

# Geospatial Imaging – Data Management and Archiving

## *Storing the Earth without breaking the bank*

### **Overview**

There is nothing more important for understanding our environment than understanding the Earth we live on. Today's technological breakthroughs have allowed us to study and capture more information about the earth than ever before. There are a multitude of methods for capturing this information including seismological information, weather models, and a variety of remote sensing techniques. While all of these earth observation techniques provide unique data about the earth, they all have one thing in common: they consume vast amounts of data. This data's value varies over long periods of time with the most current data universally delivering the most value as it relates to current events. As time progresses, the data becomes less current and, hence, less valuable. However, when a customer is interested in tracking changes over time, suddenly the older information is just as valuable as the new data. The combination of old and new data is even more valuable as the two pieces are merged to collectively answer customer's questions. This paper focuses on the management challenges in archiving and repurposing data by utilizing examples in geospatial imaging environments, most notably with aerial and satellite remote sensing. The challenges, and suggested solutions, however, apply to any environment generating and analyzing large quantities of data.

### **Business Needs**

More and more corporations are "data driven." A company's success is often driven by how it manages its data. This is particularly true for businesses whose primary product is data based such as companies capturing, manipulating, and creating unique assets from geospatial imaging information. These organizations specialize in creating airborne or space-based remote sensing platforms that capture images of the Earth using sensors including infrared, electro-optical (visible), and radar.

As with any of today's technology-driven businesses, the drive to improve the technological base, stay ahead of the competition, and satisfy customer requirements is unrelenting. For these geospatial companies, this translates to higher resolutions and larger scan areas resulting in significant increases in file sizes. And with improvements in communication technologies, airborne and space-based platforms transfer more imagery data, at higher data rates to ground stations for processing, analysis, and delivery to customers.

This sheer volume of data presents significant challenges to geospatial imagery businesses. These businesses need to maintain focus on their core business of capturing and providing value added imagery data to their customers but storage and data management issues are consuming a larger portion of information technology

resources and budgets. These organizations are challenged to adequately design and economically operate long-term state of the art storage systems.

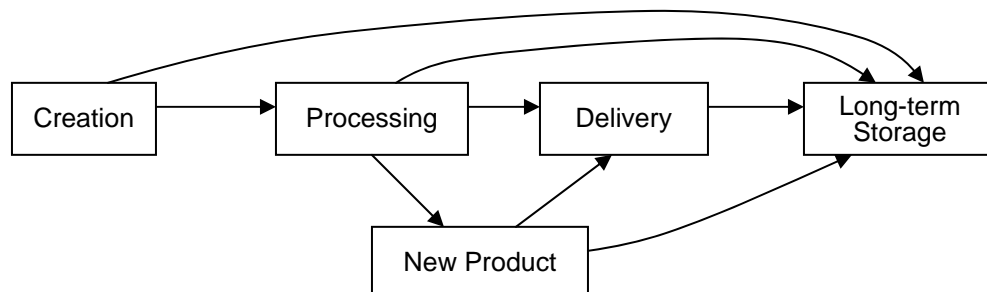
## Solutions

Today's storage devices only provide point solutions to specific portions of large data repository issues. For example, high-end Fibre Channel based RAID solutions provide excellent performance profiles to meet demanding throughput needs of geospatial acquisition and processing systems. Serial ATA (SATA) based RAID systems provide large capacity needs for longer-term storage. Tape storage provides enormous capacity without power consumption. Each of these storage technologies also has unique downsides to their utilization for certain data types. Utilization of Fibre Channel RAID for all data types is cost prohibitive. SATA-based RAID systems do not have the performance or reliability to stand up to high processing loads. Tape technology does not facilitate transparent, random access for data processing.

Architecting a solution leveraging the aspects of each device's capabilities into a seamless, scalable, managed data repository is the focus of the remainder of this white paper. Quantum's StorNext software, which provides high speed data sharing and long term archiving, is detailed to help readers identify how to architect a final solution.

## Workflow

To properly architect a comprehensive data storage solution, it is important to understand how data is ingested, stored, processed, protected, and maintained over a period of time. This is often referred to as the "data lifecycle." While the term data lifecycle may imply a fixed, linear path for a particular digital asset, the reality is that each business, and often each data type, has unique lifecycle requirements. Furthermore, as business environments change over time, so may the data's lifecycle as well as the characteristics of each lifecycle stage. A data management solution must be adaptable to changing requirements. In general, a digital asset follows the following path:



Each stage has unique processing and storage requirements which must be considered both independently to meet the specific requirements of the lifecycle stage, as well as broadly to ensure that the process flow is efficient

and cost-effective. It is also important to note that the value of the digital asset at each stage in its life may necessitate the long-term storage or protection of that asset as indicated by the arrows to “Long-term Storage.”

## ***Creation***

A geospatial image is collected using sophisticated remote sensing techniques from airborne and space-based platforms. These raw images can range in size from tens of megabytes to tens of gigabytes with a single day’s ingest ranging into the terabytes. After acquisition, images are transferred from the collection platform to a ground-based storage system - often at high speeds with real-time requirements. These raw images are one of the most valuable assets an organization can have as multiple customer products can be created from a single raw image through different processing algorithms. As a result, there is a critical demand for collecting and storing data within a shared, high speed repository concurrently accessible by multiple systems.

StorNext provides a unique capability for satisfying these requirements for high-speed data ingest from multiple sources. StorNext’s core component is a shared file system that allows multiple servers, running a mix of operating systems to directly access a single content repository. StorNext File System, SNFS, supports a breadth of operating environments allowing customers to choose the best ingest system available for their ground system interfaces and add different platforms as best of breed technologies change. Furthermore, multiple ingest systems may simultaneously store imagery to the same file system at speeds into the gigabytes per second. For example, a ground station may capture satellite imagery through SGI systems and airborne imagery through Linux systems – both to the same repository.

SNFS also provides the ability to completely de-couple presented file location from the underlying RAID storage systems by guiding data to the correct storage type. In effect, applications see a single disk volume, but underneath individual directory structures are tied to FC or SATA disk. In the case of raw imagery storage, it often makes sense to utilize high-speed full Fibre-Channel RAID systems because its performance characteristics are optimized for high-bandwidth and real-time environments. For browse images and other files, SATA storage would be used. Applications would see a single disk volume, and StorNext would steer files to the correct disk type.

## ***Processing***

After the raw imagery is stored on a high-speed RAID, the images must be processed in order for them to be useful for imagery analysts or value-added processing. Since the raw imagery is often in a proprietary format, a conversion step is necessary to transform the raw image into a format that may be utilized by commercially available software. Given the large file sizes, it is impractical to transfer these files to processing servers via a local area network. Even at gigabit Ethernet rates (up to 60MB/s), a 5GB file will take at least 83 seconds. For this reason, it is useful to employ a shared file system which allows every processing system to directly access the same pool where raw images were stored. Concurrent access is critical so that multiple customer products

can be created at the same time. Using StorNext and a Fibre-Channel infrastructure, data transfers are made directly from the raw-image file system, mounted on the processing server, directly to the application processing the data at full Fibre-Channel rates. StorNext typically utilizes 90% of the available wire rate of a Fibre Channel connection and scales linearly with the addition of multiple FC interfaces up to the speed of the underlying RAID storage system. This means that the 5GB file is transferred over a single 4Gbit FC interface at up to 360MB/s or under 14 seconds. Two FC interfaces bring the data transfer time down to under 7 seconds. Multiplying this time savings across data sets spanning multiple terabytes yields time savings of an hour with only 47 raw images (235GB). In a competitive landscape for current imagery, this timesaving equates to increased revenues and customer satisfaction. Processed images are then, typically, stored into another repository for distribution and long term storage.

## ***Delivery***

Electronic information distribution is commonplace in today's business environment. Responding to customer demands, businesses have enacted electronic distribution mechanisms to quickly transfer data, both large and small, through a variety of mechanisms including web sites and standard UNIX delivery mechanisms such as ftp. This allows customers to quickly and efficiently receive their imagery products as well as make the whole process self-service, eliminating expensive human interaction to manage the transaction.

By building the delivery repository with StorNext File System, traditional delivery packaging systems (e.g. DVD burners) as well as web and internet accessible servers can provide secure data access from a single content repository without having to replicate the data to multiple servers.

## ***Archiving and the Time Value of Data***

Every digital asset has an intrinsic value to the individual or organization storing it. In many organizations, the value of a digital asset depreciates over time. As a result, storing these less frequently used files on high end disk is not economically feasible. The following graphic, Figure 1, depicts the cost of storing the digital assets on disk or tape. The yellow and blue boxes represent the fixed cost for the storage platform; the red line represents the value of an asset over time. At time 0, an asset has its highest value which exceeds the cost of the storage platform. As the asset's daily value decreases, the cost of storing the asset on primary storage becomes higher than the value of the asset itself and it makes economical sense to move the asset to a cheaper repository.

Determining when to move a file off primary storage is based on the file's asset value curve, the red line. Some assets have a prolonged value to an organization and never leave primary storage, while others will be moved to a secondary storage tier very quickly. Building an asset value curve can be relatively simple. For most organizations, when an asset has been idle for 30 days, it is ready to be moved off primary storage.

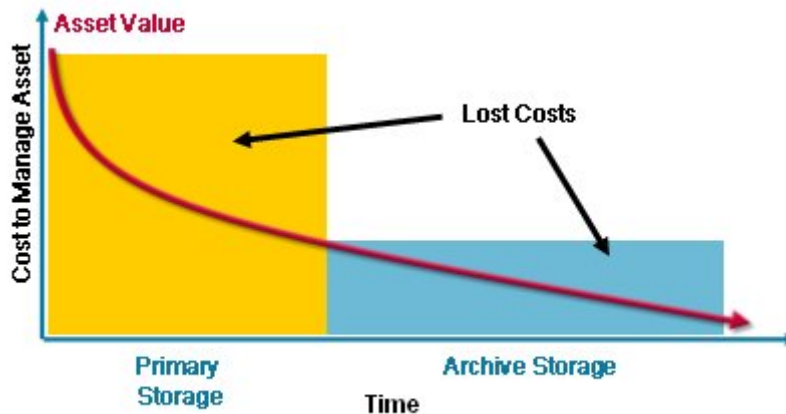


Figure 1

The crux becomes how to move the asset. Organizations cannot afford to simply remove the file from primary disk and store it on a shelf, or worse, offsite. Offline storage of material assets has caused many companies lost revenue because a file was either misplaced or unavailable within a timely fashion to build a customer deliverable. To mitigate overhead associated with offsite asset tracking and retrieval, archiving solutions were developed. Traditional data archives were software-based and composed, at best, by two types of storage: spinning disk and tape storage. The software would remove data from disk and store it on secondary storage at a specified interval. If needed, data was restored manually back to disk. The process was cumbersome but avoided the need to “shelve” tapes.

### ***Roles of Tiered Storage***

Looking at Figure 1, one other item is apparent. Even after moving an asset to an archive storage tier, there are still times when the daily asset value is lower than the storage cost and the company is paying too much to store the asset. To decrease archive costs and provide a closer match to the asset value, advanced archiving technologies now utilize multiple types of storage arranged in tiers. The types of storage may include but are not limited to Fibre Channel (FC) RAID disk, Serial ATA (SATA) RAID disk, tape libraries, and tape vaults. Figure 2 illustrates the use of multiple tiers of storage to match costs versus asset value or retrieval time. The geometry of the asset value curve is determined by the organization when considering the value of the asset against the costs of access time and the storage medium. If storage costs were negligible and independent of the type of medium, most organizations would store all assets in perpetuity on the highest performance and most reliable technology available. However, storage media can vary greatly in cost, performance, reliability, and availability. Thus new archives are embracing tiered storage as the optimum solution to economically deliver maximum storage capacity and bandwidth.

The components of a tiered storage solution vary in cost, performance, reliability, and functionality. These innate qualities of the storage medium must be evaluated against the archives required size, performance requirements, and budget to define the type and size of each tier.




Figure 2

## Storage Tier Technology Options

Identifying, sizing, and selecting the right technology option for a new infrastructure or for improving an existing infrastructure can be a daunting task. It is often advisable to engage an organization, such as Keeper Technology, that specializes in the architecture, deployment, and support of large storage environments. Making the wrong decision can have lasting effects on a business's near-term acquisition costs as well as long-term operational costs and future technology insertion difficulties.

### Spinning Disk

Spinning disk is currently the preferred front end for archives. Spinning disk offers fast random access for the storage and retrieval of digital information. Typically, Redundant Arrays of Inexpensive Disks (RAID) are used to protect against disk failures. Multiple RAID controllers and logical units (LUNs) should also be architected under an advanced file system to create a very high performance archive front end.

When considering what type of disk to use, the two prevalent disk technologies today are Fibre Channel (FC) disk and Serial Advanced Technology Attachment (SATA) disk. FC is the leader in scalability, performance, and reliability. However, SATA technology is generally less expensive to implement than FC. By carefully reviewing the requirements, a compromise between FC RAID's performance and SATA RAID's economics can be made. For large archives an initial front end tier of high performance FC RAID and a second larger tier for SATA RAID may provide the most economical solution. The size of these tiers should be based upon the temporal bandwidth requirements of the asset. However, RAID systems are designed and built in units. Thus, a minimum required bandwidth for a tier may dictate a specific volume of storage required to deliver this bandwidth and will vary with between manufacturers.

Disk systems provide a valuable high performance front end for an archive. However, disk solutions are not necessarily the preferred solution for the entire archive. Disk solutions are costly to purchase and operate for large long-term archives when compared to other technologies such as tape. Spinning disks require more space, infrastructure, power and cooling than an equivalent tape storage system. Thus, tape systems are an economic addition to store amounts of data in an archive. Additionally, tape systems can provide added reliability and disaster recovery (DR).

## ***Tape Systems***

Tape systems provide a logical tier in large long term archives for several reasons. Tape systems are currently the best economic medium for the storage of digital assets. The operational cost of tape systems is further enhanced as they require approximately a tenth of the power of low duty disk systems and less floor space. Thus, a tape library tier provides a low cost extension to the spinning disk tiers and allows maximum utilization of high performance disk by placing high access data on disk and low access data on tape. Tape libraries can provide a virtually infinite capacity as cartridges can be added and removed from the library. Additionally, tape is a transportable medium not vulnerable to shock. Thus, tape can be easily transported to off-site locations for disaster recovery.

When choosing a tape system, there are several factors to consider. First and foremost is drive technology. While LTO is currently the industry leader, other products are available with varying capacity, speed, and value added feature sets (e.g. WORM – Write Once Read Many, and encryption). A balance should be struck that allows the best \$/GB cost ratio, while providing the functionality required.

An interesting point to consider is that the highest capacity or fastest drive is not always the best choice. Depending on the frequency of media movement to offsite locations a smaller capacity may be useful so that no empty space is left on a tape when it is stored in a DR facility. In addition, if data movement is occurring from a low speed array, it may not be possible to keep the tape drive at rated speed– resulting in wear to the drive and media. Customers should choose a drive that offers some type of performance “step” functionality that allows drives to run smoothly at a variety of speeds. Long term roadmap and multi-generational support is also critical. As new revisions of a technology come to market, they should offer some level of backwards compatibility so that previous generation media can be used in new drives. Beyond drive technology, customers should also consider the service functionality of their library. Much as with disk arrays, new technology libraries include email alerting and other functions to limit downtime. Since the tape array is the long term repository for reusable content, it is critical that it be available.

## ***Integrated Solutions***

As was pointed out earlier in this document, it’s very easy to become lost in the technological details of the breadth of technology solutions available and the inherent struggles that are involved when working with vendors promoting a particular point solution as a comprehensive strategy. The truth is that

businesses need to maintain the flexibility to match-up the value of their data over time with the underlying storage infrastructure. Too often, the storage infrastructure drives business workflow instead of business workflow driving storage infrastructure. The technologies listed above are utilized in various forms to solve real business problems. Each deployment, however, is unique as the mix and utilization of storage technologies varies from business to business.

The other key component that must be considered is how to move data between these tiers. When choosing archiving software, organizations should consider three primary drivers: transparency of movement, depth of policies, and ability to attach new storage mediums. Transparency is perhaps the most critical component of an archiving solution. As data is migrated to lower cost storage and off of primary disk, how will applications locate and access the file. If the data movement is transparent, the file always appears to be in the same location incurring minimal disruption to business processes and new product generation (see following section). In addition to transparency, there is a need for customization of data movement.

Every environment has unique requirements for when and how data movement occurs. An archiving solution should allow appropriate customization of data movement to meet specific needs and, for enterprise organizations, an API (Application Programming Interface) that provides direct hooks into the data movement code. Finally, there is the need for attachment of new storage mediums. A common misconception is that tape has to be the final technology in a particular deployment. With StorNext's Storage Disk technology, a wide variety of disk technologies, from legacy RAID to NAS storage, can be a target for archive storage.

## ***New Product Generation***

Perhaps one of the most compelling business cases for archiving vast amounts of imagery data is the possibility to re-process and re-utilize imagery in new ways and in new markets. As the initial acquisition of any given image is considered a "sunk cost," selling this information to other customers, potentially in a slightly different format or as part of a trending analysis provides significant margin opportunities. Margin, however, is quickly eroded when the image you need to access is not easily located. For this reason, archiving is a critical component of data storage architectures. Archiving systems are designed to provide long term storage and continuous access to data that may be repurposed at any point and time. As mentioned above, it is important for an archiving solution to provide transparency and automated data movement. When a file is retrieved from an archive system for repurposing, the application and analyst should not need to search for the file among multiple storage devices. The file's location on day one to day one thousand should appear "static."

To this end, StorNext's archiving component, StorNext Storage Manager, is tightly integrated with the StorNext File System, StorNext Storage Manager "expands" the metadata layer of the file system so that any IO request for data is transparently handed off to the archive engine. StorNext then automatically retrieves the file back to

disk. Any StorNext client accessing a file will be unaware of the fact that the file was on a different storage tier except for a pause during retrieval from the lower cost storage medium.

## ***Long-term Data Protection***

As with any form of asset management system, cost control is only one of the key concerns. Data protection is paramount, especially in geospatial activities, because images cannot be recreated. The ability of an archive system to keep a digital asset safe in the long term starts with the ability to protect it in the short term. With file systems, even in relatively modest environments, growing into the 10s of terabytes, the traditional backup and recovery paradigms break down. Traditional backups require a nightly file system scan to identify new and changed files. If a file system has millions or 10s of millions of files stored within it, the time to scan that entire directory structure spans hours. Only after the scan has completed will data be copied to another location for safekeeping. Traditional backup paradigms also require periodic (weekly or monthly) full file system backups. With 10s of terabytes, this too can take hours or days. And since geospatial data is relatively static – once created it is unlikely to change – repeated backups of this class of data add little value and instead increase costs as more and more media are used to stored the same, unchanging data.

Using a traditional backup strategy for this class of data often results in throwing technology at the problem: deploying very high speed RAID and a sufficient number of tape drives to meet shrinking backup windows. While this strategy can work, it is quite expensive and contradicts the business strategy of deploying technology to meet business needs. After all, backup copies are merely insurance policies against a problem with the primary copy. Recovery procedures are equally difficult as the process of restoring millions of files and terabytes of data can take days or weeks.

StorNext takes a different approach to data protection. As each file arrives into a StorNext volume, it is placed on a queue for the archiving component of StorNext, Storage Manager, to determine how to protect that file. The archiving component uses data movement and protection policies, defined by the system administrator, which define how new or modified files should be replicated to another tier of storage, either disk or tape. When replication occurs, multiple copies of a file can be generated. Typically, this will be one copy for onsite storage (the copy for repurposing) and one copy for offsite storage (a disaster recovery copy). Through this process archiving occurs and there is never a need to perform a full backup. Instead, the organization automatically fulfills cost reduction strategies along with data protection. Additionally, StorNext offers data integrity checks, performing a checksum on files as they are written to tape and then verified when restored back to disk. Through this process, StorNext offers an additional measure of data protection to prevent accessing damaged data. Restoring a large storage system after a disaster is equally efficient. StorNext is able to restore the file system name space, the mechanism that allows users to access their data, at a rate of 20 million files per hour. Once the name space is recovered, the data is fully accessible. As users access their data, it is automatically pulled back

from the replicate copy on the other storage tier. If the file is not accessed, it is not pulled back, eliminating any unnecessary data transfers.

## ***Conclusion***

The definition of a “large” storage system has changed over the years. Only ten years ago, a 1-terabyte storage system was considered quite large. Today, an individual home may have a terabyte of storage under its roof considering computer systems, iPods, and DVRs. Businesses face the same challenge with increasing capacity needs and more difficult data management challenges. These issues are further magnified for geospatial imaging organizations like GeoEye (see Success Story) whose data needs are daunting, even by today’s standards. The good news is that storage technologies such as RAID and Tape systems have grown in size and sophistication to meet today’s data storage needs and data management solutions like StorNext have provided automated mechanisms to match business needs with technology.

## ***About Keeper Technology***

Keeper Technology provides design and implementation services focused on data storage and data management architectures. Keeper’s consultants have been architecting and deploying mission-critical storage solutions in commercial and department of defense enterprises for nearly two decades. Visit us on the web at [www.keepertech.com](http://www.keepertech.com) or via email at [solutions@keepertech.com](mailto:solutions@keepertech.com).

WP00096