October 2005 - Volume 1, Issue 4

Chairman's Corner

By a timely coincidence, the latest edition of the ASNP Digest has the theme of disaster recovery. With the terrible consequences of this year's hurricane season still upon us, it has become very apparent just how important it is to protect the entire storage environment using DR and business continuation safeguards.

Accordingly, you will be reading the views of experts in this arena, as well as end users like yourselves who have taken steps to proof themselves up against disaster. Read what they have to say and see how it might apply in your organization.

If you need further assistance in your DR efforts, feel free to ask a question on our discussion forums available at www.asnp.org. By posting a request, you can gain access to the combined wealth of ASNP user experience.

Additionally, look for further papers and case studies on DR to be posted on our site in the coming months, as well as chapter meetings in your area where subjects such as disaster preparedness are sure to be discussed.

One of the big purposes of the ASNP is to help you understand storage technology and help you implement it. That's why a team of ASNP members who

are also industry veterans are right now compiling storage best practices for the benefit of the membership as a whole.

And if you need anything that isn't covered in the above paragraphs, don't hesitate to call your chapter officers or ASNP headquarters. Good reading!

Daniel Delshad, Chairman and Executive Director of the ASNP

Back to Top

Maximizing Security and Manageability for Business Continuance and Disaster Recovery
By Tom Nosella, Director of Engineering, Cisco Systems

ASNP Contributed Article Submitted by Cisco Systems, Inc.

Introduction

Today's businesses depend on the continued availability of mission-critical applications and their associated data. Networks, both storage and data, continue to play an increasing role, especially related to their ability to minimize business interruption and data loss during any type of disaster. It's essential for businesses to ensure that the right failover mechanisms are in place, most often in geographically dispersed locations, so that data access can continue uninterrupted if one location is disabled.

It is up to each organization to create a strategic plan to best implement disaster recovery and business continuance plans. Companies must consider a wide range of issues, including application performance, the effects of distance on existing applications, availability needs, and which applications and data to address within the scope of the business continuance solution. For example, not all data will be backed up; some may be synchronously replicated.

Among the top challenges are network security and administration. As more businesses turn to technologies such as Small Computer System Interface over IP (iSCSI) and Fibre Channel over IP (FCIP), which use TCP/IP for transport, the need for storage area networking (SAN) security is increasing with sensitive information passing over common data networks and often extending outside the traditional data center. The growth and diversity of data networks is also making network management a key concern. With an intelligent, highly secure, multiprotocol solution, organizations can help ensure that their disaster recovery and business continuance systems will deliver the performance they demand.

Applying In-Depth Security

As networked applications become an integral part of business, organizations are extending their network infrastructures beyond the enterprise data center or the campus, crossing into service provider networks, and increasing their network security concerns. A closely related issue, the need to provide protection for data transmitted over a network, is becoming more urgent. For example, new industry-specific regulations such as the Healthcare Insurance Portability and Accountability Act (HIPAA) and the Gramm-Leach-Bliley Act (GLBA) are making data security mandatory. To help ensure that SAN data is secure and protected from violation, SAN administrators must address numerous concerns as they deploy and manage their disaster recovery and business continuance systems.

Authentication and Authorization—Robust network authentication and authorization is critical for organizations that wish to avoid accidental corruption and/or malicious attacks on the data stored on their SANs. Authentication and

authorization enables only certified users and devices to connect to the network, provides an audit trail of those users and devices that have connected or attempted to connect to the network, and helps prevent one host from having access to another host's data.

Fibre Channel Security Protocol (FCSP), part of the ANSI T11 FC-SP draft standard, can provide robust authentication in a SAN environment, as well as data integrity for both host-to-switch and switch-to-switch communication. Organizations can perform authentication locally in the switch or remotely through a centralized authentication, authorization, and accounting (AAA) server. FCSP is supported by a variety of SAN switch vendors, including Cisco Systems, and by all major host bus adaptor (HBA) vendors.

Maintaining Data Integrity and Encryption—Data encryption is also important for preventing intruders from viewing or modifying confidential information. The popular IPSec encryption protocol helps ensure confidentiality, integrity, and authentication. The protocol is often used to build secure tunnels between data centers, and is transparent to applications. Cisco MDS 9000 switches include integrated hardware-based IPSec support that provides wire-rate IPSec encryption and decryption with Advanced Encryption Standard (AES), Data Encryption Standard (DES), and Triple Data Encryption Standard (3DES) for FCIP and iSCSI storage traffic.

To reduce the potential performance impact associated with FCIP transmission over the wide-area network (WAN), businesses can also employ FCIP compression (also supported in Cisco MDS 9000 Family SAN solutions), which increases effective WAN bandwidth. By integrating data compression as part of their IP storage services modules, organizations can implement more efficient FCIP-based business continuity and disaster recovery solutions, maximizing both network security and performance. Using this technique, a Gigabit Ethernet port on a Cisco MDS 9000 Family switch used for IP storage services can achieve as high as a 30:1 compression ratio, depending on the data "compressibility," with typical ratios for common database traffic of 2:1 to 3:1.

Secured Management Infrastructure— The management of network devices in a data center needs to be secured as well to thwart unauthorized access, because malicious users with access to the console of a network device can easily alter the network configuration. The Cisco MDS 9000 Family provides a range of highly secure management functions, including Secure Sockets Layer (SSL) and Secure Shell (SSH) Protocol Version 2, which provides secure remote access through the use of authentication and encryption. SSHv2 can be used in conjunction with authentication protocols such as TACACS+ and RADIUS. Secure File Transfer Protocol (SFTP) provides a highly secure, encrypted TCP-based method for moving files. In addition, full Simple Network Management Protocol version 3 (SNMPv3) support provides authentication and authorization services for SNMP Management Information Bases (MIBs).

Using a storage solution that includes intelligent services to perform authorization, authentication, encryption, and administration tasks, organizations can help ensure that their business continuance and disaster recovery initiatives will meet today's security concerns.

Managing a Diverse Environment

Network management is another key concern for businesses deploying a business continuance solution. As storage network environments continue to grow, organizations need more comprehensive network administration services. Many are deploying storage solutions using equipment from multiple vendors, which further complicates management tasks. Companies need the ability to

effectively manage these heterogeneous storage solutions to ensure maximum performance and cost-effectiveness.

The Cisco MDS 9000 Family of multiprotocol storage switches was designed for diverse environments, and can support Fibre Channel, iSCSI, Fibre Connection (FICON), and FCIP simultaneously.

To provide extensive visibility into the environment, the Cisco Fabric Manager lets administrators view and manage the fabric as a collection or network of devices. This GUI-based application recreates a fabric topology and represents it as a customizable map. Third-party devices, such as hosts and storage devices, and other fabric switches that support the FC-GS-3 standard, can also be discovered and mapped as part of the topology.

The Cisco Fabric Manager consists of three major components, including a topology window, inventory window, and tools window. The topology window displays the discovered topology in a form that administrators can customize and navigate. The inventory window displays a tree-like structure of all elements, both physical (such as fabric switches), and virtual (such as zones and virtual SANs). Finally, the tools window displays a series of fabric-oriented tools that organizations can use to configure, monitor, or troubleshoot network devices.

Effective SAN management should not only provide a set of tools to manage switches, but also open interfaces with access to raw performance and configuration information within the switch that can be used by third-party applications. The Cisco MDS 9000 Family of multilayer directors and fabric switches provides an industry-standard application programming interface (API) using the Storage Management Initiative Specification (SMI-S). SMI-S facilitates SAN management in a multivendor environment, and is based on the Common Information Model (CIM), an object-oriented information model that describes management information in a network or enterprise environment. By using a standardized architecture, SMI-S provides support for common, extensible applications that work across multiple vendors' SAN products.

Flexible Virtual SANs

Virtual SAN (VSAN) technology can give organizations additional options as they seek to improve the overall health and scalability of their business continuance and disaster recovery initiatives. As part of the new ANSI T11 FC-FS-2 draft standard, VSANs provide a method of building separate, virtually isolated fabrics on top of the same redundant physical infrastructure, and can secure and isolate SAN domains without the need to build out multiple separate and costly physical infrastructures.

Cisco MDS 9000 Family devices can create up to 256 isolated VSAN topologies (the hardware supports the expansion to 4096 topologies) or layouts within the same physical infrastructure. Similar to virtual LANs, VSANs use the proven technology and ease of configuration of Ethernet networks to provide features such as traffic isolation and security—in a SAN environment.

For example, administrators can use simple zoning to restrict the access and traffic flow between devices within the fabric by securing access at the edge. But with the storage network domains created by VSANs, businesses can segregate even a single switch into multiple virtual environments. They can completely separate different VSANs to help ensure that a device outage or fabric instability is isolated within a single VSAN and doesn't cause a fabric-wide interruption. This eliminates the need to purchase expensive solutions that require multiple, physically isolated fabric switches. VSANs also lower the storage network cost of ownership by maximizing port utilization and reducing the effective per-port cost.

A well-planned VSAN architecture reduces the total number of SANs or fabrics deployed in the data center, while enabling businesses to separate their backup, recovery, and remote data mirroring domains from application-specific SANs.

Conclusion

Cisco storage networking solutions deliver industry-leading, standards-based, intelligent services that help ensure optimum manageability for a wide range of diverse environments. At the same time, they provide advanced security to safeguard sensitive data at all levels of the network. A rich array of performance and latency technologies enable organizations to cost-effectively implement these intelligent services without sacrificing network performance.

Together with its storage networking partners, Cisco Systems delivers robust solutions that help organizations build end-to-end backup and recovery solutions and disaster recovery solutions that are scalable, cost-effective, and highly secure

Back to Top

Taking a More Comprehensive View of DR: Best Practices in Disaster Recovery
Written by Robb Dennis

Disaster Recovery (DR) means different things to different people. Some regard it as a purely a matter of technology. You buy the latest and greatest box, plug it in and, Hey Presto! You have a DR plan. Others regard it a little more deeply. They believe that IT coordinates its systems and equipment in conjunction with a written plan to keep organizational data protected. Some people, however, recommend going much further. They look upon business continuity (BC) as a more encompassing label, with DR being only one element of the BC strategy. Further, they recommend that BC be taken out of the hands of IT and placed under the control of top management.

Depending on who you ask, then, you get a different take on DR. “The two keys to DR are redundancy (in storage) and the ability/willingness to test what you have done,” said Mike Karp, an analyst with Enterprise Management Associates.

But technology is only one aspect of DR. Some experts hardly even mention technology when they discuss being prepared for a disaster. “Most organizations I know about put BC/DR under the IT umbrella,” said John Glenn, a BC consultant based in Clearwater, Florida. “My preference is to put BC - of which DR is a subset - under the CFO, CEO, COO. BC is not about technology, it is about processes."

Roberta Witty, an analyst with Gartner Group also believes that DR effectiveness is hampered by an IT centric view. “One of the big No-No’s is thinking that the IT department is the only department needed to develop, test and recover the business,” said Witty.

So what is exactly is DR, then, how does it fit in with BC, and what are considered best practices in the industry? Let’s take a closer look.

Disaster 101

It’s good to start any DR discussion with a definition of what DR is:

"Disaster Recovery is the coordinated process of restoring systems, data, and infrastructure required to support key ongoing business operations,” said Chip

Nickolett, a disaster recovery specialist with Comprehensive Solutions of Brookfield, Wisconsin.

While this may seem basic and possibly even obvious, there is usually a big disconnect between the initial DR project goals and the ultimate DR plan developed by many companies. That's why it's important to perform a business process analysis to truly understand how the business operates, the dependencies that one area has on another, as well as the dependencies on or requirements for external data sources and interfaces. This basic understanding helps define the detailed requirements and drive the technical plan development. Output from this effort includes:

- Recovery Time Objective (RTO), or the goal to have business operations restored
- Recovery Point Objective (RPO), or how much data can be lost from the point of the disaster going backwards.

The resulting recovery plan needs to address the unique requirements of the business. If it is important to the business then the DR plan must address it in detail. Note, though, that you can't just do this linearly. You have to also map the dependencies between plan components. That way, if one part fails, you can quickly determine the downstream impact of that failure. Then apply even more scrutiny to those critical areas to minimize the risk of failure or problems, and thereby maximize the probability of overall success.

Let's take a look at an example. There may be an understanding of a process that generates a file for external processing, for instance, but detailed knowledge about that external processing is often very limited. This type of gap can create delays in recovery. Another example is the lack of understanding about hardware specific limitations.

"Time and time again we have seen problems due to tapes not being readable by another tape drive, file system "dumps" or "ghost" images that cannot be restored on different hardware (platform or disk size), older operating systems that are not supported on newer hardware, or restoration platforms that are not binary compatible," said Nickolett.

DR and BC

Many people are confused by the terms DR and BC, not to mention failover. The day-to-day management of critical business functions that utilize mission critical information is often built with failover functionality. This ability provides users with access to information needed to keep a business running. This is essential but it is not DR. However, this can be looked at as a sub-function of disaster recovery and business continuity. Backup and recovery can also be viewed as a sub-set of DR.

Actual DR plans are built to restore full functionality in relation to infrastructure, including facility and IT at time of crisis. That might mean fire, flood, a plague of locusts, or any other interruption that denies the ability to run in a normal operational mode.

Business continuity plans takes things from a higher and more broadly encompassing level. Whereas you need a disaster to activate a DR plan, BC includes anything and everything that is required to keep the business running. Therefore, BC address all functions of a business -- from personnel to facility and first aid to food/shelter. It addresses every aspect of the organization, not just technology. Increasingly, for example, organizations are giving more emphasis to employee "health" and welfare during an event.

“In a regional outage, you can't expect personnel to show up for business recovery if they are having problems at home related to the event, so you must maintain ongoing support so they are prepared at home,” said Roberta Witty of the Garner Group. “The American Red Cross is used a lot for this part of the education and awareness training for BC.”

IT, then, is just one element of an organizational BC view. As well as information and IT systems, it looks at having an alternate site for the company to operate, extra personnel in times of emergency, perhaps even the ability of staff to work from home if headquarters is inaccessible. IT obviously plays a pivotal role in BC, but IT is junior to BC and must be treated accordingly.

What steps do companies need to do to get from where they are to having state of the art DR systems and processes? Start with business continuity. Understand the business context and business value of IT. It is critically important, therefore, that the DR plan is based on a solid business continuity plan that has taken into account the reality of the business requirements for recovery.

“If the DR plan cannot meet the requirements of the business units, it is of no value,” said Michael Croy, director of business continuity solutions at Forsythe Technologies Inc., a Chicago-based IT consultancy and infrastructure firm.

He gives an example of an often overlooked aspect of DR – the criticality of tying change management into the recovery plan. The infrastructure is in a constant state of flux due to new technology, new or changing business processes and changes in the application mix. Thus the DR plan must be regularly monitored and managed.

DR Technology

Redundancy is one of the primary areas of DR technology i.e. the idea if one method of saving the IT shop doesn't work, you have another as backup.

Remote SAN Mirroring and Replication are leading-edge and expensive technologies that can keep a “hot site” up to date, minimizing data loss (RPO) and the time to recover (RTO),” says Nickolett.

To reduce costs, it might be possible to replicate from high price storage to cheaper disk arrays such as moving data from EMC Symmetrix to a remotely located Serial ATA (SATA) array. There are even examples of collaboration between companies for replication purposes i.e. to cut costs they provide failover capabilities for each other. This is a potentially attractive solution for Small and Mid-sized Enterprises (SME) that can't afford a DR site or a second data center.

“If you need to have an active-active environment when both sites are handing off to one another, get equipment of similar capabilities at both ends,” suggests Karp. “Alternatively, if you have active-passive, the remote site just replicates the data center using much more economical gear.”

However, one thing to pay attention to is the network. If the communication channels go down, you can’t replicate. So it makes sense to also have to add a redundant network connection. You have to implement this sensibly, though. One company, for instance, took careful steps to have redundant lines to a remote site. Unfortunately, the cables went through the wall at the same place. So a rogue backhoe took out both cables and there went communication redundancy.

Recovery Challenges

The acid test of any DR strategy comes when data is lost and you need to recover it.

There are always challenges in recovery. In Nickolett's experience, he has never failed to solve these challenges for UNIX. For Windows servers, though, he says it's been a different story. Problems crop up in trying to restore backups or ghost images to equipment that is not identical, for example. Similarly, software conflicts with service packs, patches, and other software packages are also common.

"This is not to say that recovery can't be done on Windows, but these types of systems require a greater degree of control, configuration management, and ongoing testing than other systems," says Nickolett.

Another challenge in recovery is cost. DR initiatives are frequently pushed to the bottom of the spending pile due to other more pressing priorities. Unfortunately, it typically takes a real-life disaster on site or near to home to change this mindset. And by then the damage could already be catastrophic. The best way to sell DR, then, is not in terms of ROI but in terms of risk management. Would the company consider operating without an up-to-date insurance agreement? Then why on earth would anyone consider running IT without DR?

"We suggest to our clients and prospects that they view the cost of DR planning as insurance," says Nickolett. "After all, what is the cost of downtime on a daily basis and how many days could you be down and not go out of business?

Look at the impact of the loss of data to the company and use that to build a business case? Yes, it is expensive. But if you can demonstrate that any DR system would pay for itself by eliminating one day, one week or one month of company non-operation, you can gain wider support for approving the PO. In simple terms, the cost of being prepared for a disaster is less than the cost of being unprepared for a disaster. When placed in this context the focus on doing things right often outweighs a budget target.

But the value of DR varies from company to company. Some can't tolerate an hour's downtime without catastrophic losses. Others can get by with perhaps a few days offline and live to tell the tale. Work out your companies tolerances, or lack thereof, and price out the DR strategy and technology infrastructure accordingly.

"What is best for the recovery of the business is what is critical," says Croy. "If it is zero loss of data then newer technologies, probably at a higher cost, need to be utilized. However, if it is not mission critical or business critical data, perhaps less sophisticated technologies may be used that are much less expensive."

Nickolett suggests that this analysis include a categorized list of "systems" (hardware, software, and network access), "data" (true data, audit information, business process procedures and documentation, support procedures and documentation, and other "business knowledge"), infrastructure (office space, phones, network access, remote access, websites, etc.), and "business operations" (unique ways that each company generates revenue and stays in business). Categories for items should be listed in tiers based on particular business requirements. For example, Tier 1 might be "recover within 24 hours", Tier 2 might be "recover within 72 hours", and Tier 3 might be "recover within 10 business days".

One last thought on the subject of DR challenges. During a disaster, people will be under a great deal of stress. Something that might normally seem routine and easy may all of a sudden become very difficult. Having very detailed procedures that eliminate the need for the person executing the plan to make decisions is extremely helpful. Having tools to help work through problems and support decisions made under stress will help keep the process moving. Having clear, frequent, and ongoing communication will help identify problems early and may

stop someone from doing the wrong thing. All of this is truly important. But just as important is getting people in this mindset during a test. It is not “just a test”. You cannot get information from a production system. Apply a reasonable amount of pressure. Make sure that if something is not brought to the DR site or publicly available that it isn’t used. And keep a detailed log of everything that occurs, when it occurred, what the symptoms were, what the process was to identify the root cause, and what the fix was. This will help ensure that nothing is lost and provide valuable timing metrics.

Despite the obvious benefits of DR testing, it is shocking how rarely it is done. A survey of your peers will reveal that relatively few organizations have tested their DR plans. Most offer the excuse that they just don’t have the time. That’s a little like never having time to do routine maintenance in your car. After a while you will grind to a halt.

“We are dealing largely with event driven people,” said Karp. They have to be bitten in the ass by an alligator, before they decide its important to drain the swamp.”

State of the Art

Some seek to cut costs by purchasing good-enough. Other prefer to find state-of-the-art DR and BC technologies and stay on the leading edge. But what is state-of-the-art?

“State-of-the-art recovery is engineering a system and infrastructure that always provides access to critical data and systems,” says David Palermo, Vice President, Marketing, SunGard Availability Services. “The challenge that companies face is keeping up with the 2x costs and the level of redundancy that data centers require to achieve true information availability.”

Some experts, though, think anxiety over the latest and greatest technology should be junior to a thorough strategy.

“State of the art is a myth when it comes to BC,” says Croy.

Back to Top

Top Ten Things Not to Do if you Want to Recover from a Disaster

1. Not involving federal, state and local authorities in the planning process - it can be hard in some cases to get them to test with you
2. Groups taking on too large of an initial project without having the experience or expertise to make it work
3. Plans that do not provide very detailed procedures (including troubleshooting procedures and vendor support procedures and information)
4. Plans contained in a huge book that are difficult to use, that are unavailable offsite when needed, or that are out of date
5. Don’t assume that a plan will work unless it is tested on a regular basis and kept up to date
6. Don’t presume data center staff know fully the requirements of the business units and, conversely, don’t assume the business units understand the technology solutions that are in place
7. Thinking that the IT department is the only department needed to develop, test and recover the business
8. Not having senior management sponsorship
9. Thinking that one scenario is all that you need to plan for and that this one

scenario is completely under your control
10. Not planning for an external event that could impact your own business

Back to Top

Harsh DR Realities – Funding Delays at Hospital Site

Jeffrey Pelot, CTO of Denver Health realized that his building's layout was a disaster waiting to happen. It is in the basement below street level. One floor above the data center are six hot tubs, each containing 1000 gallons of water. Then one day a small yet significant systems failure highlighted the problems inherent in having no disaster recovery infrastructure.

Thus Pelot got to work, figuring out how to fully protect the hospital's vital patient records and medical systems in the event of a disaster. He worked out a location to build a nearby disaster recovery facility that can duplicate the various campus configurations offsite. This center will utilize snapshot, remote IP copy and asynchronous replication to accomplish this. Only one snag. The PO was not approved. The DR site was not funded for 2005 and has been pushed to 2006. "The second site will house half of the network core, and all critical servers will have cluster elements located there," says Pelot. "Where servers are not clustered copies of the servers will be implemented on virtual servers and turned up as required in the event of a disaster."

How does it mirror the data and what types of data does it mirror. All data on the IP SAN is mirroed via technology protocol called chain de-clustering (details on LeftHand Networks web site).

Back to Top

Using a Third Party DR Site

Glenmede Trust of Philadelphia deals in wealth management for high net worth clients. Its IT environment is based mainly on Windows 2000. It uses Cisco switches, Dell servers and desktops (XP), CheckPoint firewalls and intrusion detection by VeriSign. It also harnesses a hot site at Sungard Availability Services as its recovery center in case of disaster. "This hot site is a microcosm of our network," says Nick Voutsakis Glenmede's CTO.

Glenmede has a contract with Sungard to use their recovery facilities. It has 30 desktops available at that location as part of the minimum resources needed to function in the Philly area. There is also a high bandwidth connection between the DR site and the Glenmede HQ which is used to replicate data all day long. "We can be up in the Sungard site very fast without backups," says Voutsakis.

In the event of heavy snow, for example, staff can work from home and connect to the corporate systems without any delays. In the event of a more serious event such as an Anthrax scare that closed the building for a week or more, the company would execute the DR plan and come into Sungard. The may systems would be left running at HQ, but staff would be using Sungard desktops. Alternatively, if the building blew up, mission critical servers would be activated at Sungard to run the business. "We have a subset of our mission critical systems at the DR site," says Voutsakis. "This arrangement gives us a lot of flexibility and that's what working during disasters is all about as no one can predict. Every disaster takes a different face."

Power Outage

Glenmede had a chance to test its DR plan during an unpredicted event. It lost power due to a power company problem. Even though the building is on two different grids, when one went down, a repair team inadvertently severed the second connection. That dropped out electricity to the building for almost a week. 90 percent of the building occupants went home – this is a 61-story building. Glenmede stayed behind, shut its servers down properly and assembled at two locations in Philly. The CEO declared a disaster at 8.30 a.m. and Sungard was told to activate the hot site. Non essential employees were sent home. By 10.15 phones were working. By 11, 30 workstations were active. By 11.30, the mission critical servers were going strong. “We had never done this before so weren’t sure what to expect,” says Voutsakis. “We greatly exceeded our service level agreement.”

He notes that essential services like the trust system, trading systems, Bloomberg, CRM, email, Internet and other applications were ready in two hours. To do this from scratch, he says, might have taken two weeks. DR at Glenmede, however, is not a case of everything being treated equally. Replication is done on some systems minute to minute while others are only replicated nightly. This is a function of their overall value to the company.

Best Practices

Voutsakis has some advice for others based on his hard won experience. Make plans as simple as possible, he says. A 100-page document won’t ever be read or understood so keep it simple, understandable yet thorough and accurate. Further, you have to have a copy at home, a copy on line, a hard copy in red cross bags used during disasters (that include flashlights, masks, water, food, etc.) And test, test, test. “We have seven tests a year that cover different levels of disaster such as losing our phones, a building evacuation, and so on,” says Voutsakis.

In addition, he recommends a solid BC committee be put in place that reports directly to the board. It has cross-functional representation with a core BC committee being comprised of the head of office services, risk management staff and the CTO. Voutsakis stresses that it is NOT just an IT activity. Then there is an extended BC committee that includes reps from 20 departments. Each rep has a backup. It is these reps that write the plans, collaborate with business units. Using this framework, the BC committee gets the various business units to execute the plan, trains them on how to do it, organizes the tests and grades them.

One time, for instance, they sent staff home for the day and rated their ability to get their jobs done. The grading system was 1 for being able to do the job with no trouble, 2 if workarounds had to be done to function, and 3 – could not function. Glenmede scores an average of 1.2 “We were very pleased with such a high score,” says Voutsakis

Another test was run at Sungard – a large test with all business units involved. It was done on a Saturday to simulate all servers being down. This proved to be an excellent way to evolve and mature the DR plan. “You can have the most brilliant IT infrastructure for DR, but if you don’t have people who understand and can execute the plan you are in serious trouble,” says Voutsakis.

Small Price

In summary, he sums up the view of DR at Glenmede. “We are trading millions of dollars and can risk no downtime, so it is a small price to pay for high availability and risk management,” says Voutsakis. “Our clientele are high net worth

individuals and they expect the highest level of service."

Back to Top

OpenVMS in a German Financial House

Deutsche Börse Systems AG is the systems house for Deutsche Börse AG, which is the German exchange for stocks and derivates. Its IT infrastructure is based on OpenVMS running on AlphaServers – ranging from DS10's to GS1280's. It has HP EVA storage running over Brocade switches. "We running our OpenVMS Cluster over two sites which are about five kilometers away from each other," says Michael Gruth, Head of System and Network support at Deutsche Börse Systems. "All data is mirrored on RAID 5."

What best practices does he recommend for implementing a DR plan?

"We recommend the base operating system be OpenVMS which already has all the features included for a disaster tolerant environment," says Gruth. In addition, he stresses the importance of having two separate sites using cluster technology which is online on both sites – that means no cold or warm backup, just active and online. "The technologies we are using are the cluster functionality of OpenVMS, Fiber Channel switches from Brocade, and Cisco switches/routers for the network," says Gruth.

What are the absolutely crucial aspects of a state-of-the art disaster recovery plan?

Gruth believes in having every component available at both sites rather than skimping on the DR site to save a few dollars. Also, he says that IT managers should be careful not to forget things like having a physical office space for remote management. Further, like so many others who are immersed in DR, he urges tests, tests, and more tests. "The first step for any company not active in DR is thinking about it, and realizing it is an absolute necessity that has to be paid for by the company," says Gruth. "Then think about additional hardware, not just splitting up the hardware you already have over two sites. You have to add capacity so you can live with only one site for more than half a year if necessary."

What lessons has be learned that he would like to pass on to others?

"Planning and realizing the project is a lot of work, and we needed about a year to finish up with it," says Gruth. "After planning the technical things, also plan the financial side in the same depth as it will be expensive."

Back to Top

The Members Group

The Members Group is an Iowa-based company that provides a wide variety of products and services for credit unions, including card processing and mortgage services. The company is most of the way way through implementation of StoneFly's IP SAN-based Backup Advantage (SBA) solution. The main focus has been on the disaster recovery component of SBA, which entails use of StoneFly's Replicator product, a remote asynchronous disaster recovery solution. (For more info on Replicator, see http://www.stonefly.com/products/product.asp?key=70; for more information on SBA, which is a completely integrated software and hardware IP SAN-based backup solution, see

http://www.stonefly.com/products/product_category.asp?key=23).

Jeff Russell, CIO of The Members Group says the company has an IP SAN in a production environment. Its primary site is in Des Moines, and replication is done to another site in Minneapolis – about 4 hours away by car. Why 300 miles away? That is where the network service provider has its hosting center. The Members Group effectively rents space there and uses an IBM iSeries (AS/400) to run several core applications. These are proprietary in-house applications that perform various transaction processing and settlement functions. Most apps are Internet based.

The company's Intel processor-based Dell servers run Windows 2000 and are being migrated to Windows Server 2003. Cisco gear is used for networking. The company has WAN for financial customers. An IBM storage array is connrected to the company's IBM iSeries server. A total of 60 servers are present, of which 25 Windows servers are connected to the IP SAN. The rest use direct attached storage (DAS). Only key servers are part of the SAN. The server used for virus signatures, for example, not being mission critical, is DAS only.

"From a business standpoint, those that servers aren't required to run the business are not on the IP SAN," says Russell. "It was a financial decision to not have everything in the SAN. That also set priorities on what mission critical applications would be replicated."

Prior to implementing an IP SAN, The Members Group had a pre-staged site that didn't offer real time replication. This site was 15 miles away. It consisted of little more than backing up the production servers to tape and then sending tape offsite to a place where there were pre-staged Windows servers. In the event of a disaster, it was a case of driving to that other facility, raking through the backup tapes and doing a restore to get the servers up and running. It could take days to get production rolling once again.

"When we analyzed our business, we realized that the recovery window was too long," says Russell. When we received greater scrutiny internally and from customers about our BC plan, we dived in and realized our old approach was inadequate."

The iSeries was backed up via a third party to a site in Chicago. That meant the iSeries people had to fly to Chicago to bring up another iSeries in Chicago. The Members Group shared that machine with several other companies.

Now, however, the replicated SAN is connected to the production SAN in real time. That means that if the building is down and the production SAN is not available, the remote SAN can take over and is then attached to pre-staged servers.

"We can power up the remote servers remotely and have our core applications running within 30 minutes," says Russell.

While key applications are gotten running as priority, Russell says some other applications are recovered at a more leisurely pace. It might take 24 to 48 hours to get certain things restored. Essentially, Russell and his colleagues evaluated each application against business requirements. Some applications had to be available in real time, some 4 hours, some 8 hours, etc.

"We were through all the applications and looked at the amount of time we could afford to be down," he says. "We didn't have the capital investment to have our own alternative data centers and so we also didn't have to buy all the telecom infrastructure and other supporting gear."

With the IP SAN, though, bandwidth is a constraint. The company's large SQL database for example, is constrained by bandwidth. The company is investing in more bandwidth to support the IP SAN.

"We did some estimates for traffic to try to understand the need and we thought 1 T1 would be enough and it wasn't," says Russell. "We can replicate only selected applications due to these bandwidth limitations."

This problem was resolved when management decided to move to a new facility. That new site will add more bandwidth.

Best Practices

What best practices does Russell recommend for implementing a DR plan? He believes that the business drivers are the key to any plan's success. It has to be more than a tech plan.

"We have gotten away from DR term as DR assumes the facility is not available," says Russell. "BC asks how are you going to continue despite a business interruption."

He also emphasizes being ready for anything. A cut in the data lines, for example, is quite different from a tornado.

View the biz impact of downtime. What service levels and windows can we afford to be down.

Another important nuance is the details of the different applications. A blanket, 4-hour or 8-hour or 24-hour recovery period may not be applicable. Take a payroll process, for example. If today is the first Monday in the payroll period, then a two-week recovery window would work. But if it is the last day of the payroll period, then you better have things running within 4-hours.

"You have to work with the business unit heads to understand the drivers of each application," says Russell.

In terms of cost, he estimates that the hardware, software and other equipment purchases comes to about two thirds of the total cost. The Member Group saves about one third that it would otherwise have to spend to have its own facility and network in a remote location.

"That amount was the make-break of us having a viable DR solution," says Russell. "Renting made the project possible."

Back to Top

DR in Medicine

The Cancer Therapy & Research Center (CTRC) in San Antonio, TX has four locations in town feeding data into two data centers. It has two EMC Clariion FC4700 arrays (one at the medical center's data center and the other at the research center 22 miles away) and an EMC CX 400 (at the medical center). It also has two Cisco SN 5428 iSCSI routers and an MDS 9506 Cisco switch. The two data centers are connected over a 1 Gb Ethernet Metropolitan Area Network (MAN). Each Clariion initially carried 1 TB and that has expanded rapidly to3 TB. The CX 400 contains 7 TB.

"The CX 400 was purchases as we outgrew the Clariions," says CTRC CTO Mike Luter. "We are also buying a CX 500 this summer so we can mirror it with the CX 400CX 400

This architecture allows the medical facility to have complete redundancy and a 10-minute recovery window for any of its 25 servers. 21 are at the medical center and 4 at the research facility, all running mainly Windows 2000/2003 and some Linux. Each site houses backup and storage servers for the other site using host-based mirroring.

“We went for data mirroring rather than a hot site due to the fact that if something happened to the building, it would be very hard to treat anywhere else,” says Luter. So to us, BC is far more important than DR.”

That says, CTRC has a three-year plan to implement a hot site.

Image Storage

The operation uses a lot of 3D images from MRI’s etc. These take up a huge amount of space. One MRI is 170 to 250 MB per patient. As the facility does a patient treatment every 10 minutes, it can’t afford any system downtime.

“An hour of system downtime can be a matter of life and death in patient care, so BC is vital,” says Luter. As I can boot remotely, I can recover servers in about 10 minutes. I’ve done it in test many times successfully.”

Back to Top

No BS – Are list prices a starting point, a journey or a state of mind?

An analyst recently told me that anybody who pays list price for storage gear ought to have his head examined. When I talked to ASNP members about the subject, they voice a serious discrepancy between list price and what they end up paying. The question is by how much?

Some say that they receive anywhere from 25 percent to 50 percent off list on various storage hardware and software products. Others say they get less, and a few big companies might tell you in private that they get a lot more. Bottom line: the price you pay varies tremendously depending on who you are, who you are buying from, the state of the market, and the slings and arrows of whatever sales targets have to be met by when.

Now is it just me, or does this state of affairs in the storage world seem a little nutty? Let’s compare it to a more everyday purchase – a cup of coffee. Pretend for a moment that you receive the same sales experience from Starbucks as you might get from a major storage vendor:

You walk in to buy a latte. You look at the price list and it says $10. You tell the assistant that $10 seems a little high so he mentions the discount program – if you sign a long-term contract to buy one cup a day for three years, and that you will never shop at Coffee Bean or Peet’s, you can have it for $7. Meanwhile, a rep from a rival chain offers you coffee at $6. Starbucks goes lower and a price war ensues. You go home and boast to the wife what a killer negotiator you are – only paying $4.75. Until she tells you that her friend’s boss, who happens to work for a much larger company only pays $3.25.

Most users, then, look upon a list price as a place to start from. Invariably, they expect that the price will plummet sharply during negotiations.

In one way, this is a sign of a healthy industry. There are so many competitors out there with so many new products and technologies, that its natural they cut their prices to get a sale. Fibre Channel vendors have had to completely revise

their cost structures in the face of the IP SAN onslaught. Other vendors are introducing “lite” or low-cost SME editions of standard products.

When you talk to the vendors about all this, you really don’t gain much insight. It’s hardly any different than a political argument. The FC vendors give a completely different view from the iSCSI camp. The iSCSI’s say they are standardizing storage hardware so dropping costs are inevitable. They accuse the FC guys of trying to keep prices as high as possible as long as possible. But from the FC perspective, they have proprietary technology that they feel should cost more.

Perhaps in an ideal world we might eventually get pricing like the Saturn no-haggle deal. You walk into the car showroom, the sticker says $9,999 and that’s what you pay. No haggling, no surprises, etc. But that dream is a long way off. From the vendor perspective, there are too many variables to be able to set up one straightforward rate and standard discounts that apply across the board.

Back to Top

Chapter Spotlight - Michigan

Diane Favrow, president of the Michigan Chapter of the ASNP, is a Senior Engineer on the Digital Storage and Solutions Team at Comerica Bank.

Why did you join ASNP?

When I originally joined ASNP, I was on the Server Engineering Team at Comerica Bank, but I was involved in designing SAN solutions for large projects. Our company did not yet have an Enterprise Digital Storage Team, and I was looking for a SAN “support group” when my search lead me to ASNP. I was looking for the camaraderie and the exchange of ideas with a group of individuals working with Storage Networking technologies. At this time, there was no Michigan Chapter, and I was recruited into the role of President, so that we could start a local chapter. I agreed to do so. I liked the idea of a Storage Networking group focused on the needs and interests of the consumer of storage networking products. The ASNP chapter meetings give members the opportunity to hear lectures on trends and emerging technologies as they relate to Storage Networks. The officers select a meeting topic based on member interest and then recruit a vendor to sponsor the meeting. The vendor sponsor provides the guest lecturer and the meeting venue. The guidelines require the lecture remain non-vendor specific. After the lecture, based on the interest and questions of the members, the discussion may turn towards specific vendor technologies. This is a member driven discussion. The meetings are a great opportunity to talk shop and get ideas from other members.

What have you gained from it so far?

For starters, I am meeting many interesting people who are very knowledgeable on Storage Networking technologies. It’s enlightening to learn about the real life storage experiences of others and comforting to learn that we are not alone in many of the challenges we face. This industry is ever evolving, and it’s difficult to keep up through reading alone. Meeting individuals from other companies working with similar technologies brings a sense of reality to what is what in the world of SANS. Also, participation is giving me an opportunity to work in a leadership role as President.

What benefits do other members report?

One of the greatest benefits to attending the meetings is the opportunity to

network with peers. As a matter of fact, on one of our post meeting surveys, many of our members commented on how important the networking and open discussion time was to them, and they would like to see us make it longer for future meetings even if it meant longer meetings! Also, it's an opportunity to hear an expert speak on a specific Storage Networking related topic and then ask questions. It's an overall great place to meet technical folks and exchange information. Not to mention, it's a good reason to spend some time away from the office once a quarter!

How many members do you have?

The Michigan Chapter has about 67 members from Michigan and Ohio. We also have a great group of officers who work hard to incorporate members' input into the content of our meetings. We are fortunate to have all our officer positions filled.

What sort of activities do you engage in?

We've conducted four chapter meetings thus far. We typically have a dozen or so members attend each meeting. Each of our meetings has a vendor presentation followed by a "members-only chalk talk". This gives the members a chance to discuss and seek feedback regarding issues that are puzzling them in their own environments. We've received a lot of positive feedback on our meetings, and we are looking for ways to get more members involved, so that this can truly be a chapter that represents our diverse group of storage end-users!

Tell me about a recent chapter meeting?

Our last Michigan Chapter meeting held in the Media Room on the University of Michigan's North Campus. We had an exciting program lined up for our members. It was a panel discussion on the emerging trends in Information Life Cycle Management, and it included vendor participants from: Computer Associates, NetApp, EMC, and Compellant. The panel was asked five questions which were selected through a poll of our members which was posted on the ASNP website. Each panel member had four minutes to respond to the question. After each panel member responded, the audience had an opportunity to ask additional questions specific to the topic being discussed. The questions asked were:

- What do you see as the most important starting point for an ILM solution?
- What is the most critical component of an ILM solution?
- What methods/tools do you recommend for estimating the cost and ROI of an ILM solution?
- How important is data classification to an ILM solution?
- How many tiers of storage should be implemented in an ILM solution?

The result was an exciting discussion on the complicated topic of Information Life Cycle Management. All present felt it was a worthwhile presentation, and a rare opportunity to have such a knowledgeable panel openly discussing the various aspects of ILM.

What can we look for from the Michigan Chapter in the near future?

Look for our next chapter meeting on October 28th at the Integrators of New Systems office in Livonia. The topic for this meeting is Storage Virtualization. Michigan Chapter members can go to the ASNP website for more details and to register.

Back to Top

A Day in the Life - Karl Lewis

My name is Karl Lewis and I am Storage Administrator with the College of Engineering at the University of Michigan. I am responsible for the various storage platforms used to provision services to approximately 8000 students, faculty and staff of the College of Engineering in Ann Arbor, MI. I have been a member of ASNP for over a year and am the Events Chair of the Michigan Chapter of the ASNP.

I manage an 8TB SAN on a Dell/EMC CX700 array, front-ended by an EMC Celerra NS700G NAS gateway. I also manage roughly 5TB of DAS storage on many entry-level SUN servers (Netra and SUN Fire). Lastly, I manage several small tape libraries attached to the SUN servers. We're primarily a SUN Solaris shop but we're migrating many services to Windows Advanced Server and Linux.

We use automated freeware tools (such as Nagios) to email daily status and error reports from servers and storage systems in our data centers. I am the only person in my group tasked with investigating, evaluating and managing storage, so products that are feature-packed and easy to administer are of vital importance to us. Here's a typical work day for me:

0700 Check email from home before starting the daily commute.

0800 Arrive at the office and continue to check email.

0830 Update spreadsheet with expected storage consumption for the semester.

0900 Call the freight company to schedule a pickup of the new SATA array I evaluated last week. Leave sticky note monitor to remind me to finish writing all the product reviews I need to finish this semester.

1000 Meet with storage vendors from a startup company to investigate new products. I'm always interested in products that do more, but simply cost less. Many products from startups only offer one or two features that I need – which is a lot better than products with seven or eight features I don't. If possible, I'll get the vendor to issue an evaluation license for the product and I'll give it a try.

1200 Lunch

1230 Conference call with system administrators from other academic units on Campus. Discuss initiative to standardize on similar servers, storage and backup technologies to improve pricing leverage with vendors.

1400 Meet with my part-time student employee. Review progress on the web-based tool he's writing to display NAS space consumption by user, department and degree program.

1500 Meet with our customer help desk staff. Update customer web page with instructions for self-service restores from the NAS.

1600 Meet with my supervisor. Review various data center and host consolidation issues.

1730 Start the daily commute home.

1830 Get in a quick game of Counter-Strike on the PC.

1845 Start cooking dinner before my wife gets home and thinks that I've been playing Counter-Strike all afternoon.

Back to Top