

Fibre Channel Connectivity in Modern Content Creation Workflows

Jim McKenna

The Dilemma - IT Infrastructure Convergence with Rich Media Production

For over a decade now, content creation infrastructure has been penetrated by IP-based networks and network-attached storage products. Shared storage, once the bastion of a handful of high-end facilities, has become commonplace with the help of network file systems and lower cost 1/10Gb Ethernet hardware. The most popular workstation operating systems have decent compatibility with these systems, and mapping network shares to the desktop for access to project files is a popular workflow.

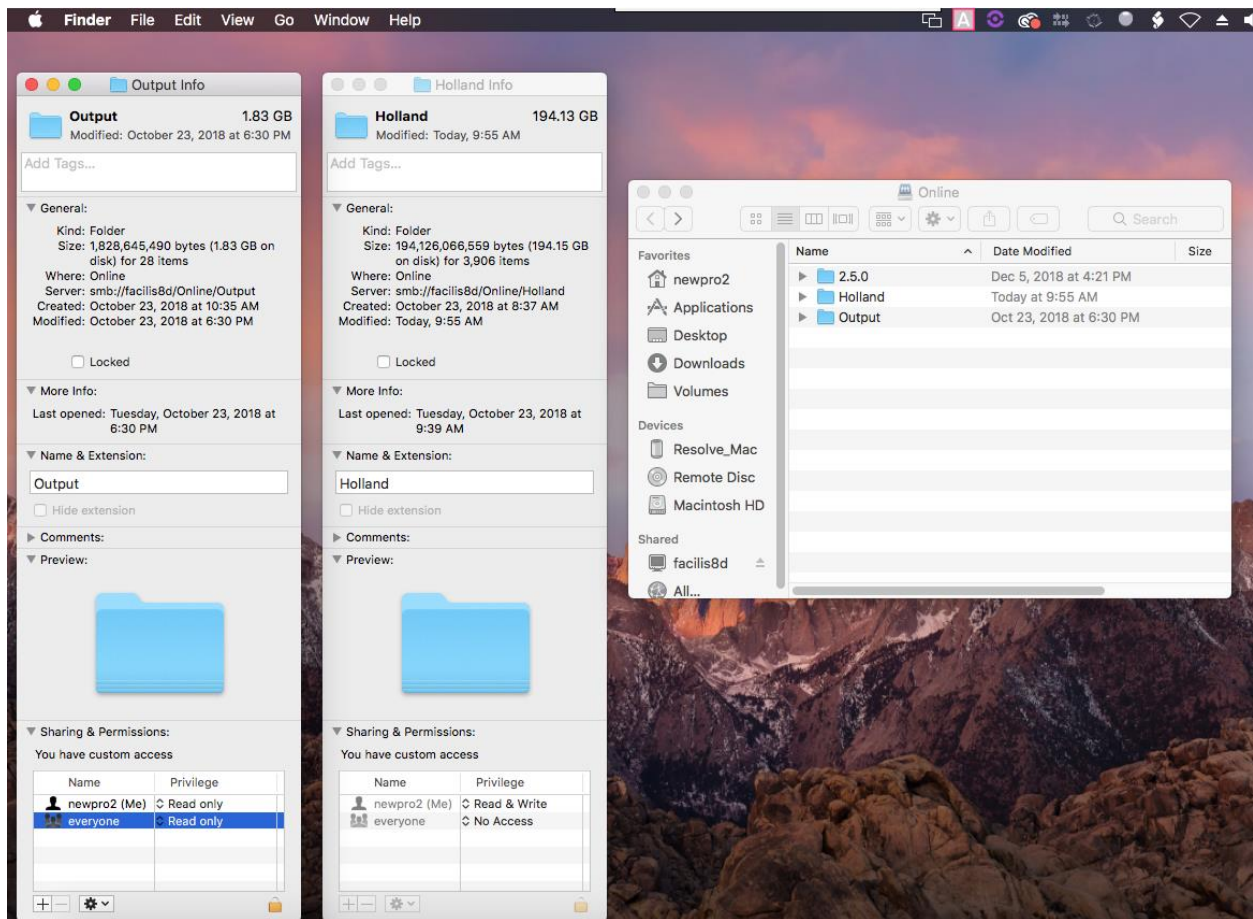
It's time for content creation professionals to pose the question, what do we lose in the move toward NAS topology, and what is better achieved through Fibre Channel with a custom shared file system?

It's true that operators and administrators have learned to deal with challenges in the IP-based workflow that would not occur in traditional SAN-based environments. These include dependency on external IT resources (DHCP, DNS, Active Directory), dependency on desktop OS file system support, lack of isolation causing traffic disruptions and security concerns, and higher resource usage on desktop PCs when processing Ethernet transactions. Are these limitations necessary?

The Impact of Permissions & Dependency on External Resources

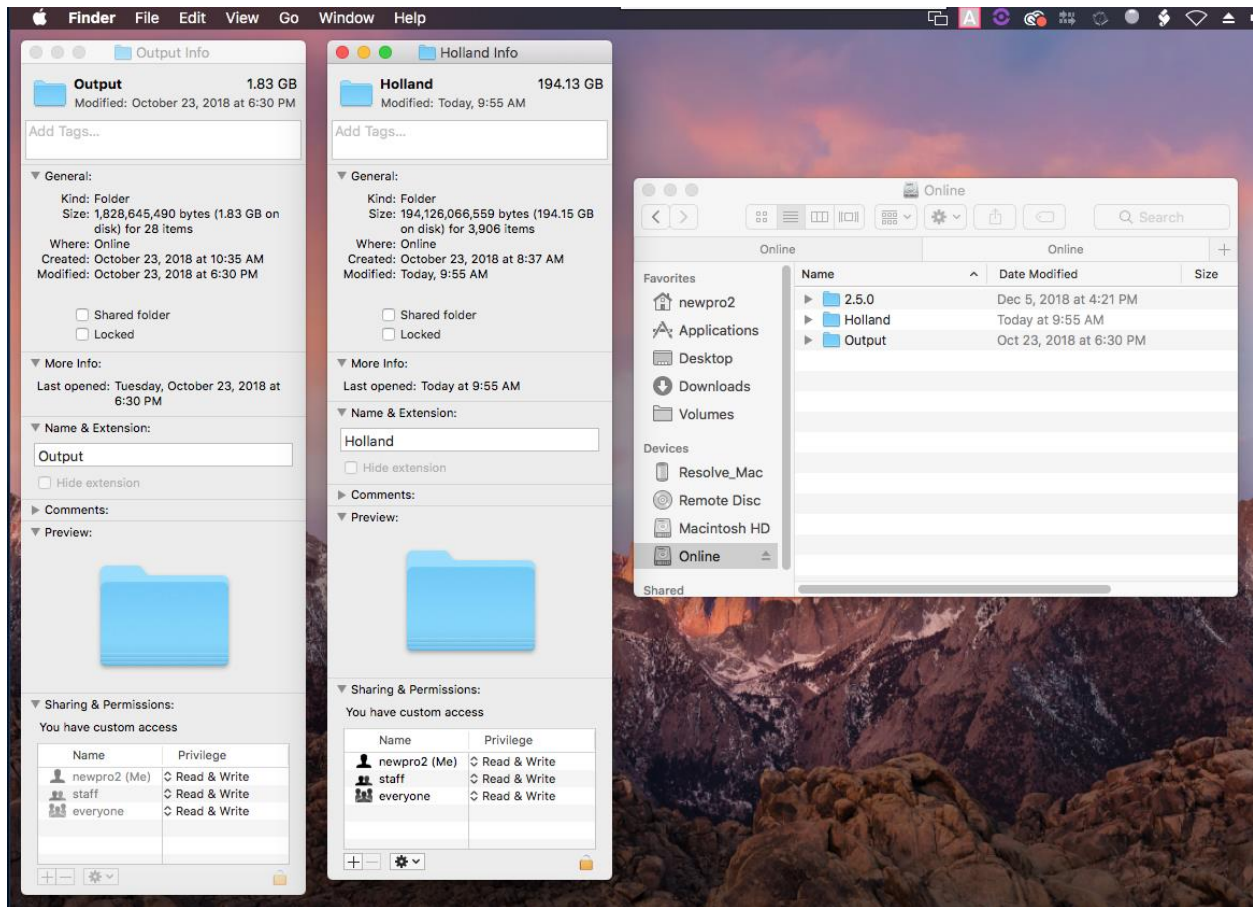
In the world of IP and NAS-based workgroups, solutions are designed to deploy within a set of services that manage permissions and authentication for the network. These include directory services, with which user accounts are assigned to the network resources, and IP handling DHCP services that manage network addresses. If there is any disruption in these services, network storage can become unavailable. Small facilities relying upon prosumer router technology and having little network experience can find problems with IP address management and certainly will have problems with user account setup and permissions. Login to local user accounts that are not recognized as part of a domain group can exclude the user from files, and the user may be unable to modify or delete files that are owned by user accounts that no longer exist. In addition, the state of a file on another client workstation may determine what level of access you have to the file.

Permissions on NAS resources are variable folder to folder and file to file. These permissions are not apparent by looking at the file or folder, so mistakes can be made by operators, mistaking a read-only location for a writable one. By the time the administrator fixes the problem, you've lost valuable time. Alternatively, a custom shared file system may be deployed with a Fibre Channel collaborative network, to help alleviate this confusion. As you can see below, read-only and writeable folders on NAS storage are identical to the end-user.



Also notice that the network drive is not shown in the Mac sidebar by default.

With a custom shared file system, the volume can be set to display as a local drive, appearing in the sidebar for ease of navigation. Permissions are set to full read/write access, or full read-only access depending on the permissions of the volume itself, to avoid mistakenly selecting a read-only folder in an otherwise writable location.



OS Interoperability Dependency is Often Overlooked


As operating systems advance, and take different approaches to network interoperability, network storage in the facility is at risk of being left out of the compatibility matrix. Mac OS X has a checkered past with SMB support. Since the change to avoid the GPLv3 licensing in 10.7 there were various SMB-related connectivity issues that have arisen, which put facilities using NAS storage and SMB shares into problematic situations.

This is especially damaging when facility workstations use different versions of the Mac OS, some supporting newer, more stable versions of the SMB code, some with older and broken versions. The workflow of a facility sometimes hinges upon older systems, with hardware that may not be supported in new OS versions but is still integral to the operation of the workstation. Artists have their preference of software applications, and sometimes that application is only supported on older OS revisions. Take for example that many facilities still use the Final Cut Pro Studio product in 2018, nine years after the final release of that software.

When using a custom shared file system, whether connected through Fibre Channel or Ethernet, the list of compatible OS versions is more complete, sometimes reaching back several years of OS revisions. The client operating system is not the owner of the storage volume, nor does it play as integral a part in the exposure of the volume to the desktop. This helps maintain a consistency of functionality across multiple workstations running various OS versions.

Performance is More Than Just Benchmarks

Performance in general is a measure of how well a job is being done. High-performance NAS systems may be fast in benchmark tests, but that doesn't tell the whole story. The job may still be done poorly, even with systems that display high-speed metrics. The reasons for this can take on many forms, including issues described in the sections above with permissions and interoperability. Many times, the performance of a network storage system is limited by the network protocol itself.




TCP-IP protocol consists of several layers, the actual data payload among them, which also includes additional header containing client IP and TCP data

These layers must be processed by the client and server whenever there's a request for data. The max size of a payload in standard Ethernet protocol is 1500 bytes, or 1.5KB. Jumbo frames can produce payloads up to 9000 bytes each, or 9KB, but Jumbo frames are not compatible with every switch infrastructure. It only takes one switch in the facility that fails compatibility for the connectivity to drop.

In default setup, for a common request size for video playback of 1-2MB, over 1,333 frames must be sent, and that 2MB request may happen 30 times per second for each video track on an edit system timeline. In a 4-way multi-camera timeline, up to 160,000 Ethernet frames per second must be processed. This affects CPU usage at the client and server. When the client count on a server is scaled, the processing required to service all the client workstations becomes a limiting factor.

Client workstation CPU load is a common cause of poor high-bitrate video playback. If you consider all the intensive decompression and display processing needed to play video, you can see why the CPU level is constantly a concern on content creation workstations. If it's possible to lower the CPU usage, video will play more smoothly, and applications may not require draft mode settings for video display. Fibre Channel block-mode offers this type of low-CPU overhead and that can be a real benefit.

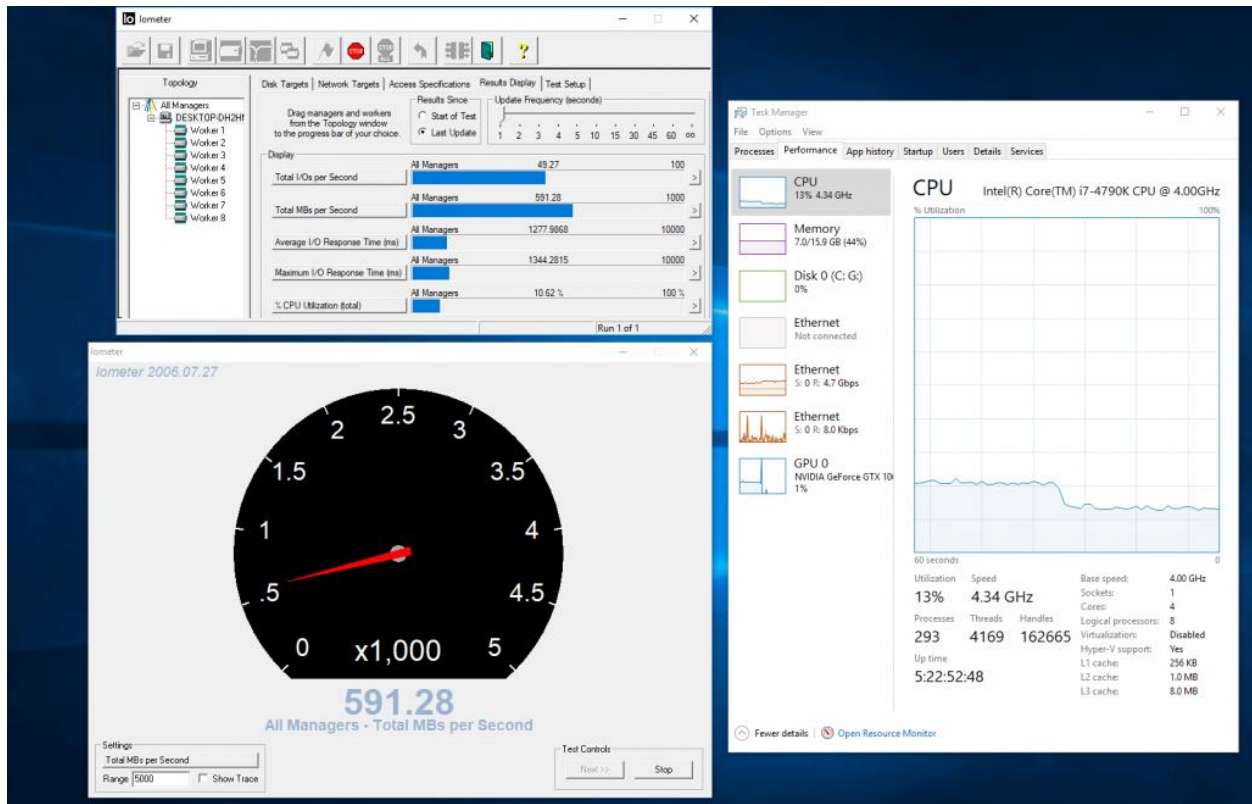
The Fibre Channel Block-mode Protocol – Built for Speed



The most common language used over Fibre Channel is SCSI, the block-level interface for direct-attached storage devices. The SCSI frame is held in the Fibre Channel data field. The standard data field in a FC frame is larger than the TCP-IP payload by 30%, and it doesn't contain additional IP and TCP header information like an Ethernet payload. In addition, the Fibre Channel protocol is a credit system per

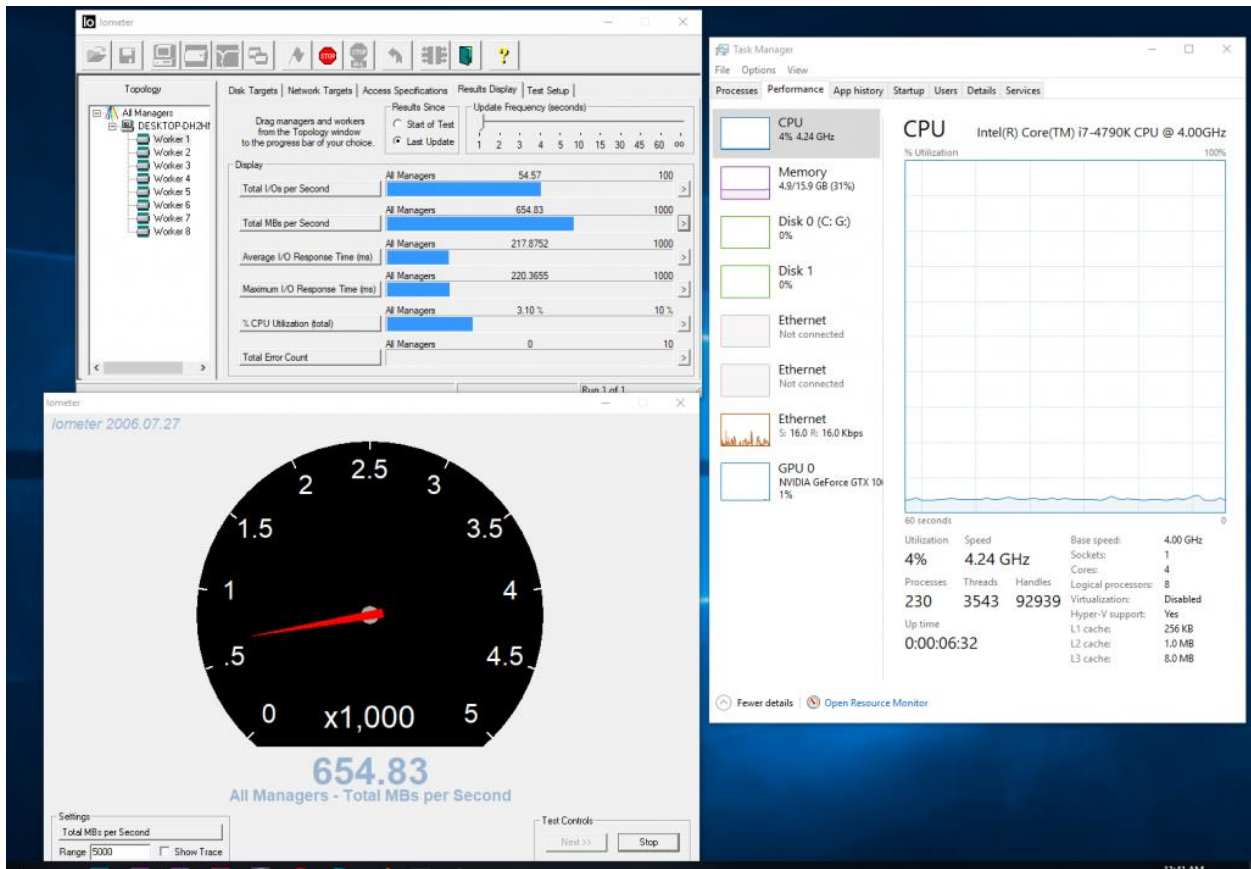
device. There are memory buffers associated with each attached client, and there is a constant running count of the memory buffers available for that client connection. This is different than TCP-IP, where memory buffers are accounted for, but not on a per-client basis. Therefore, when traffic is introduced, buffers may run dry without TCP-IP clients knowing about it, causing congestion.

The result of the increased efficiencies of block-mode Fibre Channel is an increase in consistency of bandwidth, lower latency, and lower CPU usage on the server and client workstations. Below are images of benchmark tests that show the differences in processing data on the client CPU. These benchmark applications are sending large request sizes, as would be required for high-bitrate 4K video formats. The first image is from a Windows 10 client with an Intel i7 4Ghz 8-core CPU, on a 10Gb optical Ethernet connection through the NAS SMB filesystem. For each of these, the CPU time at idle was 2%.

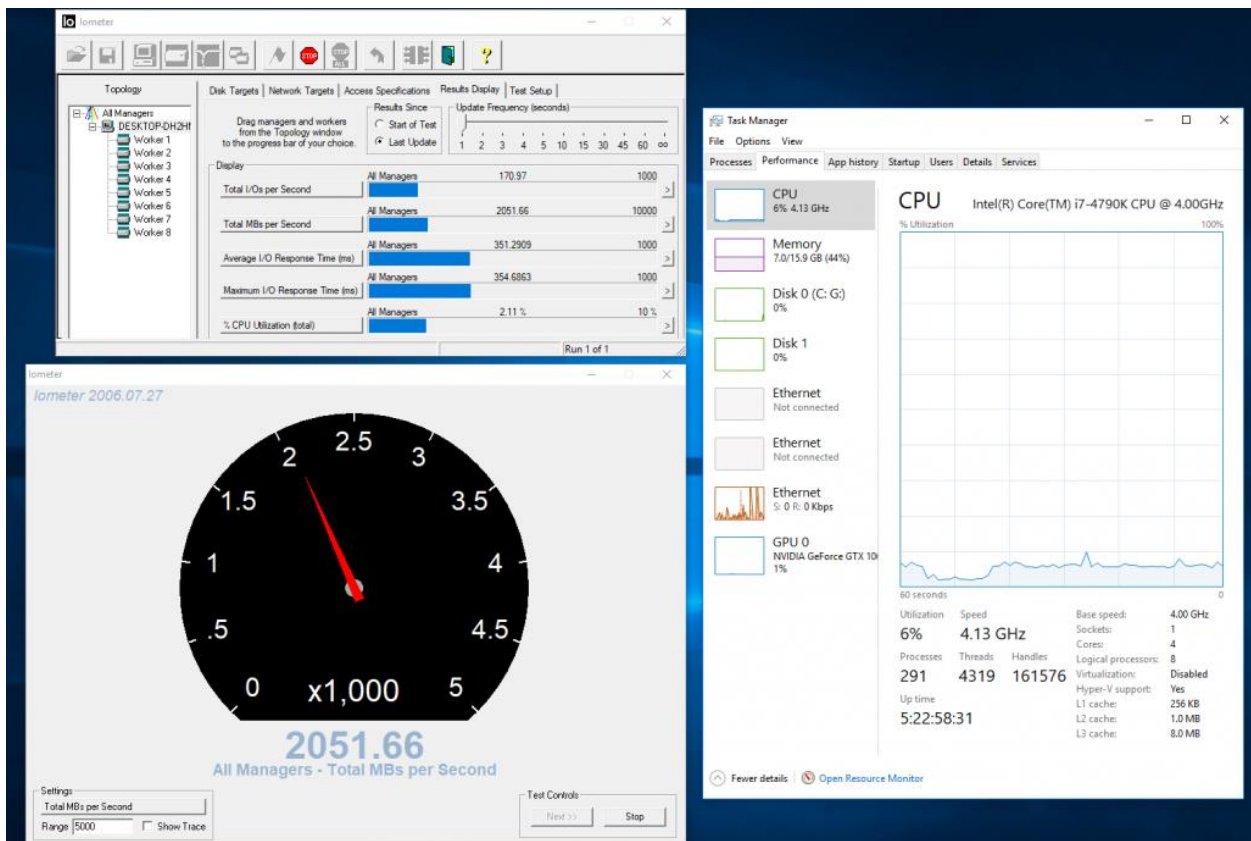


The 10Gb benchmark shows 590MB/sec of sustained bandwidth, and 13% CPU. Taking idle CPU at 2% into account, the processing of the frames is consuming 11% CPU for 600MB/sec, or 1% for every 55MB/sec.

The next image is the same Windows 10 client, using the same storage volume and optical cable, but attaching to a custom shared file system over 8Gb Fibre Channel.

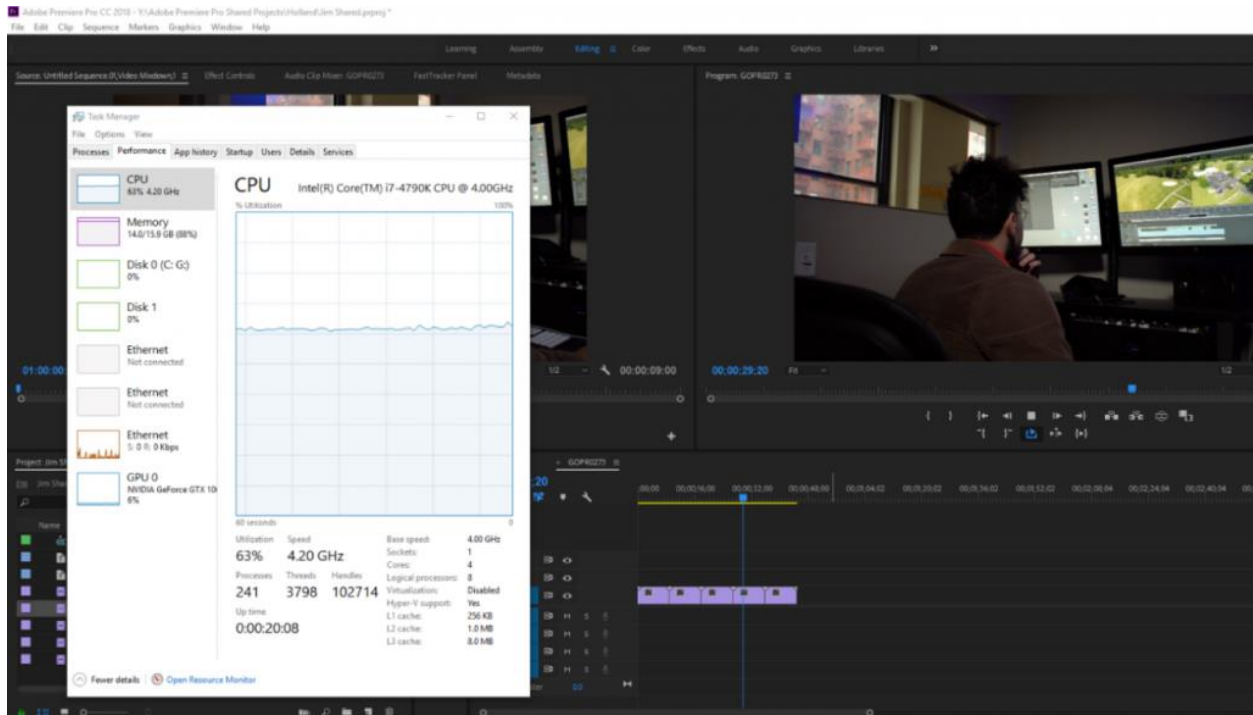


The speed of the same volume through 8Gb Fibre Channel is about 650MB/sec, and the overall CPU usage is 4%, or about 2% load. The client CPU is only spending 1% for every 325MB/sec.



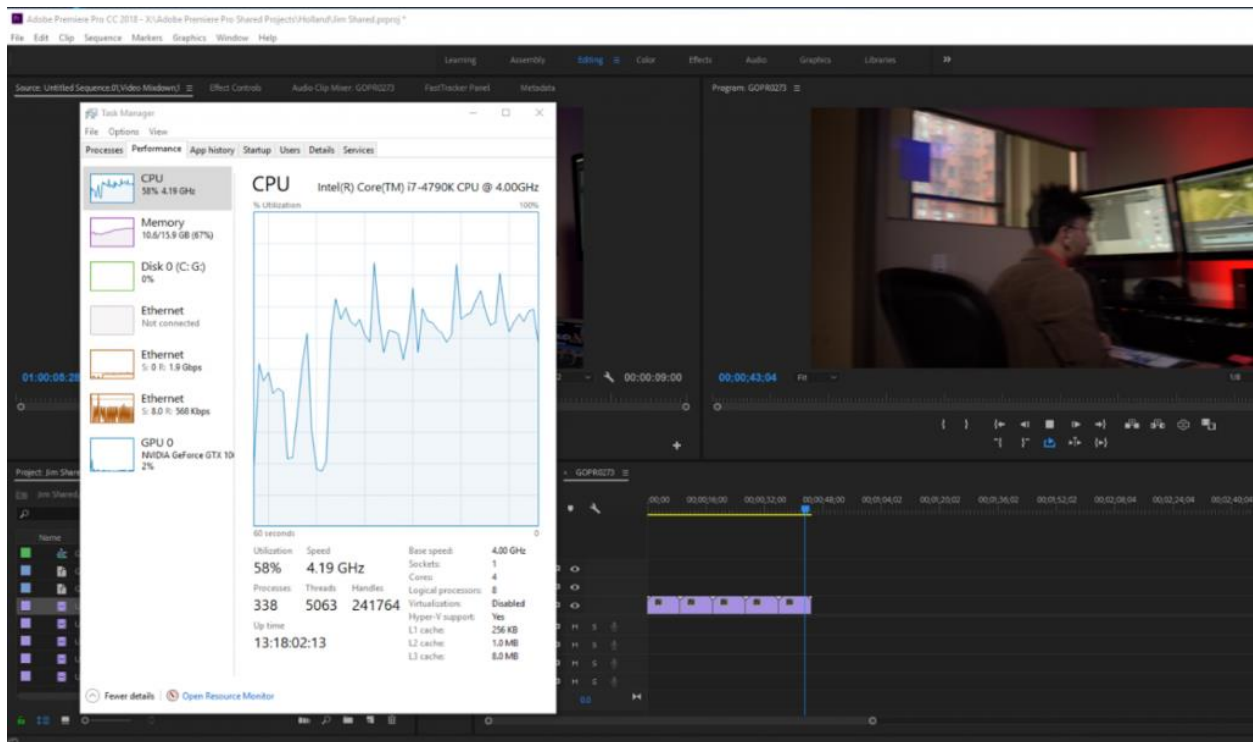
This image shows the speed of the same volume through 32Gb Fibre Channel is about 2000MB/sec, and the CPU usage is 6%. This equates to a 4% load for the test, or 1% CPU time for every 500MB/sec. This is 9 times the efficiency of the TCP-IP connection.

The 10Gb benchmark numbers above may change with different networking cards, but the ratio of MB/sec to CPU usage should remain consistent. To show how greater CPU usage can affect a workflow, here is an image of playback within Adobe Premiere Pro of 8K (UHD2) compressed video. Results will vary with memory, motherboard and GPU, but this is a very simple way to compare the two connection methods on an average workstation.



Using the Fibre Channel connection, CPU usage is solid at 63% to maintain full speed playback.

When using the 10Gb Ethernet connection, CPU was unstable from 53% to 82%, and the video playback could not be maintained at full speed. The video format being read required 430MB/sec for a single stream. This is far less than the 10Gb benchmark, but when the application is fighting for CPU time due to TCP-IP processing, the available bandwidth is deficient.



From what we now about %CPU per MB/sec, if the full speed video was played successfully, we can expect a much higher CPU usage. Higher bitrate formats requiring more MB/sec will eventually bind the CPU during decompression and display of the video.

Hardware Infrastructure Now & In the Future

Ethernet is so ubiquitous that it is believed that any speed can be transmitted over any wire, but this isn't the case. Higher bandwidth IP links are appearing in data centers, but these require that specific drivers, network adapters and switches be deployed in scaled environments. Even with the newest cable formats, the limitations on connectivity can be prohibitive. The most common connectivity options for Ethernet are below.

1Gb	OM1-OM4 MMF SFP up to 100m; Cat5e/6/6a Up to 100m
10Gb	OM3-OM4 MMF SFP+ up to 400m; Cat6/6a up to 100m
25Gb	OM4 MMF SFP+ up to 100m; Cat8 Up to 30m (in development)

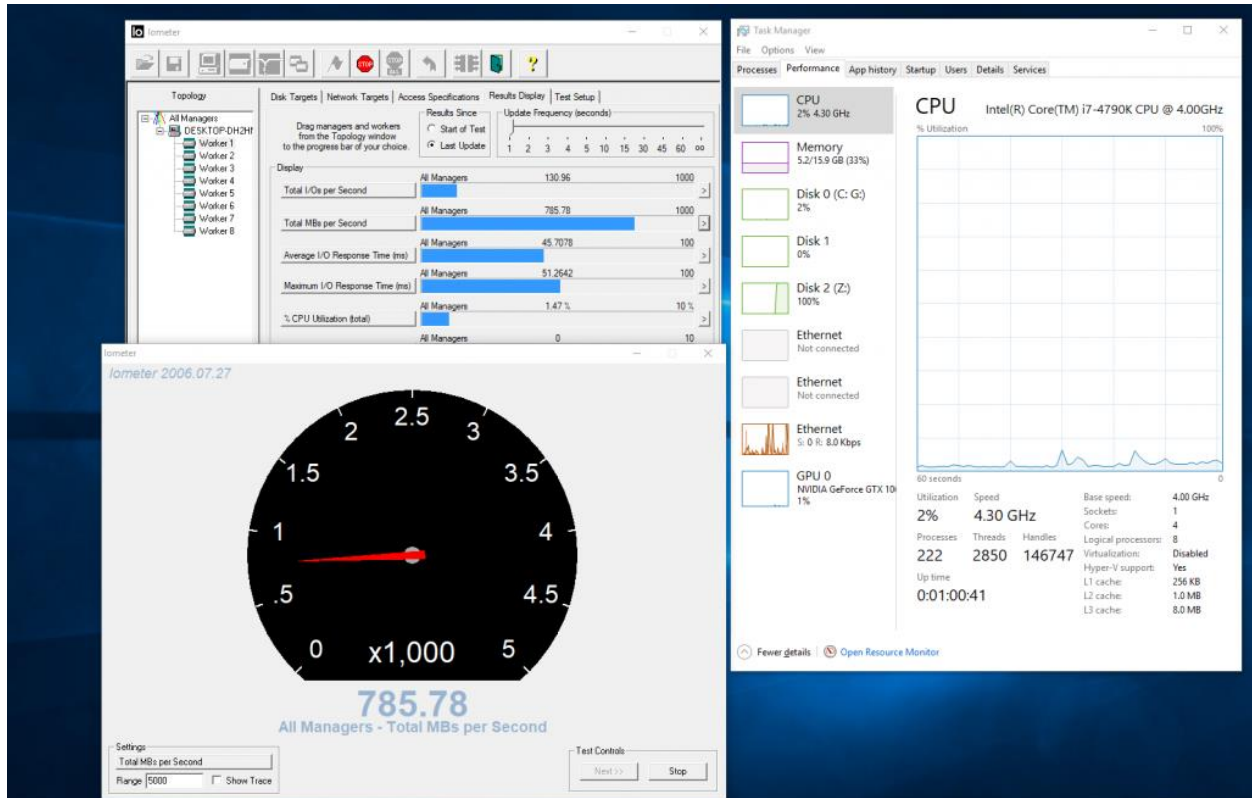
The most commonly used options for Fibre Channel are below.

8Gb	OM1-OM4 MMF SFP 21m to 190m
16Gb	OM1-OM4 MMF SFP+ 15m to 125m
32Gb	OM2-OM4 MMF SFP+ 20m to 100m

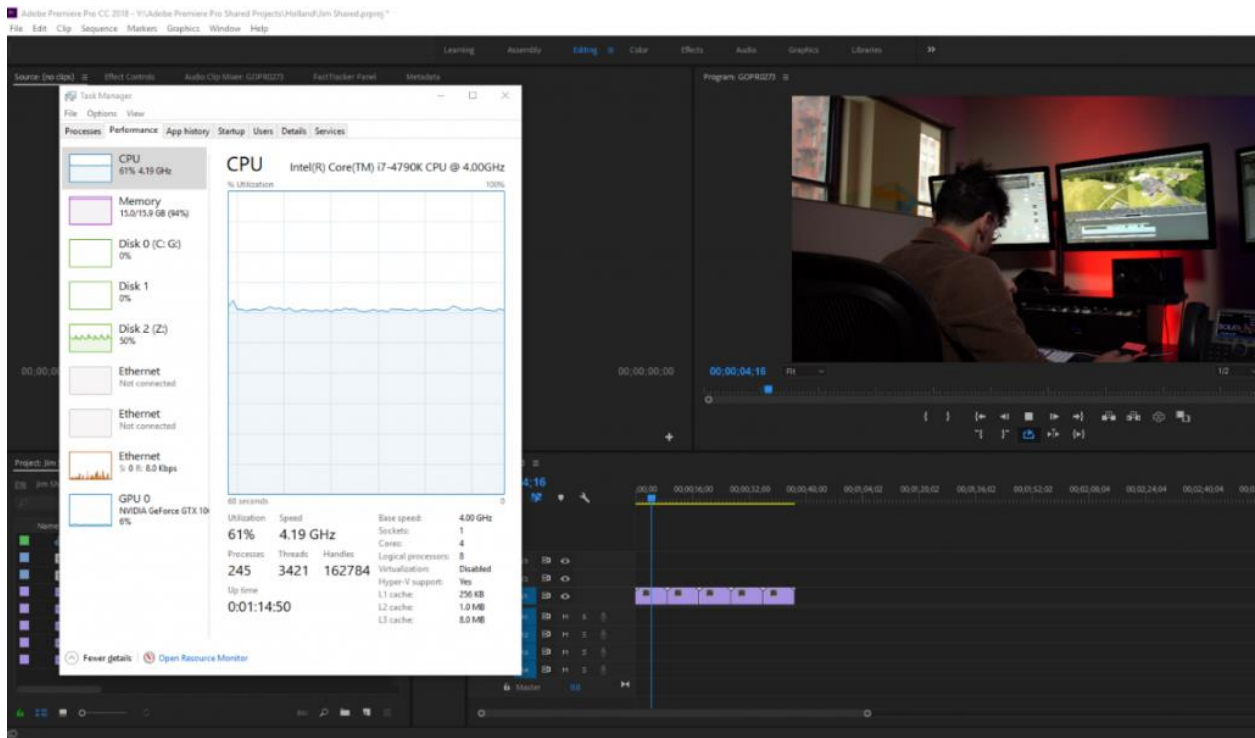
If your facility started out with 8Gb Fibre Channel in 2008, your growth to 32Gb FC is possible with as little as OM2 cabling. If you started out with 1Gb on Cat5e or 10Gb on OM3, your respective path to 10Gb or 25Gb may be cost-prohibitive. A facility may deploy 16Gb Fibre Channel within infrastructure built for 4Gb over a dozen years ago. Even with the most common OM3 cable today, a path to 64 and 128GB Fibre Channel link speeds through parallel optics will be possible.

Block-level Access Provides Additional Value

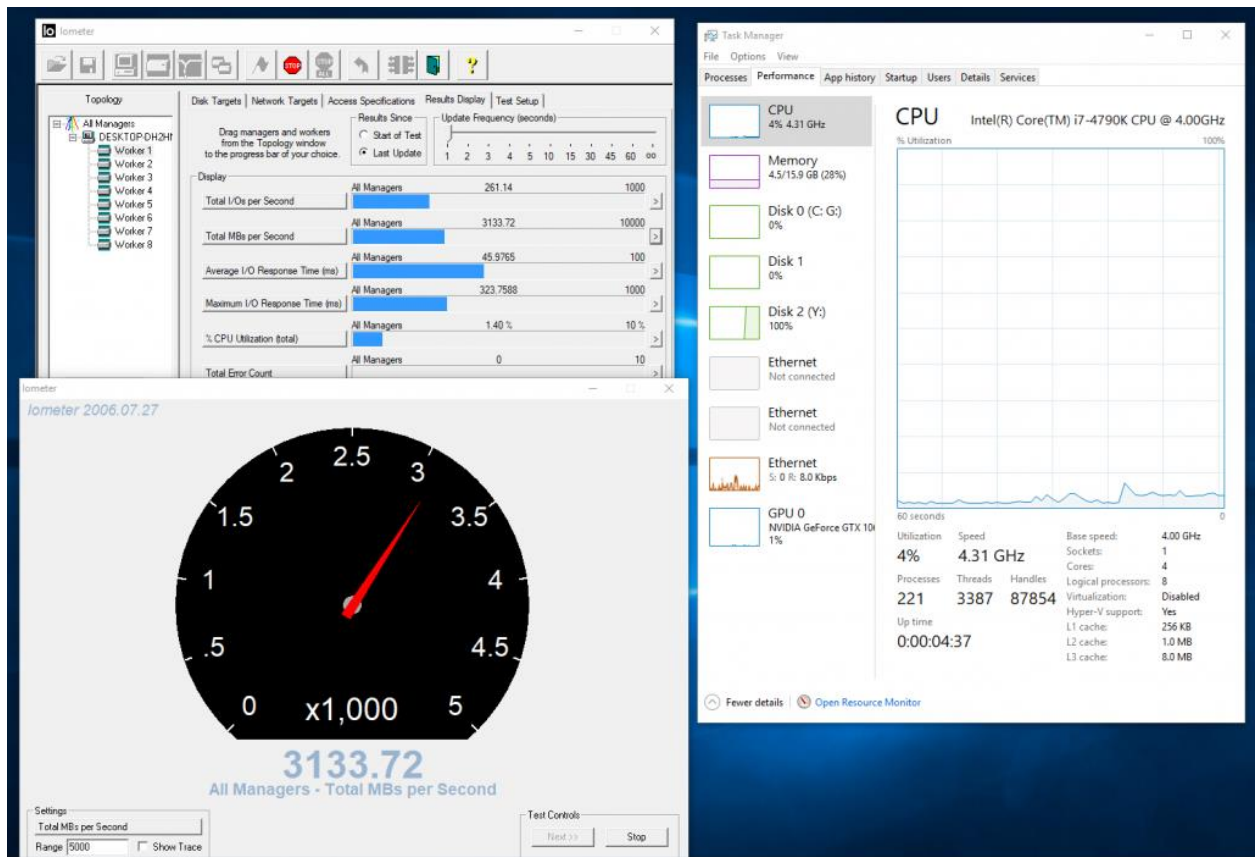
Since Fibre Channel speaks the language of SCSI, it's possible to expose raw SCSI devices over Fibre Channel, for block-level access. These devices will act and perform like a direct-attached device, complete with native filesystem for each client OS. One advantage is ultra-compatibility - not only is the device appearing like a local HDD, but it's seen in a native file system that every application can use. Another key advantage is speed. By removing the network management layer that allows for multi-user write permissions, the data is directly accessible at the block level of the hard drives. The image below is a benchmark of an 8Gb Fibre Channel-connected client with a Single-user (block-level) volume mounted.



This is the most efficient way to deliver data to the desktop. Practical saturation of 8Gb link speeds at 750-800MB/sec. CPU usage is near idle at 2%. In the Premiere Pro playback test, CPU usage is even lower than the Fibre Channel managed volume, at 61%.



When required for extreme bandwidth jobs, like 4K 60P uncompressed and EXR 16bit workflows, a 32Gb Fibre Channel block-level drive can produce speeds above 3000MB/sec.



CPU load for this bandwidth is 2%, or 1% CPU for every 1500MB/sec.

Conclusion

The technology of collaborative storage systems continues to follow the trend of IP-based topologies because of the wide accessibility and low cost of scale-out hardware and systems. However, for facilities looking to maximize their spending on future-proof connectivity, the value of dedicated, purpose-built networks for high-bandwidth data access can't be overlooked. As important as NAS can be for unlimited access to company data, investing in NAS alone can leave a facility spending more in supplementary storage to offload network traffic.

In partnership with ATTO Technology and using their 6th Generation Fibre Channel HBAs, Facilis Technology designs systems that take advantage of both the scale-out architecture of NAS, and the dedicated low-latency design of Fibre Channel SAN on the same turn-key server. To Facilis, these connectivity methods are complimentary - they can be used in parallel, and even in failover scenarios. ATTO Technology connectivity hardware and drivers enable Facilis to differentiate their shared storage solutions for the benefit of the content creator, facility owner and engineer.

As client-side bitrates increase due to higher quality imaging and aggressive broadcast requirements, more limitations will appear in NAS architecture that have already been eliminated by Facilis' custom Shared File System and ATTO Technology's Fibre Channel connectivity options.

For more info visit www.facilis.com or email sales@facilis.com.

Learn how ATTO and Facilis work together <https://www.atto.com/solutions/facilis/>

ATTO Fibre Channel solutions <https://www.atto.com/products/fc-adapters>