

VMware vCloud® Architecture Toolkit™
for Service Providers

Architecting a Hybrid Database Strategy with Microsoft SQL Server and the VMware Cloud Provider™ Program

Version 2.9
January 2018

Martin Hosken





© 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

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



Contents

- Overview 5
- Target Audience 5
- Use Case 5
- Introduction to Microsoft SQL AlwaysOn 6
- Employing SQL Server 2016 AlwaysOn Availability Groups for Disaster Recovery 10
- Designing the Solution 11
 - 6.1 Distance and Latency Considerations 13
 - 6.2 Sample Availability Group Design..... 14
- Architecting a Robust Technical Platform 15
 - 7.1 Virtual Machine Placement 15
 - 7.2 Disk Provisioning..... 15
 - 7.3 Network Considerations 15
 - 7.4 Database Recovery Mode 16
 - 7.5 The Impacts of a Readable Secondary Replica 16
 - 7.6 tempdb Design Considerations 16
 - 7.7 Client Connectivity 17
 - 7.8 Considerations for Quorum Mode..... 18
- Conclusion 19
- Assumptions and Caveats 20
- Reference Documents 20



List of Figures

- Figure 1. Windows Server Failover Cluster with AlwaysOn Availability Group 7
- Figure 2. Windows Server Failover Cluster with AlwaysOn Availability Group Configuration..... 8
- Figure 3. Solution Logical Architecture 9
- Figure 4. SQL Server 2012 AlwaysOn Availability Group for HA and DR 10
- Figure 5. Sample Architecture Database Functions 12
- Figure 6. Sample Availability Group Design 14
- Figure 7. Availability Group Listener Client Connectivity 17

List of Tables

- Table 1. Microsoft SQL Server High Availability Technologies..... 6
- Table 2. Solution Prerequisites and Requirements 11
- Table 3. Architecture Virtual Machine Function 12
- Table 4. Distance and Estimated Link Latency..... 13



Overview

Today's dynamic business environment requires continuous uptime and fast recovery from disaster scenarios to meet business requirements and service level agreements (SLAs), with one of the most important requirements being availability of databases. The "being prepared" approach to providing application high availability is aimed at meeting customer agreed to SLAs, while reducing risk of revenue losses and maintaining compliance.

Planning for disasters and minimizing their impact is critical for all enterprises and government agencies. VMware Cloud Providers™ can assist in providing data center resources and services to leverage application layer high availability solutions for Microsoft and other software vendors. Designing and deploying a database solution such as Microsoft SQL Server 2016, in conjunction with a highly available hybrid cloud solution from a VMware Cloud Provider, can help businesses mitigate risks and still control deployment and operational costs.

Many different approaches to architecting a VMware based hybrid cloud exist, depending on the specific use case requirements and technologies employed. However, the end goal is always the same. A hybrid cloud is a cloud computing environment in which an organization provides and manages some resources in-house and has others provided externally. This vCloud Architecture Toolkit for Service Providers (vCAT-SP) solution architecture document describes a VMware Cloud Provider Program solution for supporting a geographically dispersed SQL Server 2016 environment protected by Microsoft's AlwaysOn technology.

Target Audience

This paper is for architects, engineers, application owners, and business leaders involved in the decision making process, and anyone else interested in guidance on designing a VMware Cloud Provider Program hybrid cloud solution on VMware technologies.

While the design outlined in this document is targeted toward a medium-sized or large-sized customer with enterprise application workloads, this type of solution can easily be leveraged for smaller environments.

Use Case

Businesses of all sizes are demanding improved levels of availability for virtualized business critical applications to meet service needs. IT organizations are upgrading or migrating to Microsoft Tier 1 applications that can leverage application availability across multiple physical data centers. One such offering from Microsoft is SQL Server 2016 with its AlwaysOn feature that supports for high availability and disaster recovery solutions. The Microsoft SQL Server 2016 AlwaysOn technology is an integrated, flexible, and cost-efficient solution that can provide redundancy within and across data centers. In addition, AlwaysOn provides fast application failover for maximum availability, and data protection of business critical applications. The migration to Microsoft SQL Server 2016 can also replace older physical infrastructures, which are difficult to manage and expensive to support.

VMware Cloud Providers who offer consumers hybrid cloud-based virtualized on-premises and off-premises database solutions, like SQL Server, are supporting the rapidly transforming way organizations deploy and manage their IT assets and resources. Application-based hybrid cloud service offerings can provide both consumer and provider with a flexible solution that provides significant advantages in cost, efficiency, and availability. With fewer servers consuming less space, power, cooling, and physical facilities, IT organizations can realize significant capital and operational expense savings by moving the IT infrastructure that provides high availability to Tier 1 applications to a VMware Cloud Provider Program data center.

This vCAT-SP solution architecture focuses on deploying Microsoft SQL Server 2016, taking advantage of its AlwaysOn functionality to provide a hybrid cloud approach to application availability with a VMware Cloud Provider.



Introduction to Microsoft SQL AlwaysOn

Traditionally, running Microsoft SQL Server on the VMware vSphere® platform offered multiple database availability and disaster recovery options using VMware features, such as VMware vSphere High Availability and VMware vSphere Distributed Resource Scheduler™ (DRS), with Microsoft features, such as failover clustering. The Microsoft SQL Server high availability features have also been combined with vSphere to create a flexible availability and recovery solution, applying the most efficient and appropriate tools for each specific use case.

All of Microsoft SQL Server availability technologies are supported on the core vSphere platform, including:

- Log shipping
- Database mirroring
- AlwaysOn failover cluster instances
- AlwaysOn Availability Groups

It is important to recognize that SQL database availability involves more than the technologies just described. When architecting this type of solution, you must also consider all the design factors for local site and remote site availability options, backup and restore options, and best practices and deployment considerations.

While the objective of this paper is to address use cases available with the AlwaysOn Availability Groups mechanism, to provide a complete picture of the available options with AlwaysOn technology, the following table lists all of Microsoft SQL Server 2016 availability options and their potential ability to meet Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs).

Table 1. Microsoft SQL Server High Availability Technologies

Technology	Granularity	Storage Type	RPO	RTO
Log shipping	Per database	Non-shared storage architecture	Last shipped transaction log	DBA recovery only
Database mirroring	Per database	Non-shared storage architecture	None*	< 3 seconds or administrator recovery
AlwaysOn Failover Cluster Instances	Per instance	Shared storage only	None	~30 seconds
AlwaysOn Availability Groups	Per database	Non-shared storage architecture	None**	~3 seconds or administrator recovery

* With high safety mode.

** With synchronous commit mode.

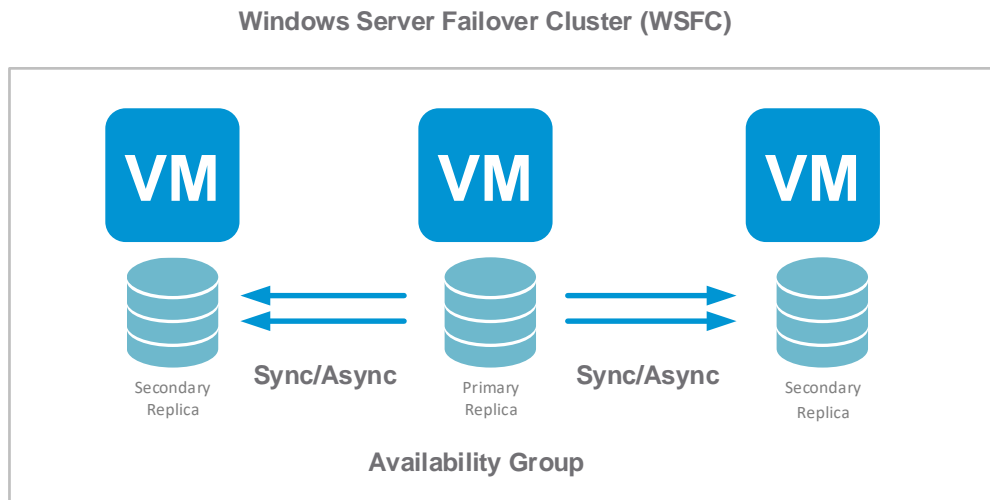
The potential RPO and RTO values shown in the previous table are typical examples of achievable values in real world implementations. However, architecture, hardware, and configuration are significant factors that can skew these estimations significantly.

The focus of this document is the SQL Server 2016 AlwaysOn Availability Groups feature, which provides flexible design choices for selecting an appropriate high availability and disaster recovery solution for



applications, and supports availability at either the application database or instance level. The AlwaysOn Availability Groups feature was developed for applications that require high uptime, need protection against failures within a data center, and require adequate redundancy against failures across data centers. This solution architecture helps protect application databases from planned and unplanned downtime without the need for shared storage between nodes.

Figure 1. Windows Server Failover Cluster with AlwaysOn Availability Group



The Availability Groups feature provides an enterprise-level alternative to database mirroring. The availability group supports the failover environment for a set of read-write primary databases, and one to four sets of corresponding secondary databases. Additionally, you can make secondary databases available for read-only access or backup operations. With AlwaysOn Availability Groups automatic failover mode enabled, if the primary replica fails, a secondary replica automatically takes over servicing client requests and becomes the new primary replica.

The AlwaysOn Availability Groups feature provides support for the following functions:

- The ability to offload backup jobs to a secondary node
- The ability to provide read-only access to secondary replica nodes
- Provide a single primary database replica and up to four secondary database replica targets
- Provide a mixture of synchronous and asynchronous data replication between primary and multiple secondary replicas
- Faster application failover is provided by availability group listeners
- Multiple databases can be configured in a single failover unit
- Flexible failover policies through the Windows Server Failover Clustering (WSFC) feature

This AlwaysOn Availability Group mechanism is also different from traditional SQL failover cluster Instances because a failover cluster can provide availability only at the entire Microsoft SQL Server instance level, not for individual databases.

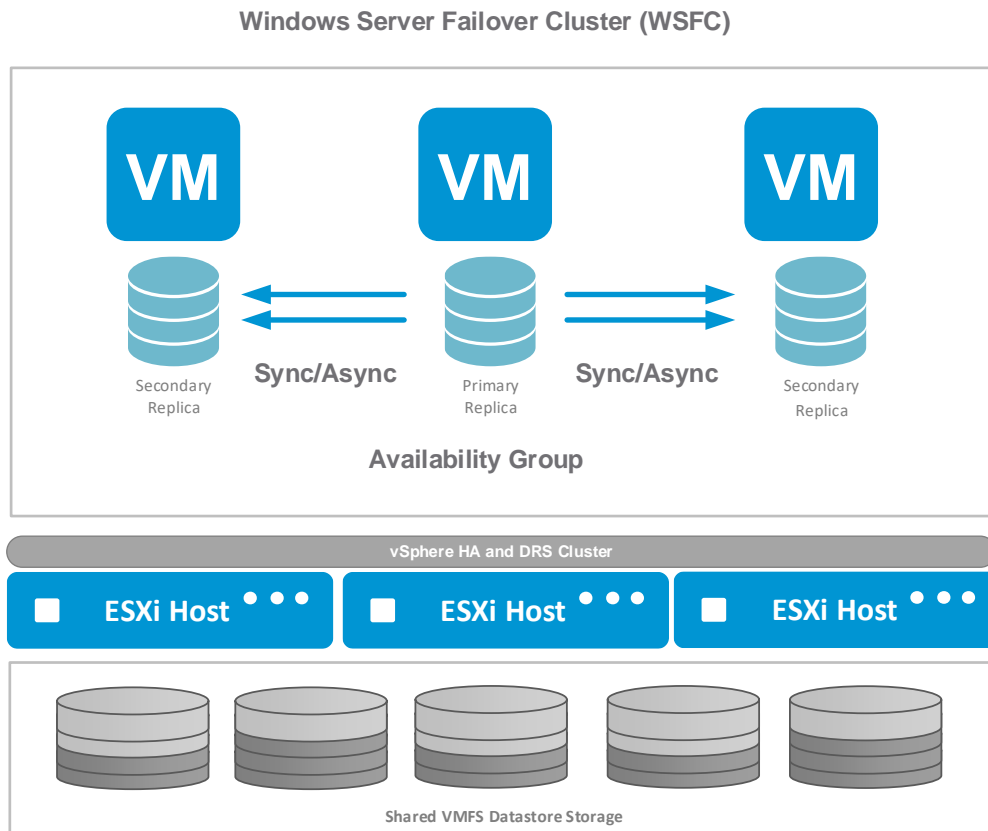
As mentioned previously, the AlwaysOn Availability Groups feature is built on a non-shared disk architecture. Each availability group replica has its own copy of the database, which can be deployed on VMFS or RDMS. Log streams are used to replicate data from the primary to the secondary replica nodes synchronously or asynchronously, with up to two synchronous replicas allowed. Availability groups sit on top of WSFC technology to provide a failover policy or quorum management. WSFC is required by



AlwaysOn Availability Groups, with each availability group replica also being a WSFC node, unlike AlwaysOn Failover Cluster Instances, for which there is no requirement to use shared disks.

While VMware does not support VMware vSphere vMotion® or DRS on clustered SQL Server virtual machines within a shared disk architecture (such as AlwaysOn Failover Cluster Instances), this restriction does not apply to AlwaysOn Availability Groups that are built on a non-shared disk architecture. VMware fully supports using vSphere HA, vSphere vMotion, and DRS with AlwaysOn Availability Groups.

Figure 2. Windows Server Failover Cluster with AlwaysOn Availability Group Configuration



With vSphere vMotion, a VMware ESXi™ host can be powered down for planned maintenance at any time without interruption to the client requests. vSphere HA can quickly reboot a Microsoft SQL Server virtual machine, which can rejoin the AlwaysOn Availability Group session if a hardware failure occurs.

In this planned maintenance scenario, vSphere vMotion is employed to proactively live migrate an availability group replica to a different host to allow the servicing of the hardware to take place without requiring an AlwaysOn Availability Group failover event. There is no disruption to the Microsoft SQL Server service during the vSphere vMotion operation and no interruption to the client's application connections or any in-flight transactions. By coupling together vSphere vMotion with AlwaysOn Availability Groups technologies, you can eliminate the need to fail over the Microsoft cluster and reduce service interruptions for operational hardware maintenance or renewal.

In the unplanned failure scenario, the SQL Server environment could be vulnerable if further host failures were to occur during the time between the loss of a replica and its restoration because, depending on available bandwidth and network conditions, the resynchronization of the replica could take a significant period of time to complete. vSphere HA helps to alleviate this issue by restarting the failed replica virtual machine on another available host in the vSphere cluster. This facilitates a significantly quicker restore to full protection of the database, and reduces the amount of time spent by the availability group in the

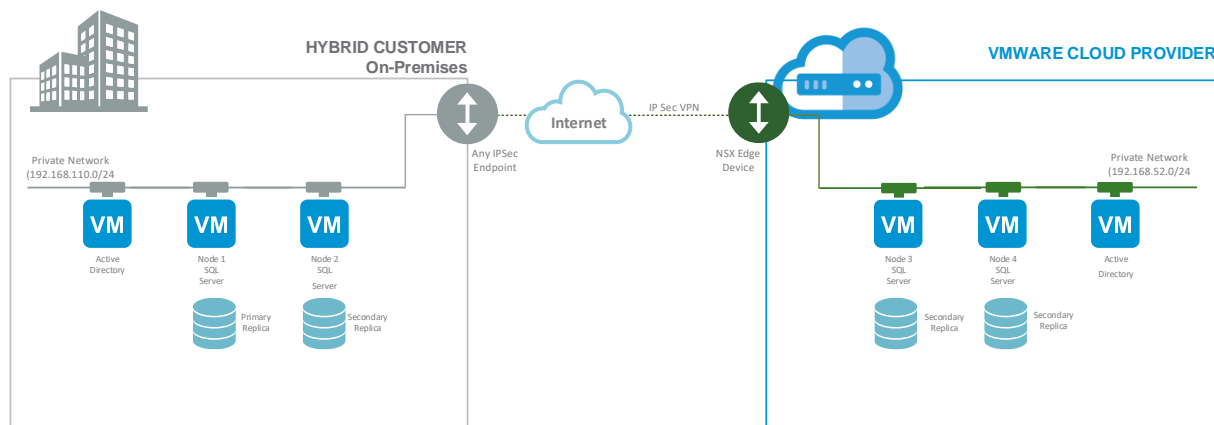


failover state. In the event of an unplanned physical host failure, you do not need to wait for the physical host to be serviced and brought back online to bring back the AlwaysOn Availability Group replica. Instead, vSphere HA automatically detects the host failure and immediately reboots the AlwaysOn Availability Group replica virtual machine onto a different available ESXi host.

With these integrated VMware mechanisms for high availability, it is perfectly possible, if not expected, to achieve better levels of uptime running AlwaysOn Availability Groups with vSphere than on physical hardware. This provides a high availability and disaster recovery solution for Tier 1 virtualized business critical databases to maximize continuous availability of the application it services and for business continuity during disaster scenarios. Furthermore, this solution architecture:

- Demonstrates business-critical levels of high performance and availability between the IT organizations private data centers and VMware Cloud Providers' facilities.
- Promotes resiliency to help meet recovery time and point objectives when faced with application, storage, network, or compute node failures.
- Demonstrates how to achieve business continuity SLAs in partnership with a VMware Cloud Provider to lower risk and operational costs.

Figure 3. Solution Logical Architecture



For the solution shown in the previous figure, which includes a cross-site IPsec VPN and Active Directory Sites and Services, consider the following:

- For an IPsec VPN, the endpoint networks cannot be the same.
- The IP Sec VPN mapping must use the correct VMware Cloud Provider Program edge gateway firewall rules to allow traffic to pass.
- Networks segments in the VMware Cloud Provider Program data center VM must be defined as a new site with the Sites and Services configuration of an Active Directory domain controller.

SQL Server 2016 AlwaysOn provides flexible design choices for selecting an appropriate solution for your application with protection against failures within a data center and adequate redundancy across data centers. For a further overview of the AlwaysOn design options for high availability and disaster recovery, see *AlwaysOn Availability Groups (SQL Server)* at <https://technet.microsoft.com/en-us/library/hh510230>.



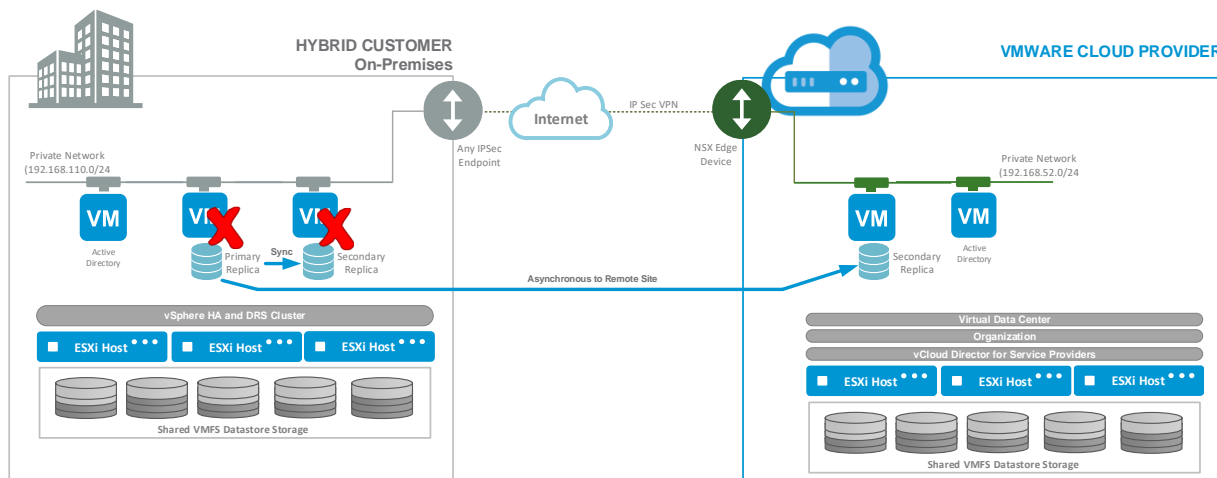
Employing SQL Server 2016 AlwaysOn Availability Groups for Disaster Recovery

The Microsoft SQL Server 2016 AlwaysOn Availability Groups ability to support multiple replicas with a mixture of synchronous and asynchronous data replication between primary and multiple mirrors, makes this technology an ideal solution for building a disaster recovery solution on top of the existing high availability solution.

In the sample design shown in the following figure, the on-premises data center hosts two availability replicas in synchronous commit mode for high availability. The VMware Cloud Provider Program facility hosts one replica, running in asynchronous mode for disaster recovery protection. In addition, each data center facility employs vSphere HA to protect virtual machines from ESXi host failure and to facilitate the rapid recovery caused by a locally contained outage.

This sample architecture allows for variations to this topology using multiple VMware Cloud Provider data centers and multiple replicas (up to five) within the solution architecture.

Figure 4. SQL Server 2012 AlwaysOn Availability Group for HA and DR





Designing the Solution

Configuring AlwaysOn Availability Groups requires:

- A minimum of two virtual servers (one primary and one secondary/replica) running, in the sample solution, Microsoft Windows Server 2012 R2 Enterprise Edition
- Microsoft SQL Server 2016 Enterprise Edition
- WSFC configured between the server nodes

Additional Microsoft prerequisites and requirements for this configuration are listed in the following table. For more information, see *Prerequisites, Restrictions, and Recommendations for AlwaysOn Availability Groups (SQL Server)* at <https://msdn.microsoft.com/en-us/library/ff878487.aspx>.

Table 2. Solution Prerequisites and Requirements

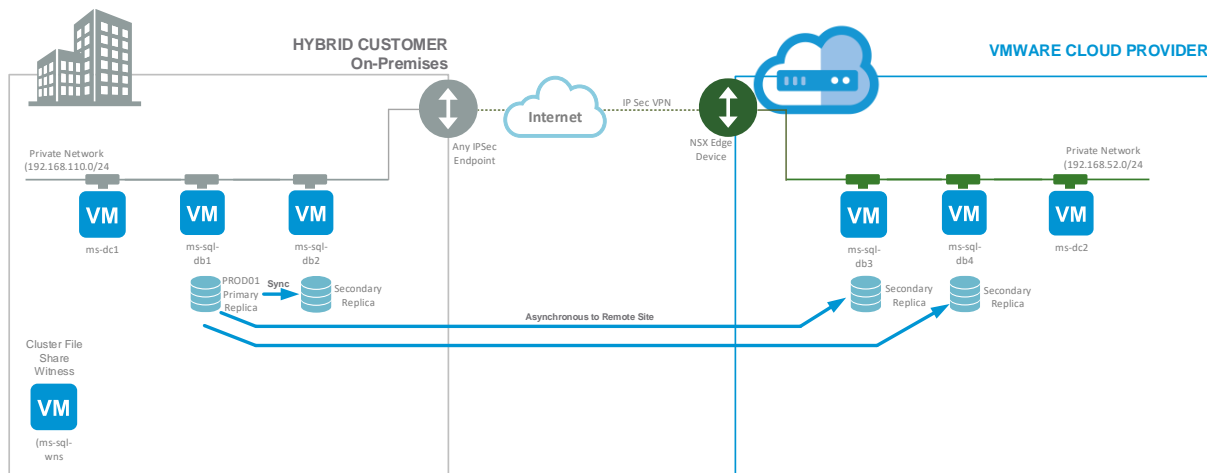
Component	Prerequisites or Requirements
Windows Server 2012 R2	<ul style="list-style-type: none"> • Availability groups are not supported on Active Directory domain controllers. • Each SQL compute node must be part of a Windows Server Failover Cluster. • Each server that hosts an availability replica must be provisioned with sufficient disk space to host all the databases in the availability group. • A WSFC node can only host one availability replica for a given availability group. On a given WSFC node, one or more instances of SQL Server can host availability replicas for many availability groups.
SQL Server 2016	<ul style="list-style-type: none"> • The instances of SQL Server that host availability replicas for a given availability group must reside on separate nodes of a single WSFC cluster. • Each server instance must be configured with the Enterprise Edition of SQL Server 2016. • The AlwaysOn Availability Groups feature must be enabled on each server node that will host an availability replica.
Availability groups	<ul style="list-style-type: none"> • The number of databases and availability groups you can put on a physical host or individual virtual machine depends on the virtual hardware provisioned and workload of the environment. • Each AlwaysOn Availability replica must be hosted by SQL server instances running on different nodes in the same WSFC cluster. • Each AlwaysOn Availability Group can support one primary replica and up to four secondary replicas. • Be sure to operationally test the environment, including failover scenarios, with a production-like peak-time workload, to validate the environment's ability to handle the capacity required, and to meet or exceed the application's SLAs.



Component	Prerequisites or Requirements
Availability database	<ul style="list-style-type: none"> • Must be a multi-user database. • Must be a read-write database. A read-only database cannot be added to AlwaysOn Availability Groups. • Cannot be configured for database mirroring. • Cannot belong to any existing availability group.

The sample solution architecture shown in the following figure is built on the Windows Server 2012 R2 Enterprise Edition operating system platform, with a four-node WSFC, consisting of a two-node cluster in the primary on-premises data center and a two-node cluster at the secondary VMware Cloud Provider Program provider facility using a stretched, active-active failover cluster configuration.

Figure 5. Sample Architecture Database Functions



The four SQL Server 2016 instances are configured with a single AlwaysOn Availability Group. This availability group was created with one primary replica and three secondary replicas. For more information about creating Availability Groups, see *Creation and Configuration of Availability Groups (SQL Server)* at <http://msdn.microsoft.com/en-us/library/ff878265.aspx>.

The following table describes the virtual machine functions in this sample solution architecture.

Table 3. Architecture Virtual Machine Function

Data Center Site	Virtual Machine	Function
On-Premises	Windows Server 2012 R2 (ms-dc1)	Provides DNS and Active Directory, services
On-Premises	SQL Server 2016 (ms-sql-db1)	Provides database services
On-Premises	SQL Server 2016 (ms-sql-db2)	Provides database services



Data Center Site	Virtual Machine	Function
On-Premises	SQL Server 2016 (ms-sql-wns)	Provides Cluster File Share Witness services
VMware Cloud Provider Program facility	Windows Server 2012 R2 (ms-dc2)	Provides DNS and Active Directory, services
VMware Cloud Provider Program facility	SQL Server 2016 (ms-sql-db3)	Provides database services
VMware Cloud Provider Program facility	SQL Server 2012 (ms-sql-db4)	Provides database services
VMware Cloud Provider Program facility	nsx-edge-02	VMware NSX Edge™ device

6.1 Distance and Latency Considerations

The data center link between the customer's on-premises and VMware Cloud Provider facility is a significant factor in this solution architecture. Therefore, it is important to have an understanding of factors such as link capacity, latency, and utilization when designing a high availability and disaster recovery solution. Typically, network latency corresponds to distance, and Round-Trip Time (RTT) can vary from 0 ms to 250 ms or higher. The following table provides distances and their corresponding link latencies. The values shown here are estimates and can vary significantly depending on interconnect quality, network hops, and physical backbone.

Table 4. Distance and Estimated Link Latency

Approximate Distance	One-Way Latency (ms)	Round-Trip Latency (ms)
50 km	0.25	0.5
500 km	2.5	5
1,000 km	5	10
5,000 km	25	50
10,000 km	50	100
20,000 km	100	200
22,500 km	125	250



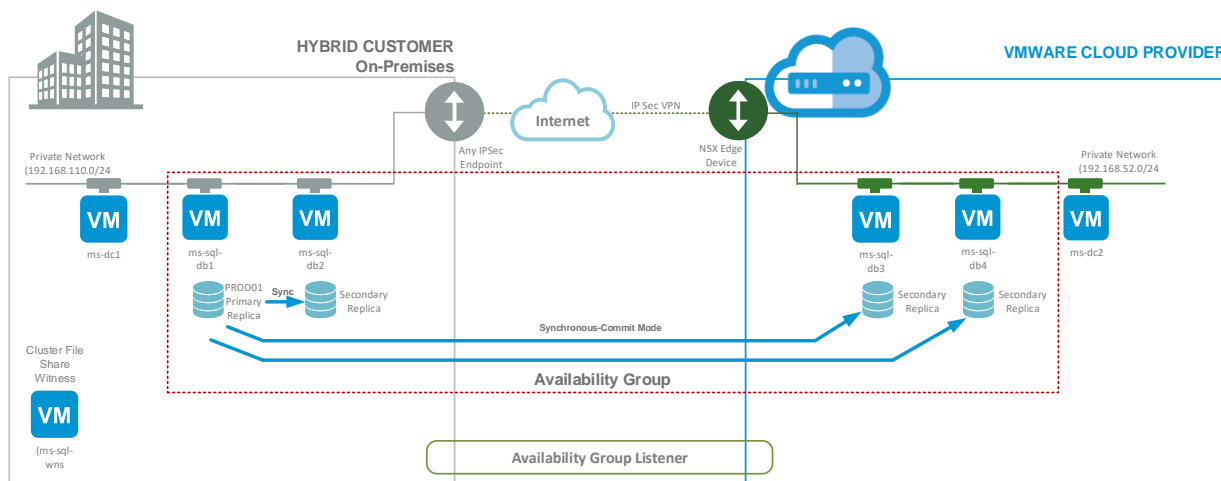
6.2 Sample Availability Group Design

AlwaysOn Availability Groups are configured with the availability group wizard. In the sample design shown in the following figure, the availability group is configured with a primary replica database (MS-SQL-DB1\PROD01) and one local secondary read-only replica (MS-SQL-DB2\PROD01) at the customer's on-premises data center. Two secondary replicas (MS-SQL-DB3\PROD01 and MS-SQL-DB4\PROD01) are configured at the VMware Cloud Provider facility for high availability and disaster recovery purposes. This architecture provides redundancy in multiple failure scenarios at both the on-premises and provider data centers.

The automatic failover mode option provides high availability by causing the database to fail over quickly and making it available after the loss of the primary replica. In the event of an unplanned outage, the failover takes place without manual intervention by operational teams or database administrators.

To configure an AlwaysOn Availability Group for automatic failover, the current primary replica of the database and one secondary replica must be set to synchronous-commit mode with automatic failover. Synchronized mode provides for little or no data loss, but requires more available network bandwidth, and therefore the environment must be within the minimum latency requirement limits set by Microsoft. In a typical architecture, to maximize performance and minimize data loss for disaster recovery purposes, configure all replicas for synchronous-commit mode with automatic failover.

Figure 6. Sample Availability Group Design





Architecting a Robust Technical Platform

When deploying Microsoft SQL Server AlwaysOn Availability Groups on a vSphere platform, consider the guidelines in this section to increase the solution availability and performance by architecting a robust technical platform.

7.1 Virtual Machine Placement

vSphere HA and DRS can automatically perform virtual machine placement for initial power on or when the cluster is under resource contention. To avoid a single point of failure when running AlwaysOn Availability Groups across virtual machines in a vSphere HA or DRS cluster, configure WSFC virtual machines on different physical hosts. The following practices can help to achieve this requirement:

- Create a DRS anti-affinity rule to keep virtual machines on different hosts.
- Enable strict enforcement of affinity rules.
- Set the DRS automation level for WSFC virtual machines to Partially Automatic.

If the advanced option for vSphere DRS, ForceAffinePoweron, is set to 1, strict enforcement of the anti-affinity rule is adhered to. VMware recommends using multiple server racks for the ESXi hosts and distributing the vSphere clusters across the server cabinets to minimize the impact of a single component failure.

Refer to *Setup for Failover Clustering and Microsoft Cluster Service* at <https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-setup-mscs.pdf> for detailed setup procedures.

To set up anti-affinity rules to keep virtual machines on separate hosts, follow the instructions found in *Affinity or anti-affinity DRS rules are not applied during a virtual machine power-on (2003128)* at <http://kb.vmware.com/kb/2003128>.

7.2 Disk Provisioning

Virtual machine disk files (VMDKs) can be deployed in three different formats—thin, thick, or eagerzeroedthick. Both thin and thick disk files use lazy zeroing, where the initial zeroing of the disk blocks is delayed until the first write. With eagerzeroedthick disks, the disk blocks are preallocated with zeros at the time of disk provisioning, and it is unnecessary to zero the disk on a first write basis during normal running operations. This provides approximately a 10–20 percent performance improvement over the other two disk formats.

Because Microsoft SQL Server high availability features are highly sensitive to system response time, the additional small overhead of disk zeroing during normal operations can cause unnecessary disk latency and potentially cause a false cluster failover event. If you are deploying AlwaysOn Availability Groups on any vSphere platform, define an operational standard to employ eagerzeroedthick disks for SQL Server data, transaction log, and tempdb files to mitigate the risks associated with non-eagerzeroedthick disks.

7.3 Network Considerations

The network is a critical component that is required for cluster node communication locally and across sites. In addition to normal network communications, Microsoft SQL Server in a non-shared disk, high availability solution, such as AlwaysOn Availability Groups, also utilizes the network for data replication between replicas. Especially in synchronous replication mode, the performance of the solution is highly dependent on the network bandwidth and latency.

Consider the following network configuration guidelines to achieve a robust technical platform and optimal performance:

- Configure appropriate network adapters on the host and virtual machines to separate networks used for different vSphere, virtual machine, and application purposes. For instance, you might want a



separate network for data replication, heartbeat, vSphere vMotion, management VMkernel, and so forth.

- If using iSCSI at either the host level or in-guest, the network adapters must be dedicated.
- Use the VMXNET3 paravirtualized NIC with all Windows Server 2012 R2 instances. The VMXNET3 is optimized for virtual environments and is designed to provide best performance.
- Enable jumbo frames for the iSCSI and vSphere vMotion network.
- Always employ static IP addresses for network interfaces 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.
- If deploying AlwaysOn Availability Groups in synchronous commit mode, use a high-speed network for replication traffic. Confirm that the bandwidth and latency of the network is sufficient to support the amount of Microsoft SQL Server transaction traffic.

7.4 Database Recovery Mode

Microsoft SQL Server non-shared disk, high availability solutions, such as AlwaysOn Availability Groups, database mirroring, and log shipping, use log stream or log record replication technologies for database redundancy. For this type of solution architecture, full recovery mode must be enabled on the database because log truncation does not occur until a log backup is taken. Confirm that sufficient disk space is planned for to allow for transaction log growth.

7.5 The Impacts of a Readable Secondary Replica

When Microsoft SQL Server 2016 AlwaysOn Availability Groups are employed, it is possible that a secondary replica can be made readable. The read-only access to a secondary replica is useful if your read-only workloads can tolerate some data latency. In situations where no data latency is acceptable, the most appropriate option is to run read-only workloads against 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. Take this into account when planning the overall design from an ESXi host workload perspective.

The workload on the secondary node is often completely different from the workload on the primary node. Typically, read-only queries are more analytical in nature. They usually generate a high number of sequential large block I/Os and do more aggregation operations, which can consume a large amount of CPU and memory. Database analytic type queries can also benefit from parallel processing threads, and therefore, additional vCPUs.

Due to the workload on a read-only secondary replica often being completely different than the workload on the primary, it might be necessary to configure additional indexes to get to optimal query plans. If additional indexes are required, they must be created on the primary replica to be transferred subsequently to the secondary replica. Plan for additional disk space, if needed, to accommodate the new indexes.

Finally, consider the impact of enabling read access on a secondary replica, as well as the characteristics of a read-only workload when capacity planning. Depending on your database characteristics, a readable secondary replica could demand more resources than the primary replica.

7.6 tempdb Design Considerations

The tempdb is a temporary work space for the Microsoft SQL Server itself. The tempdb is recreated every time the Microsoft SQL Server instance is started, so the system always starts with a clean copy of the database. There is never any data in tempdb that should be saved from one session of Microsoft SQL Server to another, so there is no need to include tempdb in a virtual machine level backup, or to replicate



tempdb in a multisite disaster recovery solution. Consider placing tempdb on its own dedicated VMDK to exclude it from any unnecessary operations.

When deploying AlwaysOn Availability Groups with a readable secondary replica, consider that tempdb space will increase due to the following:

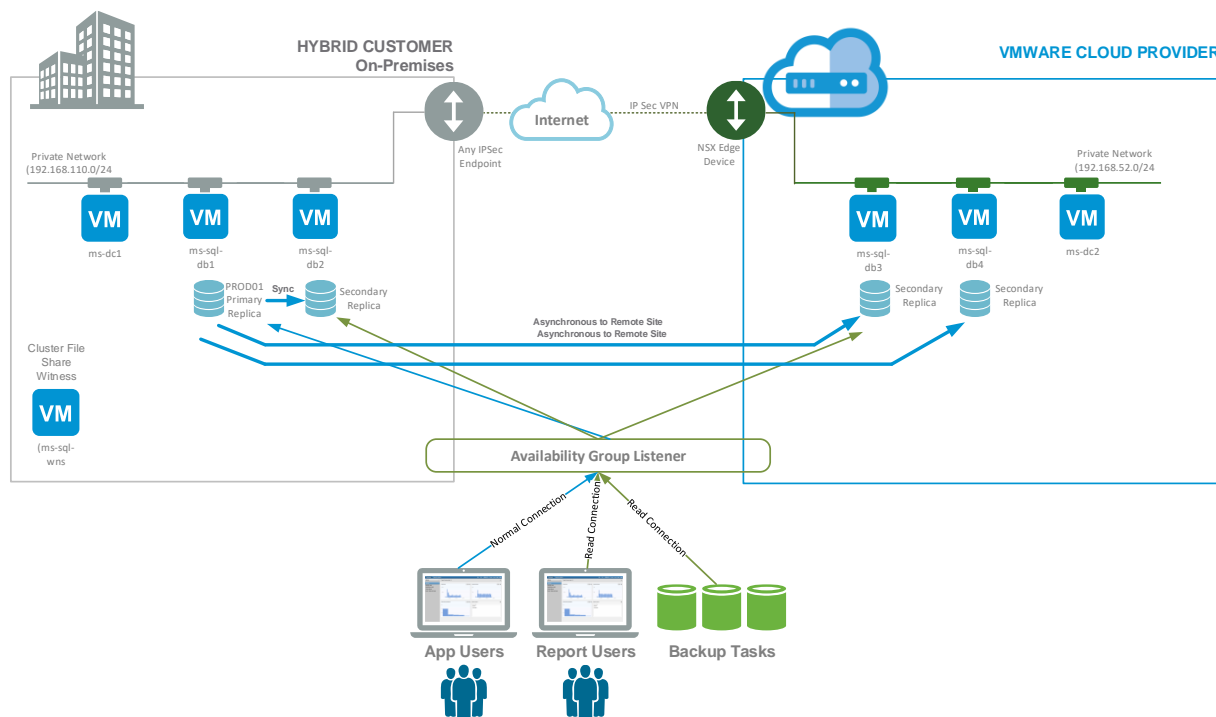
- Snapshot isolation level copies row versions into tempdb
- Temporary statistics that are created by the Microsoft SQL Server on secondary databases

Tempdb is a Microsoft SQL Server instance level resource. If you configure multiple databases in a single failover solution, the capacity of the tempdb must be able to accommodate demands from the aggregate of databases.

7.7 Client Connectivity

When deploying Microsoft SQL Server 2016 AlwaysOn Availability Groups, you can connect your application or users to the appropriate database replica using an availability group listener. An availability group listener is a virtual network name that is able to direct read-write requests to the primary replica and, if configured, read-only requests to the read-only secondary replica. VMware recommends always creating an availability group listener when deploying AlwaysOn Availability Groups on a vSphere platform. The availability group listener enables a client's connections to connect to an appropriate replica without knowing the name of the physical instance of the SQL Server installation. Therefore, the client connection string does not need to be modified in case of a failover.

Figure 7. Availability Group Listener Client Connectivity



For more information about creating an availability group listener, see *Prerequisites, Restrictions, and Recommendations for AlwaysOn Availability Groups (SQL Server)* at <http://technet.microsoft.com/en-us/library/ff878487.aspx>.



7.8 Considerations for Quorum Mode

The quorum is an important component for any WSFC high availability solution. If you deploy any of the Microsoft SQL Server 2016 AlwaysOn solutions, you can take advantage of the newer quorum models available in Windows Server 2008 or later. The quorum models are as follows:

- Node Majority – This model requires an odd number of nodes and is less common.
- Node and Disk Majority – This is a combination of Node and Quorum disk. This quorum model requires using an RDM for the quorum disk. This is not ideal due to the limitations of RDM disks.
- 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 traditional Windows Server 2003 quorum disk model. VMware recommends discontinuing use of this model.

Because each customer architecture is different, consider the specific use case when deciding which quorum model is most appropriate for a specific solution. The Node Majority and the Node and File Majority models do not require an RDM disk for the quorum. This could provide additional flexibility and compatibility with vSphere, because vSphere vMotion, DRS, and HA features are fully supported when deploying AlwaysOn Availability Groups with these quorum modes.

See *Understanding Quorum Configurations in a Failover Cluster* at <https://technet.microsoft.com/en-us/library/cc731739.aspx> for additional information about choosing the appropriate quorum model for the cluster architecture.



Conclusion

The goal of this solution architecture document is to describe how to design a highly available and disaster recovery solution for virtualized SQL business crucial databases. This document provided a validated business continuity solution for Microsoft SQL Server 2016 using the high availability and disaster recovery capabilities of AlwaysOn Availability Groups across a distributed architecture, combining both an on-premises and VMware Cloud Provider data center facility.

The AlwaysOn Availability Groups feature is the ideal solution for deploying highly available Microsoft SQL Server 2016 instances on a vSphere platform. AlwaysOn Availability Groups provide out-of-the-box protection for hardware, software, and data failure as well as additional capabilities to offload backups and read-only reporting queries to a secondary replica. With its non-shared disk architecture, AlwaysOn Availability Groups can be used safely with vSphere vMotion, DRS, and vSphere HA to reduce downtime and improve flexibility in your application architecture, while lowering costs and minimizing the need to perform an SQL failover. This solution can significantly reduce Microsoft SQL Server recovery time from hours or days to minutes or even seconds. Consider the following key architectural points when designing a solution:

- Remember that AlwaysOn Availability Groups and database mirroring can be used in combination with vSphere vMotion, vSphere HA, and DRS to maximize database availability.
- ESXi 5.1, 5.5 and 6.0 support up to five-node clusters for Windows Server 2008 SP2 and later, whereas earlier ESXi versions support only two-node clusters.
- To avoid single points of failure, use vSphere anti-affinity rules to run AlwaysOn Availability Group replica virtual machines on separate hosts.
- SQL Server AlwaysOn Availability Groups on vSphere are supported for non-shared disk configurations, except when the system disk's VMDK is located on an NFS datastore.

With the combined power of the VMware Cloud Provider Program facility and virtualized Microsoft SQL Server 2016, an IT organization can demonstrate business-critical levels of performance and availability of database platforms, while accelerating the time to deployment and simplifying IT operations.



Assumptions and Caveats

VMware, Microsoft and other third-party hardware and software information provided in this document is based on the current performance estimates and feature capabilities of the versions of code indicated. These are subject to change by their respective vendors.

Reference Documents

The following documents are available for additional information:

Document Title	Link or URL
<i>Virtualizing Microsoft SQL on VMware vSphere</i>	https://www.vmware.com/solutions/business-critical-apps/sql-virtualization.html
<i>Microsoft Clustering on VMware vSphere: Guidelines for supported configurations (1037959)</i>	http://kb.vmware.com/kb/1037959
<i>Active Secondaries: Readable Secondary Replicas (AlwaysOn Availability Groups)</i>	https://msdn.microsoft.com/en-GB/library/ff878253.aspx
<i>Understanding Quorum Configurations in a Failover Cluster</i>	https://technet.microsoft.com/en-us/library/cc731739.aspx
<i>Create or Configure an Availability Group Listener (SQL Server)</i>	https://technet.microsoft.com/en-us/library/hh213080.aspx
<i>Setup for Failover Clustering and Microsoft Cluster Service</i>	https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-setup-mscs.pdf
<i>Creation and Configuration of Availability Groups (SQL Server)</i>	http://msdn.microsoft.com/en-us/library/ff878265.aspx
<i>Prerequisites, Restrictions, and Recommendations for AlwaysOn Availability Groups (SQL Server)</i>	https://msdn.microsoft.com/en-us/library/ff878487.aspx
<i>AlwaysOn Availability Groups (SQL Server)</i>	https://technet.microsoft.com/en-us/library/hh510230