NEXSAN

SOLUTION BRIEF
KEY CONSIDERATIONS FOR LONG-TERM, BULK STORAGE

IT organizations must store exponentially increasing amounts of data for long periods while ensuring its accessibility. The expense of keeping large amounts of data on very expensive tier-1 disk is untenable. Virtualization technologies help to partially address this challenge by consolidating data storage. However, organizations that carefully analyze their storage usage and performance needs, and implement a tiered storage architecture, can optimally align this storage infrastructure with their business requirements.

NEXSAN

TOP FIVE LONG-TERM BULK STORAGE BEST PRACTICES

1. **MAXIMIZE ALLOCATION AND UTILIZATION**—Most Storage has allocation waste and utilization Waste. Implement a Storage Resource Manager that can put wasted capacity back to work and reduce the costs of tier-1 storage.
2. **VIRTUALIZE STORAGE**—Storage virtualization enables thin provisioning and storage resource pool sharing. This significantly decreases idle and wasted storage.
3. **ADOPT TIERED STORAGE**—Tiered storage aligns usage and performance requirements with the capabilities and cost of the storage device. Use a classification exercise to drive migration into tiered storage. Storage Resource Manager greatly simplifies the adoption of a tiered management strategy.
4. **MANAGE FROM THE STACK**—Virtualization Hypervisors and excellent Storage Resource managers available today in the server stack have made advanced management capabilities in the array redundant, complex to manage and expensive. Managing storage from the stack dramatically streamlines management and reduces costs.
5. **USE ENTERPRISE-CLASS STORAGE**—Storage that is purpose-built for a stack-managed environment eliminates costs for redundant management capabilities. This storage must be high performance, ultra-reliable, easy to manage and efficient to operate.

[1] Source: IDC - The Expanding Digital Universe

NEW TECHNOLOGIES MEET GROWING STORAGE REQUIREMENTS

One of the biggest challenges for IT professionals is storing large amounts of data efficiently and safely over long periods while maintaining quick and easy accessibility.

The major factor contributing to this IT concern is the explosive growth in data that needs to be stored. According to "The Digital Universe," IDC forecasts, "Between now and 2020, the amount of digital information created and replicated in the world will grow to an almost inconceivable 35 trillion gigabytes as all major forms of media – voice, TV, radio, print – are converted from analog to digital. This year, the Digital Universe will grow almost as fast to 1.2 million petabytes, or 1.2 zettabytes. This explosive growth means that by 2020, our Digital Universe will be 44 TIMES as large as it was in 2009. Cloud-delivered services and embedded systems that are part of logistics item tracking, text messaging, email, documents, pictures, video, social networks, and so on; have caused the amount of files that we must manage to grow by a factor of 67."

While individuals at home, at work or on mobile devices create an estimated 70% of all digital information, most of that content ends up in storage managed by IT departments[1]. Thus, IT management teams become responsible for building the storage management infrastructure as well as managing and protecting the content.

Yet, while the exponential growth of files adds considerably to storage administrators' already heavy and complex workload, the number of administrators available to perform this work is unlikely to meet the demand. According to IDC, the number of IT professionals worldwide will only grow by a factor of 1.4 through 2020[1].

To manage more data with a slow growing workforce, IT organizations will need to improve operational efficiency. The most significant efficiency technology to emerge so far has been the Cloud, which takes advantage of virtualization technologies. Virtualization consolidates workloads to increase resource utilization and minimize expensive idle time. IT can also dynamically allocate hardware resources when and where they're needed, eliminating the need for over-provisioning to meet peak load demand. Consolidated systems improve IT efficiency by reducing the time needed to maintain and manage storage systems.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | **p.** 866.4.NEXSAN **f.** 805.418.2799 | www.nexsan.com

While virtualization efforts have historically focused on servers, more recently, IT organizations anxious to align business needs with infrastructure costs have begun to look at storage virtualization. Storage virtualization extends the advantages of consolidation and optimization to the storage infrastructure and has the potential to double the data center cost savings organizations achieve through server virtualization alone. One emerging storage virtualization technology that has proven particularly efficient is stack (server) management, which scales by adding enterprise-class storage arrays to the storage network.

This type of storage management scheme takes advantage of storage management capabilities built into the virtualization hypervisor. This means rather than purchasing arrays that include redundant storage management capabilities, a storage array can include a thin layer of software, leaving the bulk of storage management capabilities to the virtual server stack. This allows administrators to efficiently manage larger amounts of storage and reduce the cost of purchasing and managing storage.

However, in order to make the most of storage virtualization technology, administrators need to fully understand the storage needs of their environment. In particular, they must determine the nature of the data stored, reference behaviors and performance requirements, then use this information to build a classification scheme that defines the value of the data and the service levels required to support application performance. Only by creating an IT infrastructure that aligns accurate data storage and performance classifications with the capabilities of the storage hardware used can organizations efficiently and cost effectively house and protect the massive amounts of data needed to store for the long term.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | **p.** 866.4.NEXSAN **f.** 805.418.2799 | www.nexsan.com

NEXSAN

CLASSIFICATION

Protection Tier	Classification	Performance
1	**Mission Critical Data** • Critical to an enterprise, continuous access • Highest performance, near zero downtime	Highest
1-2	**Business Critical Data** • Very important to the enterprise, frequently accessed • High performance, high availability, less than four-hour recovery	High
2	**Accessible Online Data** • Necessary to the enterprise, infrequently accessed, cost sensitive • Online performance, high availability, less than eight hours of recovery	Medium
	• Non-changing data, backup recovery - unmanaged archive, cost sensitive • Disk performance, automated recovery	

DESIGNING YOUR LONG-TERM BULK ARCHITECTURE

While organizations take advantage of many types of storage, this paper will focus on long-term, bulk storage as it represents the vast majority of data — and optimal management of this type of storage will have the greatest impact on meeting an organization's long-term storage requirements.

Typically, long-term, bulk storage is used for unstructured data (e.g. audio, video, email or documents) which represents 80% of all data. Organizations also have a financial justification to keep this data on disk for fast retrieval. Therefore, the storage infrastructure must allow operating systems, applications or users to actively reference the data. Long-term, bulk storage excludes data stored in databases, object oriented data or data that must be retained for very long periods (or archived for legal conformance or discovery).

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | p. 866.4.NEXSAN f. 805.418.2799 | www.nexsan.com

TIERED STORAGE

NEXSAN E60X

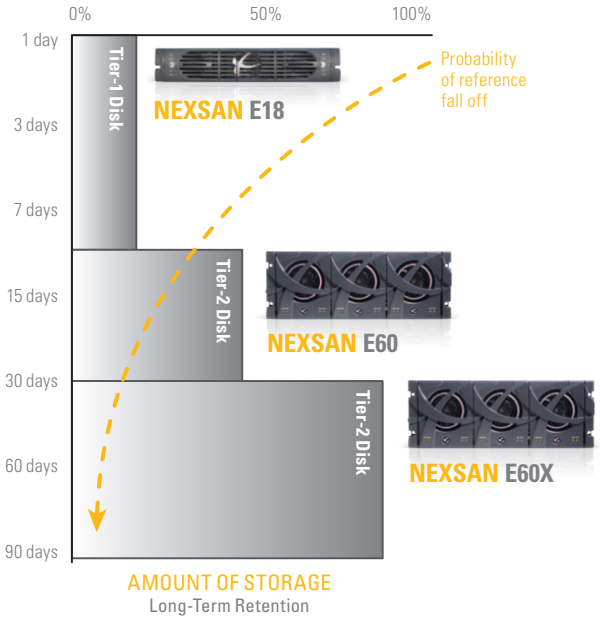

Age in Days	Probability of reference
1	75%
3	50%
8	25%
30	2%
90	Near 0%

How does one build an efficient bulk data repository that optimizes the use of stack managed virtualized storage?

The first step is to classify the performance service level for the types of applications the repository will support and the unstructured data that it will store. Unstructured data can have highly variable performance requirements. For instance, some businesses may consider an Exchange server to be a mission critical application, requiring very high performance service levels. Conversely, documents may have a very low reference probability and therefore, require a low service level for performance. Classifying the data according to required performance levels not only helps to define the technical requirements for the infrastructure, it will also expose the optimal management methodology.

The second step is to understand how data will be used initially and over time. This helps to determine how the data should be stored and whether any business-justified reasons exist to migrate the data to a different tier of storage over time. Typically, the reference behavior of most user data falls steeply over time. After the first day, the probability that the record will be referenced again is about 75%; after 90 days, the probability of reference is near zero.

After the analysis, an organization might determine that some data should be captured onto tier-2 storage and remain there for its entire lifecycle. Other data might originate on high-performance, tier-1 storage, which is the most feature-rich and expensive type of storage, then migrate it to tier-2 storage (which has terrific performance and reliability at less than half the cost) as its access requirements decline.

Information Lifecycle Management is the term used to describe a tiered storage infrastructure that is managed by hierarchical storage management capabilities in the stack. A tiered storage architecture allows organizations to optimally align the performance and availability characteristics of their storage hardware with their business requirements.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | p. 866.4.NEXSAN f. 805.418.2799 | www.nexsan.com

When the IT organization does not actively classify and manage the location of unstructured data within their infrastructure, storage systems typically contain considerable inert data that has not been referenced in six months or more. Inert data occupies, but essentially wastes, inordinate amounts of capacity. It makes far better sense to move low reference probability data, or inert data, off expensive tier-1 disks and return the capacity for use by high performance data. This allows the organization to "grow" its storage capacity without cost.

Storage classification also improves the efficiency of the protection architecture. When administrators attempt to restart or recover a failed storage volume, they often find multiple copies of disk volumes, along with all of the inert data on them, replicated multiple times. Commonly, too many copies are made and captured on expensive tier-1 storage when they could be housed on tier-2 long-term, bulk storage.

The biggest objection storage administrators typically have to adopting a cost-effective tiered storage architecture is the time and effort necessary to set up the management controls. The good news is that cost effective software greatly simplifies this process. Storage administrators can use a migration manager running in the stack to apply simple classification guidelines to locate, move and manage hierarchical tiered storage.

BEST PRACTICES: IMPLEMENTING TIERED STORAGE WITH NEXSAN

When organizations properly classify data, design tiered management based on expected usage and performance requirements, and store data on a stack-managed virtualized storage array, they can improve storage management efficiency and meet their storage growth challenges.

But as organizations increase their virtualization, they have one additional concern - increasing virtualization drives higher I/O demands. Running databases and long-term bulk, storage on virtualized systems, such as those managed by Symantec's Storage Foundation, VMware vSphere or Microsoft Hyper-V, increases I/O load because virtualized systems push more data over fewer I/O channels. Fortunately, these virtualized environments include management capabilities that manage I/O paths to maintain high performance.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | **p.** 866.4.NEXSAN **f.** 805.418.2799 | www.nexsan.com

Nexsan storage takes advantage with built-in I/O path management capabilities for SAN storage, or iSCSI for those that do not have or want SAN connectivity. Nexsan supplies modular, stack-managed arrays for tier-2 long-term, bulk storage. The Nexsan modular approach not only adds scalable capacity but also performance. Nexsan arrays have four 8 Gb Fibre Channel and four 10GbE iSCSI ports, allowing great performance to the individual array and additional performance as arrays are added. This, along with highly available controllers and built-in virtualization performance features, makes Nexsan a perfect solution for efficiently storing and managing the expanding digital universe. Nexsan arrays are also very dense, allowing more drives per rack unit, while built-in AutoMAID capabilities selectively reduce power consumption rates on the array, ultimately improving energy efficiency by 85%.

Too often, IT departments are forced to purchase, house, power and cool more hardware than needed. Administrators waste considerable time and effort managing complicated storage systems with dead-end growth paths. The Nexsan E-Series has been purpose-built for the midmarket to drive down the cost and time associated with managing data by delivering the leading solution to long, term bulk storage.

Even organizations that have not taken the virtualization plunge will find that features such as built-in, easy management and performance, scalability and ultra reliability are a perfect complement to their storage strategy. For example, a client might start off with a Nexsan E18™, which holds eighteen 3TB drives and provides 54TBs of raw capacity in 2U of rack space. As their business grows, they can easily install a Nexsan E60X™ expansion chassis to add a whopping 180TB of storage in 4U of rack space to grow into.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | **p.** 866.4.NEXSAN **f.** 805.418.2799 | www.nexsan.com

CONCLUSIONS

The amount of data that needs to be stored is growing exponentially and that growth is expected to continue. Vendors have responded with storage management systems, storage capacities, file management capabilities and performance that would have been unimaginable just a few years ago. Purpose–built, enterprise-class arrays managed from the stack for larger and more efficient capabilities, or within the array for smaller challenges, are an important step toward meeting the need to store data that must be both actively accessible and held on disk for extended periods.

Organizations with complex storage and server infrastructures will find that the sooner they develop a classification scheme to manage changing data reference behaviors using a tiered infrastructure, the sooner they will dramatically reduce overall expenses.

For those that have chosen server virtualized environments, storage virtualization is now built-in to those products. Nexsan arrays take advantage of these virtualized management capabilities and complement them with high performance, scalable, purpose-built arrays for the midmarket with the right feature set.

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.**

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | **p.** 866.4.NEXSAN **f.** 805.418.2799 | www.nexsan.com