THE CLIPPER GROUP
Explorer™

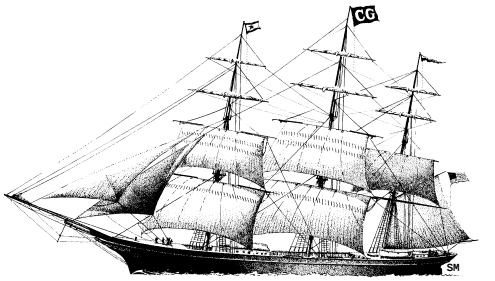

Navigating Information Technology Horizons℠

Published Since 1996 Report #TCG2003064 November 27, 2003

Disk-Based Replication Enhances Resiliency For Microsoft Exchange

Analyst: Michael Fisch

Management Summary

A parachutist is hurtling toward the earth at 124 miles per hour, pulls the rip cord, and – nothing happens. The chute does not open. Meanwhile, the ground below draws nearer by the second. He reaches to open the reserve chute and, thankfully, it quickly deploys. The parachutist floats harmlessly back to earth, though somewhat shaken from the close call. *That's disaster recovery.*

Now, picture a *Microsoft Exchange* server supporting hundreds or thousands of enterprise e-mail users on a busy workday. The database becomes corrupted and the application goes down. E-mail stops and, with it, perhaps 80+% of internal and external communication for the affected users. Management is not happy. You think that you have a reserve parachute to recover from the disaster, and you turn to your tape backup of the database, which is already stored off site. It includes the last full backup plus daily incrementals until the previous day. The IT staff loads each tape, rewinds it to the appropriate spot, and transfers the data to disk. During the process, they discover that one of the incremental tapes had failed, requiring them to spend additional time replaying transaction logs to bring the database up to present. It is into the evening by now, and the other workers have gone home. Their e-mail will be available again when they return in the morning. Result: one day lost!

Is that disaster recovery? No, that's just a disaster. **E-mail is too important, and too closely tied to business productivity to be subject to this level of uncertainty.** For many workers, e-mail is more critical than the telephone. What enterprise would accept telephone outages lasting hours or days? That would be like a backup parachute that takes 30 minutes to open – too slow for what the situation requires. And the same urgency applies to email.

Compounding the situation are rapid e-mail growth and the trend to consolidate and migrate Exchange onto fewer servers. This is akin to putting one's eggs in fewer baskets, while the number of eggs continues to multiply. Therefore, Exchange has to be resilient and scalable.

Smart storage and data management are an essential part of delivering this requirement. **In particular, disk-based replication plays a critical role for maintaining appropriate availability and recovery times.** Point-in-time disk copies (or snapshots) allow non-disruptive backup and fast recovery from data corruption or deletion. Furthermore, backing up to disk/recovering from disk, and archiving to tape can deliver fast recovery from major system crashes while preserving the cost-effectiveness of tape for long-term storage. Read on for details.

IN THIS ISSUE

Vital Exchange

A discussion about *Microsoft Exchange* begins at the desk of an enterprise worker. It could be anyone – a sales representative, production line manager, or senior executive. To do his or her job, the worker must be in continuous contact with other employees, customers, suppliers, and partners. If the flow of communication stops, then productive activity slows down, frustration rises, and opportunities are lost. This link to the rest of the enterprise and broader industry is vital. The preferred method of communication in times past may have been the postal service, wire service, or telephone, but *e-mail* has become that link for many workers today. So, it needs to be reliable. **Occasional downtime or even data loss is not acceptable anymore because e-mail is mission-critical.**

Dynamic E-mail Demands

The challenge is to attain and maintain a high level of resiliency in the dynamic environment of enterprise e-mail. Rising message volumes, large attachments (e.g., graphics, video, audio), spam, and viruses place an ever-increasing demand on computer processing and storage. Scalability and management complexities are a given.

Enterprises have been turning to Exchange consolidation and migration to cope. *Consolidation* implies hosting the application on fewer, larger servers and storage systems and reducing the number of physical sites to improve economies of scale. *Migration* means upgrading from an earlier version to *Exchange 2000* or *2003*[1] to take advantage of more advanced features and consolidation capabilities. These efforts can deliver a streamlined operation that is easier to manage, uses resources more efficiently, and is more cost-effective, especially as it scales – that is, simpler to administer *at scale* and faster to provision new users and resources in order *to scale*.

Focus on Storage

Storage is a key component of the consolidation and migration process. Application software, servers, and networking equipment are also part of a consolidated infrastructure. But storage is as important – and maybe more so – because it delivers a large portion of the benefit and risk mitigation. Much of the simplification in a consolidated environment is found in storage and data management. Better utilization of storage infrastructure is a major part of cost reduction. Furthermore, storage has a large and direct impact on the level of Exchange availability and performance experienced by users.

Storage should be separated from servers and consolidated in its own right. Moving from dedicated internal or direct-attach storage (DAS) to consolidated, networked storage delivers a similar set of simplification, utilization, service level, and cost benefits as server consolidation.[2] It decouples servers and storage, so they can scale and be deployed independently of one another. No more upgrading servers because of a lack of storage capacity, or vice-versa. It also speeds and facilitates the migration process itself. So, storage ought to be a distinct and strategic infrastructure decision, not an afterthought or a component that is lumped together with servers.

As such, an enterprise should manage storage and the data itself for high availability. Consolidation naturally results in Exchange depending on a fewer number of servers, storage systems, databases, and so forth. A system failure or data corruption incident has the potential to cause downtime for even more users. Larger databases take longer for tape backup and restore. Therefore, enterprises should manage the consolidated assets properly for resiliency, to meet business expectations. It requires that the right technologies to be in place, and many of these are in the storage arena.

[1] In particular, Exchange 2003 can allow administrators to reduce the number of sites running Exchange servers without impacting user performance.

[2] See *Networked Storage – A Buyer's Guide to Pain Relief* in **The Clipper Group Explorer** dated April 25, 2003, at www.clipper.com/research/TCG2003017.pdf.

Disk-Based Replication for Resiliency

In particular, disk-based data replication is an increasingly prominent tool for delivering resiliency. With Exchange or any IT system, resiliency is a function of the three R's: *redundancy, remoteness*, and *recoverability*. *Redundancy* implies extra components or systems to fall back upon, like a reserve parachute. *Remoteness* adds distance to protect from local disasters, and *recoverability* describes how quickly and completely a system can return to normalcy when a failure occurs. (*See sidebar on the right.*) Disk-based replication can enhance resiliency in all three dimensions.

Point-in-Time Copy

Point-in-time (PIT) copy is a feature that takes a "snapshot" of data on disk at an instance in time.[3] The copy can be an exact replica of a volume, which is also known as a *clone*. Alternatively, it can be a copy of an index of data locations, which is called a *differential copy*. As the original data is overwritten, a differential copy makes a copy of the modifications to preserve the integrity of the original (i.e., copy-on-write). A clone offers greater flexibility and performance for reading and writing, while a differential copy conserves disk space, as long as data does not change too much for the duration of the copy.

A PIT copy happens instantaneously and, if properly integrated, without disrupting the application. Since it is a disk copy rather than to tape, access time and transfer speeds are both very fast. These copies can be used for recovery to a prior point in time and/or for staging data for other purposes, including backup. The use of these redundant copies confers several advantages:

- **Rapid restore** – If an Exchange database becomes corrupted or files accidentally deleted, a system can quickly revert to a prior "clean" copy on disk and use transaction logs to rebuild the data back to the present. Recovery time is fast and basically independent of database size.
- **More frequent replication** – PIT copies can be performed on a frequent basis, such as hourly. In case of a recovery, the

The Three R's of Resiliency

For enterprise IT systems and data, resiliency or high availability is enabled by three critical characteristics:

- **Redundancy** – Like the reserve parachute, redundant hardware components and multiple copies of data provide a means for fast or instantaneous recovery. Depending on the system requirement, redundant configurations can mean N+1 (i.e., one more than the minimum needed to function), dual, or many.
- **Remoteness** – Remote isolation of backup hardware and data from the primary system adds distance to the equation and protects against local disasters like floods or fires. The primary and secondary systems can be located at different sites on a campus, within a metro area or region, cross-country, or even between countries.
- **Recoverability** – This ability to recover an application quickly and continue operations after a system failure, operator error, data corruption, or other disaster is critical. The degree of restoration (i.e., recovery point objective) and allowable time (i.e., recovery time objective) varies by the business process, application, and available budget.

The key tradeoff in implementing greater resiliency is cost. The more redundant components, the farther the remote distance, and the faster the recovery time – the greater the financial investment. An enterprise must determine the relative importance of its applications to ongoing operations and what resources it is willing to deploy to ensure high availability.

[3] As opposed to a *mirror copy*, which is continuously updated in real-time (synchronous) or near real-time (asynchronous).

system would not need to go as far back in time for a clean copy, as it would for a traditional tape backup.

- **Non-disruptive backup** – A PIT copy can be used to stage data for backup, preferably over a SAN to a backup server. This means that backups can occur during business hours without taking Exchange offline, resulting in no more off-hours staffing or backup windows that conflict with business operations, and it also takes backup traffic off the LAN. A copy feature that resides off of the host, such as in a storage array, can avoid degrading application performance. Moreover, the system can perform a consistency check to ensure that the copy is good before backing it up.
- **Expedited testing and development** – IT can use PIT copies for application testing and development with real production data. They are also useful to stage data for data mining.

Backup to Disk

Backup to disk is also gaining popularity because of its advantages over tape in speed and media reliability. The high bandwidth of a RAID array and its ability to read and write random, intermittent streams of data make it favorable for fast backups and restores. Backup to disk can also provide remoteness if the array is located at a remote site and connected to the primary data center over a wide-area link. While both tape and disk media can fail, a RAID array, by its nature, stores data redundantly and knows immediately when a failure occurs. Some can even predict failures in advance by monitoring for abnormal drive behaviors. With tapes, however, one does not necessarily know a media failure exists until the tape is needed for restore – not a good time to find out. The recent use of ATA disks in enterprise storage arrays also offers a significantly lower cost point for disk storage than high-performance Fibre Channel or SCSI drives. However, tape still delivers the lowest cost per unit of capacity, especially when storing very large quantities of data. **Therefore, many advocate a**

Volume Shadow Copy Services

Microsoft *Windows Server 2003* has a new feature called *Volume Shadow Copy Services* (VSS) that facilitates the creation and use of PIT copies for backup, recovery, and other purposes. It is application-aware and allows copies to be made non-disruptively while an application and its files are in use. It supports both clones and differential copies. The VSS feature consists of:

- **VSS Coordinator** – The service itself that handles interaction among the other components.
- **Writers** – Plug-ins for specific production applications that ensure the integrity and consistency of a copy. Exchange 2003 is a supported application.
- **Requestors** – Data management applications such as backup that utilize the application-aware copies.
- **Providers** – The facility that actually performs the PIT copy. Windows 2003 comes with a software provider for differential copies that runs on the server itself. It can make a copy for backup purposes, though the backup process still must run on the application server, which impacts performance. This copy also cannot be used directly for a system-level restore, nor can it be exported to another server for purposes like testing or data mining. However, these kinds of advanced features are available with third-party hardware providers that run on storage arrays or network devices. Microsoft is working with a number of vendors to integrate their hardware providers with VSS.

If you are migrating to Exchange 2003 and want to take full advantage of PIT copy capabilities, look for data management software and hardware providers that specifically support VSS.

combination of disk and tape for backup. Initial backups are stored on disk for quick recovery, and tape is used for archiving and possibly remote storage.

Look Like a Hero

PIT copies and disk backup are not exclusive of one another. They can work together synergistically to deliver a level of availability and recoverability that neither could by itself. In any case, the smart use of disk-based replication will enhance the resiliency of Exchange. Business managers concerned about productivity and e-mail users frustrated by downtime will be pleased with the results. **And IT administrators can look like heroes – or at least avoid the blame for inordinate downtime of this critical communications tool.**

Solution Checklist

If you are interested in the additional resiliency afforded by disk-based replication, the next question is what to look for in a solution for an Exchange environment. Backup applications can be configured to support disk as a target, though some are easier and more flexible than others.[4] You will want to ask the vendor about how it is performed. For a backup application, the interface with Exchange occurs at the application level, so check to see if it has a specific integration package for Exchange.

As for PIT copy, there are a number of characteristics to consider. Some may be more important than others, depending on your particular needs:

- **Integration with Exchange** – An application-aware solution can make consistent copies of right set of databases and files with virtually no impact on Exchange availability. It is also easier for e-mail administrators and non-storage experts to use the solution, if it is well integrated into the Microsoft management environment. If you have or will migrate to Exchange 2003, look for integration with VSS. (*See sidebar on the previous page.*)
- **Offloading application server** – Copy functionality that runs in a storage array or network-resident platform avoids negatively impacting the performance of the Exchange server.
- **Ease of deployment** – An automated installation utility or professional services installation makes it easier to deploy.
- **Automated operation** – The ability to automatically make PIT copies according to a programmed schedule simplifies administration.
- **Comprehensive support** – Customer support that includes not only the copy software but also its integration with Exchange helps avoid finger-pointing among vendors if something goes awry. A cooperative support agreement with Microsoft also would be helpful.
- **Less expensive disks for replicas –** PIT copies do not necessarily require the performance or robustness of Fibre Channel or SCSI disks. If the functionality is array-based, support for ATA disks in the array can lower costs.
- **Reasonable pricing** – You get what you pay for, of course, but make sure the value is there.

Conclusion

As you engage in activities like consolidation and migration to keep up with rising e-mail demands, be sure to build resiliency into your Exchange solution. **Storage is a key part of this effort, and disk-based replication can make a big difference in the availability of e-mail to your enterprise. For many, it would be a smart investment in user satisfaction and business productivity.**

[4] This applies to general-purpose RAID arrays. Another option is a virtual tape library, which is basically a RAID array running special software that makes it appear like a tape library to the backup server.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.

About the Author

Michael Fisch **is Director of Storage and Networking for The Clipper Group.** He brings over seven years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at Mike.Fisch@clipper.com or at 781-235-0085 Ext. 25. (Please dial "1-25" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.