

ATTO TECHNOLOGY FASTFRAME 10 GIGABIT ETHERNET NETWORK ADAPTER CONFIGURATION GUIDE

Using the ATTO FastFrame NS11, NS12, and NS14 with
EMC Isilon OneFS scale-out NAS

- Installation guidelines
- Infrastructure considerations
- EMC Isilon OneFS and client OS performance tuning

EMC Solutions Group

Abstract

This white paper reviews the integration requirements for the ATTO FastFrame NS11, NS12, and NS14 10 Gigabit Ethernet network adapters when designing media and entertainment workflows for EMC Isilon scale-out NAS.

April 2012



Copyright © 2012 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All trademarks used herein are the property of their respective owners.

Part Number H10643

Table of contents

Executive summary	4
Business case.....	4
Solution overview.....	4
Key results/ recommendations	5
Introduction	6
Purpose.....	6
Scope.....	6
Audience	6
Guidelines	7
Client configuration	7
Network design considerations.....	8
EMC Isilon OneFS configuration	10
Conclusion	11
Summary.....	11
Findings.....	11
References	13
Other documentation.....	13

Executive summary

Business case

The ATTO Technology FastFrame NS 10 Gigabit Ethernet network adapter is an ideal complement to media and entertainment workflows featuring the EMC® Isilon® family of Network Attached Storage (NAS) devices. EMC Isilon scale-out NAS is easily deployed into existing media and entertainment workflows utilizing standard file-sharing protocols over existing 1 gigabit and 10 gigabit network infrastructure. The ATTO FastFrame Ethernet network adapter provides high-performance 10 gigabit connectivity for Mac OS X, Windows, and Linux clients using standard file-sharing protocols.

Broadcasters, motion picture studios, and media professionals worldwide trust EMC Isilon to provide the storage foundation for their media workflows. The X200 and S200 hardware platforms are the building blocks of a single scale-out NAS volume. EMC Isilon customers leverage the power of a single OneFS volume to consolidate the storage requirements for editing, media asset management, transcoding, and other workflows. A single EMC Isilon OneFS volume eliminates the need for costly silos of proprietary, application-specific storage at each stage of the workflow. EMC Isilon clients use open protocols such as SMB, NFS, iSCSI, FTP, and HTTP to access a OneFS volume, eliminating the need for costly per-client file-system licenses.

Media and entertainment workflows require flexible tools that protect your storage investment. EMC and ATTO Technology provide flexible tools based on open protocols. A workflow based on open network file-sharing protocols and the effortless scalability of EMC Isilon OneFS allows media and entertainment professionals to rapidly adapt to changing production requirements, ever-increasing content storage requirements, and increasingly complicated workflows. EMC Isilon easily adapts to your needs, without introducing costly downtime to your tight production schedules.

Solution overview

The ATTO Technology FastFrame NS 10 Gigabit Ethernet adapter is your gateway to high-performance EMC Isilon media storage. Properly tuned 10 Gigabit Ethernet networks allow EMC Isilon users the ease of use or a shared network volume and throughput previously restricted to a complicated Fibre Channel storage area network (SAN). When combined with a low-latency 10 Gigabit Ethernet switch from the Cisco Nexus 5000 or Arista 7000 product families and EMC Isilon scale-out NAS, the ATTO Technology FastFrame NS 10 Gigabit Ethernet network adapter provides reliable performance for even the most demanding media and entertainment workflows.

**Key results/
recommendations**

The ATTO Technology FastFrame NS is a PCIe 2.0 compliant network interface card optimized for high-performance access to media storage workflows on a local area network. ATTO's proprietary Advanced Data Streaming (ADS™) technology manages network latency to provide high-performance media file transfers and reliable playback of streaming media.

Introduction

Purpose

This guide covers performance tips related to installation of the ATTO Technology FastFrame NS 10 Gigabit Ethernet network adapter, client OS tweaks for 10 Gigabit Ethernet performance, and high-performance 10 Gigabit Ethernet network design considerations. This guide also reviews some of the many configuration changes available in EMC Isilon OneFS designed specifically to support the playback of high-resolution video in media and entertainment workflows.

Scope

This guide is intended for customers who wish to configure the ATTO Technology FastFrame NS11, NS12, and NS14 10 Gigabit Ethernet network adapter in a media and entertainment workflow containing both a 10 gigabit enabled EMC Isilon cluster and a 10 Gigabit Ethernet switch.

Audience

This guide is intended for storage engineers and network engineers supporting EMC Isilon storage in media and entertainment workflows. This guide also covers server and workstation configuration parameters that will prove useful to technical staff supporting clients of EMC Isilon storage. For additional information about performance tuning parameters in 10 Gigabit Ethernet workflows, please contact [EMC Isilon support](#).

Guidelines

Client configuration

For best performance, install the ATTO Technology FastFrame NS 10 Gigabit Ethernet network adapter in a PCIe 2.0 slot that is electrically wired for an 8x lane width or greater. Please consult your server or workstation documentation for additional information regarding the host computer PCI slot configuration. Table 1 demonstrates the bandwidth limitations of PCIe 1.0 versus 2.0 at varying lane widths.

Table 1. PCIe bandwidth

Lanes	Bandwidth Per Direction PCIe 1.0	Bandwidth Per Direction PCIe 2.0
1	250 MB/s	500 MB/s
2	500 MB/s	1 GB/s
4	1 GB/s	2 GB/s
8	2 GB/s	4 GB/s

Ensure that your clients are running current OS patch levels and the latest version of the ATTO Technology FastFrame 10 GbE NIC drivers.

For optimal 10 Gigabit Ethernet performance, the following entries are added to the file `/etc/sysctl.conf` on clients running Linux or Mac OS 10.6 and greater:

```
net.inet.tcp.sendspace=2097152
net.inet.tcp.recvspace=2097152
net.inet.tcp.maxseg_unacked=32
net.inet.tcp.delayed_ack=2
```

The `net.inet.tcp.sendspace` and `net.inet.tcp.recvspace` system control values determine the client OS transmit and receive buffer sizes. The minimum recommended transmit and receive buffer sizes for a client's network connectivity to EMC Isilon scale-out NAS may be determined using the following equation:

$$\text{Network bandwidth (bytes per second)} \times \text{round-trip delay (seconds)} = \text{buffer size (bytes)}$$

The `net.inet.tcp.maxseg_unacked` system control determines the maximum number of outstanding TCP segments left unacknowledged.

The `net.inet.tcp.delayed_ack` system control determines when TCP delayed acknowledgment is enabled. The system control is changed from the default value of 0 (automatically detect) to 2 (compatibility mode).

Jumbo frames support for the ATTO Technology FastFrame NS 10 Gigabit Ethernet network adapter is enabled via the Mac OS X network preferences pane. Please see the "Network design considerations" section of this document for guidelines on how to deploy a network with jumbo frames enabled.

Please see [EMC Isilon knowledgebase](#) article 3124, *Using Mac OS X Clients with Isilon OneFS 6.5*, for additional considerations when using EMC Isilon to support

high-performance Mac OS X clients, recommended NFS mount parameters, and directory service integration guidelines.

The Windows TCP stack is largely self-tuning and requires little tuning when the ATTO Technology FastFrame NS 10 Gigabit Ethernet network adapter is installed. A 64-bit operating system such as Windows 7 x64 or Windows Server 2008 R2 is recommended when deploying the ATTO Technology FastFrame NS 10 Gigabit Ethernet network adapter. In the Advanced tab of the network adapter properties, ensure that the Receive Side Scaling (RSS) Queues value matches the number of available CPU cores in the host computer. If the host computer runs a particularly CPU-intensive application, the RSS Queues are optionally left at the default value of 1. In the Performance Options section of the ATTO FastFrame network adapter properties, increase the Transmit and Receive Buffers from the default values of 512 to 4096. Larger Transmit and Receive Buffers improve the reliability of certain media playback and ingest applications by caching data in the host computer system memory during periods of network latency. Ensure that Large Send Offloads and RSS are not only enabled in the advanced tab of the network adapter properties, but also enabled in Windows via the command `netsh int tcp show global` in the Windows command prompt.

Network design considerations

EMC Isilon scale-out NAS clients demanding the highest possible 10 Gigabit Ethernet performance use 2 physically isolated networks, as shown below in the example in figure 1.

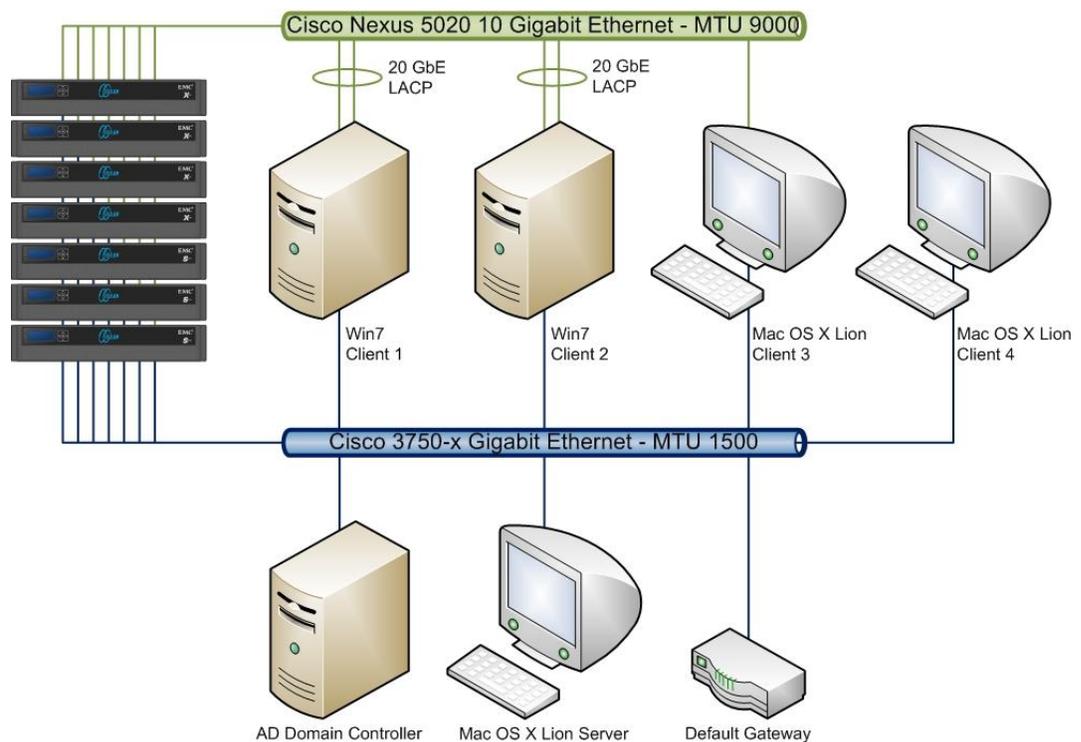


Figure 1. Multi-homed jumbo frames network topology

The primary network in this design is a standard Gigabit Ethernet network, shown at the bottom of figure 1. The Gigabit Ethernet network is common to all devices on the LAN using each device's on-board Gigabit Ethernet port(s). All devices on the gigabit LAN segment utilize a standard Maximum Transmission Unit (MTU) of 1500 bytes per packet. The Gigabit Ethernet network is a Layer 3 network, and each host on the network is configured with a default gateway address. Clients on this network with ATTO Technology FastFrame NS 10 Gigabit Ethernet network adapters installed are configured for a dual-homed network configuration. The 10 Gigabit Ethernet network shown at the top of figure 1 is common to only 10 gigabit enabled clients. Unlike the Gigabit Ethernet network, the 10 Gigabit Ethernet network is a Layer 2 Ethernet switch configured to pass jumbo frames or packets with an MTU size of 9000. Each host interface on the 10 Gigabit Ethernet network is configured with an MTU of 9000. The 10 gigabit switch supporting this network is configured with an MTU of 9216 bytes in order to prevent the retransmission of packets exceeding 9000 bytes. Every client on the 10 gigabit LAN and every EMC Isilon 10 Gigabit Ethernet interface is configured with an MTU of exactly 9000 bytes per packet in order to prevent retransmission of packets due to fragmentation. EMC Isilon scale-out NAS clients utilizing a jumbo frame enabled LAN segment keep all interfaces using an MTU of 9000 bytes per packet isolated to a dedicated jumbo frames network. The only interface that is allowed to route traffic to other subnets is the Gigabit Ethernet interface configured with an MTU of 1500 bytes per packet. The MTU 1500 interface is also configured as the primary interface on each of the client operating systems. A "jumbo frames" enabled 10 Gigabit Ethernet network yields up to 10–20 percent better throughput in certain EMC Isilon workflows by more efficiently packetizing data for network file transfers. EMC Isilon scale-out NAS customers wishing to simplify their network infrastructure configure their 10 Gigabit Ethernet switches in a Layer 3/4 topology with an MTU of 1500 bytes per packet.

For optimal performance, 10 Gigabit Ethernet clients are connected to the same switch as the EMC Isilon 10 Gigabit Ethernet connections, as each "hop" on a network of tiered Ethernet switches adds latency to file transfers.

Workflows that require real-time video ingest and playback benefit greatly from large Ethernet switch buffers. Low-latency Ethernet switches with large shared buffers or dedicated buffers of 1 megabyte per port are highly recommended in a media network design. The Cisco Nexus 5000 or Arista 7000 product families are excellent choices when purchasing a 10 Gigabit Ethernet switch for media and entertainment applications. Note that certain quality of service (QoS) policies may create buffer partitions, reducing the available buffer size for individual ports supporting real-time video applications. When possible, it is recommended to use multi-homed network interfaces rather than QoS in media and entertainment network designs.

10 gigabit Windows 7 and Windows Server 2008 clients see improved read performance in certain workflows when bonding one or more 10 Gigabit Ethernet ports of the ATTO Technology FastFrame NS12 or NS14 10 Gigabit Ethernet network adapters using Link Aggregation Control Protocol (LACP). LACP must be enabled on the switch when using this protocol. Additionally, the network administrator must configure a unique LACP port channel group for each set of switch ports used in the aggregated network connection. For more information, please consult your Ethernet switch documentation. There is no performance benefit derived from LACP port aggregation in the EMC Isilon OneFS or Apple Mac OS X operating systems at this time.

EMC Isilon OneFS configuration

Most performance tuning for OneFS is performed at the folder or file level within a single file system. File system performance can be optimized for streaming video playback at one folder level, a nested subfolder one level down can be optimized for concurrent client access in a broadcast automation or transcoding workflow, and another subfolder can be optimized as an archive folder with up to an N+4 protection level. Workflows on smaller EMC Isilon clusters requiring real-time video playback performance benefit from setting an N+1 protection level on a portion of the OneFS file system and a “streaming” data access pattern versus the default “concurrent” data access pattern. Each of these workflow-specific file system modifications may be altered on the fly, without disrupting access to the file system and related workflows. For more information on file system optimization settings, please see the [OneFS data protection](#) and [Modify I/O optimization settings](#) sections of the EMC Isilon OneFS User Guide.

In environments where DPX, TARGA, or other image sequence playback is the dominant workflow, the file name prefetch system control is used to increase read performance. The file name prefetch feature automatically identifies access to a range of sequentially numbered files and proactively loads the rest of the image sequence into the server-side cache. This feature is set on a per-node basis, not cluster-wide. The file name prefetch feature is recommended only for nodes in EMC Isilon SmartConnect™ load-balancing zones dedicated to image sequence workflows. The following commands enable the feature on ALL files using the respective data access pattern:

```
sysctl efs.bam.fnprefetch.enable_streaming=1
sysctl efs.bam.fnprefetch.enable_concurrent=1
```

The following sysctls determine the number of files ahead in the sequence to prefetch:

```
efs.bam.fnprefetch.lookahead_streaming: 9
efs.bam.fnprefetch.lookahead_concurrent: 9
```

When enabled (1), the following sysctl determines if the entire file is loaded into cache. When disabled (0), the prefetch will simulate loading the first 32K of the file:

```
efs.bam.fnprefetch.prefetch_ll_all: 1
```

If the `efs.bam.fnprefetch.prefetch_ll_all` sysctl is enabled, the following sysctl will set a limit on the size of files that are prefetched into cache:

```
efs.bam.fnprefetch.prefetch_ll_size_limit: 67108864
```

The following commands will disable the feature:

```
sysctl efs.bam.fnprefetch.enable_streaming=1
sysctl efs.bam.fnprefetch.enable_concurrent=1
```

Again, the system control commands must be run on a per-node basis. If the node in question is subjected to workflows requiring random data access patterns, these settings may have a negative impact on client performance. Please contact the [EMC Isilon Global Customer Support Center](#) before making any modifications to the EMC Isilon OneFS system control settings.

Conclusion

Summary

Modern file-based media and entertainment workflows create significant challenges for data centers utilizing traditional storage technologies. EMC Isilon scale-out NAS provides a flexible, easy-to-manage storage platform. From ingest to edit, transcoding, streaming media, and archive, EMC Isilon provides a storage solution for every workflow within a single file system.

Findings

In tables 2 through 7, we see the single stream test results from the popular Windows Blackmagic Design® Disk Speed Test® and Mac OS X AJA® System Test® utilities. The tests are conducted with folders configured for N+1 data protection, streaming mode data access pattern. As you can see, the ATTO FastFrame 10 Gigabit Ethernet adapter provides a significant performance advantage over a standard on-board Gigabit Ethernet interface. The single stream performance of a properly tuned 10 Gigabit Ethernet interface easily accommodates some of the most challenging HD editing workflows. The simple, scale-out architecture of EMC Isilon allows this single stream performance to be replicated to dozens or hundreds of clients with no downtime. Please contact [EMC Isilon sales](#) to arrange a demo.

Table 2. 10 GbE MTU 9000 Windows 7 Single Stream SMB2 Performance

	Read	Write
Qty. 4 EMC Isilon X200 nodes	300 MB/s	250 MB/s
Qty. 3 EMC Isilon S200 nodes	320 MB/s	320 MB/s
Qty. 3 EMC Isilon S200 nodes ¹	320 MB/s	400 MB/s

Table 3. 1 GbE MTU 1500 Windows 7 Single Stream SMB2 Performance

	Read	Write
Qty. 4 EMC Isilon X200 nodes	90 MB/s	100 MB/s
Qty. 3 EMC Isilon S200 nodes	90 MB/s	100 MB/s

Table 4. 10 GbE MTU 9000 Mac OS X Lion Single Stream SMBX Performance

	Read	Write
Qty. 4 EMC Isilon X200 nodes	95 MB/s	135 MB/s
Qty. 3 EMC Isilon S200 nodes	115 MB/s	260 MB/s

¹ LACP disabled on cluster side, enabled on Windows 7 client side

Table 5. 1 GbE MTU 1500 Mac OS X Lion Single Stream SMBX Performance

	Read	Write
Qty. 4 EMC Isilon X200 nodes	45 MB/s	45 MB/s
Qty. 3 EMC Isilon S200 nodes	50 MB/s	50 MB/s

Table 6. 10 GbE MTU 9000 Mac OS X Lion Single Stream NFSv3 Performance

	Read	Write
Qty. 4 EMC Isilon X200 nodes	565 MB/s	240 MB/s
Qty. 3 EMC Isilon S200 nodes	575 MB/s	150 MB/s

Table 7. 1 GbE MTU 1500 Mac OS X Lion Single Stream NFSv3 Performance

	Read	Write
Qty. 4 EMC Isilon X200 nodes	100 MB/s	100 MB/s
Qty. 3 EMC Isilon S200 nodes	110 MB/s	105 MB/s

References

Other documentation

For additional information, see the documents listed below.

- EMC Isilon OneFS Operating System, [EMC](#)
- Suggested Tuning for ATTO 10GbE Adapters, [ATTO Technology Product Support](#)
- Using Mac OS X Clients with Isilon OneFS 6.5, [EMC Isilon KB article 3124](#)