

# Dot Hill AssuredSAN™

8Gb Fibre Channel SAN Storage

## Apple Systems Best Practice Guide

*Including Recommendations for Video Post Production Environments*



## Introduction

Dot Hill Systems has added support for the OS X operating system from Apple, Inc. with the AssuredSAN 3000 Series Storage System products. This means that customers of Apple products, such as the Xserve server and the Mac Pro workstation can now benefit from the features, performance and reliability provided by Dot Hill AssuredSAN storage arrays.

This whitepaper will outline the connectivity options available to customers, as well as best practice guidelines to assist customers in optimizing the performance and reliability of their investment.

In addition, special consideration will be given to a popular configuration in Video Post Production Editing environments. This paper will outline the components, configuration, and best practices needed to establish an editing studio that is fast, reliable and cost effective.

## Connectivity Guide

This guide addresses the Fibre Channel connectivity between the Dot Hill Systems AssuredSAN 3000 Series Storage Arrays, and Apple workstations, servers and software. Specifically, the following products are covered:

- 1) The Dot Hill Systems AssuredSAN 3000 Series RAID Array
  - a. DH 3930                      Fibre Channel & iSCSI host interface, 3.5" disk drives
  - b. DH 3920                      Fibre Channel & iSCSI host interface , 2.5" disk drives
  - c. DH 3730                      Fibre Channel host interface, 3.5" disk drives
  - d. DH 3720                      Fibre Channel host interface, 2.5" disk drives



AssuredSAN 3920 / 3720



AssuredSAN 3930 / 3730

- 2) Apple Computers and Software
  - a. Mac Pro Workstation 4.1
  - b. Xserve "Xeon" Server
  - c. Mac OS X Snow Leopard Operating System (version 10.6)
  - d. Apple Xsan File system (version 2.2)



Apple Mac Pro Workstation 4.1

- 3) The ATTO Celerity FC-8xEN Fibre Channel HBA
- a. ATTO Celerity FC-81EN      Single port
  - b. ATTO Celerity FC-82EN      Dual port
  - c. ATTO Celerity FC-84EN      Quad port

ATTO Celerity  
FC-82EN

## Configurations & Best Practices

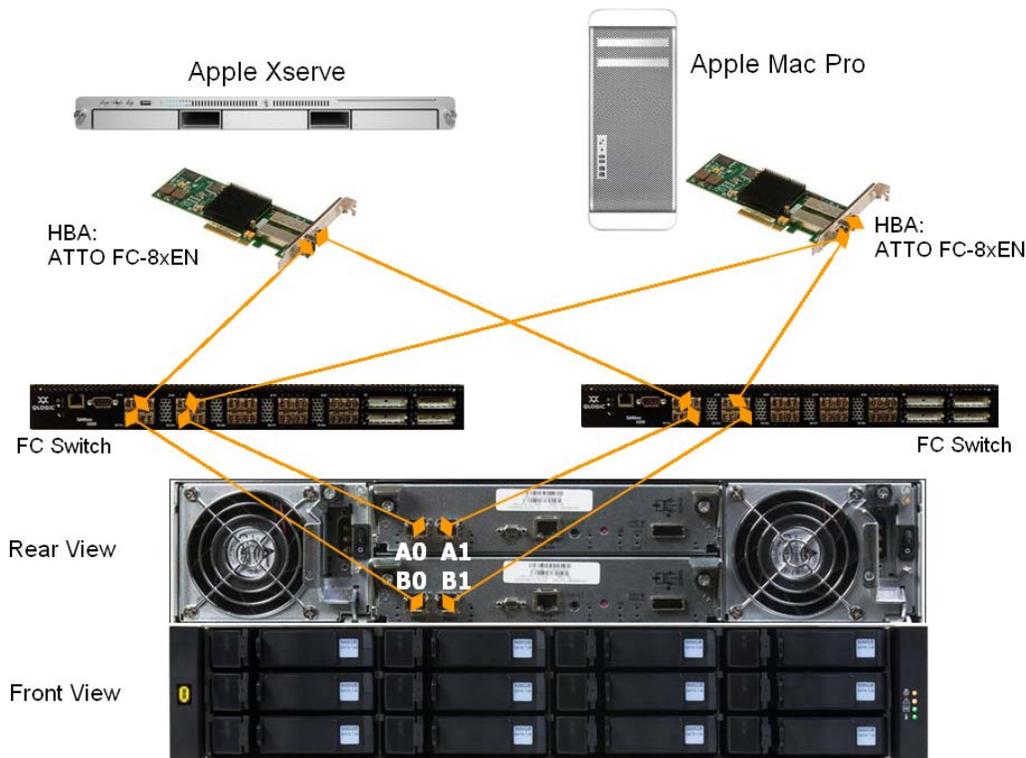
The following sections will review recommended configurations, Best Practices, and known issues with Apple systems and Dot Hill storage arrays.

We will begin with a general overview of a basic hardware configuration and generic Best Practices. These guidelines will apply to virtually all configurations and deployments of Dot Hill storage in an Apple environment.

In addition, a specific configuration for Video Post Production environments will be examined. The distinctive feature of this configuration is the Xsan File System, and its ability to distribute access to the file system among multiple client systems. The distinguishing traits of this vertical segment will require some special software and configuration considerations.

## Basic Configuration

The Basic Configuration is a generic SAN infrastructure connecting a Dot Hill AssuredSAN RAID storage array to one or more Apple Xserve servers and/or Apple Mac Pro workstations. This is a very common SAN configuration, used in most storage networks to provide optimum use of bandwidth while maintaining redundant connectivity to storage. An example of this configuration is shown in Figure 1.



**Figure 1**

Apple / AssuredSAN Reference Configuration

## Storage Configuration Best Practices on Apple

- 1) Make sure you are optimized for throughput and redundancy

Redundant components in a storage network are necessary if continuous access to storage is of paramount importance. The Dot Hill AssuredSAN storage system typically comes pre-configured with redundant controllers and redundant power supplies. The loss of any one of these components will not compromise access to storage. In addition, deploying multiple switches in the network is necessary to eliminate the switch as a single point of failure.

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.

## 2) Distribute the Load

The DH 3730 RAID Array typically comes with dual Active-Active controllers. With two controllers, it is important to optimize the processing power of both controllers in order to realize the fastest throughput to the system. To do this, the RAID groups that are defined on the array should be evenly distributed among the two controllers (each RAID group is assigned to an “owning” controller). For example, if there are 12 drives in the array, then it makes sense to create two RAID groups of 6 drives each; one RAID group assigned to controller A, and the other RAID group assigned to controller B.

## 3) Be aware of limitations to expanded volumes

The Apple Mac OS Extended File System (HFS+) has limited capability in addressing the new space of a Volume (LUN) that has been expanded at the array level. The system will recognize the existence of the new space, but new partitions cannot be created in the new space, nor can existing partitions be expanded into the new space. The only mechanism for utilizing the space is to first destroy all existing partitions on the Volume.

# ATTO Celerity FC-8xEN HBA Configuration Best Practices

## 1. Install the best Driver

Download and install the vendor specific Dot Hill 1.31MP (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.

## 2. Select an appropriate MPIO policy.

The Celerity HBA provides 3 different load balancing policies that can be selected on a LUN basis or Target basis. The default policy is “Pressure”. This algorithm balances the load based on the amount of throughput each port is currently handling. This default policy will provide the optimum performance for most applications, especially video editing and streaming.

## 3. Be aware of LUN 0

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.

## 4. Set the **Link Down Timeout** value

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 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.

## Video Post Production Configuration

One of the most compelling solutions utilizing networked storage in an Apple environment is the Video Post Production 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:

- 1) Capacity. Video has an insatiable appetite for Megabytes
- 2) Cost. Value is an important aspect of the solution. This factor is more important than providing extra redundancy and reliability into the solution.
- 3) Speed. High speed access to data allows for streams to be uninterrupted (lost frames) during editing sessions.
- 4) Availability. This means that the storage should be accessible from any workstation at any time.

Since the editor needs to have access to any video file from any editing workstation, it is necessary to provide networked storage for this purpose.

A clustered 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 3730 Storage Arrays, and 8Gbit Fibre Channel infrastructure, the Xsan file system can deliver more than 1000 MB/s to the Apple Mac Pro Workstations used for video editing. In video terms, that is more than 5 uncompressed High Definition video streams being rendered at the same time.

### Video Post Production Configuration – Components

Metadata Server	It is absolutely necessary to have at least one metadata server in the Xsan environment. If the budget is really tight, then one of the Xsan client systems can serve in this role. If redundancy and high availability is a concern, then two metadata servers should be deployed. The Apple Mac Pro Workstation can be easily deployed as a Metadata Server.
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.

Xsan	The Xsan distributed file system allows access to video files from any workstation. The software must be installed and licensed on each metadata server and each client workstation.
Fibre Channel HBA	The selection of a Fibre Channel interface in the Apple servers and workstations will have a significant impact on performance and efficiency. The ATTO Celerity FC-8xEN 8Gbit Fibre Channel HBA is the only supported HBA that provides 8Gbit speeds on its ports. In addition, this HBA offers an MPIO driver that implements the ALUA specification. This means that the driver will be able to distinguish between paths that are preferred, and paths that are not.
FC 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. With a single switch, it is recommended that the switch be configured into two equal zones (two virtual switches) to facilitate simple management of the paths to storage.
Storage Array	The DH 3730 is the RAID array storage subsystem for the solution. It features dual, active-active RAID controllers with four 8Gbit Fibre Channel interfaces. The base configuration supports up to 12 SAS or SATA disk drives. The configuration can be expanded to include more drives by adding DH 3130 JBODs via SAS interconnects.
Connections	All of these components are connected with fiber optic cables. In the terminology of cables and connectors, the type you will need are 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 DH 3730 storage system, and the Fibre Channel Switch. In figure 1, each orange line represents one Fibre Channel fiber optic cable. <b>IMPORTANT NOTE:</b> These cables plug into a module called a “transceiver” on the port. Most components already provide all the transceivers pre-installed on all the ports. However, some Fibre Channel switches do not. You may have to purchase transceivers separately for some of the ports on the switch you purchase, along with separate licenses to use those ports.

Refer to Figure 1 for an example of how to deploy and connect workstations and storage in an Xsan environment. One of the workstations in the figure will act as the Xsan Metadata server. All the rest of the workstations will be deployed as Video Editing stations. If only one switch is to be used, then a best practice is to configure the switch into two equal zones. In this case, the two switches shown in the figure represent the two zones in one switch.

## Video Post Production Configuration – Best Practices

All of the Best Practices outlined for the Basic Configuration also apply to the Xsan Configuration for Video Post Production environments. In addition, the following specific guidelines apply:

- 1) Where can you scrimp and save
  - a. Eliminate redundancy. Most Video Post Production environments are not adversely impacted if the data is temporarily unavailable due to a component failure. Deploy a single Fibre Channel Switch, and a single Xsan Metadata Server.
  - b. Use SATA or Nearline SAS drives. These drives will provide the best MB/\$ value and also provide the highest storage density. Use RAID 6 to get extra protection from drive failure, since these drives are not as reliable as SAS drives.
  - c. Consider fewer paths to storage. Dual ported HBAs provide extra throughput and redundancy to storage. However if redundancy is not a priority, then single ported HBAs may be used. These HBA's can still provide enough throughput to handle almost any video editing workload.

- 2) Dedicate an Ethernet subnet for metadata

Every Apple Mac Pro Workstation and Apple Xserve Server comes with at least two 1Gbit Ethernet ports built into the motherboard. One of these ports on each Xsan client and metadata server 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.

- 3) 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).

- 4) Make an intelligent choice of RAID level

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. RAID 6 provides two physical disks worth of redundancy, allowing for two drive failures in a RAID group without losing data. This RAID level is recommended in SATA configurations, where individual drives are less reliable.

#### 5) 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.

#### 6) 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 3000 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.

## What about the LSI7x04EP FC HBA?

The official Apple branded and supported Fibre Channel HBA’s are the MB842G/A and the MB843G/A. These products are re-branded from LSI and are the OEM equivalents to the LSI7204EP and the LSI7404EP. Dot Hill does not recommend or support these HBA’s with AssuredSAN storage for the following reasons:

- 1) Advanced LUN management features, such as ALUA are not implemented in the driver. This means that all paths are treated equal, even though some paths are optimum, and other paths are not.
- 2) The HBA does not seamlessly failover LUNs (failover can cause the driver to lock up).
- 3) Support from Apple and LSI is limited.
- 4) The LSI 7x04EP is older 4Gbit Fibre Channel technology.

In contrast, Dot Hill and ATTO Technologies offer the newest LUN management features and Fibre Channel speeds in their solutions, along with full support for Apple Workstations, Servers and distributed file systems.

## What about iSCSI?

Up to this point, we have gone on and on about how wonderful Dot Hill Systems storage arrays work with Apple workstations and servers using a Fibre Channel network. But what about iSCSI? There has been no mention of it, even though iSCSI offers a very compelling alternative in many environments. Besides the low cost and ubiquitous nature of Ethernet equipment, the iSCSI protocol also offers relatively simple management and support for common storage networking features such as MPIO. Before we conclude our guide on Apple solutions, we will now address the issues around using the iSCSI protocol.

While iSCSI connectivity is within the realm of possibilities in connecting Apple systems to Dot Hill storage, there are some limitations to consider.

- 1) The currently deployed infrastructure of Ethernet is limited to 1Gbit of speed on each port. While this is plenty fast for most e-mail and file serving applications, it is not nearly fast enough for the demands of video streaming environments. Since video streaming consists primarily of sequential throughput to and from a storage device, the bandwidth of the connection is of paramount importance.
- 2) Apple support for the iSCSI protocol has thus far been somewhat lacking. Apple's OS X Snow Leopard Operating System does not offer an iSCSI initiator built-in to the Operating System (unlike Windows Server, which includes an iSCSI initiator by default). For customers wanting to take advantage of iSCSI connectivity on Apple systems, the best approach is to obtain the Xtend SAN iSCSI initiator software from ATTO technology. This initiator is specifically designed for the Apple Mac OS X Operating System, is designed to industry standards and is target (i.e. storage array) agnostic.
- 3) The ATTO Xtend SAN iSCSI initiator is expected to function properly with the Dot Hill Systems 3000 series storage arrays that provide iSCSI connectivity. However, as of this writing, Dot Hill Systems storage arrays have not been qualified with the ATTO Xtend SAN iSCSI initiator (or any other iSCSI initiator on OS X). Therefore, customer support for Apple solutions with iSCSI connectivity will be limited. Please check with the Dot Hill Systems Customer Support site for up-to-date information on supported configurations.

## Conclusion

The Dot Hill AssuredSAN storage array is an ideal storage solution in Apple computing environments. The AssuredSAN has been demonstrated to work seamlessly with Apple Workstations and Servers, 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 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.