Big Unstructured Data Archives

Highly durable, scalable and cost efficient disk storage for large volumes of unstructured data

AmpliStor Solutions Paper
March 2012
V 3.0

Unlocking Business Value in Big Unstructured Data

Today, many mainstream businesses realize that there is unparalleled value in storing and maintaining very large and essentially unbounded volumes of unstructured data. Big unstructured data takes all shapes based on specific industries, such as the healthcare industry with medical images, travel and hospitality industries capturing surveillance video footage, the retail and manufacturing industries that require design data and product images, as well as other industries. Massive amounts of this unstructured data have traditionally been stored in host-based file systems and on tape, but now a new type of scalable storage system is required to replace these decades-old solutions. This data can then serve as the basis for next generation analytics that can help businesses make more informed decisions related to product strategy, marketing, research and historical trends.

Data growth is occurring in a multitude of application domains. In the past, most electronic corporate data was highly structured, and could easily be represented as fields within a tabular format, such as the rows & columns of a relational database. The next evolution was semi-structured data, such as email repositories including simple messages plus file attachments. Now, with the growth of digital media, we're witnessing a new growth in large-scale *unstructured* data, Big Unstructured Data. The key characteristics of this data are unique in a number of dimensions. The first is that individual data objects are much larger, typically measuring Megabytes to Gigabytes in size. Video streams are now even reaching into tens and hundreds of Gigabytes, especially with the advent of high-definition video resolutions. Second, these applications can generate enormous numbers of these objects, with the need to manage billions of such objects seen as the next inflection point. Finally, all of this data in aggregate certainly registers into hundreds of Terabytes and petabytes today. For unstructured data applications in digital asset management for media and entertainment, digital archives, medical imaging and online media applications, the day of Exabyte scale data is certainly around the corner.

Storage Reliability at Petabyte Scale

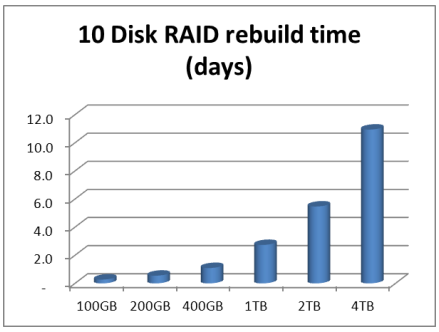

In media & entertainment companies, aside from people, media assets are their most valuable and monetized assets. The need to protect this key Intellectual Property from corruption or loss is the ultimate concern; therefore storing these large media assets clearly require highly reliable storage systems. In the past, this was possible with traditional storage systems employing RAID technology protecting relatively small disk drives. Today, our opportunity is to leverage low-cost, high-density (multi-terabyte) disk drives to make large-scale storage very cost effective. But this is exactly where RAID technology is now hitting crucial limitations. RAID groups comprised of multi-terabyte drives are exposing applications to very high-probabilities of data loss or silent corruption. This is due to multiple factors, including the excessively long rebuild times incurred when a terabyte-level disk drive fails, and must then undergo a rebuild operation. Rebuild times are commonly exceeding a day, with multi-day and week-long rebuild times becoming a reality, especially in systems that must prioritize access performance over rebuild times. Silent data corruption occurs due to bit rot (media decay), and the resulting unrecoverable read errors. The truth is this: petabyte-scale RAID systems built on top of terabyte density disk drives, will indeed corrupt and lose data. It is a simple matter of statistics before the system experiences multiple disk drives in succession, or experiences an unrecoverable read error – both of which can cause data loss.

The recent solution to this weakness of RAID has been simple: enable the application to store multiple copies of these large unstructured objects. This allows one copy to be lost, while still ensuring that the second copy is available. In many cases, applications will maintain three (3) copies of these objects so as to provide full assurance against data loss, even during routine maintenance and upgrade operations. While this course of action is acceptable for smaller data sets, or small documents – it is a highly expensive proposition for large object data. The notion of having to triple the underlying storage capacity in order to make an object reliable makes sense for small data objects or a few Terabytes, but when the scale is changed to petabytes, it clearly makes no economic sense to incur the cost of 3PB of underlying storage for every 1PB of actual unstructured data being stored.

Management of Petabyte Scale Unstructured Data

The problems of managing storage at petabyte scale are entirely different than managing Terabytes. Statistically, the number of failures that can occur when dealing with thousands, or ten thousands of disk drives certainly creates a management nightmare. The need to incrementally scale system capacity over time, as well as matching performance (throughput in the case of large objects) to the increasing capacity leads to difficult problems. How do we scale a system that contains petabytes of data to Tens of petabytes without impacting daily operations? Will it be possible to do this without hitting system limitations in capacity (for example, the limit of a storage system frame), and moreover does the system need to be reconfigured to utilize the new capacity? Data migration becomes a key concern as well: if we are storing petabytes of data on today's technology what happens when this technology is unsupported or becomes end-of-life by the manufacturer – will all of this data need to be migrated to an entirely new system after a forklift upgrade?

These problems have lead to the search for alternative, highly reliable and efficient storage solutions for large unstructured data in a wide range of applications and industries.

AmpliStor Optimized Object Storage for Big Unstructured Data

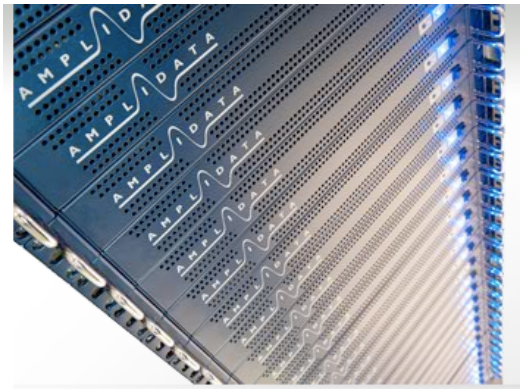

The AmpliStor system has been purpose-built to address the needs of *petabyte-scale, Big Unstructured Data* applications. It is the only storage system to employ patent-pending BitSpread technology that enables storage of large-scale archival data with any level of reliability and availability with the lowest possible overhead, but without the access performance and management complexity issues of tape, plus assured data integrity and the highest levels of reliability. Moreover, the system provides very low costs both from a capital and operational perspective. For large data archive applications, this provides a number of distinct advantages:

- **Unbreakable Storage**: AmpliStor provides any desired level of data reliability of availability – by tolerating any (user-specified) number of failures. The system eliminates and solves the reliability issues of RAID on multi-terabyte disk drives. Data integrity is pro-actively monitored, verified and assured through checksums and disk scrubbing, to provide full protection against multiple disk drive failures, unrecoverable read errors and bit rot. Component level failures such as disk failures are automatically resolved by the self-healing design.

- **The lowest storage overhead:** For petabyte-scale data archives composed from millions of big files, AmpliStor utilizes BitSpread to provide super-reliable protection from data loss without the lowest possible overhead. AmpliStor will require 50-70% fewer disk drives and data center rack / floor space than RAID and replication based storage systems.

- **Scalability with low management effort:** AmpliStor can be scaled with additional AmpliStor storage nodes on the fly. The system auto-detects and utilizes new capacity, to scale the system without reduced manual intervention. Self-monitoring of disk & node-health plus self-healing reduces the need for intervention as the system grows.

- **90% less power consumption:** AmpliStor has the lowest storage overhead, leverages low-power disk drives, plus efficient storage enclosures to provide the most reliable storage for only 7 Watts per TB. Future versions of the AmpliStor software will enable disk drive and node-level power-down, to further decrease power cost to the level of offline systems such as tape.

- **AmpliStor provides 50-70% lower Total Cost of Ownership:** all of these capabilities are provided at very low Capex and Opex (data center floor space, power and cooling costs), and lowest management overhead to provide an incredibly low TCO.

The AmpliStor system architecture consists of AmpliStor storage nodes and AmpliStor controller appliances, connected over 10 Gigabit or one Gigabit Ethernet networks. The system can be scaled from a small configuration of a few storage nodes up to thousands of nodes, accessed by hundreds of controllers to serve large groups of concurrent users. This enables the system to be deployed in a range of backup and archival applications with a broad spectrum of requirements. For additional information on the AmpliStor system, see us at www.amplidata.com, email info@amplidata.com or contact your local sales office.

Amplidata European Headquarters

Antwerpse Steenweg 19
9080 Lochristi, Belgium
Tel: +32 9 324 25 90
Fax: +32 9 324 25 99

Amplidata US Headquarters

2100 Seaport blvd, Suite 400
Redwood City, CA 94063 USA
Tel: 650-367-9683