

Availability Guide for Deploying SharePoint 2007 on VMware® vSphere

August 2009



Contents

Introduction.....

SharePoint 2007 with VMware HA/DRS

SharePoint 2007 with VMware HA and Database Mirroring.....

SharePoint 2007 with Microsoft Failover Clustering.....

SharePoint 2007 with Log Shipping

Summary

References.....

1

2

4

10

13

18

18

Introduction

When planning a SharePoint Farm deployment, defining measures to provide system availability and redundancy is a very important step. The right availability architecture will ensure that all critical elements of your SharePoint deployment can withstand various failure modes that could otherwise affect users accessing SharePoint services.

This document describes various options to provide SharePoint farm availability in a virtualized environment using VMware vSphere™ and VMware HA/DRS, using redundant SharePoint server roles, and implementing other availability/redundancy options such as clustering, database mirroring, and log shipping. SharePoint designs follow three-tiered web-application architecture with Web Server, Application Server and Database Server roles hosted across multiple server(s)/host(s). SharePoint supports running Web and Application (except Index) roles on multiple servers within a farm to increase capacity, improve performance and provide basic availability. Web Servers can use software or hardware load balancing techniques to provide redundancy. Database availability requirements within a SharePoint Farm can be met to varying degrees by using VMware HA or SQL Server clustering, high-availability mirroring, or log shipping.

Before choosing any one option, evaluate your own business requirements to determine which solutions best meets your specific needs. The remaining sections in this document describe overall procedures and also provide key steps specific to implementing chosen SharePoint availability options. This document assumes you have already configured VMware vSphere using best practices and have enabled VMware DRS, HA, and any other system functionality required by specific solutions.

Specific details of implementing high availability for SQL Server are provided in a separate document titled *Availability Guide for Deploying SQL Server on VMware® vSphere*, available at:

<http://www.vmware.com/solutions/business-critical-apps/sql/resources.html>

This document makes references to some details described in the SQL Server availability document, but focuses primarily on SharePoint specific issues and configuration related to availability.

SharePoint 2007 with VMware HA/DRS

The following diagram shows a sample SharePoint farm deployment using virtual infrastructure provided by VMware vSphere with VMware HA/DRS.




Figure 1. SharePoint 2007 Farm Implemented with VMware vSphere and VMware HA/DRS

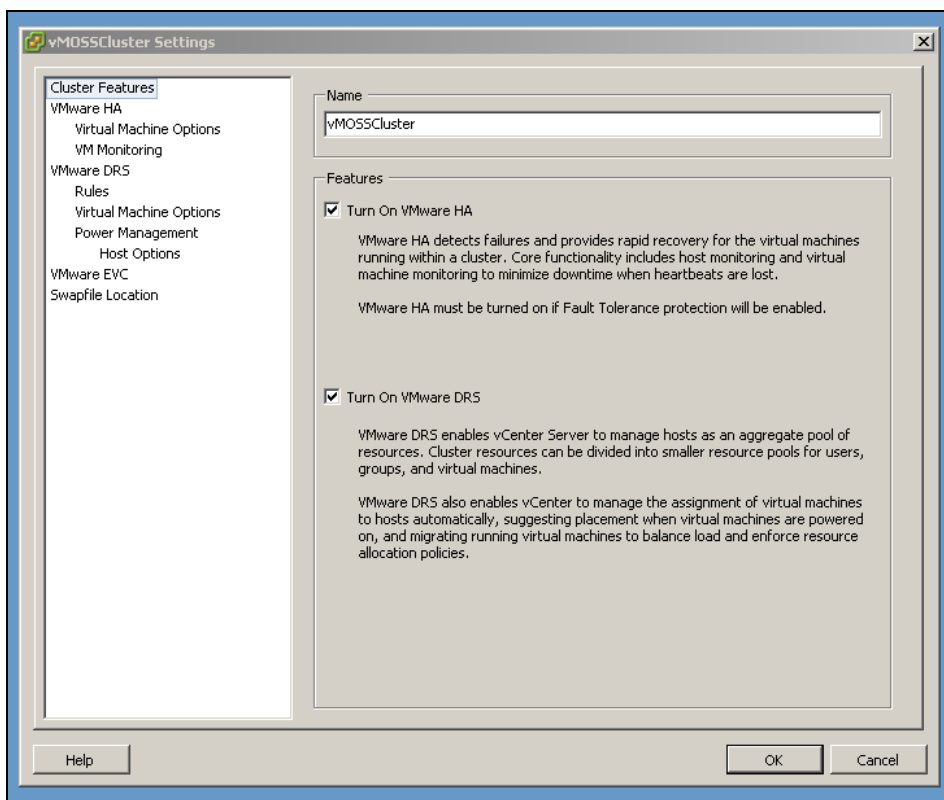
VMware vSphere with VMware HA/DRS provides high availability and host-level protection against hardware failure. VMware HA provides high availability for virtual machines by pooling them and the hosts they reside on into a cluster. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts.

Use VMware HA when your business requirements dictate that you can have limited downtime and you don't want to tackle the complexities of configuring failover clustering, log shipping, or database mirroring. VMware HA allows you recover from host outages by restarting SharePoint server virtual machines on other surviving nodes in a VMware HA/DRS cluster. Crash consistency of the virtual machines is ensured during a host outage.

Using VMware HA in combination with VMware Distributed Resource Scheduler (DRS) facilitates automatic restart of virtual machines as well as intelligent load balancing of the entire VMware HA/DRS cluster. After a host failure, the virtual machines that were running on that host are automatically restarted on other surviving hosts in the cluster. During the restart of the virtual machines, VMware DRS will dynamically re-balance the cluster by intelligently relocating the virtual machines, effectively optimizing the cluster resource utilization, all while providing the required resources to the virtual machines. Note however that VMware HA/DRS does not provide application-aware monitoring and only protects against hardware failures.

Here are specific steps you can use to configure VMware HA/DRS option for a typical SharePoint 2007 configuration:

1. Create a VMware vSphere Cluster combining two or more ESX hosts.
2. Configure vSphere Cluster Features as shown below:



For more information on configuring a vSphere cluster and configuring VMware HA/DRS to provide high availability, see the VMware [Availability Guide](#).

3. Place your SharePoint virtual machines across the ESX hosts based on your load profile.
4. Use Software/Hardware Load Balancer to make your Web Front Ends redundant.

SharePoint 2007 with VMware HA and Database Mirroring

The following diagram shows a sample SharePoint farm deployment with VMware vSphere virtual infrastructure using Database Mirroring to provide availability.

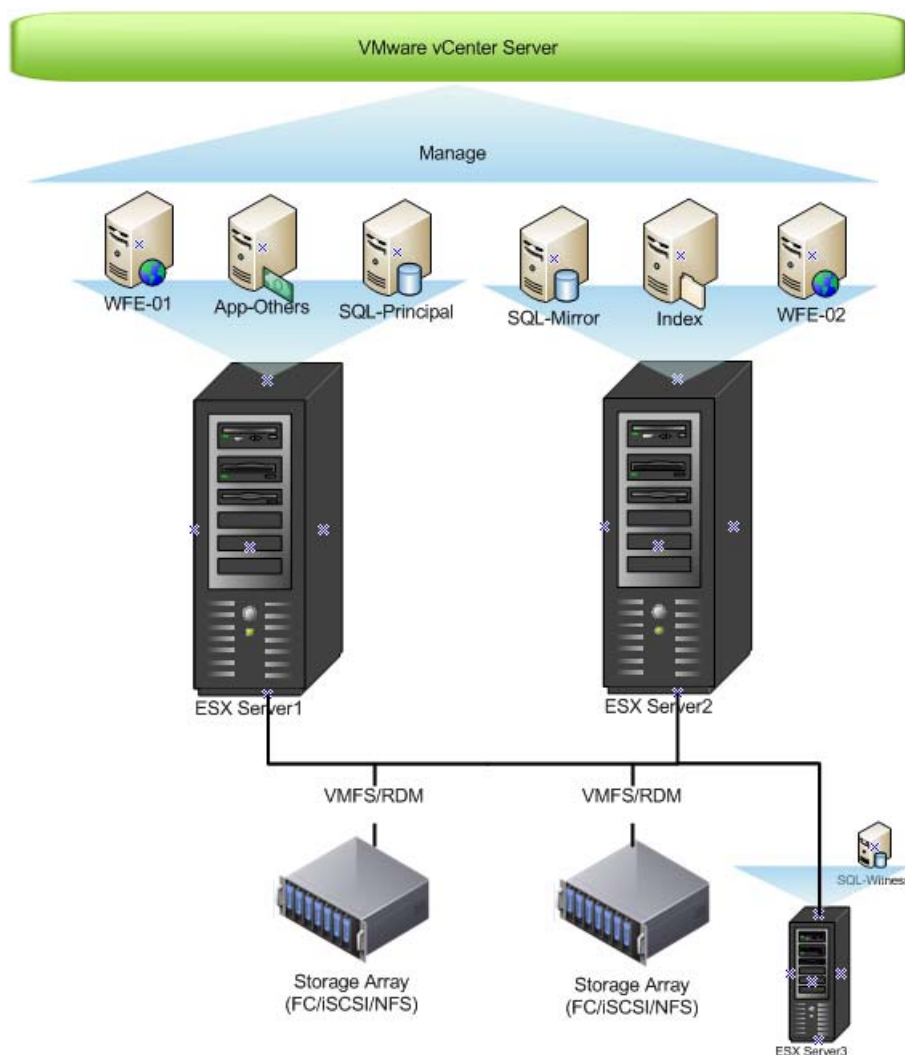


Figure 2. SharePoint 2007 with VMware HA/DRS + DB Mirroring (High Availability Mode)

The Database Mirroring availability option provides complete protection against data loss and fast recovery through automatic failover. Database mirroring maintains a redundant database by shipping log files to mirror servers when the transactions are committed on the principal server.

Use database mirroring when you need to maintain multiple copies of your data and to create a warm standby of your SQL Server database environment. You should note that SQL Server database mirroring is performed at the database level. Only user databases can be mirrored. This can give you flexibility over SQL Server failover clustering in that you don't need to include all of your databases in the cluster. You can pick and choose which database(s) to mirror. There is also no single point of failure, as in SQL Server failover clustering, in which the storage can become unavailable or corrupted.

There are three modes in which database mirroring operates:

- **High Availability mode (High Safety)** – A principal server, mirror server, and witness server are required. This mode sets the transaction safety to 'FULL' thereby using synchronous I/O between the principal and mirror servers. The principal server will wait for an acknowledgement from the mirror server that the shipped transaction has been applied on the mirror server. In the event of failure at the principal server, the witness server will form a quorum and fail over to the mirror server. You do not lose any data as a result of the transaction safety mode being configured as "FULL".
- **High Protection mode** – A principal server and mirror server are required. Transaction safety mode is set to 'FULL'. This uses synchronous I/O between the principal and mirror servers. In the event of a failure of the principal server, you will need to manually fail over to the mirror server as there is no witness server to perform quorum and automate failover.
- **High Performance mode** – A principal server and a mirror server are required. Transaction safety is set to 'OFF' and as a result, asynchronous I/O is used between the principal and mirror servers. In this mode, the principal server does not wait for an acknowledgement from the mirror server that all transactions have been recorded on the mirror server. The client application gets a confirmation that a transaction has been committed as soon as the principal server writes the transaction to the log. In the event of failure or unavailability of the principal server, you will need to manually fail over to the mirror server. There is the possibility of data loss as some transactions might be 'in flight' during the failure of the principle server.

To implement SQL Server database mirroring:

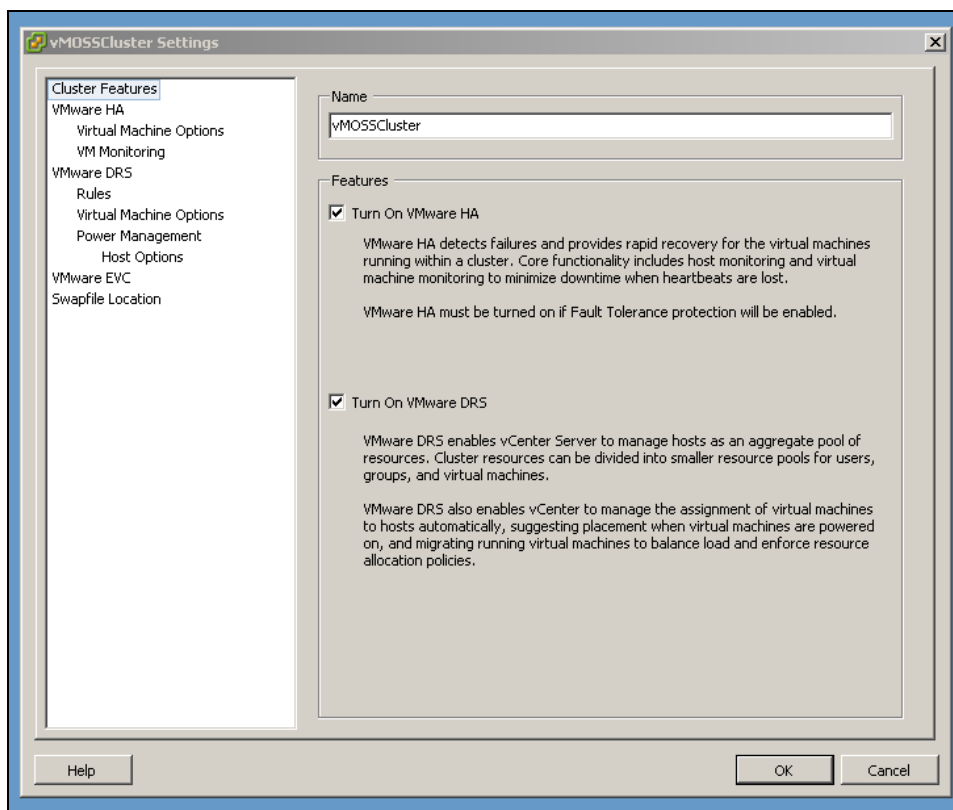
1. First install SQL Server on two servers and create your database(s) on the principal server.
2. You must then perform a full backup of the principal server database(s) and transaction logs.
3. After you have completed your backup, you must then restore your backup to the mirror SQL Server and configure the database for mirroring.
4. For this scenario, we will configure a SQL Server database mirroring configuration that is in High Safety mode. During this process, you configure security for the principal, mirror and witness servers (if you will be using "High Safety Mode".)

In High Safety mode, all updates to the database are done synchronously. There is a small performance hit for using this option, but using the high safety mode ensures that your database recovery point objective (RPO) is as close to real-time as possible.

5. After you have completed database mirroring configuration, you can then establish your mirroring session.

Here are details for a typical SharePoint 2007 Farm with DB Mirroring configuration:

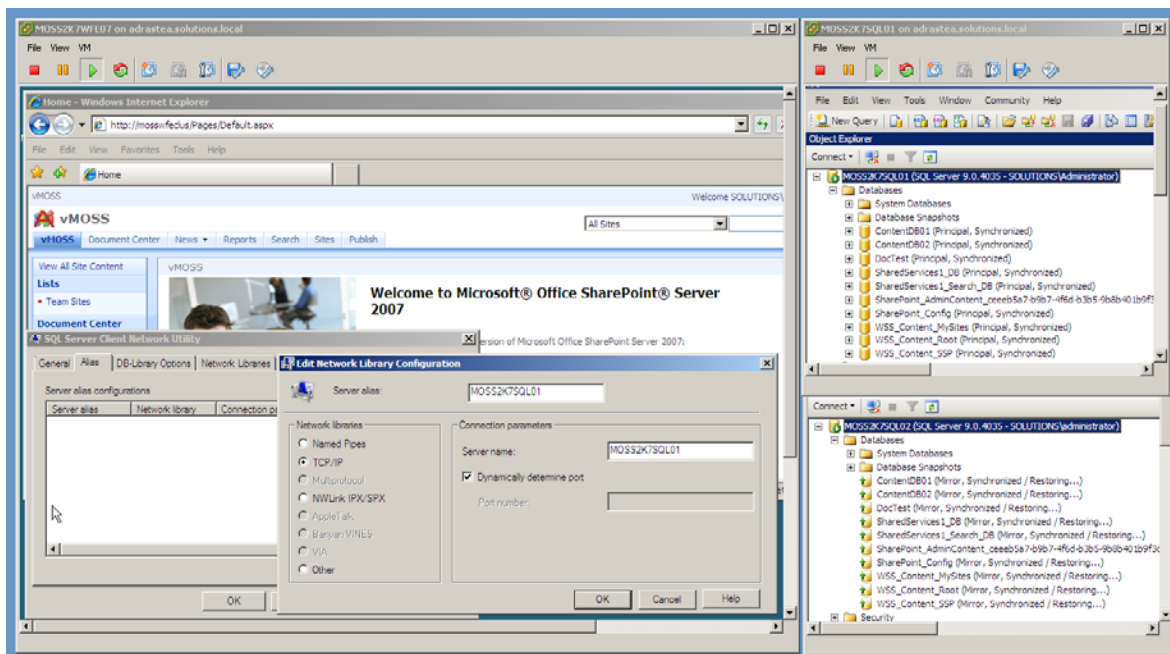
1. Create a vSphere Cluster combining two or more ESX hosts.
2. Configure Cluster Features as shown below:



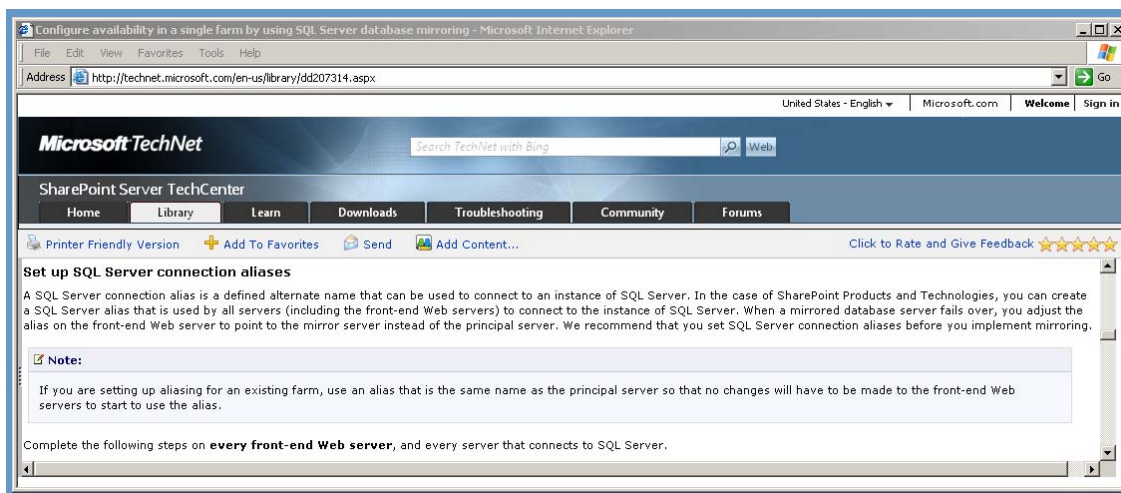
For more information on configuring a vSphere cluster and configuring VMware HA/DRS to provide high availability, see the VMware [Availability Guide](#).

3. Place your SharePoint virtual machines across the ESX hosts based on your load profile.
4. Use Software/Hardware Load Balancer to make your Web Front Ends redundant.
5. To configure database mirroring on your SharePoint Database, follow the instructions in the Database Mirroring section of the "Availability Guide for Deploying SQL Server on VMware® vSphere" white paper located at:
<http://www.vmware.com/solutions/business-critical-apps/sql/resources.html>
6. Since Mirroring is done on a per database level, make sure it is configured on all SharePoint databases (on both the Principal and Mirror SQL Server instances), and follow Microsoft SharePoint best practices for using SQL Server Alias as shown in the following screens.

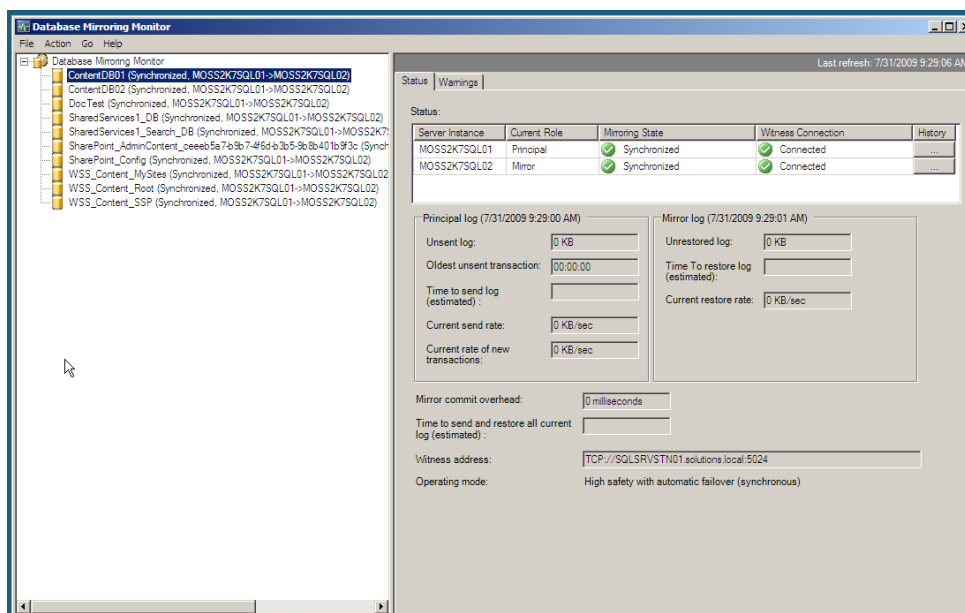
Here is a sample configuration of Principal (MOSS2K7SQL01), Mirror (MOSS2K7SQL02) and SQL Server Alias (MOSS2K7SQL01) pointing to Principal Server:



NOTE: As a best practice, it is recommended to create SQL Server Alias before configuring SharePoint. If you are setting up aliasing for an existing SharePoint farm, use an alias that is the same name as the principal server so that no changes will have to be made to the front-end Web servers to start to use the alias.

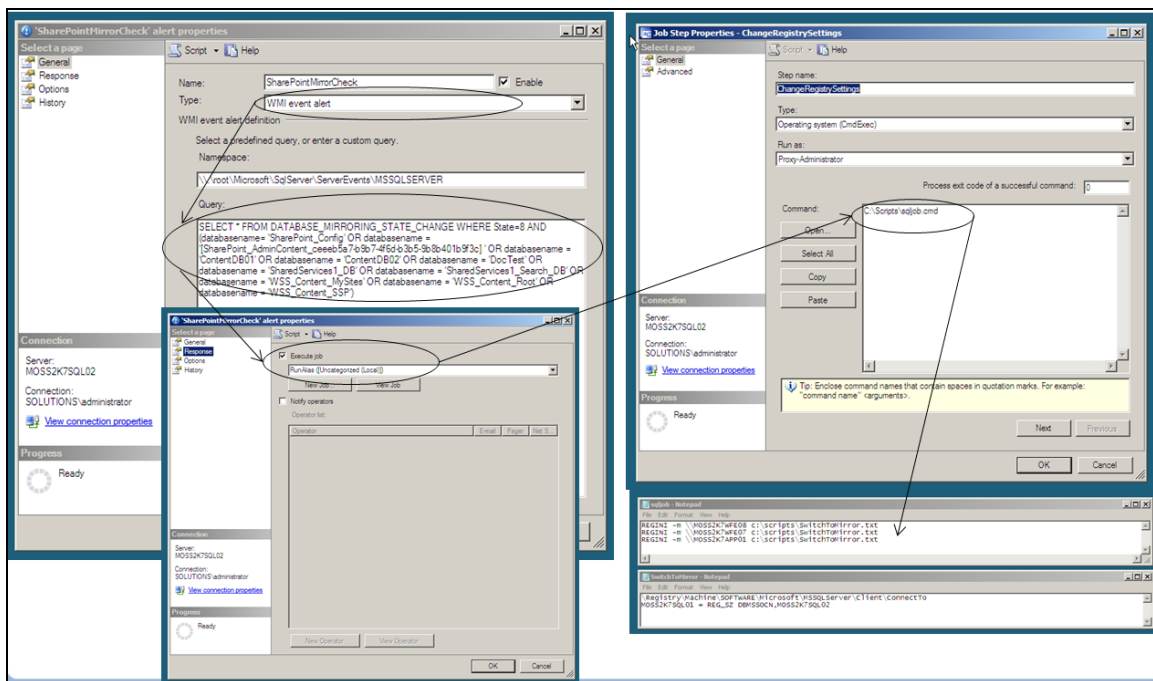


7. Open the Mirror Monitor to check Principal/Mirror/Witness server configurations as shown below:

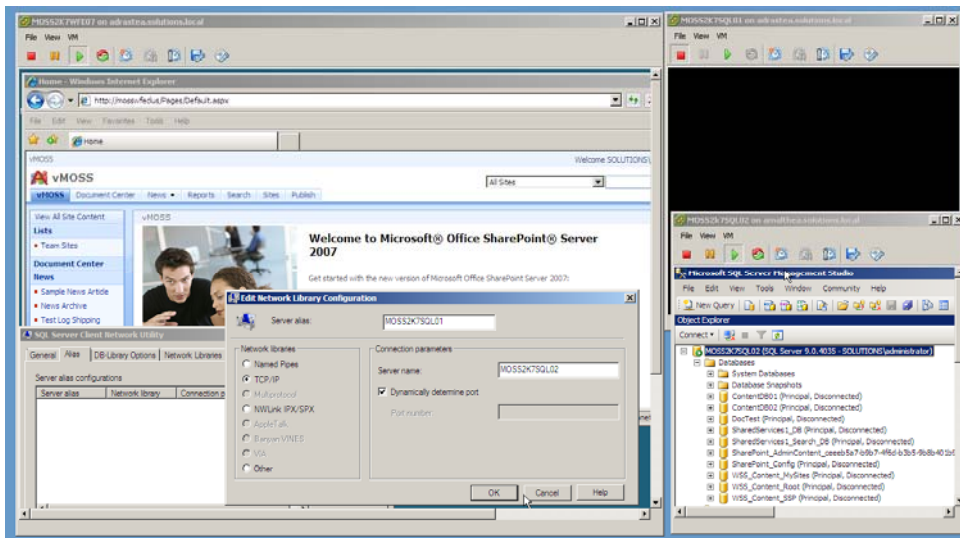


NOTE: As a best practice, configure your Witness Server on an ESX host which does not run Principal or Mirror.

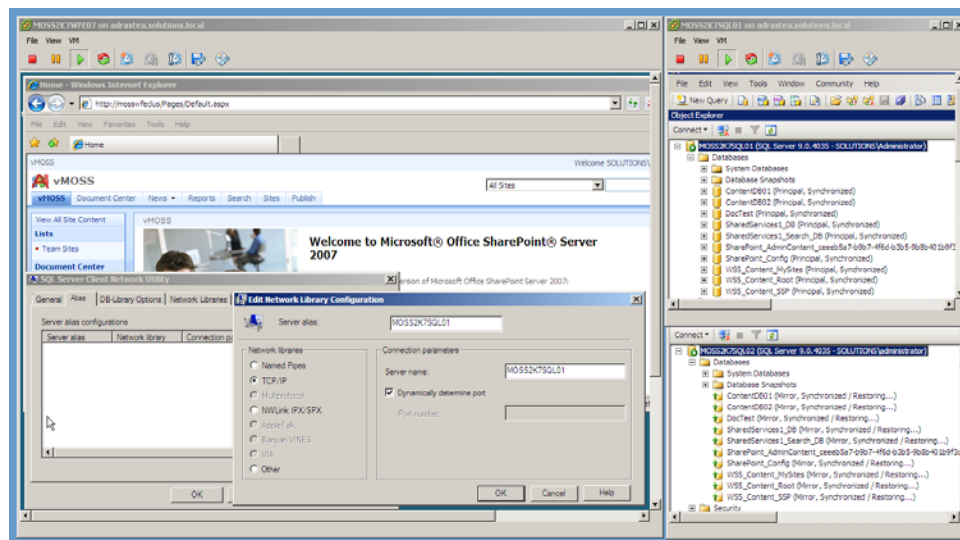
8. Follow the steps for scripting mirror failover from <http://technet.microsoft.com/en-us/library/cc531332.aspx> and create a SQL WMI Alert on the Principal and Mirror servers for monitoring failover and executing SQL jobs on failover, as shown below:



9. On completion of the above scripts, test failover from Principal to Mirror as shown below:



10. Also, test failback from Mirror to Principal as shown below:



SharePoint 2007 with Microsoft Failover Clustering

The following diagram shows a sample SharePoint farm deployment with VMware vSphere virtual infrastructure using Microsoft Failover Clustering to provide availability.




Figure 3. SharePoint 2007 with Microsoft Failover Clustering

The Failover Clustering option is a more expensive solution for providing SQL Server availability. It provides very good protection in the case of hardware failure. Failover to a passive node is fairly quick (anywhere between 1-5 minutes depending upon the state of the cluster and database). Failover Clustering provides service availability but does not provide data redundancy like database mirroring and log shipping. Data protection has to be provided at the storage level or by combining with other solutions.

Failover Clustering provides instance-level protection built on Windows Failover Clustering. Nodes of the cluster are typically co-located within the same site or datacenter to provide local availability but can also be deployed regionally. Disaster Recovery (DR) and Business Continuity (BC) are possible using geographically dispersed clustering solutions. These solutions can be difficult to design, implement, and maintain from a cost and operational standpoint.

Implementing SQL Server failover clustering involves a number of steps:

1. Create a virtual machine and install all tested and critical updates and service packs. You should template this virtual machine using sysprep and deploy your virtual machines from templates. This will significantly reduce the time needed to provision your virtual machines.
2. Create a new VMware vSphere virtual switch or port group and add a second network adapter to the cluster's primary node virtual machine to carry cluster heartbeat communications. If using vSphere, select the vmxnet-3 adapters for your public and heartbeat adapters.
3. After completing the network configuration, you map your storage array LUNs as Raw Device Mappings (RDM) to the primary virtual machine. Keep in mind you need to set up a separate vSCSI controller for clustered disks as clustered disks cannot reside on the same vSCSI controller as the OS boot drive. You also should map your database and log LUNs on separate vSCSI adapters.

NOTE: Details on virtual machine setup for use in Microsoft clusters is provide in the VMware [Setup for Failover Clustering and Microsoft Cluster Services](#) guide.

4. Following setup of the SQL Server 2008 primary node virtual machine, you can configure the virtual machine of the secondary SQL Server node.
5. When you are finished with the secondary node, you need to install Microsoft .NET application server roles, configure associated role services, and add the failover cluster feature to both primary and secondary virtual machines in the cluster.
 - a. On the primary node, open the Disk Manager and bring all newly presented disks online and, using DiskPart, align your disks to prevent split I/Os.
 - b. On the secondary virtual machine, open the Disk Manager utility to bring disks online to the secondary node virtual machine. Map the disk letters exactly as they are mapped on the primary virtual machine.
 - c. Use the Microsoft Cluster validation wizard to validate a cluster configuration. When validating a cluster for the first time, you MUST run ALL tests and all tests must pass to be a Microsoft supported failover cluster solution.
6. After completion of the validation, use Microsoft Failover Cluster Management options to create a Windows Server failover cluster to support a clustered SQL Server 2008 instance(s). You will then use the Failover Cluster Management tools to configure the Cluster Quorum or "witness" drive and settings to specify failover conditions for the cluster.
7. With the cluster now fully configured and functional, configure the Microsoft Distributed Transaction Coordinator (MSDTC). SQL Server uses the MSDTC for distributed queries, replication, and two-phase commit transactions.

To configure the MSDTC, you open the Microsoft Failover Cluster Management console to configure a service or application. The console provides a High Availability Wizard to configure MSDTC to perform the SQL Server cluster failover.

8. The last step is to actually install SQL Server 2008 on the failover cluster primary node. The procedure to add an additional cluster node to a clustered SQL Server instance is similar to the failover cluster installation of SQL Server. You can perform this task, from either the primary or second node of the Windows Server failover cluster.
9. With SQL Server node installation complete, perform instance failover tests.

Here are details for a typical SharePoint 2007 Farm with Failover Clustering configuration:

1. Place your SharePoint virtual machines across the ESX hosts based on your load profile.
2. Use Software/Hardware Load Balancer to make your Web Front Ends redundant.
3. To configure Microsoft Failover Clustering on your SQL Servers, follow the instructions in the Failover Clustering section of the *"Availability Guide for Deploying SQL Server on VMware® vSphere"* white paper located at:

<http://www.vmware.com/solutions/business-critical-apps/sql/resources.html>

SharePoint 2007 with Log Shipping

The following diagram shows a sample SharePoint farm deployment with VMware vSphere virtual infrastructure using database transaction log shipping to provide availability.




Figure 4. SharePoint 2007 with 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. SQL Server 2008 provides a sub-minute scheduling interval providing the ability to do quick backup and restores. Backups are performed on the primary SQL Server instance and restore performed on secondary SQL Server instance(s). Scheduling is done through SQL Server Agent jobs.

As an availability strategy, the Log Shipping option does not provide any automatic failover capability and may allow some data loss, however the time synchronization interval is configurable, which gives users some level of control. (This data loss can occur due to corrupted, missing and possibly transaction log data

that was not sent to the backup log file on the secondary server.) Log shipping can be used in conjunction with failover clustering to provide a 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 purposes, or to provide a database that can be used for reporting and to offload reporting and query functions off of the primary server, thereby improving overall performance and providing a better end user experience of the production SQL Server.

To implement log shipping availability, you need to do the following:

1. Configure the primary SQL Server 2008 instance to ship transaction logs to the secondary server location. (You can use the SQL Server Management Studio, Transact-SQL statements, or you can run scripts using Transact-SQL to accomplish these tasks.)
2. After configuring Log Shipping, you need to verify correct operation of the Log Shipping configuration. You should also periodically monitor SQL Server Agent job activity to confirm that backup operations are being performed correctly.

Here are details of the log shipping option for a typical SharePoint 2007 Farm configuration:

1. To configure log shipping on your SharePoint Database, follow the instructions in the Log Shipping section of the "Availability Guide for Deploying SQL Server on VMware® vSphere" white paper located at: <http://www.vmware.com/solutions/business-critical-apps/sql/resources.html>
2. Create a new SharePoint Farm on the secondary server with its own Config/Admin/SSP databases as shown below:

Central Administration > Operations > Servers in Farm

Servers in Farm

Farm Information

Version: 12.0.0.6421

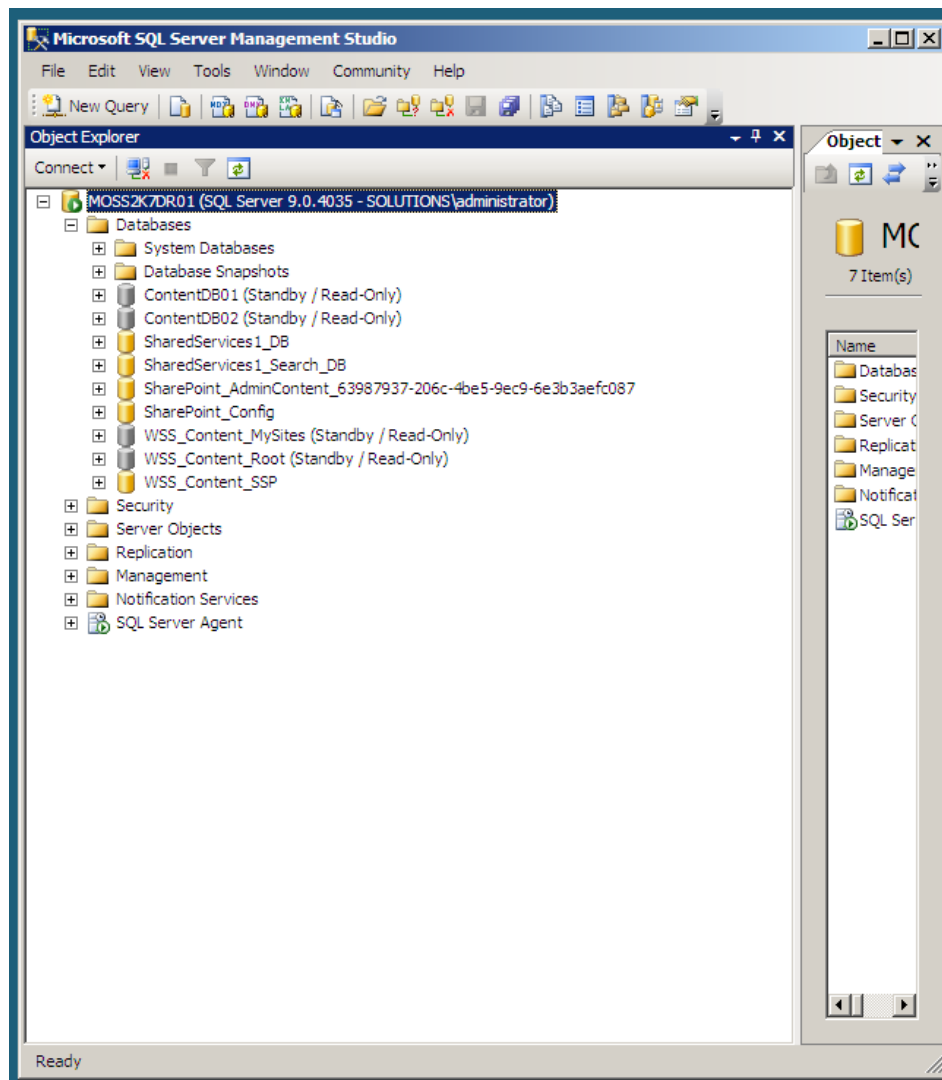
Configuration database server: MOSS2K7DR01

Configuration database name: SharePoint_Config

Server	Services Running	Version	Remove Server
MOSS2K7DR01	Central Administration Office SharePoint Server Search Windows SharePoint Services Database Windows SharePoint Services Incoming E-Mail Windows SharePoint Services Web Application	12.0.0.6421	Remove Server

NOTE: As a best practice it is recommended to configure SharePoint Web/App/Database on separate servers. In this study we have taken the approach of configuring all tiers of SharePoint in a box for the Secondary Farm.

3. Attach the read-only content database to the secondary SharePoint Farm as shown below:



4. Screen below shows both Primary and Secondary Server displaying identical contents after the content databases are added to the secondary server.

Central Administration > Application Management > Content Databases

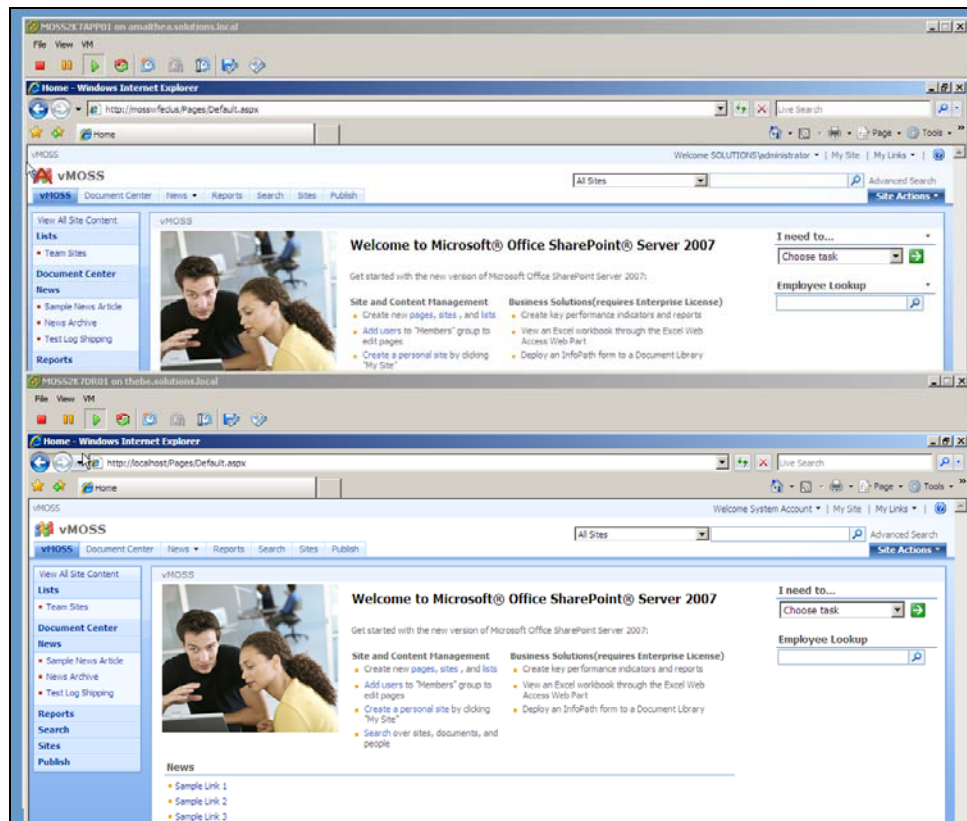
Manage Content Databases

Use this page to manage content databases for this web application. Click a content database name to change its properties.

Add a content database Web Application: <http://moss2k7dr01/>

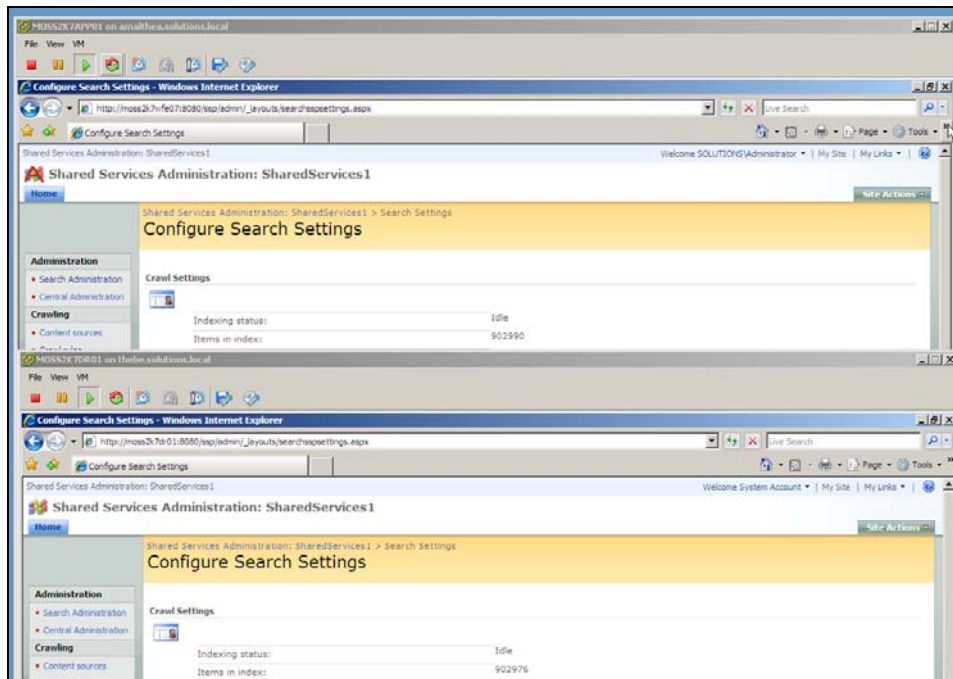
Database Name	Database Status	Current Number of Sites	Site Level Warning	Maximum Number of Sites
ContentDB01	Started	1	9000	15000
ContentDB02	Started	1	9000	15000
WSS_Content_Root	Started	2	9000	15000

SharePoint view on both Primary and Secondary Servers:

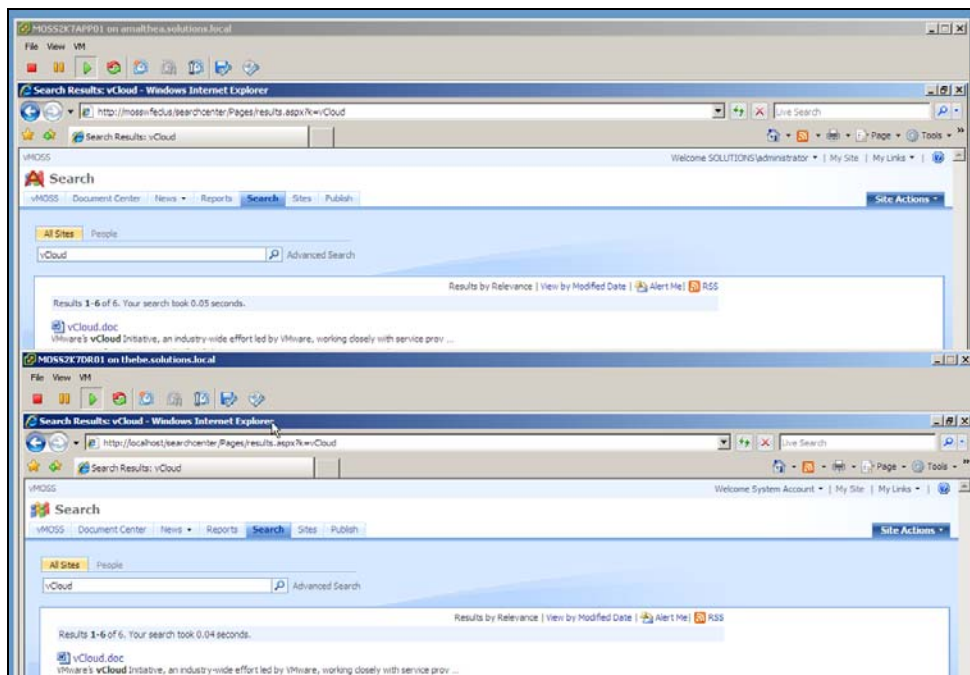


- Run Full index crawl on secondary server after attaching the content database.

Screen below shows similar number of items in index after full crawl and the search displaying similar results.



The following screen shows similar results for search:



Summary

When designing a SharePoint Farm deployment, you should also think about taking measures to provide system availability and redundancy. This document described various options to provide SharePoint Farm availability in a virtualized environment using VMware vSphere and VMware HA/DRS, using redundant server roles, and implementing other availability/redundancy options such as clustering, database mirroring, and log shipping.

Before choosing any one availability option to implement in your environment, evaluate your own business requirements to determine which scenario best meets your specific needs. The option(s) you choose may also take in account complexity and cost, available equipment and network resources, and the type of SharePoint software licenses you have available at your site.

References

Documents listed here provide additional information relevant to implementing high availability for SQL Server 2008 that is available on the VMware web site or from other third parties.

- VMware Solutions web site located at <http://www.vmware.com/solutions>
- SharePoint availability and redundancy definitions from Microsoft:
<http://technet.microsoft.com/en-us/library/cc748824.aspx>
<http://technet.microsoft.com/en-us/library/cc263044.aspx>
- Windows NLB and Clustering configuration from VMware:
http://www.vmware.com/files/pdf/implmenting_ms_network_load_balancing.pdf
http://www.vmware.com/pdf/vsphere4/r40/vsp_40_mscs.pdf
- SharePoint Mirroring and Log Shipping best practices from Microsoft:
<http://technet.microsoft.com/en-us/library/cc262910.aspx>
<http://technet.microsoft.com/en-us/library/dd890507.aspx>

Additional product information and documentation for using VMware Infrastructure is located at <http://www.vmware.com/support/pubs/>.



VMware, Inc., 3401 Hillview Ave. Palo Alto, CA 94304, USA. Tel: 877-486-9273. Fax: 650-427-5001. www.vmware.com
Copyright 2009 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

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.

