symantec.

Confidence in the connected world.

Enhancing Microsoft® SQL Server 2005 Availability with Veritas Storage Foundation™ for Windows®

Saveen Pakala, Product Manager
Paul Belk, Technical Product Manager
Storage and Availability Management Group

White Paper: Symantec Technical

Enhancing Microsoft® SQL Server 2005 availability with Veritas Storage Foundation™ for Windows®

Content

Enhancing Microsoft SQL Server 2005 availability with Veritas Storage
Foundation HA for Windows from Symantec

Executive Summary

Historically, administrators have protected servers and applications by backing up data to tape.
Although regular tape backups are essential, recovery time objectives and recovery point
objectives cannot always be met by traditional backup solutions alone. Today, Microsoft SQL
Server administrators must choose between the many built-in tools contained in Microsoft
Windows Server 2003 and Microsoft Windows Server 2008 in conjunction with SQL Server 2005
or invest in tools from other vendors. In some cases, a combination of technologies is the best
way to achieve business goals because many tools that meet one type of business requirement
don't necessarily meet the requirements of another. Finally, not all versions have the necessary
features included.

Veritas™ Storage Foundation HA for Windows from Symantec is designed to support the Standard
Edition and the Enterprise Edition of SQL Server; all versions of the Microsoft Windows 2003
Server and Windows Server 2008 operating systems, including both 32-bit and 64-bit versions;
and VMware and Microsoft Virtual Server environments. SQL Server 2005 is available in four
different versions: Express, Workgroup, Standard, and Enterprise. For organizations running SQL
Standard or Enterprise, data availability is an issue that must be explored.

The goal of this document is to discuss the threats to availability and the tools available to
address them. For organizations that don't have stringent uptime requirements, many of the
features described may not be useful, but for others, they may prevent the next calamity from
negatively affecting their business.

Threats to data availability

The following list describes the threats to availability. Any one of these could seriously affect an
organization for an indefinite period in the absence of proper planning and plan execution. While
most downtime is caused by human error, sabotage and natural disasters are risks that cannot
be ignored.

- **Data corruption**—Accidental or malicious updates can cause physical or logical data
  corruption.
- **Component failures**—A single piece of hardware can bring down a server or an entire
  organization.
- **Application failure**—If an application crashes or a system reboots, critical data will be
  unavailable.
- **Maintenance**—Routine maintenance and patching can disrupt the availability of data and
  applications.

• **Site outage**—Even when you fully protect against the previous threats, if a complete site goes
  down, the entire business may come to a screeching halt.

How do you plan to bring your applications online after suffering any one of the previous events?
The following chart indicates the products or functionalities that will best protect against the
threats identified in each row.

Threat	Backup/restore	Recovery from snapshot	Clustering	Replication/DR
Data corruption	✓	✓	✗	✗
Component failure	✗	✗	✓	✓
Application failure	✓	✓	✓	✓
Maintenance	✗	✗	✓	✓
Site outage	✗	✗	✗	✓

Feature sets are available that incorporate the basic functionality to protect against these
threats. While tape recovery can eventually recover any protected server, quick recovery from
database or full copy volume snapshots should be the first line of defense against data
corruption. Clustering and replication cannot protect against corruption because whatever
corruption exists on the shared disk will be failed over or replicated to the failover target. The
same issues will propagate, eventually forcing recovery from snapshot or tape.

How much data loss or downtime is acceptable?

While each threat can have a different root cause, the impact on the business is measured in
terms of data loss or length of outage, regardless of the cause. The real cost incurred by the
organization depends on the loss of sales or productivity as a result of the downtime. The
maximum time that an outage can be tolerated is called the Recovery Time Objective (RTO). The
amount of data loss that can be tolerated is the Recovery Point Objective (RPO). RTO and RPO are
used to define service level agreements (SLAs).

If you can meet both your RTO and RPO with tape backup and restore, this paper may not yet be
relevant to your organization. However, if you cannot meet the objectives with traditional tape
backup or you are approaching the limits of the current solution, this document describes the
suite of tools available so that you can evaluate and select the appropriate options for your
situation.

Enhancing Microsoft SQL Server 2005 availability with Veritas Storage
Foundation HA for Windows from Symantec

Microsoft has released a number of powerful features within the Windows Server 2003 and
Windows Server 2008 operating systems. When these are combined with the latest features in
SQL Server 2005, many environments will be protected well beyond their primary objectives.
There are some areas where both Microsoft and Symantec provide equivalent capability, and
other factors such as corporate standards may dictate the solution of choice. However, some
environments impose stringent availability requirements and cannot tolerate any unplanned
downtime for any reason. These are the environments for which Veritas Storage Foundation HA
for Windows was created.

Determining availability requirements
It is important to answer the following questions to determine availability requirements:

• What are your anticipated SQL Server scalability and memory requirements?

• How much data will reside on the server, and how long will it take to back up and restore the
   data?

• What are the behavior characteristics of your database workload (OLTP, decision support, etc.)?

• What are the uptime and performance requirements for this application or the line of business
   it serves?

• After a failure, can the application be restored to production using tape backup alone in the
   time required?

• What other applications does this application depend on?

• What applications depend on this application?

• Is there a specific order in which applications must be brought online or offline?

• Could a second data center also host this application?

• If a second data center is available, what type of connectivity exists between sites?

Assessing and meeting the requirements
To meet your requirements, it is also important to determine the following:

• How to eliminate single points of failure

• Hardware—number of systems, arrays, sites, etc.

• Redundant storage—number of copies of data that must be maintained for recovery

• Frequency of backups and snapshots

• Redundant SAN—number of fabrics, paths, controllers, etc.

• Redundant network infrastructure

• Number of failures that can be tolerated

• Number of failover nodes

• Distance between systems or data centers

• Network infrastructure and bandwidth

Technology overview

The following sections provide an overview of and comparison between the built-in features of
SQL Server and Storage Foundation HA for Windows. The discussion and charts make it easy to
see the features that are included with each. Additionally, there are many perceived similarities
between snapshot technologies, clustering, and replication, so it is important to make sure that
the required features are available in the choice that you make for your organization.

New availability features in SQL Server 2005

The following availability features are new in SQL Server 2005:

• Partitioning across file groups for very large databases (VLDBs)

• Online database restore

• Increased number of instances from 16 to 50 per server (leading to greater impact if a failure
 occurs)

• Support for high-performance 64-bit architectures

• Database snapshots using copy-on-write (COW) technology to create database views and
 recovery options

• Database mirroring and replication enhancements

• Backup and recovery enhancements

Database snapshot vs. full copy volume snapshot

After ensuring that a successful backup strategy is in place, it is necessary for most organizations
to plan for quick recovery snapshots so that recovery from data corruption can be achieved

without restoring from tape. Microsoft's Volume Shadow Copy Services (VSS) and Volume Disk
Interface (VDI) are the supported methods for creating SQL snapshots, and through these
interfaces, backups can be obtained by taking online snapshots of the production user
databases. Microsoft has provided VSS so that all vendors can take proper snapshots of the SQL
database while it remains online. SQL 2005 database snapshots use COW technology to make a
copy of the data prior to change; the unchanged data remains in the original volume. Symantec's
FlashSnap option takes full point-in-time snapshots of the entire volume where the database
resides. There are advantages and disadvantages to both approaches, so it is important to use
the technology that best meets the objectives of your enterprise environment.

SQL Server 2005

• Snapshot and source database must be on the same instance.

• If any database snapshots exist on a source database, the source database cannot be dropped
  or restored (must delete all of the database snapshots first).

• Backups are not allowed on the snapshot; however, source database backups are unaffected by
  database snapshots.

• Cannot be used on log shipping secondary.

• Database snapshots cannot be created for the model, master, and temp databases.

• Database snapshots use COW technology to conserve disk space, but the data is exposed to a
  single point of failure, as it relies on the original volume.

Storage Foundation HA for Windows

• Performs persistent full split-mirror snapshots of SQL Server database volumes.

• Snapshots can be used for off-host for processing, reporting, backup, or fire drill.

• Does not require access to original data volume.

• Loss of original volume does not affect snapshot and hence is not a single point of failure.

• Requires additional space equal to original data volume.

• Resynchronization from snapshot volume can be done even if original volume is lost/corrupt.

• Faster performance than on-host COW provider.

• Snapshots are readable and writeable.

Replication technologies for SQL

SQL contains a number a technologies that involve replication. For disaster recovery (DR)
purposes, log shipping has been used for many years. Log shipping has been enhanced and
automated in the new Database Mirroring functionality that will be released in SQL Server 2005.
SQL also has three other modes of replication. Peer to peer is the most likely technology to be

used for disaster recovery. Storage Foundation for Windows offers the Veritas™ Volume Replicator Option, which performs replication and automatic role reversal when used in conjunction with Symantec's clustering product Veritas™ Cluster Server, which is included with Storage Foundation HA for Windows. The single largest advantage that volume-based replication has over database replication is that data from the system databases is replicated to the secondary data center. This is important because structural or security changes that have been made to the database would be lost in the other methods, slowing down SQL recovery and requiring manual intervention.

Clustering technologies

Microsoft Cluster Server now supports clustering SQL 2005 Standard Edition on Windows Server 2003, Enterprise Edition. This creates an advantage for smaller but still important databases that don't require the extra capabilities provided by the more robust operating system or SQL version. Storage Foundation HA for Windows provides support for any server version of the SQL application and operating system. In supporting up to 32 nodes, larger environments can take advantage of fewer failover targets per cluster, reducing the expense of idle servers. Additionally, the 50-instance support now offered in SQL Server 2005 makes it possible to create very large clusters with many instances. And the adaptive workload management features of Veritas Cluster Server enable intelligent dynamic failover policies to help ensure that resources are not overwhelmed if multiple failovers occur in a short period of time. Veritas Cluster Server is also the foundation for a solid DR solution and includes a cluster simulator and a feature called "Fire Drill" that provides automated DR testing without affecting the production environment.

Clustering SQL Server 2005 servers for protection from component failures

Organizations can protect SQL environments from a wide range of component failures by implementing local clustering for availability. Storage Foundation HA for Windows integrates award-winning Veritas Cluster Server technology, which provides highly flexible, scalable failover clustering with workload management capabilities. In a Veritas Cluster Server cluster, multiple servers are linked with shared storage and private Ethernet heartbeats. Each system in the cluster can access the storage of any other system. Application-specific agents monitor the various components and dependent application resources for failures.

Although it's possible to share storage with switched SCSI devices, most clusters implement SAN-based storage or iSCSI to support a larger number of nodes. Veritas Cluster Server implements a policy-based response to failures, attempting a restart and/or automatically moving a failed application to another resource in the cluster. It has a built-in sense of cluster load and can relocate application components based on available capacity augmented by a rules-based policy engine.

Veritas Cluster Server offers a highly scalable, flexible alternative to Microsoft's native clustering
capabilities and includes:

• Support for up to 32 nodes and 4 global cluster sites

• Policy-based failover

• Workload management

• Automated DNS updates

• Support for failover across subnets

• Support for mount points

• No restrictions on the SQL Server configuration.

Cluster configurations

Storage Foundation HA for Windows supports either active/passive or any-to-any cluster
configurations for SQL Server.

• In an active/passive cluster, one "spare" system serves as the "standby" or passive system for a
   production server.

• In an any-to-any cluster, SQL Server 2005 can fail over to any server in the node that is not
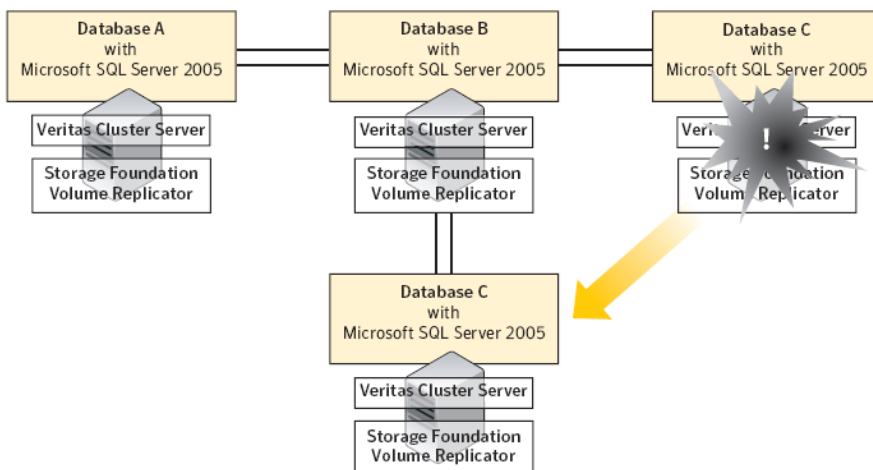   already running SQL Server and that has adequate capacity (see Figure 1).

Figure 1. A SQL Server database instance fails over to another server not running SQL Server that has available capacity.

It is important to note that with Veritas Cluster Server, the cluster can include numerous other
servers running different applications, consolidating management of the high-availability
environment and providing spare capacity for critical applications. Veritas Cluster Server clusters
can support up to 32 nodes.

9

Cluster management

The Veritas Cluster Server Web-based Cluster Management Console (VCS MC) simplifies the task
of managing multiple clusters (see Figure 2). VCS MC provides a centralized GUI to monitor,
manage, and configure every Veritas Cluster Server cluster running in multiple data centers. With
VCS MC, it is possible to set cluster attributes for a single cluster or to update values globally
across all clusters with the click of a button. Also, users can run detailed historical reports that
measure SLAs in aggregate and track results over time. Notification policies across all clusters
can be easily modified. Administrator rights and roles can be set for different users in the IT
organization. Most importantly, VCS MC provides a single place to instantly view the health of all
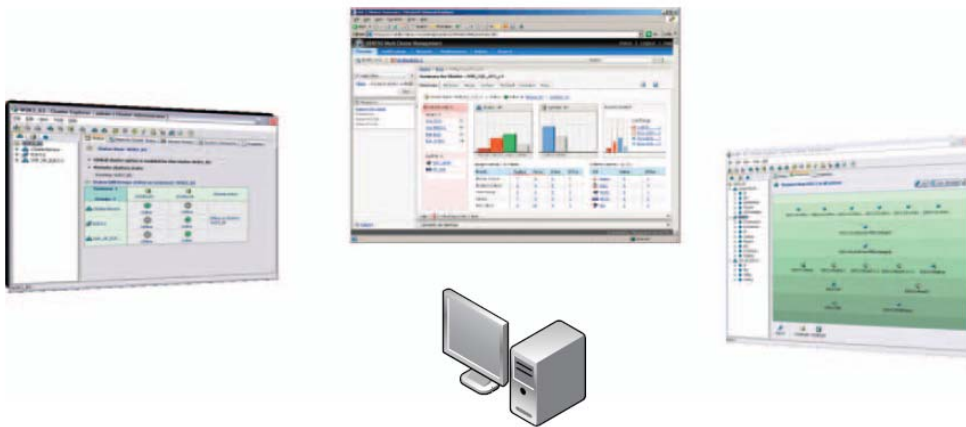clustered applications in the data center.

Figure 2. Cluster Management Console: Managing multiple Cluster nodes across sub-nets
and providing a single consolidated view.

Cluster service groups

Veritas Cluster Server encapsulates the resources required for each application into a service
group, creating a virtualized application that can be moved between the nodes of the cluster. Each
SQL service group contains a set of dependent resources—the lower level components that an
application requires to operate successfully. Resources include disk groups, disk volumes, file
systems, IP addresses, and NIC cards. A typical SQL configuration has many components from
storage to hold the SQL databases, network identity, and SQL components, such as the SQL server
service and SQL Server Agent. Veritas Cluster Server has agents that are designed to manage and
monitor these components, helping to ensure high availability. Note that there are several layers
of components (see Figure 3). The lowest layer will come online before any of the upper layers. So
in the case of SQL, the SQL components cannot come online unless the storage and network

components are ready and available. You do not want to be in a position in which SQL attempts to come online and there are no databases to mount because the storage is not available. This is an intuitive way to view the service group and how it behaves within the cluster.
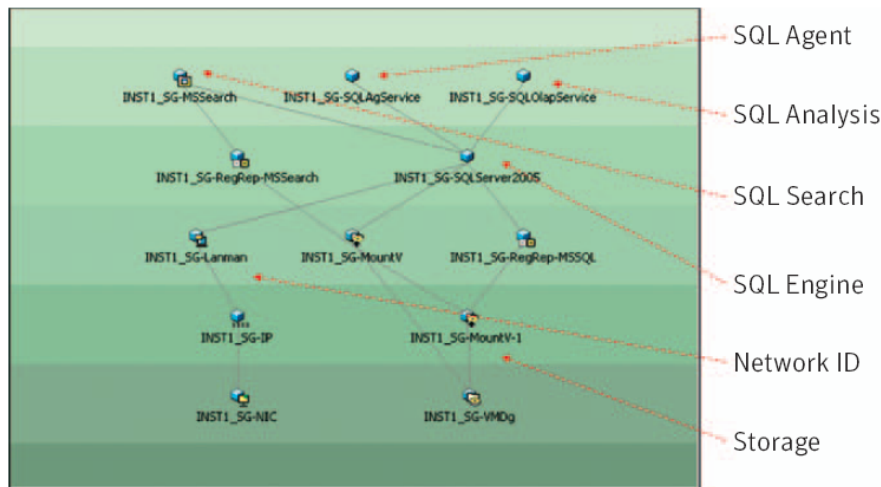

Figure 3. Veritas Cluster Server dependency graph: Veritas Cluster Server works in multi-tiered applications.
It is application and resource aware and enables orderly/sequenced failover and failback.

Once the cluster is configured, administrators can define failover policies for critical applications. Veritas Cluster Server supports flexible, customized failover policies that combine rules with workload management. Policies are configured using a graphical interface or command line interface, and they may be supplemented by custom scripts, if required. Storage Foundation HA for Windows also offers a simulator for testing failover. For example, policies might specify which applications have failover priority. Storage Foundation HA for Windows provides the opportunity to test failover to make sure it occurs as you expect it to—before an actual crisis occurs. Local clustering with Storage Foundation HA for Windows addresses many of the challenges of SQL administration:

• The software automatically monitors all of the application components and responds appropriately in case of a problem, failing over to other available resources, if necessary.

• Administrators can proactively switch applications to other resources to perform routine maintenance or upgrades on application components, such as server upgrades or OS patch applications.

Veritas Cluster Server supports the 32-bit and 64-bit versions of SQL 2000 and SQL 2005 and all components of SQL. A very significant feature of the SQL agent is Detailed Monitoring. Standard monitoring of SQL monitors SQL services, but there may be a concern about database access.

Detailed monitoring allows the agent to open the database and perform any operation via a SQL script.

If the script fails, the agent reports it to Veritas Cluster Server. A sample script is included with the agent, so no SQL scripting is necessary. Storage Foundation HA for Windows strives to make the process as easy as possible and provides intuitive, easy-to-use wizards to deploy the SQL solution.

Protecting SQL Server from localized failures with campus clusters

Local clusters are, by definition, local. Although they can survive the failure of individual components, they cannot protect you from the loss of an entire site. A power system failure or fire at the data center would disrupt service.

Many large organizations enhance recoverability by stretching the clustering across different facilities that are nearby but in the same general area. These campus clusters, or metropolitan clusters, combine high availability with an intermediate level of disaster recovery. They can protect against many localized disasters, but not against major area-wide disturbances like hurricanes and earthquakes.

The limiting factor in the campus cluster is the distance between sites, in particular, the storage arrays. Since a campus cluster uses a volume (or volumes) that is comprised of disks in different locations, disk latency is a critical factor. The data is mirrored across the cluster nodes using a SAN infrastructure. For example, a single cluster could span two distinct sites, with a SAN infrastructure in place between them. The SQL database administrator uses Storage Foundation HA for Windows to mirror the SQL data volumes automatically between the two locations. Once the distances start to impact application performance then the better solution is a "replicated data cluster" (This approach is described in the next section). The data is mirrored across the cluster nodes using a storage area network (SAN) infrastructure (see Figure 4). For example, a single cluster could span two distinct sites, with a SAN infrastructure in place between them. The SQL database administrator uses Storage Foundation HA for Windows to mirror the SQL data volumes automatically between the two locations.
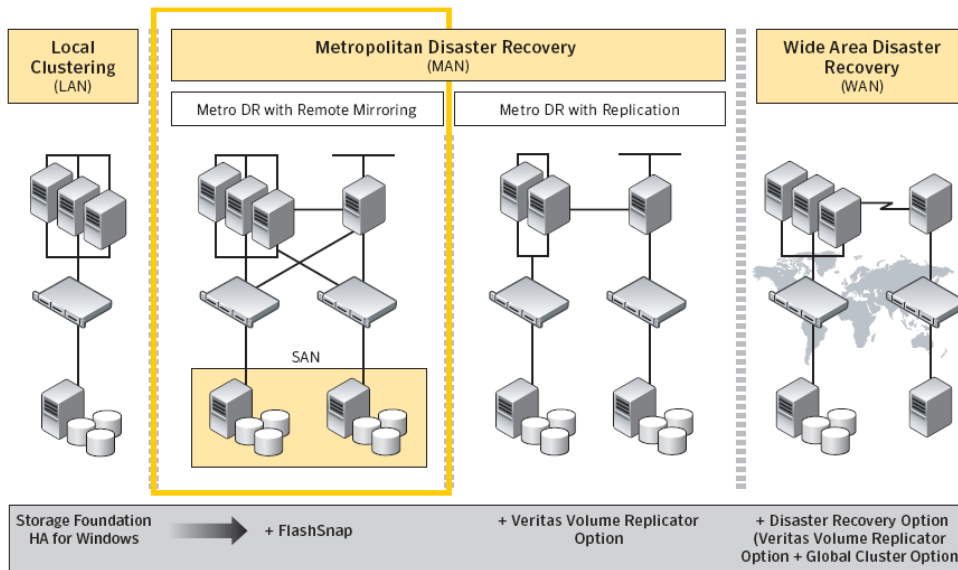
Figure 4. Metropolitan area network with remote mirroring.

For example, a financial company in Manhattan might maintain fiber connections to a data center
in New Jersey and mirror data between the two sites. In the event of a failure at the Manhattan
facility, the application automatically switches over to the New Jersey site and service is restored
within minutes, with no data loss. This configuration works with Storage Foundation software's
volume-based mirroring or third-party arrays that offer data mirroring.

Although the network infrastructure is needed to support communications between servers and
between mirrored volumes, campus clustering with mirroring requires no additional software
licenses beyond those available with Storage Foundation HA for Windows.

Protecting SQL Server with replicated data clusters

For organizations that do not have the SAN infrastructure in place, it's also possible to create a
campus cluster using data replication—either volume replication through Storage Foundation or
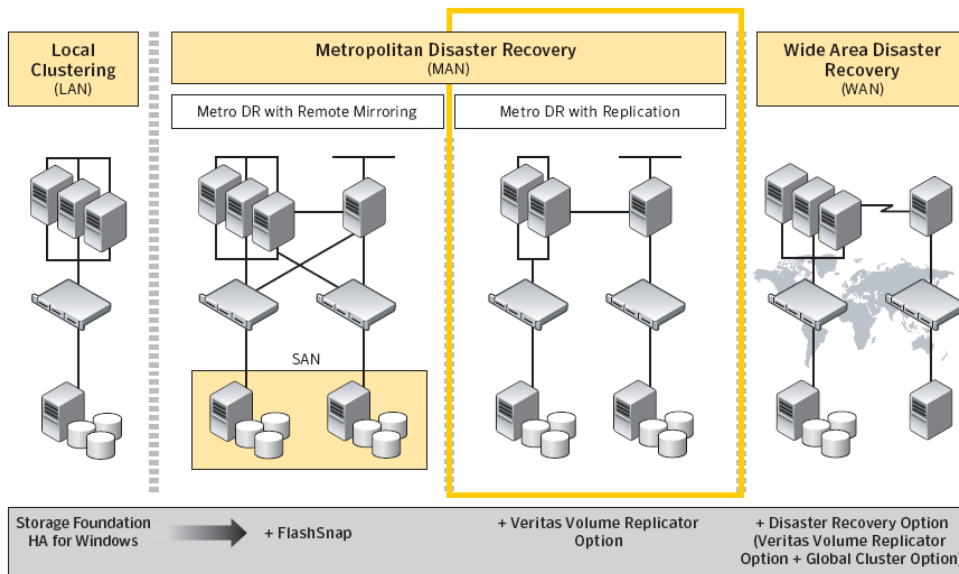replication provided by a third-party array (see Figure 5).

Figure 5. Metropolitan area network with replication.

The data is replicated across sites, which eliminates the need for the SAN infrastructure across
sites. Synchronous replication provides the best data protection, although Veritas Volume
Replicator supports an option that switches to asynchronous replication when network latency is
a problem.

For organizations with primary data centers in lower risk areas (without the threat of widespread
problems from earthquakes or hurricanes), maintaining a disaster recovery facility in close
proximity to the primary site is an easy and cost-effective approach. Most disasters are localized,
so campus clusters offer rapid disaster recovery for a wide range of potential problems. If you
need a large geographic separation between sites, then a wide area clustering solution is
preferable.

Disaster recovery with global clustering

For business continuity in the face of a broader disruption, the best protection is the ability to
resume operations, nearly instantly and without data loss, at an alternate site located a
significant distance from the primary site. Using Storage Foundation with the Global Cluster
Option (GCO) and the Veritas Volume Replicator Option, you can replicate data between two
separated sites and switch application services between them with a single mouse click (see
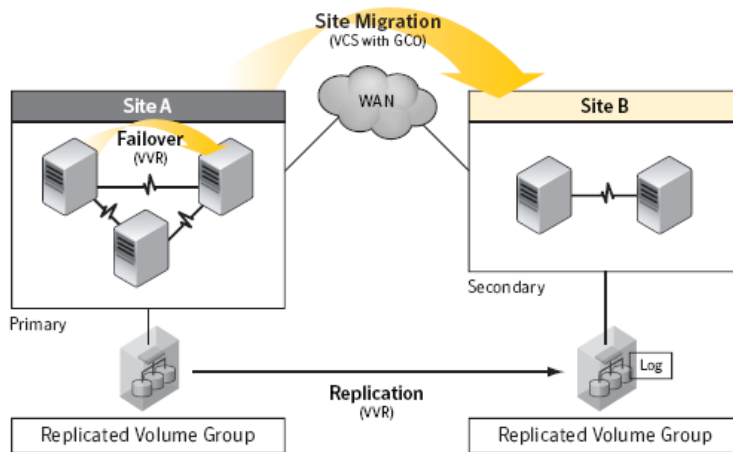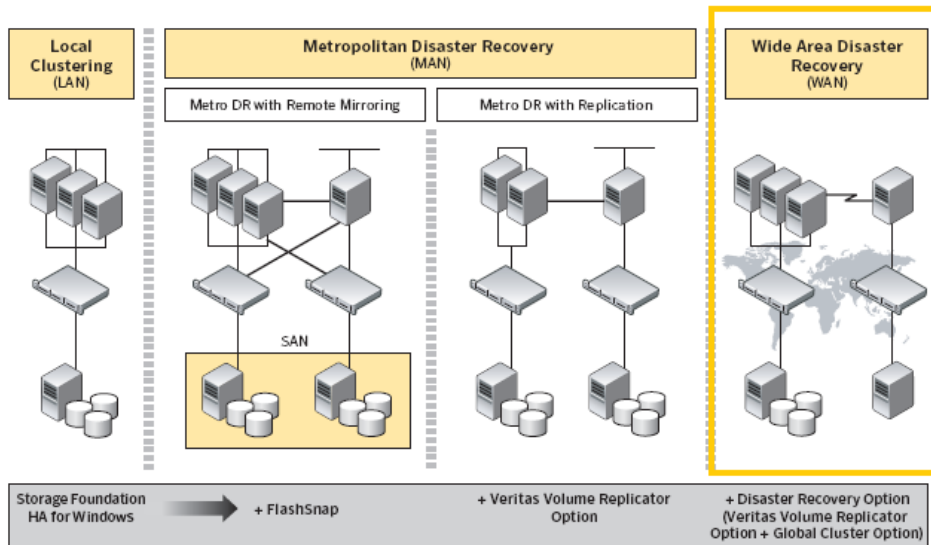Figure 6).

Figure 6. With Storage Foundation HA for Windows and the GCO and Veritas Volume Replicator Options, you can
replicate data between two separated sites Veritas Volume Replicator, and switch application services
between them (Veritas Cluster Server, GCO) with a single mouse click.

GCO creates a "cluster of clusters." The local cluster remains at one location. Should an
application fail, the clustering software will first try to restart it within its local cluster. If the local
cluster fails, then the Veritas Cluster Server software notifies the administrator, who can verify
the problem and move the application with the touch of a single button. Failover and replication
can occur over any distance.

You can replicate the data between sites using either the Veritas Volume Replicator Option
(which offers hardware-independent wide area replication over IP) or third-party hardware-based
replication.

Primary and secondary sites

Global clustering depends on the existence of an alternate site. The cost of the alternate site is a
major concern for many organizations—how many can afford to have a duplicate data center,
somewhere distant from the primary data center, ready to accept production processing at a
moment's notice in case of a disaster?

Using Storage Foundation HA for Windows helps organizations manage alternate sites without
damaging the IT budget. It offers:

• Server savings—The secondary site doesn't need to be identical in computing resources to the
   primary site. The DR site might have less capacity but still provide enough computing
   resources to run SQL Server should a disaster strike.

• Storage savings—The secondary site can use the same storage from a single vendor, different
   storage from a single vendor, or storage from multiple vendors, as long as the overall capacity
   is sufficient.

• Efficient and effective operation—Storage Foundation HA for Windows and the Veritas Volume
   Replicator Option do not require the same disk array on both sites of the replication link and
   work effectively in a heterogeneous storage environment.

• Alternate uses for the DR environment—Many organizations actively use the DR site for other
   purposes that could be stopped in the case of a true disaster. For example, the secondary site
   could host development and QA processes. You would then configure, as part of the failover or
   switchover process, the graceful shutdown of development processes before bringing the
   production instance online.

• Flexible global configurations—In a company with multiple data center locations, one site could
   serve as the DR site for multiple other data centers. In many cases, you can use existing
   resources to achieve a level of protection not previously available at a manageable cost.

Volume Replicator Option

To support application migration, the production data must be resident and up to date at the
secondary site. It's possible to use hardware-based replication, particularly if you already have
the solution in-house. The Veritas Volume Replicator Option offers the ability to consistently and
reliably replicate data across IP networks using any available storage.

The Storage Foundation HA for Windows Veritas Volume Replicator Option replicates the
contents of each volume across a wide area network to the secondary site. It is completely
transparent to the application components. Unlike traditional block-based approaches, Veritas
Volume Replicator replicates I/O instead of disk tracks to help ensure the data is always
replicated in a consistent fashion.

Veritas Volume Replicator supports both synchronous and asynchronous replication:

• Using synchronous replication, the initial write is not committed until the data has been
  replicated successfully. This guarantees that there will be no data loss in case of a site failure
  but affects application performance because the application at the primary site waits for the
  transaction to travel to the secondary site and back before the transaction is committed at the
  primary site. Over long distances, synchronous replication may introduce unacceptable write
  delays in production systems.

• Using asynchronous replication, Veritas Volume Replicator commits the data at the primary site
  immediately and then queues replication operations for network availability. Asynchronous
  replication does not affect the application performance at the primary site, but there may be
  some potential for data loss. Typical data loss between sites over long distances can be
  measured in milliseconds. Veritas Volume Replicator is unique in the market because it
  enforces write order on the replicated site, helping to ensure data integrity and consistency
  and guarantee that the data will be recoverable at the secondary site.

Most sites select asynchronous replication for long-distance, global failover scenarios. Veritas
Volume Replicator can dynamically switch to asynchronous replication if network latency is a
problem during synchronous replication.

Scenario: Site migration

To illustrate how the integrated Symantec solution can migrate the SQL Server environment
across sites for disaster recovery, following is a hypothetical failover scenario.

Consider the case of a company with a production facility running SQL Server in Dallas and a DR
failover site in Charlotte (see Figure 7). The Charlotte site could run other services, which would
then be shut down, in case of a failover. This example highlights the SQL application.
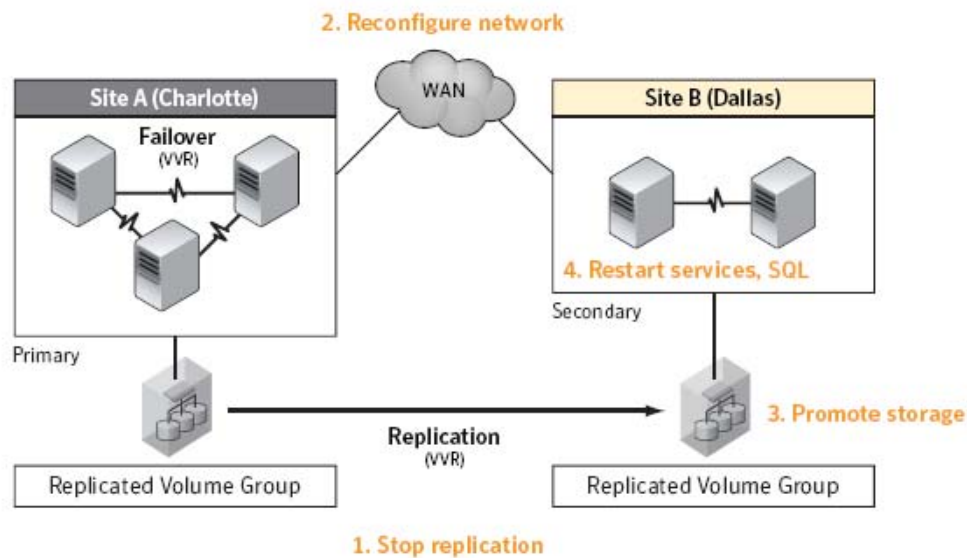
Figure 7. Steps that Veritas Cluster Server, GCO, and Veritas Volume Replicator would take to fail over
SQL from the primary site in Dallas to the disaster recovery site in Charlotte.

The Dallas site has a major problem, such as a power failure or flooding, and an operator is
notified that the site in Dallas is down. The operator logs on to a Web GUI from anywhere in the
world to verify the site is not available. At that point, the operator initiates the failover with a
single click to fail over SQL Server services.

The solution automatically performs the following steps:

1. **Stop replication:** In a failover, the first event is termination of replication between the primary
   and secondary sites. (If the primary site has experienced a dramatic failure, then replication has
   already stopped.)
2. **Reconfigure the network:** GCO reconfigures the logical network, migrating the virtual IP
   address for the application to the Charlotte site's systems.
3. **Promote the replicated data:** GCO promotes the secondary data to primary status. If a
   secondary site is still available, Veritas Volume Replicator can replicate to another site to
   maintain data availability.
4. **Restart services:** GCO starts the Cluster Server service groups, bringing SQL Server back
   online.

At this point, SQL is now running successfully at Charlotte—typically in a matter of minutes.
Users can access email without being aware that the application is running in a secondary site.

When the disaster has passed and the primary data center is ready, the Dallas storage can be updated, replication restarted, and the application switched back to its original location using a central console. The only data sent back to the Dallas site is the data that changed while operations were occurring in Charlotte; a complete reinitialization is not necessary.

Compare this scenario to the processes companies manage today. First, companies maintain the DR site at the same update and patch levels as the primary site—difficult when the site is physically removed from the primary data center. Coordinating change control for multiple sites for each production application presents a significant challenge. Then, when it's time to switch an application to the DR site, IT has to move tapes and staff between facilities, rebuild servers, and load operating systems and backup software while manually tracking and maintaining what's available at each site and relying on experts to bring up the environment at the DR site appropriately and consistently. The process can be very time consuming and expensive as well as subject to human error in the event of a complete site outage.

The Symantec solution eliminates change control problems, as all changes to the primary application (and its subsidiary components) are replicated to the DR site automatically. The failover is automated and nearly instantaneous, reducing risks while improving application availability.

A single integrated solution

Storage Foundation HA for Windows offers organizations that depend on SQL Server a single solution for protecting their business-critical data from a variety of problems, ranging from storage-caused outages to server failures and site-wide disruptions.

The Symantec solution is Microsoft certified and well integrated into the Windows environment. And it complies with key standards and frameworks such as MPIO, VSS, and VDS that enable tight end-to-end integration from application to backup systems and services. Microsoft and Veritas worked together for years to improve storage manageability on the Windows platform, and Veritas was selected to develop the disk management software for both Windows 2000 Server and Windows Server 2003. The Veritas Storage Foundation for Windows solution, now from Symantec, builds on the dynamic disk and dynamic volume capabilities now native to the Windows platform.

Whether your organization is most concerned with data corruption or business continuity, Storage Foundation HA for Windows provides a packaged, integrated solution tailored to the SQL Server environment, with the potential for expanding the solution through integrated agents and options.

Feature comparison

These tables provide an overview of the features to consider when architecting your SQL
implementation.

Feature	SQL Server 2005 on Windows Server 2003	SQL Server 2005 on Windows Server 2008	Storage Foundation HA for Windows
Tape backup and recovery	✓	✓	X
Dynamic disks on the SAN	X**	X**	✓
Multipath HBA support	X*	X*	✓
Drag-and-drop online volume growth	X	X	✓
Volume capacity monitoring and auto growth	X	X	✓
Drag-and-drop online data migration from one disk or array to another	X	X	✓
Hot spot detection	X	X	✓
Full copy volume snapshot	X	X	✓
Copy on write snapshot	✓	✓	X
Off-host backup/snapshot integration	X	X	✓
Support for mounted volumes (mount points)	✓	✓	✓

* Requires hardware vendor support for MPIO.

** MSCS plus Storage Foundation HA for Windows support dynamic disk groups. MSCS only supports basic
disks.

Table 1. Protection at the storage level with Veritas Storage Foundation HA for Windows

Feature	SQL Server 2005 on Windows Server 2003	SQL Server 2005 on Windows Server 2008	Storage Foundation HA for Windows
High-availability clustering up to 8 nodes	✔	✔	✔
High-availability clustering up to 32 nodes	X	X	✔
Customized behavior or custom applications	X	X	✔
Multi-tier application control	X	X	✔
Failover across subnets	X	X**	✔
Automatic failover between sites in campus cluster?	X	X	✔
Dynamic workload management?	X	X	✔
Data replication – user data only Log shipping, database mirroring, peer to peer	✔	✔	X
Data replication – system and user data	X	X	✔
Wide area disaster recovery with array-based replication	X	X	✔
Automated disaster recovery with testing and validation	X*	X*	✔
Troubleshooting SQL outside cluster control	✔*	✔*	✔

* Requires manual workaround
** Limited subnet failover capability

Table 2. High availability and wide area disaster recovery with Veritas Storage Foundation HA for Windows

Conclusion

Many factors both planned and unplanned affect information availability. By implementing
Veritas Storage Foundation HA for Windows from Symantec, organizations can benefit from tools
that are designed to increase availability and optimize SQL Server performance. Symantec is
working together with our many partners to enable a complete set of fully tested, integrated
hardware and software solutions that are designed to provide high availability, quick recovery,
and verifiable disaster recovery for Microsoft SQL Server 2005 environments. This scalable, easy-
to-manage approach—which includes Veritas Cluster Server, the key component of Veritas
Storage Foundation HA for Windows—is well suited for medium to very large enterprises that are
implementing mission-critical SQL Server servers in their data centers. Onsite planning and
assistance are available through Symantec professional services. Additional information can be
found on the following Web sites.

Symantec Technology Network home page

http://www.symantec.com/stn/index.jsp

Microsoft SQL Server home page

http://www.microsoft.com/sql/2005

For support or general information

http://www.symantec.com

http://support.veritas.com

Contributions:

Saveen Pakala, Product Manager, Storage Foundation HA for Windows

Paul Belk, Senior Technical Product Manager, Windows High Availability Platform

Nicole Kim, Senior Manager, Product Marketing, Storage Foundation for Windows

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