



# STONEFLY

## **Implementing a Storage Concentrator™ in a Microsoft Exchange® Network**

**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.

## **Implementing a Storage Concentrator in a Microsoft Exchange Network**

### **Executive Overview**

Microsoft Exchange presents a unique storage challenge to IT management. The ever-increasing volume of e-mail, and the need to store messages long term, presents a growing demand for storage – even if the organization's size and revenues remain constant.

E-mail has become the preferred method of saving and storing files and crucial data. From a storage standpoint, IDC found that storage demand from email had a 300% CAGR (compound annual growth rate) due to the increase in the increasing number of attachments, the trend towards sending rich media, and the general reluctance of users to delete email. To accommodate this compound growth, Microsoft Exchange storage must be frequently enlarged, requiring downtime. Furthermore, the demand for resource administration and management including, user administration, directory synchronization, database maintenance, and server management, also grows.

Developing the right storage infrastructure is a strategic business opportunity. The traditional approach of deploying large servers each with direct attached storage (DAS) is no longer feasible because this architecture does not provide the scalability and reliability required for mission-critical operations such as Microsoft Exchange.

Traditional Fibre Channel SANs (Storage Area Networks) and NAS (Network Attached Storage) solutions address certain challenges. However, the cost and complexity of Fibre Channel SANs can often delay any return on investment. The file system structure of NAS systems introduces performance issues. Microsoft does not recommend NAS systems for Microsoft Exchange environments unless block level services are layered on top of the file system structure. This creates additional performance issues.

New storage networking solutions are required to effect a significant reduction in the cost of installing and maintaining SANs, while establishing a foundation for managing growth at an affordable cost. This paper introduces an innovative storage networking solution from StoneFly that is simple, affordable and easily integrated into an existing Ethernet network.

### **New Advances in Storage Technology**

The traditional approach to information technology – 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 achieved by implementing a SAN. In addition, SANs are the only method of

## Implementing a Storage Concentrator in a Microsoft Exchange Network

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. Now, the IP SAN is a viable alternative that offers many of the same benefits of fibre channel-based SANs, but in addition, the IP SAN substantially reduces the cost and complexity of maintaining a SAN.

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. Implementing a SAN based on Fibre Channel architecture solely for Microsoft Exchange storage creates a much greater impact on IT management than the resulting relief offered by simplifying Microsoft Exchange storage management. This is because Fibre Channel SANs require the purchase and implementation of a new network along with extensive training or hiring of personnel experienced in fibre channel implementation and management.

Relocating storage to a NAS system is often considered as an alternative to SANs, but this actually increases the burden on the server's processor. The demand for data stored on NAS volumes by I/O intensive applications such as Microsoft Exchange increases server file access overhead resulting in incremental performance degradation. NAS systems are accessed at the file level, requiring that all disk operations transverse both file and network operating system drivers on both the server and the NAS system, increasing the risk of data corruption and performance decay.

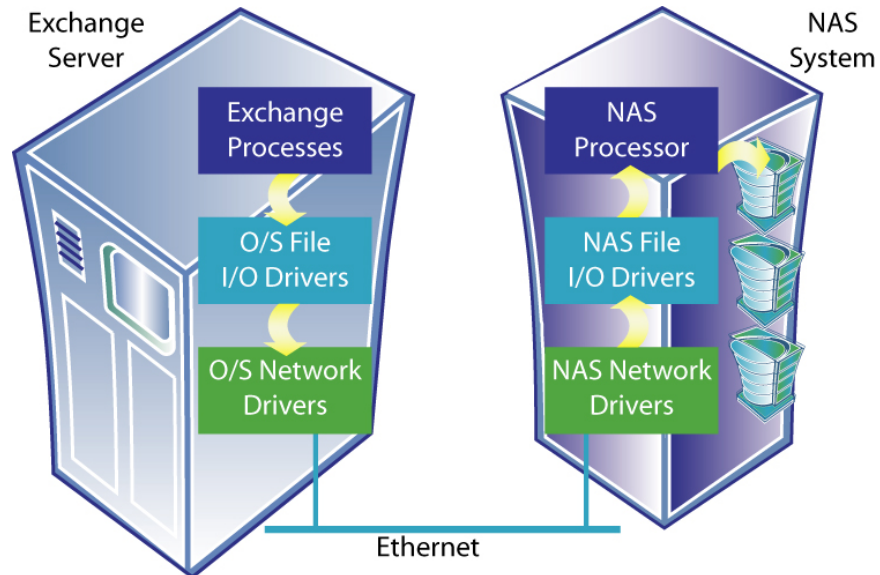


Figure 1 – NAS File Access Overhead

## Implementing a Storage Concentrator in a Microsoft Exchange Network

Microsoft Exchange also has local data access and I/O bandwidth requirements that NAS products do not generally meet. Microsoft Exchange uses an Installable File System (IFS) driver, which requires access to physical disk characteristics that are reported back by block mode storage devices. Previous versions of Microsoft Exchange did not include an IFS driver and did not require block mode storage. As a workaround, some NAS systems replace the e-mail message in the local Microsoft Exchange database with a link to the actual data on the NAS product, but I/O bandwidth is still limited.




Figure 2 – A Storage Concentrator enabled SAN

A Storage Concentrator-enabled IP SAN presents the most 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 stand-alone storage. 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.

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 downtime in the transition. Finally, a Storage Concentrator uses the existing Ethernet networking infrastructure and IT knowledge base, eliminating the need to install and train for a fibre channel network.

## The Challenges Facing Microsoft Exchange Storage Administrators

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.

## **Implementing a Storage Concentrator in a Microsoft Exchange Network**

Microsoft recommends each mail store reside on a separate disk/volume for performance and capacity reasons. In addition, Microsoft recommends that each volume is 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.

The 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 must be spread amongst 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.

One way to attempt to control the growing storage problem is to apply mailbox quotas, which actually can make things worse than better. Quotas force the end-user to apply archiving to their desktop and/or to personal folder files. Microsoft Outlook, for example, sees these files as ".PST" files. IT administrators may encourage their staff to use .PST files to reduce required storage capacity. However, .PST files create security problems because they aren't part of the master database and are not subjected to the same 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. Archived files may even take up more space than since Exchange can store. Additionally, tracking and finding past emails in archives do not provide as valuable, reliable or easily accessible record of sent or received mail.

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 Concentrator Technology**

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.

## **Implementing a Storage Concentrator in a Microsoft Exchange Network**

A Storage Concentrator enables 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.

Technology that creates an independent storage resource, such as a Storage Concentrator, enables sophisticated backup procedures such as disk-to-disk backup and restore. Backup and recovery no longer cause performance degradation of the server as they are processed by other systems. Disk-to-disk backup, which would significantly impact a directly attached server, provides a high level of data protection and the fastest recovery.

A single Storage Concentrator is a highly available device, with redundant critical components that include hot-swappable power supplies and fans, multiple power connections, and mirrored operating system drives. If any redundant component fails, it can be replaced while the system remains online; this provides uninterrupted service to the Microsoft Exchange system. For higher levels of availability, Storage Concentrators can be set up in clustered configurations. If a substantial failure takes down an individual Storage Concentrator, fail-over to a standby Storage Concentrator will occur automatically to ensure continuous access to the Microsoft Exchange data store.

For mission-critical applications such as Microsoft Exchange, the Storage Concentrator uses Microsoft's Cluster Service (MSCS) to enable clustering at the host. For its Microsoft Exchange databases, each host 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.

## Implementing a Storage Concentrator in a Microsoft Exchange Network

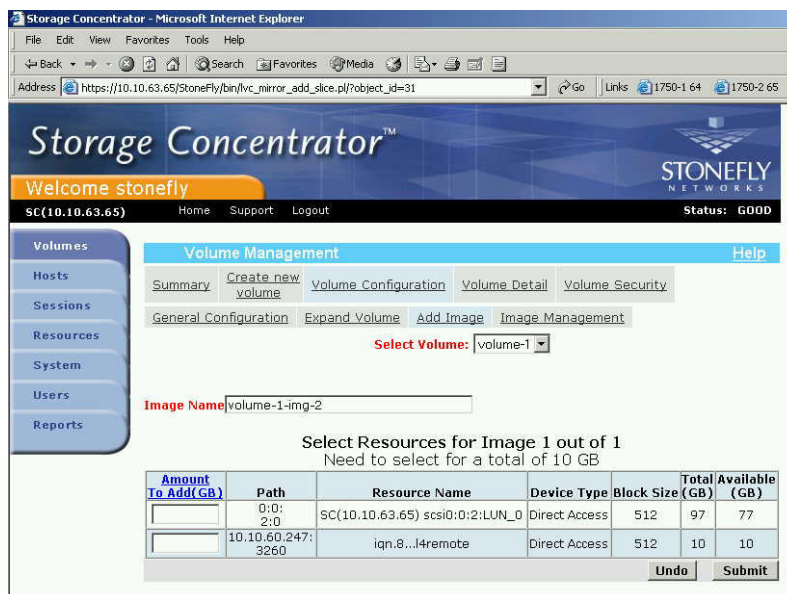


Figure 3 – Sample screen from the management GUI

The Storage Concentrator may be managed from any system attached to the Ethernet network via a secure and easy-to-use graphical user interface (GUI). Initial setup and configuration, performance optimization, and storage management are all accomplished with just a few mouse clicks. Simplifying storage management by consolidating the storage of multiple arrays managed by multiple administrators into a centrally administered resource frees administrators to be available for other tasks.

## Implementing a Storage Concentrator in a Microsoft Exchange Environment

Migrating Microsoft Exchange storage to a Storage Concentrator is a simple process and can be accomplished in one session. Most of the effort is dedicated to installing and configuring the Storage Concentrator – relocating the Microsoft Exchange Information Store is the final step, and is easily accomplished in a matter of minutes. For more information about migrating an existing Microsoft Exchange configuration to an IP SAN, see StoneFly's *Migrating Microsoft Exchange Data to a Storage Concentrator*.

The Microsoft Exchange server communicates with the IP SAN through an iSCSI interface, called an initiator, which enables the server to send and receive SCSI commands over an Ethernet network. The simplest way to enable a server to communicate with an IP SAN is to load a software-based iSCSI driver that communicates over the server's existing network interface; this implementation is appropriate for servers that have excess processing power. The server's excess processing power is used to process translation of SCSI commands into TCP/IP commands that can traverse the network.

In situations where the server does not have excess processing power, or where the amount of data traffic is expected to be very heavy, cards that act as initiators and offload the translation processes from the server's CPU(s), can be installed for better overall performance. There are two categories of cards: TCP/IP Offload Engines

## **Implementing a Storage Concentrator in a Microsoft Exchange Network**

(TOEs) and Storage Network Interface Cards (SNICs). Both of these technologies provide significant benefits in situations where CPU resources are at a premium.

For increased resiliency, the server can be configured with multiple network interfaces. In the event of a failure of any link, both the normal network traffic and the iSCSI storage traffic will continue to function across the other network interfaces in the server.

For the highest availability, multiple servers may be configured as a cluster – a group of servers managed as a single system. Should a server fail, its resources and workload are automatically transferred to a designated fail-over server allowing processing to resume with little user inconvenience. A cluster allows IT management to move resources and work between the servers, enabling load balancing. Software upgrades and maintenance can be performed one server at a time, known as a “rolling update”, as the other servers temporarily assume the off-line system’s workload.

MSCS organizes a server’s virtual server name, network address, storage and other resources into “resource groups”. The on-line and off-line state, as well as server ownership of resource groups can be dynamically controlled. The logical volumes managed by the Storage Concentrator are assigned to MSCS resource groups, along with the server’s IP addresses and other resources. Each resource group is then assigned a “preferred owner” server for normal operation and “possible owner” servers for fail-over. In the event of a failure, or the need to bring a server off-line for maintenance, MSCS will automatically transfer its resources, including the Storage Concentrator logical volumes, and workload to one of the “possible owner” servers.

Once the Storage Concentrator is installed and verified as operational, the Microsoft Exchange data can be migrated to IP SAN volumes in one of two ways. The StoneFly Direct Path™ feature can be used to physically move and attach Microsoft Exchange data to storage controllers that are connected to the Storage Concentrator. This procedure does not require data migration and offers almost instantaneous access to the Microsoft Exchange volumes that have been relocated onto the IP SAN. Alternatively, the Microsoft Exchange data can be migrated from its existing location to a logical volume located on new storage that is attached to the Storage Concentrator.

The Storage Concentrator uses secure Access Control Lists to determine which administrators have access to each logical storage volume. It is critical that new administrator accounts and passwords be established at this stage to ensure security. Additional storage management accounts may be created for other administration tasks at a later time.

Given the critical importance of Microsoft Exchange data to most organizations, it is highly recommended that the Microsoft Exchange data stores be placed on volumes configured with appropriate RAID capabilities; this provides a higher level of resiliency in the case of a disk failure. Microsoft Exchange data stores fall into two categories: random access data stores that include databases and mailboxes and sequentially accessed data stores that include transaction logs. Optimal performance can be achieved by locating the random access data stores on volumes with drives configured at RAID level 5 and locating the sequentially accessed data stores on volumes with drives configured at RAID level 0+1.

## **Implementing a Storage Concentrator in a Microsoft Exchange Network**

Microsoft Exchange supports up to four storage groups. Each group maintains its own transaction logs for disaster recovery, and can support up to five databases. Generally each storage group should be located on separate storage arrays to isolate faults and optimize performance. For clustered server implementations, one of the five available storage groups is used to support the cluster, leaving four storage groups available.

The final step, configuring Microsoft Exchange to use the new storage, requires only a few mouse clicks. Within the Microsoft Exchange System Manager tool, simply move each storage group, transaction log, and mailbox store to point to the newly created virtual drives on the Storage Concentrator. Exchange services may then be restarted.

## **Exchange Performance on a Storage Concentrator IP SAN**

StoneFly tested Microsoft 2003 Exchange Server under simulated email traffic loads using Microsoft LOADSIM, first with directly attached storage (DAS) connected via Ultra320 SCSI, and then compared with running the same storage subsystem accessed over a IP SAN powered by a Storage Concentrator.

### ***Test Environment***

The Microsoft Exchange server used in the lab tests was from a leading manufacturer of industry-standard business servers. The server configuration used was based upon the manufacturer's recommended configuration for installations with five hundred (500) user mailboxes. The Exchange server had two Pentium 4 Xeon processors, 1 GB of RAM and an on-board Ultra320 SCSI storage interface. An Intel Pro/1000 T Storage Adapter Gigabit Ethernet storage network interface card (SNIC) was used as an iSCSI initiator during the IP SAN portion of the test. The storage used was a DNFstorage Enterprise 12322 Ultra320 SCSI-attached Serial ATA RAID sub-system with twelve (12) 250GB, 7200 RPM SATA drives. Refer to Exhibit I for a more complete description of the test environment employed for this study.

### ***Test Methodology***

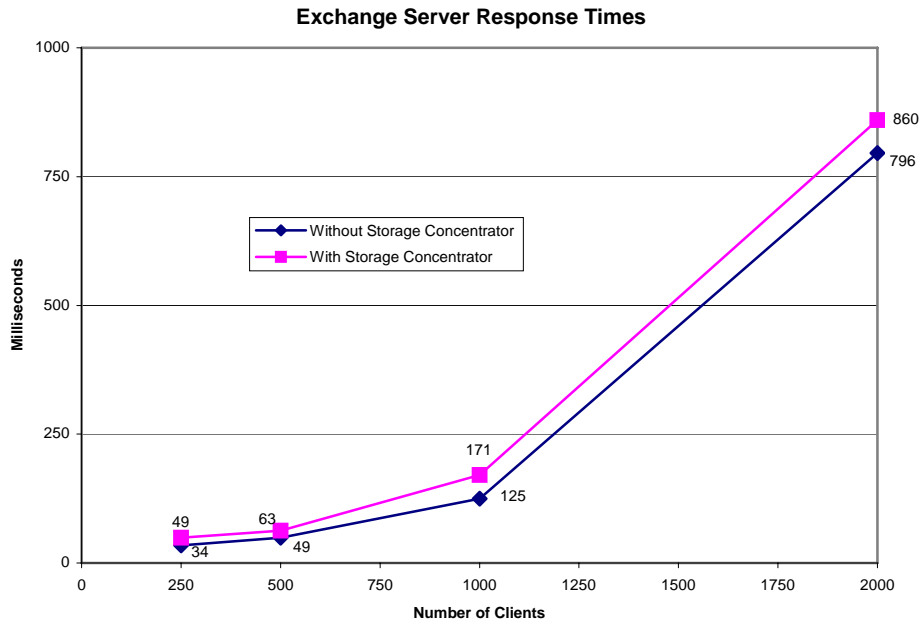
Microsoft's LOADSIM program was used to simulate mail messaging traffic for different numbers of simulated users (MMB2 rating) ranging from two hundred fifty (250) to two thousand (2000) mail boxes. As indicated in the LOADSIM documentation, the MMB2 number of simulated mailboxes does not correlate directly with the number of actual mail users that a particular configuration can support. However, the tests simulated the traffic for a relevant range of users for the Exchange server configuration used in the test as well as the number of mailboxes typically deployed in a medium-size enterprise or workgroup. Microsoft recommends that access times do not exceed 950 milliseconds (the 95<sup>th</sup> percentile of 1 second). The Storage Concentrator performed within the range recommended by Microsoft for all tests.

Measurements were first taken in a direct attached storage (DAS) configuration with the DNFstorage RAID system connected directly to the Exchange server via parallel SCSI. Keeping all other variables constant, the tests were then repeated with the same DNFstorage RAID system connected to the Ultra320 SCSI connection on the Storage Concentrator. The Exchange server was connected to the Storage

## Implementing a Storage Concentrator in a Microsoft Exchange Network

Concentrator using standard GbE and Intel's SNIC. Exchange server volumes were carved out of the 3TB storage pool. Separate virtual volumes were used for the storage group data files and transaction log files as recommended by Microsoft.

### Test Results



The Storage Concentrator-based IP SAN performed well under simulated mail traffic load and delivered access times very close to that of directly attached storage, a margin that would not be perceptible by clients using Exchange mail services.

CPU Utilization of the Storage Concentrator and the bandwidth utilization of the Gigabit Ethernet links were monitored throughout the tests. The peak bandwidth utilization for all tests was 10%. Similarly, the peak CPU utilization for all tests was just under 35%. Average utilizations were considerably lower.

These results indicate that one Storage Concentrator can easily support the Microsoft Exchange server storage needs for a medium-size company or a workgroup within a large enterprise while still leaving significant system resources available to serve the storage needs of multiple additional servers on the IP SAN.

### Value Proposition

Storage Area Networks enhance the Microsoft Exchange environment by consolidating storage into a centrally managed pool of resources. This creates a central storage repository that facilitates powerful information sharing, protection, and management across the enterprise. Unlike Fibre Channel SANs, iSCSI SANs use existing company resources, including existing Ethernet infrastructures and network management tools, and works with existing staff that is already trained and

## **Implementing a Storage Concentrator in a Microsoft Exchange Network**

comfortable implementing new Ethernet network infrastructures. In addition, because of the proliferation of Ethernet components, IP SANs have a significant up front cost advantage over Fibre Channel networks. This cost advantage is even more significant in disaster-tolerant and disaster-recovery scenarios where remote sites can be set up quickly and easily without the need for expensive translation devices between remote IP SANs.

Storage Concentrators extend the benefit of IP SANs by moving many of the host-based storage management and advanced storage system capabilities into the network. This intelligent managed network storage architecture simplifies implementation and ongoing maintenance, and increases the return on investment over traditional SAN architectures.

Microsoft Exchange platforms that use this architecture deliver enterprise-level availability, scalability, and management to customers that could not justify the investment in traditional fibre channel Storage Area Networks.

## **Support Commitment**

StoneFly is committed to maximizing each customer's return on investment in StoneFly solutions. Recognizing the critical nature of Microsoft Exchange downtime, StoneFly provides 24 hour, seven day a week telephone support for Storage Concentrator solutions, staffed by IT professionals experienced in deployment and optimization. Should telephone assisted troubleshooting indicate a component failure, an experienced technician is immediately dispatched to arrive onsite within four hours. Customer replaceable parts are couriered within four hours from one of over eighty local parts depots located nationwide.

### **Conclusion**

StoneFly Storage Concentrators provide the optimal solution for improving Microsoft Exchange reliability, management and performance. A Storage Concentrator enabled Microsoft Exchange IP SAN delivers all the benefits of a traditional Fibre Channel SAN at a fraction of the cost, complexity and effort. By reducing Microsoft Exchange management overhead, enabling clustering, strategic data protection and disaster recovery efforts, optimizing the use of existing and future storage, and providing higher availability, a Storage Concentrator significantly increases the return on the overall messaging application investment.