NEXSAN

SOLUTION BRIEF: KEY CONSIDERATIONS FOR DISASTER RECOVERY

A Disaster Recovery and Business Continuity plan is specific to the circumstances, priorities and expense versus the value decisions of the organization, which presents a broad range of unique challenges for the mid-market. This brief outlines mid-market best practices in developing a disaster recovery strategy for their top, middle and bottom tiers.

TOP FIVE DISASTER RECOVERY CONSIDERATIONS

1. **PLAN FOR SUCCESS** - Whether by fire, power outage, hardware or human failure, or an act of God, setting rigorous standards for recovery service levels that are cost-efficient is the first order of business. Procedures must be well documented, detailed and up-to-date, as a critical step.
2. **VIRTUALIZATION** - Although virtualization is not required to have a tenable Disaster Recovery Business Continuity plan, it simplifies it by providing a robust, reliable and secure platform that isolates applications and operating systems from their underlying hardware. This dramatically reduces the complexity of implementing and testing DRBC strategies.
3. **AUTOMATION** - Reliability in the backup and recovery process is increased proportionally to the extent human intervention is reduced. Automation also enables the unstoppable progression of changes in the primary data center to be mirrored synchronously in the backup location.
4. **COST EFFICIENCY** - The protection architecture is often over-built with a premium solution applied across the board. That creates sticker shock and impacts an appropriate protection architecture at both the primary and secondary sites. Classify the necessary protection for each site to achieve a cost-efficient overall solution.
5. **TEST-OFTEN** - Not only do you need an effective plan for a broad spectrum of failures, you need to TEST it... often! Several analysts have reported that many organizations never test their DRBC plans, or test them so infrequently as to be ineffective. The more frequently you test, the higher the probability of a successful recovery. Test at least once a quarter.

Recent tornados, hurricanes, earthquakes and tsunamis have refocused attention on disaster recovery among business leaders and IT managers. While the broader scale of disaster recovery planning includes facilities, power, cooling, communications and people, data recovery remains key to business continuity. The tasks associated with the data center are a specialized and complex discipline that requires unique planning and management.

Large scale, regional and localized disasters require comprehensive business continuity plans that include the use of a secondary data center located far enough away so that it will not be impacted by the disaster at the primary site. A secondary site can obviously be very expensive, causing some to minimize the expense by using external vendors that offer disaster recovery services. In either case; whether you choose to have a secondary site or co-located services, the need to organize and move data from the primary site to the secondary site remains if you want your business to survive. When it comes to disaster recovery planning, without a plan, you can plan on your business failing.

THE DISASTER RECOVERY ARCHITECTURE

DRBC FOR THE MID-MARKET

A Disaster Recovery and Business Continuity plan is specific to the circumstances, priorities and expense versus the value decisions of the organization. This presents a broad range of unique challenges for the mid-market. While the “mid-market” is not strictly defined, it is not difficult to categorize. In an effort to provide context and guidance for how strategies must adapt to meet mid-market needs, we will look at the mid-market in three segments; Bottom Tier, Middle Tier and Top Tier.

Bottom Tier

This mid-market tier is comprised of organizations that are generally able to serve all their data needs with a single small-to–medium size storage array. They typically do not use advanced storage management capabilities but, like all businesses, are in need of highly available, high-performance storage that is easy to manage.

Middle Tier

The middle tier of the midmarket is represented by businesses that have multiple, sophisticated applications with varying demands (in terms of performance, availability, protection, recovery, etc.) to serve their users. Typically, multiple processes and procedures will be used for storage management in a storage infrastructure that is itself often tiered. This tier will likely have multiple virtualized servers managing capacities in the hundreds of terabytes. Performance, reliability and manageability are critical to their business objectives.

DEFINING THE TIERS OF THE MID-MARKET

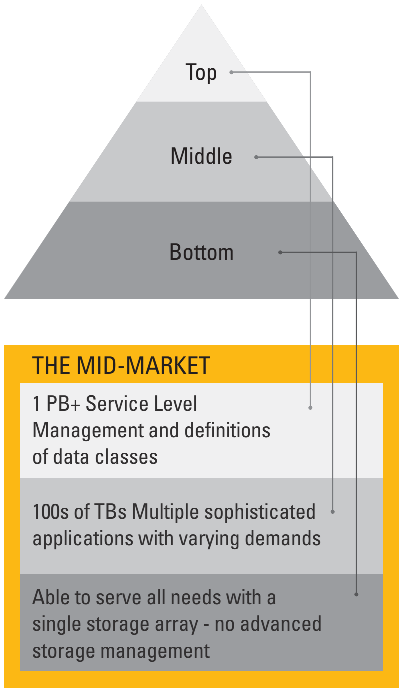

THE MID-MARKET

1 PB+ Service Level Management and definitions of data classes

100s of TBs Multiple sophisticated applications with varying demands

Able to serve all needs with a single storage array - no advanced storage management

Top Tier

These are organizations that are more proactive in designing infrastructure with service level management, resulting in the definition of standard service levels. They will define data classes and policies in the management structure. Policy-based storage management processes are usually standardized. Compliance is managed and enterprise content management is linked to storage optimization. This tier frequently has multiple consolidated and virtualized servers with extended virtualization and Storage Resource Management (SRM). They will commonly have a tiered storage infrastructure to meet performance and protection requirements, including an archive for legal and business protection. At this tier, performance, reliability and manageability are at the peak of the requirements. These organizations will typically manage hundreds of terabyte or even petabytes of data.

THE DRBC PROTECTION ARCHITECTURE

Building a data protection architecture appropriate for DRBC objectives is an extension of building a backup and recovery plan as covered in the Solutions Brief titled "Backup and Recovery." Suffice it to say, the principles are the same - understand the business value of the applications being protected and align that with technology, resulting in a cost-justified level of protection. Not every application deserves the highest level of protection money can buy. An organization must develop a classification schema to organize the value of an application's data and the performance levels for recovery. However, a DRBC plan extends beyond storage and additional consideration must be given to the servers and networks.

Disaster recovery implies the need to restore full operation after extensive repair to an existing center, or to build a full recovery at a secondary site. Assuming that time is of the essence, a DRBC recovery strategy should be targeted for a secondary site to achieve the fastest possible recovery to maintain business continuity. When thinking about an optimal storage architecture strategy, organizations must also consider their tier within the mid-market. Each tier faces increasingly complex challenges in achieving a recovery. The least complicated bottom tier may simply have a bucket of tapes or may write to a secondary array at a remote location. This approach could include the higher tiers as well. Below is an introduction to the basic protection architecture to be used at a primary location along with the strategy used to synchronize operating systems, applications and data at a secondary location.

CLASSIFICATION

Protection Tier	Classification	Availability	RTO	RPO
1	**Mission Critical Data (RAID 5 RAID 10)** • Critical to an enterprise, continuous access • Highest performance, near zero downtime	99.999%	1 min	0
1-2	**Business Critical Data (RAID5)** • Very important to the enterprise, frequently accessed • High performance, high availability, less than four-hour recovery	99.99%	1 hr	15 min
2-3	**Accessible Online Data (RAID 5 or RAID 6)** • Necessary to the enterprise, infrequently accessed, cost sensitive • Online performance, high availability, less than eight hours of recovery	99%	3 hrs	1 hr
2-3	**Nearline Data (RAID 6)** • Non-Changing Data, Backup/Recovery - Unmanaged archive, cost sensitive • Disk performance, automated retrieval	96%	24 hrs	8 hrs
3-4	**Compliance Data (RAID 6)** • Managed Archive • Enforced record retention and verifiable data integrity Discovery • CAS Classification, index and search capabilities • Audit	100%	48 hrs	0%

THE BASICS OF A PROTECTION ARCHITECTURE

The protection architecture is an infrastructural organization composed of layers of protection as a tradeoff against time and money. As an example, a subsystem RAID protection scheme is great for an individual disk failure, but it has no value if a water main breaks on top of the disk array and the entire subsystem is lost. Also, if a virus invades the subsystem, or if critical files are lost, deleted or corrupted, there must be a way to recover. There are four major Protection Tiers in an effective architecture with some areas of differentiation. Protection Tier-1 provides volume failover. If the primary volume fails, the server is able to recognize and "failover" to a surviving volume. Protection Tier-2 allows for rapid restart in case of a hardware failure, data corruption or data loss by using mirrored or “point-in-time” copies of volumes. Protection Tier-3 considers backup and recovery when “point-in-time” copies have also failed and recovery is necessary versus a restart. Protection Tier-4 is a logical layer of protection used for compliance regulations, governance and long-term archiving.

Each tier features tools that are used to provide local layers of protection. Architectural adjustments must be made to accomplish DRBC. This is differentiated from what is done locally by adding the ability to move data in each of the Protection Tiers to a secondary site. There are tools and techniques available to move data to a remote location for each of the tiers in the protection architecture

COLD, WARM, OR HOT SITE

There are also ways of looking at the various Protections Tiers largely based on the frequency with which you update data. These views of the Protection Tiers are measured as Recovery Point Objectives (RP). RPO determines how much data is willing to be put at risk between updates. This is sometimes referred to as “Cold,” “Warm” or “Hot” sites. A “Cold” site environment is turned on once to have software installed and configured then turned off until needed. This view of Protection Tier-1 assumes you are bringing data with you.

A “Warm” site is a backup site that is turned on periodically to receive updates from the server being backed up, thus reducing the amount of data at risk. You still have to bring data with you, or have it located nearby for a recovery. “Warm” servers are often used for replication and mirroring as “point-in-time” copies of data, mapping to Protection Tier-2 in the protection architecture.

A “Hot” backup site receives constant updates and is on “Hot” standby, ready to take over immediately in the event of the need for a failover. This view of Protection Tier-3 offers little to no risk to data. If the “Hot” site is nearby, or if certain technology options are used, it can even be configured to share workloads and thus, simply picks up the workload in the event of a primary failure. In this event, the user is completely unaware of a failure and business continuity is absolute. Protection Tier-4 adds logical layer protection to this type of “Hot” site architecture to ensure compliance.

AMOUNT OF DATA

RECOVERY TIME OBJECTIVE

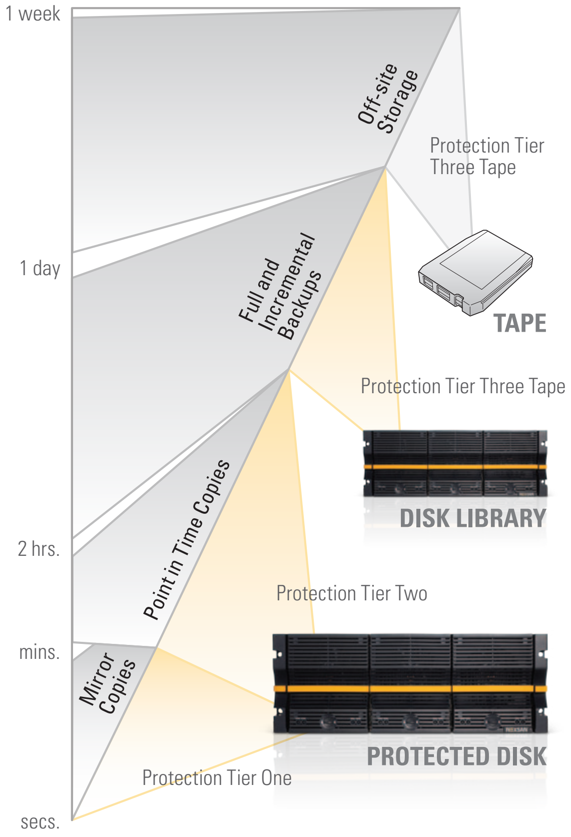

BEST PRACTICES

A well-founded DRBC protection architecture starts as a classification process to determine the value of an application and its data when it is running and the impact of an application when it stops. This determination is used to set the Recovery Time Objectives and Recovery Point Objectives. The DRBC plan uses various software and hardware technologies to move data from a local Protection Tier level to the same remote tier level. Once that is done, the system is configured as a Highly Available Clustered System. A recovery at the remote site is equivalent to what it looks like locally, including the appearance that no failure has occurred whatsoever.

From a local point of view, Protection Tier-1 would typically represent the smallest amount of overall data. In this case, mirroring data with a hardware RAID or software solution for failover is cost justified. That is good news as it reduces the amount of mission critical change data that must be written to a remote site. Even so, if the remote site is farther than 6.2 miles away, using Fibre Channel is out of the question. iSCSI must be considered for synchronous or asynchronous data transfer at distances greater than 6.2 miles. Beyond the physical interface, there are various methods that can asynchronously move applications and file systems data great distances from a system point of view.

Tier-3 of the protection architecture would use anything from copies on a disk library, for reliable, high-performance recovery that ranges from minutes to hours, to offsite copies on tape allowing recovery that would be measured in days or weeks.

What is important about this architecture and strategy is the idea that you must first map the value of data, which offers a recover point and time objective. That can then be mapped into a solution capability requirement used to choose the right technology that will provide an effective DRBC architecture with an effective business rationalization plan.

DISK BEST PRACTICES

Consider some best practices using a Nexsan disk solution within the protection architecture and for DRBC. Nexsan's approach to the protection architecture is grounded in proven principles of matching the right technology to deliver the right data, at the right time, at the right cost. Because Nexsan designs and builds Easy, Efficient, Enterprise-Class storage, users can depend on data meeting Service Levels for Protection and Recovery objectives at a price point that will please the business.

Since many of today's mid-market data centers have, or will soon deploy, virtualized servers, Nexsan has focused on implementing integrated virtualization capabilities for management that also aid in the objectives of DRBC. For example, Microsoft offers a number of capabilities for managing virtual machines and virtual storage. Nexsan management software integrates with these Microsoft capabilities. It is through this integration with Microsoft virtual storage management capabilities that Nexsan supports not only Microsoft's Hyper-V virtualization system but also virtualization systems from other vendors including VMware, Citrix and Symantec.

Systems with multiple controllers need to handle the dual issues of array ownership and SAN load balancing in order to optimize performance. Host systems with advanced MPIO software, such as VMware's ESX and ESXi, along with Windows Server 2008, can access a Nexsan storage array and discern the subtle but important difference between an active disk port and an active port on a service processor - thanks to Nexsan's implementation of ALUA. Because I/O requests are sent only to active service processors, this enables optimal performance and avoids the overhead of switching controllers.

Nexsan has achieved an unparalleled level of reliability, adding tremendous value at both local and secondary recovery sites. Nexsan can be used as storage building blocks when matched within a larger scheme of technologies such as virtualization, advanced file systems and high-performance applications. Together, these create a complete solution that can be used to plan a rock-solid DRBC for any tier within the mid-market. DRBC applications that involve distances outside the reach of fibre channel, Nexsan offers a storage array specifically for iSCSI networked storage. This enables a simple implementation of a DRBC plan without the hassle of requiring advanced file systems and applications. Top and middle tier mid-sized organizations may already have advanced file systems and applications in place and may not need Nexsan's iSeries product for iSCSI connectivity. On the other hand, some middle and bottom tier mid-market customers may find Nexsan's iSeries ideal.

Top-tier, and many middle-tier, mid-market customers will already have virtualization, clustered high availability file systems and applications in place. Nexsan provides a variety of purpose-built, ultra-high availability disk arrays that serve multitudes of customers seeking best-in-class solutions that are cost-aligned with the needs of constrained business practices common in the market.

Customers looking to include a long-term archive as part of the DRBC plan find that Nexsan offers the very best technology available anywhere: the Assureon Archive Storage System. Assureon enables organizations to meet regulatory demands while ensuring data does not corrupt or worse yet, deleted before its time. If Assureon discovers lost or corrupted data during regular background maintenance sweeps, it will fix it first then let you know that a problem has been handled.

CONCLUSIONS

Achieving the right protection architecture for your Disaster Recovery Business Continuity plan requires serious attention to a myriad of details. IT professionals should never have to worry about their storage. Nexsan offers reliable solutions as a trusted partner to help build a DRBC architecture regardless of where your organization fits in the mid-market. Nexsan is known for Easy, Efficient and Enterprise-class storage - an important consideration when the survivability of data is the fulcrum that balances the survivability of your business.

ABOUT NEXSAN

Nexsan® is a leading provider of innovative data storage systems with over 10,000 customers worldwide. Nexsan's pioneering hybrid storage systems combine solid-state technologies, spinning disk storage and advanced software to deliver radical new levels of performance and capacity at lower cost. The company's advanced technologies enable organizations to optimize traditional, virtual and cloud computing environments for increased productivity and business agility. With more than 28,000 systems deployed since 1999, the company delivers its data storage systems through a worldwide network of solution providers, VARs and system integrators. Nexsan is based in Thousand Oaks, Calif. For more information, visit **www.nexsan.com.**