

STONEFLY

**The Secret to Affordably
Protecting Critical Data
Using StoneFly Backup Advantage™**



All trademark names are the property of their respective companies. This publication contains opinions of StoneFly, Inc. which are subject to change from time to time.

This publication is Copyright© by StoneFly, Inc. All rights reserved. All information contained in this document is subject to change without notice.

This document is intended for use only by recipients authorized by StoneFly, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of StoneFly, Inc., is in violation of U.S. copyright law.

Introduction

Whenever IT managers assess the value of their company's data, they should ask two basic questions:

1. What is the immediate impact if we lose access to the data?
2. How quickly must we regain that access?

The answers to these questions will determine the type of required data protection. Fundamentally, data that is essential to the business and represents the company's lifeblood requires the highest level of protection, since losing essential data could financially cripple or kill the business. Therefore, restoring access to this critical data needs to be immediate and as close to the point of loss as possible. The catch is that the cost to acquire this level of protection can be financially crippling as well.

Fortunately, not all data requires the highest level of protection. Companies that classify and segregate their data based on access patterns and value can lower the cost of data recovery by using a tiered data protection strategy.

Today, many IT managers are looking to innovative new practices using tiered storage and information classification to optimize access and retrieval while balancing regulatory retention requirements. The ultimate goal is to develop storage and operational practices that not only maximize storage investments, but also ensure that data is protected. Often it requires integrating multiple heterogeneous storage methods with backup and restore, local mirroring, remote data replication, security, and archival applications with information life-cycle management (ILM) policies (ILM is the process of aligning data in tiered storage according to its value).

When it comes to choosing technology, storage and systems administrators know that one-size-fits-all storage practices are not cost-effective. For example, high-end, high-performance storage is not practical for older or less frequently accessed information while storing frequently accessed information on tape is operationally impractical. The secret to affordably protecting critical data is designing a multi-layered storage solution that best meets the data recovery requirements of a mid-size organization or department of a larger enterprise.

A Conceptual Framework

The first step in developing a tiered data storage strategy is to examine the types of information you store and the time required to restore the different data classes to full operation in the event of a disaster. In data recovery, there are two main measures that determine the required level of service:

- **Recovery Point Objective (RPO)** This is the point to which data must be recovered using the data recovery process. The primary goal for most companies is to recover to the point of the last "good" data point, such as the last completed transaction, or the last complete email message received.
- **Recovery Time Objective (RTO)** This is the period of time within which systems and applications must be recovered after an outage. This includes application recovery time.

To better understand RPO and RTO, here are a few examples:

- With an Ecommerce transaction processing system, the RPO objective may be to recover the data to the last completed transaction; the RTO may be to recover within a

Secrets to Affordably Protecting Critical Data

few seconds of the interruption at the primary site. The best data recovery method may be synchronous mirroring.

- Recovering an email system, the RTO may be to restore operation at a distant site within a few minutes or an hour; the RPO may be to restore the last received email and the method used may be asynchronous mirroring.
- Without any sort of data replication, the RPO for a deleted user file may be to restore the last known good copy, and the RTO may be within a few hours or a day and the method may be to restore from backup tape or a Point in Time (PIT) copy on disk

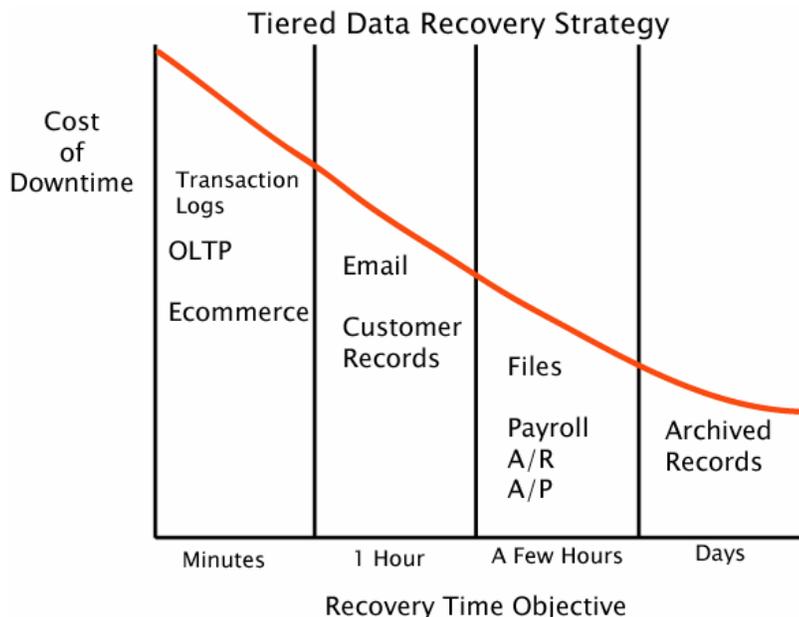
Applications often have differing RPO and RTO parameters depending on the nature of the application and its criticality to the business. For example, the SQL data for an order entry system will have a much different RPO and RTO than historical payroll records. An airline reservation system or a bank's financial tracking system will have a much different RPO and RTO than retailer's stock ordering system.

The type of data protection method required is completely dependent on the RPO/RTO objective, which in turn is driven by the application and business requirements.

Implementing a data protection strategy can be expensive, and it is very important to match the implementation to the application and business requirements. The tendency for many IT managers has been to simplify matters and apply the same data recovery solution to a very wide range of applications. However, a single solution can be extremely expensive and may in fact not actually meet the business needs.

There are still applications where a nightly tape backup sent to a remote data recovery facility is appropriate. Applications such as payroll, monthly sales results, quarterly financial results are all applications that could potentially fall into this category. As long as the data can be re-created and the application run within 2 or 3 days this is probably adequate. However, for some companies, even these applications may be critical at certain times.

The evidence is that the business need must be understood and the right data recovery solution applied to match the business need in various disaster scenarios. As a rule, the shorter the RTO and RPO, the more resource intensive the Data Recovery solution will be.



Defining Data Recovery and ILM Objectives

The next step is to define how your company's information should be categorized and classified not only for matching them to the appropriate data recovery solution, but also for security and business continuity purposes. The Information Lifecycle Management rankings listed below have been developed by the Data Management Forum of the Storage Networking Industry Association (SNIA).

Information Classifications:

- Mission-critical
- Business vital
- Business important
- Important productive
- Not important
- Discard

To determine how your information should be categorized, classify your data by how frequently the information is used, the availability requirement, the financial impact if the data is lost, and the operational impact.

Usage pattern of the information:

- Frequent
- Regular
- Periodic
- Occasional
- Rare
- On demand or request
- Never

Information availability requirement:

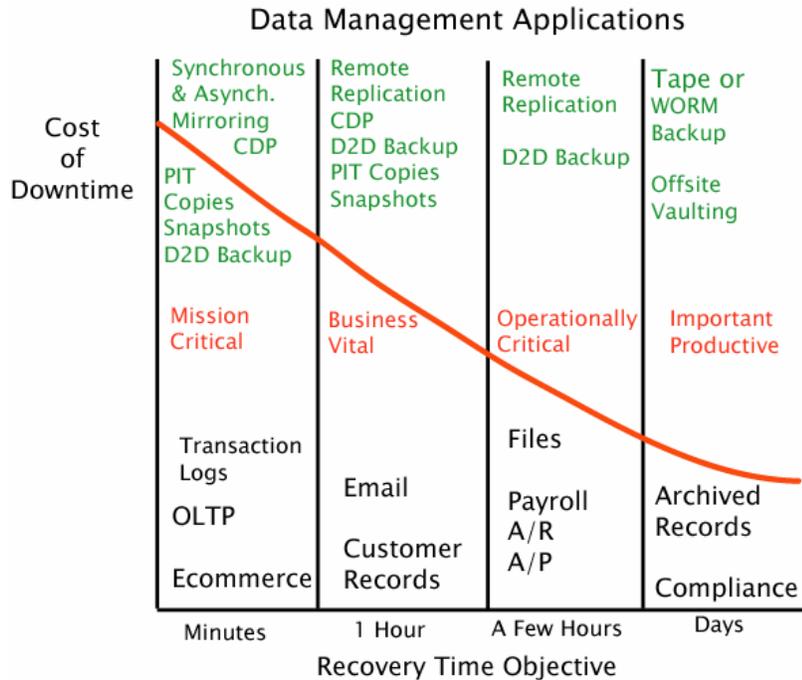
- Immediate
- Reasonable
- Defined time frame
- Extended time frames
- Limited
- Not defined or unnecessary

Financial impact of information unavailability:

- Significant and immediate
- Significant long- and/or short-term
- Potential long-term
- Possible impact
- Unlikely or no impact

Operational impact of information unavailability:

- Significant and immediate
- Significant over time
- Probable over time
- Possible over time
- Doubtful impact



After examining the Usage Patterns, Availability Requirement, Financial Impact and Operational Impact, it is easy to see how different types of data recovery applications fall into the Tiered Data Recovery Strategy matrix.

Data Recovery Solutions

The term data recovery or disaster recovery can mean many things to many people. Understanding what data recovery methods there are and how to meet data recovery goals is the next step in developing a successful data recovery strategy.

Information technologists are now starting to look at disk as both a local and remote layer in their storage hierarchy. Disk is being used for point-in-time copies of data, mirroring, and remote replication for fast application recovery. Companies faced with SEC and HIPAA compliance are also using disk for both primary and secondary data repositories. Effective use of multiple levels of disk types, locations and judicious use of tape and WORM storage can positively impact the corporate bottom line.

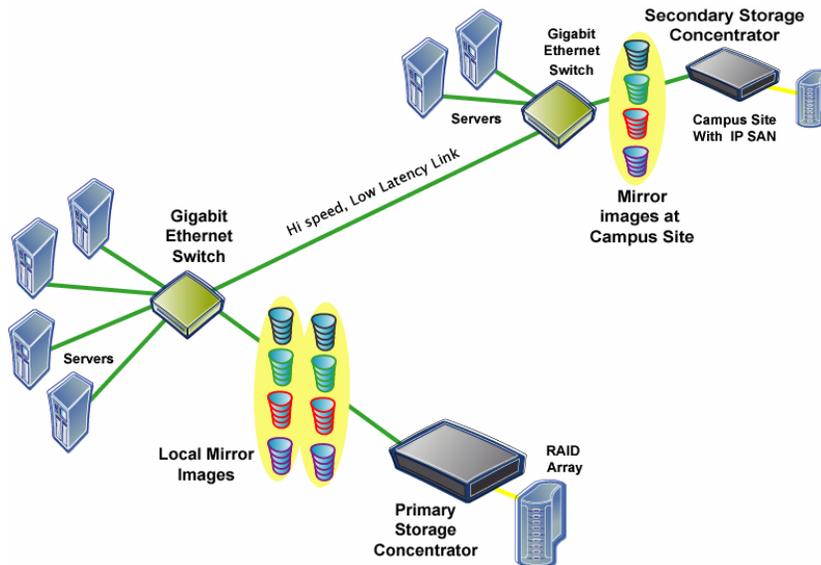
StoneFly Backup Advantage™ (SBA) is a complete turnkey data protection and recovery solution based on the advanced features of the innovative and award-winning SBA software. This integrated solution provides disk-to-disk backup using an IP SAN, split mirror backups (which allows users to automate backup operations while leaving production servers on-line), synchronous mirroring and remote data replication. SBA provides customers with different recovery point options for different classes of data.

Synchronous Mirroring

Today the type of replication that keeps data most up to date is synchronous mirroring, which has been implemented by many IT shops as a method of securing data in the event of a disaster. Synchronous mirroring provides immediate recovery to meet the most stringent RTO objectives.

Secrets to Affordably Protecting Critical Data

Synchronous mirroring can be local, campus or remote. Local mirroring protects against basic hardware failure or corruption due to bad disks. Campus mirroring, such as StoneFly Reflection™, can provide cost-effective protection against a disaster such as a fire in the primary site. Because of the high cost of high-speed communications links, remote synchronous mirroring can be an expensive solution and is typically only cost effective for critical applications at Fortune 500 companies in “glasshouse” installations.



With StoneFly Reflection, each mirror image contains an identical copy of the data, and can be substituted for one another in case of data loss or errors. Mirror images can be either local mirrors or they can be campus mirrors.

The downside to synchronous mirroring is that if there is data corruption, it can be copied to the mirror rendering all copies useless. The best practice is to combine synchronous mirroring with the most appropriate RPO method found in the StoneFly Backup Advantage solutions suite to ensure data protection objectives.

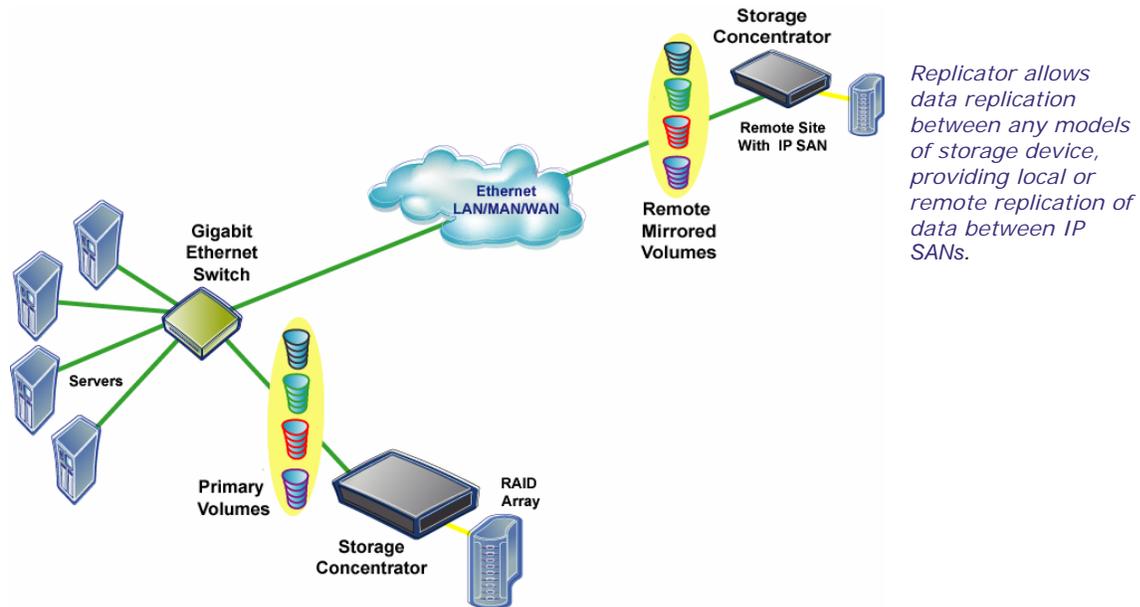
Asynchronous Mirroring to a Remote Site

A data recovery solution attractive to many IT managers for its cost effectiveness is asynchronous volume-based replication, a solution that can be applied to a much wider range of business applications because of its low cost. In addition, host-based asynchronous replication can leverage existing infrastructure and does not require specialized hardware to implement. Moreover, it does not impact application performance as synchronous replication does.

StoneFly Replicator™ asynchronous mirroring software used in conjunction with an IP SAN provides an affordable solution to keep critical data protected, highly available, and can leverage existing investments in hardware and personnel. StoneFly Replicator asynchronous mirroring supports storage-independent, local or remote, asynchronous data replication over unlimited distances using iSCSI on an Ethernet network.

With asynchronous mirroring, RPO is slightly behind an RPO for synchronous mirroring. There can be a number of writes lost in asynchronous mirroring because it doesn't wait for write-complete acknowledgements from the source. The ability to replicate to a remote site affordably is highly beneficial in the overall data protection program since asynchronous replication is less costly to implement than synchronous mirroring.

Secrets to Affordably Protecting Critical Data



Continuous Data Protection (CDP)

CDP is an integrated data recovery solution that replicates not only files, but also databases and entire applications. It is ideal for meeting data recovery point objectives (RPO) using advanced Continuous Data Protection. StoneFly Replicator[®] CDP provides quick and easy data recovery from any catastrophe. Rewind technology provides a comprehensive undo capability that allows system administrators to instantly rollback damaged data resources to a previous, valid state. The CDP software maintains fully functional replicas of file, database and complete application servers with real-time replication carried out over the WAN. Should a server fall prey to unexpected failure, its synchronized counterpart may be brought on-line to maintain high data availability, and users will have continued access to information resources.

Snapshots and Point-in-Time (PIT) SnapCopy

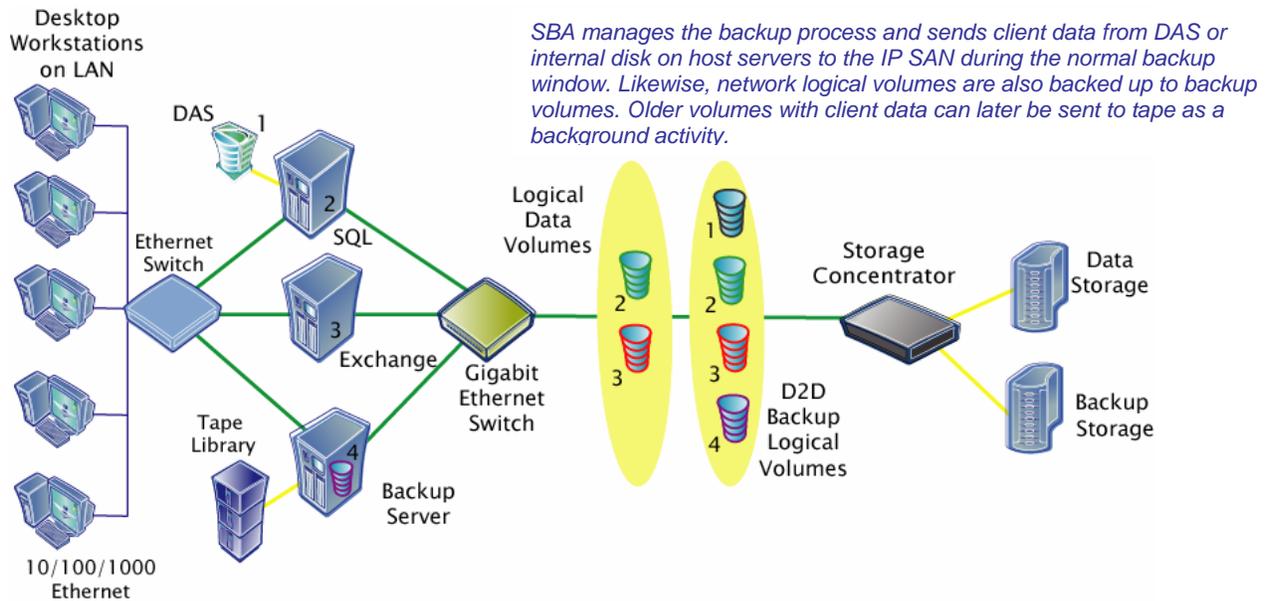
PIT copies and Snapshots are an efficient alternative to mirroring. Snapshots are virtual, temporary, and perishable point-in-time images of an active live volume. Snapshots can be created nearly instantaneously requiring less time to quiesce applications and far less server overhead for backups. You can use StoneFly Snapshot for Zero backup windows, restoring deleted files, using a volume for testing purposes, etc.

With a PIT SnapCopy, a separate complete copy of data is saved. Since a SnapCopy is completely redundant, it can't be affected by subsequent activity or failures on the original. It also takes up as much storage space as the original data.

Disk-to-Disk-to-Tape Backup

Gartner Group has stated that the most costly storage administration task is backup/data protection. Tape expenses have not reduced in cost on par with the reduction in disk prices. In addition, backing up to tape has become more, not less, complex as data volumes continue to grow exponentially. Backup windows have ceased to exist because critical applications must run near continuously. More importantly, downtime is not an option, whether planned or unplanned.

Secrets to Affordably Protecting Critical Data



The cost of disk is diminishing quickly, especially in relation to the cost of tape. Implementing a D2D backup system does not eliminate tape---the combination offers more recovery options. D2D is more than just faster throughput versus current tape solutions. It's about efficient data management and maximizing effective use of all storage types in all locations. Each storage solution is meant for a particular phase in the lifecycle of a piece of critical information.

The main benefits of StoneFly Backup Advantage D2D backup include:

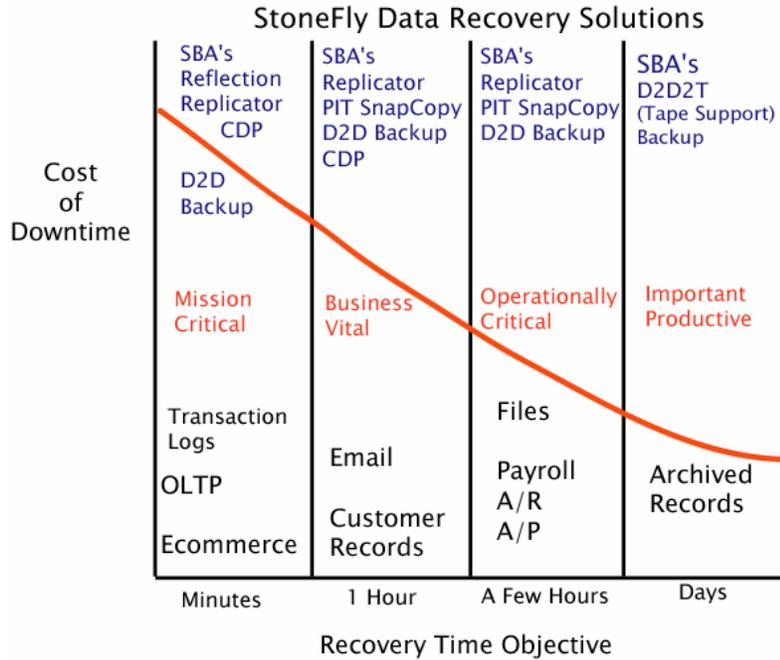
- Disks are much faster than tapes, especially when mount and seek times for tape are factored in. Disks are random access devices and are optimized for rapid access – tapes are sequential and are much slower for random file recalls. Disk offers much faster restore options
- Uses cost-effective Serial ATA disk arrays for fast, near-term backup and restore, which can lower cost of ownership, improve ease of use, and achieve higher performance
- Maximizes your investment in people, equipment, and infrastructure while minimizing costs and simplifying operations
- Offloads LAN-attached backups to the IP SAN creating a "LAN-free" backup environment
- Replicate backup data to a remote IP SAN to provide effective disaster recovery and rapid business resumption at a minimal cost
- Achieves the benefits of disk-to-disk backup, including faster restores, lower labor costs (50% lower according to the Metagroup), lower operating costs, and increased efficiency
- Reduces the impact of tape-related failures and human error in handling tapes.
According to *Diogenes Analytical Laboratories*, "Tape management is such an intensive manual operation that it can consume 60% to 70% of the effort associated with storage management."

Three trends are making IP SAN-based D2D2T solutions, such as StoneFly Backup Advantage, not only feasible, but also viable for mid-size enterprises and departments of larger corporations. The first is the availability of inexpensive Serial ATA drive, providing low cost disks that can be employed specifically for backup and restore. Second, the iSCSI standard (Internet SCSI) allows companies to deploy low-cost IP SANs (Internet Protocol

Secrets to Affordably Protecting Critical Data

Storage Area Networks) for disk-to-disk backup. A new generation of backup and restore software directly supports disk-to-disk backup and restore methods. Additionally, large-capacity optical storage technologies can be employed instead of tape to meet the new archiving and compliance regulations.

By mapping possible data management applications, such as those found in StoneFly Backup Advantage (including synchronous and asynchronous mirroring, PIT copies, D2D, etc.), it is easy to determine which data recovery applications might apply to specific classes of information as shown in the example below.



An integrated solution, such as StoneFly Backup Advantage, allows you to select the data recovery solution that best meets your business needs for each RTO and RPO level. The ultimate goal is to develop your own storage and operational practices that streamline storage investments, ensure that data is adequately protected and that lower the total cost of operation.