STORAGE SWITZERLAND

BUILDING THE BRIDGE TO TAPE WITH SCALABLE STORAGE INFRASTRUCTURES

Storage Switzerland, LLC

George Crump, Senior Analyst

Media and Entertainment (M&E) IT departments and those within organizations trying to process massive amounts of analytical data, also known as “Big Data”, have something in common. They have to store massive capacities of information and keep it as close to online as possible without breaking the budget. While many companies are trying to propose a disk-only approach, this is simply not realistic for most of these data sets, given the cost. While there are active archive products that seamlessly move data between primary storage, secondary disk storage and tape storage, identifying what disk to use for that middle tier of storage that will act as the bridge to tape, is critical for the solution’s success.

Media and Entertainment Storage Management Challenges

As the capacity requirements of modern film making have increased so have the demands on the M&E IT department. Capacity needs have increased thanks to higher resolution and 3D. There has also been an increase on the need to quickly move data between types of media as quickly as possible. Often this means shooting the live action directly to solid state media in a camera, then as soon as there is a break the data on solid state is transferred to either a tape or hard drive shuttle type of device. That media is often shipped back to the studio and it is loaded onto primary storage for editing.

Typically the entire project cannot be stored on primary storage and it needs to be transferred between primary storage and a middle tier of storage and different sections of the project need to be worked on. Then over time more of the data is archived to a tape based tier for longer term and least expensive storage.

Requirements For The Middle Tier

Performance

While much of the attention in the middle storage tier is often on cost savings, this tier must also provide excellent performance. The middle tier disk platform has to quickly promote data to the primary storage tier and in many cases, be able to perform well enough for data to be served directly from that tier. Even the move to and from tape has performance pressures as data has to either be read or written very quickly. This is challenging enough for a “normal” environment but in high capacity M&E or Analytical environments it’s even more challenging.

Flexibility

The middle tier platform also has to be flexible as the environment grows, as it could be the storage area for some of this data for years, if not decades to come. That means the underlying technology of the middle tier has to change with the times, adopting new hard drive capacities and the denser packaging of higher performance processors. Systems that can't provide this level of flexibility place the IT administrator at a severe disadvantage and require a re-buy of technology every few years, plus the Herculean task of migrating a very large data set.

Cost Containment

Of course another key requirement is to contain costs. While the tape tier is available for the ultimate in economics, the middle disk tier must also be very cost effective in order to maximize the data stored on this quicker access tier. Achieving cost containment requires high capacity drives and very dense packaging, but should also address the requirements of flexibility and performance.

Reliability

Reliability is critical in this tier. While the data may also be stored on tape, a data failure can still lead to a very long recovery effort. Given the high capacity drives that are typically used in cost sensitive applications, something other than RAID 5 needs to be leveraged as the time to rebuild 2TB+ hard drives is unacceptable when storing this type of data.

Using Scalable Storage Infrastructures as the Bridge to Tape

The most viable candidates for providing this middle tier of storage and potentially the primary tier of storage are a scalable storage infrastructure like those offered by Amplidata's Amplistor systems. These systems use scale-out designs that leverage commodity hardware servers clustered into a single storage object. Each of these servers represents a node in the cluster and comes complete with 20TBs of capacity, plus appropriate network I/O bandwidth and processing power to drive the storage. The aggregation of these nodes addresses each of the three requirements for the middle tier of disk storage and provides an ideal bridge to tape.

Scalable Performance

From a performance perspective the combined capabilities of the nodes means that each time capacity is added to the system additional network bandwidth and processing power come with it. The system actually becomes faster as the capacity increases. This means that very large, multi-year or multi-decade storage areas can be created while maintaining performance. There is plenty of I/O and processing horsepower to handle the movement of TBs of data from tape to the scalable storage infrastructure or to the primary storage tier. As mentioned earlier, depending on the environment, the performance of this middle tier may be great enough that a primary tier of storage isn't needed at all.

Scalable Flexibility

A key challenge of this middle tier is not just how long data may need to reside on it but also the amount of data that will be involved. With potentially 100s of TBs of data on this tier simply upgrading the storage infrastructure after a typical storage refresh cycle and migrating data to a new storage platform would be all but impossible. It would just take too long even across the fastest of network segments.

The advantage of a scalable storage infrastructure is that it can be upgraded with not only data 'in place' as many traditional storage systems do but with "system in place". Scalable storage infrastructures can mix node sizes and capabilities within those nodes and then use those appropriately. The system in a sense is 'self upgrading'. Newer nodes containing the latest hard drive capacities and processing capabilities are added when needed. Data is automatically rebalanced across the new nodes. Then as older nodes age, they can be decommissioned from the cluster and data will once again be rebalanced to the remaining nodes.

Scaleable Cost Containment

This scalable architecture allows for the use of more cost effective nodes. Since the cluster will benefit from the combined performance of all the nodes, a more expensive, high performance processor is not needed. Instead, something like Intel's ATOM processor can be used. Not only does this reduce the hard cost of the processor it also reduces power consumption. This means that denser packaging can be used allowing more storage to be implemented in less space, which of course also saves data center floor space.

Scalable Reliability

Finally, reliability is addressed partly by the architecture of the scalable infrastructure itself. In this design there are multiple points of redundancy instead of single points of failure, meaning that more than one thing can go wrong and data is still accessible.

Further, Amplidata has improved on the standard RAID5 style data protection and implemented a modern variant of an erasure coding technique, termed "BitSpread". BitSpread encodes reliability into data at a check-block level, so that if a drive fails the system needs to generate just the new check-blocks, and not recalculate parity on the entire drive. BitSpread can also protect against any number of disk failures and provides full protection against media errors, such as unrecoverable read errors also known as "bit rot". This technique also extends efficiency, since the capacity lost to create this redundancy is less than with traditional RAID technologies. Better efficiency equals more scalable cost containment.

Summary

Scalable storage infrastructures can solve many of the storage problems that IT is faced with today. One of the biggest is how to properly balance the use of tape and primary disk. Scalable storage infrastructures create a middle tier of disk that provides a cost effective and reliable bridge to tape. Combined with technologies like active archive it can be part of a seamless interaction of the three technologies moving the discussion away from disk vs. tape to a harmony of storage assets solving today's storage challenges.