



**Your company relies on its databases.
How are you protecting them?**

Protecting Microsoft® SQL Server™

White Paper
Published: May 2005

Microsoft®
GOLD CERTIFIED
Partner



Executive Summary

Database Management Systems (DBMS) are the hidden engines behind some of your most critical information applications including:

- Enterprise Resource Management (ERP)
- Customer Relationship Management (CRM)
- Accounting
- Sales force automation
- Order tracking
- Inventory management
- Technical support and customer service help desks

In most cases, all of the data for these business systems is stored in a DBMS. If the DBMS server fails, not only is the application unavailable, but all the current and historical data is at risk. Protecting this data asset is critical to the success of an IT organization.

One of the most popular DBMS products in the world is Microsoft SQL Server™ (MS-SQL). Like most database management systems, MS-SQL stores all of the data in a handful of database containers, or files. If one of these containers is damaged or corrupted, all of the data it contains is lost. ***The end result - a serious failure of one system could impair your ability to do business for hours, days, even permanently.***

This white paper discusses data protection strategies for Microsoft SQL Server and answers the following questions:

- What does it really mean to protect a database?
- What are your options for database protection solutions?
- How do I compare the costs and advantages of different solutions?
- Why should you consider NSI Software?

What does it really mean to protect a DBMS?

There are two phases to a DBMS protection strategy:

1. Data Protection - Ensure you have a second copy of the data stored outside of the production system
2. Data Availability - Prepare the second copy of the data, DBMS and associated applications to bring on-line in case the production system fails

DBMS protection can be easy or difficult, depending on your existing technology, procedures, and objectives.

For example, most MS-SQL administrators already back-up their databases, typically to tape. For most, it may appear that this satisfies phase one, Data Protection, however, with further discussion it is likely that this will not meet their particular business needs. Most DBMS backup strategies involve a nightly full backup of the data. However, this method only protects yesterday's work, any data that has been entered into the database since the nightly backup probably does not exist anywhere else and would have to be manually recreated if the production system failed.

To determine if a particular data protection strategy is appropriate you first need to understand the key recovery metrics and then define your recovery objectives using these metrics.

The first key metric is Recovery Point Objective (RPO).

RPO defines how much data will be lost in the recovery procedure. If a single disk drive in a RAID 5 array fails, you would lose no data - this is an RPO of zero. If the entire disk array fails and you have to restore from tape, you will lose all the data added to the database since the last backup. This is an RPO of 12-24 hours. This also assumes that your full and all incremental backups are usable.

The second critical metric is Recovery Time Objective (RTO). RTO measures how long the entire recovery process takes before users can reconnect to their applications and continue working. If a virus corrupts a database, and you simply need to restore it from tape, you may be able to accomplish this in an hour or two, depending on how large the database files are and how fast the all of the associated backup media can be located and restored. This is an RTO measured in hours or days, during which users will be unable to use the application. If an entire production server fails and must be rebuilt, the RTO may be greater than 24 hours, depending on how long it takes to repair or locate new hardware, reinstall the operating system and applications, retrieve the tapes, and restore the data.

According to Forrester Research, most large, online businesses average one to five hours of downtime every month, and lose \$8,000 or more per hr.

Once you have quantified your recovery objectives, you are likely to discover that traditional tape backup will not be good enough to achieve your RTO and RPO goals for your critical DBMS applications. Especially when you consider that, you need to protect against multiple types of failure scenarios, including individual server or device failure and site-wide disaster.

- **Individual server failure scenario.** Even if the repair is simple, the RTO from tapes, especially if they are stored off-site, may be unacceptably long. In addition, the RPO in this scenario is likely to be very poor, as you will lose all data since the last successful backup. If your business cannot tolerate the down time or the data loss, then you need a High Availability (HA) solution. HA solutions are designed to recover from a single server or device failure with very low RTO measured in minutes and RPO near zero. This solution ensures that data and applications remain seamlessly available.
- **Site-wide disaster scenario.** A disaster may come in the form of a natural disaster, power outage, industrial accident, or even peripheral emergencies that may not directly affect your facility, but prevent access to it. Disaster Recovery (DR) solutions should provide minimal data loss and quick recovery of the applications at an off-site location. DR solutions are especially critical if remote offices or external business partners rely on access to the production information systems.

What are the alternatives when low RPO and RTO are required?

When the RPO and RTO must be kept low, hardware-based synchronous replication is a potential data protection alternative. Synchronous replication capabilities are available by connecting two SAN disk subsystems via dedicated fiber optic cable. Synchronous replication or mirroring can provide zero data loss (RPO), and with proper configuration, may provide an RTO measured in minutes. However, these hardware-based solutions are very expensive to purchase and operate, and have distance limitations of about 10 miles between the two SANs. The distance limitation means these solutions are typically inadequate for Disaster Recovery since both locations could be impacted by the same crisis. Most business data does not require this level of protection; absolutely zero data loss. While some may, generally this level of RPO cannot be cost-justified.

“Less than 5% of the data stored in the enterprise requires synchronous protection”
– Gartner Group

Replication for the Masses

NSI® Software Double-Take® provides patented asynchronous replication technology that fills the vast middle ground between tape backup and synchronous mirroring solutions. Double-Take solutions can be configured to provide RTO and RPO measured in minutes, and often in seconds, with low purchase and operational costs. Without a distance limitation, Double-Take is equally capable of providing High Availability and Disaster Recovery solutions. In figure 1, you can see that Double-Take provides dramatically improved protection with only a moderate additional cost compared to tape solutions.

A Yankee Group and Sunbelt Software survey of 362 IT executives in March 2004 found **42 percent of respondents had been unable to recover data from tape in the last year as a result of tape unreliability.**




Figure 1 – RPO/RTO Relative to Cost

What about MS-SQL's built-in replication capabilities?

MS-SQL's built-in replication feature is designed to publish periodically updated, read-only copies of selected data. It is not designed to provide data protection or application availability to the entire database.

The built-in disaster recovery capability provided by MS-SQL is called 'log shipping'. This technique assumes that a recent copy of the database exists, possibly on tape, and can be restored on another server to create the baseline copy of the database. Once the baseline is established, Microsoft SQL Server periodically sends a log of recent transactions. These transactions can be applied to the copy of the database to bring it up-to-date. This updated copy of the database can then be used to replace the failed production system if needed.

Typical challenges with log shipping include:

- Potentially significant processing overhead during production hours
- Logs cannot be shipped more often than once per minute, and are typically shipped less often - increasing the amount of lost data after an outage (RPO)
- Un-logged transactions are not included
- Advanced administrative skill is required
- No automatic failover mechanism
- Tracks logged user data changes only, not schema, security, or bulk inserts
- No log copy or bandwidth management functionality

In contrast, Double-Take protects user data and automatically includes all changes to the entire MS-SQL instance, including:

- Complete transactional integrity for database state
- Schema updates
- Security configuration
- Bulk transactions
- Real-time replication to the replica
- Automatic, transparent failover
- ***Plus, no changes to the MS-SQL data schema are required to get complete protection using Double-Take***

How NSI Software technologies deliver MS-SQL HA and DR solutions

Complete protection of MS-SQL Server requires a High Availability (HA) solution and a Disaster Recovery (DR) solution. Each solution must have a defined RTO and RPO, and the solution architect must balance the cost of achieving those values against the cost of downtime and lost data. This section examines several scenarios, and identifies the type of NSI solution for each.


Scenario 1 - High Availability using Double-Take

Double-Take is designed for "many-to-one" replication between Windows server platforms across an IP network. A single target server on the same LAN can provide High Availability for one large production server, or for several smaller production servers. With the continuous replication and built-in failover capabilities within Double-Take, it is possible for the target server to stand in for a failed server(s) within seconds of a server failure. The target server will appear on the network with the same identity, same MS-SQL configuration, and with data that is within a few seconds of the failed server. This solution provides users with near-instant access to their data often without ever knowing there was a failure.

"A rolling blackout or earthquake could cost us \$145,000 a day in lost productivity if our people do not have access to a mission-critical application on our server. With NSI Software's Double-Take, even disaster can't keep us down for more than 5 minutes; access is uninterrupted."

**Robb Good, Vice President,
Director of Information
Systems, Sundt
Construction, Inc.**

Scenario #1 - High Availability Using Double-Take




Scenario 2 - High Availability using GeoCluster®

Traditional hardware clusters use multiple nodes and a shared disk subsystem to provide excellent application availability, but they still have a single point of failure in the disk storage sub-system. Should this disk subsystem fail, the entire cluster is rendered unavailable. NSI Software created GeoCluster® to enable a Microsoft Cluster Service (MSCS) cluster to have separate redundant storage - providing both application availability and data protection. GeoCluster utilizes the patented Double-Take replication engine to ensure each node maintains an exact copy of the active data. Each node in the cluster can now have its own exact copy of the data on its own storage device. Should the active node fail, MSCS will automatically fail over the failed server to one of the other cluster nodes that has maintained an exact replica of the failed disks data. GeoCluster also provides the ability to configure Microsoft clusters using any non-similar storage devices. One node could be SAN-attached while connecting the other node to a SCSI disk array. This reduces overall costs of deploying Microsoft clusters and simplifies management.

"With GeoCluster, we are saving thousands over the alternate solutions."

**Steve Wagner,
Manager, Abbey Press
Information Systems
and Services**


Scenario 2 – High Availability Using GeoCluster



Scenario 3 - Disaster Recovery using Double-Take

Placing a server in a remote location is the most fundamental requirement of a Disaster Recovery solution. DR solutions based on Double-Take can protect against an entire datacenter or regional failure by allowing data to be copied many miles away. By using Double-Take with existing facilities and WAN links, you can implement and operate a high quality DR solution at a very low cost. This DR solution has no distance limitations, so the recovery center can be placed as far away as necessary to avoid common risks.


Scenario 3 – Disaster Recovery Using Double-Take



Scenario 4 - High Availability and Disaster Recovery using Double-Take

Double-Take supports combinations of "many-to-one" and chained configurations to deliver a combined HA and DR solution. The HA server at Site A provides low RPO and RTO protection against an individual server failure, while the DR solution provides remote replication to one or more servers at a remote facility that is accessible via the WAN, and provides protection against an entire site failure.

Scenario 4 - High Availability and Disaster Recovery using Double-Take



Scenario 5 - High Availability and Disaster Recovery using GeoCluster and Double-Take

NSI Double-Take provides the same level of disaster recovery protection for Microsoft clusters as it does for standalone servers. Although a cluster is typically protected from an individual server or disk failure, it is not protected from a site failure and should still be considered a candidate for remote replication.

Double-Take is designed to work within Microsoft cluster environments to ensure appropriate data from active resources is replicated to an offsite DR location, regardless of which cluster node controls the resource. With GeoCluster ensuring that your Microsoft cluster is well protected from a disk or server failure,

Double-Take ensures it is protected from a complete cluster or site failure.

How Double-Take can meet other data protection, availability, migration, and distribution needs

In addition to protecting your MS-SQL data, Double-Take can provide powerful solutions for a number of data protection, data availability, data migration, and data distribution challenges.

Scenario 6 - NAS to SAN storage migration. As more environments move from direct attached storage to NAS or SAN, the question of how the actual migration will occur becomes more frequent. Using the same techniques outlined above for data protection, one can migrate from local storage to a NAS, a server utilizing a SAN, or even a NAS-gateway to the SAN. In all cases, the fundamental requirement is that the data is moved from one Windows platform (with local storage) to another Windows platform (with more manageable storage), like that of a Windows-Powered NAS (or Storage Server). A huge advantage of using Double-Take for migrations is that the data is continuously updated on the new server/storage, allowing users to be cut over at any time. It is not necessary to stop all user access in order to move the data.

Scenario 7 - Branch office server to Centralized Data Center. Enterprise branch offices tend to be isolated because of WAN performance, hardware and management costs. Many times this requires non-IT personnel to be responsible for tape rotations and cleanings. The result is higher labor costs and lower restore reliability. By efficiently replicating the byte-level changes within the data at the branch offices using Double-Take, one can bring the branch's data back to a centralized data center. This provides disaster recovery for the branches, and allows backups to be managed at the centralized facility by trained IT personnel using more advanced tape technologies. Backup processes can also run 24 hours a day without impact on the production servers.

Scenario 8 - Small office Server-to-Server Protection. Large enterprises may have multiple data centers and a myriad of server technologies, but the typical small office relies heavily on few servers, often with limited IT resources or personnel. When the primary server fails, the office productivity can halt. Double-Take provides a simple and cost-effective way to "fail over" to a second machine, either in the same office or perhaps even at an employee's home. The result is rapid availability of the data - and the small office continues doing business.

"Double-Take was the optimum solution that integrated smoothly with our core application. In fact, Double-Take's flexibility made it the only option."

**Christina Surmenean,
Senior Vice President and
CIO for CNA Trust**

Scenario 9 - Consolidate Backup Operations. Today's corporations are increasing their business day, as geographic and national boundaries no longer limit effective commerce. This results in an ever-shrinking backup window. However, the redundant copies of files on Double-Take target servers can be backed up at any time, even when the original copy of data is in use. Without expensive and application-specific backup agents, the second copy of the data is archived to tape using existing tape technology attached to the redundant server. In addition, the backup occurs at local disk/tape speeds, instead of a media server scheduling the backup up multiple application servers. This increases the backup window to a full 24 hours and ensures there are not missed or inconsistent files because they are actively open and updated.

Scenario 10 - Data Distribution. Data protection and availability are not the only uses of a replication solution. Like the data migration solutions discussed above, some businesses require that data is sent to an alternate location for local access. Double-Take can provide a corporation with a master-content server and then ensure that all regional locations and branch offices receive replicated data for reporting and analysis applications.

Why NSI Software?

For organizations that depend on information, NSI Software, Inc. (NSI®) provides the worlds most relied upon solutions for continuous data protection and application availability. Our flagship product, Double-Take is the most technically advanced software in the market and is the solution of choice for thousands of customers, from SMEs to the Fortune 500.

NSI has been protecting Microsoft SQL Server databases since 1996. We have deployed production solutions on MS-SQL 6.5, 7.0, and 2000, for leading healthcare, legal, financial, energy and government organizations worldwide. With over 15,000 licenses shipped to protect SQL Server, we are by far the largest supplier of replication software for this solution in the industry.

NSI has strategic technical and marketing relationships with industry leaders including Dell®, Hewlett-Packard®, IBM®, Microsoft® and SunGard®, and works with its OEM partners an reseller channels to deliver comprehensive solutions and support to business users worldwide. Founded in 1991, NSI is headquartered in Hoboken, NJ.

For more information, please contact NSI Software at www.nsisoftware.com.



NSI Software, Inc. - Corporate Office

Two Hudson Plaza, Suite 700
Hoboken, NJ 07030
800-775-4674 or 201-656-2121
Fax: 201-656-2727

NSI Software, Inc. – Inside Sales

8470 Allison Pointe Blvd. Suite 300
Indianapolis, IN 46250
800-674-9495
Fax: 317-598-0187

© 2005 NSI Software, Inc. All rights reserved.

Double-Take®, NSI® and GeoCluster® are registered trademarks of NSI Software, Inc. Balance™ is a trademark of NSI Software, Inc. and all are used with permission of the trademark owner. All other trademarks are properties of their respective companies.

Microsoft, Windows Powered, Windows, Exchange, and Microsoft SQL Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Although we try to provide quality information, NSI makes no claims, promises or guarantees about the accuracy, completeness, or adequacy of the information contained in this document. Companies, names and data used in examples herein are hypothetical and/or fictitious unless otherwise stated.