



Achieving High Availability with Symantec Enterprise Vault™

Chris Dooley
January 3, 2007

Achieving High Availability with Symantec Enterprise Vault

Contents

Executive summary	4
Introduction to high availability	5
Definition of a cluster	5
Common failover configurations	6
Application-level high availability with Update Service Locations	7
Supported failover configurations	8
Requirements for Update Service Locations	8
Installation	10
Failover process	10
Failback process	11
Veritas Cluster Server integration	12
Supported failover configurations	12
Requirements for Veritas Cluster Server integration	13
Installation	14
Failover process	17
Failback process	17
Microsoft Cluster Server integration	18
Supported failover configurations	18
Requirements for MSCS integration	19
Installation	19
Failover process	21
Failback process	22
Conclusion	22

Executive summary

In recent years, email has become a business-critical tool as more and more organizations come to rely on it for their daily operations. Without email, many companies would not be able to conduct business. As email use continues to grow exponentially, organizations face the challenges of complying with regulatory and corporate policy, reducing the cost of e-discovery, and reducing storage costs. These requirements continue to fuel the archiving market. Symantec offers a leading solution for this rapidly growing market. Symantec is positioned in the Leaders Quadrant for Email Active Archiving according to Gartner's "Magic Quadrant for Email Active-Archiving Market, 2006" research note.^{1,2} And according to Gartner's "Archiving Software Market to Experience Strong Growth Through 2010" research note, Symantec led the email archiving market in 2005 based on customers and new license revenue.³

Symantec Enterprise Vault software provides a unique archiving framework that enables the discovery of content held within email, file systems, and collaborative environments, while helping to reduce storage costs and simplify management. With Enterprise Vault, organizations can ensure compliance with retention and discovery policies because it acts as a secure repository for electronic information while helping to reduce the cost of content storage, retrieval, recovery, and administration.

Not only do organizations need to retain and manage this data, but also access to this information in real time is crucial. Any outage that affects access to data could be disastrous to a business. Therefore, a clear disaster recovery plan should be an integral component of the design of any archive solution. It is essential that companies implement reliable strategies for maintaining availability. The proven strategy for achieving this level of availability is by deploying the applications in a cluster environment.

Veritas™ Cluster Server from Symantec is an industry-leading cluster solution for reducing application downtime. The integration of Enterprise Vault in a Veritas Cluster Server environment provides a robust and highly available archiving solution. Enterprise Vault also supports Microsoft® Cluster Server, as well as a built-in, application-level high availability solution. All three high availability solutions will be described in this paper.

While this paper focuses specifically on high-availability options for the Enterprise Vault server, the reader should note that Enterprise Vault also supports the archiving of clustered data sources, such as clustered Microsoft Exchange servers and Microsoft Windows® file systems. Also,

¹ Gartner, "Magic Quadrant for E-Mail Active Archiving, 2006"; Carolyn DiCenzo & Kenneth Chin; May 12, 2006.

² The Magic Quadrant is copyrighted 2006 by Gartner, Inc. and is reused with permission. The Magic Quadrant is a graphical representation of a marketplace at and for a specific time period. It depicts Gartner's analysis of how certain vendors measure against criteria for that marketplace, as defined by Gartner. Gartner does not endorse any vendor, product or service depicted in the Magic Quadrant, and does not advise technology users to select only those vendors placed in the "Leaders" quadrant. The Magic Quadrant is intended solely as a research tool, and is not meant to be a specific guide to action. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

³ Gartner, "Archiving Software Market to Experience Strong Growth Through 2010," #G00142798, September 2006, based on customers and new license revenue.

Achieving High Availability with Symantec Enterprise Vault

it should be noted that clustering for Compliance Accelerator or Discovery Accelerator is not currently supported. However, a nonclustered Compliance Accelerator or Discovery Accelerator server can reference a clustered Enterprise Vault virtual server.

Introduction to high availability

When forming a disaster recovery plan, the two main factors that need to be considered are the Recovery Time Objective (RTO) and the Recovery Point Objective (RPO). The RTO refers to the amount of time the application is down, while the RPO refers to the amount of data lost. Organizations should have an RTO (per application) that must be satisfied as well as an RPO that must be met. Most organizations tend to focus on the RTO, or how much downtime is acceptable. However, the amount of data loss an organization can tolerate is just as important—just a few minutes' worth of lost transactions can have a far-reaching negative impact on the business. Many companies in the banking and financial industries are actually legally bound to ensure that take proactive steps to protect against data loss.

Depending on whether the RTO and RPO are measured in minutes, hours, or days, there are different technologies available to meet those requirements. Many companies today still rely primarily on tape backup and restoration as the centerpiece of their disaster recovery plan. This usually means at least a day of lost data and a few days of downtime after a disaster. This is acceptable if it meets the business's needs, but most organizations have applications that require a shorter RPO and RTO that is measured in minutes or seconds. In these cases, a comprehensive disaster recovery plan should include data replication or continuous data protection, application clusters, and traditional tape backups. Replication or continuous data protection of the data to a remote site eliminates data loss, while clustering can manage the local and remote failover of applications to minimize downtime. The integration of replication with application clustering, supplemented with tape backups, can satisfy even the most stringent recovery time and recovery point requirements.

Definition of a cluster

The term cluster may have different meanings for different people. In the context of high availability, cluster refers to a combination of two or more servers with shared storage that are interconnected for failover to achieve high levels of system reliability.

In a cluster configuration, each server is called a node. Each node in a cluster functions either as a primary node or a failover node. The primary node is the server where an application initially starts. Each application has its own set of resources that are combined to provide the services for that application. These resources can include the Network Interface Card (NIC), the virtual IP for

that NIC, the storage, the file systems on that storage, the mount points for those file systems, and the application data and binaries residing on those file systems. These resources are allocated per application and are monitored and managed as a group. In the event any of the resources fail, the cluster software initiates the controlled switchover of the remaining resources if possible. If a controlled switchover is not possible, then a failover of the entire application is initiated.

Common failover configurations

Typical cluster deployments are architected in one of four configurations: Active/Passive, Active/Active, N+1, or N+M. These configurations have been in use for a number of years in all major cluster platforms and represent the majority of current implementations.

Active/Passive (asymmetric)

In an Active/Passive configuration, an application runs on a primary node, and each primary node has its own spare node. The spare node is not configured to perform any other functions, and it cannot take over the functions of any other node except the primary node to which it is dedicated. From the perspective of the application, this configuration is traditionally referred to as an *asymmetric configuration*.

The asymmetric configuration is the simplest and most reliable. The secondary node is standing by with full performance capability, and no other applications run on it that might present compatibility issues. This configuration is also perceived to be the most expensive, because the hardware is unused when the primary node is functioning normally, and every active node has its own dedicated passive node.

Active/Active (symmetric)

In the Active/Active configuration of a two-node cluster, each node is configured to run a specific application or service and to provide redundancy for its peer if needed. In the event of a failure, the surviving node would temporarily host both applications. From the perspective of the application, this configuration is traditionally referred to as a *symmetric configuration*.

This symmetric configuration may appear to be a far more beneficial configuration in terms of hardware utilization, since many IT organizations object to having a valuable system sitting idle. There is, however, a serious flaw in this line of reasoning. In the Active/Passive (asymmetric) configuration described above, the passive node would only need as much processing power as its peer. On failover, performance would theoretically remain the same. In the Active/Active (symmetric) configuration, though, the server taking over would need sufficient processing power to not only

Achieving High Availability with Symantec Enterprise Vault

run its own, existing application, but also enough for the new application it takes over. To put it another way, if a single application needs one processor to run properly, an asymmetric cluster would likely require two dual-processor nodes.

Further difficulties can arise in symmetrical configurations when multiple applications running on the same system do not coexist well. Some applications work perfectly well with multiple copies started on the same system, while some will fail. Even more difficult may be an environment that has two different applications with vastly different I/O and memory requirements running on the same system.

N+1 (hot standby)

The capabilities brought about by storage area networks (SANs) enable not only much larger multinode clusters, but even more important, allow multiple servers to be connected to the same storage. In an N+1 configuration, one extra node is added to the cluster to act as a “spare” node for any of the multiple active nodes. When a node fails, the application restarts on this spare node. When the original node is repaired, it could either take back its responsibilities, or become the new spare for the other nodes. Any node can provide redundancy for any other node. This allows for eight or more node clusters with one spare server. By choosing an N+1 design, the architect can safely plan for each server hosting an application to run at near capacity, and allocate one spare server of the same size to act as a spare node for all applications.

N+M

In larger organizations where there is an increased risk that more than one node may fail at a time, an N+M configuration can be used. This is an extension of the N+1 configuration that allows the use of multiple spare nodes. An organization can designate several spare nodes, and multiple failures can be accommodated at one time.

Application-level high availability with Update Service Locations

For customers that prefer not to invest in clustering hardware and software, the Symantec Enterprise Vault solution provides an application-level high-availability feature called Update Service Locations (USL) that allows a system administrator to quickly fail over the services from one Enterprise Vault server to another, without taking the time to restore data from point-in-time backup media.

Achieving High Availability with Symantec Enterprise Vault

The main drawback of this option is that it is a manual process, and therefore would involve a slightly longer time for resolution after a system failure. However, this problem is partially mitigated by the Enterprise Vault solution's unique Offline Vault feature, which displays archived content from a local cache on the client computer, even if the Enterprise Vault server is unavailable.

Supported failover configurations

The USL feature supports all four failover configurations described earlier. Again, care should be used when designing an Active/Active Enterprise Vault configuration because, in the event of a failure, one server would temporarily handle the responsibilities of two different servers, often effectively doubling the normal workload. To compensate for this, consider using more advanced server hardware, designing a solution where both servers do not run at maximum capacity (such as a journaling server failing over to a normally dedicated Discovery Accelerator server), or simply switching the system into read-only mode during a failure when the workload of an active server is doubled. During a failover, the most important task is usually the retrieval of archived content, not the ability to continue archiving additional content on a nightly basis.

In the other three configurations that include passive servers, there can be any number of active and passive Enterprise Vault servers, and there is no permanent distinction between an active and passive server. The roles of the servers depend solely on the status of the name resolution. After a failover, the newly activated server can continue its role indefinitely, and the repaired server can become a passive server for future use. While this configuration is more costly in terms of unused hardware, it avoids potential performance problems during a failover, when continued archiving during the failure is required.

Requirements for Update Service Locations

Although the USL functionality is built into the Enterprise Vault application, there are some important requirements necessary in order to take advantage of this feature. The requirements involve the way the initial Enterprise Vault implementation is designed, the type of storage used, and name resolution between client and server.

Building-block design methodology

When designing a new Enterprise Vault implementation that may take advantage of this built-in USL feature, it is mandatory that each Enterprise Vault server have a complete set of core services as a self-contained unit, or building block. In other words, each Enterprise Vault server should have Storage, Indexing, Shopping, and Task Controller Services. This way, when one server fails, it does

Achieving High Availability with Symantec Enterprise Vault

not affect the functionality of other Enterprise Vault servers. In addition, this building-block design allows the transfer of responsibilities from a failed server to another server that has the same core services in place (within the same Enterprise Vault logical site).

Remote storage

To maintain the availability of all Enterprise Vault content even after a server failure, the archived content, indexes, and database metadata must not be stored on direct-attached disks that would become unavailable after a server failure. Instead, SANs or network attached storage (NAS) should be used so that all Enterprise Vault servers could potentially access all content if needed. Note that this also includes using an SQL Server environment that is not installed on an Enterprise Vault server. While it is important to use storage technology that can be accessed by all Enterprise Vault servers if needed, each vault store partition and index location should be accessible from only one Enterprise Vault server at a time.

Some Enterprise Vault customers have successfully implemented site-to-site failover using the Update Service Locations feature, combined with storage replication software such as Symantec's own Veritas™ Volume Replicator product. As long as the secondary Enterprise Vault server can access all archived content, indexes, and SQL databases using the same paths that the original server used, the server can be located anywhere, and any storage replication software can be used to make this content available to the secondary Enterprise Vault server.

Dynamic name resolution

One critical component of any high availability solution is the use of permanent virtual server names that can be modified behind the scenes to point to different servers as needed. This way, client computers can connect to one consistent, virtual server name at all times, but actually communicate with different physical servers as necessary.

The USL feature supports the use of DNS aliases (CNAMEs) or network load balancers to provide this dynamic name resolution between client and server. When using DNS aliases, a DNS administrator must modify the alias record to point to a different, working Enterprise Vault server when one server fails. One potential issue with this solution is that, by default, Windows DNS clients cache previously resolved names and IP addresses for 15 minutes. Therefore, there could be a window of 15 minutes when clients would continue pointing to the failed Enterprise Vault server, even after the DNS administrator modifies the alias to point to the new Enterprise Vault server. A Windows Group Policy can be used to adjust this cache interval setting, and at any time the client's DNS cache can be flushed manually by running the `ipconfig /flushdns` command.

Achieving High Availability with Symantec Enterprise Vault

Using a third-party network load balancer, while more costly, would eliminate the potential downtime caused by name resolution cache. In this scenario, DNS records would always direct traffic to the network load balancer, which would then redirect the clients to the appropriate physical server based on its configuration. Network load balancers are most commonly used to distribute the user traffic to multiple servers simultaneously, using percentages to dictate how much of the overall traffic should be sent to each server. However, for Enterprise Vault server high availability, each virtual network load balancer server identity must always point to only one physical server at a time, meaning the percentages must be configured so that 100% of the traffic goes to the active Enterprise Vault server, and 0% of the traffic goes to all other servers managed by that network load balancer. In the event of a failure of the active Enterprise Vault server, the system administrator would change the 100% setting to 0%, and direct 100% of the traffic to a working Enterprise Vault server that can take over the responsibilities. Therefore, the network load balancer would be used only for failover, and not traditional load balancing. Enterprise Vault load balancing is achieved by properly designing a multiserver implementation that divides the daily workload and retrieval across servers.

Installation

One of the main benefits of the USL high-availability option is that there are no additional hardware or software requirements beyond the typical requirements for a standard Enterprise Vault server. The USL feature is included and automatically installed with the base Enterprise Vault server, and can be run on any licensed edition of Windows 2000 or Microsoft Windows Server® 2003.

Failover process

The USL failover process, while manual, is a relatively straightforward process, and common across all four failover configurations, including Active/Active.

1. The system administrator must first redirect the name resolution (via DNS or network load balancer) so the alias for the failed server points to the physical address of the new server.
2. The administrator should then log into the new server and, if using DNS for dynamic name resolution, flush the DNS cache by running the `ipconfig /flushdns` command. This ensures that the working server will be able to connect to itself, rather than the failed server, using the alias.

Achieving High Availability with Symantec Enterprise Vault

3. The administrator then opens the Enterprise Vault Administration Console from the working server and runs the Update Service Locations command by right-clicking on the “Enterprise Vault Servers” container. This process follows the name resolution process to notice that the alias now points to a different physical server, and transfers the responsibilities (and missing services if necessary) to the new server.
4. At this point, clients can continue communicating to their server’s alias and use Enterprise Vault as if nothing changed (unless their DNS cache is temporarily out of date and still pointing to the physical address of the failed server).

Failback process

The USL failback process is similar to the failover process:

1. The system administrator repairs the problem with the failed server.
2. If the administrator wants to transfer the server’s responsibilities back to the repaired server, the name resolution change (DNS or network load balancer) would be reversed to its state prior to the failure.
3. If using DNS for name resolution, the administrator should then flush the DNS on the temporarily active server.
4. The administrator opens the Enterprise Vault Administration Console from the temporarily active server and reruns the Update Service Locations command. This process notices, once again, the resolution has changed, and moves the responsibilities and services to the physical server to which the name resolution is now pointing the alias.
5. At this point, clients start communicating to the original server again (unless their DNS cache is temporarily out of date and still pointing to the temporary server, which no longer handles their archived data).

Veritas Cluster Server integration

Veritas Cluster Server is an industry-leading clustering solution for reducing application downtime. It is a component of a comprehensive Symantec availability solution called Veritas Storage Foundation™ HA for Windows. Symantec is the market leader in cross-platform server clustering, according to the 2006 edition of the IDC Worldwide Clustering and Availability Software report.⁴ A few highlights of Veritas Cluster Server include availability across wide area networks and different IP subnets, centralized management from a single Web console, support for multiple operating systems and arrays, and integration with multiple disk replication solutions, including Symantec's own Volume Replicator.

Support for Veritas Cluster Server was introduced in Enterprise Vault 6.0 Service Pack 2. Veritas Cluster Server has the ability to monitor the health of the Enterprise Vault services and storage, and automatically trigger a failover if necessary to maintain system availability. Veritas Cluster Server uses the Generic Service agent to monitor the Enterprise Vault services on different nodes, based on the information in the Enterprise Vault Directory database. The resources that are monitored within the Enterprise Vault service group include the IP resource, LANMAN resource for the virtual hostname, the Microsoft Message Queue (MSMQ) resource, storage resources, and service resources. The agent brings the services online, monitors their status, and takes them offline as needed.

Supported failover configurations

The configurations currently supported for Enterprise Vault and Veritas Cluster Server are Active/Passive, N+1, and N+M. Active/Active configurations are not yet supported.

Figure 1 illustrates an Active/Passive failover configuration. Here, the volumes for the Enterprise Vault data are configured in a cluster disk group on shared storage. The Enterprise Vault virtual server is configured on the active node (System 1). If System 1 fails, System 2 becomes the active node, and the Enterprise Vault virtual server comes online on System 2.

⁴ IDC, "Worldwide Clustering and Availability Software 2005 Vendor Shares," IDC #203676, October 2006.




Figure 1. Example of an Active/Passive configuration.

Requirements for Veritas Cluster Server integration

The first step in configuring any cluster solution is to eliminate single points of failure in the hardware. Verify that the physical connectivity of all nodes in the cluster go through separate paths. This should include separate or redundant network, storage, and power connections.

The storage capacity on each server must be identical and must have sufficient connectivity to all the failover target servers or nodes. All Enterprise Vault storage must be under Veritas Cluster Server control and must be configured as dynamic disks.

The software requirements include:

- Windows Server 2003
- Veritas Storage Foundation HA *for Windows* 4.3 MP1 or later
- Enterprise Vault 6.0 SP2 or later

There is no specific Enterprise Vault agent necessary for Veritas Cluster Server integration. Instead, Enterprise Vault uses the Generic Service agent.

All Enterprise Vault services must be set to manual startup in the Windows Services control panel. Once clustered, the Enterprise Vault application prevents the manual starting and stopping of the Enterprise Vault services and only accepts service control requests from Veritas Cluster Server, because manually stopping a running Windows service will trigger a Veritas Cluster Server failover.

The Veritas Cluster Server shared volumes must be configured to store the Indexing Service data, Shopping Service data, Vault Store Partitions, .PST holding folders, and EMC® Centera staging areas.

Achieving High Availability with Symantec Enterprise Vault

Also, DNS aliases are still used for client connections to Enterprise Vault, but these aliases should be pointed to the virtual cluster name rather than to specific nodes.

Installation

There are several options for installing and configuring Veritas Cluster Server and Enterprise Vault. If this is a new installation of Enterprise Vault, then install Veritas Cluster Server on all the nodes before installing and configuring Enterprise Vault. If this is an existing installation, then Enterprise Vault can be upgraded to a cluster configuration.

The order in which the administrator installs and configures the various components in a clustered environment is important. The steps are outlined here and given in more detail in the Enterprise Vault documentation called *Installing_and_Configuring.pdf*, which is available in the Enterprise Vault media kit.

New Enterprise Vault installations

1. Ensure that all prerequisite components have been installed on each of the cluster nodes.
2. Install Veritas Storage Foundation HA *for Windows* 4.3 MP1 or higher on each Enterprise Vault node to be clustered. Note: MP1 is applied after installing the base version of Veritas Storage Foundation HA *for Windows* 4.3.
3. Configure the cluster by running the Veritas Cluster Server configuration wizard.
4. Install Enterprise Vault Server components on all the nodes in the cluster.
5. Configure the disk group and volumes from the first node. These must store the following Enterprise Vault storage components:
 - Indexing Service data
 - Shopping Service data
 - Vault Store Partitions
 - .PST holding folders
 - EMC Centera staging areas

Note: Symantec also recommends creating separate volumes to store the MSMQ and registry replication data.
6. Mount the volumes on the system where the Enterprise Vault service group will be configured.

Achieving High Availability with Symantec Enterprise Vault

7. Configure the Enterprise Vault service group. These resources include:

- IP address
- Computer name (LANMAN resource)
- MSMQ
- Disk/storage (MountV and DiskGroup resources)
- Service resources

8. Run the Enterprise Vault configuration wizard on each node in the cluster. The configuration wizard will automatically detect the presence of Veritas Cluster Server on the server and lead the administrator through a specialized sequence of configuration steps for clustered implementations (see Figure 2). The options selected in the wizard will vary depending on the chosen failover configuration (Active/Passive, N+1, and so on). The [Installing_and_Configuring.pdf](#) document contains thorough documentation, with screen shots, about the appropriate wizard options for each failover configuration.

9. Test that the nodes in the cluster fail over correctly.

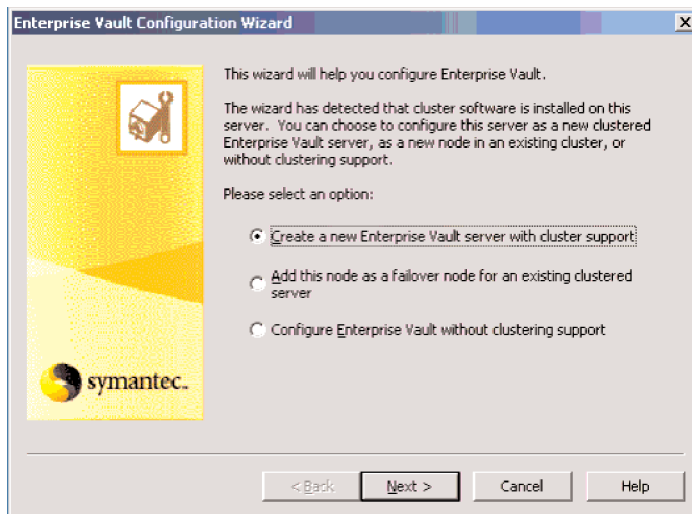


Figure 2. Enterprise Vault configuration wizard when Veritas Cluster Server is detected.

Upgrading existing Enterprise Vault servers to a cluster

There are two types of Enterprise Vault installations that can be upgraded to a cluster: a single, nonclustered Enterprise Vault server, and a building-block configuration that contains multiple Enterprise Vault servers. To be eligible for upgrade to a cluster, the Enterprise Vault installation must meet the following conditions:

- Enterprise Vault should already be configured in a nonclustered configuration.
- Enterprise Vault must be configured using DNS aliases rather than fully qualified server names.
- The Enterprise Vault server must have a full set of Indexing, Shopping, Task Controller, and Storage Services. However, it must not contain the Microsoft SharePoint® Portal Server 2001 service, because this is not supported in a cluster.
- In a building-block environment, an Enterprise Vault server that is hosting services must not be running in failover mode.

To upgrade an existing installation to a clustered Enterprise Vault environment, perform the following steps (again, more detail is available in the *Installing_and_Configuring.pdf* document):

1. First, verify that all Enterprise Vault storage locations (Indexing locations, Storage partitions, Shopping locations, .PST holding folders, and EMC Centera staging areas) are all on highly available storage devices. Follow the instructions in the product documentation to move these locations to different storage if necessary.
2. Run the Enterprise Vault Cluster Setup wizard to create an Enterprise Vault cluster service group and add to the group the server to be configured.
3. Run the “Convert to Cluster” wizard found in the Start Menu.
4. When the wizard prompts the choice of a service group in which to create the cluster resources for each Enterprise Vault service, select the group created earlier.
5. Modify the existing DNS alias to point to the virtual server name rather than the local name.
6. Use Veritas Cluster Manager to bring the resources in the cluster online.

Failover process

As mentioned previously, Veritas Cluster Server uses the Generic Service agent to monitor the Enterprise Vault services on different nodes based on the information in the Enterprise Vault Directory database. The agent brings the following services online, monitors their status, and takes them offline:

- Admin Service
- Directory Service
- Indexing Service
- Shopping Service
- Storage Service
- Task Controller Service

Note that Veritas Cluster Server currently manages Enterprise Vault *services*, not *tasks*.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for detailed information on the Generic Service agent, including the resource type definitions, attribute definitions, and sample configurations.⁵

The Generic Service agent detects an application failure if a configured Enterprise Vault service is not running. When this happens, the Enterprise Vault service group (including the virtual identity and IP address) is failed over to the next available system in the service group's system list, and the services are started on the new system. No manual intervention is required. Veritas Cluster Server also supports SNMP/SMTP and can be configured to send alerts to system administrators automatically in the event errors or failures are detected.

Failback process

Another significant advantage of managing failovers with Veritas Cluster Server is that there is no longer a need to fail back Enterprise Vault. Once the failed node is repaired or replaced, it will be configured as the new spare. There is no need to incur additional downtime in order to fail back the Enterprise Vault server. If for some reason the administrator would like to fail the application back to the original node, this process would be performed according to standard Veritas Cluster Server documentation.

⁵ http://iwww.veritas.com/home/pubs/wxrt_train/VSFW_5.0_Docs/to_mfg/VCS_BundledAgents.pdf

Microsoft Cluster Server integration

Since Microsoft Cluster Server (MSCS) is the number one clustering solution used today for Microsoft Exchange and Windows file servers, many organizations may wish to use the same clustering technology to provide a high-availability solution when archiving content from those applications.

Like the Veritas Cluster Server solution, high availability is provided for MSCS by creating an Enterprise Vault virtual server that can fail over between physical nodes in the cluster. When Enterprise Vault services are running on a virtual server, they operate with virtual IP addresses, a virtual computer name, virtual Microsoft Message Queues, and highly available shared disks. When a failure occurs, the cluster software can automatically move the virtual server's resources to a different physical node in the cluster.

Supported failover configurations

Like the Veritas Cluster Server solution, the failover configurations currently supported for Enterprise Vault and MSCS are Active/Passive, N+1 and N+M. Active/Active configurations are not yet supported.

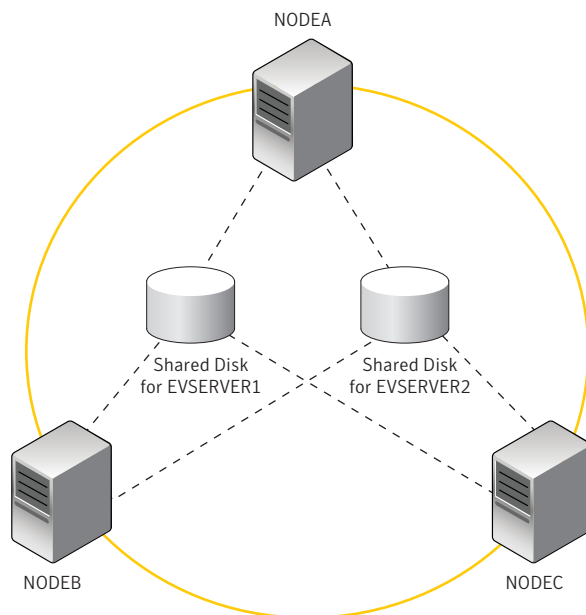


Figure 3. Example of an N+M configuration.

Achieving High Availability with Symantec Enterprise Vault

Requirements for MSCS integration

The following software is required on each primary and failover node:

- Microsoft Windows 2003 Enterprise Edition or Datacenter Edition
- Microsoft Exchange System Manager, if the cluster will archive from Microsoft Exchange
- Symantec Enterprise Vault 7.0 or later

Additionally, all Enterprise Vault storage must be under the control of MSCS. The shared volumes must be configured to store the MSMQ data, Indexing Service data, Shopping Service data, Vault Store Partitions, .PST holding folders, and EMC Centera staging areas. It is good practice for MSMQ data, Indexing Service data, and Storage Service data to each have a separate physical disk resource. Placing them on the same drives may result in degraded performance.

All Enterprise Vault services will be set by the configuration wizard to manual startup in the Windows Services control panel. In a clustered configuration, code in Enterprise Vault prevents a manual start or stop of Enterprise Vault services using the Administration Console, Services MMC snap-in, or the EVService utility. To safely start or stop Enterprise Vault services, use only Cluster Administrator or the Windows command-line utility `cluster.exe`.

Also, DNS aliases are still used for client connections to Enterprise Vault, but these aliases should be pointed to the virtual cluster name, rather than to specific nodes.

Installation

As with the Veritas Cluster Server configuration, the installation order is very important, and fully documented in the `Installing_and_Configuring.pdf` document. If possible, ensure that Exchange System Manager is installed on each node before the MSCS cluster is created. If the cluster has been created already, then the Microsoft Distributed Transaction Coordinator (MSDTC) must be present as a cluster resource before the Exchange System Manager can be installed (see the Microsoft Knowledge Base article 312316). Once Exchange System Manager is installed, the MSDTC resource can be removed.

New Enterprise Vault installations

1. Ensure that all prerequisite components for both MSCS and Enterprise Vault have been installed on each of the cluster nodes.
2. Set up the shared disks and volumes for the cluster.

Achieving High Availability with Symantec Enterprise Vault

3. Use Cluster Administrator to create the cluster and to add the primary and failover nodes.
4. Set up a resource group, including the prerequisite resources, for each Enterprise Vault virtual server required. MSCS resources include IP Address, Network Name, Physical Disk, and Message Queues.
5. Create a static DNS host entry and an alias entry for each Enterprise Vault virtual server.
6. Install Enterprise Vault server components on all the nodes in the cluster.
7. Run the Enterprise Vault configuration wizard on each node in the cluster. The configuration wizard will automatically detect the presence of MSCS on the server and lead the administrator through a specialized sequence of configuration steps for clustered implementations. The options selected in the wizard will vary depending on the chosen failover configuration (Active/Passive, N+1, and so on). The *Installing_and_Configuring.pdf* document contains thorough documentation, with screen shots, about the appropriate wizard options for each failover configuration. At a high level, the wizard includes the following steps:
 - a. Decide whether this is a new primary Enterprise Vault server, or a passive server for an existing cluster.
 - b. Specify the resource group to use for this server.
 - c. Establish the DNS identity for this server (aliases should point to virtual server name).
 - d. Configure storage paths used for Indexing Service and Shopping Service data (and validate that shared disks are used).
 - e. Enterprise Vault resources and dependencies are then automatically created using MSCS APIs.
 - f. Generate a report.
8. After configuration, the cluster should be tested to verify that the nodes fail over correctly.

Upgrading existing Enterprise Vault servers to a cluster

If an organization has an existing Enterprise Vault 7.0 or later installation then, subject to certain restrictions, the administrator can use the Enterprise Vault “Convert to Cluster” wizard to convert the Enterprise Vault servers to servers with cluster support. The conversion requires the administrator to move the Enterprise Vault data manually to highly available locations.

Achieving High Availability with Symantec Enterprise Vault

To be eligible for conversion to a cluster, the existing Enterprise Vault installation must meet the following conditions:

- Enterprise Vault should already be configured in a nonclustered configuration.
- Enterprise Vault must be configured using DNS aliases rather than fully qualified node names.
- The Enterprise Vault server must have a full set of Indexing, Shopping, Task Controller, and Storage Services.

Note that the administrator can configure a combination of new and existing Enterprise Vault servers as virtual servers, if required, and a new installation of Enterprise Vault must be deployed on the nodes that are to act as failover nodes.

The process to convert existing standalone Enterprise Vault servers into an MSCS cluster is very similar to the steps outlined above for Veritas Cluster Server. A suitable resource group must exist and be online on the Enterprise Vault server node, and then the administrator can run the “Convert to Cluster” wizard from the Start Menu. The wizard performs a number of checks to make sure the existing installation is cluster-ready, such as having all the core services present on the existing server, and using only highly available storage devices for all Enterprise Vault storage locations. The wizard then creates the necessary resources, updates the Enterprise Vault services to manual startup, and updates the Directory database tables to remove the local computer name from the computer entry table and the message queue names. Finally, DNS aliases must be updated to point to the virtual server name rather than the local node names.

Failover process

If an active node fails, the Enterprise Vault virtual server attempts to fail over to the next available node in the resource group's preferred node list, assuming all the resources have that node as a possible owner. The Server Instance resource fails over first, provided the failover node is not already running an Enterprise Vault virtual server. The remaining resources then fail over in order of dependency. The resources start the Enterprise Vault services on the failover node, ensuring continuing availability for the data that Enterprise Vault is managing and archiving.

Achieving High Availability with Symantec Enterprise Vault

Failback process

As with the Veritas Cluster Server and USL configuration options, there is no need to fail back Enterprise Vault after a failover, because the repaired node can be reconfigured as either an active or passive node for future use. After repairing the failed node, Enterprise Vault should be installed on the node and the system should be configured as an active or passive node in the cluster. If the administrator chooses to configure the repaired node as active again, he or she would use Cluster Administrator to change the ownership of the resource group to this repaired server.

Conclusion

Symantec's mission is to provide the essential tools to help its customers protect the security and availability of their information. We are proud to provide "cluster aware" support for both Veritas Cluster Server and MSCS, in addition to providing an application-level high availability solution that can reduce downtime to only a few minutes. This is one of the many reasons why the Symantec Enterprise Vault solution is regarded as the clear leader in the email archiving market, according to both market share and analysts' opinions. As both the volume and importance of an organization's archive grow over time, Enterprise Vault customers can find peace of mind in the knowledge that their archiving platform excels in the scalability and high availability they need to fulfill their business requirements.

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world.

The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries.

More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

Copyright © 2007 Symantec Corporation. All rights reserved. Symantec, the Symantec Logo, Enterprise Vault, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Microsoft, SharePoint, Windows, and Windows Server are registered trademarks of Microsoft Corporation in the United States and other countries. Other names may be trademarks of their respective owners. Printed in the U.S.A.
03/07 11856752