



STONEFLY

Seven Steps to Cost-Effective Disaster Recovery With IP SAN Technologies

**Disk-to-Disk Backup & Restore
and Off-Site Replication
Now Within Reach of SMEs**

May 2006

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 © 2006 by StoneFly, Inc. and 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

Small and medium-sized enterprises (SMEs) face the same storage management headaches as Fortune 1000 companies, but they often do not have either the time or the resources to become SAN architects or implement costly and complex Fibre Channel SANs.

Implementing an effective backup/restore strategy is difficult enough in today's server-centric network architectures, but the additional requirement for offsite replication and/or backup can be overwhelming for smaller companies. Downtime can result from causes far short of a disaster, including the accidental deletion of a key file or the failure of a disk drive. With the gold standard of electronic business at 99.999% uptime, even planned outages for necessary jobs like reconfiguring servers and loading new software take a significant toll on a business.

Currently, most IT managers at SMEs are forced to separately evaluate and choose infrastructure components, backup, migration, compliance and disaster recovery products from many disparate sources. What IT managers truly need are integrated, optimized, customizable solutions that go beyond storage infrastructure and provide a complete data management platform on which to build and deploy their business applications. It's no wonder that an *INFOWORLD* research report on Storage, cited data protection, disaster recovery, archiving and email as the top drivers for IT spending and that networking storage is a top priority.

Fortunately, IP SAN technology offers the cost-conscious IT organization searching for a reliable SAN-based approach to enterprise backup and disaster recovery a simple and inexpensive alternative to the cost and complexity of Fibre Channel-based solutions. IP SANs also offer a pre-integrated alternative to the assemble-it-yourself piece-part approach of network attached storage (NAS) vendors.

This white paper describes, step-by-step, how to use a StoneFly Intelligent Networking Platform to set up a LAN-free disk-to-disk and disk-to-tape backup and restore environment for 2-25 servers or more with offsite replication.

Why Now, More than Ever?

Most IT managers have seen first hand that digital data is growing at unprecedented rates and that semi-structured data, such as email, is fueling enormous capacity growth. According to International Data Corporation (IDC), the demand for storage capacity is growing 75% every year. IDC also found that storage demand from email alone had a 300% CAGR (compound annual growth rate) due to the increase in the number of attachments, the trend towards sending rich media, and the general reluctance of users to delete email.

This deluge of digital data is driving the need for reliable, high-performance storage that is easier to implement and manage. These requirements favor network-based, centrally managed storage over alternative storage solutions and infrastructures. Internet Protocol (IP) has become firmly established as the predominant general-purpose networking protocol. The IP storage concept naturally followed as a way to meet growing networked storage requirements.

The second major storage driver is mandatory compliance with the host of new regulatory requirements enacted in the past year. These regulations dramatically impact data retention and storage requirements. These regulations include, but are not limited to:

- HIPPA for insurance, healthcare and medical providers
- SEC 17a-3 and 17a-4 for brokers and dealers in the financial services industry
- The Sarbanes-Oxley Act for public companies and institutions to regulate corporate and public accounting practices to ensure that financial statements are accurate.

This new regulatory environment is personified in the INFOWORLD survey: the second most cited motivating factor for IT managers to buying storage and backup or archiving solutions arises from the desire to save everything in case some obscure new regulation or other legal circumstance dictates that certain documents must, at some future point in time, be retrievable.

The prohibitive cost of using traditional methods for retaining, managing, and disposing of regulatory compliance records, as well as the sheer volume of newly generated records, has resulted in the need for integrated, easy-to-use disaster recovery and backup/restore solutions.

Why New Solutions Are Needed

To date, Fibre Channel SAN deployments have been driven by an insatiable demand for storage at large companies and the user benefits delivered by networked storage. Similarly, NAS vendors and commodity server vendors have been suggesting, with increasing stridence of late, that their technologies can form the basis for a reliable and cost-effective backup and disaster recovery infrastructure. But these solutions suffer from a number of significant drawbacks, including:

Fibre Channel Remains Expensive, Difficult, and Time Consuming to Deploy

Adding a FC SAN requires a second network based on a new, unfamiliar technology. Deploying a Fibre Channel network requires a significant overhaul for existing computing environments. It is also time consuming and adds significant cost for maintaining separate network management systems, training and education for IT personnel, and separate sourcing of new networking equipment.

Fibre Channel Requires Far Too Much Manual Intervention

While standards exist for Fibre Channel technology, the standards are interpreted differently by each vendor, precluding interoperability among vendor products. This leads to manual processes with vendor-specific management tools. Most FC SANs require round the clock operators to ensure that backups complete. But what most small and medium enterprises need is automated solutions that work seamlessly with lights out efficiency, all at a price they can afford.

Software Solutions Alone Are Insufficient

Implementing a simple backup and restore software solution without changing the hardware environment doesn't solve the root infrastructure problem. Since most backup software applications require client servers to be backed up one at a time, scheduling backups during non-peak hours — typically between 8:00 pm to 6:00 am — may not provide sufficient time to back up all data on all servers. Moreover, this scenario is not viable for businesses that operate across multiple or international time zones.

NAS Solutions Offer Too Little

NAS (Network Attached Storage) does not support direct access to block-level data residing on SCSI-attached devices or FC SANs. Rather, such access is only available by routing the connection through a file system. Some newer NAS devices may allow block access by bypassing the file system, but these devices are still optimized for file I/O and not block I/O. An application running on a client that accesses storage on a NAS device must go down the entire seven layers of the networking protocol, across the LAN wire, up the seven layers, fetch the data, and transmit back the data. The steps required to support raw I/O prevent NAS from being an effective solution with data-intensive server applications, such as direct connections to SQL databases and Microsoft Exchange server applications.

The IP SAN Solution

Three trends are making IP SAN-based disaster recovery solutions not only feasible, but also viable for small and medium size enterprises to consider. The first is the availability of inexpensive Serial ATA drives, providing low cost disks that can be employed specifically for backup and restore. Second, the new iSCSI standard (Internet SCSI) allows companies to develop cost-effective IP SANs (Internet Protocol Storage Area Networks) for disk-to-disk backup. Finally, a new generation of backup and restore software supports disk-to-disk backup and restore methods. Additionally, optical storage technologies can be employed instead of tape to meet the new archiving and compliance regulations.




Figure One: A StoneFly IP SAN uses iSCSI to communicate between the application servers and the ATA RAID storage. It operates over standard Gigabit Ethernet infrastructure.

Key advantages of IP-based storage include:

- Its ability to “universalize” storage networking by putting both message/file and storage I/O onto an Ethernet/IP network and ultimately “converging” networking and storage architectures.
- Its natural fit within the prevalent IP/Ethernet infrastructures deployed in today’s business environments.
- Its support for block storage over widely deployed IP-based networks, allowing easy access to storage over long distances.
- Its use of a familiar networking technology and management, with virtually no additional training and support required.

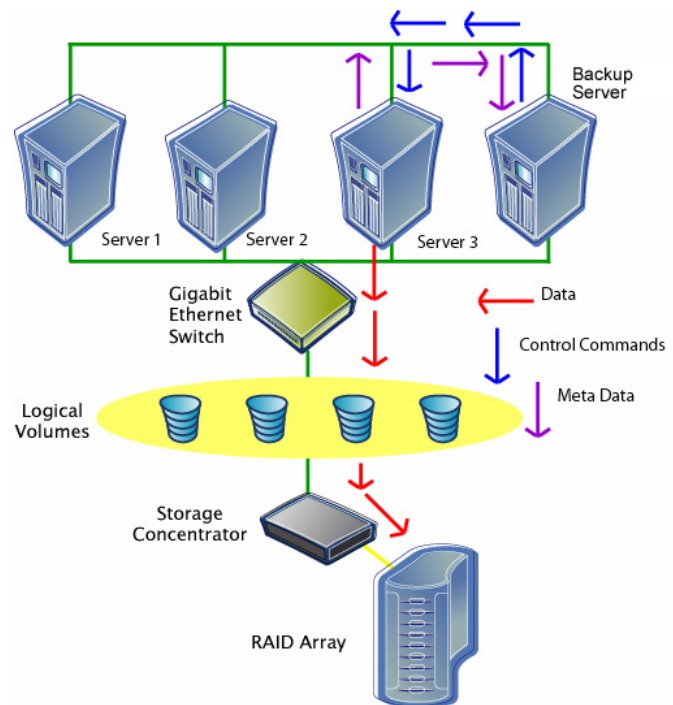
- Its built-in awareness of the impending transition from one gigabit Ethernet to 10 Gbps and beyond, thereby protecting investments with standards-based performance upgrades.

How IP SAN Disk-to-Disk Backup Software Works

Disk-to-Disk backup using an IP SAN resolves the conflict between data backups and shrinking backup windows. Instead of backing up data one server at a time, users can backup data from multiple servers simultaneously onto a central disk storage repository, slashing backup times.

StoneFly Backup Advantage™ (SBA) is a complete disk-to-disk backup solution that provides a cost-effective way of offloading backup and restore operations from the company LAN onto a dedicated StoneFly IP SAN. With this integrated hardware and software approach, the typical constraints imposed by the LAN are removed from the backup process and the burden of backup traffic placed on the IP SAN. See Figure 1.

*Figure Two:
The SBA backup server instructs another server to backup its logical volume onto the IP SAN. With only minimal control commands and meta data required as communication between the two servers, it provides a virtually LAN-free disk-to-disk backup environment.*



In a StoneFly SBA IP SAN configuration, the backup server instructs client servers to send backup data to logical volumes presented by the IP SAN over a dedicated Gigabit Ethernet SAN (See Figure 2). The IP SAN makes these logical volumes appear as a local disk to each client server. The backup operation is a simple disk copy operation that is performed at disk-to-disk speeds. The entire operation is as fast and seamless as writing backup data to a local drive.

Asynchronous Mirroring to a Remote Site

Asynchronous mirroring software supports storage-independent, local or remote, asynchronous data replication over unlimited distances using iSCSI on an Ethernet network. Asynchronous mirroring software used in conjunction with an IP SAN provides an affordable

solution to keep critical data protected and highly available and leverages existing investments in hardware and personnel.

Benefits of Integrated SBA IP SAN for Disk-to-Disk Backup/Restore

- Uses disk for fast, near-term backup and restore with a lower cost of ownership, ease of use, and 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
- Solves longstanding problems with tape-based backup, such as performance, recoverability, and manageability, at a price point well within budget
- 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 management efficiency
- Reduces the impact of tape-related failures

Seven Steps to an IP SAN Disaster Recovery Solution

Step 1: Convert Direct Attached Storage to IP SAN Storage

First, create logical volumes for copying the data on DAS to a StoneFly IP SAN. Then, use Microsoft’s built-in logical volume management or other data migration tools to mirror your old direct-attached storage (DAS) to the logical volumes on the new IP SAN. Once you mirror the data from the DAS onto the networked RAID, break the mirror and move the old DAS behind the *IP SAN* and use it as additional networked storage. Follow the directions below to move the DAS from other application servers onto the repurposed storage.

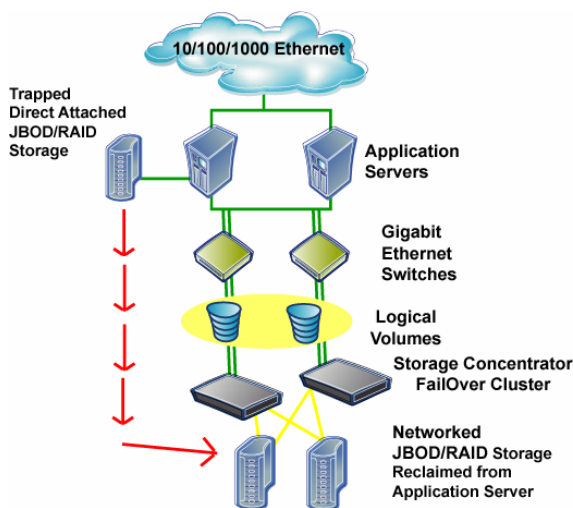


Figure Three: The first step in creating a robust IP SAN Disaster Recovery system is to backup direct attached storage from the servers onto the IP SAN and then move it so it can be reused. Some sites choose to upgrade to new storage in addition to repurposing direct attached storage.

Step 2: Install IP SAN Backup & Restore Software

With the StoneFly Backup Advantage software, backup operations on an IP SAN are enhanced by Ethernet’s inherent ability to support concurrent server data transfers on the network. This intrinsic feature enables multiple client servers to backup data to their own logical disks at the same time, without having to wait for other backup jobs to finish. Running multiple backup jobs simultaneously dramatically reduces the time required to

backup data. Laboratory tests demonstrate nearly 60% timesavings backing up disk-to-disk rather than backing up to tape.

Once the client servers complete their data backups to the IP SAN volumes, the SBA backup server can direct the host to move the data directly from the IP SAN logical volume and place it on tape or an archival media such as WORM or optical drives. This transfer can take place without regard to the backup window since all of the client servers have already completed their respective backups and returned to normal operation.

As an added bonus, the system administrator has complete control over configuring the size and number of logical volumes presented to each client server, thus providing the ultimate flexibility in determining how much and how long backup data remains accessible at disk speeds. For example a system administrator could choose to have as little as yesterday's backups, or as much as last month's daily backups accessible via the IP SAN.

Additional StoneFly Backup Advantage Disk-to-Disk features include:

- If a backup fails for any reason, when the backup is restarted, backup software can pick up at the point of interruption rather than all the way at the beginning, saving time
- An auxiliary copy feature with the ability to backup data from one server to the IP SAN and then, based on policies and without user intervention, create a second (or third) copy of the backup, storing it to another local or remote IP SAN with additional disk or tape storage devices

Step 3: Setup and Test Disk-to-Disk Restores

In order to facilitate the quick recovery of recently backed-up data, the system administrator needs to make choices on how backups are stored on each of those logical volumes. In the rare instance that a client server needs to restore this data, its backup software is pointed to the correct logical volume on the IP SAN to copy the required data. The entire transaction is conducted over a dedicated Gigabit Ethernet storage network using a simple copy operation. This restores the data at disk-to-disk speeds that are as fast and seamless as accessing data from local drives in the client server. If the client requests backup data that is more than what is stored on-line, it can be restored from tape or WORM by the backup server and then sent over the LAN, to the requesting server.

Ethernet's inherent ability to support concurrent server data transfers on the network enables a client server to restore data from its own logical volume without impacting other restore operations that may be occurring at the same time on other servers. SBA backup software is even able to provide message level restores from Microsoft Exchange in just minutes.

As a company's storage requirements increase, the IP SAN allows client servers and their associated logical volumes to be added, without increasing backup or restore times.

Step Four: Set Up Storage Policies

Storage Policies are a fundamental element of the backup software's logical view of storage resources. Through Storage Policies, all key parameters are defined. Such guidelines include where the data is stored, how long it is retained, the number of copies to create, the location of the copies, the retention of the copies and when to make the additional copies.

Once Storage Policies are defined, data is "assigned" to a Storage Policy for management according to those guidelines. It is easy, point and click, to assign data to a different Storage Policy, changing the way the data is managed. There is no reconfiguration of

hardware, any re-cabling or re-networking to accomplish these changes. It is all handled within backup software and is derived from the Storage Policy definitions. This feature reduces the cost and complexity faced in many IT shops when setting up and maintaining the infrastructure and management schema for safeguarding the data.

Key Automation Features of StoneFly Backup Advantage:

- Automated movement of data from disk to tape, no manual intervention required
- Data path traceability enables administrators to more easily discovery and remedy configuration issues for backup, restore, archiving and migration operations
- Auto Discovery of storage devices by software - removes human intervention
 - Less chance of incorrect configuration
- Automatic Default storage policy includes all data - no unprotected data
- The data aging operation is a clean-up routine that deletes data after it exceeds its data retention period. This frees up valuable media for reuse and removes unnecessary data that impacts performance or takes up valuable media space

Step Five: Setup Local (Synchronous Mirroring)

By mirroring business-critical information to a local or remote location, StoneFly Mirroring Solutions provide disaster recovery protection for vital business information that, if lost, could deal a severe financial blow to a company.

Built into a StoneFly Storage Concentrator, StoneFly Reflection™ is designed for IT professionals who manage environments where business operations can suffer severe financial blows if critical datasets are lost or offline for even a moderate period of time. StoneFly Replicator supports backup/restore operations and data migration to campus IP SANs. Mirrored images are identical to each other, and are interchangeable, and entirely transparent to applications. For example, if the primary volume is broken, the application will continue to run from a different mirrored volume in the mirror set.

StoneFly Backup Advantage™ uses a process called Split Mirror Backup based on StoneFly Reflection™ synchronous mirroring technology to automatically backup the data on application servers without taking them offline. In Split Mirror Backup, a mirror image of the volume is detached from a mirror volume and used for the backup operations leaving the production image data on-line. The detached image is mounted to the backup server. When the backup is complete, it is unmounted and then reattached to the StoneFly Reflection mirror volume. StoneFly Reflection automatically rebuilds the image to match the current mirror volume.

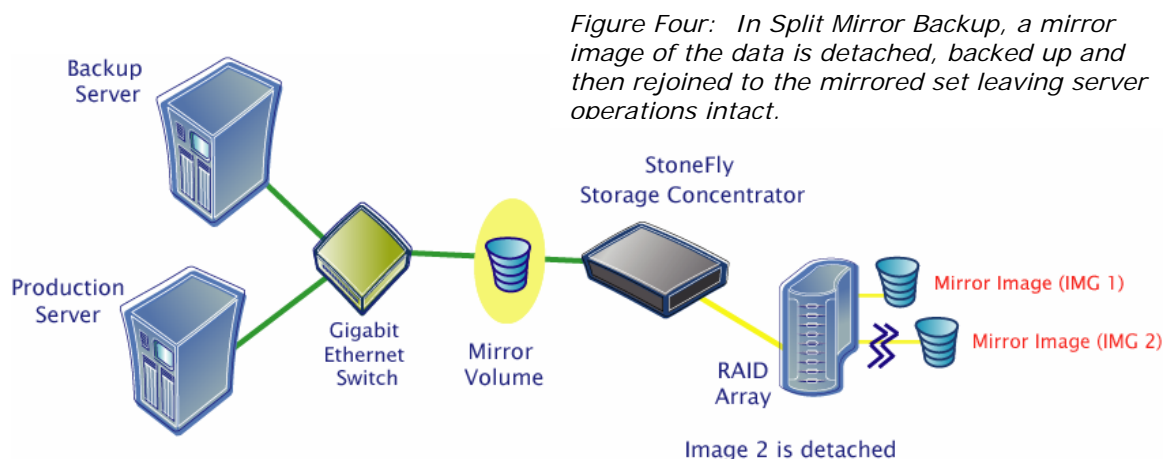


Figure Four: In Split Mirror Backup, a mirror image of the data is detached, backed up and then rejoined to the mirrored set leaving server operations intact.

Step Six: Setup Offsite Storage

Whether you decide to use tape or disk, you need to set up a process for offsite storage of backups also known as tape vaulting. Some companies may prefer to skip to Step Seven, Remote Replication or asynchronous mirroring.

Offsite storage of tapes is as easy as calling a tape vaulting service that picks up copies of backups on a daily or weekly basis and storing them at a facility far enough away that it is outside the immediate disaster area. For Sarbanes-Oxley, this is 200 miles. Most companies prefer the vaulting company to pickup the tapes, but for some, Federal Express will do.

The most important thing is to include a process where the data can be restored at an offsite facility and accessed via the IP SAN. Here is a place where Auxiliary Copy can really provide some real protection, so make sure your backup software supports it.

Step Seven: Set Up Remote Replication

Asynchronous mirroring software supports storage-independent, local or remote, asynchronous data replication over unlimited distances using iSCSI on an Ethernet network. Asynchronous mirroring software used in conjunction with a local and remote IP SAN provides an affordable solution to keep critical data protected and highly available and leverages existing investments in hardware and personnel. The most important thing is to ensure seamless integration between your backup and recovery software and your asynchronous mirroring application.

StoneFly Replicator includes the following features:

- Continuously replicates data to a local or remote site to preserve data integrity
- Synchronization between remote sites for enhanced disaster recovery planning
- If a failure occurs on the primary system, the software must provide immediate access to business-critical data at a remote site
- Recovers automatically from network outages
- Scales to multiple replication sites over any distance
- Enables data migration to an IP SAN
- Supports business continuity planning, protecting against site disaster or loss




Figure Five: Backup volumes are replicated locally for backup and to a remote site where they can be accessed in the event of a disaster.

- Supports Microsoft Exchange, SQL, and other transactional applications, maintaining coherency on primary and secondary data sets
- Maximizes data availability to the source files during replication
- Accelerates data replication to save time and shorten backup windows
- Allows a Quick Sync between mirrors to minimize replication overhead
- Supports application scripts and APIs for automatic operation

Dramatic Cost Savings

IP SANs not only enhance data backup operations, but also they help businesses reduce costs significantly by consolidating existing and new disk storage. IP SANs provide the optimal solution for improving reliability, management and performance. An StoneFly Backup Advantage backup and remote asynchronous mirroring solution delivers all the benefits of a traditional Fibre Channel SAN at a fraction of the cost, complexity and effort. By reducing management overhead, enabling clustering, strategic data protection and disaster recovery efforts, optimizing the use of existing and future storage, and providing higher availability, an IP SAN significantly increases the return on the overall investment.

An IP SAN simply plugs into a business' existing Ethernet infrastructure, without requiring proprietary network equipment or specialized personnel training. With backup operations now performed centrally, businesses can re-deploy IT personnel and eliminate tape drives/libraries on individual servers and at remote locations.

In addition, implementing D2D solutions using an IP SAN enables business to use readily available IP network security technologies, such as firewalls, encryption, and authentication tools, to prevent unauthorized access into the business' storage.