



Better Backups through Replication



2 Hudson Place – suite 700
Hoboken, NJ 07030

800-674-9495
www.nsisoftware.com

Powered by  **Double-Take**

When many people consider replication software, their assumption is “High Availability” or “Disaster Recovery”. And while traditional Business Continuance is an important focus, the routine and automated protection of data will always include a tape aspect – for long term archival of data. The industry is learning that the two technologies are not mutually exclusive, and in fact complimentary.

There are several aspects of combining tape and replication technologies that will be explored in this paper:

- The typical backup / restore solution
- What are the issues with tape alone?
- How Replication adds to an existing Tape Solution
- What about key O/S information, like the System State?
- What about application-specific data, configurations and large “scrubs”?
- How about snapshots, like VSS from Windows 2003?
- How does Restore work?

The purpose of this document is to educate a “Technical Decision Maker” on how replication can add value and reliability to one’s existing data protection strategy. For more information on in-depth “how-to” issues related to tape backup or Double-Take®, please visit www.nsisoftware.com/backup.

Strategic Planning Assumption

By 2003, 75 percent of large enterprises will combine data replication and tape technology for rapid application recovery (0.7 probability).

Bottom Line : Organizations that are running 24x7 operations and are confronting shrinking backup windows, or that foresee an inability to meet service-level agreements, should plan and budget for deploying data replication technologies.

Gartner Designing to Restore From Disk: Backup Futures

The typical tape backup / restore solution

The primary flaw with tape backup is that it only occurs (for most environments), once per day. This means that one must always measure data-loss windows and recovery times in “days”.

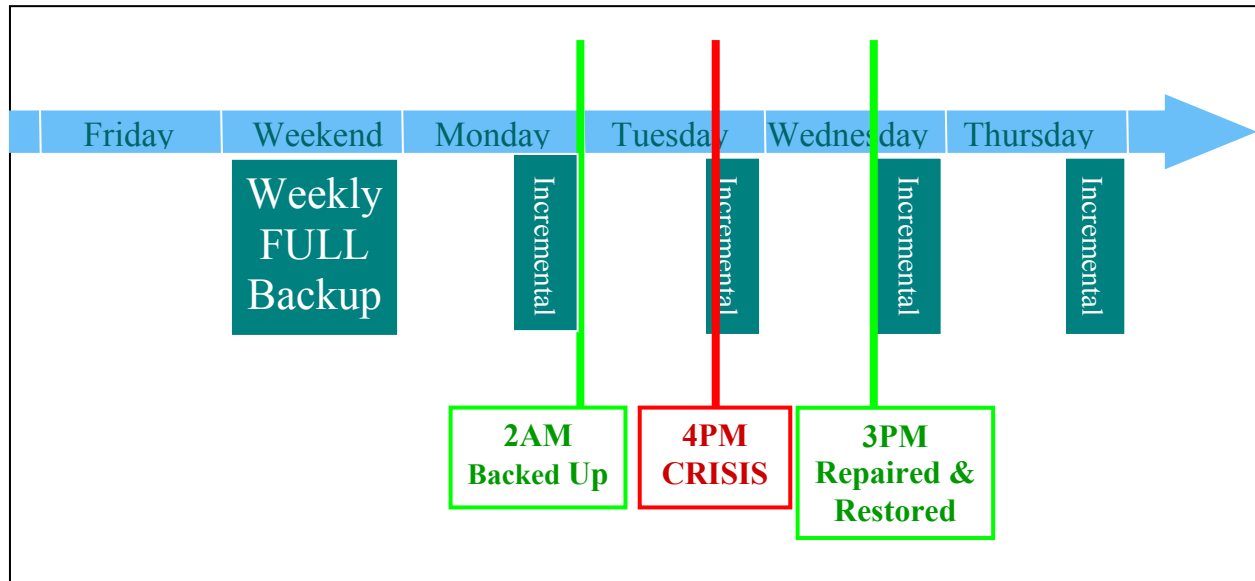


Figure 1 – Typical data loss and productivity loss with tape

As an example, consider if a primary server failed at 4 PM on Tuesday afternoon. If all parts were on hand, the server could be rebuilt on Tuesday evening and the data restored. When users return on Wednesday morning, their data will be as it was on Monday night's backup. Tuesday's data was lost.

But for most of us, an extra server isn't sitting on the shelf. So, the components could be expedited on Wednesday morning and the restoration began in the afternoon, users would begin working by Thursday morning - with data as it existed on Monday night's backup. Tuesday's data would be lost. And Wednesday would have had limited productivity at best.

If off-site tape couriers are used, Monday's tape might have been off-site, which adds an additional day before the restore could have begun. To imagine the story even worse, consider if the environment does a full backup only on weekends (and incrementals during the week). If something were wrong with that weekend's backup, the data loss and restore effort both reach back through the entire previous week.

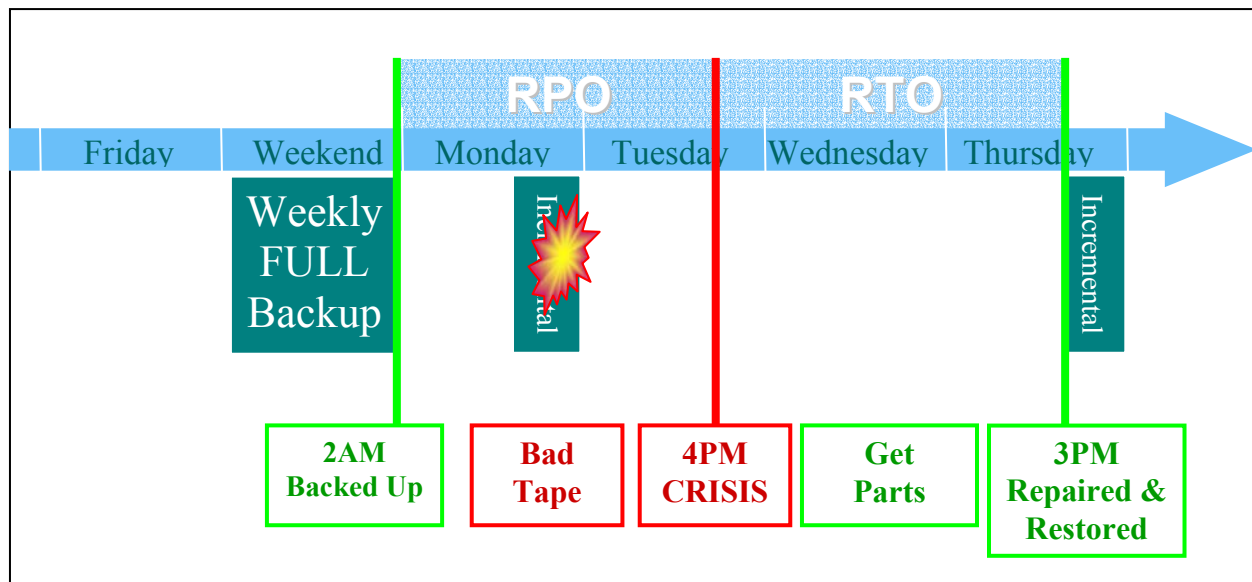


Figure 2 – Realistic view of potential data loss and productivity loss with tape

The only way to reduce the window of exposure for data-loss (RPO) and lost-productivity (RTO) is to use a technology that occurs more than once per day – replication.

What are the issues with Tape alone?

When considering conventional backups there are several caveats that collectively force most companies to having the dreaded “Backup Window”. Backup windows are in place because of a few problems with most tape approaches:

High CPU & I/O during backups – While a server is being backed up, it tends to use excessively more CPU and I/O – which results in overall slower performance. This is a common reason why backups are done during “off-hours”.

High Network Usage during backups – A majority of environments no longer have tape devices on every server. The benefit is better use of tapes and a more manageable software solution. Unfortunately, this does require that all files go across a network (from the production servers to the backup platform). In worst cases, these files traverse the production network segment – slowing down user traffic. In better cases, there is a “backbone segment” for the backup traffic – but the Disk and Network I/O on the production servers will still suffer. To reduce the impact, most backups are still done after hours.

Open Files – By far, the most common reason for performing backup’s during non-production time is open files – those files held locked by applications (e.g. SQL or Exchange), as well as by users’ (e.g. Office documents and spreadsheets). To combat this, one can deploy application-specific agents or open-file handlers. But application

agents are expensive, specific to a version of a particular application, and are only available for a handful of common back-office applications. And all agents and handlers still require the production CPU to do extra work to circumvent the file-locks, which inherently takes cycles away from production tasks.

Collectively, to eliminate some open files, reduce CPU and I/O impact and minimize network traffic – many of us have been forced to back up our data only between 2AM-6AM (or some other obscure window).

Replication resolves these issues (and eliminates the “backup window”) by providing a second copy of the data to be backed up:

- The second copy of the data is not locked, so there are no open file issues.
- CPU and I/O issues are irrelevant on the redundant server, because the users are running from the production platform and are therefore unaffected.
- And because replication provides a second real-time copy (by continually sending small updates), the backup platform already has access to the files without reaching across the network to the production servers.

How Replication adds to an existing Tape Solution

NSI ® Double-Take ® replicates data at a byte-level, meaning that if an application writes a string of 12-bytes within a file; then the actual 12-bytes plus header is transmitted across the existing IP infrastructure. At the target location, the 12-byte string is applied to the same location within the second copy of the file. This provides for multiple copies of one’s data for fault tolerance.

Because the data files on the target server(s) are only loosely coupled to the production copy, the target data is not locked or in use; even when the production copy is. This allows our customers to utilize any backup software / hardware solution on the target data volumes, without open-file or application-oriented agents. This also allows for backup jobs to run during the day, even when the primary copy of the files (on the production servers) is in use. Removing the consideration for open files, NSI customers are able to back up the redundant copy of data 7x24, eliminating the backup window.

D – 2 – D – 2 – T ... Disk to Disk to Tape

To further enhance data protection strategies, a growing number of enterprises are bringing data from their remote sites to their corporate site, before performing tape backups.

Instead of relying on weekly/daily rotations of tapes at remote offices, where human error can affect the reliability of tape handling (and increased manpower cost), tape backups can occur at the data centers. This is enabled by continuously replicating the data changes from the remote sites to the data center with NSI Double-Take.

For remote sites without a local admin staff, file data can be replicated to an upstream hub site using Double-Take and backed up as part of the normal hub site backup process. By centralizing file backup, NSI clients are able to increase the reliability of tapes, reduce manpower costs, and eliminate hardware and backup software in the remote sites.

This architecture can also be complimented by the use of snapshot technologies like those by PSM (by Columbia Data Systems) or VSS (new from Microsoft® – described later).

The broader view on Business Continuity

To achieve larger business continuity goals, NSI Software's Double-Take software is used to replicate data between facilities.

- Data is replicated between hub data-centers to allow for disaster recovery between geographies.
- Data is replicated from remote branches to the nearest data center (and optionally replicated again to the alternate data center) for data protection and resilience of the branches.

Double-Take also provides failover capability, whereby a target platform can be configured to stand-in or “fail-over” for a production server; offering the name, IP and file shares (as needed). This allows remote users to utilize the data-center copy of the data, if the branch server were to fail.

What will have to change for me to implement?

Implementing Double-Take does not require changes to the existing AD user environment. Permissions to production files will also be applied to the replicated copies. To protect the replication environment, Double-Take creates a local machine group on each machine that runs the software. By adding an AD-domain group to this local machine group, authentication for management will be automatically bestowed.

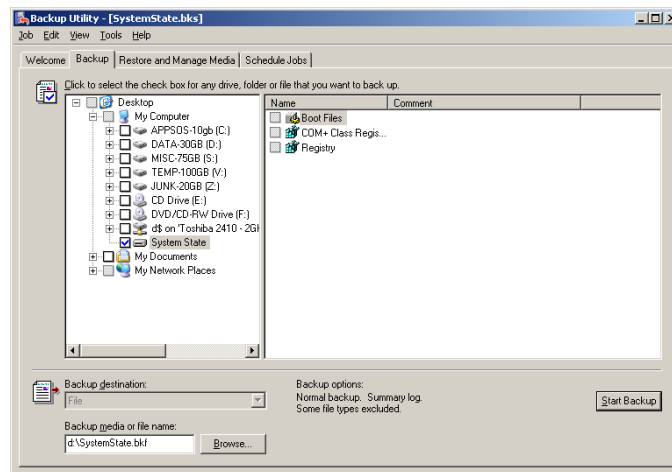
Double-Take does use the existing infrastructure, requiring merely native IP connectivity between sources and targets. Specifically, two defined TCP/UDP ports are used for all Double-Take traffic; thus allowing network management and monitoring, as well as “quality of service” or packet-prioritization to be optionally used.

In addition, Double-Take does not rely on any particular backup technology (although using one that is supported by Microsoft is suggested). Without changing one’s tape backup software or hardware, the backup process can be enhanced – simply by pointing the backup solution at a Double-Take target server, instead of the production server(s).


What about the O/S specific information?

If the primary goal is a predictable and reliable restoration of the production environment, one simple step is to automatically back up the system state of each production server. According to Microsoft, the System State “*comprises the registry, COM+ Class Registration database, files under Windows File Protection, and system boot files*” – and can also include information from Active Directory, DNS, IIS, and the Cluster Service. In short, one can completely restore a failed server by doing a clean OS installation followed by restoring the system state.

Thankfully, Microsoft server OSES provide a backup utility that can be run while users are active on the machine. The default location for this utility is `c:\windows\system32\ntbackup.exe` and it can be initiated from StartMenu / Programs / Accessories / SystemTools / Backup. It can also be executed from the command line or scheduler.



First-time users of the backup utility should consider using the GUI to configure a backup of the system state (plus additional key files, such as INI's within program directories). During configuration of the backup job, one can schedule the job to run routinely - with a best practice being at least weekly.



The results will be an individual file (*.BKF) instead of using actual tape or other media. By selecting the directory (where the backup file will reside) as part of the Double-Take replication set, this BKF file will also be replicated to the target server. During any recovery, the BKF can be used to restore the system state (including registry and other in-use files).

As part of the recovery process, one should configure the new production server with the same O/S. Then, if the various system drive directories (e.g. Windows and Program Files) have been replicated, those can be copied to the new server. If you are using a third-party backup package, one might consider backing up the remote source server's O/S volume (including system state) monthly. This will cause some network congestion, but once per month is typically tolerable. To restore, one would restore from tape, and then still replay the latest System State backup that was replicated via Double-Take to the target server.

In either model, after the "new" server has a functioning O/S and application directory, then the only restoration is the data set (from Double-Take), which will be seconds old. This results in near zero loss of data, including the precious registry information.

For more information about how to protect and restore the registry and other System State components, please visit the Microsoft website.

What about Snapshots ... like VSS in Windows 2003?

One of the most exciting enhancements to data protection beyond tape backup is the built-in feature of Windows Server 2003 called Volume Shadow Copy Service (VSS) which allows administrators to create a point-in-time snapshot of a file server volume. A snapshot can be taken at any time even if files are still open and can be configured automatically at intervals up to every two hours. Each time a snapshot is taken the current contents of a file are frozen and any future changes are tracked and saved to a different part of the disk. This process is transparent to the user but provides them with the ability to restore a file to a previous version on their own without the need to restore from tape, reducing the number of support calls. For more detailed information about the Volume Shadow Copy feature please see

<http://www.microsoft.com/windowsserver2003/docs/SCR.doc>.

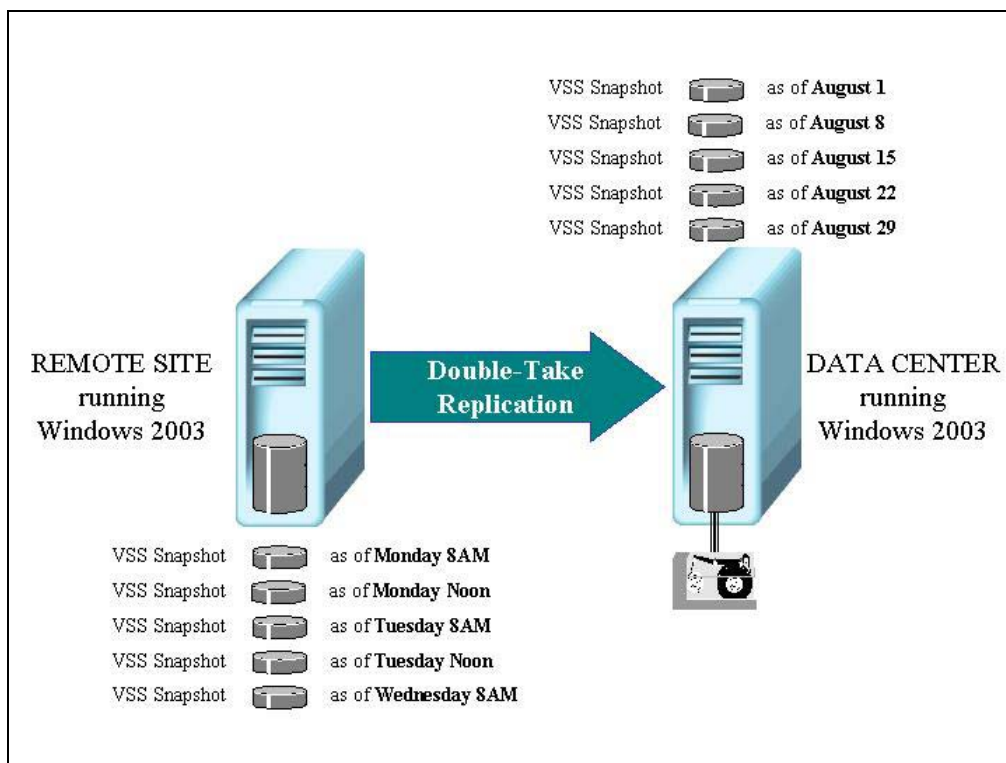


Figure 2 – Data Protection via Snapshots, Replication and Backup

Real Data Protection = Snapshots + Replication + Backup

NSI Double-Take operates in compliment to Windows 2003's volume shadow copy service. VSS can be used in conjunction with the replication technology between servers.

If one snapshots the production (source) server, then users are able to have historical access to older copies of their local data. Transparent to this, the current data will continue replicate from the local server to the remote data center.

If one snapshots the redundant (target) server, then the IT team at the corporate data center will have the same historical access to the data. This is preferable, if storage space is limited at the branch, but multiple copies are desired.

Used together, one might provide 14 daily snapshots within a local branch. This would allow the users to do self-directed restores of data for two complete weeks. In addition, one might do weekly snapshots of the data center copy. This might allow for upwards of 60-90 days of online restorability from the data center (all without ever mounting a tape).

How does Restore work ?

By this point, it should be clear that there are several advantages to using replicated data for “better backups”. But, alas, backing up is simply preparing to restore. So, it is important to understand how restorations work in this solution.

Whole Dataset Recovery – For the scenario where a data volume (or disk set) have been damaged and need to be restored, the Double-Take mirroring and replication processes can be put “in reverse” – pushing the data from the target to the source. One simply repairs/replaces the storage on the production source server. The Double-Take database is aware of where the various target data files came from. So, one can use the Double-Take Restoration Manager to select a set of files and then use Double-Take engine to put the files back where they came from.

The difference between using the replicated files for restoration and last night’s tape is the currency of the restore. Double-Take copy of the files will be seconds away from what the production source had at the moment of failure. Last night’s tape would have lost all the files that had been changed during the entire business day.

Individual File Recovery – For the scenario where the source server has simply lost a few files, there are a few options.

- Double-Take can be configured to “burst” the changes (instead of real-time replication). The result is a copy of the files on the target, which can be minutes to hours behind the source server. This allows a redundant copy to quickly restore from.
- Tape (or Disk snapshots) can be configured to protect the files on the target server – even while the production source files are in use. With this approach, one can go to a snapshot from this morning or a differential tape from two hours ago – and

recover the file (all without impacting the production users). Restoring the errant file directly to the source server (via snapshot UI or backup console) will provide the recovered file to the users – and it will be immediately replicated back to the target, to provide consistency for all copies.

And in all cases, the inherent tools of Double-Take provide for easy restoration. As an example, the Double-Take patented “partial difference mirror” allows large files to be restored by only restoring the partial sections of the file that have actually been changed. The unchanged sections are untouched, which significantly reduces bandwidth requirements and restoration times.

What application-specific data and configurations ?

For those applications that store parts of their configuration information outside of the data set, replication can still be used to provide a “better backup”.

If an application stores its configuration information as flat files (e.g. INI's), then one simply configures the replication set to protect those directories, with optional filters to include the configuration files and/or exclude large binaries.

If an application stores its configuration information in the registry, one can use REGDMP to routinely secure those particular registry hives to a flat file (which would be replicated to the target server, similar to the SystemState information discussed earlier). REGINI would be used to restore those registry hives during a restoration activity.

Some applications do a month-end “scrub” (wholesale changes, compaction, etc). For those environments, three additional benefits of Double-Take come into play.

Extended Queuing – Double-Take provides for a queuing model to cache up to 4TB of byte-level changes, so that even the most dramatic data changes can be propagated to the target server. Since most environments do these types of operations on weekends, a properly configured queue and infrastructure will ensure that both copies are maintained by Monday morning.

Scheduled Verification and Scripting – Some customers choose to temporarily disable replication, when the same large data areas will be repeatedly scrubbed within a short window. Instead, using DTCL (Double-Take Command Language), the real-time replication of Double-Take is turned off while the data is modified. Upon completion of the data compaction, DTCL can be used to re-enable the replication and initiate a “scheduled verify”. Double-Take verify will compare strings within the source and target files, and only send those sections that are determined to be different. This can reduce the bandwidth impact during large repetitive scrub operations.

In-band Command Processing – Many NSI customers wish to do other activities to the target copy of the data, after they are assured that all of the scrubs are complete.

Common examples include a fresh backup and/or invoking a snapshot. To accomplish this, NSI provides the ability to insert a flag immediately behind the replication traffic of some operation. So, an application can perform its scrub and then use DTCL to insert the flag. Upon the target server receiving the flag, one can be assured that the target has also received all of the data from the scrub operation. At that point, a script is invoked – for a backup or snapshot.

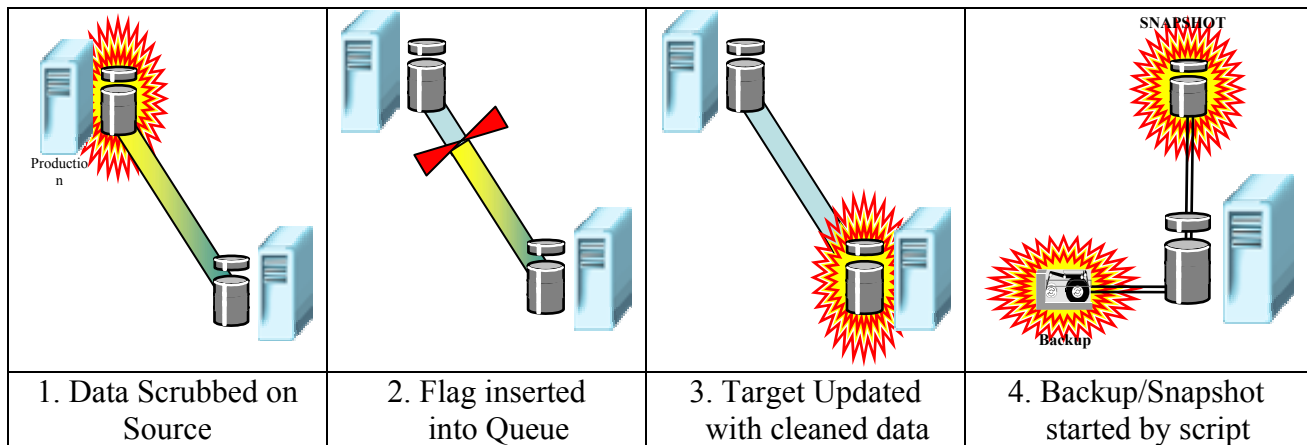


Figure 4 – How In-Band Command Processing works

These benefits are based on using replication to enhance one’s existing (or planned) backup solution. And this is because the target’s data is simply a “Better Copy to Back Up”.

All of these solutions are based around the NSI fundamental philosophy that all business continuity efforts start with protecting the data. From there, it is simply a matter of what you want to do with it. And while we continue to be the leader in Windows® Business Continuity, High Availability and Disaster Recovery – the same software can be leveraged for the more daily protection of Windows file systems by complimenting your existing tape backup solution.

NSI Software knows how to protect applications running on Windows file systems. "Business Continuity through Replication" is the single focus of every person in our company. That focus, and the quality of our products, has helped NSI forge relationships with HP ®, IBM ®, Dell ®, SunGard ®, Microsoft ® and probably your preferred reseller-integrator.



For more information regarding Microsoft's partnership with NSI for business continuity, please visit <http://www.microsoft.com/windowsserversystem/storage/businesscontinuity/default.msp> .

For over 10 years, NSI has been providing the products, services, and support to help you be successful in protecting your most critical applications...

We'd like the chance to prove it to you.

© 2004 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, Volume Shadow Copy, VSS, VSC, Exchange, and SQL Server, are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Our Vision

To be the leading provider of data protection & high availability software solutions for 24x7 business operations

Our Offer to You

We would like to become your partner in ensuring the continuous operations of your business.

Please allow us the opportunity to talk to you about your specific data protection needs and to discuss our products and services that may apply.

Products that Protect your Data

Services that Ensure your Success

For more information on NSI's products and services, please contact NSI.

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



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

Or visit us on the web at WWW.NSISOFTWARE.COM