



ethernet alliance

## **Storage Traffic in a Converged 10 Gigabit Ethernet Network**

**A Demonstration of a Reliable, Lossless Ethernet  
Fabric with Simultaneous Transport of FCoE, iSCSI,  
LAN and Multicast Traffic**

Version 1.0 May 2009

Authors:

NetApp  
Finisar Corporation  
Fulcrum Microsystems



## 1.0 Introduction

The advantages of IEEE 802.3™ Ethernet are well documented and established for implementation in traditional Local Area Networks (LAN). The role of Ethernet in the LAN data center is evolving beyond the interconnection of clients, servers and switches. The modern Ethernet, specifically 10 Gigabit Ethernet (10 GbE), is becoming the new interconnection technology of choice for Storage Area Networks (SAN) using the Internet Small Computer System Interface (iSCSI) protocol on an IP-based network. The usage of Ethernet infrastructure is being further extended by a relatively new protocol that maps Fibre Channel over Ethernet (FCoE). This allows servers, storage, and management devices to communicate over Ethernet links in a converged data center fabric.

## 2.0 Demonstration Highlights

This demonstration of FCoE traffic in a Converged Ethernet network features the real-time operation and interaction between two Ethernet-based storage networks - an FCoE SAN and an iSCSI SAN - utilizing some of the enhanced capabilities of Converged Ethernet, specifically IEEE P802.1Qbb also known as Priority-based Flow Control (PFC) and IEEE P802.1Qaz also known as Enhanced Transmission Selection (ETS) to manage critical storage data reliably. Participants in the demonstration include NetApp, Finisar Corporation, and Fulcrum Microsystems.



## 3.0 FCoE and Converged Ethernet Technology Background

FCoE and Converged Ethernet are the key enabling technologies driving convergence in next-generation Enterprise data centers. FCoE, which is in development by the INCITS T11-BB-5 Work Group, enables Fibre Channel data to be transported reliably over a Converged Ethernet fabric. With its simple one-to-one mapping structure, FCoE provides a seamless connection between the FC SAN and Ethernet host without requiring any change of FC applications. In this way, network administrators can leverage the ubiquity and superior economics of Ethernet technology while preserving the infrastructure, strengths, and tools of their existing Fibre Channel storage management framework.

The underlying foundation of FCoE is Converged Ethernet, an enhanced version of conventional Ethernet which provides the reliability and guaranteed delivery of Fibre Channel. These important enhancements, defined by the IEEE 802.1 Data Center Bridging (DCB) Working Group (WG), have been developed for the purpose of enabling convergence of various applications in data centers (LAN, SAN, and HPC) onto a single interconnect technology. By adding additional bandwidth control and management capabilities, Converged Ethernet achieves higher performance and lower latency.

Two critical components of Converged Ethernet being demonstrated include Priority Flow Control and Enhanced Transmission Selection. PFC is at the heart of the lossless capabilities of Converged Ethernet, guaranteeing zero packet loss through link level traffic flow control technology. ETS optimizes bandwidth efficiency and enables premium latency performance for mission-critical data through dynamic provisioning of bandwidth per traffic priority and assignment of minimum guaranteed bandwidth for specific traffic groups.

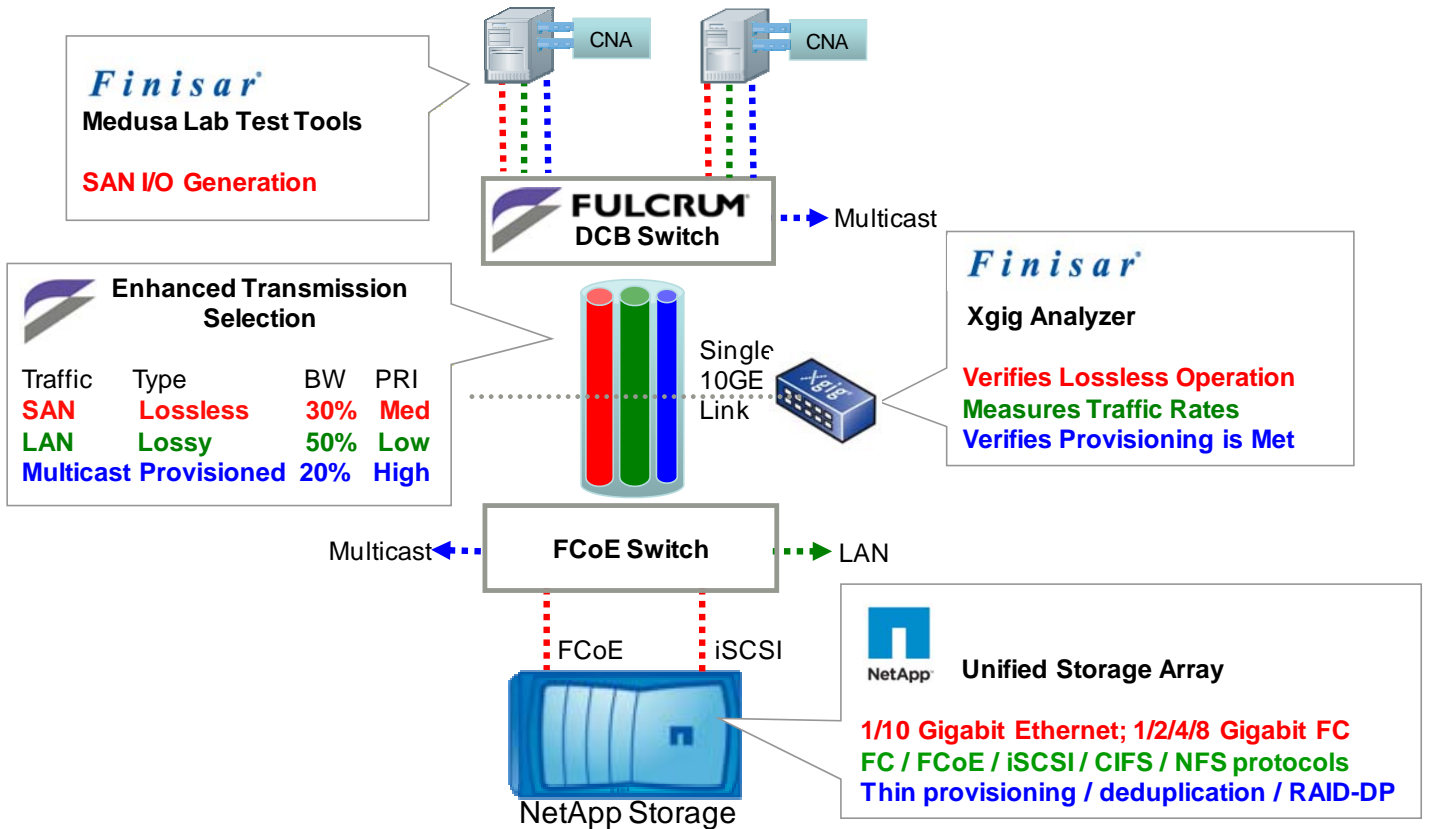


Figure 1—Demonstration Setup

The converged 10GbE demonstration network is comprised of:

- One Fulcrum DCB switch and one FCoE switch serving as the network fabric and implementing the enhancements required for Converged Ethernet
- Two storage networks based on FCoE and iSCSI technology to serve as source and destination for storage traffic across the Converged Ethernet fabric
- NetApp’s FAS 3140 providing unified Ethernet, Fibre Channel and Fibre Channel over Ethernet protocol support simultaneously within a single array to further demonstrate storage system convergence



- Various protocol test tools from Finisar to verify network functionality, measure network performance, and testify to the effectiveness of the converged storage network throughout the demonstration.

Three types of traffic - SAN, LAN, and Multicast - will be simultaneously transported across the network to establish a converged Ethernet traffic scenario and demonstrate effective traffic convergence:

*FCoE and iSCSI SAN traffic:* The red dashed line in the diagram indicates high bandwidth SAN traffic generated by storage I/O application tools. The Converged Network Adapters (CNAs) can initiate both FCoE and iSCSI traffic. The DCB switch and FCoE switch connect the initiator to the unified storage, which has both 10 Gigabit FCoE and iSCSI interfaces. The FCoE switch also forwards FCoE traffic.

*LAN traffic:* In addition to running SAN traffic, the network also carries client-server application LAN traffic on the same wire. LAN traffic is generated by running an interactive game between two servers and is represented by the green dashed line.

*Multicast traffic:* Indicated by the blue dashed line, Multicast traffic is added as a third type of converged Ethernet traffic between the two switches.

A major goal of this demonstration is to show how the PFC and ETS protocols enable Converged Ethernet. Specifically, the effectiveness of how ETS manages allowable bandwidth is shown by generating Multicast data representing traffic with multiple priority groups provisioned with different maximum throughput settings. Whenever a priority group exceeds the maximum allowable bandwidth, the PFC protocol automatically pauses the specific priority group to throttle down the traffic without packet loss. The impact on performance, throughput, zero packet loss, and latency of PFC and ETS on all three traffic types are shown using a protocol analyzer monitoring the throughput of individual traffic groups.



## 4.0 Results and Conclusions

In this live demonstration, the participating Ethernet Alliance members have shown how SAN and LAN traffic can be aggregated onto a converged 10GbE network. The SAN data is carried via FCoE and iSCSI protocols. The converged Ethernet network provides lossless properties with PFC while the minimum bandwidth for each traffic type is managed through ETS to achieve the optimized latency performance. We observed that the bandwidth provisioning of SAN, LAN and multicast traffic follows the ETS assignment and when a selected traffic type exceeds its allowable bandwidth, the stream is throttling down by PFC pauses.

## 5.0 About the Products Used in this Demonstration

### NetApp's FAS 3140 Unified Storage Array

The NetApp FAS system is a midrange, networked storage system that comes as a single-controller system or in a high availability (HA) configuration with dual, active-active controllers. NetApp FAS systems can be configured with a combination of high performance disk drives and economical, large-capacity disk drives. All FAS systems share a single storage architecture built on the platform of Data ONTAP, one set of management software, and a single family of value-added software products. The FAS3140 used in this instance can be easily upgraded to a model with higher performance and capacity should needs change in the future.

NetApp FAS systems provide simultaneous support for block based FC-SAN and IP-SAN (iSCSI) as well as NFS and CIFS file services via Ethernet and Fibre Channel interfaces. NetApp Fibre Channel over Ethernet (FCoE) is a logical progression of this unified storage approach of offering Fibre Channel, iSCSI, and NAS in our enterprise systems, and provides an evolutionary path for Fibre Channel SAN customers to migrate over time. FCoE will extend Fibre Channel



into the Ethernet environment, combining two leading technologies, the Fibre Channel protocol and Ethernet, to provide more options to end users for SAN connectivity and networking.

### Finisar's Xgig® Monitoring and Analysis Tools

Xgig is the industry's leading multi-function, multi-protocol testing and analysis platform and supports testing of FCoE and DCB protocols with a diversified suite of capabilities:

- The 10GbE/FCoE Analyzer provides unparalleled visibility into network transactions with 100% traffic capture at full line rate
- The 10GbE/FCoE Load Tester creates customizable test traffic streams designed to stress network equipment and verify FCoE functionality and performance
- The 10GbE/FCoE Jammer verifies system robustness by stress testing the performance of priority traffic flows and exercising error recovery capabilities through the manipulation of live traffic flows.

Additionally, the multi-protocol capabilities of the Xgig platform enable developers to evaluate and analyze network traffic as it crosses protocol domains. For example, developers can follow traffic as it crosses the FCoE network and links to Fibre Channel storage equipment.

### Finisar's Medusa Labs Test Tools Suite

The Medusa Labs Test Tools (MLTT) suite is a series of speed benchmark, data integrity, and application-based stress test tools that allow test and interoperability engineers to better develop and validate quality products in real SAN environments. Extensive analysis capabilities and automated operation produce results quickly, accelerating development, troubleshooting, and time-to-market.



## Fulcrum Microsystem's Monaco Reference Design and FM4000 10GbE Switch Chip

Fulcrum's Monaco platform is the first 10GbE OEM switch to offer lossless fabric switching, allowing Ethernet switch manufacturers to deliver Ethernet throughout the data center that can carry storage, clustering and networking traffic and contains all the necessary features to enable Ethernet as the single, converged data center fabric.

Monaco features 24 SFP+ 10 GbE switch ports and can deliver enhanced congestion management performance because it combines the low 300ns L3 latency provided by Fulcrum's FocalPoint silicon, the FM4000, with VantagePoint™, a congestion management technology incorporated into the Monaco platform. The VantagePoint technology enhances the FocalPoint switch features by providing a network-wide view of congestion, contention and traffic loads. It then uses QCN, ETS and FocalPoint's IEEE 802.3x class-based pause congestion management capabilities to dynamically react to this global information.

The FocalPoint FM4224 is a member of the FM4000 series and is layer-three switch/router 10GbE chip is the heart of the Monaco platform. To support comprehensive congestion management, the FM4000 series supports IEEE P802.1Qbb priority flow control (PFC) prevents FCoE packets from being dropped on oversubscribed ports. PFC allows 10 GbE to reliably serve as a viable, low latency fabric technology for converged FCoE and Ethernet environments.

The FM4000 series also supports ETS and is compliant with the most current draft of the IEEE specification. ETS support enables priority groups at each egress port of the switch. For more information on Fulcrum Microsystems, please visit [www.fulcrummicro.com](http://www.fulcrummicro.com).





## 6.0 About the Ethernet Alliance

The Ethernet Alliance is a consortium that includes system and component vendors, industry experts and university and government professionals who are committed to the continued success and expansion of Ethernet technology. The Ethernet Alliance takes IEEE 802 Ethernet standards to market by supporting activities that span from incubation of new Ethernet technologies to interoperability demonstrations, certification and education.

For more information, visit [www.ethernetalliance.org](http://www.ethernetalliance.org).