

Post Production Video Editing Solution Guide with Apple Xsan File System **AssuredSAN™ 4000**



INTRODUCTION

Dot Hill Systems offers high performance network storage products that are ideally suited for environments where distributed access to data is desired, and high speed streaming access to the data is an absolute must. In particular, the AssuredSAN™ 4000 series storage arrays are specifically designed for networked access, 99.999% availability and streaming performance of over 5000 MB/s.

One of the more common data streaming environments is the Video Post-Production Editing studio using Apple Mac Pro Workstations and a professional editing suite such as Final Cut Pro or Adobe Premiere. Apple has done an exemplary job over the years at providing simple, intuitive interfaces and tools on their platforms that specifically address the needs of Video Editors. However, video editing is one of the most data intensive applications on a computer. This means that the requirements of the system demand fast access to large amounts of data, usually in a sequential (streaming) fashion. At the same time, modern workflows require that the video files be available at any one of many Workstations, since the same video may be edited in different ways by different individuals at different Workstations. These two pervasive characteristics of an efficient Video Editing studio tend to be at odds with each other:

- 1) High speed access means using local storage, or high speed access to networked storage.
- 2) Distributed access means using File Serving in order to access a file from any Workstation. Local storage isn't an option here. However, most file serving solutions are far too slow for the streaming needs of an Editing Workstation.

These two requirements are commonly addressed with the Xsan File System, from Apple. The distinctive characteristics of this solution is the distributed nature of the file system (files can be accessed from any Workstation), and the high speed access to the storage. When matched with the AssuredSAN 4000 storage array, all the pieces are in place for a professional video editing environment.

Unfortunately, the task of identifying all the necessary pieces and putting them together can be confusing and intimidating to the un-initiated. And storage vendors don't help matters by offering a confounding mix of competing and overlapping protocols, form factors, drive types and interfaces. Coupled with Apple's exit from the server market and the landscape is more puzzling than ever. It is clear that simple guidelines are needed to point the way.

This whitepaper will outline the elements of an Apple Mac-centric Video Post-Production Editing environment that includes high speed access to all data files from every Workstation. It will include recommendations for equipment, connectivity, configurations, and options, along with best practice guidelines to assist customers in optimizing the performance and reliability of their investment. The intent is to guide the customer or the integrator along the path to success in deploying a distributed storage subsystem for Mac Workstations.

THE VIDEO POST-PRODUCTION SOLUTION

One of the most compelling solutions utilizing networked storage in an Apple environment is the Video Post-Production editing studio. In this environment, multiple editing Workstations are used for manipulating one or more video streams, and then rendering and saving the results to disk. Following are the critical needs of a Video Post-Production environment:

- Capacity Video has an insatiable appetite for Megabytes.
- Cost Value is an important aspect of the solution.
- Speed High speed access to data allows for streams to be uninterrupted (no dropped frames) during editing sessions.
- Availability This means that the video files should be accessible from any Workstation at any time. It also may refer to the robustness of the solution. A highly available solution can sustain the failure of any single component without losing access to any data.

Since the editor needs to have access to any video file from any Editing Workstation, it is necessary to provide a networked file system for this purpose.

A distributed file system, such as the Apple Xsan file system, can provide the fast, streaming access to data that is necessary for video post-production studios. At the same time, the file system is designed to share the resources of the file volumes among multiple client systems.

When combined with Dot Hill Systems 4000 Storage Arrays, and 8Gbit Fibre Channel infrastructure, this video editing solution can deliver more than 4500 MB/s to the Apple Mac Pro Workstations used for video editing. In video terms, that is more than 20 uncompressed High Definition video streams being rendered at the same time.

The key to constructing a high performance video editing studio is in the selection of components, the proper network configuration, and the proper storage configuration. The following sections include an overview of recommended equipment, configuration diagrams, and a list of best practices.

SOLUTION COMPONENTS



Storage Array. The Dot Hill AssuredSAN 4000 is the RAID array storage subsystem for this solution. It features dual, active-active RAID controllers with eight 8Gbit Fibre Channel interfaces. The array supports SAS disk drives in the 2.5" and 3.5" form factors.



Metadata Controller. It is absolutely necessary to have at least one metadata server in the Xsan environment. We recommend the use of the Apple Mac Mini to serve this role. If high availability is a requirement, then two metadata servers must be deployed.



Fibre Channel / Thunderbolt Bridge. The Mac Mini does not support Fibre Channel interfaces, nor does it allow the addition of a Fibre Channel HBA card. It does, however, support the high speed Thunderbolt interface. Therefore, to connect the Mac Mini to the underlying storage subsystem, a bridge from Thunderbolt to Fibre Channel is needed. The recommended product for this role is the ATTO ThunderLink FC 1082 bridge.



Client Workstations. The Workstations play the role of the video editing Workstation. A video post-production environment will have one or more of these Workstations deployed, all working from the same storage pool. The Apple Mac Pro Workstation is commonly used for Video Editing Workstations.



Distributed File System. The Xsan distributed file system allows access to video files from any Workstation. This software is pre-installed and pre-licensed on the Apple OS X Mountain Lion Operating System.



Fibre Channel HBA. The ATTO Celerity FC-8xEN 8Gbit Fibre Channel HBA is the only supported HBA in Apple OS X configurations with AssuredSAN storage systems. This HBA offers 8Gbit speeds and an MPIO driver that implements the ALUA specification. This means that the driver will be able to distinguish between paths that are preferred (optimum), and paths that are not.



Fibre Channel Switching. The infrastructure of the storage network requires the deployment of at least one Fibre Channel switch. In order to realize the benefit of the 8Gbit interfaces, the switch should also be rated at 8Gbit.



Fibre Channel



Thunderbolt

Connections. Most of these components are connected with fiber optic cables. In the terminology of cables and connectors, the type you will need is LC-LC. You can purchase these cables in a variety of lengths. These cables will connect directly into the ports provided by the ATTO HBA, the ATTO FC 1082 Bridge, the AssuredSAN 4000 storage system, and the Fibre Channel Switch. In figure 1, each orange line represents one Fibre Channel fiber optic cable. **IMPORTANT NOTE:** These cables plug into a module called a “transceiver” on the port. Most components already provide the transceivers pre-installed on all the ports. However, some Fibre Channel switches do not. You may have to purchase transceivers separately for the ports on the switch you purchase, along with separate licenses to use those ports.

The only exception in the storage cabling is the connection between the Mac Mini and the ATTO Bridge. The cable for this connection is a Thunderbolt cable.

SOLUTION CONFIGURATION

Refer to Figure 1 for an example of how to deploy and connect Workstations and storage in an Xsan environment. The Mac Mini's in the figure will serve as the Xsan Metadata Controller (MDC). All the rest of the Workstations are deployed as Video Editing stations.

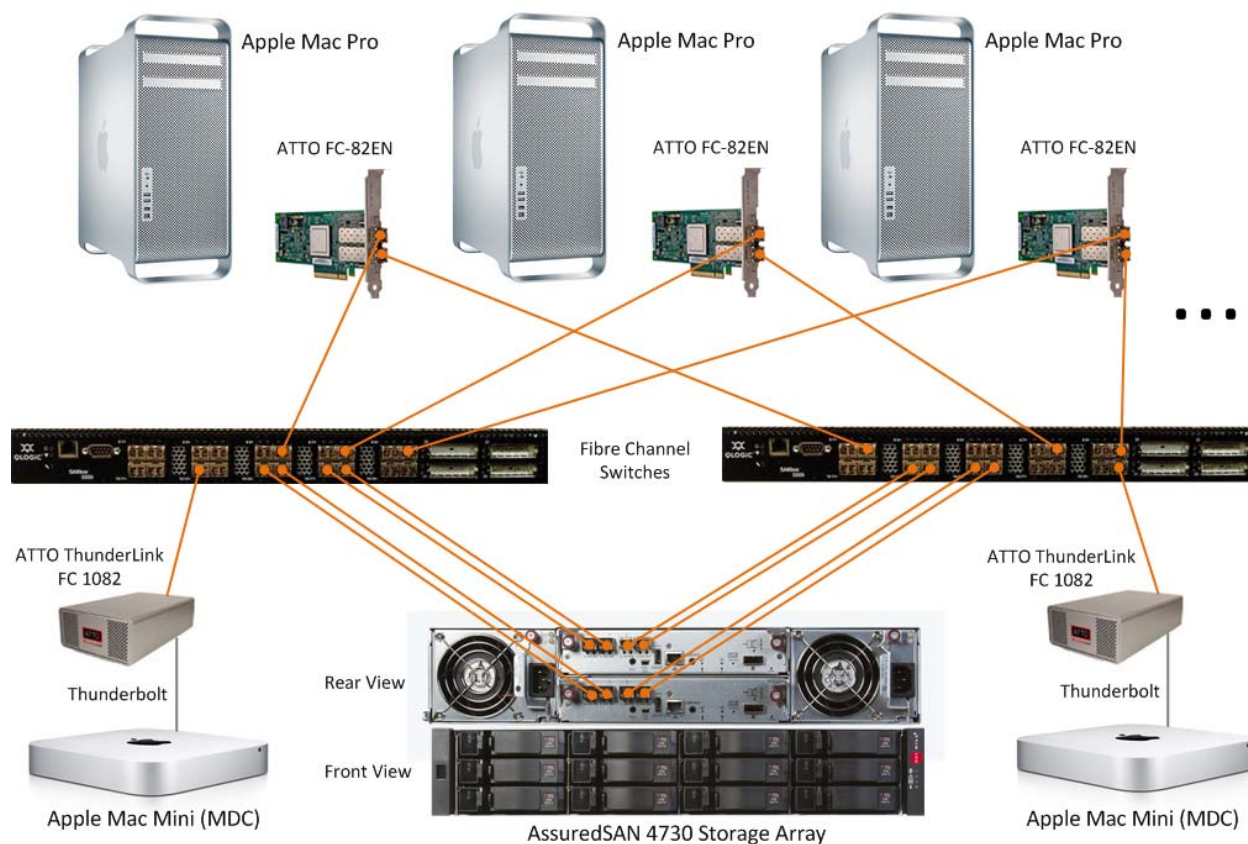


Figure 1

Video Post-Production Editing Solution – Fully Redundant High Availability Configuration

SOLUTION BEST PRACTICES

When configuring the storage subsystem for a Video Post-Production Editing environment, the following Best Practices will aide in configuring the solution for high performance and reliability.

1) ATTO Celerity FC-8xEN HBA Configuration Best Practices

- a) Download and install the vendor specific, Dot Hill 2.02MP (or later) multipathing driver from ATTO. This driver is a full featured MPIO solution that implements the ALUA specification. This means that the driver can distinguish between paths that are preferred and paths that are not. NOTE: You may need to separately upgrade the ATTO Configuration Tool to the latest version.
- b) The software driver for the Celerity HBA requires the first mapped LUN to be identified as LUN 0. This requirement may be met by simply assigning a LUN number of 0 to the first LUN that is presented to the host from the storage array.
- c) The **Link Down Timeout** value for each channel is set to zero by default. In order to allow for the temporary interruption of I/O that occurs in the case of a controller failover, this timeout value should be set should be set to something other than zero. Dot Hill recommends this timeout value be set to 60. This will allow sufficient time for controller failover in the unlikely event of a controller failure.

2) Consider the need for High Availability

Redundant components in a storage network are necessary if continuous access to storage is of paramount importance. The Dot Hill AssuredSAN storage system comes pre-configured with redundant controllers and redundant power supplies. The loss of any one of these components will not compromise access to storage.

A fully redundant storage network will also include two Fibre Channel switches and two Xsan Metadata Controllers. This configuration allows for any single component in the system to fail, and still provide un-interrupted access to all storage.

Beyond the deployment of redundant components, another important aspect of redundancy in a storage network is in the cabling. Proper cabling provides redundancy, so that the failure of any one component or path does not compromise access to any of the storage. At the same time, this redundancy should also provide a mechanism for aggregating throughput, delivering higher performance to the host systems. Refer to the configuration diagram in Figure 1 for an example of proper redundant cabling.

If budget is constrained, and High Availability is not a requirement, then the following optimizations may be considered:

- a) Deploy only a single Fibre Channel switch, with the assumption that a single switch will provide sufficient connectivity to the clients.

b) Deploy only one Mac Mini metadata Controller.

3) Use Nearline (7K RPM) SAS drives.

These drives provide the best MB/\$ value, provide the highest storage density, and provide the reliability of an enterprise class disk drive.

4) Dedicate an Ethernet subnet for metadata

Every Apple Mac Pro Workstation comes with at least two 1Gbit Ethernet ports built into the motherboard. One of these ports on each Xsan client should be connected on an isolated and dedicated subnet, and used for the primary Xsan metadata traffic. This will guarantee bandwidth for metadata traffic for Xsan, and reduce the potential for delays in delivery of data.

The Mac Mini, deployed as a Metadata Controller, has a single Ethernet port. The recommended deployment configuration uses this port in the isolated network with the other clients. This configuration negates the access of the Mac Mini from outside of the isolated subnet, which will impact remote management of the Metadata Controller.

5) One Volume (LUN) per RAID group is best

In a configuration where sequential throughput (i.e. video streaming) is important, then configuring the physical disk drives for minimal seek operations is a good idea. If multiple Volumes are created on a RAID group (VDisk), then the physical disk will require seek operations as it goes from one Volume to the other to perform I/O. To prevent unnecessary seek operations; the configuration should be designed with only one Volume per RAID group (VDisk).

6) Configuration of RAID groups is important

In Video Post-Production environments, fast sequential (streaming) access to storage is important. Also important is optimized use of disk space. These two factors combine to suggest that RAID 5 or RAID 6 is the best option for data volumes in an Xsan environment. These RAID levels provide redundancy protection against drive failure, yet also provide high speed sequential access to data.

The volumes for Xsan metadata and journal are best on RAID 10 or RAID 1 disk groups. These RAID levels provide the best performance in random workloads, which is a characteristic of how these volumes are used.

7) Present Volumes (LUNs) to all Workstations and Servers

To optimize the features of Xsan, all volumes (LUNs) should be presented to all the hosts via all the FC ports (default presentation). Depending on how ports are connected between the storage system and the hosts, there will likely be more than one path to the Volume. Not only that, but some paths may be more “optimized” or “preferred”, meaning that they will operate faster than other paths. The MPIO driver on each of the host systems will detect these

conditions and automatically configure how to use all the paths properly. The driver will balance the load ONLY across the “preferred” paths. The non-preferred paths will be held in standby in case the preferred paths fail.

8) Spend some time tuning

Published tuning guidelines for Xsan indicate that the optimum block size for a streaming application is 64K, and that the associated stripe breadth should be 16K, so that the product of the two parameters is 1MB; an optimum transfer value for OS X and storage systems.

However, in actual testing, we have determined that the general formula does not necessarily apply to Xsan file systems on AssuredSAN 4000 series storage. In the Xsan configuration we evaluated, a block size of 16 and a stripe size of 16 across 4 LUNs delivered the optimum performance in single user / single volume workloads. The customer is encouraged to evaluate various combinations of block sizes and stripe breadths to determine the optimum performance for their environment.

CONCLUSION

The Dot Hill AssuredSAN 4000 storage array is an ideal storage solution in Apple computing environments. The AssuredSAN has been demonstrated to work seamlessly with Apple Workstations, supported HBA's, and the Apple Xsan distributed file system. In addition to the inherent benefits of SAN based storage, the AssuredSAN storage array delivers:

- High Performance – 8Gbit FC speeds
- Reliability – redundant components and paths
- High Availability – Clustered controllers with failover technology
- Flexibility – a variety of drive options and configurations

The Apple Xsan distributed file system configuration with AssuredSAN 4000 storage is of special interest to the Video Post-Production Editing market. This solution offers distributed access to storage, a necessary element of a shared editing environment. Coupled with the ATTO Celerity FC HBA, the solution utilizes 8Gbit Fibre Channel speeds, path optimization and load balancing to deliver tremendous throughput for intensive streaming and rendering applications.