Centralised storage in film

The problem with centralised storage

The advent of Digital Intermediate (DI) has introduced new and powerful tools for the manipulation of film images using a digital process rather than following the traditional photo-chemical path. This process has enabled filmmakers with finer control of colour grades and has simplified the addition of computer generated visual effects.

These changes also bear new data storage challenges to companies taking advantage of this new workflow. Scanning a 90 minute feature length film at 4K resolution incurs a storage penalty of more than 7 terabytes. Network data transfers are expensively time consuming, even for applications working in 2K.

Installing a Fibre Channel (FC) Storage Area Network (SAN) is the current high-performing solution employed to provide digitised film frames to the DI process. The common goal being to avoid lengthy transfers by directly connecting systems to a shared pool of storage. This type of storage system is commonly employed for grading, review and visual effects systems.

High-cost fibre optic cabling, Fibre Channel switches and host cards are required to connect a select few systems to storage arrays. Due to disk-head seek latencies, a high number of disk (typically SCSI) spindles are needed to sustain concurrent read/write data streams.

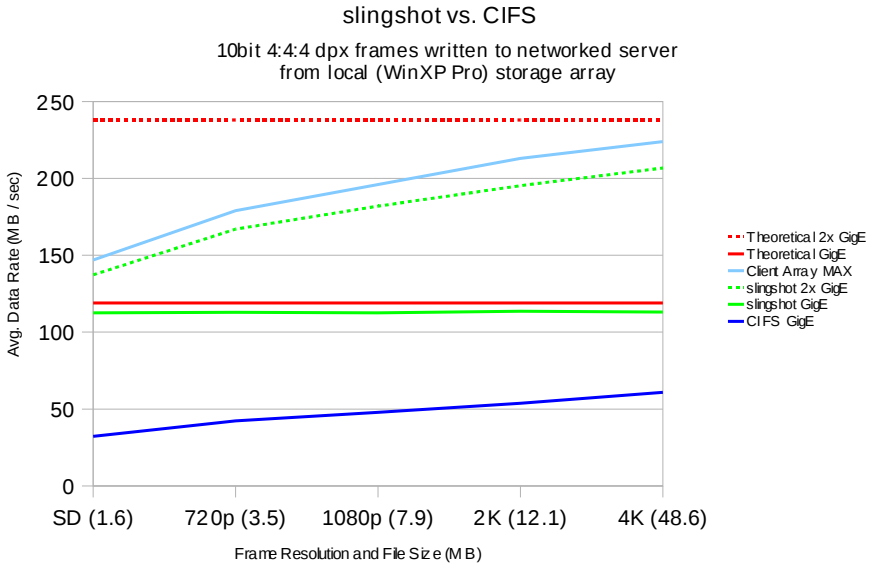

Illustration 1: slingshot vs. CIFS

Despite the generous monetary investment in deploying a SAN solution, the storage array and network infrastructure is only used to a fraction of it's streaming data-rate capabilities. This is due largely to the use of generic purpose filesystems and network protocols that do not take into account the sequence-of-frames nature of film.

Access to shared SAN storage from non-SAN systems is achieved through a Network Attached Storage (NAS) broker directly attached to the SAN. This access path provides

CIFS and NFS services over common Gigabit infrastructure to Visual Effects, Offline Editing, 3D and Render Node systems.

NAS systems inherently limit network access to a single point of entry. Access from any single system is at maximum 1 Gigabit, despite any Ethernet aggregation configuration on the NAS system. Ethernet aggregation (port trunking, 802.3ad, etherchannel) will only provide a single Gigabit path for any single point-to-point (workstation to server) data transfer. The current implementations of CIFS and NFS over TCP/IP typically utilise Gigabit Ethernet to less than fifty percent of its capability. (i.e. less than 50 MB/sec – *see Illustration 1*)

Overcoming these limitations of SAN and NAS solutions is the focus of our research. Utilising existing infrastructure more efficiently and scaling to high-end storage and network fabrics accelerates DI workflows at a significantly lower cost.

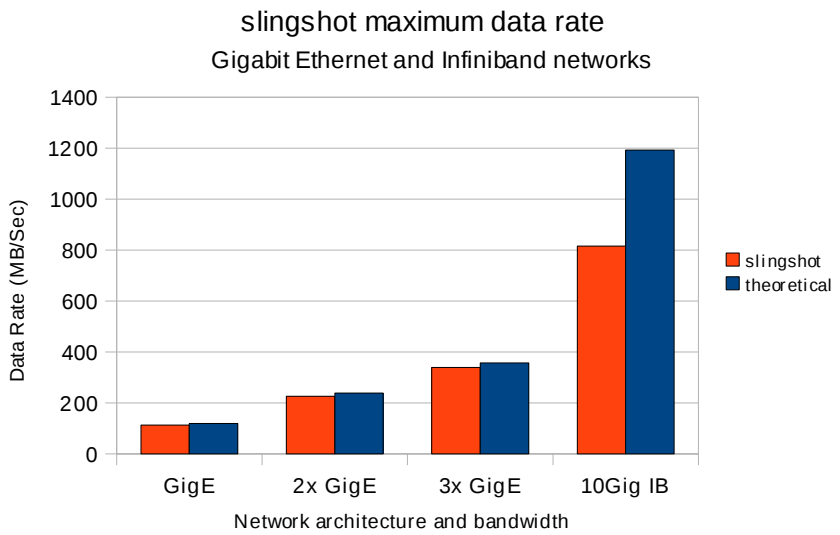

Illustration 2: slingshot over InfiniBand

With growing demand for DI workflows comes pressure to deliver in a shorter time frame with lower overheads. Studios are already pushing these limits with each new production. Increasing storage capacities on SATA hard disk drives appears attractive, but needs carefully regulated data access and large storage pools require new methods of data management.

Caveats: catapult server™ excels where file sizes are large (> 3 Megabytes), and total data transfers are large (> 100 Megabytes). catapult server™ will allow data transfer approaching the maximum speed of the bottleneck. The bottleneck will be one of: The reading data array, the writing data array, or the network connection. Some situations may only have one read or write array.