Storage Switzerland, LLC

Storage Switzerland White Paper

Storage Infrastructures for Big Data Workflows

Sponsored by:

Quantum®

Prepared by: Eric Slack, Sr. Analyst

May 2012

Introduction

Big Data is a term used to describe data sets which have grown so large that traditional storage infrastructures are ineffective at capturing, managing, accessing and retaining them in an acceptable time frame. The thing that separates Big Data from simply a large archive is the need to process these data sets or to provide file access to multiple users quickly.

Some Big Data use cases involve analytics, the computer-based analysis of large amounts of relatively small data objects, for the purpose of pulling business value from that information. Many of these involve files supporting transaction analysis or automated event processing, such as database or web analytics, which won't be addressed in this white paper.

Instead, this paper will deal with another form of Big Data that supports file processing workflows, often sequential in nature, where large files are shared by knowledge workers to create digital products, support research and perform analysis to increase productivity. Also considered will be Big Data supporting large file analytics in which files are shared by large, high performance compute clusters to support complex analysis and drive business decisions.

Big Data File Processing Workflows

Some of the industries using large file, Big Data sets in these two use cases include Media and Entertainment, Life Sciences, Healthcare, Defense/Intelligence and Oil and Gas. One of the things that makes this topic compelling for infrastructure suppliers is that extracting value from the collected data usually involves a time constraint, meaning that Big Data is often being ingested so quickly, it can't just be put into a large, traditional backup repository or archive. In addition to data protection, these infrastructures must have the performance to provide fast access and throughput to satisfy users' "just in time" needs for files a collaborative workflow.

Another compelling aspect comes from the word "big". Very large data sets require some foresight into the design of their infrastructures since petabytes of data can't be easily manipulated or moved to facilitate a major change in storage or data handling systems. Like a construction project, once the foundation has been poured, a building is very difficult to move around, if not impossible. Unknowns created by future requirements can mean risk, since structural modifications are very difficult, so Big Data systems must include flexibility in order to address those risks.

Infrastructure Requirements

The very nature of Big Data, the realities of its size and the analysis and workflows it must support, puts many demands on the storage infrastructure as well. Capacity and performance efficiency must be maintained in order to keep the costs of storing and handling such large amounts of data under control. Also, Big Data can include a record of events, such as surveillance video or involve costly data acquisition, such as oil & gas seismic exploration, and need be kept for long periods of time, bringing a need for longevity of the storage system and long term data integrity.

These and other requirements of a Big Data storage infrastructure will be examined in the rest of this report. In addition, this white paper will look at Quantum's StorNext File System and Storage Management software and how it can form the core of a Big Data storage infrastructure that addresses these requirements.

Flexibility

One of the challenges that Big Data brings is the requirement to support many different data types, or at least have that ability. Tasked with finding ways to pull business value out a given data set, users are creating more ways to cross reference those data. Mergers and acquisitions, as well as pure financial motive, may drive an organization to purchase storage or applications from one vendor one day and a different vendor in the future. This can mean combining data sets stored on systems from different vendors and sharing the resulting file pool with clients using a variety of applications running on different operating systems.

Similarly, applications used by the industries that generate Big Data are evolving, bringing new file types with them, even new platforms - and usually a need for more capacity. Advancements such as the use of 3D video in the Media & Entertainment industry and 3D images to be analyzed by geophysicists are examples. When a new application is implemented it should be able to access the current file system so that existing files can be processed while the current applications are being supported.

Whether it's new applications, different file types, multiple storage platforms or something else entirely, the task of maintaining a Big Data storage infrastructure into the future, where the requirements are largely unknown, carries some significant risks. In order to reduce those risks, this infrastructure must be extremely flexible.

Heterogeneous Environments

The StorNext file system supports truly heterogeneous environments by connecting Linux, Windows, Mac, and even UNIX hosts to the same files via a SAN, LAN. This enables the widest possible compatibility between applications or users and the files

that they're required to process today. This flexibility also reduces the risk that a future application, data type or compute platform won't be supported.

Big Data's volume may quickly outgrow existing storage, causing purchasing organizations to look for affordable capacity wherever they can. This can lead to the acquisition of storage systems from different manufacturers and a need to combine the capacity on diverse platforms. Ideally the organization may want to explore many different storage options in order to keep up with Big Data's appetite for capacity. Unfortunately, its scope and scale doesn't lend itself to nimbleness and data migration is usually out of the question for data sets this large.

Storage Virtualization

StorNext has a virtualization layer that abstracts the physical location of storage keeping those details hidden from users and applications on the front end and enabling capacity to be added on the back end transparently. This could be SAN arrays of high-performance 'tier one' disk, economical, high-capacity arrays or even storage systems or NAS filers that were decommissioned after the last refresh cycle.

This kind of flexibility is ideal in order to stay abreast of the inevitable changes in applications, storage platforms and workflows that Big Data infrastructures will see. It also supports the continual need for affordable capacity and non-disruptive upgrades that a storage system will experience as it's kept active for decades.

Big Data environments may also evolve. For example, a relatively simple file sharing workflow between clients using a single platform may grow into one that includes long term archiving and data protection with multiple OSs. In these situations the file system infrastructure must support the additional data services needed. StorNext provides multiple technology choices for protecting data such as tiering, archiving (to NAS and tape as well as bulk block storage), deduplication and replication for offsite data protection.

File Sharing and Collaboration

In addition to storage virtualization, a Big Data infrastructure must be able to support file sharing across operating systems on the front end and across storage systems on the back end. This can include workflows in which multiple users, doing diverse tasks on different platforms, with different software need to share the same files, often concurrently. A Big Data file system that excludes one application or platform can cause manual workarounds to the data flow process. These 'sneaker net' types of solutions can result in reduced productivity and an increased potential for error or data loss. To prevent this StorNext enables clients, regardless of OS, to access the same files. Currently, these include multiple variants of Windows and Linux, plus UNIX and Mac OS X.

In sequential file processing use cases, like those used for video editing, productivity can be directly related to the storage system's ability to access and transfer a set of shared files quickly. Time is money and as soon as one person is finished, another may need to start work on the same file or set of files. In environments, where users are running applications like 3D editing on high definition video the storage system may not have the horsepower to stream these very large files fast enough without dropping frames or support multiple workflows. Because StorNext enables performance of up to 90-95% of the underlying SAN or LAN infrastructure, it has built strong proof points of success in the Media & Entertainment industry where these high throughput workflows are common.

Performance

Traditional NAS devices which run CIFS and NFS protocols over an IP network can be sufficient for regular types of files but may be inadequate in these kinds of high performance environments. As a SAN file system, StorNext provides fibre channel performance (often hundreds of MB/s per single stream) to server clients. Also, a StorNext LAN client protocol that was built for large block transfers through StorNext LAN gateways, allows access at near Gb Ethernet speeds to these same files over the LAN, significantly faster than NAS storage systems running on IP-based protocols.

StorNext helps maintain performance as the infrastructure grows and storage capacities expand by separating the file system metadata controller from the data movement function that's providing access to storage archive tiers. By dividing this process across multiple dedicated servers StorNext allows the system to maintain performance and better match changing workloads.

StorNext's architecture enables another option, called the Distributed Data Mover (DDM), which can improve performance as well. DDM offloads the data movement operation from the metadata controller to an alternate dedicated compute engine. This allows for faster file retrievals during periods of heavy system activity and frees up cycles on the metadata controller which can be applied to other operations.

Scalability

Big Data environments can mean constant data growth and a requirement for the storage infrastructure to expand to what may seem like an unlimited capacity - and do this easily. In industries that deal with imagery or other visual data, resolution is continually increasing, driving file sizes up in the process. As an example, a typical 3D movie can involve multi-TB sized files. But the storage requirement for these projects doesn't just include the finished product, it may also include the copies made during the interim production steps. These must be made available to operators at multiple processing stations along the way and then archived after the project is complete.

In addition to providing the storage capacity to support these Big Data applications, the infrastructure must also support extremely large numbers of files. This requires an expandable file space which can be laid across multiple physical storage devices and extended seamlessly as data grows. StorNext's global namespace can support file systems in the multiple petabyte range and expand dynamically to file counts in the hundreds of millions.

Cost control

Cost is always an issue when storing large data sets, especially ones that can grow almost without limit, and this is certainly true with Big Data. One large genomic sequencing vendor with PBs of data studied use patterns and found that 40% of its data files had not been accessed in over eighteen months and 60% in eight months. One way to minimize costs is to leverage policy-based tiering features which only move files needed for the current projects onto the fastest (and most expensive) storage areas so they're accessible to support sequential workflow processes. Other methods include data reduction technologies, like deduplication, or using a high density, low cost recording medium like tape.

With extremely large file sizes, as is common in satellite imaging or genomics applications, storage costs can be reduced by truncating files and storing only a portion on a high performance tier, with the remainder resident on an archive disk tier, or even tape. In this way, the user or application requesting the file can get started with this initial segment and have the rest of the file streamed up to the performance tier concurrently. Since StorNext controls the file system and the data management function, it can accomplish this process transparently.

Tiered Storage

A tiered storage architecture is an effective strategy for creating affordable capacity and increasing the scalability of storage systems. It enables higher capacity systems to be added, like arrays with multi-TB SATA drives, and can include tape as well. In order to integrate these different storage platforms into the common data pool, the Big Data storage infrastructure needs a mechanism that can move files between storage tiers based on predefined policies, while maintaining the single namespace for its users. This enables the system to keep the most active data sets on the highest performing storage assets, and move the rest off to lower cost capacity or an archive.

In Big Data environments that support file sharing workflows, a tiering mechanism can move the files associated with a project off to an archive tier when their access levels indicate the project's complete and save that premium space. Files movement policies are set at the directory level to accommodate scheduled workloads. StorNext's

distributed architecture helps maintain file system performance throughout these data movement operations.

Tape

These ‘capacity tiers’ can also include tape, the most cost effective storage medium available. The current generation, LTO5, can store over 1.5 TB per cartridge, uncompressed, giving Big Data infrastructures the density to archive enormous numbers of very large files in a relatively small data center footprint. Compared to even the most cost effective disk arrays, the use of tape archives can translate into lower operational costs since idle tapes draw no power and don’t require cooling. In addition, tape makes an excellent ‘deep archive’ tier since it remains viable for longer periods of time than does magnetic disk technology.

Long term viability

Big Data archives will often have to be stored for long periods of time, maybe indefinitely. Given the amount of money and other resources that can be put into applications like genome sequencing or feature films, Big Data storage infrastructures can represent a significant investment which will need to remain viable.

Data Protection

A Big Data infrastructure should provide data protection assurance so that this investment is appropriately cared for. Obviously, a traditional backup process can be impractical, since making weekly, or even daily backup copies of large numbers of very large files could take too long. StorNext provides this protection with a process that makes file copies off of primary storage continuously.

When file accesses have stopped, a data protection copy is made (actually, up to 4 copies) and moved to a secondary storage tier, like tape. Then, after a certain number of days, a second policy marks that primary copy as a candidate for truncation. The actual truncation of the file from the primary disk tier is carried out when capacity thresholds on primary storage are reached. The result is at least one full-time, full-size copy of each file stored on long term archive media, as soon as that file has become inactive. StorNext maintains that copy on primary storage as long as possible and can restore it from the archive copy automatically, when accessed again.

Data Integrity

StorNext provides a data integrity feature that embeds checksums into archived data to help maintain viability. Quantum’s AEL Archive technology also includes tape integrity

checking that regularly tests each cartridge for wear to confirm that its data can be reproduced reliably. When a piece of media is found to contain an excessive number of data errors it can be automatically copied to a new cartridge, so it won't further degrade and risk data loss or fail during use.

Highly scalable storage products are coming out from new vendors on a regular basis, so the storage infrastructure must be able to support platform replacement if it becomes necessary. StorNext's comprehensive archive and data management functions can support the use of multiple vendors' platforms and the data handling this requires. From an industry perspective, StorNext has become a standard SAN file system solution in the Big Data space, with over 60,000 file system clients deployments and 500PB of data under license.

Summary

Big Data creates challenges for the infrastructure that it's stored on and the people who manage it. Obviously, to support growth, capacity must be available, but also performance, in order to meet the access requirements of users and applications. In use cases like those supporting file processing workflows, for example, throughput is essential to deliver large files to users in an acceptable time frame. The file system and archive infrastructure must provide this performance and capacity expansion without creating a burden on the IT staff.

These infrastructures must also be able to maintain very large data sets for a very long time, and do so cost effectively. This means providing assurance that data integrity is maintained while the physical infrastructure is upgraded, updated and expanded. It also means design and operational efficiency so that costs, especially future costs, are controlled.

Big Data storage infrastructures must also be flexible, so that they can support multiple file types and client-side operating systems on the front end and multiple storage platforms on the back end. In reality these large infrastructures can eventually become a consolidation of many disparate storage devices as they accumulate more diverse data sets. StorNext SAN and LAN File System and Storage Management suite is designed to meet these considerable requirements of Big Data today, evolve to meet the requirements of tomorrow and do so cost effectively.

This white paper is sponsored by Quantum