Clipper Notes™
VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE FOR THE INFORMATION TECHNOLOGY BUYER
Navigating Information Technology Horizons℠
New in 2007 Report #TCG2007040 March 10, 2007

Business Continuity Goes Better With SANs — The 3 Rs of Resilience

Analysts: Anne MacFarland, Joe De Natale, and Mike Fisch

Management Summary

When you take the field in an athletic event, you realize that, over the course of a game, probably several unexpected, potentially threatening things will happen, and, if Murphy's Law is in play, they won't be the threats for which you have prepared. Your well thought out offensive strategies may be neutralized quickly and the defenders may romp down the field. Your defense may prove unprepared, and there is always the chance that your most important players will get hurt. You cannot prevent bad things from happening. Instead, you must take steps to provide recovery alternatives – a bench of backup players, a playbook of alternate strategies, and practices to enable the team to reconfigure smoothly. These provide the resilience needed to continue playing to win.

Unexpected and potentially threatening events occur often in business as well. Again, resilience is the key. Business continuity practices of risk assessment and contingency planning are the way to mitigate the most likely risks, including the risk to IT systems and data that this paper will explore. **For enterprise IT systems and data, resilience is enabled by three very critical components – redundancy, remoteness, and restorability.**

Today, you may have a lot of storage attached directly to your servers. While you may have taken some precautions to recover from failure of these storage systems, such as making backup tapes, you may not be prepared for a broad or multi-stage catastrophe, and your preparations may not be adequate to recover in a meaningful timeframe to keep the enterprise running.

Storage Area Networks (SANs) – composed of intelligent switches and storage arrays, and storage management software – offer a variety of strategies to build IT resilience that an enterprise needs to survive the difficult challenges of the real world. You may think that you can't afford to go this route, but the alternatives do not deliver the resilience that you need. While there are further reasons for adopting a SAN (including storage consolidation, sharing of resources, increased performance, ease of administration, and scalability), the one that may justify new or additional investment is the resilience that it can deliver to your enterprise.

Today's SANs are more established, less expensive, and easier to install and administer. With business continuity in mind, you may not be able to afford to delay your decision. Read on to find out why.

IN THIS ISSUE

Business Continuance Today

Enterprises are vitally dependent on their information systems, and they succeed or fail based on their ability to deliver *business continuity*, in spite of day-to-day or catastrophic disruption. The financial impact of downtime of the IT system can be devastating. The hourly cost for downtime is estimated in the multi-million dollar range for brokerage houses and credit card authorization agents. Airline reservation system hourly downtime costs are near six figures. No matter where you fit in the impact of downtime spectrum, there is a cost – and the cost of doing nothing can be very expensive.

Disaster recovery is the most dramatic subset of business continuity. Almost half of companies experiencing a true disaster involving major data loss never reopen, and of those that do, over a quarter close within two years. And the more that you rely upon real-time data, the higher the risk of recovering "stale" data, and the greater the need for business-continuity measures that allow your systems to restart quickly from where they left off. **For businesses that require uninterrupted services, restarting from old data, whether stale by an hour or by a week, is unlikely to be acceptable.**

Fundamental to successful business continuity is the identification of those systems, applications, and data that are critical to the missions the enterprise. Following that, the three interrelated R's of resilience are critical: *redundancy*, *remoteness*, and *restorability*. Here's why.

- **Redundant components and multiple copies of data help avoid downtime and provide the means for rapid systems recoverability.** Implementing needed redundancy in hardware, data, and infrastructure is an important first step.
- **Remote isolation of the secondary copies from the primary system adds distance to the equation, an important consideration in today's world. The more operationally complete and accessible the remote site (an expensive capability), the quicker the recovery it can enable.**
- **This ability to restore or recover data quickly after a disaster is critical. The degree of restoration and allowable time varies by the business process, application, and available budget. Only you know which applications are fundamental to your enterprise and what resources you are willing to deploy to ensure their recoverability.**

Many solutions and techniques are used to deliver the three R's of resilience, such as server clustering, alternate paths, remote mirroring, consolidated tape backup, and electronic tape vaulting. Like pillars supporting a roof, a SAN supports and extends these methods. Though a SAN is not necessarily required for these solutions[1], it would be more difficult, or even impossible, to create a highly robust, enterprise-scaled resiliency implementation without one. To see why, it is important to take a closer look at a SAN's role in delivering each of the three R's.

SAN Technology Overview

In the continuum of technological advances in the computer industry, occasionally the introduction of a new architecture has profound impact on the way an enterprise processes information. Such is the case with Storage Area Networks (SANs). From its beginning as a means to offload storage traffic from the LAN and to consolidate storage devices, SANs continue to improve the robustness of IT systems.

Basically, a SAN is a network set apart from traditional Local Area Networks (LANs) and Wide Area Networks (WANs). It consists of host bus adapters (HBAs) connecting servers to switches, which are then connected to storage and other devices. All of this is interconnected by Fibre Channel cabling.

This architecture has many benefits. Administration is consolidated. Storage from different vendors can be scaled extensively and without interruption of information storage and retrieval. Storage capacity can be used more efficiently. Fibre Channel is inherently much faster than traditional point-to-point SCSI (Small Computer System Interface) connectivity, which it replaced, and extends to much longer distances. Fibre Channel's faster processing promotes speedy recovery from a disaster.

SAN management software monitors and manages not only devices, but also bandwidth and storage utilization, performance and server access to storage for the entire network. Good management software can dramatically reduce the time and resources needed to manage a SAN.

R #1: Redundancy

Redundancy takes many forms, depending on the user's requirements. A parachutist needs instant recovery, and therefore carries a spare parachute that can be deployed immediately on failure of the first

---

[1] The IP-enabled approaches of NAS and iSCSI both may involve LAN performance degradation, TCP/IP stack latencies, and data integrity issues that require still-expensive bandwidth to overcome. All telecommunications protocols have tradeoffs, the discussion of which is beyond the scope of this bulletin.

Redundant Copies of Data

RAID

RAID (Redundant Array of Inexpensive Disks) was developed because early disk storage and disk drives were unreliable. The many RAID schemes write data to multiple disks, so that failure of a single device will not stop the flow of data. RAID-1, also called physical mirroring, keeps a complete physical copy, and is the one that is meaningful to our discussion.

The mirroring may be local – within an array or to a redundant array – or it may be remote. Depending on the level of protection required, mirroring can be synchronous or asynchronous. Synchronous means that both the local and remote volumes are updated before the next write to the disk can occur. This provides the highest level of protection, but can slow application performance because of delays moving data across the distance between the arrays. Asynchronous means the local volume can continue to receive writes, even if the remote volume has not yet been updated; and the remote updates trickle in with some delay. This method provides higher application performance at the risk of potentially losing any updates not yet recorded on the remote volume, if there is a disaster.

Copy-on-write replication schemes share timing characteristics with mirroring, but are more limited in scope, and do not provide a complete, functional copy.

Point-in-Time Copy

Point-in-time or snapshot copy replicates either an entire physical volume or an index of the locations of the data stored at a given point in time. This copy can be used for quick recovery to a point in time (a restore) or for data staging data (to a back up, test bed or isolated branch system). Should a data file become corrupt, a system can quickly revert to a prior "clean" copy and use transaction logs to rebuild the data back to the present. This is called repair, as opposed to restore, and is the method that many enterprises use to recover from small and not-so-uncommon hiccups.

one. In situations where immediate recovery is not so critical, say in an automobile, a spare tire may be sufficient for redundancy, but you do have to stop to change the tire. In IT systems, for each component, there should be at least one spare to have full redundancy, and possibly more than one, if a high level of resilience is required. The same concept is applied to the applications and data to ensure high availability. Maintaining copies of data is a challenge in transaction-based environments, because the data is frequently being changed.

On the application side, clustering of servers[2] is the preferred means for creating redundancy for crucial applications. A SAN can also provide high-performance connections between clustered servers. In storage arrays, redundancy is achieved through redundant hardware and, importantly, through multiple copies of data enabled by the use of RAID or point in time copies. (See box above.)

How SANs Deliver Redundancy

The SAN itself delivers redundancy by providing multiple paths between servers and storage arrays. If a server or, optimally, a switch detects a failed connection, it can automatically route traffic through an alternate path, without application awareness or downtime. When the problem is solved, the fail-back, or restoration of the original route, can be similarly transparent. A well-designed SAN's redundant connections also can be enhanced by trunking[3].

However, data redundancy is not enough. Multiple copies of data are still at risk when they are stored in the same local storage array or in the same facility or general vicinity. This is only solved by adding "remoteness" to your solution. Inherent distance advantages of a Fibre Channel SAN enable this remoteness.

R #2: Remoteness

Tragic events have emphasized the need for all enterprises to include remoteness as an essential part of a well-conceived business continuity plan to keep their enterprise up and running. The largest enterprises have had this base covered for years, and some reaped the rewards of their preparedness last September. Larger and mid-sized enterprises are now scolding themselves into action to better ensure their own continuity.

The basis for remoteness is to have another facility far enough away[4] to be spared by local disasters, and preferably in a location less prone to disaster or interruption. SCSI protocol limits that remoteness to tens of meters. Fibre Channel

[2] Via software which coordinates two or more servers to work cooperatively.

[3] Trunking logically joins, in real time and when needed, separate paths for higher throughput and automatic failover.

[4] The more geographically remote the facility, the greater the expense and latency of delivering files and databases to that location. This is a factor when choosing a remote site.

Backup — The Traditional Way To Achieve Redundancy And Remoteness

The most rudimentary backup system is to methodically make backup copies of the data stored on or attached to each server to a locally-attached disk or tape drive. Then somebody gathers the media and takes it off site to a safe and preferably remote location. Because of the heavy dependence on human intervention in this process, this is often called a *sneaker net*. The drawbacks of dedicated drives, sneaker nets, and transportation of media become painfully obvious as enterprise applications and servers proliferate: too slow, too complicated, and too expensive because of the huge amount of labor involved. Often the process is incomplete. If an enterprise only does local tape backups, it is short-shrifting business continuance, both in terms of the lack of remoteness and the lack of restorability**.**

LAN-based backup, where the data to be backed up is transported over a local or wide area network to a centralized tape library, followed sneaker nets in popularity, but could slow performance of enterprise applications due to network congestion and consumption of server resources. Slow performance of production applications decreases business productivity, so a better solution was needed. Also, the dependence on tapes added a physical dimension to the quagmire; they were slow to write and to stage for recovery and, in many cases, still needed transportation for safekeeping. Thus, many enterprises turned to LAN-free and server-free backups and electronic tape vaulting to speed up the time required for both backups and recovery. Electronic tape vaulting, where tape images are sent a remote facility electronically, eliminates the manual and lengthy process of handling and retrieving tapes from the remote facility.

Through storage management software, enterprise-wide automated backup policies can easily scale in terms of the number of servers and amount of data being backed up, without linearly adding staff. Human error and lack of procedural consistency are diminished. Restoration can begin any time after it has been received at the remote location – a time usually measured in minutes, not days. The result is a more secure and more robust backup system and a lower total-cost-of-ownership (TCO). SANs enable electronic tape vaulting, fast non-disruptive off-the-LAN backups, and also server-free backup. When invoked by array, switch, or appliance, this preserves server CPUs for operational workloads.

increases this limit to 10 km, and more with repeaters or using encapsulation over WANs.

At the remote location, an up-to-date synchronous mirror, nearly current asynchronous mirror, or frozen point-in-time copy of data is kept. (See boxes on page 3 and 5.) To ensure data integrity, a copy of the data must also be isolated from corruption, either by "breaking off" a point-in-time copy, or by using backup procedures. (See box above.) A real-time copy of production data is not enough – mirrored files are subject to simultaneous corruption.

How SANs Enable Remoteness

SANs facilitate robust, enterprise-scale remote mirroring by enabling server-to-many-arrays and array-to-array copy strategies. Today, a high-performance Fibre Channel switched fabric can connect multiple local and remote volumes using a variety of campus-area, metro-area, or wide-area configurations.

Fibre Channel SANs optimize the backup process, delivering all of the benefits of LAN-based backup, while adding faster data transfer rates and reducing the impact on LAN or production applications. As a network, Fibre Channel is capable of quickly and reliably moving large blocks of data between two points with very low protocol overhead. Impact on server computational resources is, therefore, minimal and backup and restore times are faster. Furthermore, there are SAN backup solutions that bypass servers entirely. These server-less, zero-impact solutions, available for certain application and most storage array configurations, offer the ultimate in uptime and performance. Finally, the extensibility of a SAN allows servers dispersed over wide areas and in remote facilities to participate in a centralized backup system.

R #3: Recoverability

A key element of resilience is the time it takes to recover from small and large interruptions. Recoverability is rooted in the other two Rs – Redundancy (having the data with which to recover) and Remoteness (assuring the integrity of that data). Beyond those two Rs, recovery is all about a clear plan to reinitiate failed business processes, one that identifies the varying degree of immediacy of restoration needed by these processes. Various recovery techniques[5] involve different costs: **in general, the more immediate the process, the more expensive it is to enable.** The type of recovery should depend on the application and its importance to business operations. Most businesses prefer a tiered recovery process, one that channels resources first to the most critical applications. Applications enabling

[5] See box on previous page.

Recovery Techniques

Synchronous Mirroring

For applications that are critical to business continuity of an enterprise, the original and backup (or mirrored) data must be concurrent in time. This means that every transaction or write to the active disk is backed up at the same time to another disk. This synchronous writing guarantees the concurrency of the data files with no loss of data. Intuitively, this may seem simple and straightforward, but in reality, there is a slight downside to this method. Before another write to the disk can be made, both writes must be completed. Although there may be a lag in performing it in this manner, the enterprise is assured of the integrity of the back up and is, indeed, a true copy of the original. This approach provides the greatest protection of data loss in the event of a business interruption.

Asynchronous Mirroring

Other applications are not of first order criticality in maintaining business continuity. These files, databases, etc., are mirrored, but in this case, it is done asynchronously. The system issues a write to the active disk and then sends it to the mirrored disk, but does not wait until there is confirmation back from the second disk before it goes on to other processing. There is a potential of loss of the some of the later changes or transactions, because the mirrored copy is somewhat behind. How much depends on the amount of data, the distance, and the bandwidth. Therefore, there is a greater risk with asynchronous mirroring, but transactions are processed faster and less bandwidth may be required. It's a compromise.

Optimized Tape Restore

The optimized tape restore involves the organization of the restore process and the underlying tapes according to the enterprise priorities set for business continuity. In most cases, there is software available that can control this process. The major and most serious drawbacks to this method are twofold. Although recovery is somewhat automatic, it still may be a time consuming task, measured in hours and possibly days. Secondly, backup data on tape may not have full concurrency with the latest operational data. This methodology is used principally when the applications do not require concurrency. It is faster than conventional tape restoration methods.

Conventional Tape Restore

Conventional tape restoration involves systems (or, maybe, applications) in their entirety. Files that have been dumped to tape are thusly restored en masse. This method is never recommended as the way to recover quickly from business interruption, but can be used as a last resort.

core business processes and interactions with customers are most critical. Email and messaging used to be considered less-than-critical, but that is changing at many enterprises. **Only an enterprise can determine what is truly critical, near critical, or not critical.**

A tiered approach allows the overall business impact of a disruption to be minimized, without the expense of providing the highest level of recoverability across the board. Less-critical applications either take longer to recover or have potentially more data at risk of loss.

How SANs Deliver Recoverability

Such a tiered approach is immensely facilitated by the reach, connectivity, and management capabilities of a SAN. Manual invocation of such a tiered process would not only be costly and prone to errors further exacerbated by stress, but its invocation would not be as fast as the automated recovery enabled by a SAN. The speed of Fibre Channel, and the augmented bandwidth enabled by the redundant connections of a well-designed SAN, speed the process once invoked. SAN software optimizes the multiple processes and timing of a tiered restore. **Like the playbook and the players on the bench, SANs enable more alternatives for recovery while shortening the time to recover.**

Resilience For Good Times, Too

Even when things do not go wrong, resilience is needed. Common challenges to an enterprise's success include increased business activity and the need for IT integration due to mergers and reorganizations. The resilience, or agility, gained from a SAN provides the ease of management and a foundation for scaling to meet growing or changing business needs. **On a sailing ship, agility is required to compensate continually for the change in wind, seas, and weather. Without networked storage, it is difficult for enterprises to be resilient, agile, and prepared for business continuance as well as business growth. With storage consolidated around a robust SAN and with the right software tools for management, these adjustments usually can be made without "taking down the sails".**

In addition, the ability to monitor and report the details on all of the nodes involved in SAN traffic flow enables development of application-centric workload balancing and service-level fulfillment. These are necessary to preserve the health of the IT system and maximize the infrastructure investment.

Conclusion

SAN infrastructure supports, extends, speeds, and enables the following solutions for enterprise-class resilience before, during, and after disasters**.**

- Server clustering
- Alternate paths
- Point-in-time copy between storage arrays
- LAN-free or server-less tape backup
- Electronic tape vaulting
- Synchronous mirroring
- Asynchronous mirroring
- Optimized tape restore
- Conventional tape restore

Business continuance is a "bet your business" issue. The attention and dedication to insuring a well-planned methodology is directly related to the success that an enterprise will have when a business interruption occurs. This business interruption can be either a temporary one, or a longer-term incident that requires complete recovery of enterprise IT systems. **It's all about maintaining the resilience of the enterprise. Redundancy, remoteness, and recoverability, provide IT with system resilience.** When the resources of a SAN are in place, enterprises will have a flexible foundation and tools with which to craft resilience strategies.

A SAN infrastructure inherently possesses characteristics that both minimize downtime and speed recovery, in the case of major and minor disasters. SAN software optimizes administration and utilization of storage during periods of normalcy as well, and facilitates the IT growth that is a consequence of successful business continuance.

Business continuity may not have been the primary reason that many larger enterprises have already chosen to implement a SAN. They were driven by performance requirements, server and storage array consolidation, sheer size, and scalability. **Today, SANs also present those larger enterprises with a significant opportunity to further their quest for business continuity. This benefit is also available to those investing in SANs for the first time, and may be the most important justification for investing in new SAN infrastructure.**

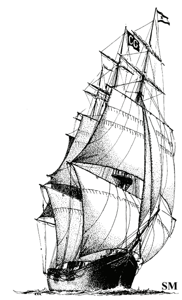

About The Clipper Group, Inc.

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 Authors

Anne MacFarland **is Director of Data Strategies and Information Solutions for The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 128. (Please dial "128" when you hear the automated attendant.)***

Joseph S. De Natale **is Director of Enterprise Systems Planning with The Clipper Group.** He brings more than forty years of experience in the data processing field with particular emphasis on systems management and application development on large-scale mainframes. Prior to joining The Clipper Group shortly after its founding, Mr. De Natale was an independent consultant where he provided expert opinion on data center management for civil cases. He later joined International Data Corporation (IDC), as a senior consultant and analyst, where he covered banking systems, data center management software, and large systems computers and storage. Formerly, Mr. De Natale spent eleven years at Citicorp Information Resources (CIR) as CIO of the Boston Data Center, where he managed the support of over 200 outsourcing contracts for thrift institutions. Earlier, he was MIS Director for the Lahey Clinic, and prior to that was a Project Manager for Computer Sciences Corporation, where he was involved with government outsourcing and applications contracts. Previously he was Director of AVCO Computer Services for fourteen years. At AVCO, he was responsible for all internal data processing, and initiated the marketing of computer services to commercial clients. Mr. De Natale began his career with Pratt & Whitney Aircraft as a programmer. During his career, Mr. De Natale was involved in the evaluation, installation, and operation of large-scale mainframe systems and for the development of commercial, scientific, and engineering application systems. He has also had successful experience in the marketing and operation of outsourcing contracts. Mr. De Natale attended Boston College's Undergraduate and Graduate schools of Mathematics.

- ***Reach Joe De Natale via e-mail at denatale@clipper.com or at 781-235-0085 Ext. 124. (Please dial "124" when you hear the automated attendant.)***

Michael Fisch **is Director of Storage and Networking for The Clipper Group.** He brings over ten 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. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

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.