



STONEFLY

Achieving High Availability & Rapid Disaster Recovery in a Microsoft Exchange® IP SAN

April 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.

Microsoft Exchange Requirements for Storage

The traditional approach to IT Networks – a powerful server directly connected to large storage – is insufficient for many Microsoft Exchange implementations. A single-server configuration with DAS storage exposes the entire server as a single point of failure – a software, hardware or storage failure will halt all Microsoft Exchange services, and potentially corrupt the Microsoft Exchange databases. In addition, increasing the size of storage and performing routine storage maintenance requires downtime, inconveniencing users. Improving the performance of DAS with new controllers or drives also requires significant downtime and reconfiguration. The Microsoft Exchange server must also process disk I/O for backup and recovery operations, which impacts performance.


Based on Microsoft's own recommendations, to improve the scalability and manageability of Microsoft Exchange, storage must be independent of the server; this is best achieved by implementing a SAN. In addition, SANs are the only method of adequately meeting the needs of a clustered Microsoft Exchange environment. In the past, Fibre Channel-based SANs have been the traditional model used to meet this need. While Fibre Channel SANs offer the scalability and reliability required for a mission-critical platform, their complexity and steep learning curve generally limit Fibre Channel SAN utility to large data center installations.

An IP SAN is a viable alternative for Exchange that offers many of the same benefits of Fibre Channel-based SANs, but substantially reduces the cost and complexity of maintaining a SAN. For example, an IP SAN based on a StoneFly Storage Concentrator presents a straightforward approach to establishing a Microsoft Exchange storage resource.

Storage Concentrators use the industry standard iSCSI protocol to transmit low-level SCSI disk commands from the server across a standard Ethernet network to storage on the IP SAN network. This allows the server to use the same block file access as DAS, which avoids the delays caused by file-level access that are typical of NAS solutions. In essence, the servers see volumes served up by Storage Concentrators as locally attached storage capacity.

Storage Concentrators reduce the cost of migrating storage to the SAN in several ways. Existing storage is attached directly to the SAN, eliminating the need to purchase new, SAN specific storage. In addition, once the Storage Concentrator is installed, access to existing storage is almost immediate, ensuring minimal down time in the transition. Finally, a Storage Concentrator uses the existing Ethernet networking infrastructure and IT knowledge base.

Figure 1 – A Storage Concentrator enabled SAN



Challenges Facing Microsoft Exchange Administrators

Rapid disaster recovery and high availability remain the most significant issues for implementers of Microsoft Exchange. Due to the mission-critical nature of messaging and collaborative applications, minimizing downtime and providing for rapid recovery are paramount.

Microsoft recommends each mail store reside on a separate disk/volume for performance and capacity reasons. In addition, Microsoft recommends that each volume be more than twice the size of the actual mail store. Exchange mail stores that grow above that 50% threshold may require migration to another larger disk volume to accommodate maintenance and recovery activities.

As the Microsoft Exchange Information Store consumes endless gigabytes, backups become time-consuming, quickly exceeding any window of system availability. Data loss or corruption can trigger a tape-based restore, resulting in lengthy downtimes.

Many administrators attempt to control the growing storage problem by applying mailbox quotas, which actually can make things worse rather than better. Quotas force the end-user to apply archiving to their desktop and/or to personal folder files. However, personal folder files create security problems because they aren't part of the master database and are not subjected to the same security and maintenance routines as the main database. In addition, many IT Managers see these personal archives as lost islands of data which actually use more storage in the whole organization than what ever could be saved by email quotas. Additionally, tracking and finding past emails in archives do not provide as valuable, reliable or easily accessible record of sent or received mail. In addition, many of the new Federal regulations require on-line access to email stores.

A single Microsoft Exchange server presents a single point of failure, encompassing the processor, memory, peripherals and power supply. In order to eliminate this single point of failure, Microsoft Exchange processing is commonly spread between several servers interconnected in a cluster. If one server fails, that server's Microsoft Exchange clients are automatically "failed over" to another designated processor, restoring operations in a matter of minutes. When the server is repaired, its clients can be "failed back" to their original server.

A Microsoft Exchange storage solution must address each of these challenges. Storage must be scalable and transparent to the applications and users, with little impact on IT administration staff. Rock solid reliability for high availability is essential, requiring the storage sub-system to support clustering. Data protection must be bullet proof, quickly recovering from loss or corruption. Fortunately, a solution exists that addresses all of these challenges: this solution is a StoneFly Storage Concentrator-enabled IP SAN.

Storage Concentrators Address Requirements of Exchange

A Storage Concentrator is part of an end-to-end IP SAN solution that is deployed with external storage arrays, offering the functionality of more costly Fibre Channel SANs. Based on industry standard TCP/IP, Ethernet and iSCSI technologies, Storage Concentrators allow existing directly attached storage to be aggregated and centrally managed as an enterprise-wide resource, reducing the cost and complexity of managing growing storage demands. Storage Concentrators provide advanced

storage management, offering lower cost of ownership and significantly reduced complexity than traditional SAN alternatives.

A Storage Concentrator enables centralized storage provisioning, also known as logical volume management. Storage users require capacity, performance and availability – the physical aspects of disk size, number of drives and configuration are irrelevant. Rather than dealing with drives and physical mapping, a Storage Concentrator provides applications and users with one or more logical volumes that appear to the host as directly connected SCSI disk drives. The logical volume points to physical storage areas on disk drives by blocks. The translation of the physical SCSI command to the appropriate logical volume is processed within the Storage Concentrator. Logical volumes free the application administrator from the details of physical storage location and management.

The Storage Concentrator transparently accommodates growth. When storage demands increase, data can be easily moved to larger volumes, and the now unused volume is returned to the storage pool for other applications. Under-utilized storage can be re-allocated to other uses. The Storage Concentrator makes it possible to scale storage for the enterprise as a whole, rather than one server at a time.

Exchange Disaster Recovery Options

It is important to keep in mind that there is not a single solution for Exchange data protection, disaster prevention, or business continuity that fits all corporate requirements. Matching technologies to individual business requirements is not a simple task. Striking the right balance between security, ease of implementation, and cost requires careful thought and a lot of know-how about a wide range of technologies. Here are some options to consider.

Disk-to-Disk Backup

Technology that creates an independent storage resource, such as an IP SAN, enables sophisticated Exchange backup procedures such as disk-to-disk backup and restore. Disk-based backup reduces backup windows by copying from multiple Exchange servers simultaneously onto the dedicated IP SAN, slashing backup times. Disk-to-disk backup, which would significantly impact a directly attached server, provides a high level of data protection and the fastest recovery.¹ Logical volume management allows an administrator to create a separate “volume” for each server, which to the server appears as a locally attached disk. Volumes can be created daily, or more often, if required.

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 Ethernet 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.

¹ For more information on implementing disk-to-disk backup in an IP SAN environment, see whitepapers on StoneFly Solutions at <http://www.stonefly.com/solutions/whitepapers.asp>




Figure 1: SBA manages the backup process and sends client data from DAS or internal disk on host servers to the IP SAN during the normal backup window. Likewise, network logical volumes are also backed up to backup volumes. Aged volumes with client data can subsequently be sent to tape as an automated background activity.

Disk-to-Disk backup using SBA 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.

SBA Benefits

- Provides lower cost of ownership, ease of use, and performance benefits superior to traditional backup solutions
- Minimizes costs and simplifies operations using integrated IP SAN technology
- Takes advantage of high-speed, reliable Disk-to-Disk (D2D) backup and restore technology
- Supports standard iSCSI protocol ensuring integration into standard Ethernet networks for high performance applications
- Offloads LAN-attached backups to the IP SAN creating a "LAN-free" backup environment
- Replicates backups to a remote IP SAN using StoneFly Replicator™, which is included in the turnkey SBA system

FailOver Clustering

Storage Concentrators can be set up in clustered configurations. If a substantial failure takes down an individual Storage Concentrator, FailOver to a standby Storage Concentrator will occur automatically to ensure continuous access to the Microsoft Exchange data store. Currently, a Storage Concentrator FailOver Cluster consists of two units---one is active and one is on standby. A cluster appears as a single entity to hosts on the network.

The Storage Concentrator FailOver Cluster forms the basis of a redundant and reliable IP-based Storage Area Network. The Storage Concentrator FailOver pair

features redundant CPUs, power supplies, hard drives, port connections and operating system. All configuration data from the active Storage Concentrator is replicated to the standby unit for unparalleled data protection. Redundant Storage Concentrators ensure storage volumes are continuously available.

Server Clustering

For mission-critical applications such as Microsoft Exchange, Microsoft's Cluster Service (MSCS) enables clustering at the server. For its Microsoft Exchange databases, each server in the cluster is assigned individual logical volumes on the Storage Concentrator and all servers are provided access rights to each other's logical volumes. If a server fails, MSCS transfers the ownership of that server's storage to a designated fail-over server, which is then able to access the failed server's data, and the affected users are transparently redirected to the fail-over server. Once the failed server is back on line, MSCS returns the ownership of its storage and the affected users are transparently returned to the original server.

Mirroring

Mirroring provides a technique for creating and maintaining identical data sets on different physical IP SANs for data migration to a new SAN or replication for disaster recovery planning. If a disk fails within a mirror, it will have an identical set in a separate IP SAN.

- Local (synchronous) mirroring is performed between two or more Storage Concentrator-based IP SANs within the same storage environment. For example, building A can keep Exchange Databases in an IP SAN synchronized with a separate IP SAN in building B.
- Remote (asynchronous) mirroring is performed between two or more IP SANs in separate (remote) storage environments across a WAN connection. The remote site typically acts as a disaster recover configuration for the primary location.

StoneFly Reflection™

StoneFly Reflection synchronous mirroring provides high speed Storage Concentrator-based mirroring to a local or campus site. This effectively moves 'intelligence to the network core' by consolidating disk management and mirroring functions at the Concentrator. This offloads these functions from servers and localizes mirroring traffic on the IP SAN.


StoneFly Reflection provides continuous access to a volume without interruption to data availability by substituting an exact mirror image, simplifying disaster recovery and reducing system downtime. StoneFly Reflection simplifies deployment, operating transparently to users, applications, databases, and host processors. The result is "Next Generation" performance, reliability and high data availability for critical applications.

StoneFly Replicator™

StoneFly Replicator provides an affordable data replication solution that allows mirroring between any model of storage device, providing local or remote (synchronous or asynchronous) replication of data between IP SANs over unlimited distances for increased data availability and disaster recovery planning. Used in conjunction with StoneFly's Storage Concentrator, Replicator offloads the mirroring

process from the storage subsystem, allowing storage-independent replication of Exchange data from any storage device to any storage device, at any location, either local or remote, at the block level over an iSCSI network. StoneFly Replicator benefits include:

- Preservation of data integrity
- Enhanced disaster recovery planning
- Immediate access to business-critical data at a remote site
- Accelerated backup windows
- Consolidated backups and restores from multiple servers
- Data exchange and synchronization between sites



StoneFly Replicator mirrors Exchange data to IP SAN Volumes, across any available network connection supporting TCP/IP and iSCSI. Data is duplicated in near real-time to ensure integrity between the two systems in the event of hardware failure, natural disaster, or human intervention. StoneFly Replicator accomplishes this through near real-time data transfers from the Source Volumes to the Target Volumes on the Storage Concentrator. Should a failure occur on the Source Volumes, the Target Volumes can provide immediate access to business-critical data.

For more information on implementing an IP SAN in a Microsoft Exchange environment, see whitepapers on StoneFly Solutions at <http://www.stonefly.com/solutions/whitepapers.asp>.