

Cisco Nexus 5000 and Emulex LP21000

**A Fibre Channel over Ethernet
Technology and Solutions Overview**

Prepared by:
Cisco Systems, Inc. and
Emulex Corporation



Introduction

Data center managers today are being challenged to keep pace with unrelenting growth—trying to deliver the applications necessary to run their businesses, while using a finite amount of resources. More and more frequently, data centers are facing floor space, power and cooling limitations. At the same time, many of the data center resources are significantly underutilized. The combination of these constraints and underutilized resources are a catalyst for increased server virtualization deployments.

The implementation of server virtualization is placing daunting demands on LAN and SAN administrators. Servers need large amounts of LAN connectivity—frequently 6 to 10 gigabit Ethernet connections per server, and these servers and workloads that weren't typically attached to the SAN now require SAN connectivity. To place things in perspective, SANs have typically been designed to service the ten to twenty percent of the most mission critical servers in the enterprise data center. Virtualization is now driving SAN administrators to connect the other ninety percent of servers in the data center not previously SAN attached. The combination of LAN and SAN bandwidth and connectivity growth, leave data center managers with an unmanageable cabling infrastructure.

In response, storage and networking leaders Emulex and Cisco are working together to design and deliver a new networking paradigm based on the creation of Unified Fabric that can convey both LAN and SAN traffic on a single network. This standards-based solution, called Fibre Channel over Ethernet or FCoE, implements the transportation of Fibre Channel SAN traffic over 10 Gigabit Ethernet. This paper provides a technology overview of Fibre Channel over Ethernet, discusses the benefits of solutions being brought to market by Cisco and Emulex, and highlights some of the deployment scenarios for today's data centers.

Addressing Growth Challenges with a Unified Fabric

Many IT organizations today are struggling to contain growth on multiple parallel networks: one for IP networking and one for storage. In the case of High Performance Computing (HPC) environments, a third network exists for clustering or Inter-Process Communication (IPC). These networks cost organizations in terms of the additional capital equipment, the cost and complexity of cabling at the rack level, administrative costs, and the additional power and cooling expense imposed by multiple redundant interfaces and transceivers for each server.

I/O consolidation promises to support all three types of traffic on a single network. One of the primary enablers is 10 Gigabit Ethernet, a technology with the bandwidth and latency characteristics sufficient to support multiple traffic flows on the same link. Recent technology advances with Ethernet in the data center enable the creation of differentiated classes of service, including lossless services over a common data center network. This foundation of Ethernet

support enables the transport of storage traffic using the new Fibre Channel over Ethernet (FCoE) protocol being standardized by the INCITS T11.3 working group. A Unified Fabric using 10 Gigabit Ethernet and Fibre Channel over Ethernet provides several key benefits including:

- The ability to transport SAN and LAN traffic on single infrastructure
- Extending the SAN to additional servers in the enterprise data center without requiring physically extending SAN fabric to servers
- Leveraging existing investments in Fibre Channel including SAN infrastructure, tools, drivers and processes.
- Decreasing cost and power consumption for SAN attachment by reducing Network Interface Card (NIC), cable and switch port requirements.

Technology Overview

Fibre Channel over Ethernet

Fibre Channel over Ethernet transports native Fibre Channel frames over a lossless Ethernet infrastructure, allowing existing Fibre Channel management modes to stay intact. By directly mapping Fibre Channel frames into Ethernet frames, FCoE fully leverages Fibre Channel investments by providing transparent access to Fibre Channel SANs using the same SAN management methodology. FCoE retains the use of proven Fibre Channel drivers on servers and Fibre Channel SAN management tools to minimize operational change.

Lossless Ethernet

One FCoE prerequisite is for the underlying network fabric to be lossless. While FCoE can operate on existing Ethernet networks, several standards-based extensions to Ethernet improve how well the consolidated I/O traffic is handled on the fabric. Priority Flow Control (PFC), sometimes called Per-Priority PAUSE, is a proposal to enable PAUSE capability on the user priorities or classes of service that are defined by the IEEE 802.1p specification. With PFC, administrators can create lossless lanes for Fibre Channel while retaining packet-drop congestion management for IP traffic. IEEE 802.1Qaz is a proposed standard that specifies enhanced transmission selection to allocate bandwidth among different traffic classes. When a given load in a traffic class doesn't fully utilize its allocated bandwidth, enhanced transmission selection allows other traffic classes to use the available bandwidth. This helps accommodate the bursty nature of some traffic classes while maintaining bandwidth guarantees. For these enhancements to work seamlessly with existing Ethernet deployments, a management protocol is being developed that allows switches and endpoints to dynamically discover the capabilities of its peers. Data Center Bridging Capability Exchange Protocol (DCBCXP) provides this needed configuration and discovery functionality.

Cisco and Emulex: Delivering Unified Fabric to the Data Center

Following a long tradition of standards leadership, Cisco and Emulex have been active in the creation of standards for Fibre Channel over Ethernet, and extensions to Ethernet. In addition to collaborating on standards, Cisco and Emulex have been working closely together to develop products that deliver a Unified Fabric for the enterprise data center. The first CNA products supporting the FCoE paradigm are the Cisco Nexus 5000 and Emulex LP21000 family. Both the Cisco Nexus 5000 series of 10 Gigabit Ethernet rack switches and the Emulex LP21000 family of Converged Network Adapters (CNAs) are fully interoperable and designed to be compliant with emerging industry standards. These products work in conjunction with each other to provide server connectivity and data center switching elements of this Unified Fabric.

Emulex LP21000 Converged Network Adapter

The Emulex LightPulse® LP21000 family of CNAs are intelligent multi-protocol adapters that provide host LAN and Fibre Channel SAN connectivity over 10Gb/s Ethernet using Fibre Channel over Ethernet (FCoE) and Enhanced Ethernet functionality. Unrivaled scalability and industry-leading virtualization support make the single-port Emulex LP21000 and dual-port LP21002 CNAs ideal solutions for server I/O consolidation.

The use of the LightPulse architecture ensures driver compatibility with Emulex's entire Fibre Channel HBA product line, providing the broadest platform support in the industry, and seamless integration of mixed Fibre Channel and FCoE deployments. Further, interoperability is guaranteed as the same Emulex drivers have been qualified by every leading server and storage vendor. The LP21000 family supports the latest versions of the Emulex management tools, and supports standards based management such as SMI-S, integrating with enterprise management tools, including the Cisco Fabric Manager.

Emulex and Cisco have performed extensive interoperability testing to ensure that the Cisco Nexus 5000 and Emulex LP21000 deliver a joint solution that:

- Enables IT organizations to begin deployment and protect current investments in equipment and process.
- Allows seamless integration of Cisco Nexus 5000 and Emulex LP21000 Data Center Ethernet capabilities.
- Offers the most comprehensive security, enforcing authentication-protected access with server-to-fabric FC-SP compliant authentication.
- Includes an SFP+ Direct Attached 10Gb/s Copper solution for cost-effective connectivity.
- Provides VM optimized services that enable workloads to migrate seamlessly between dual-fabric and Unified Fabric virtual server farms, and provide statistics such as I/O throughput and error rates to data center managers.

As a Data Center Ethernet (DCE)-enabled CNA, the Emulex LightPulse LP21000 family supports a rich set of Data Center Ethernet capabilities that simplify flow control, reduce network congestion, and coordinate switch-to-switch and switch to host parameter exchange. The LP21000 family appears as an Ethernet NIC and a Fibre Channel HBA to the system, transparently encapsulating Fibre Channel into Ethernet before passing both flows to the network. This enables host I/O consolidation, where a common interface to the network supports both LAN and SAN traffic.

Cisco Nexus 5000

The Cisco Nexus 5000 Series is a family of 10 Gigabit Ethernet rack switches intended for data centers that provides high density, low latency 10 Gigabit Ethernet, Data Center Ethernet, Fibre Channel over Ethernet and virtualization functionality. The Cisco Nexus 5000 Series is ideal for access-layer applications ranging from 10 Gigabit Ethernet to fully Unified Data Center fabric, including LAN and I/O consolidated environments with Fibre Channel over Ethernet.

The Cisco Nexus 5000 Series is available in both 28-port 1U and 56-port 2U form factors. The Cisco Nexus 5000 Series supports standard optical SFP+ transceivers, as well as introducing a new generation of low-latency, low-power transceivers with integrated cabling. The SFP+ Direct Attached 10 Gigabit Copper solution virtually eliminates adoption barriers through its low cost.

Using the same high-availability, feature-rich NX-OS operating system as the rest of the Cisco Nexus Family, the Nexus 5000 Series enables massive scale through Layer 2 Multipathing, a feature that supports large Layer 2 networks without the traditional limitations of spanning-tree protocols. The wire-speed, cut-through forwarding with full throughput and no oversubscription, gives the Cisco Nexus 5000 Series the capacity to alleviate the bandwidth demands created by full racks of high-density, multi-core servers.

The Cisco Nexus 5000 implements a number of standards based extensions to Ethernet to deliver a lossless transport for FCoE. The Cisco Nexus 5000 series also implements a rich set of Data Center Ethernet capabilities that simplify flow control, reduce network congestion, and coordinate switch-to-switch and switch to host parameter exchange, allowing seamless integration with DCE-enabled CNAs. As the first open standards-based access-layer switch to support I/O consolidation, the Nexus 5000 Series eliminates the need for separate switches, cabling, interfaces, and transceivers for storage and networking of traffic. Simpler I/O requirements translate into a need for fewer I/O interfaces (2 vs. the typical 5-7). The Nexus 5000 family preserves existing investments with expansion modules that bridge to native Fibre Channel SANs while preparing data centers for future direct connections to native FCoE-based storage systems.

The Cisco Nexus 5000 Series provides virtual machine (VM) optimized services which allow IT organizations to dynamically respond to changing business demands through rapid provisioning of application and infrastructure services from shared pools of consolidated compute, storage, and network resources. In addition, the Cisco Nexus 5000 series supports end-port virtualization, an approach which simplifies the network, supporting massive scale, high-performance active/active links, and fine-grained control over network resources.

Business Value of the Cisco/Emulex Solution

The Unified Fabric solution provided by the Cisco Nexus 5000 series and Emulex LP21000 family provides five key business benefits to IT organizations:

1. Total Cost of Ownership (TCO) reduction through data center infrastructure

simplification—The Cisco/Emulex solution delivers a Unified Fabric over Ethernet for LAN and Fibre Channel SAN traffic. This provides both lower capital infrastructure expenses and operational savings due to infrastructure simplification. The Cisco/Emulex solution simplifies cable management, allowing hosts to connect to LAN and SAN via the unified Ethernet interface of a converged network adapter, ultimately enabling faster rollout of new applications and services. This Unified Fabric provides consolidation and higher utilization of previously separate resources, reducing by as much as 50 percent the number of network adapters and cables. Such reduction can enable the deployment of smaller server form factors due to the reduction in the number of slots required to support additional NICs in the system. This cohesive infrastructure also lowers power and cooling costs by up to 30 percent. In addition, the LP21000 high performance architecture is designed to offload the processing of both Ethernet and Fibre Channel traffic, freeing up the host processor to focus on application and virtual machine management. CPU efficiency means that smaller servers can be used, further adding to the reduction of power and cooling costs.

2. Data center consolidation with investment protection for existing server, network, storage and facilities assets

—The Cisco/Emulex solution allows you to begin taking advantage of the benefits of a Unified Fabric in a gradual way that allows adoption at the pace your business requires. IT organizations are able to achieve immediate business benefits from a Unified Fabric while leveraging the investment they have made in existing servers, storage, network and facilities. Existing servers and storage devices can continue to operate and support current applications. FCoE deployed between the server and the Nexus 5000 series switch simplifies infrastructure and reduces cost without disrupting the existing storage architecture, network architecture, or management and operational processes. Part of this continuity comes from the Emulex Common Driver model being extended to support the LP21000. This enables the use of the same driver code on Emulex LightPulse FCoE CNAs and Fibre Channel HBAs. In addition management tools such as Cisco Fabric Manager and Emulex's highly scalable and effective tools can be used across both the existing Fibre Channel fabric and new Unified Fabric. This provides a consistent interface and management framework, reducing training costs, and increasing administration staff productivity.

-
- 3. Increased business agility through VM optimized services**—Leveraging existing VM optimized services such as N-Port ID Virtualization (NPIV), the Cisco Nexus 5000 Series in conjunction with the LP21000 CNA delivers Virtual Machine optimized network services, enabling applications to move between physical servers, while maintaining consistent levels of security, Quality of Service (QoS), and performance. This collection of VM optimized networking and storage services enables IT managers to tie VM mobility to policy and event based triggers, such as CPU utilization, thermal load and environmental variables, faults, etc. It supports the Data Center 3.0 vision by allowing IT organizations to dynamically respond to changing business demands through rapid provisioning of application and infrastructure services from shared pools of consolidated compute, storage, and network resources.
 - 4. Enhanced business resilience through higher levels of operational continuity**—With specific features embedded in the Nexus 5000 Series chassis hardware, operating system, and network management, the Cisco Nexus 5000 Series provides component- and system-level operational continuity. Similar features in the Emulex LP21000 provide port isolation between Fibre Channel ports and no reboot Fibre Channel firmware updates. Both the Cisco Nexus 5000 Series and the Emulex LP21000 are designed with a “zero service-loss” architecture, so upgrades can occur without any downtime. In addition, the solution is FC-SP compliant, so that IT departments can protect the SAN infrastructure from unauthorized access, WWN spoofing, host masquerading, rogue/compromised access, non-malicious and administrator caused outages.
 - 5. Ease of deployment and alignment to existing operational models and administrative domains**—The Cisco and Emulex solution enables IT architects to design a data center architecture that aligns to IT organizational structure and operational models. It provides CIOs and senior IT managers with a superior technological solution maximizing business benefits, while minimizing the impact on IT governance models and individual people’s areas of competence/expertise. I/O consolidation at the server access layer can be implemented without touching the existing network or storage environments.

Deployment Scenarios

New Server Deployment at the Edge of Network

By using the Cisco Nexus 5000 series and Emulex LP21000 family, IT organizations can begin to deploy Unified Fabric capabilities without replacing the existing data center fabrics. Figure 1 illustrates a Cisco Nexus 5000 interconnected to existing LAN and SAN equipment in the data center, and a server attached to the Nexus 5000 running FCoE over an Emulex LP21000 CNA. With this integration, the Emulex LP21000 appears transparently to the SAN as a Fibre Channel host. The Cisco Nexus 5000 can interface to the SAN as either an E port or N port, depending on whether it is desirable to have the switch as a logical part of the fabric.

Storage security and privacy regulations are critical factors for IT organizations worldwide and Cisco and Emulex have worked closely to deliver the most comprehensive authentication implementation. As more servers are attached to SANs, especially in a virtualized environment where multiple applications are running on a single physical host, there is heightened demand for greater storage security. Cisco switches and Emulex CNAs and HBAs play a critical role in protecting the right of entry into the storage system by enforcing authentication-protected access with server-to-fabric authentication. This delivers the key benefits of authentication, such as digital identity verification, which ensures back-and-forth password-protected communication between the switch and the server to validate the user's authenticity.




Figure 1 Basic Unified Fabric Integration

Once some basic piloting of a Unified Fabric has been completed, customers can choose to deploy compute resources by “racks” (Data Center in a Rack model) as shown in Figure 2, using a Unified Fabric for these new deployments. This allows a simplification of the cabling and I/O requirements within the rack, yet enabling full connectivity to all Data Center storage and network resources. The rack may then be interconnected to the existing data center 10 Gigabit Ethernet backbone and existing Fibre Channel SAN. This standardizes cabling requirements and simplifies connectivity of servers to the LAN as well as to networked storage devices, provides the ability for hosts to be wired once and ultimately enables faster rollout of new applications and servers. Such a “Wire-once” model provides greater flexibility in re-allocating workloads across the data center computing infrastructure, since common configurations are deployed across the data center.




Figure 2 Unified Fabric Using Data Center in a Rack Model

Virtualized Servers

Virtualized servers are an extension of the Data Center in a Rack model, where the same “Wire-once” model can be used. Cisco and Emulex have teamed to deliver industry-standard NPIV to provide customers with the functionality to maintain storage area network (SAN) best practices, while delivering improved QoS and data protection capabilities, particularly when supporting virtual servers. This solution includes the ability to create a unique identity for each virtual machine that can be tracked throughout the SAN. With this virtual machine identity, the data center manager can track statistics such as I/O throughput and error rates to better manage storage traffic and detect errors at the virtual machine level.

Advanced virtualization features of the Cisco Nexus 5000 series and Emulex LP21000 family ensure that individual workloads are managed consistently in a virtualized environment to how they were managed on physical servers prior to server virtualization deployments. In addition, extensive SAN interoperability testing has been done so that workloads can be migrated seamlessly between existing dual-fabric virtual server farms and new Unified Fabric virtual server farms.

Conclusion

Data Center managers are challenged to keep pace with continuous growth while managing capital, operational and facilities costs. Organizations have turned to server virtualization as a means of reigning in the costs by consolidating applications on virtual servers. Increases in SAN attachment have gone hand in hand. FCoE offers the framework for implementing a Unified Fabric as a way to get beyond today's cost and complexity challenges.

Cisco and Emulex have worked collaboratively on both standards and products that will deliver on the promise of a Unified Fabric reducing the total cost of ownership; offering investment protection of existing servers, networks, storage, and facilities; enhancing business agility; and enabling higher levels of operational continuity; while offering alignment with existing operational domains.

For more information

Please contact Emulex and Cisco if you are interested in learning more about FCoE based solutions.

- ▶ www.cisco.com
- ▶ www.emulex.com



Cisco Systems, Inc. 170 West Tasman Drive, San Jose, CA 95134-1706 USA Tel 408-526-4000 www.cisco.com

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.

Emulex Corporation, 3333 Susan Street, Costa Mesa, CA 92626 USA Tel 714-662-5600 www.emulex.com

Copyright © 2007 Emulex. All rights reserved worldwide. No part of this document may be reproduced by any means or translated to any electronic medium without the prior written consent of Emulex.

Information furnished by Emulex is believed to be accurate and reliable. However, no responsibility is assumed by Emulex for its use, or for any infringements of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright or related rights of Emulex. Emulex, the Emulex logo, LightPulse and SLI are trademarks of Emulex.