

# Enhancing Backup

Reel-to-reel tape systems were initially used to store all data but eventually became relegated to off-site storage and not actively used. As technology advanced, the tapes shrunk and became the backup tape solutions we use today. This paradigm of backup to removable tape media has worked for quite a long time, but has in recent years proven itself to have some drawbacks when used as the sole means of recovery.

## Tape Inefficiencies

Tape media is expensive and repetitive. Data is often stored for much longer than actively needed, simply due to the need to keep full backups on tape to be used for recovery in conjunction with incremental backups. This means that to back up a large server over the course of a month, dozens of tapes may be needed for each of the full and point-in-time backup operations to be stored. Once these backups age past a certain point, they must be physically moved to another location in order to meet most retention policies. This is because the tape systems must act on data physically local to them, or else require large bandwidth pipes to facilitate cross-site backup.

## Complexities

Tape backup systems must have exclusive access to data in order to commit a clean copy of that data to the tape media. This means one of two things: either the files must be allowed to be exclusively locked to the tape backup system or agents must allow for Volume Shadowcopy or other secondary data-streams. If agents are not used, then any file in use by another process (such as Microsoft Exchange Server or Oracle Database Services) will generate a fault. Most backup tools can retry the operations, but since these types of application keep files open as long as they are running, the retries will fail as well. Agents can get around this problem by leveraging either proprietary open-file technologies or native solutions like Microsoft Windows Volume Shadowcopy Services (VSS) to create temporarily open and available copies of locked files. The solution works, but typically does so at the expense of CPU cycles and other system resources. This process created the need for backup windows in most organizations, where server performance is expected to be limited during certain hours of the day (typically overnight and weekends) so that backups can complete

## Difficult Recoveries

Finally, once the data is committed to tape, few organizations perform restorations of the data until a fault occurs. This means that technical staff is unaware of the procedures or time frames for the recovery process, and that the data on the tapes is not confirmed to be useable. Without regular test restorations, data on the tapes cannot be confirmed as ready to restore, and therefore becomes suspect until such a test restore is done. Even with valid data, finding the right version of the right set of files on the right tape in the right sequence can become a nightmare during an emergency.

As a file set may be spread across multiple full/incremental backup jobs on multiple tapes in, potentially, multiple locations; just ensuring that all the correct tapes are loaded into the hardware readers can be a time consuming and difficult process. Added to this set of issues is the fact that most restorations must be performed to an offline system (i.e. a dismounted database) or special software must be pre-configured to allow live communication between the backup tool and the specific application being restored to. This leads to either additional downtime for restoration of data, or else additional complexity in the form of more agents and tool sets.

## The New Paradigm:

By using real-time replication to a single set of Repository Servers, Double-Take Backup allows you to minimize the amount of hardware you need to dedicate to the backup process for your organization. This limited number of Repository Servers can host both an up-to-date copy of all protected data along with the ability to recover from any point-in-time, all within the same set of physical or virtual machines. Tape backup equipment may still be used, but is now limited to backing up data from the Repository Servers,

# Enhancing Backup

as opposed to each production server. In addition, as the Repository Servers themselves can be housed off-site from the production systems, the backups taken are already off-site backups. This means that tape-based copies (if used) can be taken and housed in the same facility, minimizing transport costs. As Double-Take® Software real-time replication is bandwidth efficient via throttling, compression and byte-level replication technologies, the same data-set can be protected over much smaller connections between sites.

## Complexity is minimized

The Double-Take Software replication engine, which will reside on each protected production server, is designed from the ground up to protect open and locked files. This means that any application running on a Microsoft Windows or Red Hat or Suse Linux server can be properly protected, even if the files in question never close. Backup windows are also eliminated by continuous replication with minimal impact to production server resources. Where tape is still used, it is relegated to the Repository Servers, and therefore does not act on production resources, allowing your staff to manage tape backup operations at any time, while Double-Take® Backup replication protects the data from the production servers without the need to only operate during backup windows.

## Recovery, streamlined

By using Double-Take Backup, a real-time copy of all protected data and system-state information is housed on a mounted file system. This means that data integrity can be checked at any time by visually examining the Repository Servers, allowing the IT staff to prove that data which needs to be protected is indeed being backed up. Restoration of point-in-time data can be accomplished to a virtual directory on the Repository Server itself, eliminating the need to restore data to an offline production server in order to test recovery. Finally, in the event that actual recovery is needed, data can be mounted and tested on the virtual drive, then restored to either the original production server, a new server or (in the case of Microsoft Exchange 2003 and 2007) to a PST file for Outlook use. As the device that you restore to does not need to be defined ahead of time, you can dynamically recover part or all of a server to whatever hardware platform – physical or virtual – is available.

QUESTIONS?

1-888-674-9495

info@doubletake.com

www.doubletake.com



© Double-Take Software, Inc. All rights reserved. Double-Take, Balance Double-Take Cargo, Double-Take Flex, Double-Take for Hyper-V, Double-Take for Linux, Double-Take Move, Double-Take ShadowCaster, Double-Take for Virtual Systems, GeoCluster, Livewire, netBoot/i, NSI, sanFly, TimeData, TimeSpring, winBoot/i and associated logos are registered trademarks or trademarks of Double-Take Software, Inc. and/or its affiliates and subsidiaries in the United States and/or other countries. Microsoft, Hyper-V, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. Red Hat is a registered trademark of Red Hat, Inc. Novell, the Novell logo, the N logo, SUSE are registered trademarks of Novell, Inc. in the United States and other countries. All other trademarks are the property of their respective companies.