



Redundancy and the Last Mile to Recovery

Executive Overview

The push to achieve the vision for a real-time infrastructure is driving many storage administrators and IT managers to implement redundancy and other advanced storage strategies. While these technologies are improving availability, they could be actually causing exposures in the recoverability of that business-critical data.

This white paper outlines some of the most common data recoverability challenges, provides an overview of the technologies designed to help—and offers recommendations on using these tools to intelligently cope with exposures.

White Paper

Today's IT staffs are facing some significant challenges. Executives are calling for more agility and flexibility from their IT systems. Data and application availability has become a battle cry for these managers and the IT teams who support them.

The vision for a real-time infrastructure is well-documented by analysts and enterprises alike. In their 2006 analysis of the real-time infrastructure, Gartner notes that integrated systems and storage management, virtualization, and IT Information Library (ITIL) will deliver long-term benefits. Reducing data access response time and vastly improving the quality of service to end users are critical driving factors for going down the road to a real-time infrastructure.

	Basic	Standardized	Rationalized	Virtualized	Service-Based	Real-Time
	<i>React</i>	<i>Reduce Complexity</i>	<i>Economies of Scale</i>	<i>Flexibility</i>	<i>Service-Level Delivery</i>	<i>Business Agility</i>
Agility	Months to Weeks	Weeks	Weeks to Days	Weeks to Minutes	Minutes	Minutes to Seconds
Economics	Subsidized	Cost Center	Static Usage	Flexible Usage Costing	Variable Usage Costing	Variable Business Investment
Quality of Service	No SLAs	Basic SLAs	Class-of-Service SLAs	Flexible SLAs	End-to-End SLAs	Business SLAs

Gartner, Achieving Agility: The Data Center Is the Foundation, April 2006

The Retention Continuum

Certainly making data accessible to users whenever and wherever they need it is a noble effort, but it is important to remember that data recoverability goes hand in hand with availability—they are not mutually exclusive. In fact, many availability-focused storage strategies—when not properly used—can actually harm the recoverability of an organization's critical data.

How effectively you cope with daily data retention and storage management challenges can significantly impact the recoverability of your critical data. That's why it is so crucial to proactively create effective recovery strategies in the face of today's increasing data volumes, decreasing resource utilization, tightening budgets and a growing focus on retention and reporting.

Unfortunately, not saving enough data—or saving it incorrectly—are all-too-common mistakes that create vulnerabilities. In addition, saving too much data may seem like a good availability solution at first, but it can actually leave many exposures for recovery, compliance or legal discovery issues.

“Enterprises should consider replication and backup as the two core technology components of recovery, and they should design combined recovery strategies that incorporate both technologies to address planned and unplanned downtime.”

—Gartner, November 2005, Management Update: *Predicts 2006: Storage Technology Evolves Along with Demand*

Buying more storage and investing in increasingly advanced replication techniques to save growing volumes of data might be a good action plan for a retail operation, but it is not an effective plan of action for your storage management program. The latest statistics show that you'll spend \$8.00 managing every \$1.00 of disk storage you buy—which means that the more you buy, the more you must manage, and the more money you spend. Too often, the focus is on saving data “just in case” when it *should* be on quickly recovering mission-critical data and applications.

Because many of today's storage management and data retention solutions leave users with more questions than answers, it isn't always easy to understand the best ways to protect your business. But, thankfully, you can start making smarter technology choices with the right information at hand.

Coping with Challenges

A good first step in ensuring data recoverability is to understand the strengths and weaknesses of the most current data protection and recovery technologies available.

Availability Solutions

Availability solutions—such as tape options, mirroring, virtualization and information lifecycle management—ensure that data is available when and where you need it. However, these strategies assume continuous connectivity and zero data loss or defects.

>> *Virtual Tape Systems:* Integral to most disaster recovery applications, Virtual Tape Systems can greatly shorten processing time, are low cost and are widely considered a primary medium for backup. Virtual Tape Systems require few—if any—changes to the existing backup structure, and are faster than magnetic tape. But, in order to meet vaulting and archival requirements, transfer to removable media is still required with Virtual Tape Systems. What's more, additional equipment is sometimes required at the recovery site, which may be inconvenient.

Replication Solutions

Replication solutions do a fantastic job of preserving data in real time. Unfortunately, they don't account for potential data corruptions. Replicating a corrupt data set simply makes two copies of bad data—that's certainly not a recipe for recoverability!

- >> *Synchronous Replication*: Synchronous replication ensures zero data loss, but it is limited by distance (100 miles in theory, 20 miles realistically). Plus, with synchronous replication, I/O operations must be completed in order on both the local and remote disks before processing continues.
- >> *Asynchronous Replication*: Asynchronous replication removes the distance barrier so your remote site can be located anywhere in the world. With asynchronous replication, the process of creating a remote copy of your data runs independently of your processing. Unfortunately, this solution is highly subject to a loss of data integrity. What's more, failure points are unknown and recovery using the remote copy as backup may be difficult or impossible.

“Replication solutions transmit data when it is created and as it is changed, which is good for keeping the copy up to date; however this provides recovery only to the most recent state of the data, which in the case of an application error or virus would be corrupt.”

—Gartner, June 2006, Continuous Data Protection: What It Is and Why It Is Needed, Dave Russell.

Real-time Analysis for Continuous Data Protection

An exciting revolution in redundancy and recoverability is found in continuous data protection (CDP). CDP provides for multiple copies or snapshots of data, instead of only a single backup every day. Gartner notes that while many technologies have provided snapshot copies, CDP offers additional flexibility and reliability by creating snapshots at logical or critical times, such as every time a new write function is done (“True-CDP”) or during file closing or other specific policy-driven timelines (“Near-CDP”).

Gartner also explains that increased bandwidth and disk requirements makes CDP unlikely to be deployed for entire enterprise data sets, and is more likely to be implemented in situations where current backup or recovery methods expose the organization to data loss for critical applications.

The question at this juncture is: how does an organization know which data sets or applications are truly critical to the enterprise, and would require speedy recovery (with a very tight RPO and RTO). Real-time analysis and scheduling is required to truly determine the files, data and applications that are used most often and to manage the immediate backup of changes to that data to ensure recoverability.

This concept of “Application-Driven Data Management”—managing data backup and recovery based on the requirements of critical applications—should be a key consideration for IT managers and storage administrators when walking the line between redundancy and recoverability.

Conclusion: Creating a Solution that Helps Secure Data Protection

As you work to create an effective storage strategy in the face of complex data retention policies and regulations, it is critical to secure intelligent data protection. This means identifying the data you really need, saving it properly, and knowing how you’ll get it back in the event of a disaster.

It’s a good rule of thumb to never retain data simply based on the “last used” date—or a random retention policy. Thankfully, there are software solutions available on the market today that can help ensure you never save more data than you require. By only saving critical data, you can avoid overbuying storage hardware and software. Plus, you can also avert the management expenses associated with those assets.

By applying the latest cutting-edge approach to data management you can determine exactly what data requires mirroring—thus reducing the volume of data being replicated and increasing the recoverability factor of your organization. You can also establish what files are most needed by your enterprise applications, thereby uncovering which files need to be recovered first.

About 21st Century Software

21st Century Software delivers Application-Driven Data Management (ADDM) solutions that retain and recover data and files required by critical enterprise applications. 21st Century Software's VFI family of software solutions are used by organizations across many industries and of all sizes to bring intelligence to data protection.

VFI solutions are intelligent, automated, real-time data protection from 21st Century Software. VFI constantly monitors files and data used by enterprise applications to determine which files would be most critical to recover after a disaster, hardware failure, or other calamitous event.

The VFI family of solutions can be critical elements in your enterprise backup and recovery strategies by helping companies address ongoing data management challenges:

- >> *Recovery of critical files:* VFI solutions monitor data usage by enterprise applications to ensure the recoverability of critical data required to run the business.
- >> *Regulatory compliance:* VFI solutions help organizations determine which data and information needs to be retained to comply with governmental and industry regulations.
- >> *Storage utilization:* By retaining only the most critical files and optimizing data management processes, VFI solutions deliver greater insight into data management processes and help improve storage utilization.



For more information:

21st Century Software, Inc.
940 West Valley Road, Suite 1604, Wayne, PA 19087-1823

Tel. (800) 555-6845 or (610) 971-9946

Fax. (610) 964-9072

Email: sales@21stcenturysoftware.com

www.21stcenturysoftware.com