

VERITAS Storage Foundation™ HA *for Windows*

End-to-End Availability for Microsoft SQL Server

January 2005

Table of Contents

Executive Summary	1
Overview	1
The VERITAS Solution for SQL Server Availability	2
Solution Components.....	2
Creating a Resilient Storage Environment	3
Recovering Rapidly from Logical Errors	3
Protecting SQL Server from Component/Server Failures with Clustering	4
Local Cluster Configurations.....	5
Stretch or Metropolitan Clustering	6
Replicated Data Clusters	7
Replicating SQL Server Data	7
Disaster Recovery with Global Clustering	9
Summary	11

Executive Summary

Microsoft SQL Server is an integral part of many enterprise applications. But delivering end-to-end availability in a multi-tiered application environment is complex, requiring expertise beyond the database itself. Administrators must protect critical applications from loss of data, component failures, logical data errors or corruption, and even site-wide disruption.

VERITAS Storage Foundation™ HA for Windows offers an integrated, end-to-end solution for SQL Server availability. Using Storage Foundation HA with different agents and options, you can:

- Create a highly available data environment, with redundant storage configurations and online management capabilities.
- Create point-in-time images for disk-based backups and quick recoveries.
- Implement local or stretch clusters, eliminating downtime due to component failures and making it easier to maintain production systems.
- Replicate data between geographically-distributed clusters with automated failover, creating a flexible disaster recovery solution for SQL Server and the entire data center.

Overview

As Microsoft SQL Server has proven its capabilities in demanding enterprise environments, more companies are using SQL Server databases to serve data for critical applications. As a result, organizations now demand stringent levels of availability and performance from their SQL Server databases. Service level expectations are high.

Database administrators have the task of ensuring that data is accurate and consistent, database performance is adequate, and that database services are always available when needed. As the amount of data being managed and accessed grows, their task becomes ever more difficult.

Most SQL Server administrators are expert in SQL Server design and administration. But meeting end-to-end service levels requires more than database expertise. Today's database environment is characterized by a number of challenges to availability:

Complexity of application infrastructure: The SQL Server database is just one component of multi-tiered applications, which may include application and web servers. While the database administrator is responsible for the database itself, the end user cares about aggregate availability. A failure of any part of the application can interrupt access. Administrators need techniques for monitoring and managing end-to-end availability.

A multitude of risks: There are many potential risk factors that can disrupt access to database data. Components, such as storage subsystems, network routers, or power supplies can fail. Logical data problems like user or application errors can corrupt data. Localized disasters such as power failures can affect a building hosting critical infrastructure, while major disasters like a bad hurricane season can affect an entire state.

SQL Server administrators need the ability to identify, understand and respond to a large range of problems that affect service delivery – even if the source of the problems lies outside the SQL Server database itself. They need to be able to monitor and react to problems automatically and appropriately. They need to proactively maintain, upgrade and manage the SQL Server environment without disrupting user access.

In short, to deliver end-to-end application availability, you need a solution that solves a wide range of problems across multiple application components. VERITAS addresses these needs with a comprehensive, integrated solution for SQL Server availability that manages everything from local storage availability to regional disaster recovery.

The VERITAS Solution for SQL Server Availability

VERITAS offers a comprehensive, integrated solution for maintaining Microsoft SQL Server availability. The VERITAS solution reduces planned and unplanned downtime, simplifies the administration of complex environments and offers policy-based failover for disaster recovery/business continuity purposes.

VERITAS Storage Foundation HA *for Windows* combines VERITAS' industry-leading volume management technologies with flexible, policy-driven failover and application-specific agents. It expands and enhances the native volume management capabilities of Windows servers. Administrators can install the product once, then select among a variety of options that address specific availability and recovery concerns, such as global failover or online snapshots for quick recovery.

The solution includes support tailored to the SQL Server environment, such as application-specific monitoring and failover and application-aware snapshots for backups and quick recovery. Microsoft and VERITAS maintain a strategic partnership, integrating VERITAS technology in the Windows operating system and ensuring interoperability. You can find the VERITAS availability solution for Microsoft SQL Server running at Microsoft Technology Centers.

Solution Components

VERITAS Storage Foundation HA *for Windows* integrates the following software components:

- **VERITAS Storage Foundation *for Windows*** extends and enhances Windows 2000 and Windows Server 2003 Disk Management. Storage Foundation *for Windows* lets you configure and manage storage for optimal performance and availability, creating a scalable foundation for data growth.
- **VERITAS Cluster Server (VCS)** reduces planned and unplanned downtime by clustering critical applications and the resources they require. Application-specific agents monitor and manage the critical components of your IT environment for optimum application availability.

Storage Foundation HA *for Windows* offers the following options, each of which may be activated by a license key as needed:

- **Global Cluster Option** monitors and controls multiple, geographically-distributed VERITAS clusters and data replication. The Global Cluster Option lets administrators migrate entire data centers with a single click, enabling companies to survive serious local disruptions without significant interruptions to critical services.
- **Volume Replicator Option** delivers reliable, storage-independent replication over any IP network, facilitating rapid disaster recovery. The product replicates data at the logical volume level, ensuring integrity of replicated data.
- **Dynamic Multipathing Option** adds fault tolerance to disk storage by using multiple paths between a computer and storage.
- **FlashSnap Option** allows you to create point-in-time snapshots of volumes. These snapshots can be used for snapshot-assisted backups, off-host processing such as data mining and other purposes, or to provide disk-based recovery from logical errors. FlashSnap snapshots are integrated with the SQL Server Virtual Device Interface (VDI) for Microsoft supported and approved snapshots.

Complete descriptions of these products can be found at the VERITAS web site (www.veritas.com).

Creating a Resilient Storage Environment

At its most basic level, Storage Foundation *for Windows* helps the database administrator create a highly resilient storage configuration for SQL Server data.

Storage Foundation extends the disk and volume management capabilities of Windows 2000 and Windows Server 2003, enhancing dynamic disk storage to offer spanned, concatenated, striped and RAID volumes. It also provides online storage administration and drag-and-drop storage management.

Using Storage Foundation, you can create storage configurations with performance and availability characteristics that surpass the capabilities of the individual disk devices within them. For example, you can create mirrored disks for transaction logs, and mirror/striped combinations for database Filegroups. Storage Foundation works with any storage hardware included in the Microsoft Windows HCL catalog.

If you're using a Storage Area Network (SAN), Storage Foundation's cluster and private disk groups simplify the process of moving storage between servers, while protecting the integrity of the data by ensuring that only one system uses the disk group at any time. The Dynamic Multipathing Option (DMP) improves availability and performance by providing multiple paths to storage, with I/O load balancing between redundant paths. DMP supports hardware on the Hardware Compatibility List for VERITAS Storage Solutions *for Windows* Products, which includes a wide variety of popular storage arrays.

To handle growing storage needs, you can configure Storage Foundation to resize volumes as needed, based on predefined thresholds or other policies. This volume resizing can occur while data is online and available.

Finally, Storage Foundation offers a single interface for managing data across multiple servers and storage platforms, and lays the foundation for functions like quick recovery and volume replication.

Recovering Quickly from Logical Errors

Protecting SQL Server databases from data corruption or data loss is a significant part of the database administrator's job. But traditional tape backups take a long time for recovery, particularly when the problem is a dropped table or user error. VERITAS Storage Foundation *for Windows* with the FlashSnap option offers another option: quick recovery from disk-based images.

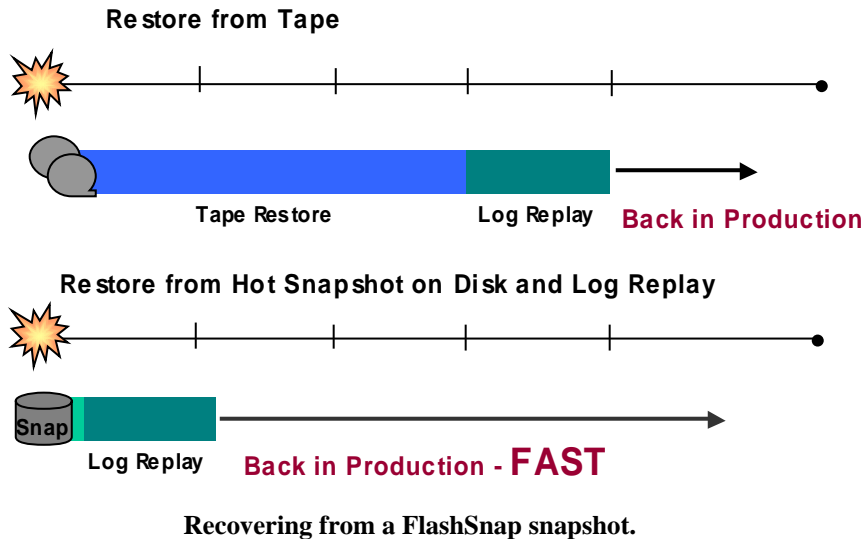
A FlashSnap snapshot is an independently-addressable volume that mirrors the production volume. By splitting the mirror, the FlashSnap option creates a point-in-time image of the data. This image can remain on the production server or be moved to an alternate location.

FlashSnap integrates with SQL Server 2000's Virtual Device Interface (VDI), which coordinates snapshots with a storage system. Using the VDI support, the database server flushes any pending writes and temporarily pauses transactions so the snapshot can get a consistent database image. You can make multiple snapshot volumes simultaneously – creating a point-in-time snapshot backup set of all of the volumes associated with a SQL Server database without taking the database offline.

You can use the snapshot itself as a backup source, offloading the backup process from the production volume or server. Or you can use snapshots as quick recovery images. The FlashSnap snapshots operate on any storage hardware except network-attached storage (NAS).

In a SQL Server environment, restoring from the snapshot is like recovering from a backup image; to reduce data loss you roll forward the transaction log from the point of the image to the point immediately *before* the failure, or any point in between. Storage Foundation automates the process, which simplifies the operation, reduces human error and provides protection for mission-critical data.

Because they have minimal impact on the production database, you can take snapshots more frequently than traditional online backups. The fresher your backup image, the fewer transactions you must replay to restore the data. The result is a much faster recovery from data corruption or logical errors than a tape restore.



Protecting SQL Server from Component/Server Failures with Clustering

To provide end-to-end high availability for applications using SQL Server data, you must protect all components of the application infrastructure. This requires redundant, intercommunicating servers, redundant application components, and automated monitoring and failover capabilities.

Storage Foundation HA *for Windows* includes VERITAS Cluster Server (VCS) technology, which provides highly flexible, scalable failover clustering with workload management capabilities. In a VCS cluster, multiple servers are linked with shared storage and private Ethernet heartbeats. Each system in the cluster shares common storage, although only one system accesses its storage at any time. Application-specific agents monitor the various components of the environment for failures. Most clusters implement SAN-based storage for high performance access to shared storage.

Although you can use SQL Server Enterprise Edition's clustering capabilities, leveraging the Microsoft Cluster Server support in Windows Server, VERITAS Cluster Server provides a highly scalable, flexible clustering solution that can protect and manage *all* application components, from network resources to application servers. VCS offers:

- Support for up to 32 nodes
- Policy-based failover
- Workload management
- Automated DNS updates
- Failover across subnets
- Support for mount points
- No restrictions on which edition of SQL Server you use

VCS implements a policy-based response to failures, attempting a restart and/or automatically moving the application to another resource in the cluster. VCS has a built-in sense of cluster load and can relocate application components based on available capacity augmented with a rules-based policy engine.

You can also cluster web services and other applications which may be vital to the total application delivery. Related, dependent resources can be grouped together so they can be started, stopped and moved appropriately. VERITAS Cluster Server offers application-specific agents to monitor and fail over a wide range of enterprise applications, including application servers and file servers.

Storage Foundation HA also offers a simulator for testing failover so you can test that failover happens as you expect it to, before an actual crisis occurs.

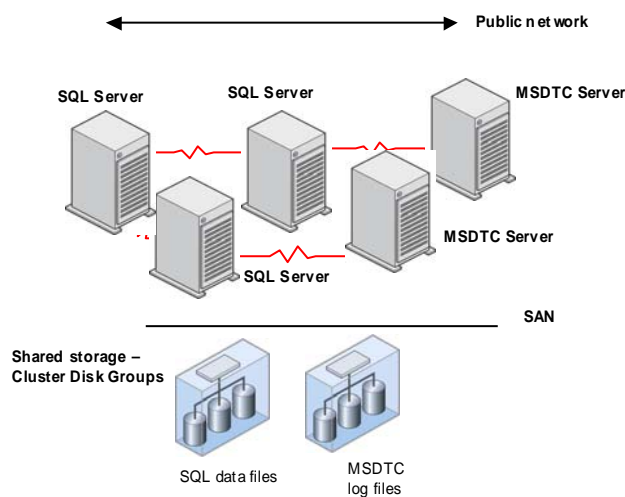
Local clustering with Storage Foundation HA addresses many of the challenges of Microsoft SQL Server administration:

- The software automatically monitors the core SQL Server services and responds appropriately in case of a problem, failing over to other 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 patches.

Local Cluster Configurations

Storage Foundation HA supports either Active/Passive or Active/Active cluster configurations for Microsoft SQL Server.

- In an Active/Passive cluster, a “spare” system serves as the “standby” or passive system for a production server, or multiple servers within a cluster.
- In an Active/Active cluster, multiple nodes may be running SQL Server and other applications, with spare capacity to host each other’s applications if necessary. Using SQL Server 2000, up to 16 instances of SQL Server 2000 may run on a single machine, while VERITAS clustering technology supports up to 32 nodes in a cluster.



An Active/Active cluster

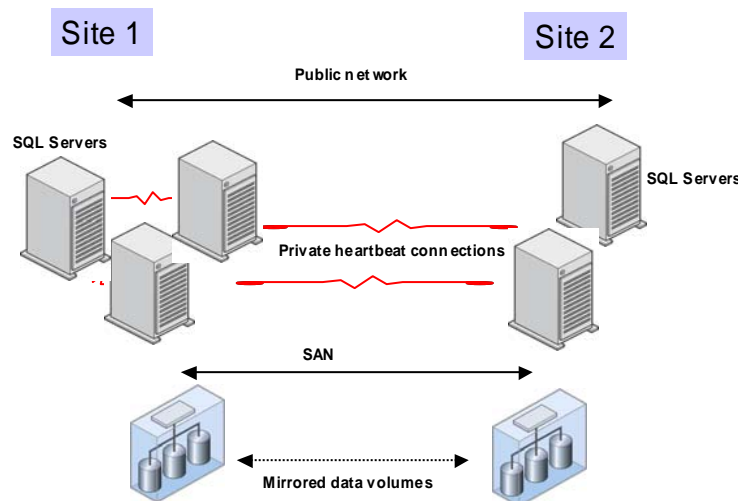
For example, a SQL Server 2000 cluster might include multiple nodes running SQL Server instances, and multiple nodes configured to run the Microsoft Distributed Transaction Coordinator. The cluster can include numerous other servers running different applications. If an entire server fails, its applications can be restarted on other servers in the cluster. VCS clusters can support up to 32 nodes.

Stretch or Metropolitan Clustering

Many large organizations enhance recoverability by clustering across different facilities that are nearby but with some geographic separation. These *stretch*, *campus* 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 the nodes that must maintain heartbeat communications. For Storage Foundation HA, the latency in the node interconnect is the key factor. VERITAS recommends that network latency not exceed 500 milliseconds between nodes. Any further, and the better solution is a global “cluster of clusters” communicating over wide area networks.

A basic stretch cluster mirrors data between two locations, using a SAN infrastructure spanning the two locations. Storage Foundation HA mirrors the SQL Server volumes automatically between the two locations, without requiring identical hardware in each location. It’s also possible to use hardware-based mirroring.

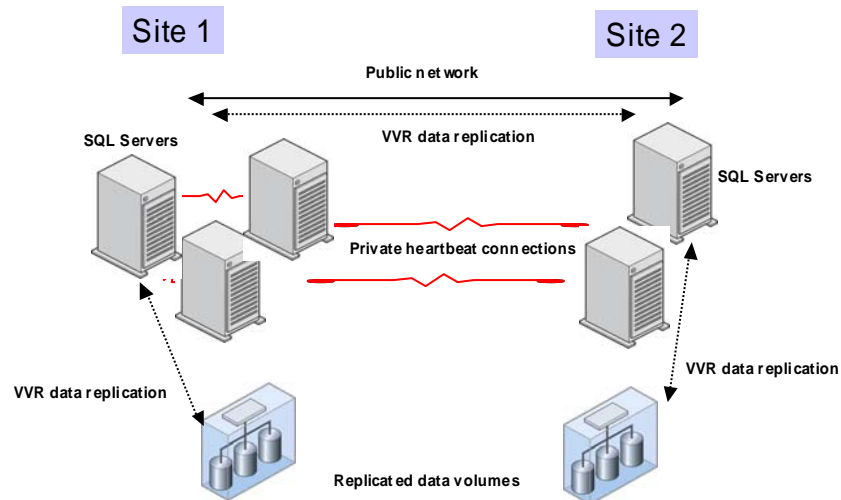


Storage Foundation mirrors data between the nodes of this stretch cluster.

This configuration works with Storage Foundation’s volume-based mirroring or third party array provides that offer data mirroring. The software-based mirroring approach does not require identical storage hardware on each side. The network infrastructure must support communications between servers and mirroring between volumes. This solution requires no additional software licensed beyond Storage Foundation HA for Windows.

Replicated Data Clusters

If you don't have the SAN infrastructure already in place, it's also possible to create a *replicated data* cluster that replicates data over the IP network between the sites. The nodes still remain in the same failover cluster, with private heartbeat communications. You can use the Storage Foundation Volume Replicator option (described below) or a third-party array to replicate the data.



Data is replicated over IP networks between different sites in the same VCS cluster

Particularly for organizations with primary data centers in lower-risk areas, maintaining a secondary facility in close proximity to the primary site is an easy and cost-effective approach to disaster recovery. Most disasters are localized, so these "stretch cluster" configurations offer rapid disaster recovery for a wide range of potential problems. However, if you need a greater geographic separation, then a wide area clustering solution with replicated data is preferable.

Replicating SQL Server Data

Data replication provides a higher level of data protection, and faster recovery from site-wide outages, than restoring a backup at an alternate location.

There are multiple approaches to replicating SQL Server data: you can choose to perform replication at the database level using SQL Server utilities, in the storage hardware itself, or at the logical volume level using Storage Foundation *for Windows*.

Database-specific replication

SQL Server 2000 offers a number of different approaches to database replication. Using snapshot replication, a "publisher" distributes a database snapshot to subscribing instances. This is time-consuming, but useful for databases with little changes. Transactional replication monitors changes to the published (primary) database and sends them to secondary databases, either immediately or periodically, and merge replication allows multiple instances to update data and use predefined conflict resolution rules to resolve conflicts on merge.

These database-specific replication approaches have limitations for disaster recovery purposes. For example, database replication does not replicate other files necessary for the database such as configuration files, extended stored procedures stored in DLL files, etc. It remains up to the DBA to maintain strict version control between sites.

Hardware-based replication

Several storage vendors offer hardware-based replication with their arrays. While this approach can provide up-to-date replicated data, it has some limitations. Most notably, you must run the same storage hardware on each side of the replication link. Since this feature tends to be offered in high-end arrays, this can add to the expense of maintaining a secondary site.

Volume-based replication. In contrast, the Volume Replicator option of the Storage Foundation *for Windows* is both hardware- and application-independent, replicating changes at the volume level. The Volume Replicator Option offers the ability to consistently and reliably replicate data across IP networks using any available storage.

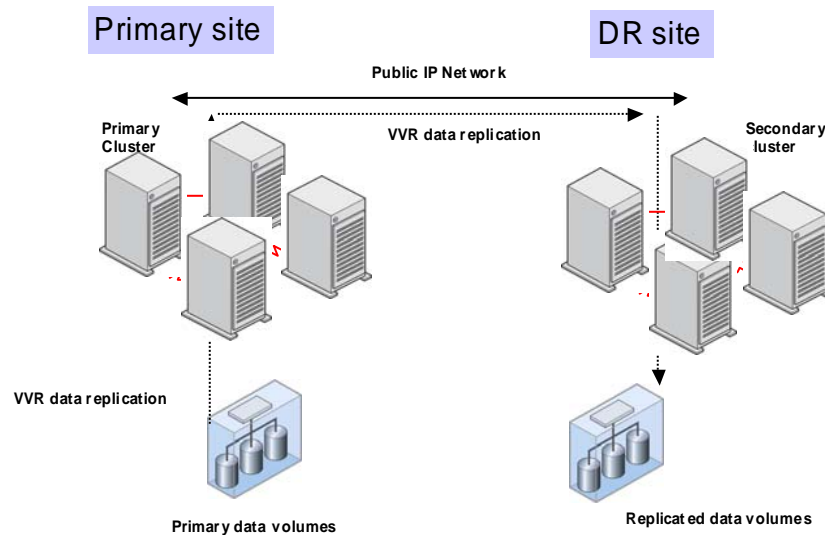
Storage Foundation Volume Replicator Option replicates the contents of each volume across a wide area network to the secondary site. It is completely transparent to the application components. Unlike traditional block-based approaches, VERITAS Volume Replicator replicates I/O instead of disk tracks so the data is always replicated in a consistent fashion – guaranteeing the recoverability of the SQL Server environment.

You can replicate all of the logical volumes required by the database, including the external file system files that are essential to the overall database. The database administrator does not have to spend any time or effort patching and updating the replicated site, as all changes are copied automatically to the secondary site.

With data replication in place, you can provide the highest levels of availability by creating a secondary site, geographically-removed from the primary site, for nearly instant failover in the case of a disaster.

Disaster Recovery with Global Clustering

Using Storage Foundation HA with the Global Cluster and Volume Replicator options, you can switch application services between widely separated sites with a single mouse click.



A global clustering solution using Volume Replicator and Global Cluster options

Global Cluster option creates a "cluster of clusters." Should an application fail, the clustering software will first try to restart it within its local cluster. If the local cluster has failed, then Storage Foundation notifies the administrator, who can verify the problem and migrate the application with the touch of a single button. Failover and replication can occur over any distance.

The Global Cluster option manages either VERITAS volume replication or third-party, hardware-based replication.

Using the Global Cluster option, you can manage multiple VERITAS Cluster Server clusters (regardless of platform) from a single console, and switchover or failover a single application, multiple clusters, or an entire site between geographically-distributed locations.

The VERITAS solution lets you create or manage secondary, "DR" sites with minimal impact on the IT budget:

- **Server savings** The secondary site doesn't need identical computing resources as 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 any available storage, as long as the overall capacity is sufficient.
- **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.

Site migration with a global cluster

To understand benefits of global clustering, compare site migration with global clustering with today's disaster recovery processes.

Using Storage Foundation HA with the Volume Replicator and Global Cluster options, you can create a DR data center that runs other processes, such as testing or development. In the event of a local disaster at the primary site, VERITAS Cluster Server detects the problem and alerts an administrator. Or, an administrator may manually start the failover if they have even minutes of advanced warning, such as the time available on an Uninterruptible Power Supply. In that case, VCS would shut down the applications at the primary site.

VCS then automatically performs the following steps:

1. Stop replication.
2. Reconfigure the network, migrating the virtual IP address for the application to the DR site's systems.
3. Promote the replicated data to primary status. If a secondary site is still available, Volume Replicator can replicate to another site to maintain data availability.
4. Restart the Cluster Server service groups, bringing SQL Server and related applications back online at the DR site.

SQL Server can be running at the DR site in a matter of minutes, with a single click; users would access the new site without being aware of the move. Furthermore, migration back to the primary site, once the disaster has passed, is just as simple.

Compare this to the processes companies manage today. First, companies need to 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, 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, while being subject to possible human error in the time of a complete site outage.

The VERITAS solution eliminates the 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 risk while improving application availability.

A "Fire Drill" feature lets you test the failover between sites without actually stopping the application at the primary site or disrupting use access. The fire drill process brings the application online at the secondary, DR site with a snapshot of production data, without interacting with outside clients or other resource instances. This enables organizations to test the disaster recovery solution without disrupting service levels.

Summary

To deliver end-to-end availability for SQL Server-based applications, you need an integrated, comprehensive way to manage availability of SQL Server data, the database instance, application components, and the entire data center.

VERITAS Storage Foundation HA *for Windows* offers a single, integrated solution that protects SQL Server availability from a wide range of threats, ranging from disk failures to data corruption to site-wide disaster. Once the solution is installed and configured, you can add options by enabling license keys. This integrated architecture lets you phase in different levels of protection and availability while paying only for what's needed today. And because the VERITAS software supports heterogeneous storage platforms, you can leave your hardware options open, leveraging storage you already have or taking advantage of new arrays as needed.

The VERITAS solution is Microsoft-certified and well integrated in the Windows environment. Microsoft and VERITAS have worked together for many years to improve storage manageability on the Windows platform. Microsoft selected VERITAS to develop the disk management software for both Windows 2000 and Windows Server 2003. The Storage Foundation solution builds on the dynamic disk and dynamic volume capabilities now native to the Windows platform.

Whether you're most concerned with data corruption or business continuity, VERITAS Storage Foundation HA *for Windows* provides a packaged, integrated solution tailored to the Microsoft SQL Server environment.

VERITAS Software Corporation

Corporate Headquarters
350 Ellis Street
Mountain View, CA 94043
650-527-8000 or 866-837-4827

For additional information about VERITAS Software, its products, or the location of an office near you, please call our corporate headquarters or visit our Web site at www.veritas.com.