

THINKING BEYOND TAPE: BETTER BACKUPS WITH REPLICATION

BUSINESS VALUE WHITEPAPER

Double-Take Software, Inc.

Published: March 2008

Abstract

Business-critical data is constantly growing and most IT Managers are responsible for protecting it. But when they consider replication software, their assumption is that they must choose between high availability and disaster recovery. While that can be true, the routine and automated protection of data will always include a tape aspect - for long term archival of data.

In recent years, the industry has learned that tape backup and replication software are not mutually exclusive. In fact, replication technologies can compliment and enhance existing software and hardware choices to provide a better backup.

Introduction

Today most companies understand that the only way to ensure data protection and business continuity in the face of the worst sorts of disasters - floods, tornados, earthquakes, terror attacks, massive power outages - is to establish a remote recovery site at a significant distance from their main and branch offices.

As a result, every night many companies are already backing up their main and branch office systems to tape, and transporting them to a site anywhere from 50 to over 1,000 miles away. What they don't understand is just how vulnerable their data, and therefore their business, remains to these threats, even after such a huge outlay of administrative effort and cost.

This paper explores the high cost, complexity and potentially dangerous shortcomings of a recovery strategy based only on traditional tape backup and demonstrates how an alternative solution - continuous data replication to a remote recovery site over existing WAN connections - provides exponentially better remote disaster protection without adding significant cost or complexity. Finally, it introduces new data acceleration technologies that can optimize the performance of remote recovery solutions and the performance of any other applications running over a WAN infrastructure.

What's Wrong With Tape Backup?

The problems with tape backup are well known; companies have been dealing with them for decades. For starters, tape backup requires a significant investment. Tape hardware and backup software are expensive, as is the labor required to set up and maintain them. Tape cartridges are a continuing cost and completing daily tape backups requires heavy administrative intervention.¹ If you have multiple branch or remote offices (a recent census bureau survey indicates the average company has 50 locations) you have to set up tape equipment and allocate administrative supervision for each. The only other option in these situations is to forego protecting data in branch office locations which becomes a problem all of its own.

Even once the equipment is in place, making backups is inconvenient – bordering on impractical. Tape backup can involve downtime, known as backup windows, since the system being backed up cannot be used during the process. Given the ever-increasing demand for around the clock data access, it gets harder and harder for companies to complete nightly backups within the time window provided. In many cases, it is so hard that the once-nightly backup goal often slips to every other night for many machines. Foregone backups are even a common problem in remote branch offices where backing up is left to non-IT staff.

Most companies don't understand how vulnerable their data and business remain to disaster – even after they've made a huge up-front and ongoing investment in tape-based disaster recovery. An article in SearchSecurity reports that in a survey of 500 IT departments, as many as 20% of routine nightly backups fail to capture all data. Among participants of another survey cited in this article, 40% of IT managers were unable to recover data from a tape when they needed it.² This is a significant concern for corporations that are regulated by industry or government requirements as they can face the risk of being out of compliance if they cannot produce required data when they need it.

Tape backup also places limits on your *recovery point objective* (RPO), the point in time to which you can recover your systems should disaster strike. Periodic tape backup guarantees hours of lost data in the event of a disaster. Suppose, for example, that a critical system fails anytime today; the best you can do is recover to yesterday's data, which will be at least twelve hours old. The later in the day disaster strikes, the older the data from which you'll recover. In addition, recovering from a disaster, any data not backed up is lost for good – unless you recreate it.

The cost of permanently lost data is high and includes the cost of the revenue that the data represents, the business value you can extract from it, and the cost to recreate it. Consider:

- *How much money would your company lose if you lost all your transaction data for the last twelve hours, or even the last ten minutes?*
- *What is the value of the knowledge contained in your company's last twelve hours worth of e-mails and e-mail attachments? What would it cost to have your engineers recreate the last twelve hours worth of original or edited CAD/CAM drawings?*
- *What's your exposure if you can't produce this data in compliance with Sarbanes-Oxley, HIPPA, SEC and other regulations?*

In *The Cost of Lost Data*, a Pepperdine University report updated in 2003 – before the advent of Sarbanes-Oxley – Dr. David Smith estimates the average cost of irrecoverably lost data at more than \$10,000 per megabyte lost.³ But if the data lost is business transaction data or data that's especially expensive to reproduce and key to your company's regulatory compliance, your costs could be much, much higher.

Cost of downtime

When a large-scale disaster strikes, with tape backup you're out of business until you can restore your systems and your data from your tapes. This kind of restoration takes a minimum of several hours, and can easily take days or even weeks.

Gartner Group estimates that the average cost of network downtime for larger corporations is \$42,000 per hour; Contingency Planning Research pegs the average hourly downtime costs for small businesses at roughly \$18,000. But the cost of downtime can be significantly higher depending on the business. In fact, it can be in the hundreds of thousands per hour for health care, consumer products and banking businesses, and in the millions per hour for brokerage, energy, manufacturing and telecommunications companies.⁴

The key to a successful disaster recovery plan is to focus not just on the data (RPO) but also on the applications that end users run to gain access to that data. Recovery Time Objective (RTO) is generally defined as the amount of time it takes to regain access to business-critical data. Solutions like tape backup, which have an RTO of hours or days, don't provide the level of recoverability that most companies require.

1 See Double-Take whitepaper "Reducing Costs and Risks of Branch Office Data Protection"

2 Regan, Keith. "Concerns Raised on Tape Backup Methods." SearchSecurity.com 15 April 2004

3 Smith, David M. "The Cost of Lost Data." Graziadio Business Report Vol. 6 No. 3

4 Meta Group 2000 data

Better Backups with Replication

Data replication has long been considered an impractical solution to the data protection problem. Historically, it required expensive hardware and large investments in bandwidth to protect data in real-time. The evolution of software-based, asynchronous replication has dispelled this long-held belief that continuous data replication isn't feasible - especially for small or medium-sized business with limited resources. And this new breed of data replication offers benefits that more traditional solutions such as tape-based periodic backup cannot:

- Data replication provides a continuously updated copy of critical data at a remote site which minimizes data loss should a recovery be necessary.
- Disk-based recovery is more reliable, less complex and takes less time, improving the RTO of the disaster recovery solution.

Even within the realm of software-based data replication, there are opposing approaches: synchronous and asynchronous replication. It's important to understand the benefits and drawbacks of each.

In synchronous replication, the replication software intercepts data being written to disk and sends it to both the primary and secondary disk arrays at the same time. Only when both arrays confirm receipt of the data does the software accept another write. Asynchronous replication can deliver recovery point objectives (RPOs) measured in minutes, and recovery times measured in seconds.

With synchronous replication, data loss approaches zero because both the primary and secondary disk arrays must contain the same data. But the confirmations required for each data write can cause performance problems, especially in applications that process lots of transactions. Acceptable performance often requires connecting the arrays with high-bandwidth fibre channel, which is very expensive and which has an effective range of about ten miles. As a result, synchronous replication is not ideal for remote disaster recovery, and is most often used to create a local backup of data in situations where having an exact copy of the data is essential.

In asynchronous replication, the replication software grabs data once it is written to disk, and rewrites it to a second array. In asynchronous replication, the application doesn't have to wait for any confirmations and can continue to operate. As a result, it has little or no impact on application performance, and can work effectively and economically over low bandwidth connections and long distances.

While it can't deliver the zero data loss available through synchronous replication, it can be configured to deliver RPOs measured in minutes, and recovery times measured in seconds, both of which are more than acceptable for most businesses. This combination of excellent data protection, minimal performance impact, long-distance effectiveness and low-cost deployment makes asynchronous replication an ideal solution for backing up data to a remote recovery site.

What to Look For in an Asynchronous Replication Solution

The asynchronous replication solution that makes the most sense for remote recovery implementation is the one that lets you implement the highest degree of data protection while making the most cost-effective use of your existing infrastructure. Specifically, you want a solution that works as-is with your existing applications and infrastructure, that poses no distance limitations, and that makes the most economical use of your existing bandwidth, enabling you to maximize data protection while minimizing the performance hit on your network overall.

One solution that clearly meets these requirements is Double-Take® from Double-Take Software. Double-Take combines patented asynchronous replication and failover technologies; it captures and replicates changes, as they happen, to a secondary array at any location, and then lets you recover from that location in seconds in the event of disaster.

Several Double-Take features combine to enable the highest level of data protection while maximizing your existing application and infrastructure investments:

- **Incremental, byte-level replication.** Double-take monitors all files and replicates only the bytes that change, as they change, which reduces replication traffic on your network to an absolute minimum.
- **Unlimited distance replication over standard IP networks.** With Double-Take you can replicate to a disaster recovery site in any location, as far away as necessary to minimize your vulnerability to natural or man-made disasters. And it replicates over any existing IP LAN, WAN, VPN or NAT.

- **Software and hardware independence.** Double-Take replicates with complete integrity from virtually any application; it runs on whatever hardware you have now and gives you the flexibility to choose whatever value-oriented hardware you prefer as your enterprise grows.
- **Bandwidth limiting.** Double-Take lets you put a cap on the bandwidth it uses, which lets you minimize or eliminate its impact on the performance of other applications, and of the network overall.
- **Proven savings.** Thousands of companies, including over half of the Fortune 500, protect their data with Double-Take; many, like Chicago-based MidAmerica Bank, use it instead of tape to provide cost-effective, up-to-the-minute remote data protection and disaster recovery for their headquarters and branch offices.

"Double-Take saves us about \$50,000 annually by eliminating the need for backup technical maintenance and assistance," says Ray Zamora, MidAmerica Bank's Vice President of Network Operations. "We also save with the simple setup of the Double-Take solution, which is less than a tenth of the cost of setting up tape drives. There's no recurring hardware maintenance cost and no loss of employee time in supervising the backup process. And during our last annual audit, our Double-Take solution for backup and recovery of branch data exceeded our expectations and requirements."

Augmenting or even replacing a tape-centric data protection scheme with more a more effective solution like continuous data replication can have a significant positive impact on your protection budget. And at the same time, it can provide a higher level of recoverability than tape. Consider the hypothetical Year 1 cost comparison of tape-based periodic backup with data replication provided by software such as Double-Take. In this example, you can look at the costs of each solution for remote disaster recovery for a company with ten offices each with 20 employees per office.

	Tape Backup	Double-Take
Fixed Costs		
Hardware, software & setup	\$120,000	\$31,570
Continuing Costs (per year)		
Maintenance	\$18,000	\$6,314
Media	\$36,000	\$0
Salary	\$150,000	\$15,000
Offsite pickup/storage	\$36,000	\$0
Year 1 Total	\$360,000	\$52,884

As you can see, replacing tape backups in branch offices with centralized backup methodologies and data replication can reduce the total costs of maintaining a recovery solution for those branch offices. Now, imagine one of these offices is struck by a disaster and needs to recover from the remote facility.

	Tape Backup	Double-Take
	Recovery to previous day's tapes in one business day - 8hrs	Recovery to 5-minute-old data with no downtime
Cost of unrecoverable lost data (\$10,000 per MB ⁵)	\$200,000	\$694.44
Cost of downtime (\$42,000 /hr ⁶)	\$336,000	\$0
Year 1 total	\$536,000	\$694.44

5 Smith, David M. "The Cost of Lost Data." Graziadio Business Report Vol. 6 No. 3

6 www.nwfusion.com/careers/2004/0105man.html

The numbers speak for themselves. Not only is relying on traditional tape backup methods costly and complex, it can negatively impact your ability to continue doing business after a disaster and can cause a company to incur additional expenses related to recreating critical lost data and employee productivity. While tape backup is the most common and cost effective method for protecting and recovering large amounts of data, it may be woefully unable to meet your established recovery goals. Tape backup is great for long-term archival and is most certainly a part of every disaster recovery plan. However, other solutions such as high availability, disk-based snapshots and data replication must become part of a company's overall data protection solution to be successful in meeting its RTO and RPO goals.

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 or daily rotations of tapes at remote offices where human error can affect the reliability of tape handling, 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 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, Double-Take Software clients are able to increase the reliability of tapes, reduce manpower costs, and eliminate hardware and backup software in the remote sites.

Business Continuity: A Broad View

To achieve larger business continuity goals, Double-Take 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 use 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 only 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.

Fortunately, Microsoft server OSs 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 INIs within program directories. During configuration of the backup job you 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, 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, 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.

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 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. TPA

Real Data Protection = Snapshots + Replication + Backup

Double-Take operates in compliment to Windows 2003 Volume Shadow copy Service. VSS can be used in conjunction with the replication technology between servers.

If you snapshot 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 you snapshot 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, you 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).

Beyond Protecting the Data – Recovering the Server

The complexity of traditional recovery solutions compounds an already difficult situation, and heightens the opportunity for human error. Speed and quality of recovery are extremely important when customers and employees are relying on access to critical data, but the average restoration takes hours at best. And with solutions like tape backup, even a successful recovery can result in the loss of any data that is new or has changed since the backup was made. The Double-Take Server Recovery Option is a whole-server data protection solution that, when combined with Double-Take real-time replication, simplifies the restoration process and reduces the time and effort involved with server recovery. Using Double-Take with the Server Recovery Option, the entire production server - its operating system, applications and data - can be protected and easily recovered to a new system quickly.

A common backup solution in a branch-office scenario is to replicate the branch servers to a central location and perform a nightly tape or disk backup. If the Branch 2 server fails, the administrator would have to provision the new server, install applications, then go through the cumbersome recovery process. On average, this would take more than one IT person and several hours. Still, when the recovery server is restored, Branch 2 is missing all of the data that is new or had changed since the night before.

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. Simply repair and then replace the storage on the production source server. The Double-Take database is aware of where the various target data files came from. Then the Restoration Manager can be used 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. A 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 two options.

1. 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 from which to quickly restore.
2. 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, you 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. It will also be immediately replicated back to the target, to provide consistency for all copies.

And in all cases, Double-Take provides 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, the replication set can be used 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, REGDMP can be used to routinely secure those particular registry hives to a flat file (which would be replicated to the target server, similar to the System State 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.

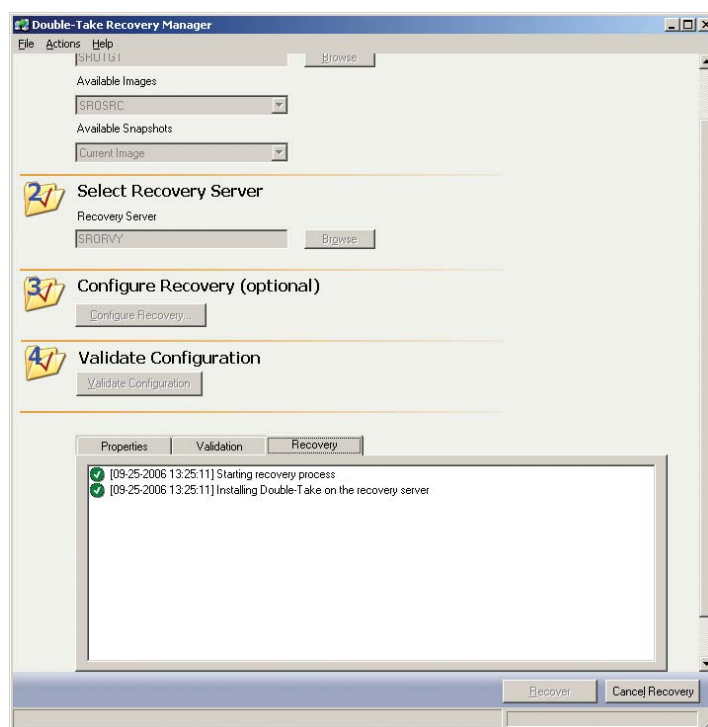
1. Extended Queuing - Double-Take provides for a queuing model to cache up to 4TB of byte-level changes, so that even the most dramatic data changes can be propagated to the target server. Since most environments do these types of operations on weekends, a properly configured queue and infrastructure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

2. 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 Double-Take Command Language (DTCL), 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". It will then verify and 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.

3. In-band command Processing - Many Double-Take 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 Double-Take Software provides the ability to insert a flag immediately behind the replication traffic of some operation. 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.

"Historically, organizations have dutifully performed their backups to tape and shipped a copy of these tapes off-site. However, as installed disk capacity has grown faster than tape performance, traditional tape-based backup solutions have fallen behind and are no longer meeting backup/restore requirements. Add the demands of 24x7 Web operations and e-commerce-based applications, and the need to supplement traditional backup and recovery methods becomes obvious."⁷

Double-Take, when combined with the Double-Take Server Recovery Option (SRO), provides a single solution to continuously protect and recover an entire server. Protection is provided by the industry-proven real-time replication of Double-Take while recovery is performed by the Server Recovery Option. The Recovery Manager, provided as part of the Double-Take Server Recovery Option, presents the task of server recovery as a series of easy-to-understand steps. Because Double-Take replication protects the entire production server - its operating system, applications and data, restoring the server encompasses as few steps as possible and provides a significantly better recovery time than existing solutions such as tape backup.



Let's look at the initial example scenario that was provided, instead utilizing the real-time, whole-server protection and streamlined recovery of Double-Take and the Double-Take Server Recovery Option. When the branch server being protected fails, the administrator in the central office uses the Recovery Manager to choose the appropriate recovery server. This recovery server resides on the network at the branch office, but only requires a baseline Windows® operating system to be installed - no additional provisioning of software is required. The Recovery Manager validates the compatibility of the recovery server, then quickly restores applications and data that have been protected, in real-time, by Double-Take. Once the restoration process has finished, business can resume at the branch office with minimal downtime and lost data.

Summary

There was a time when tape-based backup was widely believed to be the only feasible backup solution for remote offices. But tape backup alone is no longer the only realistic solution and centralized backup solutions are coming to the forefront. As a result of this shift, enterprises have re-evaluated tape-only backup solutions and don't like what they have found. According to a report by ESG⁸, nearly one-quarter (24%) of companies say that twenty percent (20%) or more of their tape-based backups fail. As such, depending on tape backup alone creates an unacceptable level of risk.

Fortunately, advances in technology have made centralized backup easier to manage and less of a drain on the WAN. A centralized backup solution takes the tape-only backup responsibilities out of the hands of non-technical resources in each remote office, and puts them in the hands of experts back at a central data center. Replicating data to a central location from branch offices using Double-Take can reduce the per-location costs associated with tape-only solutions and provide a higher level of recoverability for business-critical systems and data.

⁸ Goldworm, Barb. "Alternatives to Tapes on Trucks." DM Review. 28 October 2005

Double-Take Software Headquarters

257 Turnpike Road
Southborough, MA 01772
Phone: +1-800-964-0185 or +1-508-229-8483
Fax: +1-508-229-0866

Double-Take Software Sales

8470 Allison Pointe Blvd. Suite 300
Indianapolis, IN 46250
Phone: +1-888-674-9495 or +1-317-598-0185
Fax: +1-317-598-0187

Or visit us on the web at www.doubletake.com



Get the standard today: www.doubletake.com or 888-674-9495

Microsoft
GOLD CERTIFIED
Partner

