active sites.

This section validated the following vSAN Stretched Cluster functionalitiesand capabilities:

- vSAN Stretched Cluster setup
- vSAN Stretched Cluster performance with various network latency

9.2 vSAN Stretched Cluster Setup

We simulated a metropolitan area network in a lab environment. Figure 13 shows a diagrammatic layout of the setup. We set up the three sites in different VLANs. A Linux VM configured with three network interfaces each in one of the three VLANs acting as the gateway for in ter-VLAN routing between sites. Static routes were configured on ESXi vSAN VM kernel ports for routing between different VLANs (sites). The Linux VM leverages Netem functionality already built into Linux for simulating network latency between the sites. Furthermore, XORP (open source routing platform) in stalled on the Linux VM provided support for a multicast traffic between two vSAN fault domains. While the network latency between data sites was changed to compare performance impact, inter-site round trip latency from the Witness to data site was kept same at 200ms.

Latency Setup between Sites

With the delay introduced between the sites for vSAN kernel ports, we simulated latency between two availability groups using the Netem functionality inbuilt in Linux.

Figure 14 shows the Stretched Cluster fault domain configuration. We configured and tested inter-site round trip latencies between site A and site B: 1ms, 2ms, and 4ms. The inter-site round trip latency between site A/B and site C was fixed at 200ms.

We put the primary and secondary replica virtual machines to fault domain 1 and fault domain 2 as shown in Figure 13. The network latency induced in the inter-link between fault domain 1 and fault domain 2 would affect the AlwaysOn log transferring from one fault domain to another.

We used Benchmark Factory to generate the same TPC-E like workload (100 users without delay) in these tests.

9.3 vSAN Stretched Cluster with Network Latency

During these tests, we measured the TPS and IOPS on the two AlwaysOn Availability Groups (180GB and 350GB) on the vSAN Stretched Cluster. Before the test, the latency between the ESXi hosts of the vSAN kernel port was around 0.2ms. The inter-link network latency was set to 1ms, 2ms,and 4ms respectively, thus the actual inter-link between the sites would be 1.2ms, 2.2ms, and 4.2ms. With the increase of inter-site latency, we observed a reduction in TPS. For the test workload, the TPS reduction was proportional to the increase in inter-site round trip latency. As shown in Figure 15, for 1.2ms, 2.2ms, and 4.2ms round trip latency, the TPS reduced by around 5 percent, 19 percent, and 29 percent respectively. Similarly, we observed an IOPS reduction in vSAN Observer as shown in Figure 16. We noticed that when inter-site latency ncreased, IO and network latency increased correspondingly.

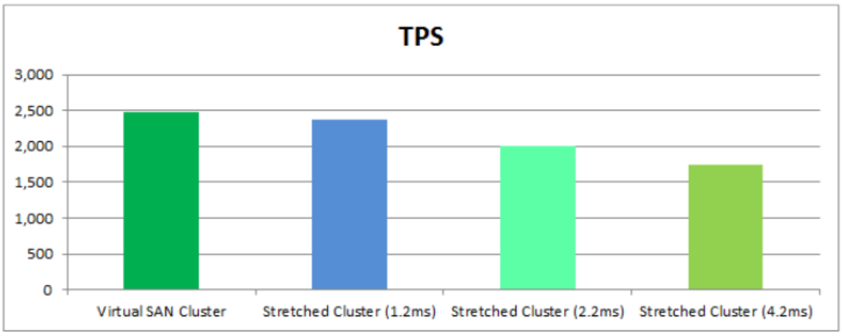

Figure 15. TPS Comparison for a TPC-E like Workload on AlwaysOn Availability Groups

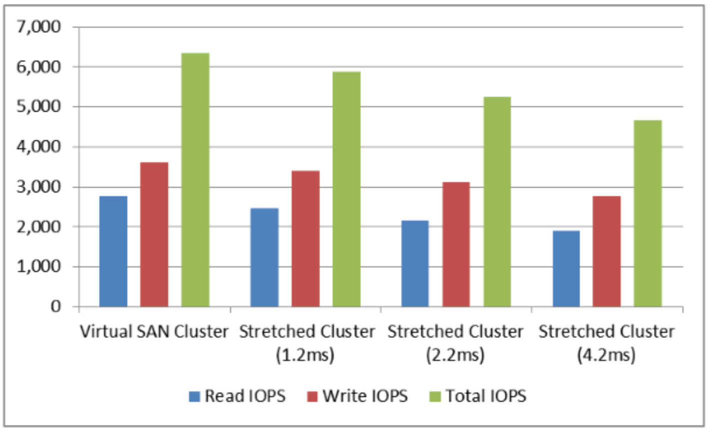

Figure 16. IOPS Comparison for a TPC-E like Workload on AlwaysOn Availability Groups

Site Failure Test Overview

This test demonstrated one of the features of the vSAN Streched Cluster, that is maintaining data availability under the whole site failure.

We ran this test on one availability group on the vSAN Stretched Cluster across two sites over the inter-link. We enabled vSphere HA and vSphere DRS on the cluster, and we generated TPC-E like workload with 100 users on the 180GB AlwaysOn Availability Group. When the test was in astable TPS and IOPS, an entire site was down as “Not esponding” by powering off two ESXi hosts in site A (vSAN Stretched Cluster preferred site) in Figure 17.

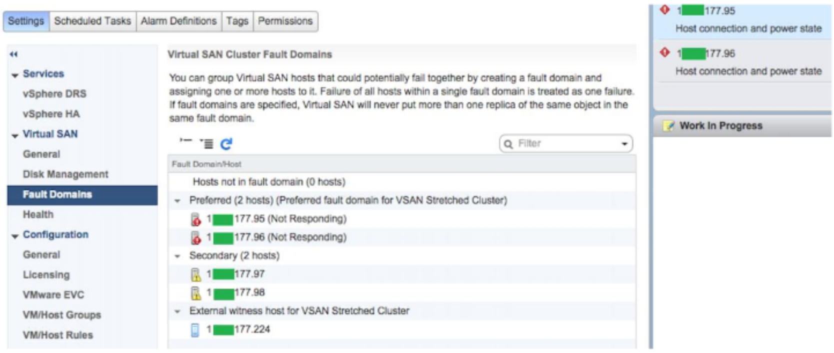

Figure 17. Virtual SAN Stretched Cluster Site Failure

9.4 Test Results

The SQL Server VM in site A was affected after the site failure. And the affected level depends on the role of the availability group with synchronous-commit mode located in the affected site(site A in this test validation).

- If the primary replica is in site A, an availability group failover occurs within one second, and the original secondary replica acts as the primary one. The existing connections to the availability group are terminated and the client needs a new connection to continue working with the same primary database.
- If the secondary replica is in site A, the existing clients can continue working without retrying the connections.

The site outage did not affect data availability because a copy of all the data in site A existed in site B. And the affected VM was automatically restarted by vSphere HA without any issues. The Stretched Cluster is designed to tolerate single site failure at a time. Combing with the AlwaysOn Availability Group ssynchronous mode, the database can automatically fail over to the secondary virtual machine within one second. The vSphere HA can restart the virtual machine affected by the site failure. The availability group can recover in two minutes with the default HA settings.

In the test, when the primary replica was in site B and the secondary replica was in site A, Figure 18 shows the Benchmark Factory TPS when site A was down: the existing clients continued working without retrying the connections.

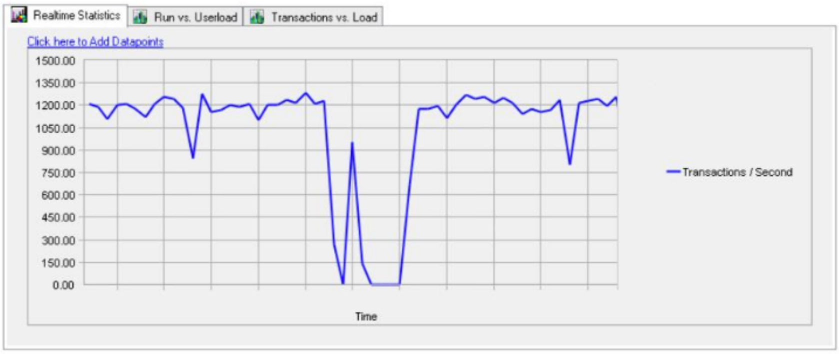

Figure 18. Benchmark Factory TPS during the Stretched Cluster Site Down

10. Disaster Recovery-VSAN Stretched Cluster

Site Recovery Manager is a real disaster recovery, site migration, and non-disruptive solution for VMware customers.

10.1 Overview

Disaster Recovery-VSAN Stretched Cluster & Site Recovery Manager

Site Recovery Manager is a real disaster recovery, site migration, and non-disruptive solution for VMware customers. vSphere Replication is an extension to the vCenter Server that provides hypervisor-based virtual machine replication and recovery.

vSAN Stretched Cluster with Site Recovery Manager provides users an option to run OLTP application in the active/active mode in one city, while having the DR plan to recover the whole production servers in the other city. When one site is down in the same city, vSAN Stretched Cluster can guarantee that the application can run continuously. When there is an unexpected disaster in the city where the Stretched Cluster cannot survive, Site Recovery Manager can bring up the virtual machines in the other city. The minimal RPO by using vSphere Replication is 15 minutes so that data unavailability can be limited within this time frame while keeping the total solution efficient and cost effective.

When using Site Recovery Manager to protect sites on vSAN, you should configure three components on each site for Site Recovery Manager: vCenter Server, Site Recovery Manager, and vSphere Replication.

Figure 19 shows the components used for the backup and recovery plan of SQL Server availability groups.

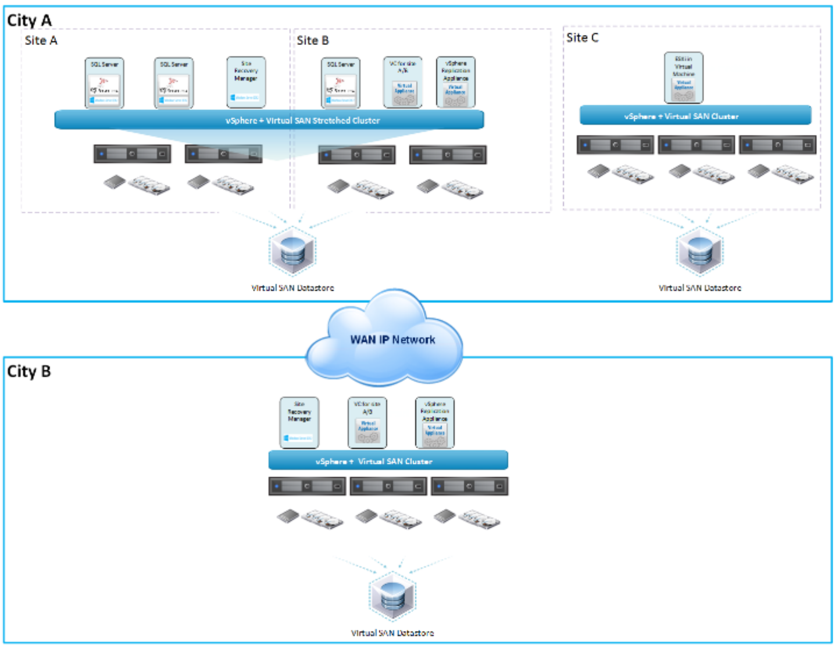

Figure 19. **Backup and Recovery Components of SQL Server Availability Groups**

10.2 Site Recovery Manager Setup

Perform the following steps to set up the DR plan for SQL Server availability groups on vSAN Stretched Cluster. After each step is set up and configured successfully, the protected virtual machines in the protection site can be recovered on the disaster recovery site:

- Install and set up Site Recovery Manager on each site (city A and city B in Figure 19). For more information about how to install and configure Site Recovery Manager, see the VMware vCenter Site Recovery Manager 6.0 Documentation Center.
- Install and configure the vSphere Replication. For more information about how to install and configure vSphere Replication, see the VMware vSphere Replication 6.0 Documentation Center.

- Register vSphere Replication in the vCenter through the https://<vr_server_address>:5480 link.
- Monitor the replication status of vSphere Replication.
  To monitor the replication status and view information about virtual machines configured for vSphere Replication, log in to the vSphere Web Client, click vSphere Replication -> Sites-> Monitor-> vSphere Replication to check three items: outgoing replication, incoming replication, and reports.

  You can observe the monitoring status from the protection site (outgoing replication) or the recovery site (incoming replication).

  The vSphere Replication report summarized the replication status, including replicated VMs, transferred bytes, RPO violations, replication count, site connectivity, and vSphere Replication server connectivity as shown in Figure 21.

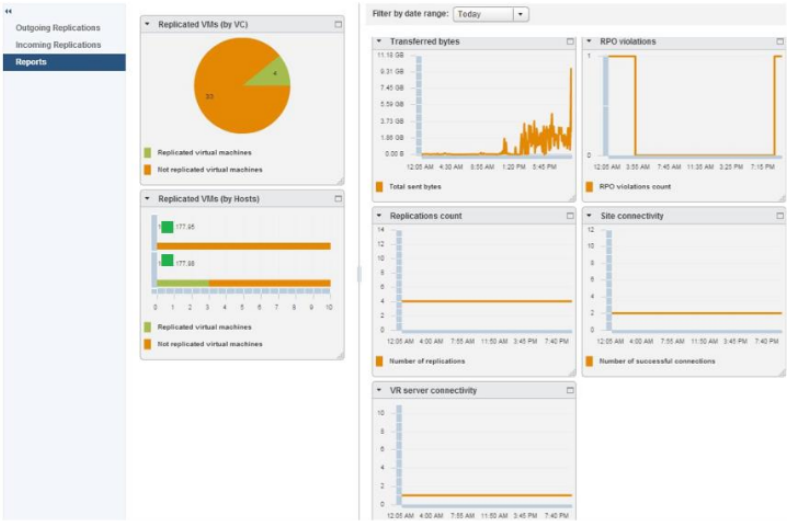

Figure 21. vSphere Replication Report

- Create the protection group and recovery plan for the specified virtual machines as shown in Figure 22 and Figure 23.

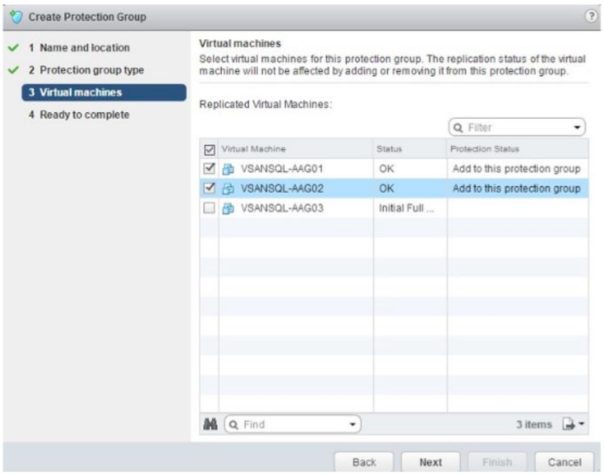

Figure 22. Create the Protection Group

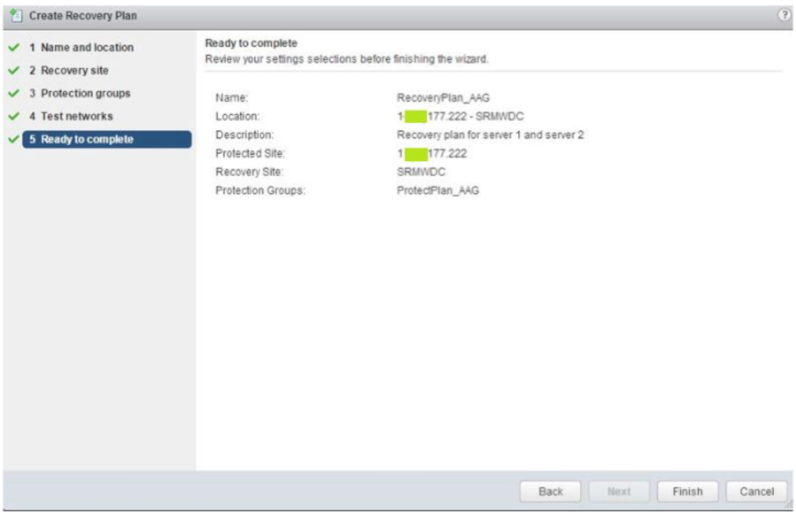

Figure 23. Create the Recovery Plan

- Run the test plan to make sure the protection site can take over the virtual machines in the protection group on the protection site. Figure 24 displays the completed test for the two SQL Server virtual machines with availability groups.

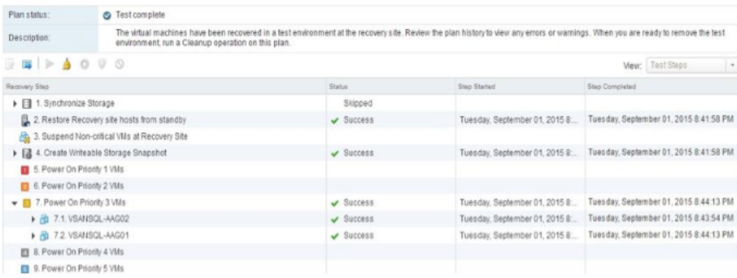

Figure 24. Test Plan Status

- Clean up the recovery plan: click Site Recovery ->select Recovery Plans on the left pane-> select the specific recovery plan -> Actions -> Cleanup.

10.3 Summary

The functional summaries are:

- vSAN provides a unified storage platform for both the protection and the recovery sites serving as a storage platform and being interoperated with Site Recovery Manager.
- The minimal RPO of the vSphere Replication can be 15 minutes to 1,440 minutes. vSphere Replication depends on the IP network to replicate the VMDK of the virtual machines, so the RPO might not be satisfactory if the network bandwidth cannot finish the replicated data within the timeframe and vSphere Replication reports an RPO violation. You should provide more bandwidth or the RPO policy to achieve the RPO goal.
- vSphere Replication replicates the virtual machines from one site or city to another, so the DNS server and IP network may change. The DNS redirection or modification is required to ensure that the virtual machines can work normally after the disaster recovery.
- vSphere Replication cannot ensure primary replica and secondary replica(s) in an availability group to be replicated at the same time. Therefore, if the primary replica cannot be synchronized with the secondary replica(s), the availability group is in out of sync and cannot be recovered. The solution is removing an availability group, then dropping the unsynchronized secondary database and recreating the availability group.

11. SQL Server Database Backup and Restoreon VSAN

vSphere Data Protection supports both vSAN and traditional SAN storage.

11.1 Overview

vSphere Data Protection supports both vSAN and traditional SAN storage. The vSAN can provide a unified data store to backup virtual machines and user databases. The backup configuration is easy to implement through vCenter web client by the vSphere Data Protection plugin.

Deploying the vSphere Data Protection appliance in a VMware environment is simple. Generally, one appliance can satisfy the SQL Server backup in most circumstances. Download the open virtualization appliance (OVA) file through the VMware portal to the tune of about 4.5GB and walk through the quick setup screens to select which datastore to deploy, and apply network settings to make it accessible in the environment. After you deploy the VM, you can set up the appliance through the wizard to complete the configuration, including connecting the appliance to the vCenter, applying the license as well as defining how much resource to allocate to the appliance in terms of CPU, RAM, and storage capacity.

11.2 Installing vSphere Data Protection for SQL Server

To support guest-level backups, you must install vSphere Data Protection for SQLServer Client on each SQL Server for backup and restore support.

11.3 Prerequisites

Before using vSphere Data Protection, you must install and configure the appliance in the VDP Installation and Configuration chapter of the vSphere Data Protection Administration Guide.

Procedures

- On each SQL Server client, access the vSphere Data Protection Web Client.
- You can download the application backup client from the vCenter Web Client to your local SQL Server virtual machine. From Downloads, select the Microsoft SQL Server 64-bit Client.

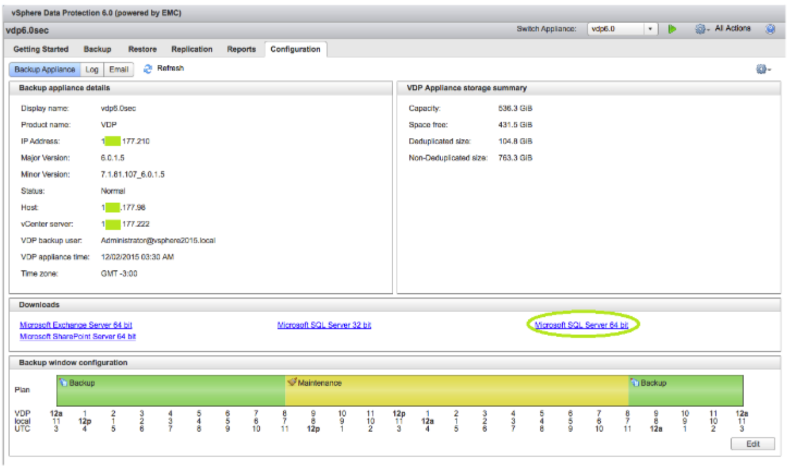

Figure 26. **Download the Application Backup Client**

- In **Job Type**, select **Applications**. In **Data Type**, choose **Selected Databases**. When creating a new backup job, select the type of backup required, which is the user database on the application servers and use the default settings in the **Backup Options**.

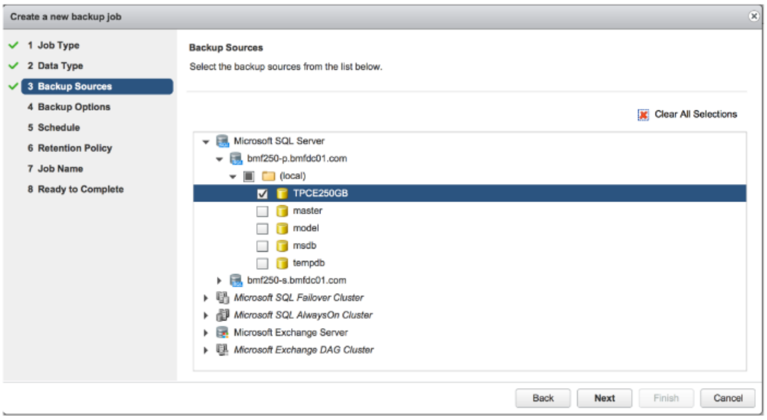

Figure 27. **Creating a Backup Job**

- vSphere Data Protection offers a wide selection of options for scheduling the backup frequency. The software includes the daily, weekly, and monthly backups with time options. From the backup retention perspective, vSphere Data Protection enables users to select a specified date or select the Forever option.

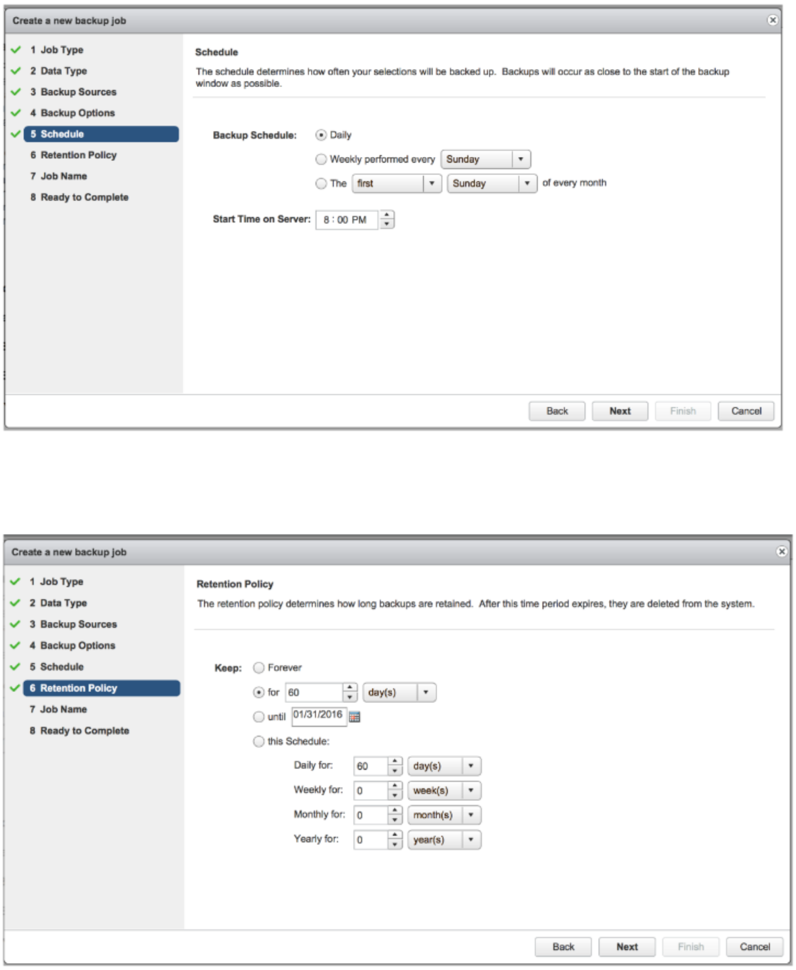

Figure 28. **Scheduling Backup Frequency**

- Name a job and click Finish in the Ready to Complete tab, a new backup job is created.

- Select the backup task and run the restore job directly as shown in Figure 29.

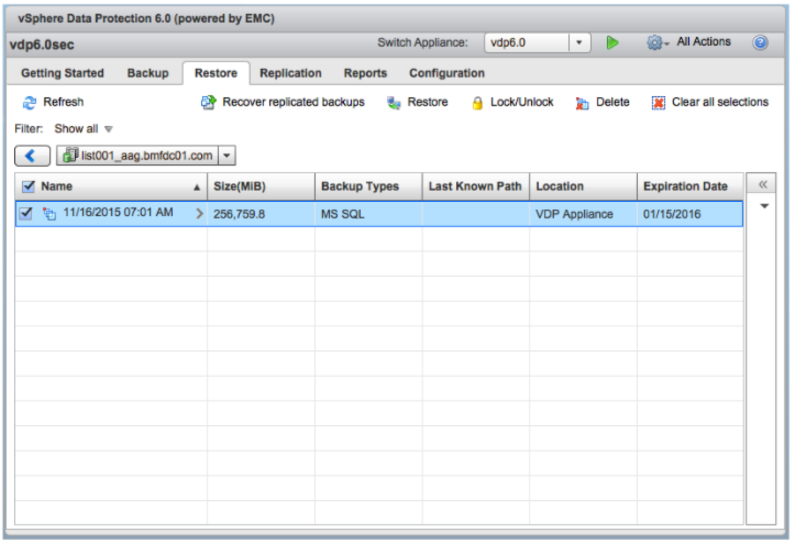

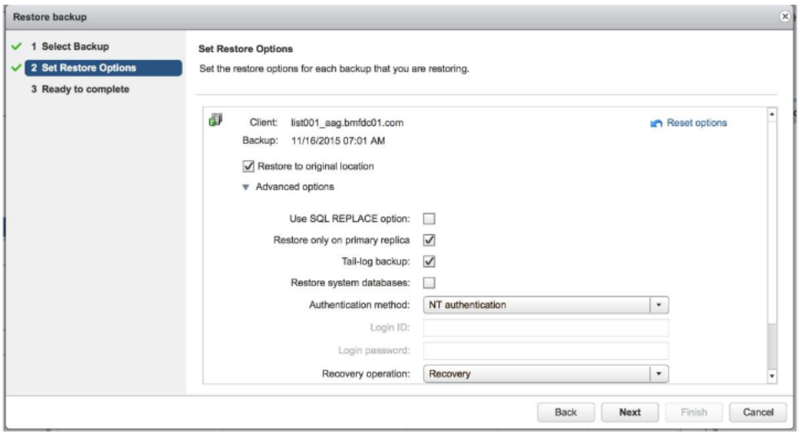

Figure 29. Restoring VMs

11.4 SQL Server Availability Group Backup on VSAN

The SQL Server availability group backup relies on the availability group listener. The listener should be created before the vSphere Data Protection client configuration. Figure 30shows the availability group name is **bmf250** and the listener name is **sqlbmf001**.

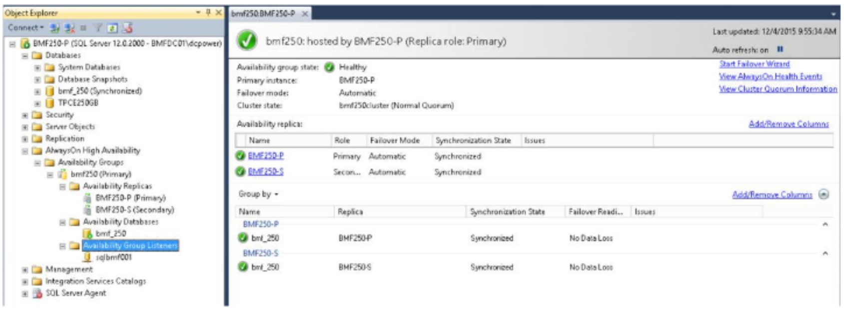

Figure 30. **Microsoft SQL Server Availability Group Dashboard in Management Studio**

To configure the SQL Server availability group backup by using vSphere Data Protection, configure the cluster client in the failover cluster.

- After downloading and installing the cluster client, configure the availability group backed up by vSphere Data Protection Windows Cluster Configuration. Choose SQL AlwaysOn and make sure each node has plugins installed as shown in Figure 31.

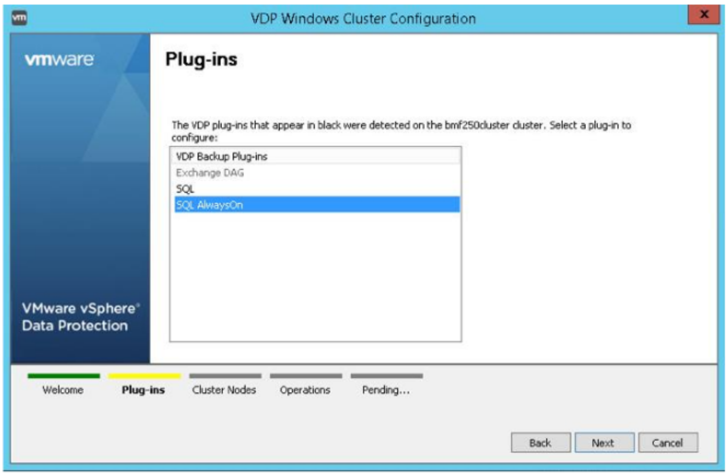

Figure 31. **Select SQL AlwaysOn as the Plugin**

- Select Configure a new cluster client for all nodes as shown in Figure 32.

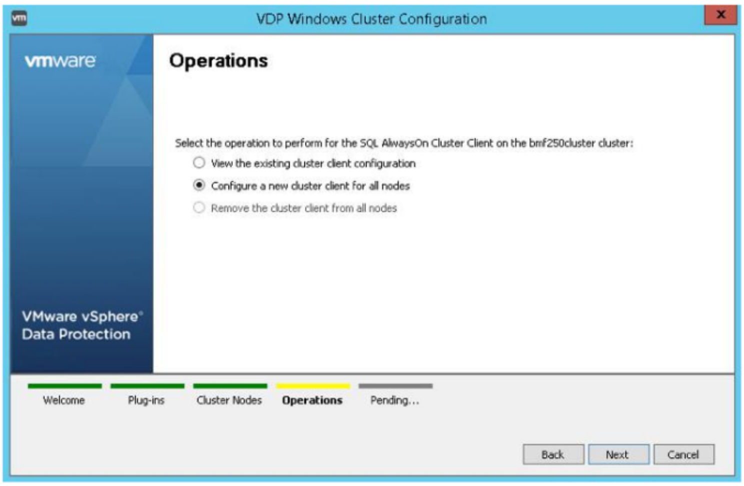

Figure 32. **Configure a New Cluster Client for All Nodes**

- Confirm the Prerequisites have been met.

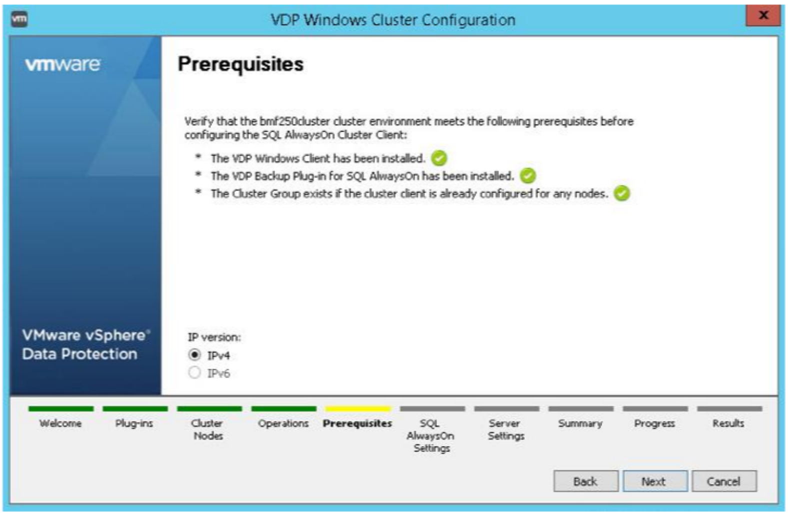

Figure 33. Prerequisites of vSphere Data Protection Windows Cluster Configuration

- Make sure the SQL Server availability group listener is created so that the vSphere Data Protection Windows Cluster Configuration can identify the SQL AlwaysOn cluster. Select the preconfigured availability group name **bmf250** as shown in Figure 34 and click **Next** to finish the configuration wizard.

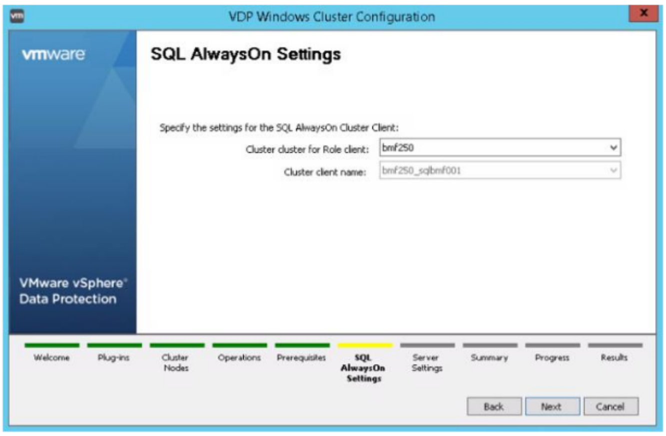

Figure 34. SQL Server AlwaysOn Settings

- Enter the IPV4 address of the vSphere Data Protection Appliance and click Next.

- Complete the client configuration: the configuration was successful as shown in Figure 35.

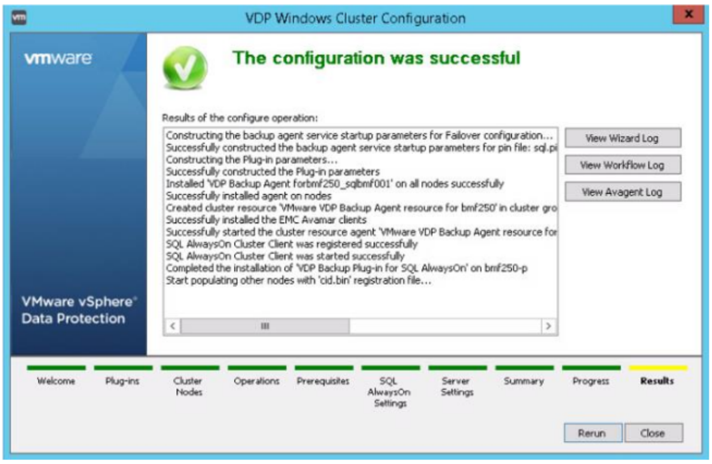

Figure 35. Complete the Client Configuration

11.5 Test Scenario

To validate the vSphere Data Protection database and availability group backup, a backup space—0.5TB with total three virtual disks allocated from the vSAN was configured for the database backup for Microsoft SQL Servers. The backup type for the availability group backup was full. The backup traffic is through the 1GB Ethernet network. After the initial configuration of the three 256GB virtual disks from the vSAN, the usable backup space is around 536GB as shown in Figure 36.

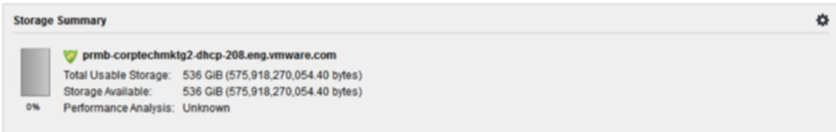

Figure 36. Storage Summary after Configuration

vSAN Storage Validation before Backup

When configuring the vSphere Data Protection Appliance, the performance test produces varying results depending on the size of the appliance being deployed.

The following table lists the minimum expectations for read, write, and seek performance by vSphere Data Protection Appliance size.

VSPHERE DATA PROTECTION APPLIANCE SIZE(TB)	DISK SIZE (GB)	READ MINIMUM (MB/S)	WRITE MINIMUM (MB/S)	SEEK MINIMUM (SEEKS/S)
0.5	256	60	30	400
1.0	512	60	30	400
2.0	1,024	60	30	400
4.0	1,024	80	40	400

6.0	1,024	80	40	400
8.0	1,024	150	120	400

After configuring the storage from vSAN for vSphere Data Protection, run the storage performance test to verify the storage capacity. The performance analysis test creates a 41 GB VMDK on the datastore, and then runs read, write, and seek tests to check the datastore performance.

Table 15 is the performance analysis report of the vSAN as a vSphere Data Protection backend storage.

Table 15. Performance Analysis Report

STORAGE PERFORMANCE RESULTS	ACTUAL	EXPECTED
Minimal Write Throughput (MB/S)	635	30
Minimal Read Throughput (MB/S)	630	60
Minimal Seek Throughput (seeks/s)	16,302	320
Total Write Throughput (MB/S)	732	30
Total Read Throughput (MB/S)	661	60
Total Seek Throughput (seeks/s)	18,491	400

You can see that the 4-node vSAN Cluster can fulfill the storage performance requirement completely.

11.6 VSAN Storage Backup and Restore Validation Result

We created a sample TPC-E like database for the vSphere Data Protection backup efficiency validation by using Benchmark Factory. To validate the backup and restore efficiency of the vSphere Data Protection, a medium-size database was created by Benchmark Factory with scale 25. The created database size was around 252.8GB, including the data, index, and the transactional log. Figure 37 is the actual space usage of the sample database.

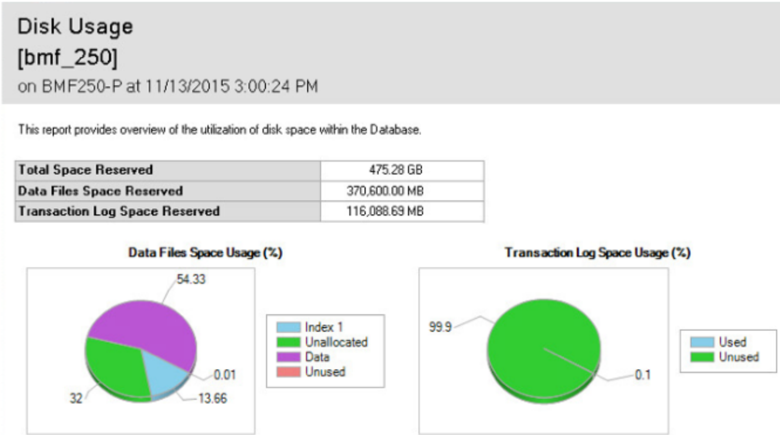

Figure 37. Space Usage of the Sample Database

Table 16 shows the backup duration for the availability groups. These backup sare offline backup without running workloads. Comparing to the SQL Server native backup method, using vSphere Data Protection can leverage dedup lication technology to reduce the actual space usage. Comparing to the native database backup, vSphere Data Protection backup can save 44 percent backup time and 78 percent space usage.

Table 16. Backup Method

BACKUP METHOD	BACKUP DURATION (MM:SS)	BACKUP SPACE(GB)
SQL native	41:42	258.8
vSphere Data Protection for database group	23:42	56.9

Table 17 shows the restore duration from the native backup file and vSphere Data Protection backup set. Comparing to the native database restore, the vSphere Data Protection based restore needs additional connection, status check, and dataset transferring time to finish.

Table 17. Restore Source

RESTORE SOURCE	RESTORE DURATION (MM:SS)
SQL backup file	47:54
vSphere Data Protection data set	20:32

SQL Server database backup from the secondary replica

There are four available backup options for AlwaysOn Availability Groups: primary, prefer secondary, secondary only,and SQL Server defined in vSphere Data Protection. To minimize the backup impact to the running workload, we choose secondary only to measure the backup duration with running workload.

For this specific test purpose, Benchmark Factory was generating workload on the 250GB availability group and weobserved that TPS was between 840 and 950. The vSphere Data Protection backup job was then initialized. Since the backup is initiated from the secondary replica of the availability group, there is no obvious impact to the resources on the primary database running workload. The backup took around 54minutes to finish and the throughput was 80MB/S.

This allows you to offload backup tasks to the secondary replica of the database group while enabling you to run read and write workload against the primary replica of it.

Incremental backup by vSphere Data Protection

TPC-E like workload causes data and log space usage growing. Also the UPDATE transaction in the TPC-E like workload updates the data in the original database. Therefore, when performing incremental backup, the space consumption calculates the differential data generated by the workload. This validation ran eight-hour TPC-E like workload with 100 users in Benchmark Factory to simulate the data changes in one business day, before the differential backup job was triggered. The data changes include the data modification and data growth by the TPC-E like workload.

Changed data set = data modification +TPC-E_like_data_growth +TPC-E_like_log_growth

We ran the backup job manually to measure the space usage. Figure 38 shows the space usage after one business day. After eight hours'run, there were more than 21.9 million transactions executed in the ~250GB database. According to the TPC-E like transaction update percentage (20.9percent), there were around4.3 million updated transactions or data changes in the database.

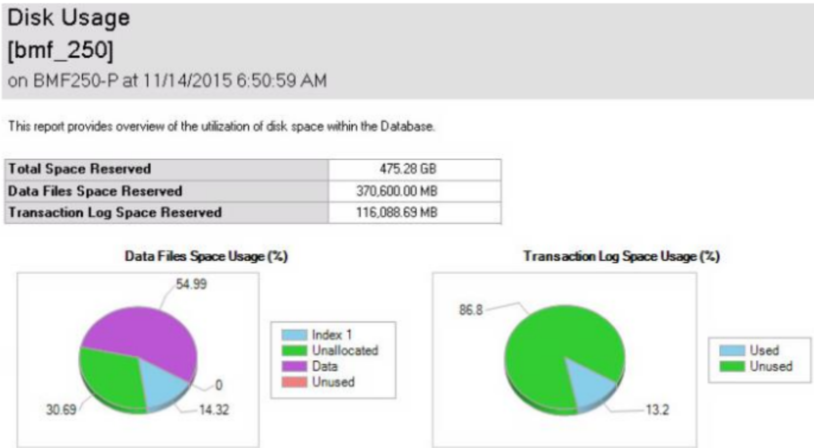

Figure 38. Disk Usage after 8-hour Benchmark Factory Test Run

Table 18 shows the data growth and backup consumption using vSphere Data Protection.

	ORIGINALDB SIZE	AFTER ONE BUSINESS DAY(8-HOUR TPC-E LIKE WORKLOAD)
Data/Index usage (GB)	252.00	256.86
Log usage (GB)	0.12	15.30
Backup space (vSphere Data Protection)	56.90	84.10

11.7 Summary

vSphere Data Protection is designed efficiently and reliably to back up and replicate virtual machines, virtual servers, SQL Server databases, and other applications such as SharePoint and Exchange. The software can perform periodically full backup with efficient space saving. Getting vSphere Data Protection downloaded and deployed by single OVF installed, you can manage vSphere Data Protection through the vCenter interface. Configure backup jobs with a wizard that helps you set policies for target, schedule, retention, and other relevant components.

12. Best Practices of Virtualized SQL Server on VSAN

This section introduces the best practices guide for detailed information about CPU, memory,and networking configurations

12.1 Overview

This section introduces the best practices guide for detailed information about CPU, memory,and networking configurations to implement virtualized SQL Server on vSAN environment.

This section introduces the best practices guide for detailed information about CPU, memory,and networking configurations to implement virtualized SQL Server on vSAN environment.

12.2 vSAN Configuration Guidelines

vSAN is distributed object-store datastore formed from locally attached devices from ESXi host. In a hybrid vSAN configuration, it uses disk groups to pool together flash disks and magnetic disks as single management constructs. It is recommended to use similarly configured and sized ESXi hosts for the vSAN Cluster:

- Plan for capacity—consider initial deployment with the capacity for the future data growth and enough flash cache to accommodate future requirement. Use multiple vSAN disk groups per server with enough magnetic spindles and SSD capacity behind each disk group. For the future capacity addition, create disk groups with similar configuration and sizing. This ensures an even balance of virtual machine storage components across the cluster of disks and hosts.
- Plan for performance—it is important to have sufficient space in the caching tier to accommodate the I/O access of the OLTP application. The general recommendation of the SSD as the caching tier for each host is to be at least 10 percent of the total storage capacity. However, in case where high performance is required for mostly random I/O access patterns, it is recommended that the SSD size should be at least two times of the working set.
  For the SQL Server mission-critical user database, follow the recommendations to design the SSD size:
  - SSD size to cache active user database

    The I/O access pattern of the TPC-E like OLTP is small (8KB dominant), random, and read-intensive. To support the possible read-only workload of the secondary and log hardening workload, it is highly recommended to have two times of the primary and secondary database. For example, for 100GB user database, design 2 x 2 x 100GB SSD size.
  - Select appropriate SSD class to support designed IOPS

    For the read-intensive OLTP workload, the supported IOPS of SSD depends on the class of SSD. A well-tuned TPC-E like workload can have ten percent write ratio. VMware Compatibility Guide specifies the following designated flash. For optimal performance, VMware recommends usinga flash-device class that meets workload performance equirements:
      - Class A: 2,500–5,000 writes per second
      - Class B: 5,000–10,000 writes per second
      - Class C: 10,000–20,000 writes per second
      - Class D: 20,000–30,000 writes per second
      - Class E: 30,000+ writes per second
- Plan for availability—Design more than three hosts and additional capacity that enable the cluster to automatically remediate in the event of a failure.

  For SQL Server mission-critical user databases, enable AlwaysOn to put the database in the high availability state when the AlwaysOn replication is in the synchronous mode. Setting FTT greater than 1 means more write copies to vSAN disks. Unless special data protection is required, FTT=1 can satisfy most of the mission-critical SQL erver databases with AlwaysOn enabled.
- Set proper SPBM—vSAN SPBM can set availability, capacity, and performance policies per virtual machine. Number of disk stripes per object and object space reservation are the storage policy that was changed from the default value for VMs in this reference architecture:
  - Set object space reservation to 100 percent and the capacity is allocated upfront from vSAN datastore.

  - Number of disk stripes per object—The number of disk stripes per object is also referred as stripe width. It is the setting of vSAN policy to define the minimum number of capacity devices across which replica of a storage objects is distributed. vSAN can create up to 12 stripes per object.

    Striping can help performance if the virtual machine is running I/O intensive application such as OLTP database. In the design of a hybrid vSAN environment for a SQL Server OLTP workload, leveraging multiple SSDs with more HDDs backed is more important than only increasing the stripe width. Consider the following conditions:
    - If more disk groups with more SSDs can be configured, setting a large number of the stripe width for a virtual disk can spread the data files to multiple disk groups and improve the disk performance.
    - Larger stripe width number can split virtual disk larger than 255GB into more disk components. However, vSAN cannot guarantee the increased disk components to be distributed across multiple disk groups, with each stored on one HDD disk. If multiple disk components are in the same disk group, the increased components are only the backed HDD number for the virtual disk. Increasing the stripe width may not improve performance unless there is a destaging performance issue.
  - Depending on the database size,it is recommended to have multiple VMDKs for one virtual machine. Multiple VMDKs spread database components across disk groups in vSAN Cluster.
- Create separate disks for data file, log file, and tempdb files on different SCSI controllers.
- Separate the SCSI controllerfor temp db and other databases.
- Use PVSCSI controller and set proper value queue depth. The large-scale workloads with intensive I/O patterns require adapter queue depths greater than the PVSCSI default values. Refer to the Knowledge Base Article 2053145 for more information.

12.3 CPU Configuration Guidelines

This section provides guidelines for CPU configuration for SQL Server database virtual machines. Understand your application requirement for CPU setting:

- Start with a thorough understanding of your workload. Database server utilization varies a lot by topology of application and deployment. If the application is commercial, make sure to follow published guidelines where appropriate. If the application is custom written, work with the application developers to determine the resource requirements. VMware Capacity Planner™ can analyze your current environment and provide resource utilization metrics in the sizing process.
- If you do not know about the exact workload, start with fewer vCPUs and increase the number later if necessary. Only allocate multiple vCPUs to a virtual machine if the anticipated SQL workload can take advantage of all the vCPUs. Over-provisioning vCPUs mightresult in higher virtualization overhead. VMware performance testing for SQL Server has shown that even a single vCPU VM can support high transaction throughput and may be enough for many SQL Server databases.
- For tightly-controlled and highly-consolidated production environments, makesure to account for some virtualization overhead (8percent to15 percent depending on the workload). If you leave enough CPU headroom to account for this additional CPU utilization, you can achieve transaction throughput rates similar to that of physical SQL Servers.
- When consolidating multiple SQL Server virtual machines on single ESXi host, proper hardware sizing is critical for optimal performance. Make sure that cumulative physical CPU resources on a host are adequate to meet the needs of the guest VMs by testing your workload in the planned virtualized environment.

Plan the CPU over-commitment carefully around actual performance data to avoid adversely affecting VM performance. For performance-critical SQL Server virtual machines (production systems), make sure that sufficient CPU resources are available, and CPU over -commitment is not the reason for any performance degradation.
- Install the latest version of VMware Tools in the guest operating systemand update VMware Tools after each ESXi upgrade.
- Use NUMA—VMware recommends keeping Non-Uniform Memory Architecture (NUMA) enabled in server hardware BIOS and at the guest-operating system level. When Creating SQL Server virtual machines,create as many virtual sockets as you request and setthe cores per socketequal to one. This enables vNUMA to select and present the best virtual NUMA topology to the SQL Server operating system.

12.4 Memory Configuration Guidelines

This section provides guidelines for memory allocation for SQL Server virtual machines.

Because SQL Server workloads can be memory-intensive and performance is often a key factor (especially in production environments), VMware recommends the following practices:

- Start with a thorough understanding of your workload. The amount of memory you need for a SQL Server virtual machine depends on the database workload you plan to host in that virtual machine.
- Database performance is heavily dependent on the amount of memory available. A common tuning technique is to increase the database buffer cache to reduce or avoid disk I/O and thus improve SQL Server performance. vSphere6.0 supports up to 4TB memory per virtual machine. This enables you to configure SQL Server databases with large in-memory caches for better performance.
- Make surethat cumulative physical memory available on a server is adequate to meet the needs of the virtual machines by testing target workloads in the virtualized nvironment. Memory over-commitment should not adversely affect virtual machine performance as long as the actual virtual machine memory requirements are less than the total memory available on the system. For performance-critical SQL Server virtual machines (production systems), make sure that sufficient memory resources are available and that memory over-commitment is not causing any performance degradation. If memory is overcommitted on the host, set reservations for performance-critical SQL Server virtual machines to guarantee that the memory for those virtual machines is not ballooned or swapped out.
- If you set the SQL Server Lock Pages in Memory parameter, set the virtual machine's reservations to match the amount of memory you set in the virtual machine configuration. This setting can interfere with the ESXi balloon driver. Setting reservations stops the balloon driver from inflating into the virtual machine's memory space.
- vSphere supports Large Page Performance in the guest operating system. If the operating system or application can benefit from large pages on a native system, operating system or application can potentially achieve a similar performance improvement in a virtual machine. The use of large pages results in reduced memory management overhead and can therefore increase hypervisor performance. Refer to Microsoft SQL Server on VMware Best Practices Guide for detailed information.

12.5 Networking Configuration Guidelines

This section covers design guidelines for the virtual networking environment. The standard VMware networking best practices apply to running SQL Server databases on vSphere:

- Separate VLANs are recommended per traffic type. Uses eparate infrastructure traffic from VM traffic for security and isolation.

- Use the VMXNET3 network adapter. The VMXNET3 adapter is optimized for virtual environments and is designed to provide high performance.
- Take advantage of Network I/O Control to converge network and storage traffic (vSAN) onto 10GbE. Network I/O Control in vSphere enables you to guarantee service levels (bandwidth) for particular vSphere traffic types: VM traffic, vSAN, management, backup, and vSphere vMotion.
- Use NIC teaming for availability and load balancing. NIC teaming occurs when multiple uplink adapters are associated with a single vSwitch to form a team.

Refer to the VMware vSAN Network Design Guide for more information.

12.6 Conclusion

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

This solution validated that vSAN is a unified storage platform supporting scalable, high performed, and AlwaysOn protected Microsoft SQL Server database groups with industry-standard OLTP workload.

Several scenarios including the single disk, diskgroup, and host failures proved that vSAN can have high availability to be the backend of Tier-1 application with the running workload.

vSAN Stretched Cluster extends the existing vSAN storage to geographically dispersed data centers through a high bandwidth and low latency link. And vSAN Stretched Cluster supports acceptable OLTP workload.

Site Recovery Manager and vSphere Replication with vSAN Cluster and Stretched Cluster provide a cost-effective DR solution for SQL Server AlwaysOn Availability Groups.

vSphere Data Protection providesa simple and easy-to-manage solution to the Microsoft SQL Server databases and can support various backup granularities including database, availability groups, application server, and image-level backup and restore. Full database group backup and restore with deduplication efficiency powered by vSphere Data Protection are validated in this solution.

13. Reference

Reference

13.1 Reference

The following documents are the reference for this solution. For more information, visit the VMware vSAN Resources Website.

White Paper

For additional information, see the following white papers:

- VMware vSAN 6.0 Design and Sizing guide
- VMware vSAN Performance with Microsoft SQL Server
- VMware vSAN Backup Using VMware vSphere Data Protection Advanced
- VMware vSAN Network Design Guide
- VMware vSAN 6.0 Performance Scalability and Best Practices
- Microsoft SQL Server on VMware Best Practices Guide

Product Documentation

For additional information, see the following product documentation:

- vSAN Hardware Quick Reference Guide
- Monitoring VMware vSAN with vSAN Observer
- vSphere Data Protection Administration Guide

Other Reference

- VMware Compatibility Guide

About the Author and Contributors

Tony Wu, Solution Architect in the vSAN Product Enablement team wrote the original version of this paper. Catherine Xu, Technical Writer in the vSAN Product Enablement team, edited this paper to ensure that the contents conform to the VMware writing style.

Contributors to this document include:

- Christian Rauber, Solution Architect in the vSAN Product Enablement team
- Bradford Garvey, Solution Architect in the vSAN Product Enablement team