

All Data is not Created Equal

All data is not of equal value to your organization; revenue recognition, daily operations and competitive advantage rely on real-time electronic transactions, not paper. Any amount of lost or corrupt data can negatively impact the bottom line no matter what the cause. An operational recovery plan details how IT administrators will handle the sometimes daily challenges of lost or corrupted files, e-mails and transactional data.

Operational recovery is a critical component in any business continuity plan. When an organization depends on its data, the challenge after operational issues such as viruses, data corruption, deleted files or other human errors is to recover the necessary data quickly and easily with the minimal amount of data loss. For the most critical data, such as file-based intellectual property, financial accounting databases and e-mail, continuous protection of the data is necessary because even restored data that is five minutes old could cost the company thousands of dollars. Reducing or even eliminating dependence on periodic tape or disk -based backup solutions and implementing a Continuous Data Protection (CDP) solution can be a practical, affordable and non-intrusive way to ensure this never happens

Backup is a Waste of Time...

...and money if the data can't be recovered. Analyst firm, Enterprise Strategy Group, estimates that 40% of recoveries fail. The fault may not lie with the backup software or the tape, but rather with the complexity of the task being performed across the application, server, storage and network layers. Unless everything lines up flawlessly at the time the backup is being performed, the procedure is vulnerable to failure. Failure root cause analysis is almost impossible, and some backup providers estimate that they expend up to 90% of their time on problems that are not directly related to their software. While protecting and recovering any kind of data with traditional backup solutions is difficult, successfully protecting transactional data like SQL or Exchange databases with traditional backup is nearly impossible. In the case of these applications, recoveries using traditional backups frequently fail because of the extended list of tasks that must be completed with total accuracy to assure restoration of the right data at the right point-in-time:

1. All applications and services that access the transactional database must be stopped.
2. The last full backup must be found, loaded and recovered.
3. Then all relevant incremental backup copies must be found, loaded and recovered in the proper sequence.
4. Finally, the database logs (assuming they are available) must be properly applied. The process typically involves at least a systems administrator and database administrator and the combined operations are generally undertaken in a stressful environment, with deadlines and time pressure looming large.

Only after this operation is successfully complete can the transactional database be brought back into operation. Using disk-based technology rather than traditional tape may speed the data recovery, but the operations flow is still laborious, complex, and prone to error. In addition to being unreliable, backup is highly intrusive to the production business environment. Windows file systems and files must be captured at a consistent point in time. This generally necessitates the deployment of an open file manager. The backup software product is tasked with reading through copious volumes of file system metadata to determine which files have changed, then proceeding to copy the entire file system and/or database, or individual files that have changed. Depending on the size of the system, this process can take hours, inhibiting use of the application server for production. Because of the time and impact of the file walking process, most corporations perform backups no more than once per day. In the vast majority of environments, full backups are performed weekly, with incremental backups performed daily. It's hardly surprising that reducing the backup window was identified as the most pressing storage issue facing companies in a recent survey conducted by SearchStorage.

TimeData™: Not Your Father's CDP

With all of the hype in the industry today over Continuous Data Protection (CDP) solutions, it is sometimes hard to separate the buzz from the benefit. For customers focused on solving real-world problems involving the recovery of business-critical data, understanding the distinction is paramount to making the right choices for safeguarding their electronic assets. To date, CDP products have been viewed as high-priced, 'niche' solutions. TimeData changes that perception by providing something really useful for organizations – the ability to recover from isolated data disasters – and delivering it as a simple, easy-to-use software product which is affordable for companies of all sizes.

TimeData is a file-based CDP product which allows organizations to recover lost or corrupted data at a granular level for Microsoft Exchange, Microsoft SQL Server, and other Windows-based applications. TimeData can quickly and easily recover data from common operational issues such as viruses, data corruption, deleted files or other human errors. It is optimized to work with Microsoft Exchange and Microsoft SQL Server, and can flag points in time when these applications are in a known good state without impact on the production system. TimeData also provides e-mail-level recovery for Microsoft Exchange, which allows organizations to easily recover individual messages or mailboxes.

File-based continuous data protection products give you the ability to go back in time to take corrective action by turning back time and capturing a data file that was lost or corrupted at the precise moment just before the deletion or corruption event occurred. TimeData's file-based CDP capabilities work by capturing file system writes as they happen and copying these changes, in real-time, to a secondary data store (the TimeData CDP Repository) running a time-enabled file system. The essential components in a TimeData system are the servers being protected, the TimeData Repository Server you configure to protect them, and the networks that connect protected servers to the repository server.

TimeData Data Servers

A data server is a computer on your corporate network that manages and performs operations on its resident data. This data may include information in files, SQL Server databases, or Exchange storage groups. TimeData protects data associated with SQL Server database servers, Exchange servers, and Windows file servers. A data server can be:

- a stand-alone computer (in other words, not part of a server cluster)
- a node in a server cluster
- a virtualized machine running in a VMware® or Microsoft® virtualized environment

TimeData Repository Server

The TimeData Repository Server is a dedicated server that stores information from one or more protected data servers. One repository server can support many data servers.

Since changes are tracked using a file system filter driver, protection is transparent to the application. There is no operational backup window because data changes can be replicated while applications are running and the data is in use. Data copied to the time-enabled repository is indexed at each write, as it occurs, so that views of the data can be directly retrieved from any prior moment in the file's life.




How TimeData Works

Backup - protection gap is typically 24 hours




Snapshot - protection gap typically 1 to 3 hours



Replication - only last version is available



TimeData Continuous Data Protection



Conceptually this process is similar to a series of snapshots or continuous replication with some major differences which make it a far better solution for operational recovery. Unlike snapshots however, TimeData has the ability to directly retrieve and restore data from any point in time and at any level of granularity, from the individual file to the full database. Unlike backup or snapshots that capture a series of discrete data views based on arbitrary, pre-defined time increments, TimeData captures all changes as they

are made to the data. TimeData software on the database or file server transparently and automatically intercepts calls to the file system in real time and sends them to a repository residing on a separate system. This effectively means that Recovery Point Objectives are no longer measured in hours or days, but are real time as data changes. TimeData file-based CDP continuously monitors the stream of data being sent to the repository and analyzes it with an application-aware event processing engine, allowing it to identify and highlight points in time that have full integrity, such as SQL/Server checkpoints, SQL named transactions, Exchange database start/stop events or file closes. Because data is collected and viewed at the file level, the administrator can select, by policy, to collect and save only important business data. Data retention, even as granular as the file level, can also be managed by policy.

Also unlike backup, file-based CDP intercepts the data change event at the file system, is unobtrusive and does not require that applications or files be closed or quiesced. The process is continuous and the impact on the data server is negligible. With TimeData, the concept and limitations of backup windows is obsolete. Since there are no schedules and no media to manage, file-based CDP is simple and not prone to the errors and complexity that diminish the reliability of traditional backup. Like replication, TimeData improves data availability by maintaining a repository in which there is a replica of the primary data set. Unlike replication, where the replica provides what is essentially a present view only, or static historical views from split mirror replicas, backup or snapshots, file based CDP offers continuous replica views. In addition to a present view, TimeData's patented "return to the moment technology" enables you to quickly view multiple past moments from its repository. Copying any desired point-in-time view to a production or secondary server is quick, simple, and intuitive. If desired, only application-consistent views are presented (e.g., database checkpoints), thereby cutting down on what has traditionally been a laborious process using tape-based or disk-based periodic backup and recovery. Because of the simplicity of recovery, TimeData can cut down Recovery Time Objectives (RTO) from hours to minutes.

TimeData software is designed to integrate seamlessly into the Windows environment, without any requirement for proprietary systems, hardware or for special operator skills or training. Although improving RPO and RTO is valuable, mid-tier companies in particular need to leverage their existing skills and infrastructure. Installing the software on the production servers to be protected is quick and easy. It takes moments and does not require a system reboot and does not disrupt operations. Installing and configuring the repository where the protected data is stored is also fast and straightforward. The system on which the TimeData Repository Server is installed can be standard Windows Server running Windows Server 2003 or Windows Server 2008 and connected to the production servers via a standard Ethernet network. Configuration and policy management is done through a management interface that is intuitive and presents a task-based view for the administrator, hiding recovery capabilities when in "configuration mode" and hiding the configuration capabilities when in "recovery mode".

For organizations that decide not to completely abandon traditional backup solutions, TimeData integrates with existing backup systems to improve data protection, while simplifying operations. Businesses can continue to run archive operations or full backups to tape and maintain their offsite vaulting procedures but eliminate the media costs, complexity and unreliability of incremental or differential backups by cutting down the frequency of tape backups to weekly or monthly tasks. Backups to tape can be performed on production data servers directly or backup can be eliminated entirely on production servers by backing up from the TimeData Repository Server instead. The completeness and simplicity of TimeData's approach to data protection and availability make it a practical, affordable, and non-intrusive approach to enhancing Operational recovery and improving an organization's Recovery Point Objective (RPO) and Recovery Time Objective (RTO).

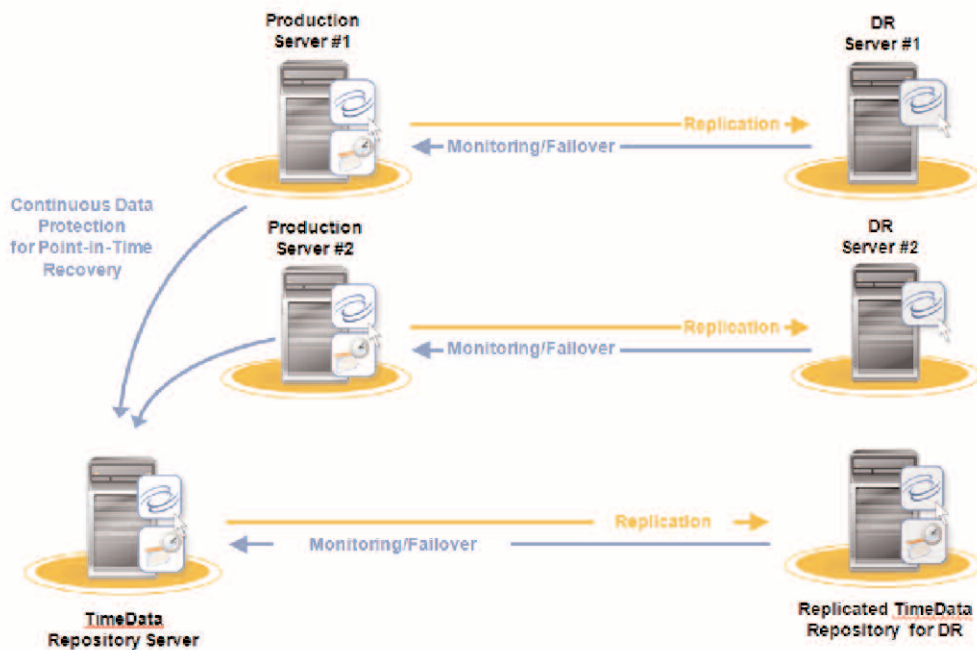
Disasters are like Pajamas...

...they come in different sizes. For site-wide disasters or server failures where the goal is to get back up and running as quickly as possible, IT departments need to focus on disaster recovery, protecting both the data and the availability of the application. For isolated data failures, which occur with more frequency, those same IT departments need to focus on improving their operational recovery capabilities. TimeData allows customers to quickly and easily recover from common operational issues such as viruses, data corruption, deleted files or other human errors. Coupled with Double-Take®, the industry's leading disaster recovery and high availability offering, TimeData extends companies' recovery capabilities to include operational recovery in addition to disaster recovery.

Double-Take for Windows is a product focused on disaster recovery. This means the focus is on remote replication, monitoring and failover. The goal is to allow the data and the applications to be recovered as quickly as possible with the most recent copy of the data available. Though changes to protected data are being continually replicated, this is very different from the industry (SNIA) definition of CDP. TimeData, however, was designed to be a true CDP solution (http://en.wikipedia.org/wiki/Continuous_data_protection). It tracks every single change that occurs to protected data and saves all of those changes in a repository so that you can recovery to ANY point in time. It provides the ability to easily recover single files, databases or even single e-mails quickly.

	TimeData	Double-Take for Windows
Continuous, Byte-level Protection	X	X
Application-aware protection for File Servers, Exchange, SQL and SharePoint	X	X
Focused on local protection (LAN)	X	
Granular 'PIT' Recovery (i.e., message, file or database –level restore)	X	
Focused on remote protection (WAN)		X
Automated Monitoring and Failover		X
Full-Server Protection for Automated Failover or "On Demand" Recovery		X

Because both products can be installed alongside each other, organizations can ensure they meet their overall business continuity needs by providing disaster recovery and operational recovery for critical data and applications.



Summary

By continuously capturing data changes as they are made to the file system, and maintaining them in a separate repository, TimeData offers complete, transparent protection of Windows data from physical and logical corruption. This file based continuous data protection improves Recovery Point Objectives for application data from hours or days to transaction time. Recovery Time Objectives can be reduced from hours to minutes, as the TimeData repository makes past, consistent moments for application data sets, file systems or even individual files, quickly available for viewing or attaching without the complexity and delay associated with tape recovery and extensive retrieval operations. Because TimeData is designed and built upon standard Microsoft technology and uses efficient, patented data capture and maintenance technology, it is easy to deploy and use, making it a very affordable alternative to complex proprietary, point solutions requiring customization, specialist skills and training. TimeData leverages standard infrastructure and skills, making it the ideal choice for mid-tier companies and departments looking to increase data protection effectiveness, not cost and complexity.

