

Business Continuity for Microsoft® SQL Server® and VMware® on Dell™ PowerEdge™ Servers and Storage

A Dell Technical White Paper

Database Solutions Engineering

By Anthony Fernandez

Dell Product Group

January 2010



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2009 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell, Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, the *DELL* badge, *PowerEdge*, *EqualLogic*, *PowerConnect* are trademarks of Dell Inc.; Intel and *Xeon* are registered trademarks of Intel Corporation in the U.S. and other countries; *Microsoft*, *SQL Server*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *VMware*, *vSphere*, *VMotion* are either trademarks or registered trademarks of VMware, Inc. in the United States and/or other jurisdictions. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Executive Summary

Virtualization technologies are a key component used by enterprises to meet their utilization goals, while reducing costs such as power and cooling associated with server sprawl. A major culprit for server sprawl historically is the Relational Database Management System or RDBMS. Microsoft® SQL Server® is a perfect example of an RDBMS that can quickly proliferate through an organization; to combat this sprawl, many organizations implement virtualization strategies at the database server layer. In the physical domain, business continuity strategies are well understood by IT organizations across all the tiers in the data center (Network, Web, Application, Database, and Storage). In the virtual domain, new challenges arise to ensure fault tolerance (FT) and high availability (HA) of Virtual Machines (VMs), especially at the database layer.

This white paper describes how to protect Microsoft® SQL Server® in a VMware® VM. It describes the challenges, techniques, and methods to provide high availability for a SQL Server VM along with best practices recommendations for small to large implementations.

CONTENTS

Executive Summary 3

Introduction..... 5

 Audience and Scope 5

Overview of Business Continuity Options 6

 Hardware Options 6

 Fault Tolerance 6

 Software Solutions 7

 High Availability 7

 VMWare Options..... 7

 Database Options 8

Define Levels of Protection 8

Pros and Cons: 11

Developing a Strategy..... 13

 Stand-alone Solution 15

 High Availability Solution..... 17

 High Availability + VMotion + VMware FT Solution 22

Conclusion 24

References 25

Introduction

Relational Database Management Systems or RDBMS are a central component of an organizations IT department. As part of a business continuity plan, implementing Fault Tolerance (FT) and High Availability (HA) techniques to help reduce database downtime is crucial to a successful enterprise.

Virtualizing database servers offers many advantages including better hardware utilization by de-coupling software from physical hardware. However in addition to the benefits gained, other challenges may arise including ways to protect databases running on virtual machines, methods for minimizing downtime, and efforts to reduce complexity.

This white paper describes an overview of technologies to help protect Microsoft SQL Server databases in a virtual environment. We explore VMware HA, VMware FT, Microsoft Cluster Server System (MSCS) and Microsoft SQL Server native options.

Utilizing the Dell Business Ready Configuration (BRC) for vSphere 4, we describe some of the key considerations to implement a highly available database server environment from a stand-alone configuration for small implementations to sizable, highly available solutions for large enterprise. In addition, these configurations are designed with ease of scale out from a stand-alone configuration to a multimode cluster configuration. This model provides the benefit of re-purposing hardware into a larger configuration.

Audience and Scope

This white paper is intended for customers, partners, solution architects, storage and database administrators, who are evaluating, sizing and planning a virtualization strategy. It provides an overview of various considerations and best practices to plan high availability for virtualized databases.

Overview of Business Continuity Options

Building a business continuity plan for a virtual environment essentially leverages already established concepts for the physical domain. However, in order to minimize planned or unplanned downtime and reduce RTO for Virtual Database servers, administrators face the challenge of designing and implementing cost-effective solutions while reducing complexity.

Below we focus on mainly three categories; hardware, hypervisor features, and database application features.

Hardware Options

Fault Tolerance

Fault tolerance (FT) refers to the concept of maintaining system uptime during a component failure. In order to maintain system uptime, redundancy is implemented at every level. Examples include: storage RAID, dual storage controllers, dual network cards, redundant data paths, dual power supplies, etc.

Server: The Dell PowerEdge R510 is a 2 socket 2U server with 4, 8 or 12 internal hard drives. For stand-alone database servers without external storage, it is recommended to use the 12 internal hard drives and redundant power supplies. For database servers with external storage, 8 drive configurations provide redundant power supplies. Note: The 4 internal drive server supports single power supply.

Internal RAID Controller: The R510 comes with PERC6/I and PERC H700. For database loads that require higher bandwidth the PERC H700 is recommended for its support of SAS/SATA 6Gb/s performance.

Memory: Up to 64GB (8 DIMM slots) 1GB/2GB/4GB/8GB DDR3 800MHz, 1066MHz or 1333MHz. Recommended to use RDIMMS for better reliability, availability and serviceability (RAS) since Single Device Data Correction (SDDC) is supported in Optimizer Mode¹.

Network: The embedded LOM Broadcom BCM5716 provides dual-port GigE in a single-chip solution. Therefore, it is recommended to add an add-on Ethernet NIC card to the PCIe Riser and enabling NIC Teaming for improved fault tolerance. The PCIe Riser contains 3 slots (Slot 1 and 2 are x4, and Slot 3 is x8). Use Slot 3 for large sequential IO

¹ Optimizer Mode is the default memory configuration. This mode permits larger total memory capacity and maximum performance. For increased RAS Lockstep Mode (Advanced ECC Mode) and Mirrored Mode are also available with reduced capacity and performance.

workloads. Redundant Switches are also recommended to ensure continuous communication.

Storage: EqualLogic PS6000XV is an iSCSI Storage Area Network (SAN) with dual controllers, Disk RAID, and multiple data paths for redundancy.

Software Solutions

High Availability

High Availability (HA) refers to the concept of improving the availability of systems during planned or unplanned downtime. HA helps organizations adhere to Service Level Agreements (SLAs) where the Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are defined. RPO defines the point in time to where data can be recovered in the event of failure and RTO defines the time required to bring the system to a running state at the point defined by the RPO.

Various 3rd party software solutions exist in the marketplace that provide High Availability; Altiris, Veritas, Marathon Technologies to name a few. For the context of this white paper we focus primarily on Microsoft Cluster Server.

Failover Clustering: High Availability is typically synonymous with clustering. Microsoft Cluster Service or MSCS provides a failover cluster solution that protects against physical hardware failures, Operating System, and Application failures. The server nodes are configured in an active/passive mode. Applications running on the cluster are registered as a clustered service. If a failure occurs, the application service is then restarted on the passive node.

VMware Options

Introducing a virtualization layer offers various benefits such as workload consolidation and the decoupling of applications from physical hardware. VMware offers two technologies to protect virtual machines.

VMware HA: Provides an easy to implement solution to continuously monitor all physical servers in a resource pool to protect Virtual Machines in the event of hardware failure and the restarting of VMs on the surviving servers depending on resource capacity.

VMware FT: is a new technology on vSphere v4. It provides fault tolerance to VMs even after a physical server failure. VMware vLockstep technology provides an active copy of the VM on a secondary server. In the event of server failure, the VM continues to run uninterrupted.

Database Options

It is recommended to continue utilizing the HA features provided by the database engine. Below is a summary of each of the options available. For more information please visit www.dell.com/sql

Replication: Provides many benefits such as multiple copies of a database at the cost of complexity. Replication is often used on geographically separated sites due to the low latency nature of the transactions. In addition, since copies of the database are available on the subscribed server(s), it allows scalability for read operations as well as real-time availability.

Log Shipping: Allows for logs for a specific database to be shipped from a primary database server to one or more secondary server(s). It is a flexible option that has many advantages such as multiple destinations on which to ship the logs simultaneously, ease of use, and simple implementation. Log shipping can easily be deployed in conjunction with Replication or Mirroring for a more robust solution.

Mirroring: Database Mirroring provides an easy to implement solution to maintain a copy of a database between two database servers; a principal and a mirror server. It provides automatic failover mechanism in high-safety mode that allows the copy database to be brought online quickly. In addition, Automatic Page Repair was introduced with SQL Server 2008 (Std/EE Ed). In the event that one of the partners is unable to read a page due to corruption, it first tries to request the page from the other partner. If it succeeds, the corrupted page is replaced automatically. For more information on Automatic Page Repair, please visit <http://msdn.microsoft.com/en-us/library/bb677167.aspx>

Define Levels of Protection

The first step in determining the type of protection required is to perform a complete assessment of database servers that are hosted in the datacenter. In the table below (table 1) we have defined three general categories for protection and a proposed a solution or combination of solutions for each category.

Level of Protection	Description	Recommendation	Observations
<p>Low Protection (Stand-Alone)</p>	<p>For Test/Development database servers. If the database server crashes due to a test case, the developer/tester requires full access to the Operating System to trouble shoot, debug and restart the application within the VM. If the OS fails (blue screen), since developers typically do not manage vCenter to do cold reboot, it requires a mechanism to restart the failed VM on the same or different node.</p>	<p>Local Site Only: 1) VMware HA</p>	<p>Single R510 can be implemented. VMware HA will restart failed VMs within the same server. May introduce Single Point of Failure if server is down. Potential for work loss and/or productivity.</p> <p>Recommended for an environment where no risk of data loss and few VMs.</p> <p>No requirement for VMotion².</p>

² VMotion is available on Standard, Enterprise and Enterprise Plus Editions.

Level of Protection	Description	Recommendation	Observations
<p>Medium Protection</p> <p>(High Availability)</p>	<p>For database servers that serve user applications that require medium to high availability. Data is critical, but can tolerate some downtime.</p>	<p>Local Site:</p> <p>1) VMware HA + SQL Server Mirroring</p> <p>With Remote Site:</p> <p>2) VMware HA + SQL Server Mirroring + External Site with Log Shipping.</p>	<p>Recommended two R510 servers.</p> <p>Requires at least two VMs for Mirror/Partner configuration. Third Witness VM can be added to provide automatic failover.</p> <p>No requirement for VMotion².</p> <p>Optional external storage SAN/DAS.</p> <p>For Sites that require Remote Site, a third server can be added to implement Log Shipping or SQL Server Replication.</p>

Level of Protection	Description	Recommendation	Observations
High Protection (High Availability + VMotion + FT)	Production environments. Requires No SPOF and HA.	Local Site: 1) VMware FT + SQL Server Mirroring Remote Site: 2) VMware FT + SQL Server Mirroring + Log Shipping.	Requires External Storage Recommended three R510 servers. Protects against physical host failures. SQL Server Mirroring with Witness protects against both FT VMs downtime due to Operating System failure. Log shipping or SQL Server Replication can be implemented for remote site.

Table 1 Defined Levels of Protection

Pros and Cons:

The table below describes some of the pros and cons of each technology.

	Pros	Cons
VMware HA	1) Easy to implement. 2) Application and OS agnostic. 3) No agent or configuration required within the VM.	1) Does not protect against application failure. OS failure can be detected with heartbeat. 2) Full VM Restart. 3) Requires Shared Storage. 4) Data on Storage SPOF.

	Pros	Cons
VMware FT	<ol style="list-style-type: none"> 1) Provides transparent failover in the event of hardware or component failure. 2) If FT VMs fail, VMware HA restarts VMs. 3) Application and OS agnostic. 4) Supports iSCSI SAN Storage. 	<ol style="list-style-type: none"> 1) If OS crashes on primary VM, the secondary will also crash. 2) No support for SMP VMs. Single vCPU only. 3) Requires uniform hardware and consistent power and load to be balanced. 4) No support for new PVSCSI or VMXNET3. 5) Data on Storage SPOF.
Microsoft Cluster Server	<ol style="list-style-type: none"> 1) Provides application based failover. No server restart. 2) Proven technology 3) Can cluster Physical servers to VMs. 	<ol style="list-style-type: none"> 1) Complex to implement. 2) No support for iSCSI SAN Storage, only Fibre Channel SAN is supported³. 3) No VMotion support. 4) No support for VMFS for shared disks across two physical servers, only Pass-through Raw Device Mapping (RDM). 5) No support for PVSCSI⁴. 6) Max two nodes per Cluster. 7) Data on Storage SPOF.

³ Microsoft Cluster Server is limited to vSphere environments with external Fibre Channel SANs.

See the Reference page for ***Setup for Failover and Microsoft Cluster Service***.

⁴ See the Reference page for additional information on Paravirtualized SCSI (PVSCSI) devices.

	Pros	Cons
SQL Server Mirroring	<ol style="list-style-type: none"> 1) Provides redundant copy of database. 2) No shared storage required. 3) Automatic Page Repair. 	<ol style="list-style-type: none"> 1) Requires additional storage capacity. 2) Automatic Failover only with Witness (ideally on separate host). 3) Application may need to be re-configured for failover.

Table 2 Technology Pros and Cons

Note: VMs that are protected with VMware FT do not support other features such as Snapshots, Storage VMotion, or DRS Load Balancing Primary or Secondary VMs. These operations should be performed during windows where these maintenance tasks can be performed without disruption to users or application activities. During these windows, VMware FT can be turned off.

Developing a Strategy

The first step in developing a high availability strategy to protect a virtual database is to determine the type of data for which the system will be utilized. Once the criticality of the data has been determined, the second step is building the hardware configuration.

The table below (table 3) describes each of the configurations that have been preconfigured to facilitate ordering.

Server Configuration	Stand-alone	High Availability	High Availability + VMotion
Dell PowerEdge Server	One PowerEdge R510	Two PowerEdge R510	Three PowerEdge R510
Processors	(2) x Intel® Xeon® 5500 Series Quad Core Processors	(2) x Intel® Xeon® 5500 Series Quad Core Processors	(2) x Intel® Xeon® 5500 Series Quad Core Processors
Memory	32 GB (8 x 4GB, DDR3 RDIMMs)	32 GB (8 x 4GB, DDR3 RDIMMs)	32 GB (8 x 4GB, DDR3 RDIMMs)
Internal Hard Drives	(12) x 300GB, 15K SAS	(2) x 146GB, 15K SAS	(2) x 146GB, 15K SAS

Server Configuration	Stand-alone	High Availability	High Availability + VMotion
Hard Drive Controller	PERC 6/I Optional Upgrade to PERC H700 6GB/s for heavy IO workloads	PERC 6/I	PERC 6/I
BIOS Settings	Performance BIOS Setting, VT enabled Execute Disable option enabled	Performance BIOS Setting, VT enabled Execute Disable option enabled	Performance BIOS Setting, VT enabled Execute Disable option enabled
Remote Management Controller	iDRAC6 Express	iDRAC6 Express	iDRAC6 Enterprise
Operating System	VMware vSphere 4.0 U1 (Essentials Plus, or Standard)	VMware vSphere 4.0 U1 (Essentials Plus, or Standard)	VMware vSphere 4.0 U1 (Advanced, Enterprise or Enterprise Plus)
Additional Network Cards	(1) x Broadcom 5709 Dual Port GbE PCI/e Card.	(2) x Broadcom 5709 Dual Port GbE PCI/e Card.	(2) x Broadcom 5709 Dual Port GbE PCI/e Card.
External Storage	None	(1) x Dell EqualLogic PS6000XV	(1+) x Dell EqualLogic PS6000XV
Drives	N/A	(16) x 300GB, 15K SAS (16) x 450GB, 15K SAS (16) x 600GB, 15K SAS	(16) x 300GB, 15K SAS (16) x 450GB, 15K SAS (16) x 600GB, 15K SAS
Raw Storage Capacity	N/A	4.2TB using 300GB SAS disk drives 6.3TB using 450GB SAS disk drives 8.4TB using 600GB SAS disk drives	4.2TB using 300GB SAS disk drives 6.3TB using 450GB SAS disk drives 8.4TB using 600GB SAS disk drives

Table 3 Business Ready Configurations

Stand-alone Solution

A stand-alone solution is built for customers who are interested in getting started with implementing a virtualization infrastructure or simply require an additional server to deploy Virtual Machines without the requirement for external storage.

The stand-alone configuration offers a quick way to start testing virtual machines in a test or development environment. It offers large internal storage capacity and up to 32GB of RAM with a single processor configuration or 64GB with dual processor configurations.

One of the main considerations for this configuration is the challenge to protect the virtual databases and or other applications that running on this system. In the event of server downtime, all the VMs become in-accessible. It is recommended to implement this configuration in an environment where application downtime does not disrupt operations. Figure 1 below shows some of the main characteristics for the stand-alone solution.




Figure 1 PowerEdge R510 Front/Back

Internal Storage Configuration

Depending on the size of database files that will reside on the system, configuring the internal storage with three Virtual Disks can offer the separation and capacity to run multiple VMs. If performance is not a requirement, RAID5 will offer the best cost per Giga Byte (\$/GB). The table below (table 4) provides applicable hard drive specifications and capacities.

	Hard Drive Specifications	Total Capacity
ESX Operating System	(2) x 146GB 15K, RAID 1	146GB
Virtual Machine Disks (VMDK)	(2) x 146GB 15K, RAID 1	146GB
Database Files (Performance)	(8) x 300GB 15K, RAID 10	1.2TB
	(8) x 450GB 15K, RAID 10	1.8TB
	(8) x 600GB 15K, RAID 10	2.4TB

	Hard Drive Specifications	Total Capacity
Database Files (Capacity)	(7+1) x 300GB 15K, RAID5	2.1TB
	(7+1) x 450GB 15K, RAID5	3.1TB
	(7+1) x 600GB 15K, RAID5	4.2TB

Table 4 Hard Drive Specifications and Capacities

Figure 2 below shows the location for each of the Virtual Disks that can be implemented in this configuration. The R510 offers a large number of internal hard drives that allows for multiple Virtual Disks without requiring external storage.




Figure 2 Logical Disk Groups

Figure 3 below shows Virtual Machines running on the stand-alone configuration.

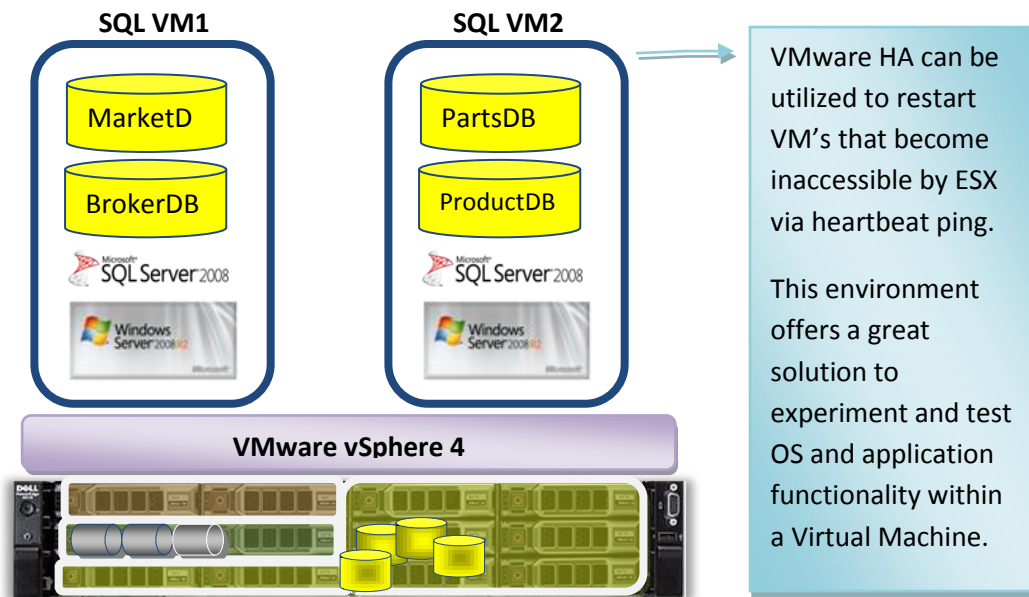


Figure 3 Stand-Alone Configuration

Another benefit of the stand-alone configuration is that it can easily be expanded to provide HA and FT features in larger configurations.

High Availability Solution

A High Availability Configuration is composed of two R510 servers running ESX Essentials Plus or Standard Edition. It can be configured as two stand-alone servers with no additional external storage or two R510 servers with external storage; these two use cases are described below. The first use case does not require Storage Area Network (SAN), but can be augmented with raw IO by adding a Direct Attached Storage (DAS) to increase the number of spindles. The second use case utilizes a SAN and fully leverages the VMware HA cluster feature to protect VMs with the added benefit of SAN based availability features.

User Case 1: Configuration without SAN

This configuration offers a quick way to mitigate the single point of failure of the Stand-Alone configuration. By simply adding a second R510 to the network, VMs can easily be moved by copying them offline to local storage and then restarted.

Considerations: VMware HA is limited to restarting the VM on the same physical server where the VMDK is located. Downtime is a factor of the time it takes to get the server up and running. SQL Server Mirroring mitigates this by placing the mirrored partner on a separate VM on the second node with its own copy of the database. The application is required to be reconfigured to recognize the failover and re-connect to the partner VM once the failure occurs. Another limitation of this configuration is scalability. The internal hard drives offer a large capacity with limited transactional performance. Due to the limited number of internal hard drives assigned to the database; 8 Hard Drives in RAID 1/0 or RAID 5, it is recommended to maintain workloads that achieve reasonable IO Latency times for the application. Ensure that sufficient memory is utilized by the database to mitigate high latency times for transactional databases.

Note: If no other feature such as VMware HA cluster, VMotion, VMware FT or SAN based protection is required, another way to mitigate the internal drive limitation is to add an external Direct Attached Storage (DAS) array to expand the raw IO capabilities of this configuration.

Figure 4 shows the High Availability Configuration with internal Storage only. In this example, each server is hosting a single VM that contains two user databases. Each database server is configured in a Mirror session, with primary and partner databases simultaneously. Two client applications are connected to each database instance; VM1 and VM2. Writes to VM2 BrokerDB are issued to the local datafile as well as on the Partner VM1 as depicted by the blue arrows.

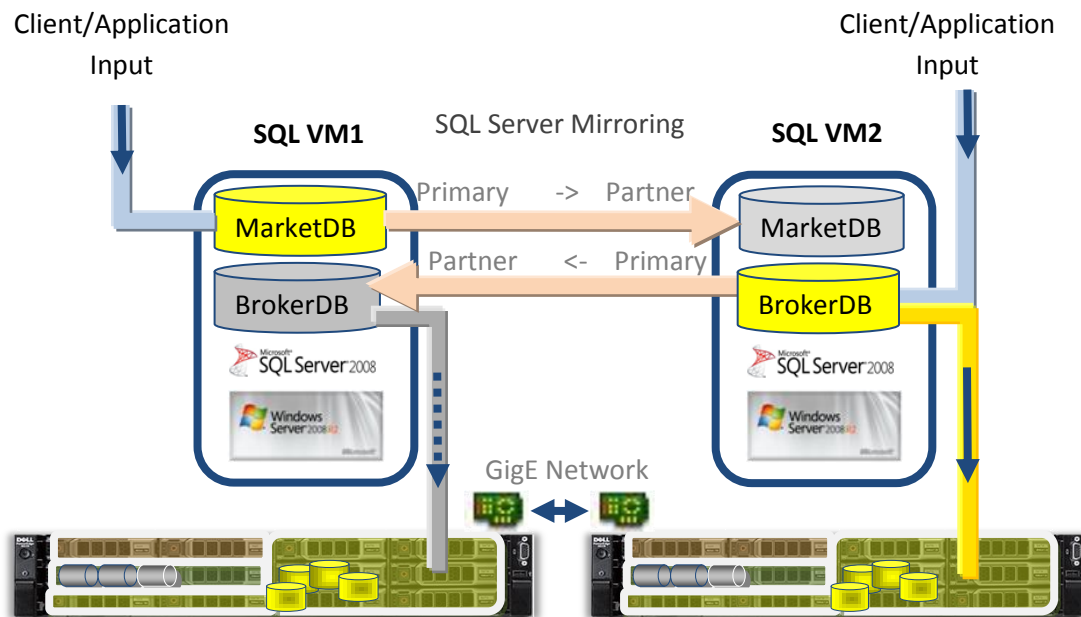


Figure 4 High Availability with SQL Server Mirroring

Figure 5 below shows the failover case. In the event VM1 goes down, the mirror session is broken. Since VM1 hosted a principal database; MarketDB and a Mirror database; BrokerDB. The Broker DB, the Principal database on VM2 is unaffected and continues operation without a mirror partner. However, for MarketDB since it was a principal, the Mirror partner on VM2 needs to be either manually set to be the principal or if automatically if deployed a third VM as a witness.

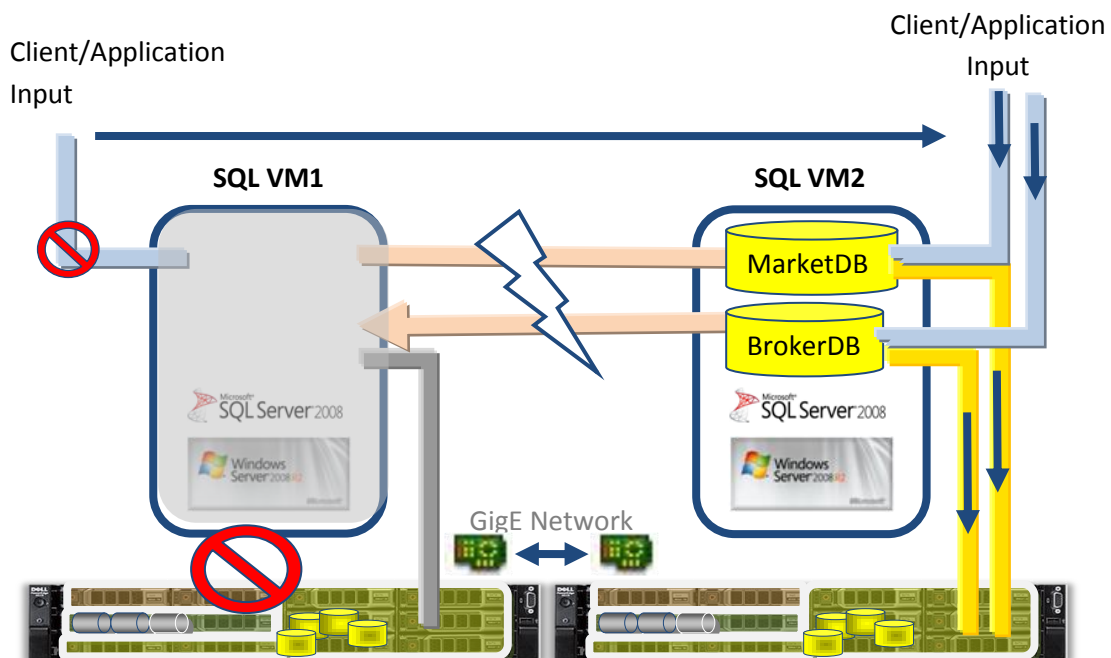


Figure 5 High Availability with VMware HA and SQL Server Mirroring - Failover

Use Case 2: Configuration with external SAN

Each node is configured as part of a VMware HA cluster. The external shared storage configuration offers the capability to restart VMs on the second node without the need to copy the VMDK files manually over the network. It also offers the benefit of decoupling data from the physical server. In the event of hardware replacement, with data residing on shared storage, databases can easily be mounted on a separate instance. In addition, shared storage offers another layer of performance, scalability and protection with features such as storage replication, snapshots and cloning.

Considerations: It is best practice to configure the connectivity to the external storage with redundant connections. In this configuration, EQL PS6000XV is used with dual PowerConnect Ethernet Switches. To maximize bandwidth, both switches are configured with Link Aggregation (LAG). In addition, each server has an equal number of NICs configured for iSCSI as the number of ports on the storage controller. Each iSCSI link from the server is placed on separate PCIe buses to ensure availability. For more EQL Network Best Practices recommendations, please see Reference page.

Figure 6 below shows the cabling configuration for the High Availability Configuration with External Storage and redundant iSCSI network.




Figure 6 High Availability Cable Diagram with SAN

In the event of physical hardware or operating system failure, VMware HA restarts the VM on the second node. To mitigate the downtime sustained by a cold restart of the VM, SQL Server Mirroring is configured with a witness server to automatically failover to the mirror partner. Figure below shows a logical diagram of a two node VMware HA cluster and two Virtual Machines configured with SQL Server Mirroring. Three storage volumes were created to hold Virtual Machine VMDK files, and database files for each VM. Note that no internal hard drives need to be configured for this configuration.

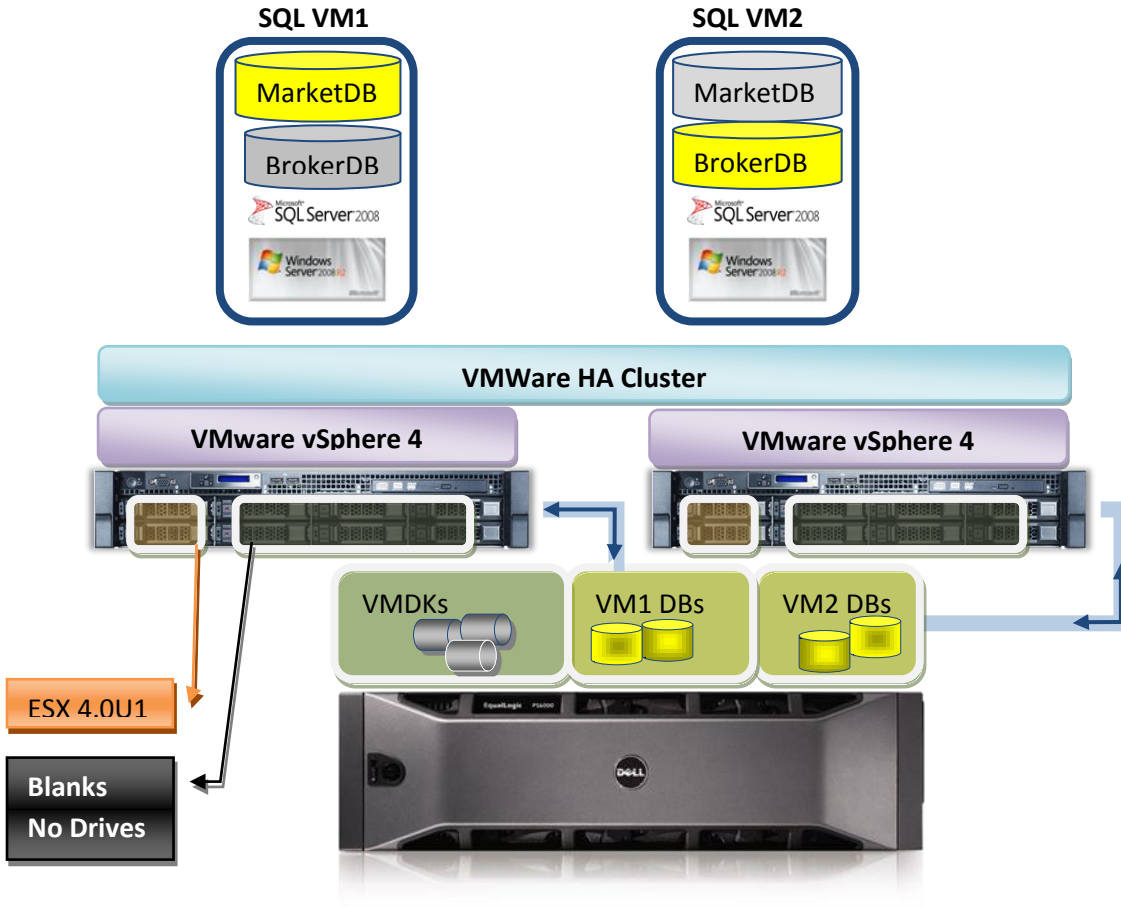


Figure 7 High Availability with SAN, VMware HA and SQL Server Mirroring

In the event of physical or VM failure, SQL Server mirroring fails over to the mirror partner. VMware HA restarts VM1 on a surviving node described in Figure 8 below.

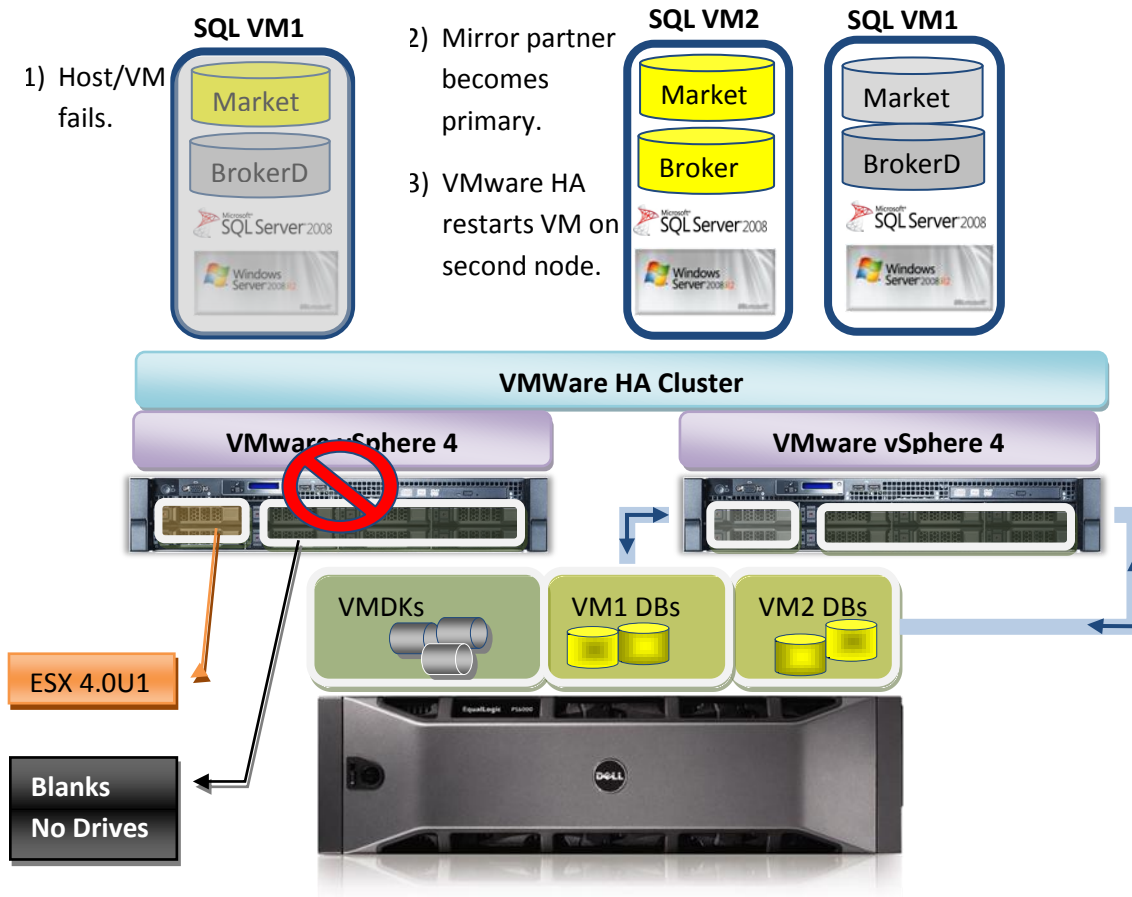


Figure 8 High Availability SAN, VMware HA and SQL Server Mirroring Failure Scenario

High Availability + VMotion + VMware FT Solution

This configuration is best suited for datacenters that require the full set of features VMware has to offer such as VMotion and VMware FT included in VMware Advanced Edition. VMotion offers VMs to preserve their state by allowing a VM to be migrated while running to another server for preventive measures. However, for critical scenarios such as component or hardware failure, VMware FT offers transparent failover capabilities.

For other advanced features such as Storage VMotion and VMware Distributed Resource Scheduler (DRS) and VMware Distributed Power Management (DPM) solutions, consider Enterprise or Enterprise Plus Editions. These features offer a robust mechanisms to protect, monitor and load balance large enterprise datacenters.

Use Case 1: Run SQL Server workloads during critical windows

Protecting a virtual SQL Server instance with VMware FT offers a robust solution during critical windows of operation. For example, e-commerce applications during peak shopping times or end of quarter financial reporting. These time frames are critical for the organization and perceived customer or user experiences are very important. A component failure during these

times that would otherwise trigger a VMware HA restart or SQL Server Mirroring failover could potentially cause some delays in user experience such as when purchasing an item online or a prolonged report execution. VMware FT offers a solution where an identical VM is created on a secondary node that is synchronized with the primary and seamlessly takes over in the event the primary VM fails.

Considerations: VMware FT only supports a single vCPU processor. Use for workloads that do not place high demands on processing power such as transactional databases. However, careful testing needs to be performed ensure there are enough resources available for the workload and peaks. VMware FT introduces some amount of overhead on the CPU for recording events on the primary host and replaying them on the secondary host (FT logging) as well as memory required on the secondary host to allocate the VM. In addition, VMware recommends having sufficient bandwidth between the two hosts dedicated for FT Logging based on workload to ensure no bottlenecks limit the number of VMs that can reside on a pair of hosts.

VMware describes the following formula as a guidance to size bandwidth required by FT logging:

FT logging bandwidth \sim [(Average disk read throughput in Mbytes/sec * 8) + Average network receives (Mbits/sec)] * 1.2 (20% additional reserved).

For instance a SQL Server for a mix workload profile. Applying the formula [(10 Mbytes/sec * 8) + 30 (Mbits/sec)] * 1.2 yields the required bandwidth for FT logging \sim 132Mbits/sec or \sim 17Mb/s per VM. The GigE link bandwidth is approx \sim 115MB/s. 5 VMs with similar profiles would saturate the link to \sim 75%. Any variation on the workload would be limited by the available bandwidth. For more information on configuring network load distribution for Fault Tolerance Logging Traffic, see Reference page. **Note:** For increased bandwidth requirements VMware recommends to implement 10GB Ethernet for FT Logging.

Another important consideration since both VMs (primary and secondary) are running in lockstep with VMware's vLockstep technology, in the event of VM OS failure such as a blue screen, both the primary and secondary VMs will fail; therefore implementing another layer of protection can be accomplished with SQL Server Mirroring. In addition, since VMware FT requires that both hosts are members of a VMware HA cluster. A physical host or VM OS failure will trigger a restart of the VM. In this case, the application connected to the database will need to retry the execution of the query on the mirror partner.

VMware FT requires that all hardware be uniform in specifications, load and power consumption to ensure the availability of the VM. This may pose a constraint if the hardware is used for other purposes and careful planning and analysis needs to be performed to ensure proper functionality.

Figure 9 below shows a logical diagram of a two node VMware HA cluster with VMs configured with VMware FT. The secondary VM is essentially a copy of the first instance that becomes active in the event of failure on the physical host.

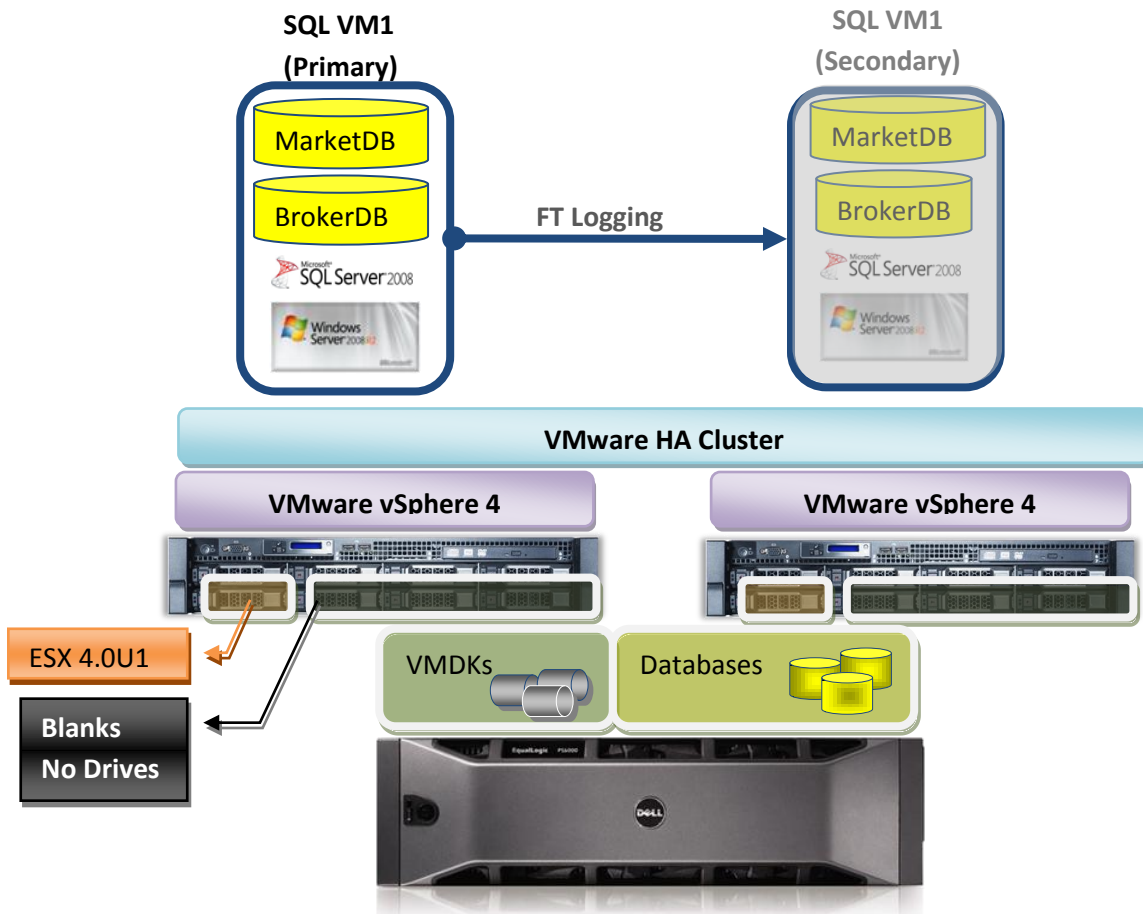


Figure 9 High Availability VMware FT and SQL Server Mirroring

Conclusion

Organizations face continuous challenges to maintain system uptime and availability to satisfy business applications needs and customers. Planning a virtualized Microsoft SQL Server databases environment requires careful planning and design of systems that can sustain component failure or even server downtime. VMware vSphere4 offers many advantages including consolidation, increased utilization levels, and power savings. In addition, VMware offers VMware HA and VMware FT. These technologies protect VMs by either by providing a stateless automatic restart of the VM(s) on a surviving node and by providing a statefull transparent failover of the VM in the event of component or physical server failure. However, it is important to add another layer of protection for SQL Server database environments due to the considerations previously presented in this paper. By combining VMware HA or VMware FT with Microsoft SQL Server mirroring, a robust, fault-tolerant, and highly-available solution is achieved.

For more information on Business Ready Configuration and Dell Services and Solutions, please visit www.dell.com/virtualization

References

Dell Services

- <http://www.dell.com/services>

Dell Database Solutions Engineering

- <http://www.dell.com/sql>

Dell Virtualization Solutions Engineering

- <http://www.dell.com/virtualization>

Dell Business Ready Configurations

- <http://content.dell.com/us/en/business/d/virtualization/Business-Ready-Small-Medium-VMware-vSphere.aspx>

Dell EqualLogic Networking Best Practices

- https://www.equallogic.com/support/tech_reports.aspx

VMware KB Article 1011966: Configuring network load distribution for Fault Tolerance logging traffic

- http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=display_KC&externalId=1011966

VMware vSphere Reference Architecture for Small and Medium Business

- <http://content.dell.com/us/en/business/d/business~solutions~engineering-docs~en/Documents~VMware-vSphere-Reference-Architecture-SMB.pdf.aspx>

VMware vSphere Availability Guide

- http://www.vmware.com/pdf/vsphere4/r40/vsp_40_availability.pdf

VMware vSphere Fault Tolerance: Architecture and Performance

- http://www.vmware.com/files/pdf/perf-vsphere-fault_tolerance.pdf

VMware Performance Best Practices for VMware vSphere 4

- http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.0.pdf

VMware Setup for Failover and Microsoft Cluster Service

- http://www.vmware.com/pdf/vsphere4/r40/vsp_40_mscs.pdf