Microsoft SQL Server on VMware vSphere® Availability and Recovery Options

Version 1.0

May 2016

VMware, Inc.
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Contents

List of Figures

1. Introduction

Running Microsoft SQL Server on VMware vSphere® offers many options for database availability and disaster recovery utilizing the best features from both VMware and Microsoft. For example, VMware vSphere vMotion® and VMware vSphere Distributed Resource Scheduler™ (DRS) can help to reduce planned downtime and balance workloads dynamically, and VMware vSphere High Availability (HA) can help to recover SQL Server databases in the case of host failure.

At the application level, all SQL Server features and techniques are supported on vSphere, including *AlwaysOn Availability Groups*, *AlwaysOn Failover Cluster Instances*, *database mirroring*, and *log shipping*. SQL Server availability features can be used inside of a virtual machine just as you would on a physical server, and combined with vSphere features to support flexible availability and recovery scenarios, applying the most efficient and appropriate tools for different use cases.

AlwaysOn is a collection of high availability and disaster recovery features introduced in SQL Server 2012 that minimize Recovery Point Objective (RPO) and Recovery Time Objective (RTO), and maximize availability of databases. SQL Server AlwaysOn branding encompasses both FCIs and AGs.

The following table lists SQL Server availability options and their ability to meet RTOs and RPOs. Before choosing any option, evaluate your business requirements to determine which ones best meet your specific needs.

Table 1. SQL Server High Availability Options

Technology	Granularity	Storage Architecture[1]	RPO – Data Loss	RTO – Downtime
AlwaysOn Availability Groups	Database	Non-shared	None (with synchronous commit mode)	Seconds or administrator initiated recovery
AlwaysOn Failover Cluster Instances	Instance	Shared	None	Seconds to minutes depending on the recovery process
Database Mirroring (Deprecated in SQL Server 2012)	Database	Non-shared	None (with high safety mode)	Seconds or administrator initiated recovery
Log Shipping	Database	Non-shared	Possible loss for yet to be transmitted transaction log	Administrator initiated recovery

1. There can be a combination of architectures, for example, AG with FCI, in which a combination of the storage architecture of both clustering technologies will apply.

1.1 Purpose

This paper details the options and guidelines for deploying highly available Microsoft SQL Server on vSphere. The recommendations included in this guide are not specific to any particular hardware or to the size and scope of any particular SQL Server implementation. The examples and considerations in this document provide guidance but do not represent strict design requirements, because the flexibility of SQL Server on vSphere allows for a wide variety of valid configurations.

1.2 Terminology

Table 2. VMware Terminology

Terminology	Description
vSphere	The VMware cloud computing virtualization platform
vSphere cluster	A collection of VMware ESXi™ hosts and associated virtual machines intended to work together as a unit. VMware features, such as Enhanced vMotion Compatibility (EVC), DRS, VMware vSphere Distributed Power Management (DPM), vSphere HA, VMware Virtual SAN™, are enabled on a per cluster basis.
VMware Site Recovery Manager™	Site Recovery Manager is a disaster recovery solution that provides automated failover, site migration, and disaster recovery testing.
vSphere HA	vSphere High Availability (HA) is a cluster-level high availability solution for applications running in virtual machines.
DRS	DRS is a feature that aggregates computing capacity in a vSphere cluster into logical resource pools and intelligently allocates available resources among the virtual machines based on predefined rules.
VMware vSphere Fault Tolerance (FT)	vSphere FT provides continuous availability for applications in the event of physical server failures by creating a live shadow instance of a virtual machine that is always up-to-date.
VMware vSphere Storage DRS™	vSphere Storage DRS continuously balances VM virtual disk storage usage and storage I/O load across a VMware vSphere VMFS storage cluster while avoiding resource bottlenecks to meet application service levels.
vSphere vMotion	vSphere VMotion enables the live migration of running virtual machines from one physical server to another with zero downtime.
VMware vSphere Storage vMotion	vSphere Storage vMotion enables live migration of virtual machine disk files within and across storage arrays with no downtime or disruption in service.
vSphere VMFS	A VMware proprietary high-performance clustered file system for ESXi virtual machines.

Table 3. SQL Server Terminology

Terminology	Description
AlwaysOn	AlwaysOn is a collection of high availability and disaster recovery functionality that includes AlwaysOn AG and AlwaysOn FCI.
AlwaysOn Availability Group (AG)	AlwaysOn AG is a high availability and disaster recovery solution. AG is the logical successor for database mirroring.
AlwaysOn Failover Cluster Instance (FCI)	AlwaysOn FCI leverages Windows Server Failover Clustering (WSFC) functionality to provide local high availability through redundancy at the instance level.
Windows Server Failover Cluster (WSFC)	Grouping of Windows Servers into a cluster to support the availability for application and services. AlwaysOn AG and FCI are built on top of WSFC.
Microsoft Cluster Services (MSCS)	MSCS is the precursor technology to WSFC.
Database Mirroring (DBM)	Database mirroring provides database-level protection by mirroring transactions from a principal to a single secondary mirror.

Table 4. General Terminology

Terminology	Description
Recovery Time Objective (RTO)	Maximum time that an application can be down in case of a disaster or a failure.
Recovery Point Objective (RPO)	Maximum data that is tolerable to lose in case of a disaster or failure, which also defines the interval between backups or replications.
Service Level Agreement (SLA)	A defined and agreed to service level for an application or a system. Typically refers to response times and availability levels.
Disaster Recovery (DR)	The ability to recover an application, system, or data center in case of a disaster. Typically involves a remote data center.
Disaster Recovery Plan (DRP)	A set of procedures for recovering an application, system, or entire data center in case of a disaster.
Disaster Avoidance (DA)	The ability to migrate seamlessly from one location to another before a disaster strikes.

2. VMware vSphere Platform Availability Benefits

Application-level clustering has been the prevalent solution chosen for most SQL Server implementations. Features of the vSphere platform can enhance the overall availability of SQL Server by providing options that help to limit both planned and unplanned downtime. In fact, for some databases, the features provided by vSphere might satisfy the availability requirements of the business without needing to follow traditional clustering approaches. For databases that have stringent SLA requirements, FCIs or AGs can be used in a virtual machine and combined with vSphere features to create an extremely flexible environment, with options for failover and recovery at both the hardware and application levels. Some availability benefits of the vSphere platform include:

- Virtual machine characteristics:
  - Made out of a set of files – The fact that virtual machines are encapsulated as files means that they can be cloned, moved, and ported easily.
  - Hardware independent – This means that your SQL Server installation is no longer bound to a particular piece of hardware, which means increased flexibility when designing both production and disaster recovery components. Both standalone as well as clustered configurations can be virtualized, eliminating the need for identical physical hardware in production and DR sites.
  - Isolation – Virtual machines are isolated from other virtual machines as well as from the host they are running on. This means that an SQL Server instance running in a virtual machine is isolated from other virtual machines on the same physical server, which allows for greater operational efficiency and security of consolidated virtual SQL Server deployments. This also provides operational efficiency because you can patch or upgrade databases without affecting other virtual SQL Server instances on the same physical server. Security is also enhanced because if one virtual machine with an SQL Server installed on it is compromised, the other SQL Server instances on the same physical machine will not be affected.

    These characteristics can enhance availability as follows:
  - Design decisions are no longer permanent – You can adjust your CPU and memory requirements on a running virtual machine with hot add or through a quick configuration change and reboot of the virtual machine. For best practices in using hot add capability, see the *Architecting Microsoft SQL Server on VMware vSphere Best Practices Guide* (https://www.vmware.com/files/pdf/solutions/SQL_Server_on_VMware-Best_Practices_Guide.pdf).
  - Easily upgrade to newer hardware – As your application requirements change, requiring new hardware, or when hardware needs to be refreshed, you can easily and seamlessly move the SQL Server virtual machine to the newer hardware to accommodate increased workloads with a click of a button.
  - Disaster avoidance – Seamlessly migrate your SQL Server virtual machine between physical servers, storage systems, networks, and data centers with vSphere vMotion and vSphere Storage vMotion, to avoid disaster with no disruption to service.
- vSphere built-in availability features called vSphere HA and vSphere FT protect your SQL Server installation from hardware failure with no additional configuration on the application side. If your physical server or any critical physical component within the physical server fails for any reason, these features will recover your SQL Server running on a virtual machine on another physical server in the following manner:
  - vSphere HA will automatically reboot the SQL Server virtual machine on another physical server, acting as a first line of defense against service outage.
  - vSphere FT protects your SQL Server continuously from hardware failure with no disruption to service.

- By combining vSphere HA or vSphere FT with traditional clustering approaches, such as SQL FCIs and AGs, you can mitigate both hardware and software failures for maximum availability.
- Utilizing vSphere DRS to balance workloads can speed recovery. As application workloads increase, DRS can move a bottlenecked virtual machine automatically (and without downtime) to another host that has more available resources. DRS can also help to recover more quickly after server hardware failure. For example, when a physical server fails, vSphere HA reboots the virtual machine on another physical server. When the failed server is replaced, DRS can migrate the virtual machine back to its original location with no downtime and no interruption to the end user.
- vSphere Storage DRS allows for intelligent initial placement and ongoing space and load balancing of virtual machines based on disk latency and storage capacity. vSphere Storage DRS continuously monitors storage space and I/O utilization across a pre-assigned pool of datastores and intelligently aligns storage resources to meet your business growth objectives. With Storage DRS, you can specify how storage resources are allocated to virtual machines using rules and policies. When one or more datastores in a datastore cluster exceeds the user-configurable space utilization or I/O latency thresholds, Storage DRS uses Storage vMotion to move one or more virtual machine disk files to achieve its goals without interruption to the SQL Server instance.

3. Increase Availability Across the SQL Server Lifecycle

Traditional, physical environment upgrades and scale-up activities require significant resources, including:

- Planning and implementation time from engineering resources involving application administration, server administration, and SAN administration.
- Sizing and acquisition of new hardware.
- Effort required to perform a hardware upgrade, which results in higher costs and risks.

In comparison, using vSphere alleviates these challenges as described in the following sections.

3.1 Design, Scale, and Upgrade your SQL Server Deployments

Scaling your environment using vSphere can be achieved by adding more SQL Server virtual machines or adding resources “on the fly” to existing SQL Server virtual machines as workload requirements increase. This also provides greater flexibility when designing a new SQL Server deployment because you no longer need to add extra resources to the SQL Server to protect from sizing errors or future growth. As scaling becomes a simpler task, the pre-deployment SQL Server design also becomes much easier.

In addition, the physical servers running virtualized SQL Server environments can be serviced without interruption to the applications using the SQL Server instances residing on that physical server. For example, when adding more memory or upgrading CPU on a physical server, all the virtual machines can be moved seamlessly to other physical server using vSphere vMotion for the duration of the maintenance operation.

3.2 Storage Maintenance and Upgrade

Physical SQL Server environments are tightly bound to the underlying storage and are less flexible than virtualized ones. This is because vSphere virtualizes storage to a simple SCSI device. That SCSI device (which is basically a file) can be migrated to another storage device with no disruption to the SQL Server virtual machine or the connections to database. The end result is that the SQL Server environment can be upgraded or migrated seamlessly without regard to underlying storage technologies. With vSphere, the storage capacity serving the SQL Server instances can also be increased dynamically with the hot add/remove storage functionality.

With built-in multipathing capability and advanced queuing techniques available for virtual machines on vSphere, virtualized SQL Server environments can leverage advanced configuration options to:

- Control, limit, and guarantee IOPS to virtual machines to provide better performance and service more clients.
- Balance the workloads of multiple SQL Server virtual machines sharing the same physical server to use multiple SAN paths and storage processor ports.

4. VMware Availability Options

When deploying SQL Server, the availability options that exist and are used in physical environments continue to be available in a virtual environment. vSphere features that provide high availability and load distribution can be used to increase availability, improve performance, and speed recovery.

These features, including vSphere HA, vSphere FT, vSphere vMotion, and vSphere Storage vMotion, provide inherent protection and availability in the local site and between sites by placing the SQL Server instance in a virtual machine. This provides low effort and low maintenance protection from planned and unplanned hardware downtime.

Note SQL Server availability options are explained further in Section 5, SQL Server Availability Options.

4.1 VMware Local Data Center Availability Options

The following sections describe using vSphere HA, vSphere DRS, and vSphere vMotion in local site high availability scenarios.

4.1.1 vSphere HA

vSphere HA provides easy-to-use and cost-effective high availability for applications running in virtual machines. In a cluster that has HA enabled and configured, in the event of an unplanned physical server failure, affected virtual machines are automatically restarted on other production physical servers that have spare capacity. Additionally, if there is an operating system failure within a virtual machine, the failure is detected by vSphere HA and the affected virtual machine is restarted on the same physical server. With vSphere HA, after a vSphere cluster is set up and HA is enabled, all that is required to protect your SQL Server virtual machine is to place it in the cluster. No additional configuration is needed on the vSphere or the SQL Server side.

Figure 1. vSphere HA

Example 1: Standalone SQL Server virtual machine with vSphere HA

vSphere HA automatically restarts your SQL Server virtual machine on another server if the current one fails, so your virtual machine can be restored to normal operation in the time that it takes to reboot the operating system and start the SQL Server services. Because no extra configuration is required, vSphere HA protects a standalone SQL Server installation without the need to configure and maintain a Microsoft application-level clustering solution.

Example 2: Virtualized AlwaysOn FCI or AG combined with vSphere HA

Implementing a clustered implementation of SQL Server (either an AG or an FCI) on vSphere is fully supported and increases the availability levels of the solution as a whole. For example, when a physical Microsoft cluster node fails due to a hardware fault, the WSFC will move the FCI to another node, and the instance will restart automatically. With an unplanned hardware failure, your SQL Server environment

could be vulnerable to further failures during the time between loss of a replica and its restoration. Resynchronizing the replica could take a substantial amount of time. vSphere HA helps to alleviate these issues by restarting the failed replica on another available host in the vSphere cluster. This quickly restores full protection to the database and reduces the amount of time spent in the failover state.

Figure 2. vSphere HA / Administrator Recovery with WSFC

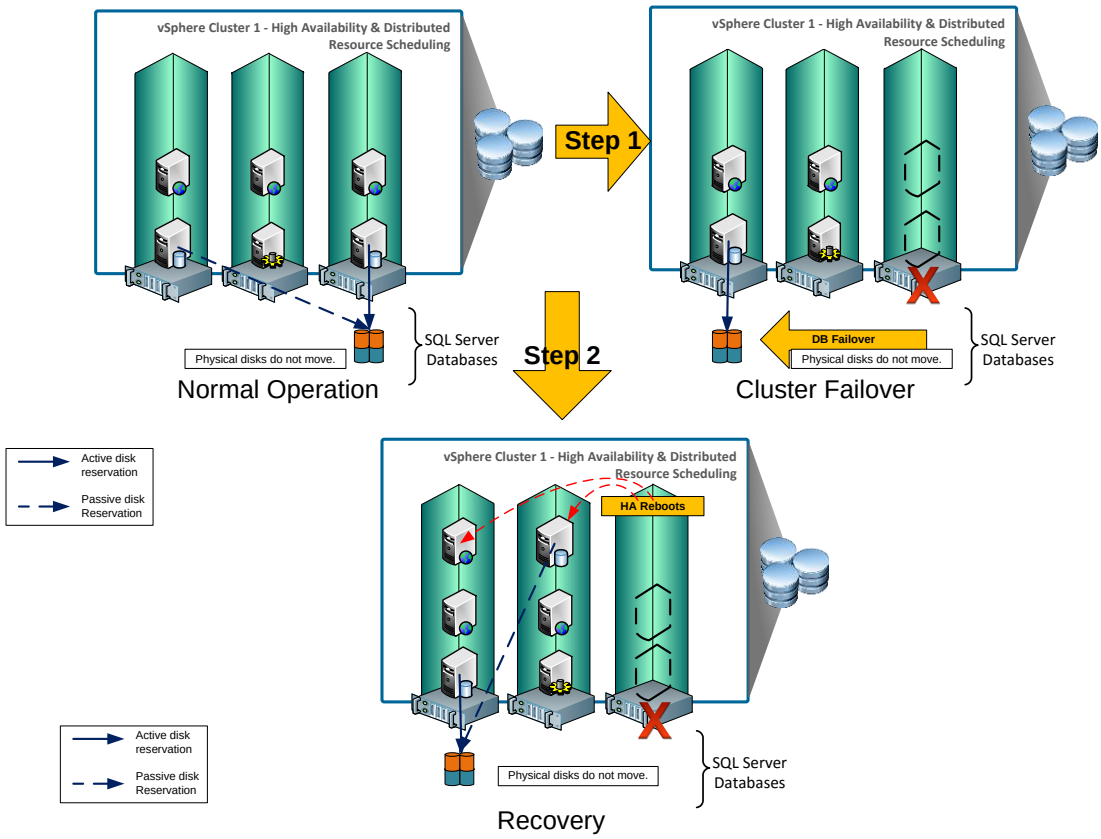

4.1.2 vSphere DRS

vSphere DRS collects resource usage information for all hosts and virtual machines, and generates recommendations for virtual machine placement. These recommendations can be applied manually or automatically. vSphere DRS can dynamically load balance all virtual machines in the environment by shifting workloads across the entire pool of vSphere hosts. This provides critical SQL Server virtual machines in the environment with the CPU and RAM resources needed to maintain optimal performance.

After the original server hardware is fixed or replaced, DRS and vSphere vMotion can be used to quickly move the virtual machine back to its original vSphere host with no additional downtime.

Figure 3. vSphere DRS with vSphere vMotion

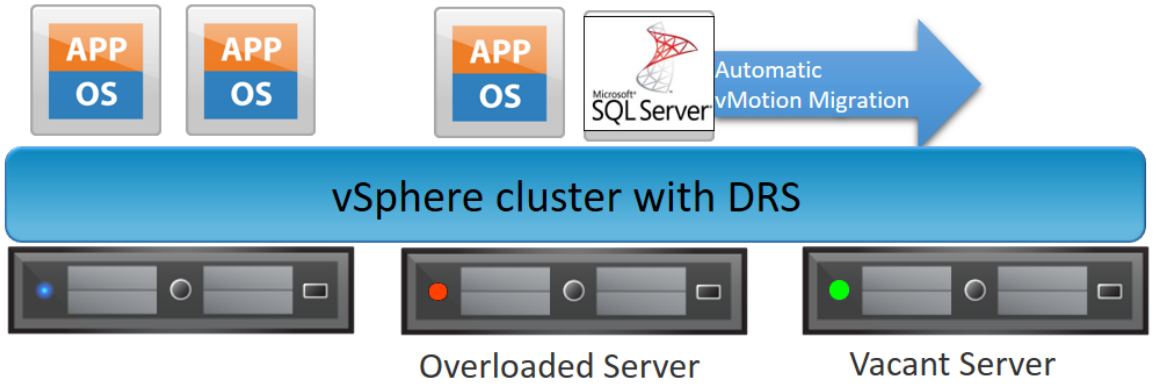

4.1.3 vSphere vMotion and vSphere Storage vMotion

vSphere vMotion leverages the complete virtualization of servers, storage, and networking to move a running virtual machine from one physical server to another. The two physical servers can be in the same data center or in remote data centers with supported RTT of up to 150 ms between source and destination. This migration is performed with no impact to running workloads or connected users. During a vSphere vMotion migration, the active memory and execution state of the virtual machine is rapidly transmitted over the network to the new physical server, while maintaining its network identity and connections. If the source and destination servers do not share storage, storage is also migrated non-disruptively.

With vSphere Storage vMotion, an administrator initiates a migration of a virtual machine's virtual disk between source and destination storage arrays without any disruption to service. Storage vMotion is fully storage-agnostic, which means the migration can be done between any supported storage systems in vSphere. When Storage vMotion is invoked, a shadow disk is created on the destination storage and an initial copy is performed. During that copy, any new block is written to both the source and destination disks. When the two disks are mirrored, a fast suspend and resume is done to transfer the control to the new virtual disk.

The following are the potential use cases for vSphere vMotion and vSphere Storage vMotion seamless migration of SQL Server:

- Migration with vSphere vMotion to allow physical server maintenance
- Migration for server load balancing purposes
- Migration to new purchased server hardware
- Migration between data centers for disaster avoidance
- Migration to a new storage system
- Migration to allow storage maintenance
- Migration for storage capacity and/or I/O load balancing

Figure 4. vSphere HA and vSphere vMotion with AlwaysOn Availability Groups

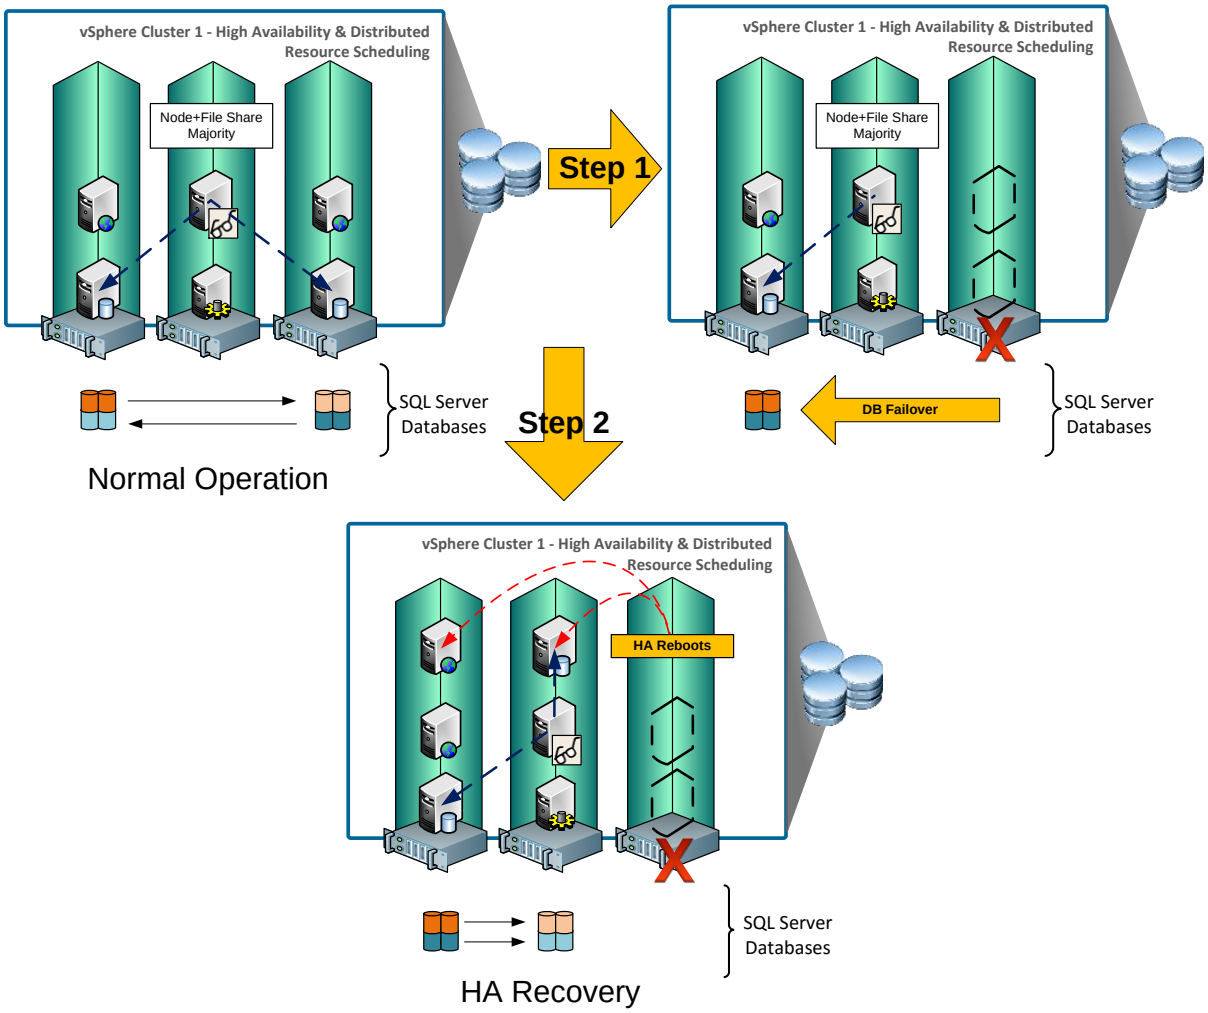

4.1.4 vSphere HA Application-Aware APIs

Starting with vSphere 4.1, an application programming interface (API) provides third-party vendors with the ability to integrate with vSphere HA. Symantec was the first partner to develop an agent for providing application awareness within a vSphere cluster. The capability to allow application monitoring agents to interact with vSphere HA is enabled per vSphere cluster, with additional configuration options available per virtual machine. When enabled, this feature allows application monitoring agents to send application heartbeats to vSphere HA. In the event of an application-level failure, the application monitoring agent can take action either to bring the application back online or to stop the application heartbeat, causing vSphere HA to initiate a restart of the virtual machine.

4.1.5 Multi-Processor vSphere FT

vSphere Fault Tolerance provides business continuity with higher levels of availability and data protection than is offered by vSphere HA. vSphere FT is built into the ESXi host platform. vSphere FT provides continuous availability for a virtual machine by creating and maintaining another VM that is identical and continuously available to replace it in the event of a failover situation. The primary VM execution is identical to the secondary VM. If the host where the protected VM fails, the secondary VM immediately

takes over with no disruption to the VM, and a new secondary VM will be immediately spawned on a third server. At the time of writing of this paper, vSphere FT supports currently up to four virtual CPUs assigned to a protected VM.

Figure 5. vSphere FT

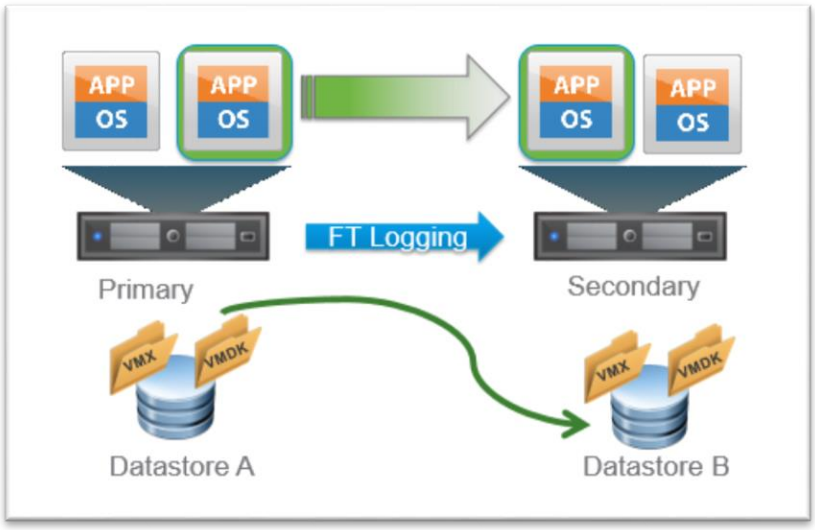

For more information about supported high availability configurations in vSphere, see the Knowledge Base article, *Supported vCenter Server High Availability Options*, at http://kb.vmware.com/kb/1024051.

4.2 VMware Remote Data Center Availability Options

Incorporating disaster recovery and disaster avoidance has become a leading business objective for new deployments and upgrades. For protecting a single business critical application from disaster, it is ideal to leverage built-in solutions. Some environments require application- and hardware-agnostic methods for protecting their mission-critical applications and data. Deploying on vSphere provides the flexibility to meet all of these requirements.

4.2.1 Long-Distance vSphere vMotion Migration

Starting with vSphere 6, vSphere vMotion supports seamless migrations not only between two servers in the same data center, but also for cross-continental distances with up to 150 ms round trip time. This capability further enhances the mobility of vSphere VMs and can be used for many new use cases, such as seamless data center migrations, data center upgrades, large-scale distributed resource management, and disaster avoidance.

DA or disaster avoidance is one prevalent use case for long-distance vSphere vMotion. For example, a hurricane that is targeted to hit the production data center, long-distance vSphere vMotion can be used to migrate the production SQL Server VMs to a remote data center with no disruption of service.

Typically, when there are multiple remote data centers, each data center is managed by a separate VMware vCenter®. So, in many cases a long-distance vMotion migration will involve cross vCenter vSphere vMotion, which is supported in vSphere 6 and later as well.

Figure 6. Long-Distance vSphere vMotion

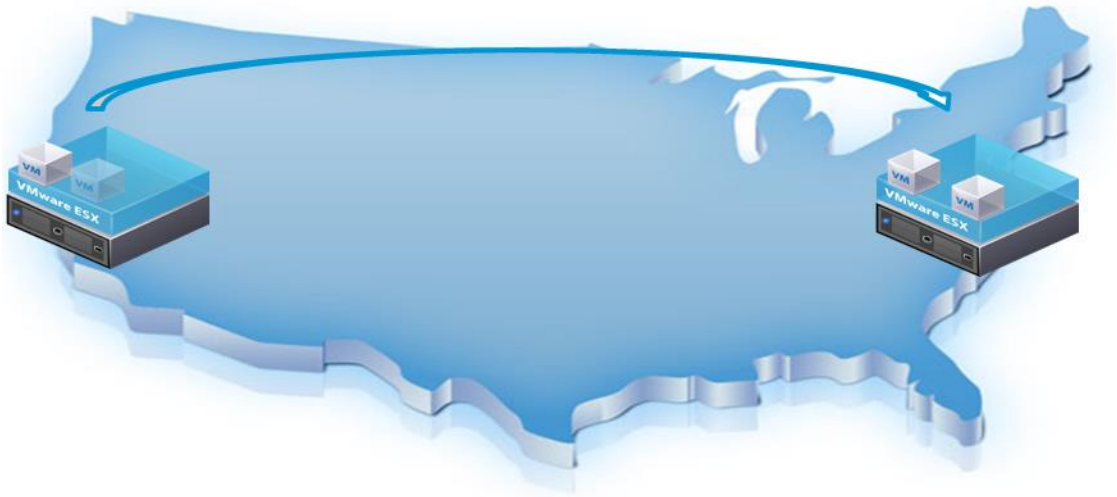

For a list of requirements for long-distance vSphere vMotion, see *Long Distance vMotion Requirements in VMware vSphere 6.0* (http://kb.vmware.com/kb/2106949).

For additional information regarding the requirements for cross vCenter vSphere vMotion, see *Cross vCenter Server vMotion requirements in VMware vSphere 6.0* (https://kb.vmware.com/kb/2106952).

4.2.2 Site Recovery Manager

Site Recovery Manager makes disaster recovery rapid, reliable, manageable, and affordable. By using VMware vSphere Replication™ or third-party storage-based replication technology, Site Recovery Manager delivers centralized management of recovery plans, enables automation of the recovery process, and dramatically improves testing of recovery plans. It transforms the complex hardcopy runbooks associated with traditional disaster recovery into an integrated element of virtual infrastructure management. Site Recovery Manager enables organizations to take the risk and worry out of disaster recovery—yet another reason the VMware virtualization platform is the safest platform for data center applications.

Figure 7. Site Recovery Manager

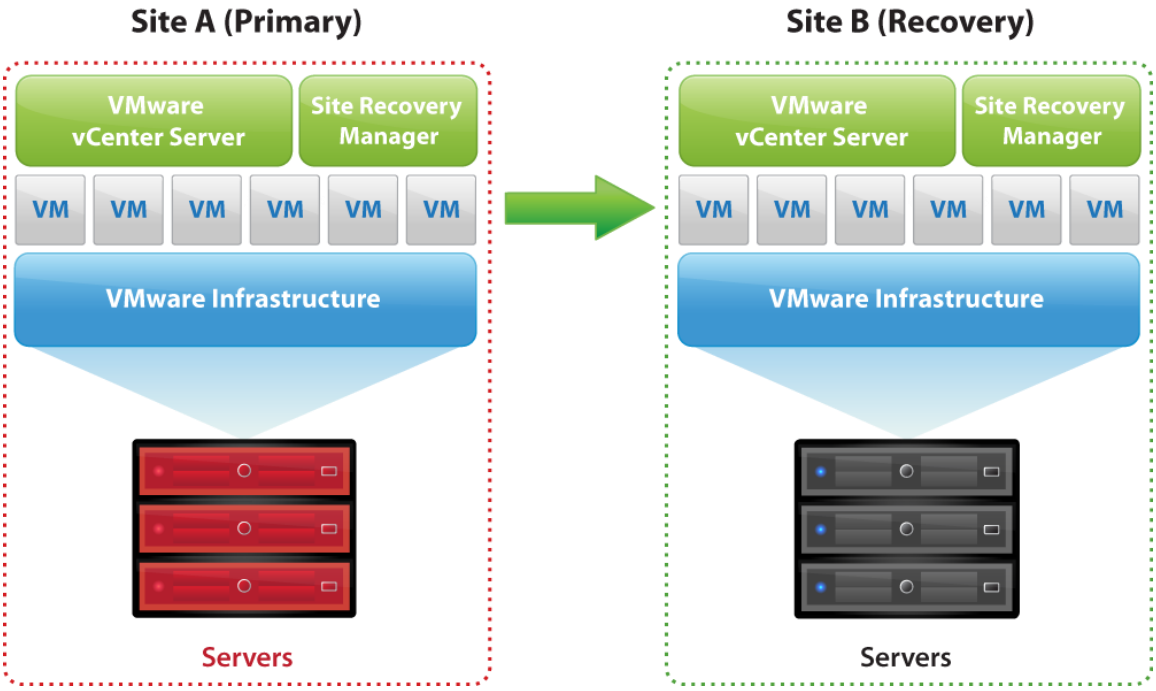

Using SQL Server AlwaysOn AG and FCI solutions within the local data center to provide high availability might meet the requirements of most organizations. FCIs provide failover at the instance level and the failure domain for an AG is the database. Failovers for both AlwaysOn features can take place within a few seconds to a few minutes of a detected failure depending on the feature and the conditions. When designing a disaster recovery solution, automated failover is usually not a desirable feature. In many cases, a DR facility is designed with a lower SLA and has a slightly delayed version of data than the production facility. Making the choice to activate the DR facility should be a conscious decision that follows an organization's change process. With a Site Recovery Manager and storage replication solution, disaster recovery can be implemented to protect the entire virtual data center, including SQL Server. Site Recovery Manager provides you with the flexibility to customize the recovery plan, and VMs running an AlwaysOn instance can be restored on a secondary data center with simple scripting. Failover testing can be accomplished with no production impact to confirm that the RTOs and RPOs will be met in case of a real disaster.

5. SQL Server Availability Options

When deploying SQL Server, the high availability options that are available and used in physical environments continue to be available in a virtual environment inside a virtual machine. This section discusses the different options and how they can be used for local site or remote site availability.

5.1 AlwaysOn Failover Cluster Instances

AlwaysOn FCI is an SQL Server high availability solution that is built on a shared storage architecture. A single instance of SQL Server is installed across multiple WSFC nodes within a local network or across multiple subnets. A single copy of the data is shared among all nodes. The SQL Server instance can run on a single node within the cluster at any point in time. FCI is built on top of a Windows Server Failover Cluster (WSFC) to provide high availability protection for SQL Server at the instance level. This means that everything in that instance including the databases, SQL Server Agent jobs, users, logins, and so on, will move automatically should there be a failure.

During a failover, the node owning the FCI relinquishes control of the resources and another node can claim it and start the FCI as shown in the following figure.

Figure 8. Microsoft Failover Clustering

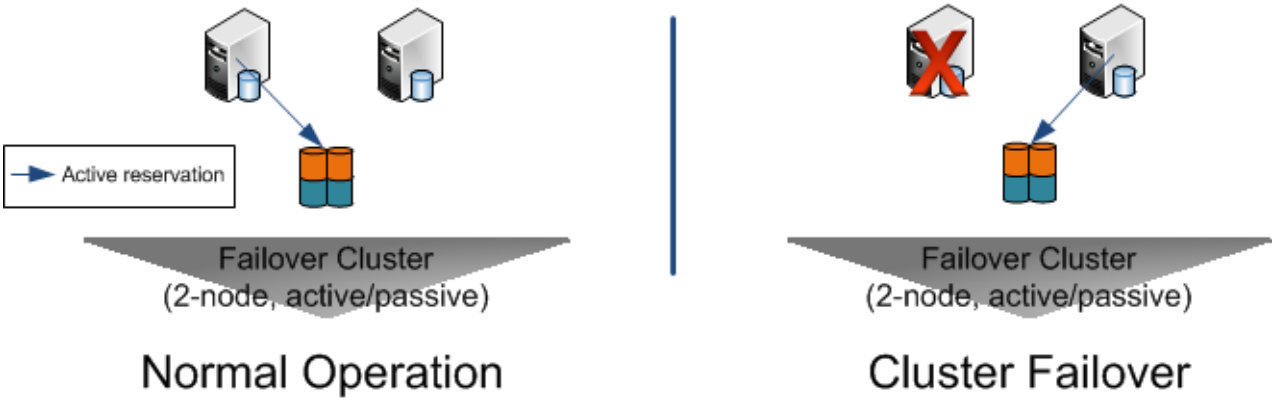

In SQL Server 2012, Microsoft introduced capabilities for FCI to support the use of SMB 3.0 shares. In SQL Server 2014, Microsoft introduced the support for Cluster Shared Volumes (CSVs), a form of shared storage. In FCIs, unlike AGs, there is always a single copy of the data.

In Microsoft AlwaysOn FCI, the first SQL Server instance serves as the primary or active node, while the other serves as a passive node. Both nodes are connected to a shared disk architecture with only one node accessing/updating the data at any given time.

Starting with vSphere 6, vSphere vMotion is supported with Microsoft failover cluster nodes with shared physical RDM. For more information, see the blog at https://blogs.vmware.com/apps/2015/02/say-hello-vmotion-compatible-shared-disks-windows-clustering-vsphere.html.

For more information about supported configuration of WSFC and AlwaysOn with vSphere, see *Microsoft Clustering on VMware vSphere: Guidelines for supported configurations* (https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1037959).

5.2 AlwaysOn Availability Groups

AlwaysOn AG provides protection on the database level rather on the instance level, allowing multiple databases to fail over as a single unit.

The AlwaysOn AG does not require shared disk architecture like FCI even though it is built on a WSFC. This means that each Availability Group replica has its own copy of the database, which removes the single point of failure from storage but adds additional storage space requirements for each AG replica.

In AlwaysOn AG, the log streams are replicated from the primary to the secondary replicas either synchronously or asynchronously where up to two synchronous replicas are allowed. AlwaysOn AG relies on WSFC at its core for failover policy or quorum management. Each Availability Group replica is an SQL Server instance running on a WSFC node.

AlwaysOn AGs have two modes of synchronization:

- Synchronous is used for high availability purposes within a single site or between sites within a metro distance. Synchronous replication allows for automatic failover or administrator initiated failover called “forced failover”.
- Asynchronous replication is used for disaster recovery purposes between two remote sites. Asynchronous replication allows for only forced manual failover and not automatic failover.

AlwaysOn AG provides support for the following use cases:

- One primary database replica and up to eight secondary database replica targets.
- Mix of synchronous and asynchronous data replication between primary and multiple secondary replicas.
- Offload backup to a secondary replica.
- Read-only secondary replicas.
- Ability to read from secondary replicas.
- Better application abstraction using the Listener Like FCI AG provide flexible failover policy with WSFC.
- Unlike database mirroring, a single AG can have multiple databases. This allows for more than one database to fail over as a unit.

VMware availability features, such as vSphere vMotion, vSphere HA, and vSphere DRS, are fully compatible with AlwaysOn AG. As detailed previously, while vSphere HA is not a replacement for AlwaysOn AG, it enhances its SLA by protecting the availability groups.

Figure 9. Windows Server Failover Cluster with AlwaysOn AG

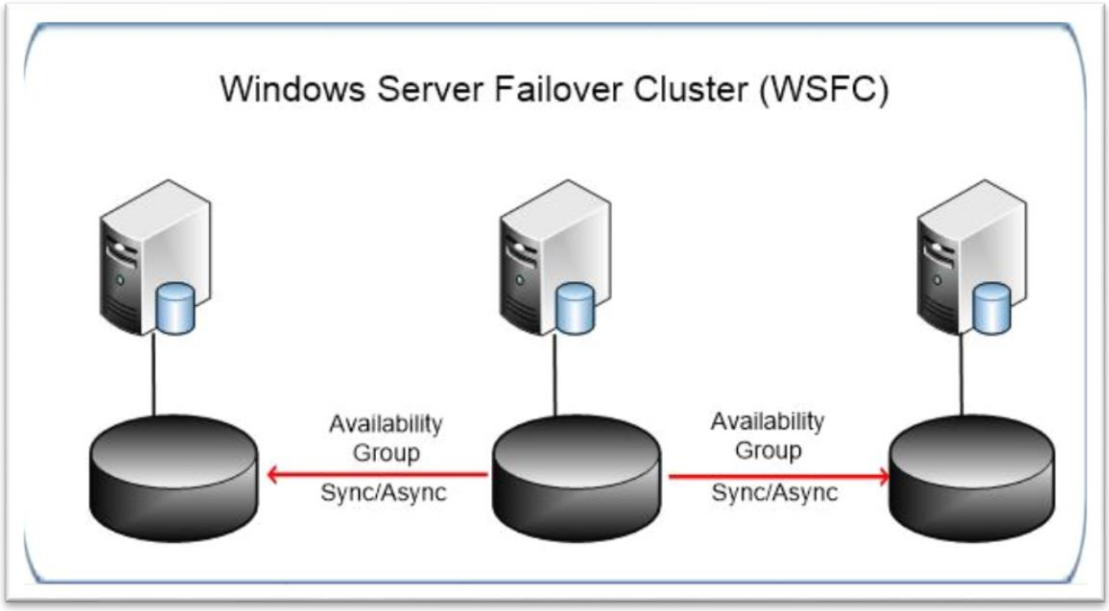

Example: SQL Server AlwaysOn AG for HA and DR

The following example is a purely non-shared storage solution. The primary data center hosts two availability replicas in synchronous commit mode for high availability. The secondary data center hosts one replica, running in asynchronous mode for disaster recovery protection. vSphere HA is used in both production and DR sites to protect virtual machines from host failure and to facilitate rapid recovery.

The architecture allows for variations to this topology using multiple data centers as well as multiple replicas up to the maximum number of supported replicas by an AG.

Figure 10. SQL Server AlwaysOn AG for HA and DR

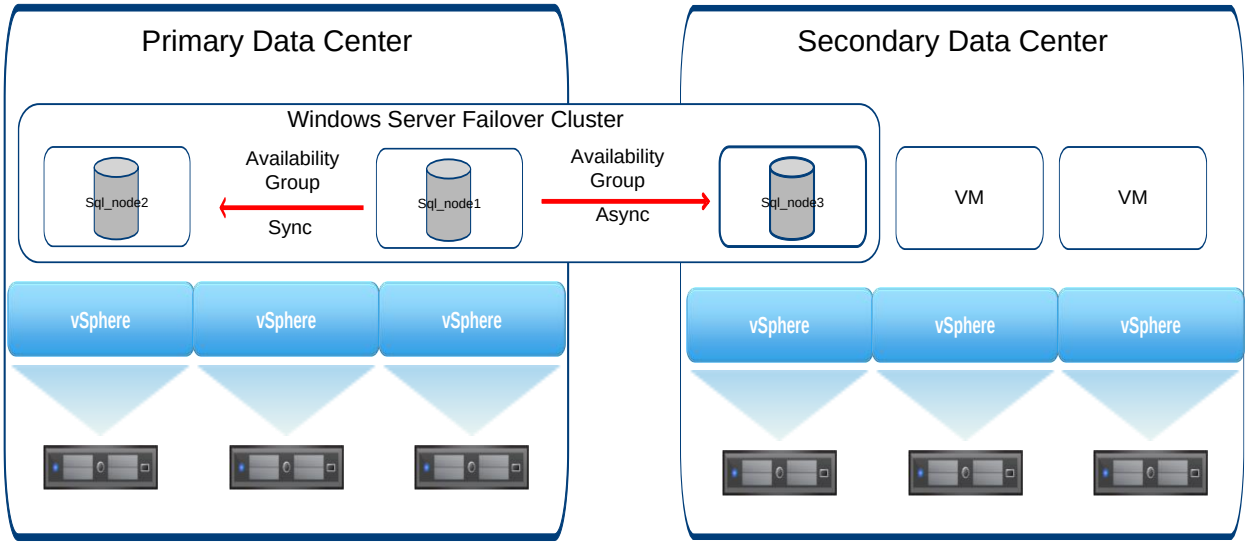

5.3 Database Mirroring

Database mirroring was introduced in SQL Server 2005 and deprecated in SQL Server 2012. It provides database-level protection by mirroring transactions from the principal to a single secondary mirror with fast recovery through automatic or manual failover.

If running a version previous to SQL Server 2016 Standard edition, use database mirroring when you require maintaining multiple copies of your data to create a warm standby of your SQL Server database environment. Only user databases can be mirrored. Database mirroring operates in the following modes:

- High safety mode with automatic failover.
- High safety mode without automatic failover.
- High performance mode.

Figure 11. SQL Server Database Mirroring

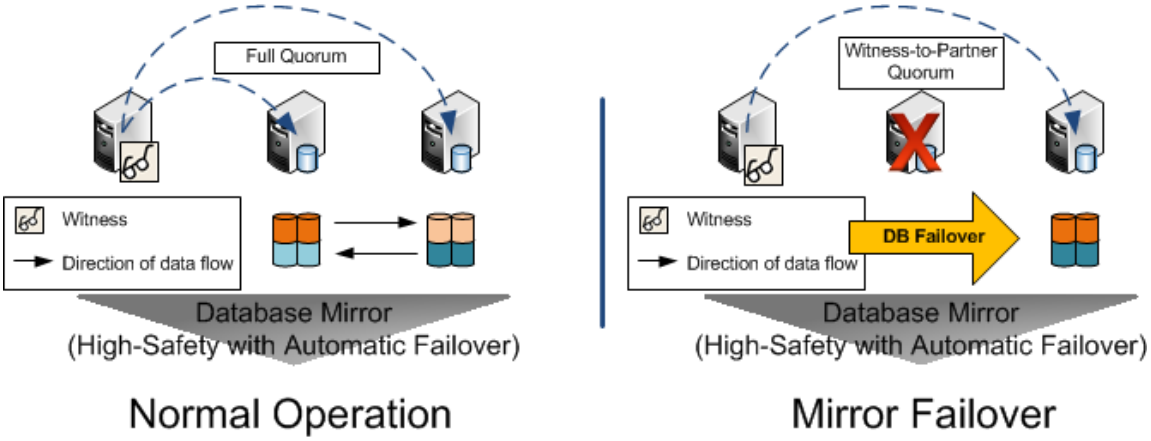

VMware availability features, such as vSphere vMotion, vSphere HA, and vSphere DRS, are fully supported in combination with database mirroring, offering the greatest level of protection for your SQL Server virtual machines. With database mirroring, you can eliminate a single point of failure in storage by locating the mirrored data on another physical array.

Example: vSphere HA with SQL Server database mirroring for faster recovery

Database mirroring can be combined with vSphere HA to provide higher levels of SQL Server protection and flexibility. If configured in high safety mode, log streams are replicated synchronously from the active database to one or more passive database copies which are kept in real-time lockstep with the active database. Upon SQL Server or operating system failure, the environment can be failed over to the mirrored database copy. If configured with automatic failover, failover to a mirrored database copy is automatic, facilitated by a server role known as the *witness*.

During the time between loss of a node and its restoration, your SQL Server environment is vulnerable to further failures. Resynchronizing the mirrored nodes could take a substantial amount of time after the failed node is restored. vSphere HA helps to alleviate this issue by restarting the failed primary node on another available host in the vSphere cluster. This quickly restores the database to full protection and reduces the amount of time spent in the failover state.

An example configuration is shown in the following figure.

Figure 12. vSphere HA with SQL Server Database Mirroring for Faster Recovery

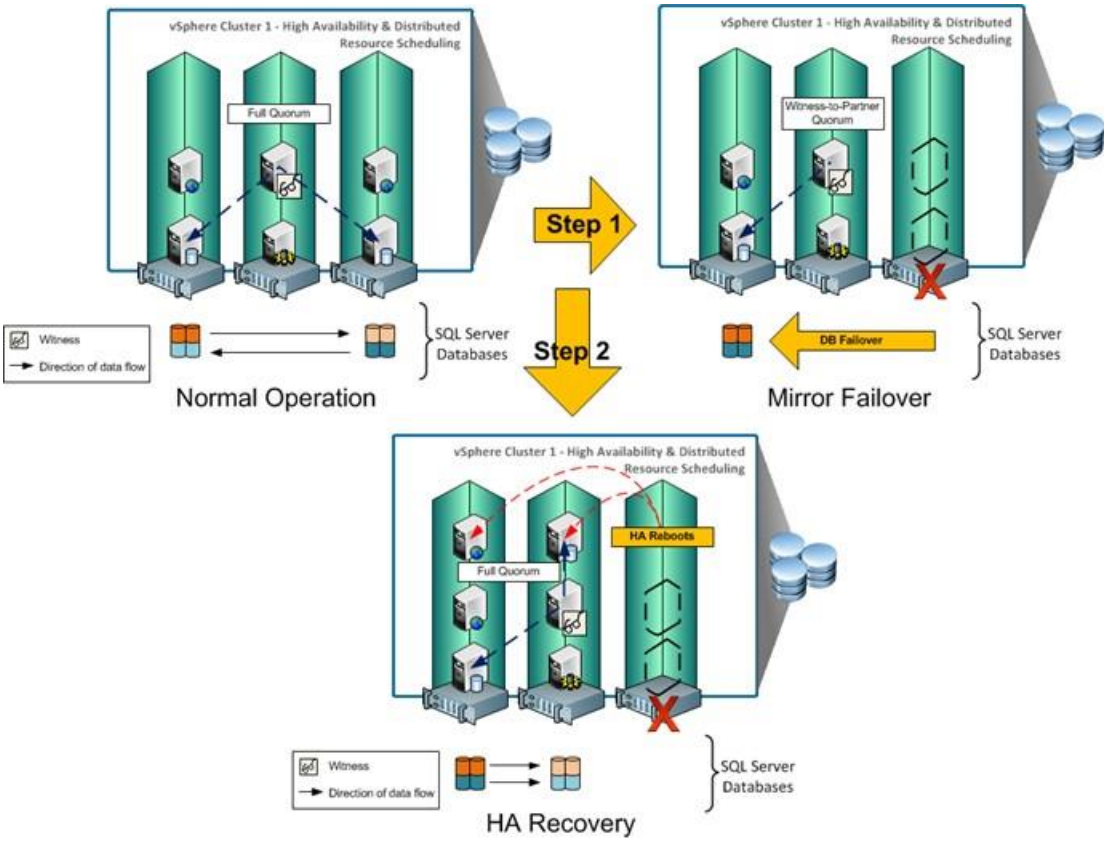

5.4 Log Shipping

The log shipping availability option typically provides lower-cost automated transaction log backup and restore functionality, which provides redundancy at the database level. SQL Server provides the underlying framework for doing automated backup, copy, and restore of transaction log files. Backups are performed on the primary SQL Server instance and restores are performed on secondary SQL Server instances. Scheduling is done through SQL Server Agent jobs.

As an availability strategy, the log shipping option does not provide any automatic failover capability and might allow some data loss, but the time synchronization interval is configurable, which gives users some level of control. (This data loss can occur due to transaction log data that was corrupted, missing, or possibly not sent to the backup log file on the secondary server.) Log shipping can be used in conjunction with failover clustering to provide good site failure redundancy with a low-cost solution.

You can use SQL Server log shipping to create multiple copies of your databases that can be used as a warm standby for recovery, or to provide a database that can be used for reporting and to offload reporting and query functions from the primary server.

6. Backup and Restore Options

The feature set available to an application, when deployed in a virtual environment, is not less than what is available with a physical deployment. In fact, there are more options available for protecting entire virtual machines. This is especially useful for applications that require extensive configuration. For SQL Server, the standard methods for backup are supported. These tend to be deployed using a third-party backup agent that uses a VSS requestor to coordinate with the VSS writer to prepare the database files for backup. Regardless of the backup solution required, VMware and VMware partners have provided solutions for most situations.

6.1 In-Guest Software Solutions

Many organizations have dedicated backup support teams or requirements that might not allow them to integrate the backup solution to the level that is available with vSphere. In these situations, traditional backup methods are used, and a virtualized environment allows that. Many of the leading backup software providers are VMware partners and provide full support for using their backup solutions within a virtualized guest operating system. Backup administrators can continue to deploy and manage the backup agents, jobs, and restores as though they were running on physical systems. This is especially helpful for AlwaysOn FCI deployments with shared disk, where in-guest backup is the only supported solution.

Using in-guest agents can also be tied to who is responsible for restoring operation, whether the DBA or IT. Some backup solutions can also provide an in-guest restore agent that allows the DBA to restore even though the backup was performed on the VM or storage level.

Example: In-guest SQL Server aware backup solution

Centralized backup management software controls the backup schedule, save set, and target location for all systems virtual and physical. The backup agent software loaded within the guest operating system allows the virtual machine guest operating system to be managed the same way as all other systems. Additional plug-ins from backup software vendors provide application-aware support.

Figure 13. In-Guest SQL Server Aware Backup Solution

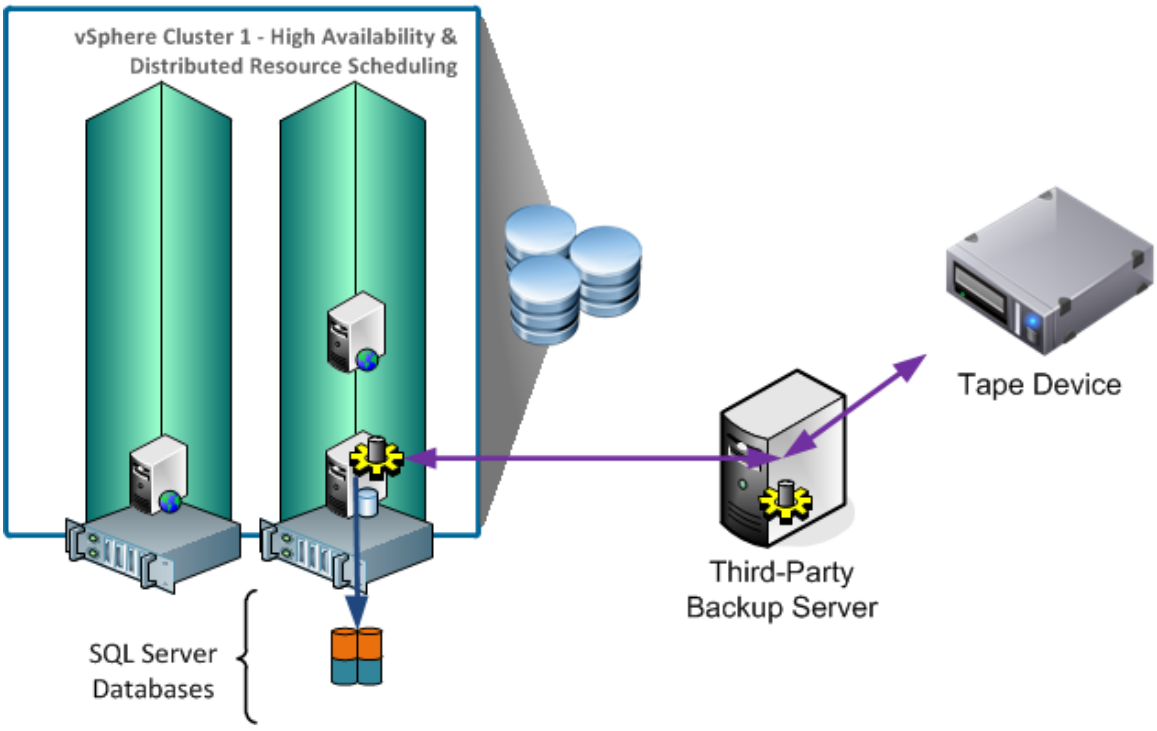

6.2 VMware Data Protection

VMware vSphere Data Protection™ is a recently available backup and recovery solution. It is fully integrated with VMware vCenter Server® and VMware vSphere Web Client, providing disk-based backup of virtual machines and applications.

vSphere Data Protection is based on the industry-leading EMC Avamar deduplication backup and recovery software. The key features provided by vSphere Data Protection are:

- Wizard-driven setup and management to quickly and easily implement a data protection solution for a vSphere virtual machine environment.
- Significantly reduced backup data disk space requirements, with the patented, variable-length Avamar deduplication technology.
- Use of VMware vSphere Storage APIs – Data Protection as well as changed block tracking (CBT) to reduce load on the vSphere host infrastructure and minimize backup window requirements.
- Agentless virtual machine backup and restore that reduces complexity and deployment time.
- Integration with EMC data domain for additional scale, efficiency, and reliability.
- Microsoft Exchange Server agent for application-consistent backup and restore of databases and mailboxes, including those protected by a database availability group (DAG).
- Microsoft SQL Server agent that leverages the Virtual Backup Device Interface (VDI) feature for proper backup and restore of databases in standalone configurations and clustered environments.
- Microsoft SharePoint agent that enables granular database backup and restore.
- Reliable, efficient replication of backup data between vSphere Data Protection appliances for redundancy and offsite data protection.
- Flexibility to restore replicated backup data at both the source and target locations.
- Automated backup verification that provides the highest level of confidence in backup data integrity.
- Secure, efficient backup data replication to Avamar for offsite data protection.
- Direct-to-host emergency restore operation that enables virtual machine recovery even when vCenter Server and vSphere Web Client are offline.
- File level restore (FLR), which enables granular file and folder restoration, without the need for an agent in Microsoft Windows and Linux virtual machines.
- Simple Web browser-based administration through vSphere Web Client.
- Appliance and backup data protection through a checkpoint-and-rollback mechanism.
- Deployment of external proxies enabling as many as 24 parallel backup operations.

vSphere Data Protection has the capability to properly back up and restore SQL Server databases. SQL Server AG is also supported. A vSphere Data Protection application agent is installed in the guest operating system of each virtual machine running SQL Server. It is also possible to install these agents on physical machines to protect physical SQL Server instances. Agents enable application-consistent backup and recovery, and provide support for other options, such as full, differential, or incremental backups, multistream backups, and database log management.

Figure 14. SQL Server AlwaysOn Cluster Backup Job Options

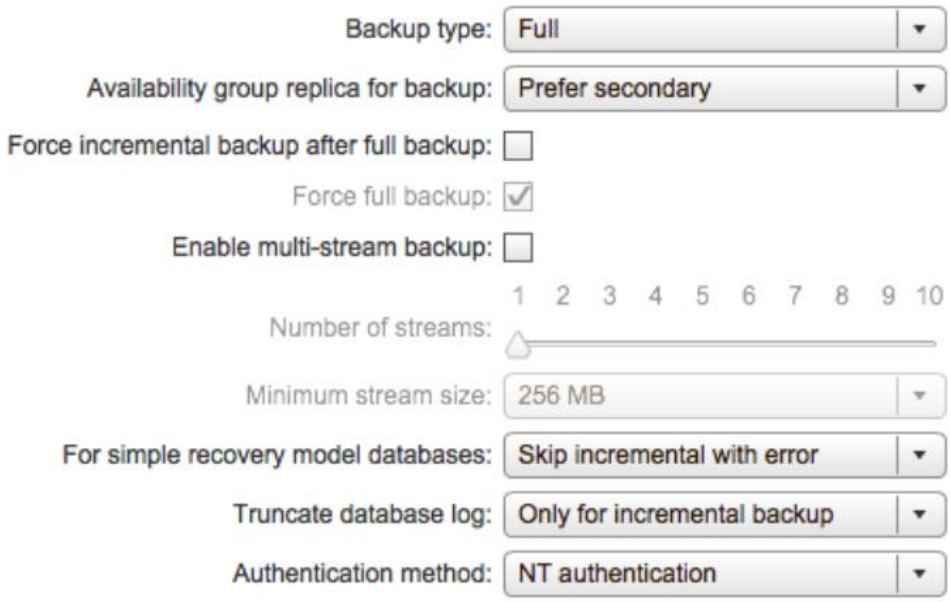

For more information about vSphere Data Protection, see *VMware vSphere Data Protection 6.1 Technical Overview* (https://www.vmware.com/files/pdf/vsphere/VMware-vSphere-Data-Protection-Overview.pdf).

6.3 Array-Based Backup Solutions

As is the case with the in-guest solutions, array-based solutions provided by many of the leading storage vendors continue to function with vSphere deployments of SQL Server. Array-based backup solutions for SQL Server use the Volume Shadow Copy Service (VSS) to produce near-instant, application-aware clones or snapshots of SQL Server databases. These local clones or snapshots can then be backed up to disk, tape, or cloned off site for disaster recovery purposes. Guidance on proper deployment methods and any additional considerations when running in a virtualized environment must be provided by the storage vendor. VMware can provide a comprehensive list of ISV partners that provide array-based replication of databases used by a SQL Server instance for backup and restore operations.

Example: Array-based SQL Server aware backup solution

An array-based backup solution provides integration with the SQL Server instance and the underlying storage solution. A software agent provided by your backup vendor coordinates with the SQL Server VSS writers to create a supported backup image of your SQL Server databases. These databases can later be streamed to tape as flat files for compliance or archive requirements with no I/O impact to the production data.

Figure 15. Array-Based SQL Server Aware Backup Solution

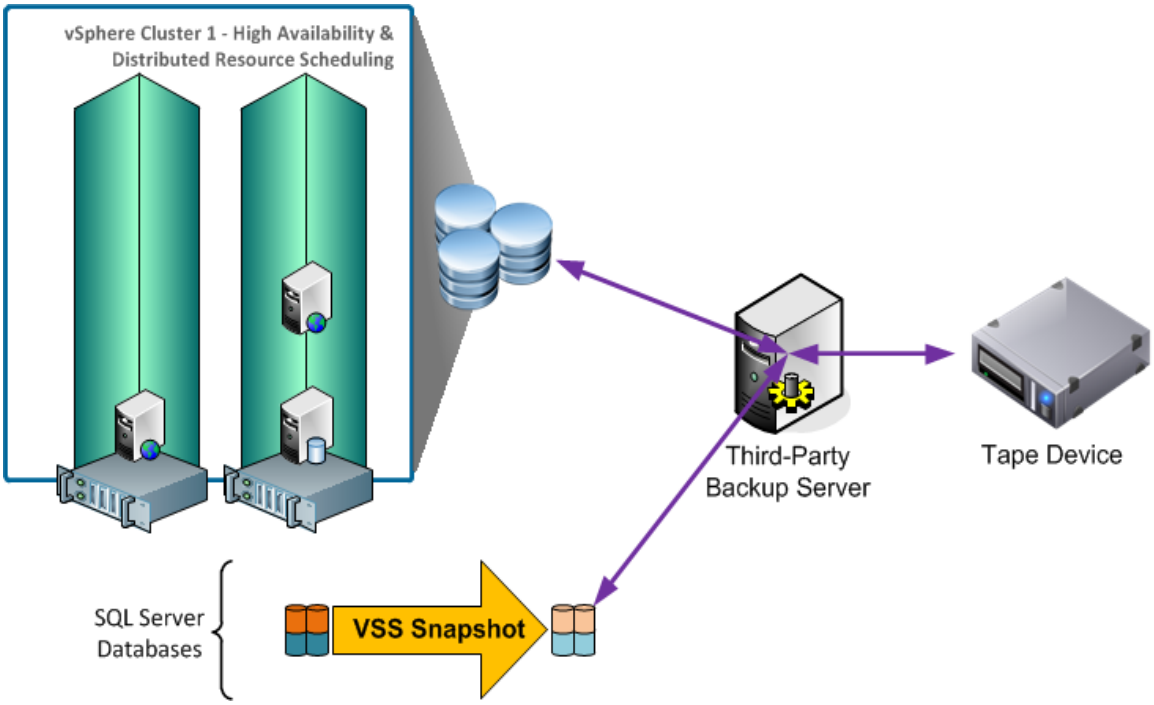

6.4 Using SQL Tools for Backup

An additional option for backup is to use native SQL Server tools to create a full backup of a specific database to a file. The backup file is a flat file that can reside on a separate VMDK. That file can be further backed up using agentless or image-level backup solutions, including vSphere Data Protection. Some organizations choose to back up their databases using this method to reduce the amount of agents deployed and managed in their environment.

Full backups can be initiated using SQL Server Management Studio, Transact-SQL, or PowerShell. For further information on how to configure a full backup using SQL Server tools, see the Microsoft article at https://msdn.microsoft.com/en-us/library/ms187510(v=sql.120).aspx.

7. Patch Management Options

VMs running SQL Server instances will require two levels of patching: for the guest OS and for each SQL Server instance. Patching is a key component of any availability strategy, because patches provide supportability and critical issues might be fixed. Some patches may require downtime to install. The following sections discuss SQL Server and vSphere capabilities that reduce downtime for patching an SQL Server instance only, not its underlying guest OS.

7.1 Rolling Upgrade Using SQL Server Native Features

Some SQL Server native availability features, such as AlwaysOn AG, AlwaysOn FCI, and DBM, support rolling upgrades and patching of SQL Server. You can apply a service pack or critical fix to the passive nodes or secondary replicas. After the installation is complete on a certain number of nodes or secondaries, you can conduct a manual failover and apply the service pack or critical fix to the primary instance or database. Refer to the following Microsoft documentation on the rolling upgrade process:

- Upgrading Mirrored Instances
  http://msdn.microsoft.com/en-us/library/bb677181.aspx
- *Upgrade a SQL Server Failover Cluster Instance (Setup)*
  https://msdn.microsoft.com/en-us/library/ms191295(v=sql.130).aspx

8. Best Practices and Deployment Considerations

When deploying SQL Server on vSphere with high availability features, consider the guidelines described in the following sections to increase the availability and performance of your deployment.

8.1 Cluster Virtual Machine Placement

vSphere HA and DRS automatically manage virtual machine placement under system failure or resource contention. When running WSFC in virtual machines on a vSphere HA or DRS enabled cluster, whether AG or FCI, to avoid a single point of failure, the clustered virtual machines should be kept apart on different physical hosts. The following practices can help to achieve this:

- If the clustered VMs reside on a vSphere cluster with DRS enabled, create DRS anti-affinity rules to keep the clustered virtual machines nodes on different hosts. See the procedure on how to configure anti-affinity rules in the vSphere documentation (http://pubs.vmware.com/vsphere-60/index.jsp#com.vmware.vsphere.resmgmt.doc/GUID-7297C302-378F-4AF2-9BD6-6EDB1E0A850A.html).
- If vSphere Storage DRS is being used with clustered VMs, create inter-VM anti-affinity rules to keep virtual machines' VMDKs on different VMFS volumes. See the procedure on how to configure anti-affinity rules in the vSphere documentation (http://pubs.vmware.com/vsphere-60/index.jsp#com.vmware.vsphere.hostclient.doc/GUID-FC84ECD9-C7DA-4D27-B5EC-0C8E396A6800.html).
- Enable strict enforcement of affinity rules to ensure that DRS anti-affinity rules are applied after an HA event. See the procedure in the vSphere documentation Enable Strict Enforcement of Affinity Rules (MSCS). (This document is currently uses the name MSCS but refers to WSFC as well).

8.2 Network Considerations

The network is a critical component required for cluster node communication. For SQL Server cluster solutions, a private heartbeat network is required for cluster "keepalive". SQL Server non-shared disk high availability solutions, such as AlwaysOn AG, also use the network to manage data replication between replicas. In synchronous replication mode, the performance of SQL Server is highly dependent on the network bandwidth and latency. Consider the following network configuration guidelines for optimal performance:

- Install enough network adapters to separate networks used for different purposes. For example, a separate network for data replication, vSphere vMotion, VMkernel, and so forth.
- Make sure that WSFC vNICS are connected to redundant physical network components such as switches and physical NICS.
- If using iSCSI, the network adapters should be dedicated to either network communication or iSCSI, not both.
- Use the VMXNET3 paravirtualized NIC. VMXNET3 is optimized for virtual environments and designed to provide high performance.
- Enable jumbo frames for the iSCSI or vSphere vMotion network.
- Use static IP addresses for network interfaces managing client connections and heartbeat links in a Microsoft Windows server cluster. Using a dynamic configuration through DHCP is not recommended, because the failure to renew a DHCP lease could disrupt cluster operations. See *Cluster May Fail if IP address Used from DHCP Server* (http://support.microsoft.com/kb/170771).
- If deploying AGs or DBM with synchronous data movement, a high-speed network should be used for replication traffic. Confirm that the bandwidth and latency of the network is sufficient for supporting the amount of SQL Server transaction traffic that might occur.

- To reduce the risk of haing dropped packets inside the guest significantly make sure to follow the following guidelines:
    - enable Receive side scaling (RSS) in the guest, for more information about RSS see section 4.1 in the “SQL Server on VMware best practices guide” https://www.vmware.com/files/pdf/solutions/SQL_Server_on_VMware-Best_Practices_Guide.pdf
    - Keep VMware tools up to date
    - Make sure that drivers and firmware of the physical NICs are up to date

8.3 vMotion Considerations

The use of the VMware resource balancing features DRS and vMotion, along with WSFC, is a configuration supported by both Microsoft and VMware, however, additional configuration options are required to provide seamless and non-disruptive interactions with SQL AG and FCI which relies on WSFC for its functionalities (or the traditional MSCS for earlier versions of SQL Server). By default, each clustered SQL node exchanges heartbeat packets with its partners every second. If the node does not receive a response from any partner after 5 consecutive probes, the probing node considers the partner to be non-responsive (unavailable). The WSFC service on the probing node then initiates corrective actions, including taking ownership of clustered resources previously owned by the non-responsive partner. This is what happens when there is an unintended and unplanned cluster resource failover. An aggressive cluster failover threshold can be problematic and disruptive. For example, a failover may not have been required because the non-responsiveness was a transient issue that went away after six seconds.

During a vMotion operation, a virtual machine’s memory pages are gradually copied from its current physical host to the target physical host and the virtual machine is subsequently switched over to the target host. During this transition procedure, the virtual machine is quiesced (VMware referes to it as stunning). Under normal operating condition, the stun period is very brief and not noticeable. However, the quiescing period might last longer than the five-second threshold under the some operating conditions, including the following:

- Size of the VM (CPU and RAM).
- Current outstanding operations, threads, and processes on the VM.
- Total CPU subscription ratio in the vSphere cluster.
- Network bandwidth or congestion state.

The net effect of a stun operation is that the stunned VM is unable to exchange heartbeat for the duration of the stun operation. If this lasts longer than five seconds, and the VM is in a cluster relationship, then its cluster partners considers it unavailable. This behavior is not peculiar to vMotion, or even virtualization in general. If your backup solution takes VSS-based snapshots of your server, it is also likely quiescing the server. This is a known behavior, and avoiding the unintended cluster failover incident (and its associated disruptive effects) when performing a vMotion operation on a clustered SQL Server node is an exact science for VMware. There are several configuration options described in the following sections that you can use to overcome these disruptive effects.

8.3.1 Cluster Heartbeat Settings

Two Windows Server Failover Cluster settings control the behavior of the cluster service to missed heartbeat probes. These are:

- *SameSubnetDelay* – Controls how often a node sends heartbeat probe packets.
- *SameSubnetThreshold* – Controls how many probe misses the node must tolerate before taking actions.

If the clustered nodes are in different subnets, the corresponding settings are:

- *CrossSubnetDelay*
- *CrossSubnetThreshold*

The following are the default values for these settings:

Figure 16. Default cluster heartbeat threashold values

```
PS C:\Windows\system32> get-cluster | fl *subnet*


CrossSubnetDelay           : 1000
CrossSubnetThreshold       : 5
PlumbAllCrossSubnetRoutes  : 0
SameSubnetDelay            : 1000
SameSubnetThreshold        : 5
```

A *SubnetDelay* value of 1000ms and a *SubnetThreshold* of 5 denotes that heartbeat probes are sent every second with a tolerance for a maximum of 5 missed heartbeats. These are the settings that enabled the unexpected failover described in the preceding section.

Adjusting either of these parameters to 10 seconds (for cluster nodes in the same subnet) and 20 seconds (for cluster nodes in different subnets) is the simplest and most effective way to prevent the unintended cluster failover issue described in the preceding section. The following are the settings to accomplish each of the recommendations:

(get-cluster).SameSubnetThreshold = 10

(get-cluster).CrossSubnetThreshold = 20

Figure 17. Changing the default heartbeat values

```
PS C:\Windows\system32> (get-cluster).SameSubnetThreshold = 10
PS C:\Windows\system32> (get-cluster).CrossSubnetThreshold = 20
PS C:\Windows\system32> get-cluster | fl *subnet*


CrossSubnetDelay           : 1000
CrossSubnetThreshold       : 20
PlumbAllCrossSubnetRoutes  : 0
SameSubnetDelay            : 1000
SameSubnetThreshold        : 10
```

Windows Server 2012 includes an additional cluster service logging component that tracks dropped heartbeat packets. The *RouteHistoryLength* setting and its default value (10) must also be modified so that it remains useful when investigating heartbeat-related issues. According to Microsoft, the recommended value for this setting must be double that of the *SubnetThreshold* setting. If the value *SameSubnetThrshold* and *CrossSubnetThreshold* follow the VMware recommendations, the value of *RouteHistoryLength* must be double the value of the *CrossSubnetThreshold* value. The VMware recommended value for *RouteHistoryLength* setting is, therefore, 40.

(get-cluster).RouteHistoryLength = 40

Figure 18. Changing RouteHistoryLength

```
PS C:\Windows\system32> get-cluster | fl RouteHistoryLength

RouteHistoryLength : 10


PS C:\Windows\system32> (get-cluster).RouteHistoryLength = 40
PS C:\Windows\system32> get-cluster | fl RouteHistoryLength

RouteHistoryLength : 40
```

8.3.2 Multiple vSphere vMotion Interfaces

Database failover due to vSphere vMotion operations can be mitigated by using multiple dedicated vSphere vMotion network interfaces. In most cases, these interfaces are also used for management traffic. Because management traffic is relatively light, this does not add significant overhead.

vSphere provides the ability to use multiple VMNIC interfaces for vSphere vMotion traffic to effectively load balance the vSphere vMotion traffic. Testing has shown up to a 25% increase in throughput achieved when multiple vSphere vMotion interfaces are used. In the test case with two 2Gbps interfaces configured for vSphere vMotion and no cluster heartbeat modifications, vSphere vMotion operations succeeded with no database failover.

Enabling multiple interfaces for vSphere vMotion requires configuring multiple VMkernel ports on different port groups. Each port group is assigned multiple VMNIC interfaces as either active or standby. See Multiple-NIC vMotion in vSphere at http://kb.vmware.com/kb/2007467 for detailed configuration procedures.

8.3.3 Enable Jumbo Frames for vSphere vMotion Interfaces

Standard Ethernet frames are limited to a length of approximately 1500 bytes. Jumbo frames can contain a payload of up to 9000 bytes. Support for jumbo frames on VMkernel ports was added to vSphere 4.0 for both ESX and ESXi. This added feature means that large frames can be used for all VMkernel traffic, including vSphere vMotion.

Using jumbo frames reduces the processing overhead to provide the best possible performance by reducing the number of frames that must be generated and transmitted by the system. During testing, VMware tested vSphere vMotion migration of clustered SQL Server nodes with and without jumbo frames enabled. Results showed that with jumbo frames enabled for all VMkernel ports and on the VMware vSphere Distributed Switch™, vSphere vMotion migrations of cluster member virtual machines completed successfully. During these migrations, no database failovers occurred, and there was no need to modify the cluster heartbeat setting.

The use of jumbo frames requires that all network hops between the vSphere hosts support the larger frame size. This includes the systems and all network equipment in between. Switches that do not support, or are not configured to accept, large frames will drop them. Routers and Layer 3 switches might fragment the large frames into smaller frames that must then be reassembled, which can degrade performance.

8.4 Database Recovery Mode

SQL Server non-shared disk high availability solutions such as AGs, DBMs, and log shipping use log stream or log record replication technologies for database redundancy. Full recovery mode is required for

each database utilizing these solutions. Log truncation does not happen until a log backup. Sufficient disk space must be planned to allow for transaction log growth.

8.5 Impacts of AlwaysOn AG Readable Secondary

With SQL Server AlwaysOn AG, a secondary replica can be made readable. Read-only access to secondary replicas is useful if your read-only workloads can tolerate some data latency. In situations where data latency is unacceptable, consider running read-only workloads on the primary replica.

When a secondary replica is configured for read access, the read-only workloads on the secondary replica consume system resources, such as CPU, memory, and I/O, from the read queries as well as the redo threads. The workload on the secondary often is completely different than the workload on the primary. The read-only queries are generally more analytical in nature. Analytics queries typically generate a high number of sequential large block I/Os transactions, and perform more aggregation operations that consume a large amount of CPU and memory. Analytics queries can also benefit from parallel processing threads.

Because the workload on a readable secondary often is completely different than the workload on the primary, additional indexes might be needed to achieve optimal query plans. If the read-only workload executed against a secondary replica requires additional indexes, those must be created on the primary replica to get transferred subsequently to the secondary replica. Sufficient disk space should be planned to accommodate the additional indexes.

Consider the impact of enabling read access on a secondary replica, as well as the characteristics of a read-only workload when capacity planning. A readable secondary replica could potentially demand more resources than the primary.

8.6 Tempdb Considerations

The name of the temporary work space for SQL Server is `tempdb`. The `tempdb` file is recreated every time SQL Server is started so that the system always starts with a clean copy of the database. There is never anything that needs to be saved in `tempdb` from one session of SQL Server to another.

Given the nature of `tempdb`, there is no need to include it in a virtual machine level backup, or to replicate it in a multisite DR solution. Consider placing `tempdb` on its own dedicated VMDK to exclude `tempdb` from any unnecessary operations. Follow Microsoft's best practices for the number of `tempdb` files required for your SQL Server instance based on the number of vCPUs the VM has (up to eight). See the following Microsoft article for more information (https://support.microsoft.com/en-us/kb/2154845).

8.6.1 Capacity Planning

When deploying SQL Server on vSphere, especially on a secondary AG copy, consider increasing `tempdb` space due to the following:

- Snapshot isolation level copies row versions into `tempdb`
- Temporary statistics created by SQL Server on secondary databases

SQL Server automatically enables the snapshot isolation level when a secondary copy is set to readable. A 14-byte overhead is added to the row on the primary if a row version was not already enabled. A 14-byte overhead is also added to the row on the secondary replica. Additionally, a row version is generated by the REDO thread in `tempdb` of the secondary database replica as it processes the update/delete operation from the primary database.

Because the workload on a readable secondary database is often completely different than the workload on the primary database, additional indexes and column statistics might be needed to achieve optimal query plans. SQL Server automatically creates additional column statistics and stores them in `tempdb` of the readable secondary database replica.

8.7 Client Connectivity

When deploying SQL Server AlwaysOn AG, you have a better way to connect your applications with the SQL Server. You can connect your application to the database replica using an availability group listener. An availability group listener is a virtual network name (VNN) that directs read-write requests to the primary replica and read-only requests to the read-only secondary replica. Always create an availability group listener when deploying AlwaysOn AG on vSphere. This enables application clients to connect to an availability replica without knowing the name of the physical instance of the SQL Server installation. The application connection string does not need to be modified in case of a failover.

Refer to *Prerequisites, Restrictions, and Recommendations for AlwaysOn Client Connectivity (SQL Server)* at http://technet.microsoft.com/en-us/library/ff878487.aspx for additional information on creating an availability group listener.

8.8 Considerations for WSFC Quorum Mode

Quorum is the mechanism a WSFC uses to help ensure it remains available. It is extremely important to choose the right quorum mode. The main thing that you must consider is the type of witness that you might need to configure depending on your WSFC configuration. A witness is an additional resource that assists in the arbitration process to keep the WSFC up and running. There are three types of witnesses: disk, file share, and (starting in Windows Server 2016), cloud. The goal is to keep the total number of votes that count towards quorum odd.

The quorum modes are as follows:

- Node Majority – In this mode, each node gets a vote. If the number of votes switches from odd to even, a witness should be added and the quorum model will be changed. This quorum model is often used for an AG with no shared storage and an odd number of nodes.
- Node majority with a witness – This is a combination of Node Majority with a witness and is the most common configuration. If you are not using shared storage, you should not use a disk witness if at all possible.
- Node and File Share Majority – This model uses a combination of node and file share as witness. This is recommended for deploying AlwaysOn on vSphere.
- No Majority: Disk Only – This is the model that was used for MSCS pre-Windows Server 2008. VMware recommends discontinuing use of this model.

See *Configure and Manage the Quorum in a Windows Server 2012 Failover cluster* (https://technet.microsoft.com/en-us/library/jj612870.aspx) for additional details about the different quorum models.

Each customer environment is different, so depending on your particular deployment scenario, you might choose a different quorum model. vSphere vMotion, DRS, and vSphere HA features are fully supported if you deploy AlwaysOn AG with these quorum modes.

9. Additional Resources

For the latest official best practices guide for SQL Server on vSphere, see *Architecting SQL on VMware vSphere Best Practices Guide* (https://www.vmware.com/files/pdf/solutions/SQL_Server_on_VMware-Best_Practices_Guide.pdf).

For a guide on how to setup WSFC/MSCS on vSphere, see the following KB *Setup for Failover Clustering and Microsoft Cluster Service* (https://kb.vmware.com/kb/1004617).

For the full list of guidelines and supported configurations for WSFC/MSCS with vSphere, see *Microsoft Clustering on VMware vSphere: Guidelines for supported configurations* (https://kb.vmware.com/kb/1037959)

Refer to *Support policy for Microsoft SQL Server products that are running in a hardware virtualization environment* (https://support.microsoft.com/en-us/kb/956893) to guarantee that ESXI-based deployments of SQL Server will be supported by Microsoft should you need to call them

Refer to the *SQL Server Failover Cluster Installation* (https://msdn.microsoft.com/en-us/library/hh231721(v=sql.120).aspx) for prerequisites on deploying a SQL Server FCI.

For additional information about virtualizing business mission critical applications, see *Virtualizing Microsoft SQL on VMware vSphere* (http://www.vmware.com/solutions/business-critical-apps/sql-virtualization).

For additional information on quorum for clustered SQL Server implementations, see Allan Hirt's PAS Summit 2014 presentation "Did You Vote Today? A DBA's Guide to Cluster Quorum" (http://www.sqlpass.org/summit/2014/PASStv.aspx?watch=o9s-Z70mizQ).

10. Acknowledgments

Author: Niran Even-Chen, Staff Solutions Architect, Microsoft Enterprise Applications

Thanks to the following people for their inputs:

- Scott Salyer – Director, Enterprise Application Architecture
- Oleg Ulyanov – Sr. Consultant, VMware PSO
- Deji Akomolafe - Staff Solutions Architect, Microsoft Enterprise Applications
- Andrea Couvrey – Senior Technical Writer, GTPS Solutions Tech Writers

Special thanks to the following SQL Server experts for their contributions and reviews:

- Allan Hirt - Managing Partner, SQLHA LLC
- David Klee - Founder and Chief Architect, Heraflux Technologies