Convergenomics™

The Guide to Network Convergence Solutions

2nd Edition

Simplify networking, reduce costs, and increase business agility with data center network convergence.

EMULEX®

Foreword

Three key factors drive data centers today. First, operational efficiency. In these tough economic times, everyone is looking at ways to lower costs. I don't know a CIO who is not being pressured to lower capital expenses and create more efficiency.

The second factor is business responsiveness. It is all about how the IT organization can help the business unit achieve its goals and drive greater levels of performance. IT managers see virtualization as a technology that allows them to drive more efficiency into their infrastructure. Virtualization has sped up the provisioning of new server resources to respond quickly to the business unit's computing needs.

Finally, companies are looking to IT to help them achieve a competitive advantage. Whether it is a drug company trying to speed discovery by finding the next cure, a financial company looking at derivative calculations or a retail company needing better business intelligence on customer needs, IT and the services it provides can play a critical role in producing a sustained competitive advantage.

Operational efficiency, business responsiveness and competitive advantage are vital for businesses moving forward in today's world. But what are some of the things that have prevented the IT vision from being realized? We certainly know that OPEX is a challenge. How do we spend less while managing an ever-increasing infrastructure? To complicate things, there are also environmental concerns, such as heating and cooling, power density and space. Even weight of the infrastructure is a concern! I recently visited a company in London and learned that the weight of the infrastructure in the buildings has become a real challenge. Then there is the big elephant in the room: increasing demand for computing resources. Today's data center needs to address security, storage, availability and speed, while keeping costs and environmental concerns under control.

These are the elements that drive consolidation. We have seen companies move from tower servers to racks, and from racks to blades. With virtualization, our customers are able to achieve higher levels of utilization, but this is also creating the need for higher bandwidth.

More and more customers want a networked infrastructure for this reason. At Emulex, we recognize the need for increased storage connectivity. If a data center manager wants to move resources from machine to machine through virtualization, it is very important that the machines are network-attached. All of these arrows point to network connectivity.

Emulex has been in networking for 30 years and has helped to create the foundation for the storage area network (SAN). Over the course of the last decade, we saw local area networks (LANs) and SANs established with Fibre Channel. In the last year, we began focusing on the next great transition to Fibre Channel over Ethernet (FCoE) and Converged Network Adapters (CNAs). With these technologies, we are bringing data and storage networks together for network convergence. Looking ahead, we really see an opportunity for converged networking to go "universal," where the new technology will allow companies to wire once and be able to deliver new capabilities, regardless of the network they want to deploy.

Convergence is game-changing and creates a new economic model for the data center. We call this new economic model Convergenomics™. This new model delivers on operational efficiency by reducing a number of networks down to a single network, and by providing cost savings on adapters, cabling, power, space and people, without throwing away the investments in existing infrastructure, processes, software tools and knowledge. Convergenomics allows for increased business responsiveness by providing the ultimate network for virtualization that will deliver on the promise of dynamic provisioning. With exceptional bandwidth, IT organizations will have a dependable, high-performance network to build competitive solutions for today's round-the-clock, video-rich, multimedia-centric, global data center.

Emulex is known throughout the industry as a company that not only innovates, but helps drive new standards. Driving standards is critical in networking to create the interoperability that data centers count on for open plug-and-play solutions. This means that we not only work closely with ecosystem partners on defining, building and testing our solutions, but it also means that we work with our customers on education for the new standards. We listen to their concerns and needs, depending on them to drive our innovation.

This is the reason we created this book. *Convergenomics – The Guide to Network Convergence Solutions* provides a comprehensive introduction to network convergence based on FCoE. It includes an assessment of the challenges facing the data center and the business case for FCoE, including use cases and applications that will likely be addressed with the first wave of FCoE deployments. Additionally, Convergenomics describes the technology evolution and standards activities leading to the FCoE paradigm. It explores networking building blocks,

such as software, storage, switches and management tools, and features contributions from key infrastructure ecosystem partners.

As the COO at Emulex, I have been impressed by the creativity, intelligence and drive of the Emulex team. I want to personally thank them for their determination and effort in helping to create this new technology that will completely revolutionize the data center as we know it.

I also want to personally thank our partners, and in particular, Bert Mauricio from Brocade Communications Systems, Graham Smith from BLADE Network Technologies, Scott Gainey from Cisco Systems, Deirdre Wassell from EMC Corporation, Gary Lee from Fulcrum Microsystems, David Colodny from Juniper Networks, Jo De Baer and Robert Wipfel from Novell, Inc., Monica Kumar and Michelle Resta from Oracle Corporation, Mark Hwang from Panduit Corporation, Chandy Nilakantan from Scalent Systems and Paul Manning from VMware, Inc. Your support and collaboration are critical to the success of this technology.

Now is the time to start planning for your future. Convergence is the innovation that will help you achieve operational efficiency, business responsiveness and competitive advantage as we move into the next decade. I hope that you enjoy reading *Convergenomics – The Guide to Network Convergence Solutions* and invite you to visit our web site, attend one of our webcasts or meet with one of our sales teams to continue that education.

Jeff Benck

Chief Operating Officer

Emulex Corporation

Contents

Chapter 6

Chapter 7

Chapter 8

Chapter 9

APPENDIX

1 Data Center Challenges

Data center challenges have always been about improving the efficiency of operations. Whether it's servers, storage or the network, the challenges have been broadly focused on improving the performance and manageability of operations in a highly cost-efficient manner.

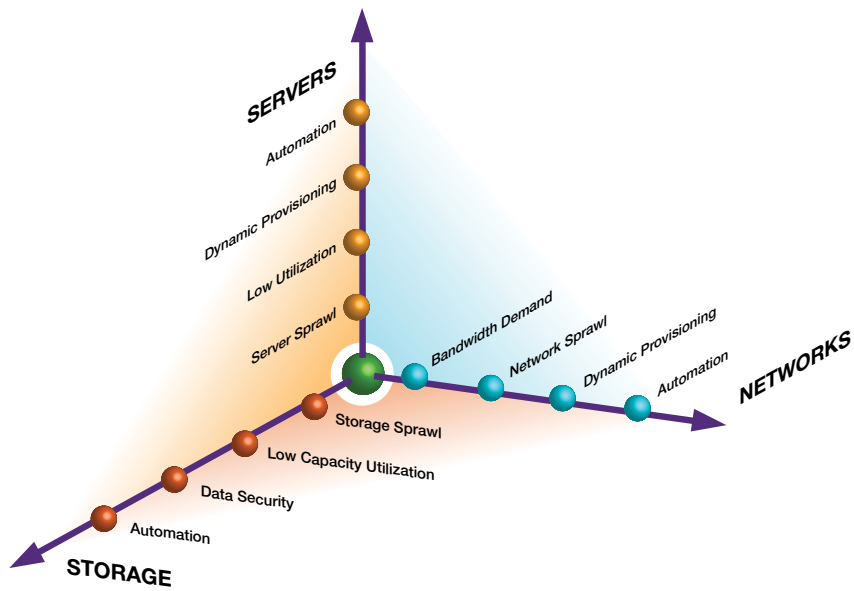

Figure 1: Data center challenges across IT segments

Data center managers have made incremental steps in improving the efficiency of operations by deploying:

- Centralized storage to effectively manage data access and backups
- Storage virtualization to improve utilization and availability
- Blade servers to contain physical server sprawl
- Server virtualization to fight low utilization, improve application availability, decrease the cost of power and cooling and minimize floor space

Improving efficiencies in one segment of the data center typically exposes inefficiencies in other segments of the infrastructure. For instance, deployment of server virtualization has increased network complexity by driving demand for multiple 1Gb Ethernet links and multiple Fibre Channel SAN connections per server. The deployment of blade servers for consolidation and flexibility has increased the need for higher networking bandwidth and increased external storage connectivity. These requirements add to complexity networking complexity and increases the power, cooling and rack space required for the network infrastructure.

The Data Center Networking Challenge

The combination of high bandwidth demand, increasing network sprawl and the need for more adaptive networking infrastructure is posing a major challenge for data center managers. Pain points in today's data center networks include:

- Multiple network fabrics, each dedicated to specific type of traffic
- High numbers of adapters and switch port deployments
- Complex cabling infrastructure
- High network provisioning times as a result of static configurations
- Complexity of managing switch and adapter firmware and associated service contracts

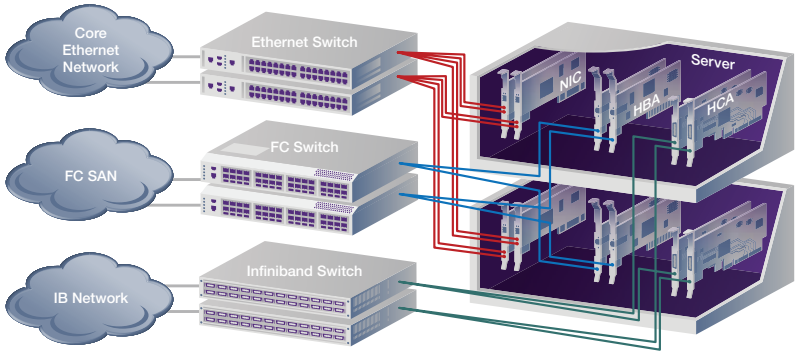

Figure 2: Dedicated networks for SAN, LAN and clustering

Data center managers are looking for solutions to transition to a more dynamically provisioned network that is highly responsive and addresses the quality and service level requirements of business applications.

Requirements for Next-Generation Data Center Networks

Data center managers are clearly in need of networking solutions that contain the sprawl of network infrastructure and enable an adaptive next-generation network. The solution for optimizing the data center network must be capable of addressing the following high-level requirements:

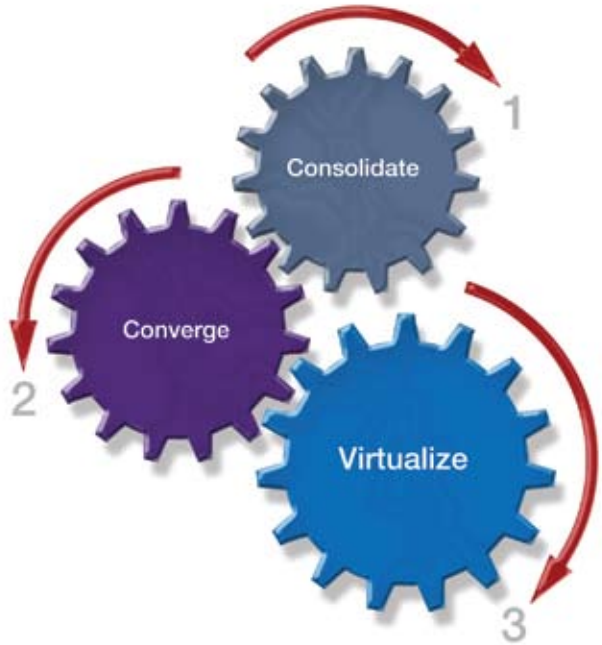

Figure 3: High-level requirements for next-generation data center network

1. **Consolidate:** The network solution must be capable of consolidating multiple low-bandwidth links into a faster high-bandwidth infrastructure and significantly reduce the number of switch and adapter ports and cables.

2. **Converge:** The network solution must be capable of converging or unifying networking, storage and clustering traffic to a single network, eliminating the need for dedicated networks for each traffic type. This functionality will further contribute towards reduction in network ports and cables, and simplify deployment and management.

3. **Virtualize:** The network solution must be capable of virtualizing the underlying physical network infrastructure and providing service level guarantees for each type of traffic. In addition, the solution must be capable of responding to dynamic changes in network services depending on the business demands of the data center applications.

10 Gigabit Ethernet: The Foundation for Network Convergence

The 10 Gigabit Ethernet networking standard (10GbE), ratified in 2002, initially found traction in specific applications such as inter-switch connectivity and high-performance applications. The Ethernet networking technology has since witnessed several key innovations:

- Enhanced physical media
  - 10Gb/s connectivity over UTP cabling
  - 10Gb/s connectivity over Direct Attach Twin-ax Copper cabling
- Optimizations in 10Gb/s transceiver technology (SFP+ form factor)
- Support for lossless Ethernet infrastructure
- New physical network designs such as top-of-rack switch architectures

These innovations are helping 10GbE networks address the requirements of consolidation, convergence, and virtualization and are paving the way for the next generation of data center networks.

A converged network based on 10GbE fully complements the data center consolidation efforts and improves the efficiency of overall operations. Leveraging 10GbE to carry data networking, storage and clustering traffic simplifies network infrastructure and reduces the number of cables, switch ports and adapters and lowers overall power, cooling and space requirements.

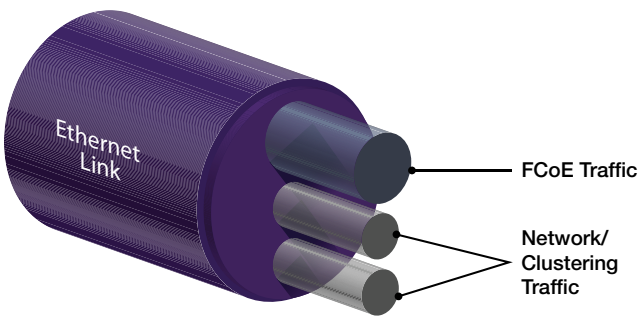

Figure 4: 10GbE enables multiple traffic type over single link

In addition to providing lowered costs, 10GbE enables much-needed scalability by providing the additional network bandwidth. 10GbE also simplifies management by reducing the number of ports and by facilitating flexible bandwidth assignments for individual traffic types.

2 Network Convergence Evolution

Ethernet is the predominant networking technology in the data center. Historically, Ethernet has outlasted many competing technologies because of its ubiquitous nature, ease-of-deployment, management tool prevalence and the eventual cost-reduction driven by high-volume classic deployments. The promise of a consolidated network has appealed to data center managers and multiple attempts have been made to achieve network convergence over a common infrastructure.

In addition to network convergence attempts over Ethernet, there were alternative convergence solutions proposed with InfiniBand as the underlying fabric. InfiniBand is a high-speed low-latency network technology that addressed the requirements of inter-process communication in high-performance computing environments (HPC). While convergence solutions over InfiniBand were creative, many perceived the solutions to be too disruptive to the existing networks and administration. As a result, the technology option never made an impact beyond HPC-environments.

Network Convergence with iSCSI and 10 Gigabit Ethernet

The iSCSI protocol ratified by Internet Engineering Task Force (IETF) in 2003, brought storage area networks within the reach of small and mid-sized businesses. The protocol encapsulates native SCSI commands using TCP/IP and transmits the packets over the Ethernet network infrastructure. The emergence of 10GbE addressed the IT manager's concerns regarding the bandwidth and latency issues of 1 Gigabit Ethernet and laid the foundation for more widespread adoption of network convergence in data centers. The 10GbE technology found immediate acceptance in HPC and soon storage arrays supported 10GbE interfaces.

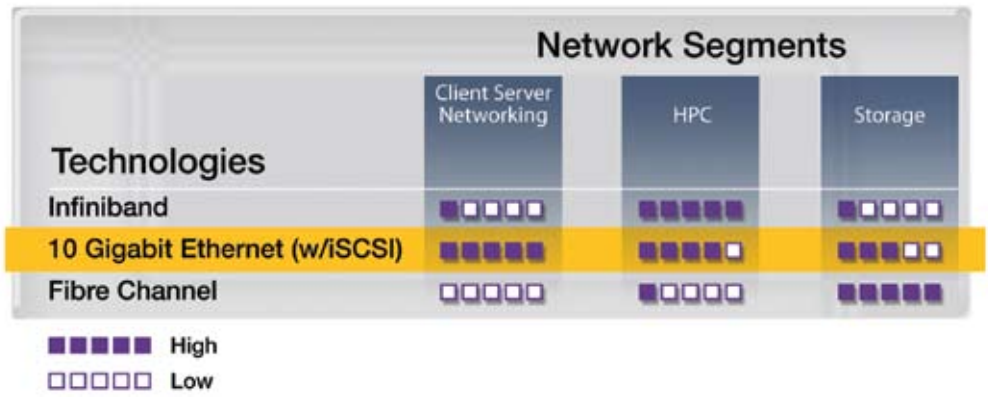

Figure 5: Technology applicability across network segments

Pros and Cons of iSCSI-enabled Network Convergence

iSCSI-enabled convergence offers several advantages:

- Highly suitable for convergence in small and medium businesses, remote offices and department-level data centers where customers are transitioning from Direct Attach Storage to SANs.
- Reduces labor and management costs while increasing reach.
- The ubiquitous nature of Ethernet means that IP networks can be deployed quickly and easily in organizations of all sizes. Ethernet is also readily understood, so IT personnel can deploy and maintain an IP environment without specialized Fibre Channel training.
- Major operating systems include an iSCSI driver in their distribution. iSCSI performance can be improved by deploying adapters which support iSCSI offload or TCP/IP offload to reduce the CPU demands for packet processing.

Although optimal for small and medium businesses, iSCSI-enabled convergence does have limitations:

- Because the underlying Ethernet network is prone to packet losses with network congestion, network designers typically recommend the use of separate Ethernet networks for storage and for data networking. This reduces some of the cost advantages of convergence.
- Large enterprise data centers have a sizable deployment of Fibre Channel SANs, managed with specific Fibre Channel tools to effectively manage storage assets. From the perspective of these customers, iSCSI is a different storage technology with different SCSI encapsulation that requires an incremental, non-trivial investment in hardware, software and training.

Enterprise data center managers have typically been conservative in adopting new technologies. Enterprise customers usually prefer that any new solution offers basic functionality such as performance and RAS (reliability, availability, serviceability), and also be robust enough to meet their enterprise application environment's demands. In addition, these customers wanted to protect their existing investments in storage networks, such as Fibre Channel SAN infrastructure (HBAs, switches, storage) and Fibre Channel SAN management tools and processes. While carrying data networking traffic with Fibre Channel traffic over 10GbE was a possible option, the unreliable nature of the underlying Ethernet was considered to be unsuitable for loss-sensitive Fibre Channel traffic.

Network Convergence with Lossless 10 Gigabit Ethernet

In order to facilitate network convergence and carry Fibre Channel traffic over 10GbE, Ethernet technology had to support "no-drop" or "lossless" behavior. Since its introduction, Ethernet has steadily evolved to keep pace with these changing requirements. In addition to higher transmission rates, Ethernet has advanced from the days of shared media to dedicated media, and from half duplex to full duplex transmissions.

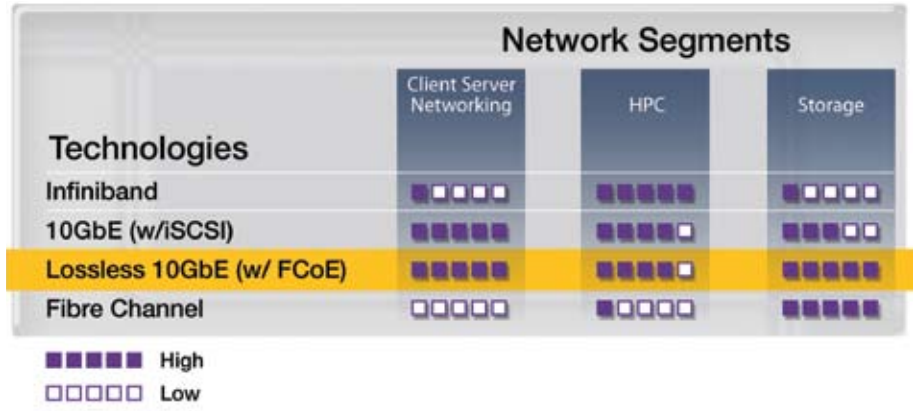

Figure 6: Technology applicability across network segments

The recent innovations that support lossless characteristic in 10GbE are the ability to:

- Isolate and prioritize different traffic types using Priority Flow Control (PFC)
- Maintain bandwidth guarantees for multiple traffic types
- Assure that end-points and switches to know about each other's capabilities through an enhanced management protocol

In parallel with the emergence of lossless 10GbE, the emergence of newer standards such as Fibre Channel over Ethernet (FCoE) standard is accelerating the adoption of Ethernet as the medium of network convergence.

Ethernet Enhancements Overview

Efforts are well underway with the Data Center Bridging Task Group of IEEE 802.1 Working Group (LANs) to provide the necessary framework for enabling 10GbE converged networking within a data center. Four key protocols are being defined or enhanced as shown in table below:

Protocols	Key Functionality	Business Value
Priority Flow Control (PFC) P802.1Qbb	Management of bursty, single traffic source on a multi-protocol link	Enables storage traffic over 10GbE link with "no-drop" in the network
Enhanced Transmission Selection (ETS) P802.1Qaz	Bandwidth management between traffic types for multi-protocol links	Enables bandwidth assignments per traffic type. Bandwidth is configurable on-demand.
Data Center Bridging Capabilities Exchange Protocol (DCBCXP) 802.1Qaz	Auto exchange of Ethernet parameters between peers (switch to NIC, switch to switch)	Facilitates interoperability by exchanging capabilities supported across the nodes.
Congestion Management (CM) P802.1Qau	Addresses problem of sustained congestion, driving corrective action to the edge	Facilitates larger end-to-end deployment of network convergence.

Table 1: Protocol standards are enabling convergence

Within the next couple of years, data center managers will be hard-pressed to find 10GbE vendors who do not support lossless Ethernet. What is now termed "lossless Ethernet" will soon become "mainstream Ethernet".

3 FCoE-enabled Enterprise Network Convergence

Fibre Channel over Ethernet (FCoE) is a standard being developed by INCITS T11 that will fully leverage the enhanced features of 10GbE for I/O consolidation in the data center. FCoE expands Fibre Channel into the Ethernet environment, combining two leading technologies, Fibre Channel and Ethernet, to provide more options to end users for SAN connectivity and networking.

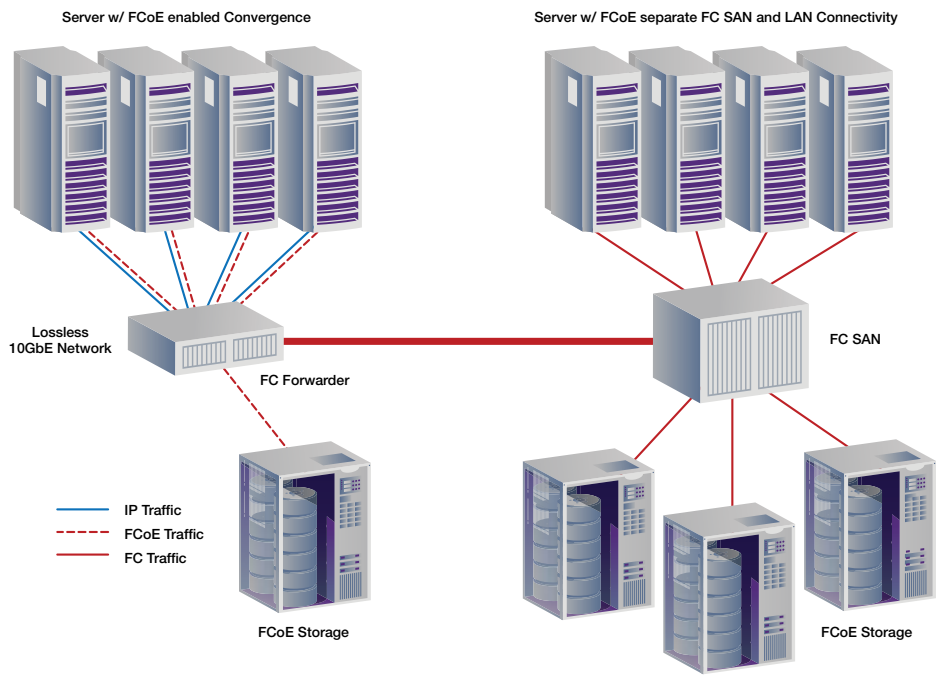

Figure 7: FCoE-enabled network convergence

The primary objectives of FCoE are to:

- Allow an evolutionary approach to convergence while providing significant cost savings benefits through I/O consolidation
- Preserve the primary characteristics of Fibre Channel
  - High performance, low latency
  - Robust security
  - NPIV support server virtualization

- Protect existing investments in Fibre Channel SANs by ensuring interoperability
- Preserves investments in Fibre Channels tools, processes and training

Evolutionary Approach

FCoE is not a replacement for conventional Fibre Channel but is an extension of Fibre Channel over a different link layer. Enabling lossless Ethernet infrastructure to carry both Fibre Channel storage data as well as other data types allows customers to simplify server connectivity and still retain the performance and reliability required for storage transactions. Instead of provisioning a server with dual-redundant Ethernet and Fibre Channel ports (a total of 4 ports), blade servers can be configured with two lossless 10GbE ports. For blade server installations in particular, this reduction in the number of interfaces greatly simplifies deployment and ongoing management of cabling. The main value proposition of FCoE is, therefore, the ability to streamline server connectivity using lossless Ethernet while retaining the channel characteristics of conventional Fibre Channel SANs.

FCoE will enhance investments in FC-based architectures and expand Fibre Channel into Ethernet environments by combining the two leading technologies to provide more options to end users for SAN connectivity and networking.

Fibre Channel Characteristics Preserved

The FCoE protocol specification maps a complete Fibre Channel frame (including checksum, framing bits) directly on to the Ethernet payload and avoids the overhead of any intermediate protocols.

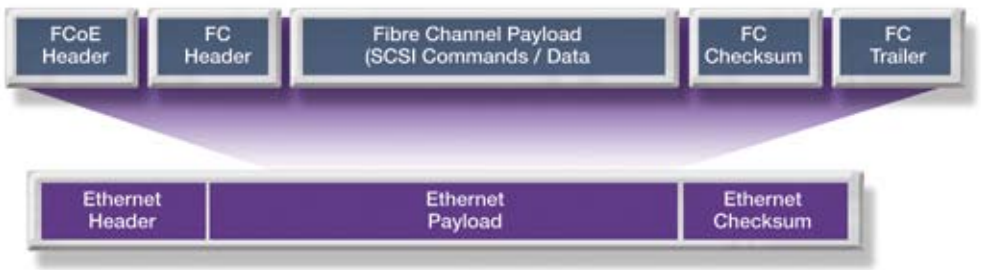

Figure 8: FCoE encapsulation in Ethernet

This light weight encapsulation ensures that FCoE-capable Ethernet switches are less compute-intensive, thus providing the high performance and low latencies of a typical Fibre Channel network. By retaining Fibre Channel as the upper layer protocol, the technology fully leverages existing Fibre Channel constructs such as fabric login, zoning and LUN masking, and ensures secure access to the networked storage.

N_Port ID Virtualization

N_Port ID Virtualization (NPIV) is an innovation in Fibre Channel that enables addresses (also referred to as multiple virtual ports) to share a single Fibre Channel port when registering with the SAN Fabric. This capability enables virtual machines to have their own dedicated virtual ports and thus limits the storage access to only the required resources. Further, the ability to reinitiate a virtual port on a different server greatly enhances virtual machine mobility for load balancing, portability and disaster recovery. FCoE retains the use of NPIV to improve the flexibility and security of virtual server deployments.

Fibre Channel SAN Investments Protected

In Fibre Channel SAN environments, isolation of storage devices is provided through zoning and LUN masking techniques based on the address (World Wide Port Name) of the Fibre Channel port. In a server virtualization environment, when multiple virtual machines share a single Fibre Channel port, the virtual machines are indistinguishable to the SAN fabric for both data isolation and Quality of Service (QoS) management.

One of the goals of FCoE, therefore, is to ensure protection of Fibre Channel investments through interoperability with existing infrastructure. The use of Fibre Channel as the upper layer protocol in FCoE ensures that the Fibre Channel SAN infrastructure is seamlessly accessed from FCoE-enabled servers. This commonality ensures that FCoE-enabled servers use field-proven Fibre Channel software and can transparently interoperate with existing Fibre Channel SAN infrastructure that includes switches, storage arrays, tape drives and data backup tools. This approach also acts as an efficient mechanism for extending Fibre Channel SAN infrastructure to the Ethernet networks and thus expands the accessibility and usage of the existing storage investments.

SAN Management Processes and Training Preserved

Fibre Channel SANs offer a single point of management for the FC network enabling more storage to be managed with fewer personnel. Each device within the SAN has a name and address that are used for management. In large FC SANs, zones may be created using those names or addresses to restrict access between particular servers and devices for improved security.

FCoE uses FC names and addresses to ensure full compatibility and seamless integration of existing FC management methods to detect failures, track performance, and determine when changes occur in the SAN. This commonality with Fibre Channel ensures that investments in management tools, storage management processes and the associated training for storage professionals are fully leveraged.

4 Converged Network Building Blocks

An end-to-end converged network deployment enables consolidation in the server access network, the core network and in storage device connectivity. The end-to end deployment of a converged network is comprised of the following components: Converged network adapter, FCoE-capable lossless Ethernet switches native FCoE devices.

Although the deployment of an end-to-end converged network may be staged and rolled out in phases as detailed in Chapter 9, the individual components for the solution remain the same. A brief overview of the components is provided here:

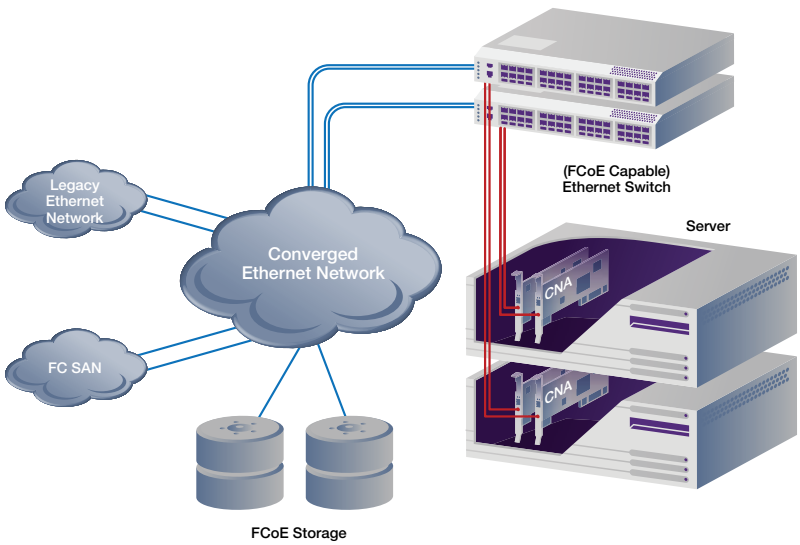

Figure 9: Converged Network Building Blocks

	Key Building Blocks
1	Converged Network Adapters (CNAs)
2	FCoE-capable lossless 10GbE switches
3	Native FCoE storage
4	Hypervisors and operating systems
5	Management tools

Table 2: Building blocks for deploying convergence

Converged Network Adapter

A new family of adapters called Converged Network Adapters (CNAs) leverages Ethernet enhancements and FCoE technology to drive convergence over the underlying lossless 10GbE infrastructure. CNAs combine the Network Interface Card (NIC) functionality and the Fibre Channel HBA functionality into a single adapter. CNAs encapsulates complete Fibre Channel frames within Ethernet frames and enable Fibre Channel traffic to be carried over Ethernet infrastructure.

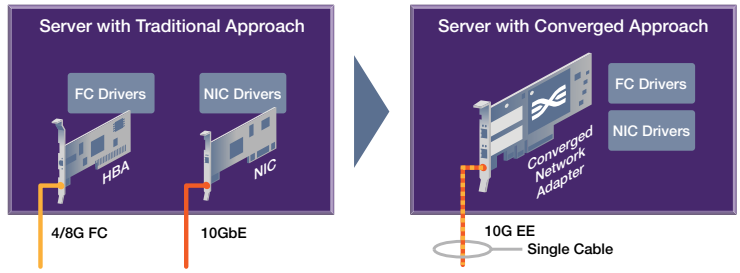

Figure 10: FCoE Converged Network Adapters are provisioned with both Fibre Channel driver and NIC drivers.

The CNA functionality is nearly transparent to the operating system since it presents both the NIC and the HBA functionality.

Dual port 10GbE adapter

Dual Port 4 Gb/s Fibre Channel HBAs

Figure 11: CNA is presented as NIC and HBA to the OS

CNAs offload Fibre Channel protocol processing from the CPU, thus providing high performance storage connectivity while improving overall CPU utilization. The additional processing power retained by the CPUs provide more CPU bandwidth for enterprise applications and in the case of server virtualization environment, enables deployment of more virtual machines per server.

Figure 12: First and second generation CNAs

Servers supporting PCI Express (x8 or higher) are capable of utilizing the full potential of dual-port 10Gb/s interfaces of the converged network adapter. By providing the capability to isolate and prioritize different traffic types on the link and by supporting bandwidth guarantees for different traffic types, CNAs enable efficient consolidation while ensuring that storage traffic is delivered with low latency and without any degradation in performance.

CNA Performance

First-generation CNAs are capable of supporting up to 4Gb/s bandwidth for Fibre Channel traffic. Given this bandwidth allocation, the FCoE CNAs are capable of supporting I/O performance that is on par with standalone 4Gb/s Fibre Channel HBAs. In laboratory tests, using the Cisco Nexus 5000 switch provisioned with a Fibre Channel module, the single port Emulex LP21000 CNA supported data throughput of 780MB/s (full duplex) and 150K IOPS, which is equivalent to the performance of the Emulex LPe11000 4Gb/s Fibre Channel HBA.

CNA vs. FCoE Software Initiators

As in the case of iSCSI software initiators, efforts are under way to develop FCoE software initiators that can work with a standard Ethernet NIC that does not support offloads. However, using the CPU resources to process Fibre Channel traffic has several disadvantages when compared with a CNA that uses a dedicated hardware engine for offloading Fibre Channel protocol processing.

Feature	*CNA*	*FCoE Software*
High Performance/Watt	✓	⊘
FCoE Offload Enables Low CPU Utilization	✓	⊘
Broad Storage Interoperability Support	✓	⊘
Unified Storage Management for CNAs and HBAs	✓	⊘

Table 3: CNAs have inherent advantages over software initiators

With server virtualization being one of the key drivers of network convergence, the use of CNAs allows the virtualization server to accommodate significantly more virtual machines per server than the FCoE software initiator option. The availability of an industry-proven Fibre Channel protocol engine in hardware also ensures that CNAs are well-positioned for interoperability with existing network storage systems and Fibre Channel networking equipment. The ecosystem support for FCoE software initiators is still nascent at best and leaves several issues unanswered for reliable deployment in product data centers.

FCoE-capable Lossless 10 Gigabit Ethernet Switch

The FCoE-capable, lossless 10 Gigabit Ethernet switches are a key component for deploying a converged network in the data center. FCoE-capable switches support the basic functionality required for lossless Ethernet implementation and provide added functionality to support processing of FCoE traffic. Besides the Ethernet enhancements and support for FCoE, the switches are typically configured with optional Fibre Channel interfaces for connectivity to existing Fibre Channel SANs. Since FCoE-capable Ethernet switches provide all the services of a Fibre Channel switch, they provides the flexibility to either be connected to an FC storage array directly or to another FC switch/director which in turn is connected to the rest of the SAN.

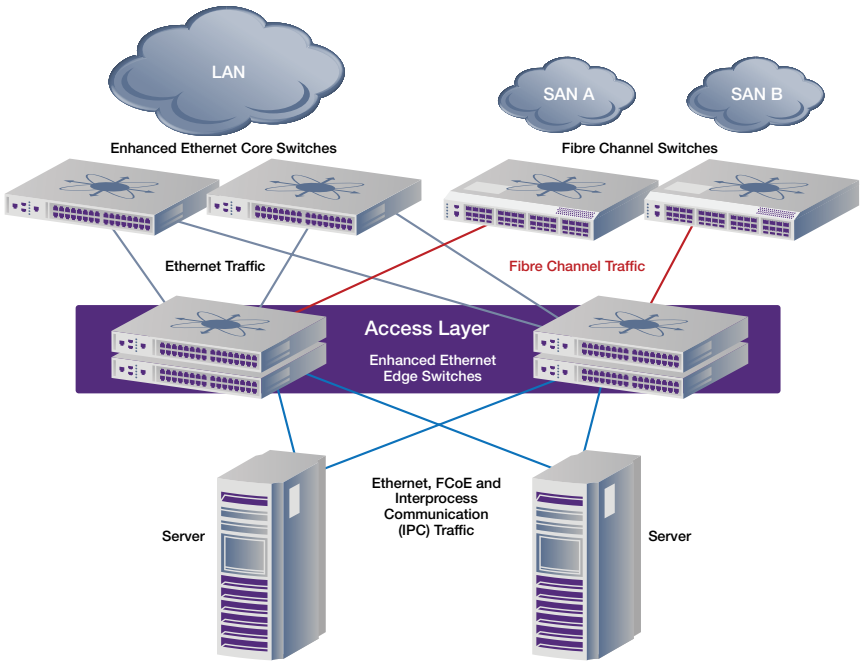

Figure 13: Role of FCoE switch in network convergence

Zoning practices used in Fibre Channel networking typically remain unaffected, and processes are transparently carried over to the FCoE-capable lossless Ethernet switch.

FCoE consolidation can initially be deployed at the edges where converged network adapters and FCoE capable lossless Ethernet switches will take advantage of the lossless characteristics of the edge fabric and standardization of the congestion management framework. The core and edge switches developed today are architected to enable a firmware upgrade to accommodate changes during the standardization process.

Native FCoE Storage

Storage arrays supporting native FCoE interfaces will enable end-to-end network convergence and are expected to be the next logical progression in the converged network environment. Besides the change in physical layer connectivity that encapsulates Fibre Channel frames over Ethernet, the functionality provided by native FCoE arrays remains equivalent to that of a Fibre Channel array. The native FCoE arrays will leverage the proven performance of Fibre Channel stack and retain the existing processes required for LUN masking and storage backup.

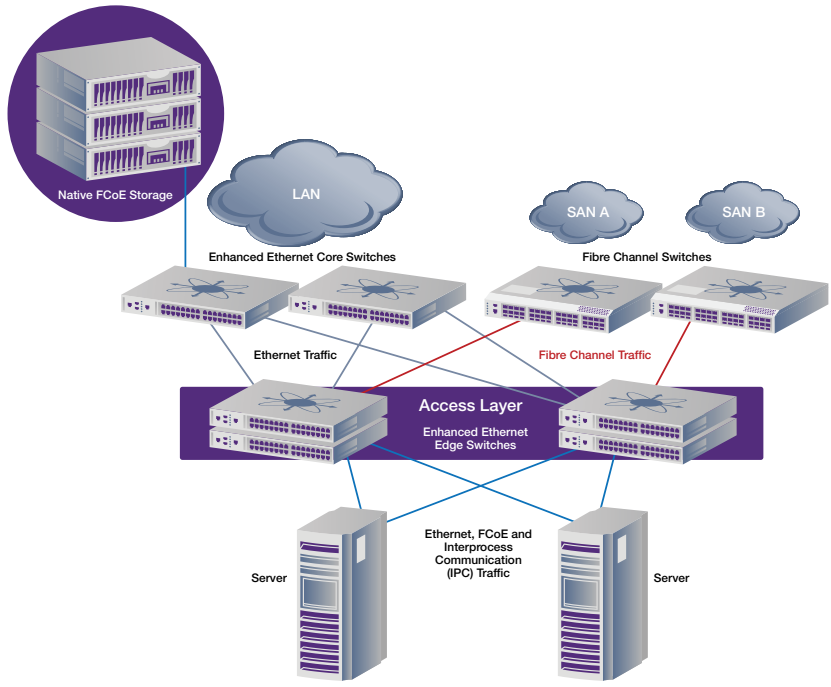

Figure 14: A native FCoE storage connected to FCoE-enabled network

As more customers implement virtual server environments, FCoE will deliver increased efficiency and cost savings, while protecting existing SAN investments.

Hypervisors and Operating Systems

LUN masking practices used by the storage administrators in Fibre Channel storage remain unaffected and processes are transparently carried over to native FCoE storage.

Host software provides the foundation on which converged services are utilized by the applications. Hypervisors and operating systems that support an underlying converged network are essential for a reliable and scalable deployment.

Management Tools

Efficient management of the data center infrastructure is a key objective for IT managers. As in the case of Fibre Channel, the management tools for a FCoE-enabled converged network must be capable of enabling IT managers to deploy and maintain highly reliable and secure infrastructures for the most critical enterprise data.

5 The Business Case for FCoE-enabled Network Convergence

Data center challenges are growing by the day. Businesses are required to maintain an IT infrastructure that is flexible and adaptive enough to meet changing customer demands and opportunities. Server consolidation initiatives such as blade servers and server virtualization have increased the demand for high bandwidth networks that can be provisioned dynamically. Pooling of multiple networks into a common infrastructure that can be shared by multiple traffic types can help achieve these objectives.

Business Benefits

A converged network environment adds flexibility to the data center network while providing the following benefits:

- Lower total cost-of-ownership through infrastructure simplification
- Increased IT efficiency and business agility
- Protection for existing investments in Fibre Channel infrastructure
- Seamless extension of Fibre Channel SANs in the data center
- Uses existing management organization

Table 4: List of key benefits to be gained by deploying convergence

Lower Total Cost-of-Ownership through Infrastructure Simplification

FCoE-enabled network convergence to a single lossless 10GbE infrastructure reduces number of cables, switch ports and adapters required to maintain both SAN and LAN connectivity, significantly lowering the capital expense incurred for the data center network.

The reduction in the number of adapters facilitates the use of smaller servers expanded and adapter offload enables higher CPU efficiency which directly affects power and cooling costs. A smaller number of ports ultimately reduces administration and management costs.

Increased IT efficiency and Business Agility

A converged network fully complements data center virtualization initiatives. IT managers will be able to adjust network bandwidth, to meet shifts in workloads. This capability helps streamline server and network provisioning with a "wire once" deployment model which enables IT managers to rapidly respond to requests for new or expanded services, new servers and new configurations.

Protection of Existing Investments in Fibre Channel Infrastructure

FCoE lets organizations phase in the converged networks roll-out. FCoE–based network convergence enables consolidation in new deployments without affecting the existing server and storage infrastructure or the processes required to manage and support existing applications.

Seamless Extension of Fibre Channel SANs in the Data Center

Fibre Channel is the predominant storage protocol deployed by enterprise data centers and the adoption of blade servers and server virtualization adds to the increased demand for Fibre Channel SANs. FCoE addresses this requirement by extending proven Fibre Channel SAN benefits a 10GbE infrastructure. The use of light weight encapsulation for FCoE also facilitates FCoE gateways that are less compute-intensive, which provides higher levels of storage networking performance.

Uses Existing Management Organization

FCoE-based converged networks gives IT architects the ability to design a data center architecture that aligns with the existing management organization for servers, networks and storage. This minimizes the need to modify operating procedures used by storage and networking IT staff.

Cost Analysis of Deploying Network Convergence

Converging the data center networks on to a single infrastructure based on 10GbE provides cost savings on capital and operating expenses. The following provides more details on how network convergence drives cost savings.

Capital Expense Savings

Examples of potential savings in capital costs are listed in the table below:

	Capital Expense Savings
1	Fewer adapters
2	Fewer switches
3	Fewer cables
4	Smaller servers with fewer PCIe slots
5	Less rack space

Table 5: Reduced capital expenses with convergence

Operating Expense Savings

Savings in operational expenses are achieved by reduced expenses for management, provisioning, power and cooling costs.

	Operating Expense Savings
1	Fewer ports reduces management time
2	Dynamics provisioning lowers provisioning time
3	Reduced power and cooling
4	Unified management across CNAs and HBAs improves efficiency and minimizes training
5	No changes to existing administration processes

Table 6: Reduced operating expenses with convergence

Cost Analysis with Sample Scenario

To help quantify the savings driven by network convergence, it helps to compare the costs of a traditional networking (LAN + SAN) environment to a converged networking environment that meets the same business requirements.

Let us consider the scenario where a data center hosts 160 servers with multiple LAN connections and Fibre Channel connections per server (Table 7). When these servers are migrated to a converged solution, the consolidation in the server access network alone can provide significant savings as listed in Table 8.

Data Center Profile	Traditional Networking	Converged Networking
Servers	160	160
LAN Connections (1GE) per Server	6	---
SAN Connections per Server	2	---
Converged Ethernet Connections per Server	---	2

Table 7: Data center profile that is analyzed for cost savings

Cost Components	Traditional	Converged	Savings %
Switches and Adapters	$925,600	$734,080	21%
Power	$23,213	$13,422	42%
Cable Installation	$270,400	$41,600	85%
Number of Inter-rack Cables	1,352	448	67%

Table 8: Estimated cost savings with converged networks

For a data center with 160 servers requiring networking and storage connectivity, converged networking provides a 21% savings in switch and adapter costs and a 85% savings in cable installation costs.

Figure 15: Estimation of cost savings with Emulex Convergenomics Calculator

6 Why Emulex?

Emulex is a trusted supplier to thousands of data centers globally. Here are some of the key reasons why Emulex is the solution provider of choice for network convergence.

Comprehensive Range of Converged Network Adapters

Emulex offers a range of network convergence solutions that include the LightPulse® LP21000 family of Converged Network Adapters (CNAs) and the OneConnect® Universal Converged Network Adapters. The LP21000 product family is in production and available for immediate deployment. The OneConnect Universal CNA products are expected to begin shipments during the second half of 2009.

Emulex LightPulse LP21000

The Emulex LightPulse LP21000 CNA family are intelligent multi-protocol adapters that provide host LAN and Fibre Channel SAN connectivity over 10 Gigabit Ethernet using FCoE and lossless Ethernet functionality. Unrivaled scalability and industry-leading virtualization support make the LP21000 CNA an ideal solution for server I/O consolidation.

Figure 16: Single and dual-port versions of Emulex LP21000 Family of CNAs

Leveraging seven generations of advanced, field-proven Fibre Channel technology, the PCI Express-based LP21000 family meets the robust interoperability and reliability requirements of corporate data centers. The design leverages Emulex's proven enterprise class drivers, firmware and hardware architectures, while delivering the sophisticated capabilities required to manage multiple types of data flow concurrently without disrupting application performance.

OneConnect Universal CNA Platform

The Emulex OneConnect™ Universal Converged Network Adapter (UCNA) is a single-chip, high-performance 10Gb Ethernet platform designed to address the key challenges of evolving data center networks and improve the overall efficiency of data center operations. Unlike first generation CNAs that only provide FCoE convergence, the Emulex OneConnect UCNA technology provides optimized performance for all protocols (TCP/IP, FCoE, iSCSI) enabling one card for all applications.

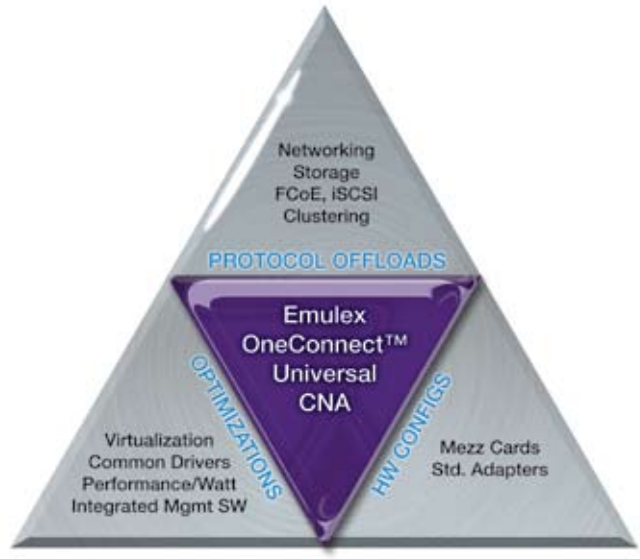

Figure 17: Emulex UCNAs provide best-platform flexibility

The Emulex UCNA platform enables data center managers to consolidate multiple 1Gb Ethernet links on to a single 10Gb link. It supports TCP/IP, FCoE and iSCSI on a single platform to meet connectivity requirements of networking and storage. This diverse applicability of the UCNA simplifies server hardware configurations and significantly reduces the number of standard server configurations deployed in the data center.

The use of a highly integrated single-chip design in the Emulex UCNA enables deployments that meet the design constraints of a server LAN-on-Motherboard, mezzanine cards for blade servers and standalone adapters for rack servers. The use of a common connectivity platform across different server hardware simplifies driver management and optimizes the number of servers managed by an administrator.

The use of multiple protocol accelerators/offload engines allows Emulex OneConnect UCNA to deliver maximum performance, regardless of the mix of protocol traffic.

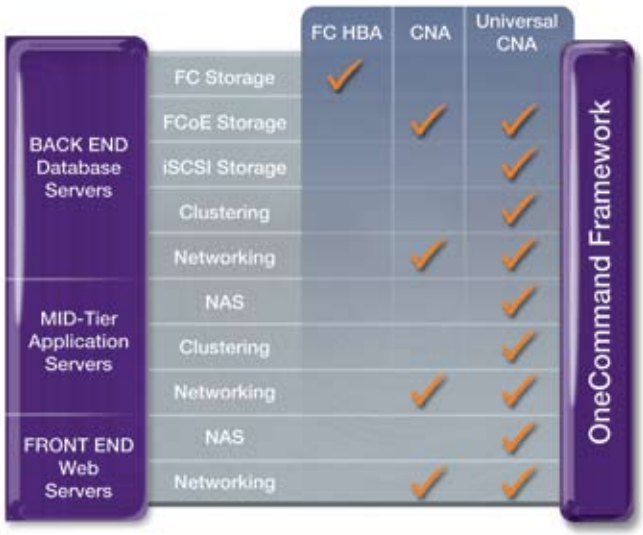

Figure 18: Emulex UCNAs provide best-application flexibility

Unified Management Interface

From installation to management, Emulex provides an extensive collection of tools to simplify the use of FCoE CNAs across a wide range of server environments. The installation utility streamlines the deployment of Emulex CNA drivers. Scripting capabilities supported by the utility enable the driver installation process to be automated for SAN-wide deployment.

The Emulex HBAnyware® management suite delivers a powerful management application, which is upwardly integrated with element management system applications such as IBM Director and HP System Insight Manager through the use of standard HBA APIs. This level of integration greatly simplifies the process of provisioning and deploying new servers with FCoE connectivity, saving IT departments both time and money.

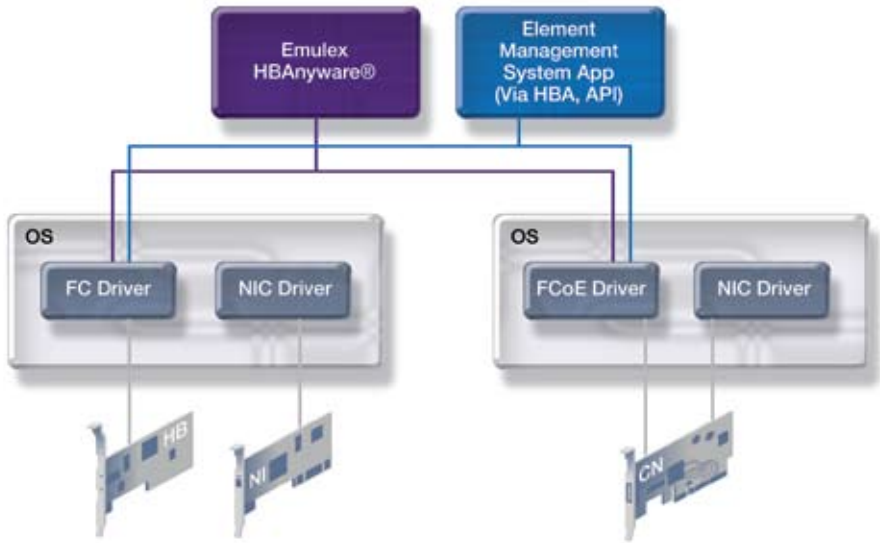

Figure 19: Use of common management tools for both Fibre Channel and FCoE enables simplified provisioning and deployment

Leadership in Enterprise SAN Connectivity

Emulex is an industry leader in providing SAN connectivity solutions for enterprise data centers. The company's products include the award-winning LightPulse family of Fibre Channel host bus adapters, with fully qualified OEM-supported drivers based on a common architecture that spans both product lines. Providing common driver architecture across Fibre Channel HBAs and CNAs allows end users to seamlessly integrate Emulex FCoE-based CNAs into existing data centers using the same software infrastructure across the SAN.

Strong Partner Ecosystem

Emulex is working with an ecosystem of partners to bring interoperable FCoE-based solutions to market.

7 Converged Network Solution Partners Ecosystem

Emulex is working with an ecosystem of partners to bring interoperable network convergence solutions to enterprise data centers. This ecosystem of partner solutions helps direct organizations on a predictable path to the broad benefits of a converged network environment.

The combination of the Emulex CNAs and partner solutions provide a standards-compliant, fully interoperable solution that facilitates faster adoption of FCoE-enabled convergence in the data center.

This chapter captures input from Emulex partners that are successfully testing and implementing network convergence and support for Emulex Converged Network Adaptors. The partner sections present a brief overview of their respective solutions, customer benefits in implementing the solution and target applications for maximizing the return on investments. Emulex expresses its thanks to the respective partner organizations and the individuals within these organizations for their contributions.

BLADE Network Technologies

I. Solution Description

BNT 10-port 10Gb Ethernet Switch Module for IBM's BladeCenter

BLADE Network Technologies (BNT) is a leading supplier of 1Gb and 10GbE network infrastructure solutions that reside in blade servers and "scale-out" server and storage racks. The BNT 10-port 10Gb Ethernet Switch Module for IBM's BladeCenter chassis provides the industry's first embedded CEE/FCoE-capable 10GbE switch for blade servers to enable fabric convergence of data and storage.

With 10-port 10Gb uplinks and 14-port 10Gb downlinks, this product is ideal for customers who require FCoE solutions for IBM BladeCenter, as well as customers with bandwidth-intensive applications such as virtualization and HPC.

10GbE is becoming an even more compelling technology for networking and storage as clients see the value of FCoE. With the BNT 10-port 10G switch, clients have the comfort of knowing the hardware is FCoE-ready and supports the following features that can guarantee lossless transmission for traffic:

- Low Latency
- Priority-based Flow Control (PFC)
- Enhanced transmissions Selestion (ETS)
- Bandwidth allocation per priority group
- FIP Snooping (FC Initialization Protocol)
- End point dicovery and fabric association
- Jumbo frames

II. Key Benefits

The BNT 10-port 10Gb Ethernet Switch Module for IBM's BladeCenter can reduce the number of adapters and switches that are required. This lowers purchase costs, power requirements and management overhead, saving both capital and operating costs. The BNT 10Gb Ethernet Switch Module provides the following benefits to data centers:

- Maintains system uptime with unmatched high availability and proven resilience.
- Consumes up to 75% less power and heat than some 10G 1U rack offerings
- Delivers as much as 7.5 times better performance per watt when compared to a 1Gb switch.

III. Target Applications

The 10Gb Ethernet Switch Module provides unmatched maximum performance for computing clusters (HPC), as well as bandwidth-intensive and time-sensitive applications. Examples include: FCoE SANs, cloud computing, financial analytics, medical imaging, virtualization, video on demand, HD IPTV and VOIP.

IT centers of all sizes can benefit from converging LAN and SAN traffic on one consolidated Ethernet wire with all of the QoS, security, and performance associated with a dedicated network. And as other applications, such as Voice over IP, become more widely used, the efficiencies associated with a single unified fabric around Ethernet become that much more attractive.

Investment Protection

The BNT switch is designed with investment protection in mind. For clients who currently use blade servers with the 1G CFFh quad port adapters, this switch supports 1G into the switch and provides 1G/10G uplinks. As data centers start to roll out upstream infrastructure, the 10G uplinks can be leveraged. And as data centers start to use the 10G adapters in the blade servers, the switch supports 10G downlinks to the server NICs.

The first FCoE products will be CNAs that provide server LAN and SAN connectivity over Ethernet, and FCoE switches that connect FCoE-capable initiators (on the server side) to existing Fibre Channel SANs. These products will enable consolidation of server interconnectivity to a single Ethernet fabric.

To the other nodes in the Fibre Channel SAN, the CNAs appear to be directly connected and can be managed and maintained with the same tools. For environments with large Fibre Channel investments, FCoE provides an ideal way to transition from one physical protocol to Ethernet without engaging in extensive training. As FCoE becomes more popular, native FCoE storage arrays will appear on the market, enabling a fully converged fabric.

IV. Management

The BNT switch modules eliminate the complexity of managing separate network fabrics by enabling consolidation onto 10GbE. Available tools simplify deployment and day-to-day management. The 10-port 10Gb switch module provides seamless integration and easy management with the following capabilities:

- Standards-based integration into Cisco and other networks to reduce downtime and learning curbs.

- Common look and feel among BNT switches to help administrators minimize the learning curve when they have a requirement for different switches.
- Support for two CLI options—the BLADEOS CLI, as well as an industry-standard CLI.
- Easy software upgrades through Web user interface, TFTP, telnet or serial download for easier adaptation to existing maintenance procedures.
- Enhanced security and traffic management.

V. Deployment Architecture

This collaboration showcases direct connectivity between an Emulex FCoE CNA and a NetApp FCoE target via BLADE's 10GigE switches in blade or rack-based form factors.

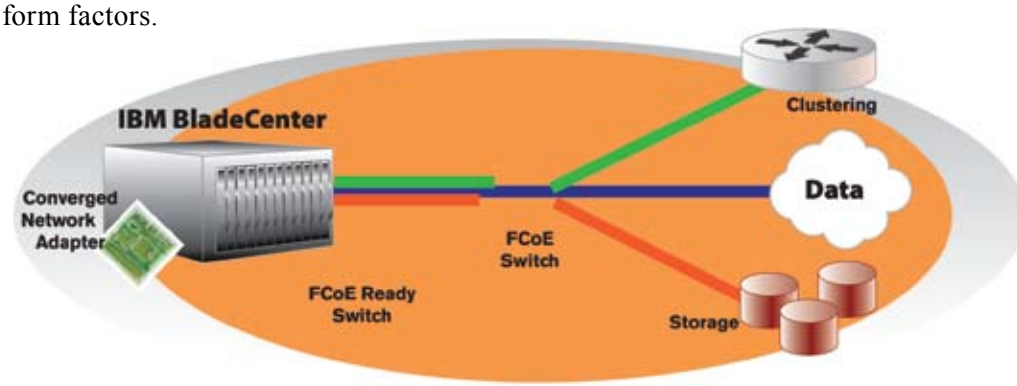

Figure 20: Direct connectivity with Emulex CNAs via BNT's 10G Ethernet switch.

Brocade

I. Solution Description

The Brocade FCoE solution includes the Brocade 8000, a line-rate top or rack CEE/ FCoE switch, and the Brocade Data Center Fabric Manager (DCFM) management application.

The Brocade 8000 switch uses non-blocking, cut-through architecture to deliver 10GbE CEE performance on each of its 24 CEE ports. It also utilizes a full FC switch that delivers 8Gb/s FC performance on each of its right (8) FC ports. Since the Brocade 8000 switch incorporates a true FC switch, its FC ports can connect directly to FC storage, saving you the trouble of investing in an additional box to bridge the FCoE and the FC storage. The Brocade 8000 supports a robust L2 Ethernet feature set and IEEE DCB enhancements for CEE including Priority Flow Control (PFC), Enhanced Transmission Selection (ETS) and Data Center Bridging Exchange (DCBx). Customers may use fiber or Twinax cables to connect the top of rack switch to their servers.

To support the most data-intensive applications, the Brocade 8000 provides best-in-class performance with a non-blocking architecture that supports port trunking for Fibre Channel and link aggregation for Ethernet. For Fibre Channel, an Inter-Switch Link (ISL) trunk can supply up to 64Gb/s of balanced data throughput. For Ethernet, the Brocade 8000 supports standards-based Link Aggregation Control Protocol (LACP) and Brocade enhanced frame-based trunking.

All the components of the Brocade FCoE solution are managed by Brocade DCFM and are based on Brocade data center-proven single Fabric OS (FOS). Customers can manage their CEE resources using a command line interface, while managing the storage side using the familiar Brocade management interface. Brocade's data center-proven expertise and strong adherence to open industry standards guarantee customers an open path to future technologies and a robust FCoE ecosystem.

II. Key Benefits

- High Availability: The Brocade 8000 provides a reliable foundation for disaster recovery and business continuance by employing enterprise-class availability features such as hot-swappable, redundant, and integrated fan and power supply assemblies. Combined with a wide range of diagnostic and monitoring functions, these capabilities help provide a highly available environment.

- Server Edge I/O Consolidation: Transporting storage and networking data over CEE reduces the number of server adapters, cables, and required switch ports—which in turn simplifies cabling and reduces power consumption and cooling costs.

- FC Investment Protection: FCoE is designed to preserve FC upper constructs, which preserves existing FC management tools and models. This enables the Brocade 8000 switch to integrate seamlessly with the existing FC environment without disruption. Customers can continue to use the familiar Brocade management tools; thus extending the value of investments in software and staff training.

- Greener Data Centers: Customers can get closer to their goals of green data centers when using the Brocade FCoE solution. The Brocade 8000 is rated for a maximum of 350W, which is much lower than comparable products. Using the Brocade 8000 abundant FC bandwidth and full FC switching capabilities eliminate an entire box needed by competing products to connect to FC storage.

III. Target Applications

The Brocade 8000 switch is a top of rack CEE/FCoE switch that is targeted at server I/O consolidation where it is used as an access layer switch connecting servers to LAN and FC storage in the corporate SAN. In this role as a TOR switch, the Brocade 8000 receives converged traffic from CNA-equipped servers and redirects it to the desired destination. LAN traffic is carried out on CEE ports, while SAN traffic makes its way to FC storage using the Brocade 8000 FC ports.

Server I/O Consolidation and Virtualization

The Brocade FCoE solution provides the benefits of server I/O consolidation while preserving data center environments where corporate LANs and SANs are connected to converged servers, without any disruption. The Brocade 8000 connects to servers through CNAs. The consolidated SAN and LAN server ports and corresponding cables simplify configuration and cabling in server cabinets to reduce acquisition costs. This design helps reduce cable clutter and operational costs associated with power consumption and cooling. It also offers a simple and lower cost server I/O solution that enables customers to deploy FCoE and CEE without disrupting existing environments and with future expansion in mind.

Additionally, FCoE and CEE 10GbE interfaces provide the scalability needed in highly demanding virtualized server environments and offer converged interfaces that facilitate virtualized application mobility to meet the dynamic needs of enterprise data centers.

IV. Management

Consistent and simple management is essential for the success of FCoE and CEE deployment. All the elements of the Brocade FCoE solution are based on the company's proven Fabric OS and DCFM management application. DCFM works with third-party tools and applications, enabling IT professionals to continue using management tools with which they are familiar. With DCFM, networking professionals can continue to use command line interface, while storage professionals can use the interface with which they are most accustomed.

V. Deployment Architecture

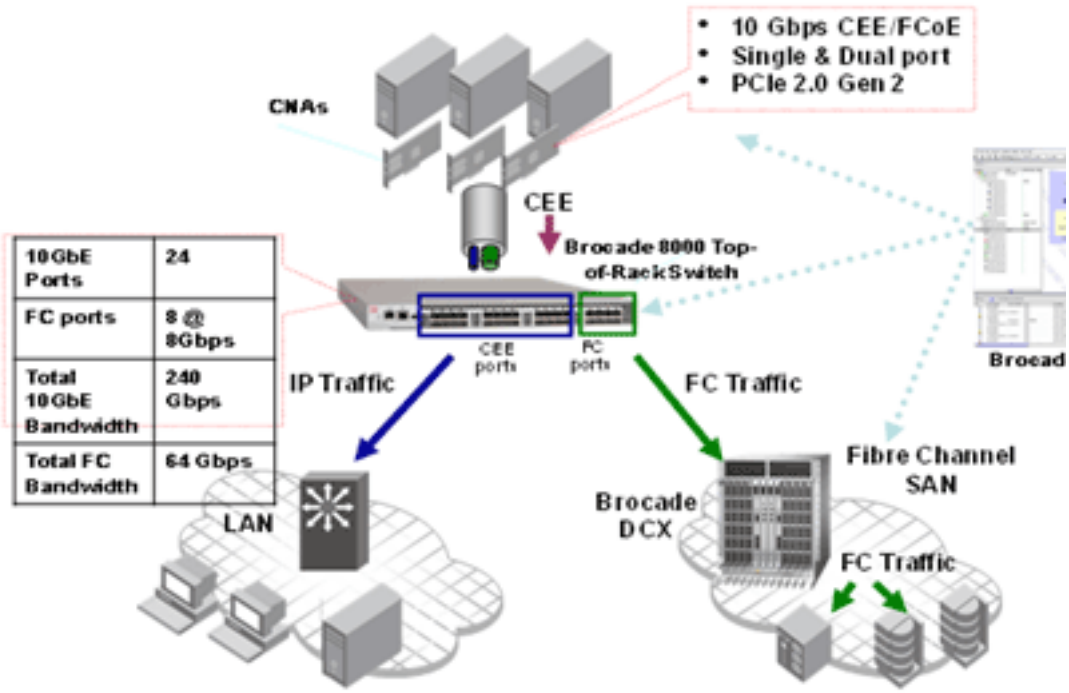

Figure 21: Brocade FCoE solution deployment

Cisco

I. Solution Description

The Cisco Nexus 5000 Series Switches comprise a family of line-rate, low-latency, lossless 10 Gigabit Ethernet, Cisco® Data Center Ethernet, and Fibre Channel over Ethernet (FCoE) switches for data center applications.

Today's data centers are increasingly filled with dense rack-mount and blade servers that host powerful multi-core processors. The rapid increase of in-rack computing density, along with the increasing use of virtualization software, combine to push the demand for 10GbE and consolidated I/O: applications for which the Cisco Nexus 5000 Series is the perfect match. With low latency, front-to-back cooling, and rear-facing ports, the Cisco Nexus 5000 Series is designed for data centers transitioning to 10GbE as well as for those ready to deploy a unified fabric that can handle their LAN, SAN, and server clusters, networking over a single link (or dual links for redundancy).

The switch family, using cut-through architecture, supports line-rate 10GbE on all ports while maintaining consistent low latency independent of packet size and services enabled. The product family supports Cisco Data Center Ethernet capabilities that increase the reliability, efficiency, and scalability of Ethernet networks. These features allow the switch to support multiple traffic classes over a lossless Ethernet fabric, thus enabling consolidation of LAN, SAN, and cluster environments. Its ability to connect FCoE to native Fibre Channel protects existing storage system investments while dramatically simplifying in-rack cabling.

In addition to supporting standard 10GbE network interface cards (NICs) on servers, the Cisco Nexus 5000 Series integrates with consolidated I/O adapters (CNAs) that combine Ethernet NICs and Fibre Channel (HBAs), making the transition to a single, unified network fabric transparent and consistent with existing practices, management software, and OS drivers. The switch family is compatible with third-party integrated transceivers and Twinax cabling solutions that deliver cost-effective connectivity for 10GbE to the servers at the rack level, eliminating expensive optical transceivers.

The Cisco Nexus 5000 Series switch fabric uses cut-through technology that enables consistent low-latency Ethernet solutions independent of packet size, regardless of the network services enabled. The product family is designed for data center environments, with front-to-back cooling and network ports in the rear, bringing switching into close proximity with servers and making cable runs short and simple. The switch family is highly serviceable, with redundant, hot-pluggable power supplies and fan modules. Its software is based on data center-class Cisco NX-OS Software for high reliability and ease of management.

The Cisco Nexus 5000 is the result of over $1 billion in data center-related research and development. This solution offers the following strengths over competing solutions:

- Integrated Fibre Channel and Ethernet: In today's data center, applications are creating increased demand for server processing and I/O consolidation at a time when budgets are shrinking. Only the Cisco Nexus 5000 provides a unified fabric by integrating Fibre Channel and Ethernet traffic over the same connection in a lossless manner, allowing you to grow to meet current demands by consolidating and sharing resources across multiple applications.

- Proven data center architecture: Cisco best-practice guidelines for architectural principles, technical design, and blueprints provide a map for evolving your data center. These principles and guidelines are based on thousands of real-world deployments, allowing companies to benefit from our expertise and experience. The result is a proven, world-class data center architecture that reduces the risk and uncertainty of "do-it-yourself" or piecemeal approaches.

- Most comprehensive solution: The Cisco Nexus 5000 Series solution is not just about world-class data center hardware and software. As part of a holistic approach, the Cisco Nexus 5000 Series solution also includes Cisco Lifecycle Services, Cisco Capital, the Data Center Assurance Program (DCAP), and an industry-leading partner ecosystem.

- Global presence: The trend towards globalization is one of the reasons companies need to upgrade their data center infrastructures. Unlike the past when regional subsidiaries handled international business, corporate data is now consolidated from around the world. No matter where your data center is located, Cisco can offer world-class support and consulting expertise.

II. Key Benefits

- The combination of high port density, lossless Ethernet, wire-speed performance, and extremely low latency makes the switch an ideal product to meet the growing demand for 10GbE with the capability to support Cisco Data Center Ethernet in enterprise and service provider data centers, protecting enterprises' investments. The switch family has sufficient port density to support single or multiple racks fully populated with blade and rack-mount servers.
- Built for today's data centers, the switches are designed just like the servers they support. Ports and power connections are at the rear, closer to server ports, helping keep cable lengths as short and efficient as possible, delivering benefits traditionally offered on blade servers to rack servers as well. Hot-swappable power and fan modules can be accessed from the front panel, where

status lights offer an at-a-glance view of switch operation. Front-to-back cooling is consistent with server designs, supporting efficient data center hot- and cold-aisle designs. Serviceability is enhanced with all customer-replaceable units accessible from the front panel. The use of SFP+ ports offers increased flexibility to use a range of interconnect solutions, including copper Twinax cable for short runs and fiber for long runs.

- Cisco Data Center Ethernet features increased network scalability, supports I/O consolidation, eases management of multiple traffic flows, and optimizes performance. Although implementing SAN consolidation requires only the lossless fabric provided by the Ethernet pause mechanism, the Cisco Nexus 5000 Series provides additional features that create an even more easily managed, high-performance, unified network fabric.

- FCoE is part of the unified network fabric that enables I/O consolidation at the rack level. It is a straightforward encapsulation of Fibre Channel within Ethernet that preserves existing Fibre Channel network management models and tools, helping protect investments in software and staff training. Additional switch features include delayed drop, which helps increase resiliency after bursts of traffic, and Ethernet Host Virtualizer (EHV), which supports higher performance. Massively scalable I/O consolidation in racks and in rows helps reduce capital and operational costs by reducing the number of server adapters, cables, and upstream switches needed to support I/O at the rack level. Rather than the overhead of a redundant pair of adapters for each of up to three networks (LAN, SAN, and cluster), I/O consolidation supports all three networks on a single link. The switch family connects to native Fibre Channel networks, protecting existing investments in storage networks. The switch family's support for FCoE also supports data center I/O consolidation. As FCoE-enabled switching becomes available across the data center, FCoE traffic can travel over a unified network fabric directly to FCoE-enabled storage devices or to native Fibre Channel SANs.

- Energy efficiency achieved through the use of the Cisco Nexus 5000 Series switches helps data centers better operate within their space, power, and cooling parameters while reducing their carbon footprints. Every network link at the rack level requires adapters, switches, and transceivers, all of which consume power. I/O consolidation reduces energy consumption by eliminating the need for separate Fibre Channel adapters, cables, and switches. In many cases, server cluster networks also can be consolidated onto 10GbE networks, especially given the low latency of the Cisco Nexus 5000 Series. The switch hardware is also designed for energy efficiency. Variable-speed fans consume only the amount of power necessary to cool the chassis . The power supplies are sized to support worst-case scenarios where inefficient SFP+ transceivers maximize power draw; however, when low-power cabling solutions are deployed, the switch family's power supplies maintain 90 percent efficiency at only 25 percent utilization, making efficient use of power in best-case scenarios.

III. Target Applications

The Cisco Nexus 5000 Series supports a number of application scenarios that use it as an access layer switch co-located with servers in data center racks or for middle-of-row deployments.

- As an access-layer switch, it can be used purely as a 10GbE switch, consolidating 10GbE connections into a smaller number of server connections trunked to the aggregation layer.
- As a rack-level I/O consolidation platform, the switch carries Ethernet traffic from servers to the aggregation layer and carries FCoE traffic to existing Fibre Channel SANs.
- As a crucial element in data center I/O consolidation, the switch paves the way with I/O consolidation at the access layer and interoperability with Cisco Nexus and other standards-based products.

The capability of the Cisco Nexus 5000 Series to function in all these capacities helps protect investment in the data center with a deployment model in which additional features are enabled as they are needed.

Rack-Level 10 Gigabit Ethernet Access-Layer Switch

The switch family is designed with the density, performance, front-to-back cooling, and rear port configuration that makes it ideal for aggregating a large number of 10GbE links from servers into a smaller number of uplinks to aggregation-layer switches. The switch port density allows each switch to support a single rack or neighboring racks using the SFP+ direct attached 10 Gigabit copper cabling option. The Cisco Nexus 5000 Series switches can be purchased with only the Ethernet capabilities enabled, allowing IT departments to deploy them in parallel with existing Fibre Channel SANs.

Unified Fabric with FCoE: I/O Consolidation

The switch family is built to consolidate multiple networks—LAN, SAN, and server cluster—onto a single unified fabric, saving the capital and operational expenses associated with multiple parallel networks, switching infrastructure, and cabling mechanisms within racks. The Cisco Nexus 5000 Series Switches are compatible with third-party consolidated I/O adapters that present separate Ethernet and Fibre Channel NICs to the server operating system, allowing existing drivers and Fibre Channel management software to work transparently with FCoE. Upstream, two different expansion modules support direct connections from the Cisco Nexus 5000 Series to existing native Fibre Channel SANs.

IV. Management

Management for Cisco products is provided through consistency of both Cisco NX-OS Software and Cisco MDS 9000 SAN-OS software management models and tools. The switch family network features can be managed using the Cisco command-line interface (CLI), and the Fibre Channel and FCoE features can be managed through Cisco Fabric Manager Suite. The capability to manage Ethernet and FCoE features independently with existing Cisco tools preserves existing management models, best practices, and investments in staff training. In addition, Simple Network Management Protocol (SNMP) MIBs, Extensible Markup Language (XML), and the Cisco CLI are made available to customers for switch management through third-party or custom-developed tools. The switch family is based on Cisco NX-OS for superior operational efficiency, pervasive security, and continuous operation even through software upgrades.

Software manageability and serviceability features include Smart Call Home and automated parameter exchange (through DCBX). Security is enhanced through role-based access control (RBAC); support for authentication, authorization, and accounting (AAA), remote TACACS+, and RADIUS servers; and Secure Shell (SSH) access.

V. Deployment Architecture

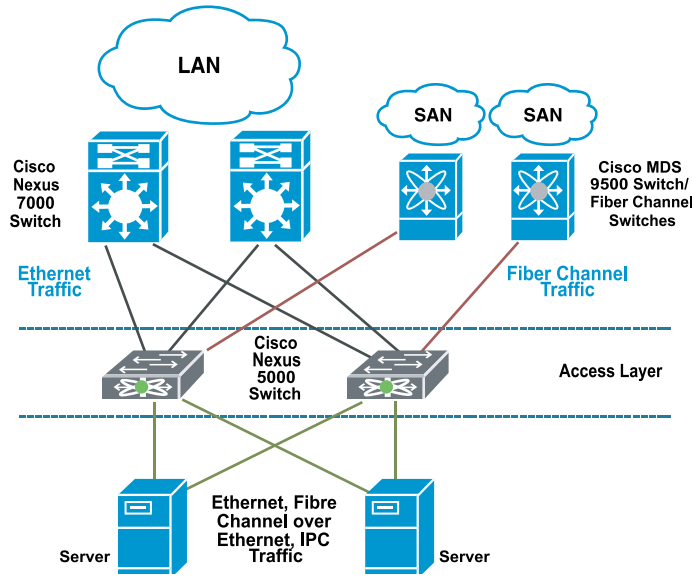

Figure 22: Converged network architecture with Nexus switches

EMC

I. Solution Description

EMC's initial FCoE offering consists of the Nexus 5000 switch family from Cisco Systems, which is sold and serviced by EMC under the EMC Connectrix brand. In addition, CNAs from Emulex are available through EMC Select. EMC Select gives customers a simple and efficient way to acquire best-in-class products, like Emulex HBAs and CNAs, that complement EMC products.

The NEX-5000 series switches support GbE ports and have expansion slots available for switching modules. For server connectivity into an existing SAN, a switching module is available to support four 4 Gb/s Fibre Channel ports and an additional four 10GbE ports. The 10GbE, 6-port, Ethernet-only module provides support for 10GbE, Fibre Channel over Ethernet, and Cisco Data Center Ethernet.

Major features and benefits of the EMC Connectrix NEX-5000 Series include:

- Line-rate 10GbE with a switching architecture that supports full-bandwidth, low latency, and predictable performance on all ports.
- Support for FCoE, which can provide server I/O consolidation because LAN and SAN traffic travels on a single 10GbE link. This can simplify cabling, reduce number of adapters, lower costs, and cut power consumption.

FCoE increases the reach of Fibre Channel SANs and allows more environments to take advantage of the operational efficiencies of networked storage. The technology will mature over time and will be delivered to market in a phased approach. With over 10 years of experience designing and deploying networked data centers, EMC is well-positioned to provide and support FCoE solutions.

II. Key Benefits

Consolidation Make Sense

FCoE enables consolidation of server cables and adapters by allowing LAN and SAN traffic to travel on a single 10GbE link. Instead of purchasing multiple NICs, HBAs, and HCAs for servers, customers can purchase a CNA, eliminating the need for separate interface cards. This reduces the number of server cables and adapters in the data center and greatly simplifies the physical infrastructure. It also reduces server power and cooling costs as well as the administrative overhead associated with managing the data center. Another advantage of FCoE is that it seamlessly integrates with existing Fibre Channel networks, management processes and workflows, offering a smooth migration path for existing Fibre Channel customers.

FCoE Differentiators

There are multiple storage networking technologies available in the market today, each of which was developed to allow individual servers to access a centralized pool of storage. Before SANs were introduced, storage was directly attached to servers with point-to-point connections, which isolated resources and limited utilization and efficiency. Fibre Channel SANs were developed to connect storage islands using a high-speed network, enabling all servers and applications to access all disks. Sharing storage greatly improved the utilization, availability, and performance of storage resources.

Other protocols emerged to meet specific business and computing requirements. Fibre Channel over IP (FCIP) and iSCSI were created to extend the SAN to isolated servers or geographically dispersed sites using traditional IP networks. Network attach storage (NAS) was introduced for file sharing, allowing many computers to access the same file system over the network. Infiniband was developed for environments that require extremely high bandwidth and low latency. Each protocol was designed to address a specific use case and each has its own pros and cons.

What makes FCoE different? FCoE combines LAN and SAN traffic over a single 10GbE connection. The biggest differentiator is that unlike other technologies, FCoE preserves existing Fibre Channel SAN investments because existing SAN management tools and concepts can be leveraged without requiring new applications. It also eliminates the overhead of TCP/IP associated with iSCSI.

III. Target Applications

Initial Target Applications/Use Cases

The main application for FCoE is in the enterprise data center. These customers likely have one or more Fibre Channel SANs in place already and are looking to connect more servers to that network. FCoE offers them another option and is well suited for dense rack-mounted and blade servers where cable simplification is the key requirement. Virtualized servers that require many I/O connections will also benefit from FCoE.

Customers that are implementing large server and storage consolidation projects are prime candidates for FCoE with today's solutions. These projects are driven by the need to save physical floor space in the data center, reduce capital expenditures, and attach more servers to the storage network. VMware will often be a component of these projects. The consolidation effort is also being driven by power and cooling constraints. Many enterprise customers are either running out of power or want to reduce the carbon footprint of their data centers. Whether it is by building an entirely new data center or redesigning an existing one, the main driver for these IT projects is cost-reduction.

Phases of FCoE Deployment

Many companies (mostly small to medium size businesses) are already running their storage traffic over Ethernet by utilizing iSCSI. FCoE is not for these customers and was not designed to replace iSCSI. For those customers that have an existing infrastructure and knowledge base of Fibre Channel, FCoE provides a path toward reaching a converged fabric. Both FCoE and iSCSI will be able to leverage 10GbE and lossless Ethernet enhancements (Figure 4). In general, iSCSI environments tend to be small configurations with five to 20 servers, while many FC customers are scaling into hundreds or thousands of nodes; FCoE can be plugged in to these existing environments.

All new technologies take time to develop a robust ecosystem, and FCoE is no exception. Since FCoE will build on existing infrastructure, it can be rolled out in a phased manner as functionality is created and hardened. EMC recognizes the first phase of FCoE deployment as the server phase, in which the separate network (NIC) and storage (HBA) interfaces on the server are replaced by the CNA that is directly attached to the Converged Network Switch (CNS). A new rack of servers can deploy these technologies while the existing storage and network environments remain unchanged.

The second phase of FCoE identified by EMC will be the network phase, in which the CNS moves out of the rack and into a unified network. In order to create this unified network, the Ethernet enhancements must be able to create a safe environment for FCoE in the same fabric as the rest of the networking traffic. In addition to the technical concerns, there may be some internal "political" discussions for customers to decide whether the storage or networking group owns the purchase, installation, and maintenance of the CNA and CNS components. All of the enhancements to Ethernet will be new to network administrators who will need to learn these functions and will require coordination with storage administrators. The overlapping domains may compel cultural adjustments, as storage networks will no longer be dedicated and network configurations can no longer be reconfigured at-will. Network and storage administrators will each have their own management interfaces to the environments, keeping tasks separate rather than converged.

And finally, EMC believes the third phase is the storage phase in which end-to-end FCoE is enabled with native FCoE storage. The ecosystem must be robust enough to be able to create a large network that reliably provides the same functionality that is available with FC today. Note that even in a totally Ethernet environment, the solution still maintains FC at the upper layer at both the host and storage.

No discussion of FCoE is complete without addressing the cabling infrastructure supporting the solution. Most data centers have not deployed 10GbE even though the standard has been complete and products have been available for many years. For 1 Gigabit Ethernet, the primary options for cabling are copper (1000Base-T with RJ-45 connectors) and optical (same physical cabling as FC); copper dominates this market with billions of ports installed while optical has historically been 1 percent

of Ethernet ports. A standard (10GBase-T) for using existing copper cabling (either CAT 6 or CAT 6a with RJ-45 connectors) is available, but currently the products are expensive and have high power requirements. A new copper option known as Twinax has become available for FCoE solutions of 10GbE. The option is based on the SFF-8431 standard and uses the SFP+ interface for a copper connection that is passive, low cost, and low power. The Twinax solution is limited to short distances; first-generation products are expected to support five meters, which is sufficient for direct attach environments in a rack. Standard multimode optical cabling will be used for environments that require longer distances such as from the rack to the core. By the time the market is ready for a unified network, it is expected that the traditional copper solutions will reach a power and price point that will allow customers to take advantage of their existing cabling plant infrastructure.

IV. Deployment Architecture

The EMC Networked Topology Guide provides a top-down view of networked storage and assists the network designer in designing a suitable networked storage infrastructure. This guide provides best practices for storage networking technologies that are created through EMC E-Lab interoperability and integration testing and authored by E-Lab engineers and product content experts. For complete detail of EMC's FCoE topology guidelines, please refer to this guide found on EMC's PowerLink website. The following is a diagram of a sample supported topology for the EMC Connectrix NEX-5000.

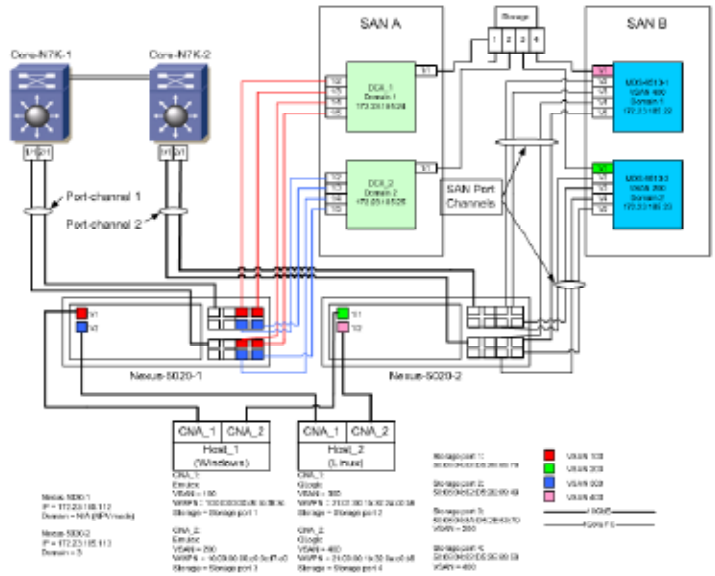

Figure 24: SAN connectivity with EMC Connectrix NEX-5000

Why EMC for SAN? Experience, Interoperability and Choice

Experience: EMC has over a decade of experience and industry-leading expertise designing, implementing, and managing the world's most critical SAN infrastructures. EMC deploys more than 7,500 business and technical specialists who help customers plan, design, implement, and manage the world's largest and most demanding Storage Area Networks. Because EMC's service and support resources are based locally, the people who architect and implement your SAN are the same people who will support you as requirements change and grow.

Interoperability and the EMC Promise: EMC delivers the highest level of interoperability assurance and leads the industry in interoperability and multi-vendor support. EMC doesn't just test devices—but looks at the big picture by testing applications, hosts, operating systems, host bus adapters, converged network adapters, Fibre Channel, backup applications, and clustering software, as well as EMC and non-EMC storage.

Choice: EMC's breadth of choice means you are not locked in to any one technology. This becomes especially important as technologies evolve—from 1Gbit to 10Gbit bandwidth…from Fibre Channel-only SANs to combined Fibre Channel, FCoE and IP SANs…from multiple, separate SAN deployments to larger, more integrated deployments. It will be equally important as SANs evolve to host network applications, such as EMC's Invista and RecoverPoint.

Fulcrum Microsystems

I. Solution Description

Fulcrum's industry-leading FocalPoint 10GbE switch chip family adds a new dimension to Ethernet—low latency. FocalPoint delivers 200 nanoseconds (ns) of switching latency, a 10X improvement over alternative solutions. This solution provides unparalleled performance critical for next-generation data centers and cloud computing architectures. Advanced congestion management and flow control techniques enable lossless operation and bounded latency for transporting and isolating high performance computing and storage networking traffic, enabling true converged fabrics that are capable of scaling to hundreds or thousands of nodes.

With its ControlPoint network OS, Fulcrum offers a complete solution along with a worldwide ecosystem of partners from which storage and clustering system designers can drive adoption of 10GbE as the fabric of choice for a fully converged data center.

Monaco Reference Design

Fulcrum and Emulex have demonstrated data center network convergence over 10GbE using the Fulcrum Monaco reference design platform and Emulex LightPulse CNAs. The combined solution enables networking and storage traffic to be carried over a single infrastructure while ensuring appropriate priorities for each.

The Monaco platform, which contains a FocalPoint FM4224 switch chip, is configured for lossless transfer of FCoE traffic by using a separate memory partition and by sending Priority Flow Control (PFC) frames to the CNA. Using this configuration, storage transactions have been successfully demonstrated to a native FCoE storage target at 4Gbs.

The FM4000 series is a family of fully integrated single chip wire-speed 10GbE switch chips, which include enhanced features such as low latency, scalability, packet classification, L3 routing, PFC, Quantized Congestion Notification, Enhanced Transmission Selection, and support for FCoE or DCBX. In addition to enhanced functionality, the FM4000 series contains up to 24 10GbE ports. All of this is accomplished while maintaining the best-in-class latency and throughput available in the FocalPoint product line.

Monaco is a complete 1U single-board reference design for the FocalPoint 10GbE switch, and can be used to rapidly and exhaustively evaluate the performance and flexibility of the FocalPoint FM4224 switch / router, either in a test environment or a live data center network. The platform contains 24 SFP+ interfaces in a 1U rack-mountable chassis, and can be ordered with software installed and ready to run.

Critical Capabilities to Support Converged Networks

Fulcrum's newest offerings include critical capabilities to successfully support converged networks in a reliable way. The first of these is the ability to truly support lossless Ethernet which is made possible via advanced congestion management, predictable low latency and flow control, all of which are important for packet order-sensitive applications.

As the size and density of data centers increase, the cost, area, power and support of multiple interconnect fabric technologies also increase; converged networks provide an excellent solution.

Storage protocols, however, require deterministic, bounded latency to ensure that time-out values are not exceeded. Traditional Ethernet switches drop packets during periods of high congestion, requiring retransmission at a higher layer. This added latency cannot be tolerated in storage applications. Therefore, it is critical that CIOs and IT managers seek solutions that offer the lowest latency possible.

II. Key Benefits

Converged networking with Emulex CNAs and the FocalPoint 10GbE switch chip family simplifies the network architecture and requires fewer pieces of equipment, resulting in cost savings and reduced real estate and power usage.

III. Target Applications

Fulcrum's data center network convergence product goes into servers that support financial services, HPC, Web 2.0, electronic commerce, video (including IPTV, VOD, and VoIP), and customer relationship management applications. FCoE is quickly becoming a reality in many applications through work on new initiatives like Datacenter Bridging (DCB) by task forces ranging from IEEE to Ethernet Alliance to the Fibre Channel Industry Association.

Fulcrum's products will further penetrate into networks such as financial services, and segments such as converged traffic in carrier networks using carrier Ethernet. It will also continue to expand within onsite data centers as well as taking advantage of applications that reside in the "cloud".

IV. Management

Converged networks do not require significant changes in hardware to enjoy its benefits. Legacy switches can be used as long as they support key features such as congestion management.

V. Deployment Architecture

Emulex CNAs enable converged fabric solutions in data center environments. By exchanging PFC frames with the FocalPoint switch fabric, storage and data traffic can be isolated, providing a lossless virtual fabric for storage with deterministic latency. This solution can be scaled to large data center installations, with up to 288 10G ports available in a two-stage FocalPoint fabric.

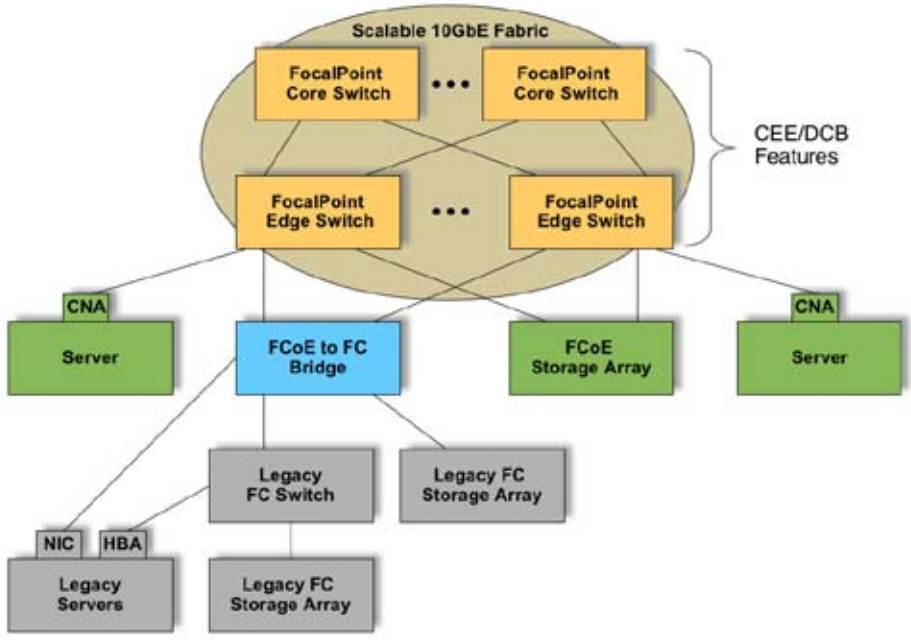

Figure 25: Fulcrum's converged fabric solution enabled with Emulex CNAs

Juniper Networks, Inc.

I. Solution Description

Juniper Networks' EX Series of Ethernet switches simultaneously advance both the economics and fundamentals of networking. Offering plug-and-play 10/100/1000BASE-T connectivity for today's converged networks, the EX Series employs industry standards, carrier-class reliability, and integrated security to deliver an open network for converged communications. The high-performance, high-density EX Series let users start small and grow incrementally, saving valuable space in crowded wiring closets while lowering recurring power and cooling costs. A common operating system and management tools shared across the Juniper portfolio of switches, routers, and security devices reduces operational and capital expenses for the network.

The EX Series can scale from small to large enterprise environments, as well as deliver the features required by largest service provider and carrier environments.

Consistent QoS Capabilities

Building a multiservice network is a strategic initiative that shouldn't be stymied by QoS pitfalls. Juniper Networks designed the EX3200, EX4200 and EX8200 switches to simplify creating a converged infrastructure. The EX Series switches enable enterprises to deliver predictable performance for any combination of traffic with the least possible administrative overhead.

Juniper allows enterprises to follow best practices for ensuring consistent end-to-end QoS by supporting a common set of QoS features from the access layer to the core. For example, robust QoS features in access devices (whether the EX3200 or EX4200) means that network traffic can be classified and policed at the source. Likewise, standard Layer 3 support on all Juniper switch and router platforms means enterprises have the option to use DSCP marking end-to-end.

II. Key Benefits

- Carrier-class Reliability: EX Series switches leverage the same field-proven Juniper technology—including high-performance ASICs, a carrier-class system architecture, and JUNOS Software—that powers the world's largest service provider networks.

- Security Risk Management: EX Series switches tightly integrate with Juniper's Unified Access Control (UAC) solution, dynamically delivering network protection, guest access and identity-based QoS based on user, device and location.

- Performance: Each of the EX Series switches deliver wire-speed performance on all ports for any packet size. For instance, the EX8200 modular switches deliver up to 128 wire-speed 10GbE ports per chassis with nearly 2 billion packets per second throughput—the ideal platform for next-generation data centers.

- Virtualization: Virtual Chassis technology, currently available on the EX4200 line, allows up to 10 switches to be interconnected and managed as a single, logical device. This delivers the reliability, scalability and manageability of traditional chassis-based systems in a cost-effective, compact platform.

- Application Control: The EX Series switches support eight QoS queues per port, ensuring proper prioritization of control plane, voice, video and multiple levels of data traffic—with room to converge other networks, such as building automation and video security systems.

- Lower Total Cost of Ownership: The innovations of the EX Series switches enable new architectures that reduce cost and complexity. Unified management and automation tools consolidate system monitoring and maintenance and save time and money.

III. Target Applications

One of the most significant IT trends is the push by enterprises to combine data, voice and video traffic on a single multiservice network. To date, the adoption of IP-based telephony (IPT) and video—including videoconferencing, video streaming and video-on-demand—has been the primary driver for converged networks.

To accommodate these mixed media applications, IT is tasked with building a network infrastructure capable of supporting converged real-time communications. Such an infrastructure must provide consistent end-to-end quality of service to ensure that each traffic type gets the bandwidth and handling it needs to meet user expectations and business objectives.

Understanding these requirements, Juniper Networks designed the EX Series Ethernet Switches with a comprehensive set of QoS features, making it easy to deploy the same QoS controls across the entire network, from the wiring closet to the WAN edge. With the Juniper Networks EX3200 fixed-configuration switches, the EX4200 switches with Virtual Chassis technology, and the EX8200 switches, IT can deliver the predictable performance needed to support today's converged voice, video and data applications, as well as building automation traffic and emerging UC-enabled applications.

IV. Management

Juniper's support for a consistent set of QoS capabilities, a consistent operating system (JUNOS) and control plane, and a single management interface across all of

its switch platforms greatly simplifies the data center's job. Rather than having to learn how to configure and manage QoS for each individual class of platform, IT can use the same configuration and management tools across all EX Series switches and achieve the same, predictable behavior from platform to platform.

V. Deployment Architecture

The Juniper Networks EX3200 line of Ethernet switches offers a simple, cost-effective solution for low-density branch and regional offices. Deployed in wiring closets to provide network access for users and other IP-enabled devices, the EX3200 switches offer plug-and-play 10/100/1000BASE-T connectivity for today's converged networks.

Designed for access and aggregation deployments, the EX4200 switches are a superset of the EX3200 switches, available in the same 24- and 48-port 10/100/1000BASE-T configurations with full and partial PoE, plus optional 1GbE and 10GbE uplink modules. The EX4200 line also offers a 24-port 100BASE-FX/1000BASE-X SFP-based platform for Gigabit aggregation deployments requiring the long distance links afforded by fiber.

The EX8200 line of modular Ethernet switches delivers a high-performance, highly available platform for today's high-density 10GbE data center, campus aggregation and core networks.

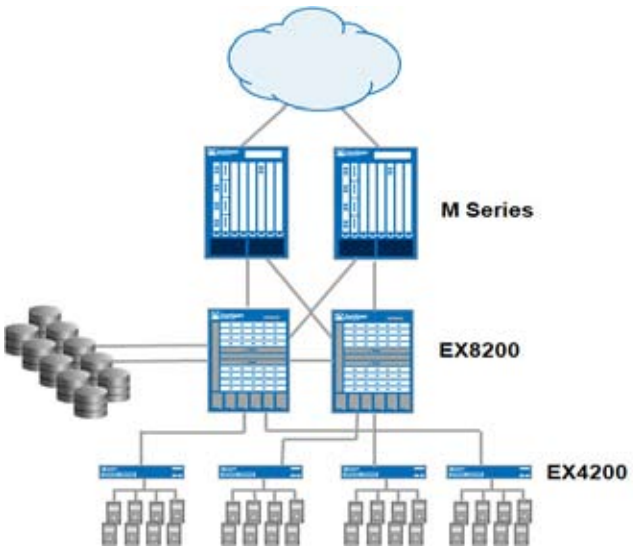

Figure 26: Converged network with Juniper's EX Series of Ethernet switches

Novell

I. Solution Description

Novell's Enterprise Linux Servers and combined with its Virtualization and Workload Management solutions allow data centers to deliver business services through a flexible, automated and cost-effective infrastructure.

Emulex works closely with Novell to provide customers with enterprise-class connectivity solutions that are optimized for virtualized and converged networking environments. Emulex CNAs support virtual port (VPort) technology, enabling customers to leverage SAN management tools and best practices at the VM level, including fabric zoning and LUN mapping/masking.

SUSE Linux Enterprise Server (SLES)

The SUSE Linux Enterprise Server (SLES) platform is designed for mission-critical computing across the enterprise, including physical, virtual, appliance and cloud deployments.

At the heart of SUSE® Linux Enterprise is the Linux kernel itself. Like Novell, Emulex has been a pioneer in driving innovative technologies into the Linux kernel that deliver real business value to users. One such example of these innovations is the Linux kernel's Fibre Channel Transport class. Designed, developed and maintained by Emulex, the FC Transport introduces a standardized substrate and development model for Fibre Channel and FCoE I/O devices in Linux. Further, the Fibre Channel Transport in the Linux kernel introduces a common userspace interface for the system administrator via Linux sysfs, helping to streamline and standardize operational aspects.

Another example of Emulex Linux leadership in innovation which addresses real business problems is N_Port ID Virtualization (NPIV). Emulex pioneered this technology and worked with partners to drive it into T11 for standardization. The NPIV capability was formally integrated with the Fibre Channel Transport class in the 2.6.23 version of the Linux kernel, providing a standardized model for implementing this Fibre Channel extension capability.

Emulex and Novell then added NPIV support to the Xen hypervisor which is built-in to SLES, , via Emulex drivers in SLES10/11, plus a jointly created Xen hotplug script ("block-npiv"). As a result of this work, virtual World Wide Node Names (WWNNs) and Port Names (WWPNs) can be allocated for Xen-based virtual machines (VMs). VMs can then connect to the Fibre Channel SAN using a unique virtual WWPN rather than the shared physical WWPNs of the host's HBAs. This enables the location transparent mobility necessary for Virtual Machine failover and live migration, without having to reconfigure LUN masking for physical hosts. Instead, LUNs are masked directly to each VM's NPIV port.

Since storage array events and statistics are reported by the I/O initiator's WWPN, this allows correlation of the storage vendor-collected I/O metrics with specific VMs regardless of the physical location of the VM.

Virtualization and Workload Management: PlateSpin Orchestrate

PlateSpin Orchestrate dramatically simplifies the management of a data center's heterogeneous virtual assets by controlling the entire lifecycle of each virtual machine. Resource usage is kept aligned with business requirements via built-in automation. This maximizes the value of physical and virtual data center resources, where the latter can be based on VMware, Xen or Microsoft Hyper-V.

PlateSpin Orchestrate also features distributed storage repository management for virtual machine images and data, together with constraint-based adaptive deployment of VMs to capability-matched physical servers. Administrators can create and test virtual machines, then automate deployment to suitable and available production servers.

Work is scheduled to managed servers in the form of scripted jobs. Jobs are units of work that are assigned to servers by a real-time resource scheduler that continuously evaluates available resources versus pending requests. Physical servers, like virtual machines, are considered to be resources that advertise their capabilities in the form of facts that describe the type and capacity of resource. For a physical server, example facts might include number and type of CPUs, memory, and direct-attached storage capacity. Static facts are attributes of a resource that don't change. An example of this might be a server with VT-capable CPUs. Dynamic facts can change over time, perhaps due to a physical hardware hot plug or memory ballooning of a virtual machine. Computed facts are calculated by the scheduler when referenced in a job control policy.

II. Key Benefits

Management of current virtual infrastructure is complex due to separate server-to-server (local area) and server-to-storage communication used by virtual machines accessing the network. By normalizing addressing, access control, and virtualization of the fabric (VLAN versus VSAN), converged networking will greatly simplify the management of the network infrastructure that supports the distributed applications on the virtual machine to interact with each other and with the shared storage.

A next-generation virtual infrastructure requires a converged network that can satisfy the demands of fail-over, live migration and fault tolerance for VMs and virtual storage, and be integrated by a common management of network access and QoS. Emulex CNAs provide SUSE Linux Enterprise Server customers with optimal performance, simplified management and seamless interoperability.

Converged networking benefits are:

- Consolidation of physical network infrastructure, equipment, experience, which lowers the total cost of ownership necessary to operate a next generation data center.

- Unification of management for service-oriented applications and storage hosted by virtual infrastructure, which results in improved quality of service (QoS).

III. Target Applications

CIOs are already consolidating servers, storage, and operating systems. Converged networking creates a consolidated data center fabric for the resulting virtualized infrastructures.

- Converged networking is being used to automate VM life cycle management across heterogeneous virtualization platforms and hosted workload operating systems.

- Converged networking is ideally matched to service-oriented applications hosted in VMs because it supports VM mobility with respect to all of the VM's network access requirements, for both application and storage level communication.

- Converged networking is the platform that can provide QoS for service-oriented applications hosted in VMs, in a common way, for all communication traffic.

IV. Management

The converged network solution is managed and integrated with existing IT infrastructure by providing comprehensive data center workload automation solutions. For example, by provisioning a VM together with 3rd party SAN storage, the infrastructure can be integrated through workflow automation that creates a SAN LUN, assigns a unique Fibre Channel NPIV to a VM, and configures the 3rd party SAN storage array to grant access from the VM (using its NPIV identity) to the LUN. Sophisticated automation helps break down management silos and reduces the complexity and risk of managing virtualized data centers. The result is an increased level of operational agility that creates a fluid enterprise; a data center operating environment that's aligned to the demands of business services hosted in VMs.

Novell's PlateSpin Orchestrate dramatically simplifies the management of a data center's heterogeneous virtual assets by controlling the entire lifecycle of each VM. Its grid-based extensible matching algorithms create a unique opportunity to match the specific requirements of VMs and their hosted distributed applications with new features of the converged network. This will enable even more streamlining in the automation of VM life cycle operations.

IV. Deployment Architecture

As an example of deployment architecture, consider deploying a virtual machine into a data center production server pool. The VM requires two VTx-enabled CPUs, 512Mb of direct attached OS image storage, 1Gb of memory, and gigabit Ethernet connectivity for access to a converged Fibre Channel SAN. These requirements are expressed as a set of deployment constraints—references to facts which are matched to available resources by Orchestrate, when scheduling the VM for deployment. Sophisticated resource allocation is made possible when combining a number of constraints into policy statements that are applied to groups of resources—matching supply with demand. The deployment of a virtual machine to a physical server is therefore unified by a general purpose (grid-based) algorithm for assigning units of work to available resources, in a manner that's respectful of competing work and shared capacity.

An example of this grid-based flexibility is PlateSpin Orchestrate 2.0's management of Fibre Channel NPIV in SLES. PlateSpin Orchestrate constrains Xen VMs with assigned NPIVs to VM hosts, via policy that requires SAN attachment as a discovered attribute of VM hosts, matched to only the VMs that declare a NPIV port.

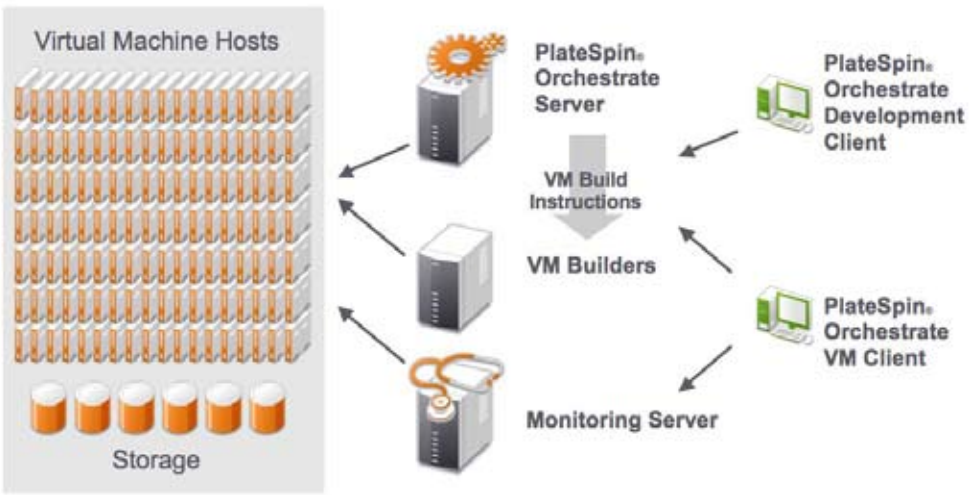

Figure 27: PlateSpin Orchestrate enables new features of the converged network

Oracle

I. Solution Description

Oracle Enterprise Linux (OEL)

Oracle Unbreakable Linux delivers enterprise-class support for Linux with premier back ports, comprehensive management, cluster software, indemnification, testing and more, all at significantly lower cost. Oracle is committed to delivering high quality, comprehensive, and integrated support solutions to help ensure that organizations succeed with the Linux operating system.

OEL is based on and fully compatible with Red Hat Enterprise Linux. It includes the exact set of packages at the same version levels with the same source code as the Red Hat distributions.

Oracle VM

Oracle VM is server virtualization software that fully supports both Oracle and non-Oracle applications. Oracle VM provides an easy-to-use graphical interface for creating and managing virtual server pools running on x86 and x86-64-based systems across an enterprise. Both Linux and Windows guests are supported. Emulex drivers for Linux are installed in the base OEL operating system to enable the seamless availability of networked storage to all of the virtual machines.

Backed by Oracle's world-class support organization, customers now have a single point of enterprise-class support for their entire virtualization environments, including the Linux operating system, Oracle Database, Fusion Middleware, and application software. All Oracle products are certified with Oracle VM, which is available for free download.

The Emulex family of LightPulse FCoE CNAs is supported for use with OEL and Oracle VM. As a member of the Oracle Partner Network and one of the first vendors to support the Oracle Unbreakable Linux Program, Emulex is a leading I/O connectivity provider for Oracle environments.

II. Key Benefits

With OEL, Oracle VM and Emulex CNAs, data centers can benefit from efficient, robust and high-performance connectivity, while protecting IT investment in existing LAN and SAN infrastructure. Oracle and Emulex make it easier to plan highly available, scalable and reliable converged mission-critical environments.

III. Target Applications

Traditionally, data centers have used specialized networks to meet individual I/O connectivity requirements for networking, storage and clustering. With increasing deployments of blades and server virtualization, data centers are inherently facing huge network sprawl, primarily driven by deployments of multiple 1Gb Ethernet links and the need for external storage connectivity. The answer is convergence of networking, storage and clustering over a single network.

"From the IT Administrator's point of view, one of the key benefits of network convergence is reduced cables. Today, Oracle customers are focusing more on clustering. But when a system within a cluster misbehaves, we have to be able to shut that system down because it could wake up again suddenly and do an I/O to storage which would then cause data corruption. With fewer cables, we can disable one port and isolate that one machine from the cluster, from the network, and from the storage."—Wim Coekaerts, VP of Linux Engineering, Oracle Corporation.

PANDUIT

I. Solution Description

PANDUIT provides complete 10Gig solutions for converged networks including 10Gig fiber connectivity, 10Gig copper cable assemblies, cabinets, cable management, routing products, as well as bonding and grounding products that are integrated into a seamless whole. These systems intelligently map physical layer solutions onto logical network architectures, and are based on Unified Physical Infrastructure (UPI) principles to achieve operational efficiencies and mitigate risk across converged infrastructure systems.

Point-to-Point vs. Structured Cabling Approach

Point-to-point cabling, common in many SAN environments, is the most available and cost-effective 10Gig solution. This solution also exhibits the lowest channel loss because there are no physical interconnect points in the channel. The main disadvantage of this approach is that it is an "unstructured" cabling model which makes incremental and organic growth and infrastructure change management difficult. A structured cabling approach with modular pre-terminated systems provides a migration path for unhindered network growth and logical organization of the cable plant, and reduces the time required for planned or unplanned moves, additions, and changes.

Connector Interface

The SFP+ connector interface offers significant cost and connector compatibility advantages for 10Gig systems. With the SFP+ interface, different reach fiber cable/optical transceivers and copper cable assemblies are freely interchangeable. SFP+ optical modules require a lower number of components, making it a cost-effective transceiver form factor. The optical connector is an LC, which enables reuse of the existing fiber infrastructure in the data center. The SFP+ direct attach copper module is a simple passive pass-through interface that provides a low cost, low power option for short reach links (less than 10m).

Generally, SFP+ direct attach copper cable assemblies are used for servers to access switch connections within and across adjacent racks, and SFP+ fiber assemblies are used for access switch to aggregation and core switch links.

Product	PANDUIT Part #
Copper Cable Assemblies	PSF1PXA1MBU (1m) PSF1PXA2MBU (2m) PSF1PXA3MBU (3m)
Opti-Core® 10Gig OM3 Fiber Optic Trunk Cable Assemblies	FSPX1211F100A
Opticom® QuickNet™ Fiber Enclosures	FCE1U, FRME4
QuickNet™ MTP Fiber Optic Cassette	FCX-12-10Y
OM3 Fiber Patch Cords	FXE10-10M1Y
Net-Access™ Cabinet	CS1
StructuredGround™ Grounding Kits	RGCBNJ660P22, RGESD2B-1

Table 9: Key PANDUIT products for enabling top-of-rack I/O consolidation over 10Gig physical infrastructure

II. Key Benefits

The SFP+ form factor offers a cost-effective approach to 10Gig connectivity which provides an interchangeable interface for either fiber or copper cabling solutions. PANDUIT offers an integrated 10Gig SFP+ connectivity solution that seamlessly integrates all physical infrastructure components which optimizes thermal characteristics and cable management for quick, effective, and reliable deployment that effectively delivers the lowest total cost of ownership.

Within the rack, direct attach copper provides a simple to use point-to-point low cost, low latency, low power connection. Constructed of a robust mechanical design and with substantial electrical margin through advanced internal design to maintain high signal integrity, PANDUIT copper connectors provide a highly reliable and durable solution for intra-rack server connections.

Uplink solutions from PANDUIT present the minimum cabling cross-sectional area which reduces vertical and horizontal pathway space. The main benefit for the customer is a modular array-based connectivity solution that allows easy access and connection to storage systems. MPO/MTP fiber connectors are sized to match the host transceiver groups configured in base 2 – 4, 8 and 16 fibers for both SAN and converged Ethernet environments. This modularity, which is native to newer host system design philosophy, allows the customer to scale host capacity by adding MTP trunk assemblies in the appropriate base fiber increments.

PANDUIT cabinets optimize thermal performance and facilitate airflow in high-density environments through the use of passive cooling systems; these systems leverage active equipment exhaust fans to move air through in-cabinet ducts, floor grommets, and overhead vertical exhaust vents to efficiently mitigate the risks of overheating.

III. Target Applications

10Gig cabling solutions are commonly deployed as server interconnects, inter-switch links, and backbone cabling for Ethernet channels in the data center. These solutions are also being deployed in pure Fibre Channel applications (switch port to HBA) as backbones and down to the equipment distribution level. Other applications include:

- Data center server consolidation
- Top-of-rack and end-of-row I/O consolidation deployments
- FCoE and data center ethernet
- Storage Area Networks (SANs), Network Attached Storage (NAS)

IV. Deployment Architecture

Fiber connectivity and media systems are agnostic with respect to transmission protocols (FC, FCOE and Ethernet). Fiber infrastructure solutions based on structured cabling in modern data centers typically are managed, with respect to moves, additions and changes, in a main distribution area (MDA). Port mapping of director electronics and correspondence to hosts is accomplished through intuitive port configuration and identification as provided by modular PANDUIT® QUICKNET™ fiber cassette solutions in the MDA. The SFP+ converged network solution will seamlessly integrate into the existing fiber cable plant.

The migration to converged networks presents minimal disruption or change in design for the fiber cable plant. Connectivity systems on converged electronics are identical to those used in non-converged systems and the fiber cable media present in these structured cabling solutions remains unchanged.

As shown in Figure 23, direct attach copper cable assemblies (CX1) are used in the equipment rack to connect servers equipped with CNA to the top-of-rack (ToR) access switch. Typically there will be redundant Network A & B connections between the server and top of rack switch, and cables are offered in two colors to identify the different networks.

Uplinks on the top of rack switch require LC patch cords to connect to a modular LC fiber cassette at the top of the rack. The equipment rack connects with the Main Distribution Area (MDA) via an MTP to MTP flat ribbon cable assembly. The MDA would contain MTP to LC fiber cassettes for cross connection purposes and the cross connections would be made complete with LC-LC patch cords.

The MDA then connects to the LAN or SAN switch rack via an MTP to MTP flat ribbon cable assembly. The LAN or SAN switch rack will house an MPO/MTP fiber adapter panel where a female MTP flat ribbon will connect to a male MTP to LC Hydra cable assembly. The LC side of the Hydra cable assembly will complete the path by connecting with the appropriate (mapped) switch ports.

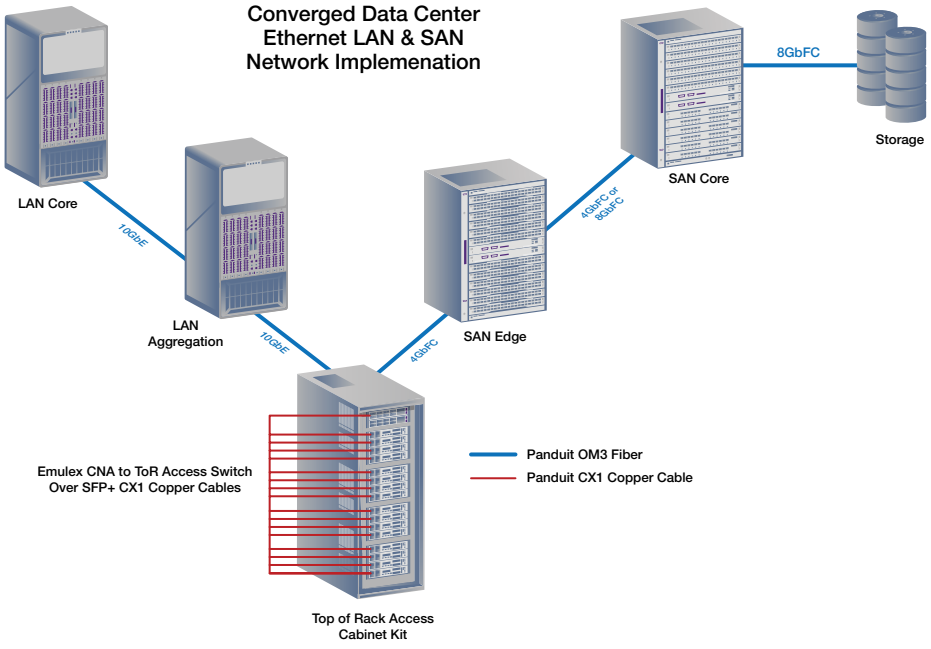

Figure 28: PANDUIT Modular Fiber Solutions and CX1 Copper Cables for Top-of-Rack Connectivity

Scalent Systems Inc.,

I. Solution Description

Scalent enables enterprises and service providers to create a dynamic data center. Scalent V/OE software allows data center operators to change entire systems and associated topologies rapidly, including which servers are running, what software is running on them, and how they're connected to networks and storage. Scalent does this in real-time, without requiring physical changes to the infrastructure. Scalent V/OE software enables customers to take their data centers from "dead bare-metal" (powered-off hardware with no software installed) to live (newly deployed or restored), connected (to network and storage) servers, in five minutes or less.

How is the Scalent V/OE Solution Different?

Currently, data centers are managed using tools that are purpose-built for specific hardware and software components in the infrastructure. Current tools do not provide the capabilities required for integrated management of legacy and converged network environment, and the seamless movement of workloads among resources in these disparate environments.

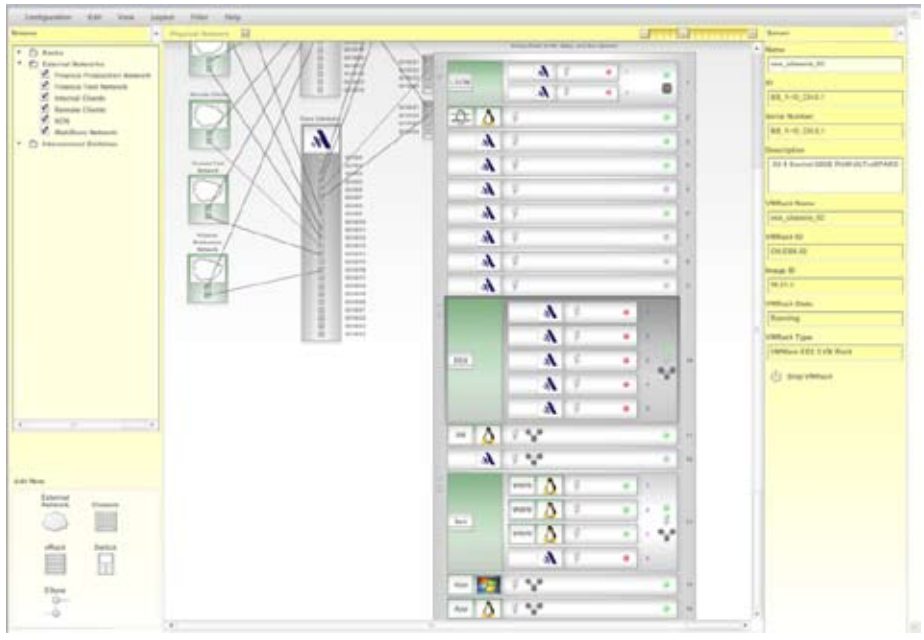

Figure 29: Scalent management application

II. Key Benefits

Lower Costs

Companies using Scalent V/OE require significantly less investment in initial hardware capital expenditure and in ongoing operational costs. By improving utilization and simplifying repurposing or reconfiguration, less hardware and administration is necessary to maintain and change your data center—and operational costs fall accordingly. Data Center infrastructure management is automated, and systems maintain performance, in a simpler, more scalable way.

Higher Responsiveness

Scalent V/OE enables real-time repurposing of existing infrastructure. With Scalent, a library of images is maintained on a central storage point, enabling physical machines to be rapidly repurposed for different business system servers, on the fly, without any physical network changes. Entire farms of servers, with associated network and storage access topologies, can be automatically changed in minutes, with less interdependence on other functional groups, and without worries about whether the underlying infrastructure is physical or virtual: P2V and V2P transitions are instantaneous.

Improved (Simpler) Processes

By using the data center abstraction provided by Scalent V/OE, stored templates, and APIs (instead of physical cable and machine movement) companies can shorten their response, reconfiguration, and deployment times—while also reducing the chance of inadvertent errors in topology recreation. With Scalent, server farm functional reality matches design, every time, consistently.

III. Target Applications

Scalent V/OE offers integrated management of the data center infrastructure, both traditional and converged. Workloads can be moved seamlessly among resources in either environment. The initial applications for network convergence in combination with Scalent V/OE include:

1. Development and test infrastructure automation
2. Multi-purpose (or multi-use) data centers
3. Disaster recovery automation
4. Automated capacity management for application grids or clusters
5. Unified management of virtualized and physical environments

VMware

I. Solution Description

The VMware Virtual Infrastructure 3 (VI3) product set offers server virtualization technologies that help optimize resource utilization and simplify systems management for the toughest enterprise data centers. VMware VI3 is the market-leading and most stable virtualization solution available today, enabling customers to achieve new heights of availability and optimization of resource pools with clustering multiple ESX servers in to a single VMware cluster. Only VMware enables non-disruptive migration of VMs across ESX servers as well as moving virtial machines between storage targets.

Network Connectivity with CNA

The key virtual networking components in a VMware Infrastructure environment are virtual Ethernet adapters and virtual switches. A virtual machine can be configured with one or more virtual Ethernet adapters. Virtual switches allow virtual machines on the same VMware ESX server to communicate with each other using the same protocols that would be used over physical switches, without the need for additional hardware. These virtual switches also support VLANS that are compatible with standard VLAN implementations from other switch vendors.

With CNAs providing identical networking functionality as that of standard NICs, there are no additional support requirements for ESX servers.

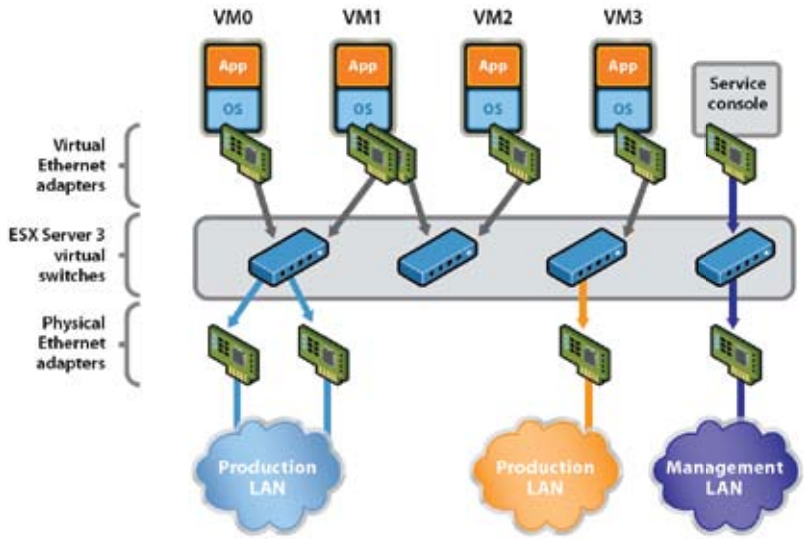

Figure 30: Network connectivity with CNA in ESX environment

Storage Connectivity with CNA

VMware Infrastructure storage architecture provides a layer of abstraction that hides and manages the complexity of and differences between physical storage subsystems. To the applications and guest operating systems inside each virtual machine, storage is presented simply as SCSI disks connected to a virtual BusLogic or LSI SCSI HBA.

The virtual SCSI disks inside the virtual machines are provisioned from datastore elements in the data center. A datastore is like a storage appliance that serves up storage space for virtual disks inside the virtual machines, and stores the virtual machine definitions themselves. Each virtual machine is stored as a set of files in its own directory in the datastore. A virtual disk (vmdk) is a file that resides in a datastore that is managed by ESX. A datastore will reside on a VMFS volume for block-based storage or a mount-point for NFS storage. The VMFS volume is typically comprised of a single LUN, but can span several LUNs. A virtual disk can be easily manipulated—i.e. copied, moved, backed-up—just like any other file. For guest OS's that support hot-adds of new disks, a new virtual disk can be added without having to shut down the VM.

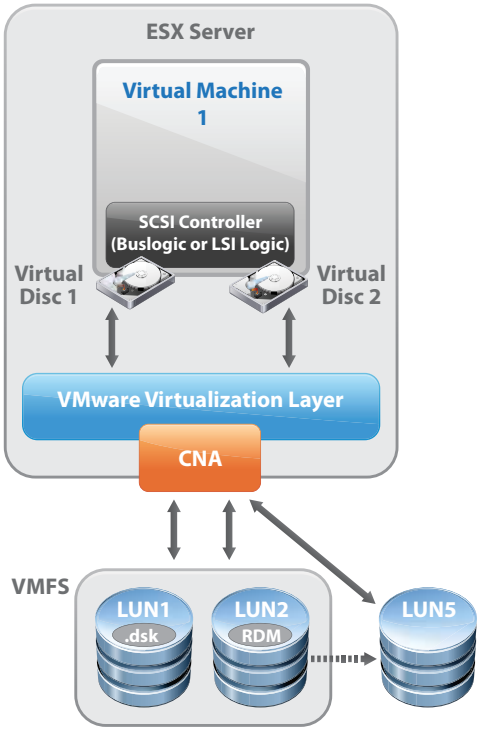

Figure 30: Storage connectivity with CNA in ESX environment

The datastore provides a simple model to allocate storage space for the virtual machines without exposing them to the complexity of the underlying storage technologies. A Fibre Channel SAN-based datastore is the most common deployment as it enables

VMotion, ESX boot from SAN, support for raw device mapping, and support high availability clusters (such as Microsoft MSCS). iSCSI and NFS protocols are also supported and support for VMotion, resource pools, VMware HA and DRS.

With release 3.5, VMware has added additional support for converged network solutions by providing support for both NPIV and FCOE. With this capability, a converged network solution such as the Emulex CNA can now provide transparent connectivity to the existing Fibre Channel storage.

II. Key Benefits

- The flexibility of a converged network in providing LAN and SAN connectivity provides enormous advantages in server virtualization environments.

- Network convergence technologies enable cost reduction by:
  - lowering the number of cables needed to support storage and network configurations.
  - increasing agility to adjust quality of service or resources to meet SLAs.
  - optimizing management of connections by lowering complexity.

- The Use of FCoE and NPIV with an ESX server provide additional control of your converged network.

III. Target Applications

The combination of VMware Virtual Infrastructure and network convergence enables data centers to deal with a large number of use cases and thus facilitate an adaptive next generation data center.

Datacenters that are migrating to virtualization are well-positioned to benefit significantly by also consolidating their storage and network traffic with FCoE. VMware VI3 supports this today and will be enhancing this with future releases as well.

8 Key Applications for Network Convergence

Server Virtualization

A converged network fully complements server virtualization and enables the roll-out of on-demand services where applications and network services are provisioned dynamically. The key advantage with FCoE is that it extends consolidation benefits of infrastructure simplification and lower costs without disrupting existing architectures and operating models.

FCoE-enabled network convergence significantly reduces the number of cables, switch ports and adapters as shown in the picture below –

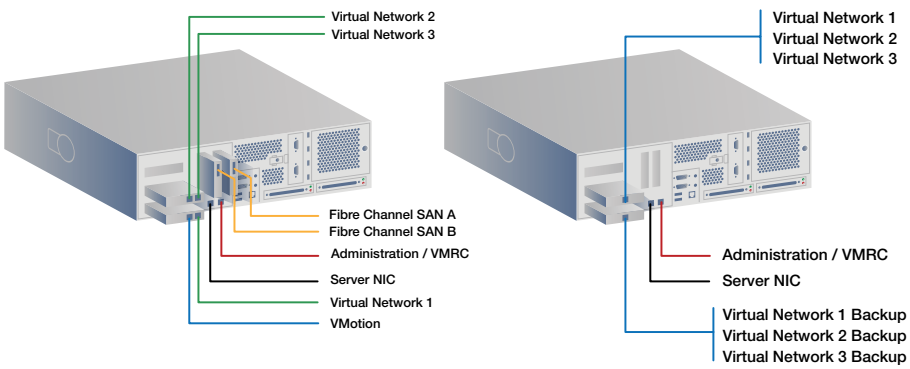

Figure 31: Significant cable reduction from deploying CNA in a rack-mount server running ESX

Although a CNA is a single physical entity, the adapter is represented to an operating system as a HBA and a NIC. This level of transparency to the operating system ensures that the CNA configuration and management practices are performed in the same manner as current standalone NIC and HBA devices.

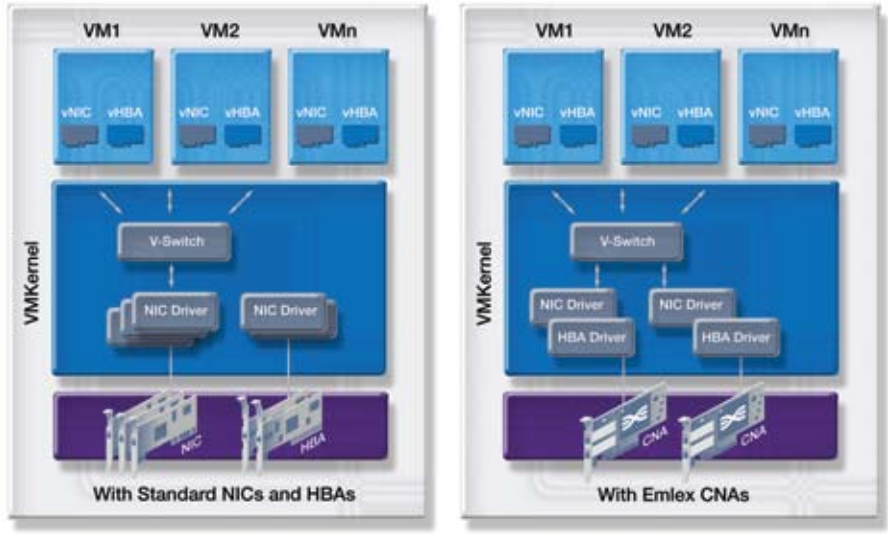

Figure 33: The hypervisor/VMKernel views the CNA as an HBA and a NIC

The Emulex CNA supports Virtual HBA technology, which provides the scalability necessary to support running more applications on virtualized servers. Emulex LightPulse Virtual HBA technology allows end users to effectively virtualize the CNA functionality so each virtual machine running on a server can have independent access to its own protected storage. Customers can also leverage SAN management tools and best practices, such as fabric zoning and LUN mapping/masking based on virtual machines.

In addition, Emulex CNA supports Message-Signaled Interrupts eXtended (MSI-X), an interrupt handling mechanism in the PCI standard for today's multi-core and multi-processor servers. MSI-X enables improved overall performance, reduced host CPU utilization and lower latencies. The overall CPU efficiency achieved by Emulex CNAs allow for more virtual machines to be accommodated per server.

Blade Servers

According to IDC, blade servers are the fastest-growing segment of the server market with 1.1 million units expected to ship in 2008, growing to 2.6 million units in 2011. By 2010, it is expected that one in four new servers will be a blade.

Blade servers have enabled organizations to tackle server sprawl, simplify manageability and improve availability of the server infrastructure. Besides consuming less space and less power than the equivalent number of rack-mount servers, some of the key advantages of blade servers include expandability, increased redundancy, greater integration of data and storage networking and flexibility of blade platform configurations.

Networked storage enables diskless blades provides, which added advantages such as increased reliability, lower total cost of ownership and acquisition cost. The use of server virtualization in blade servers also increases the need for external storage connectivity since external shared storage forms the key component for virtual server mobility.

Although the demand for external storage connectivity drastically increases the number of network connections in a rack server integrated switching modules with in the blade chassis enables internal port aggregation for Ethernet and Fibre Channel. This reduces the number of cables and simplifies management of external storage connectivity.

While a blade chassis optimizes network connectivity by aggregating the network ports and lowering the number of cables, the use of multiple networks to carry different types of traffic creates the need for higher level of network consolidation that fully complements the consolidation driven by blade server technology.

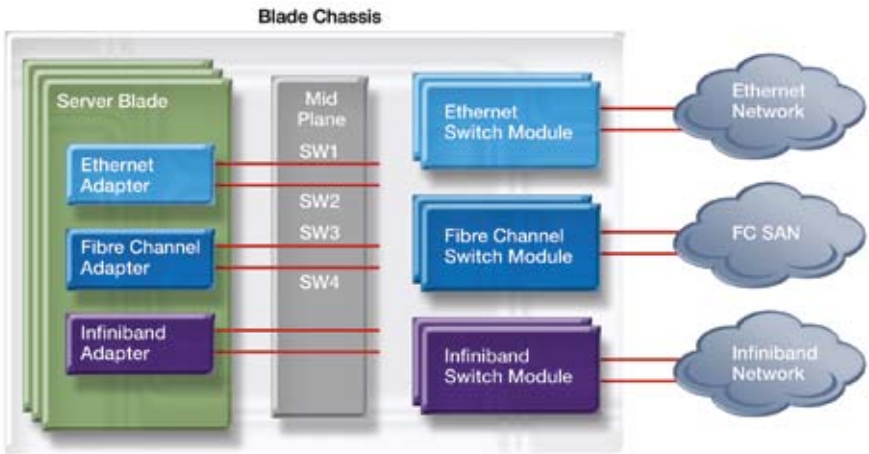

Figure 34: Blade deployments with multiple network fabrics

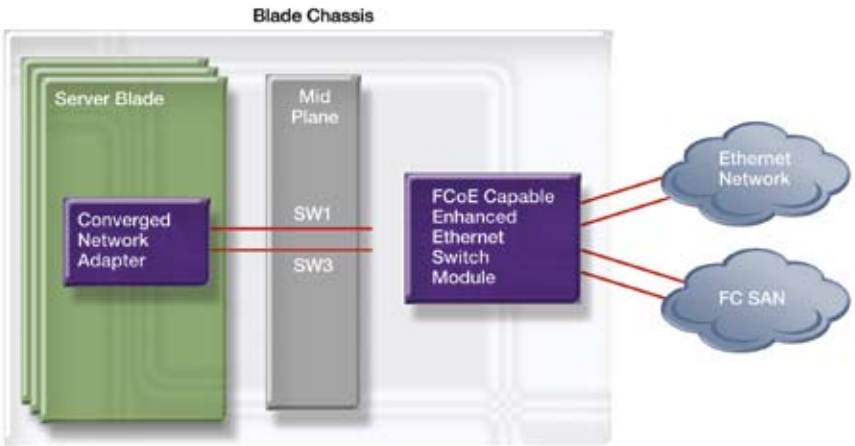

Figure 35: CNAs reduce the number of adapters required in server blades

Supported by the wire-once architecture, IT managers can deploy a standard blade configuration that can support Ethernet and storage connectivity as needed. This reduces server platform diversity for better control of capital and operating expenses.

Boot from SAN

With Boot from SAN, server OS is installed on a logical drive (LUN) that resides on external storage. The deployment of FCoE-enabled network convergence enables a wide range of servers to leverage this functionality and increase the efficiency of operating system image management.

For the Emulex CNAs, the boot from SAN process is identical to the proven boot from SAN processes followed for the Emulex HBA product lines.

By extending the server system boot BIOS, boot from SAN functionality is provided by the boot BIOS contained on the Emulex CNA. When properly configured, the CNA directs the server to boot from a LUN on the SAN as if it were a local disk.

Emulex provides the following types of boot code:

- x86 BootBIOS works with x64 and x86 systems.
- OpenBoot works with Solaris SPARC systems, Linux PowerPC systems and AIX systems.
- EFIBoot works with Intel Itanium 64-bit systems using the EFI (Extensible Firmware Interface) Shell.

Virtual Desktops

Virtual desktop infrastructure (VDI) is an extremely compelling technology because of its potential to reduce desktop costs, improve manageability and enhance security. VDI is best-suited for knowledge workers and/or organizations that have large, homogenized desktop environments in sectors such as government and educational institutions, retail, call centers, and help desks. High bandwidth connectivity is a key requirement for the virtual desktop deployment.

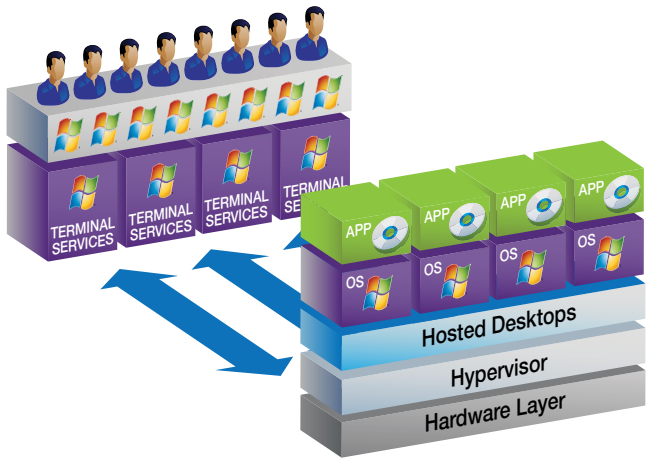

Figure 36: Virtual desktop with server virtualization

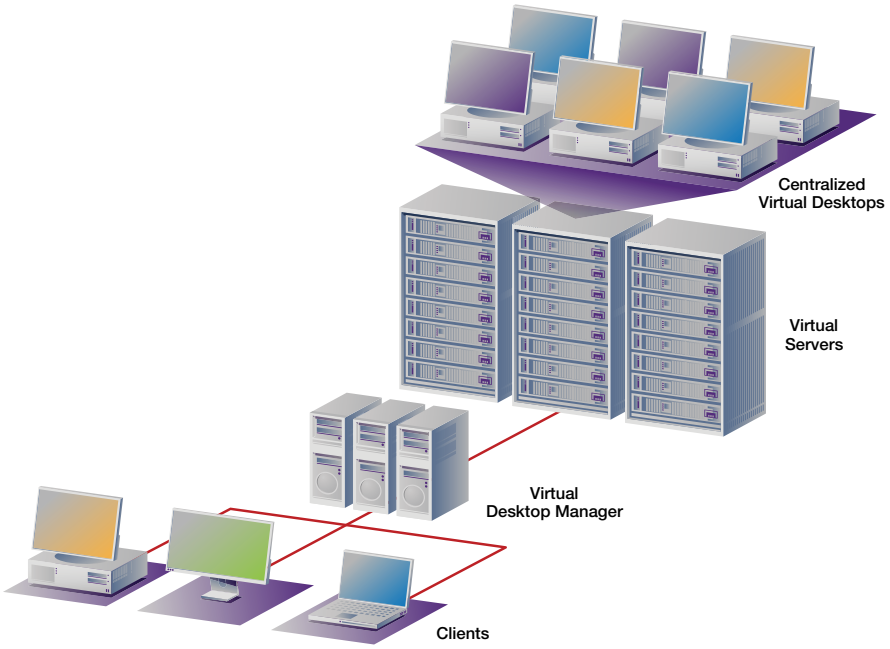

Figure 37: Virtual desktop Improves security and manageability

VDI Benefits include:

- Standard desktop OS environment
- Provisioned on-demand
- Customizable on per-user basis
- Fully isolated and secure

VDI is a relatively new technology that is likely to see increasing adoption and more production roll-outs in 2009 and beyond. As user experience and cost containment improve and more service offerings are commercially available, VDI deployments will gain broader adoption.

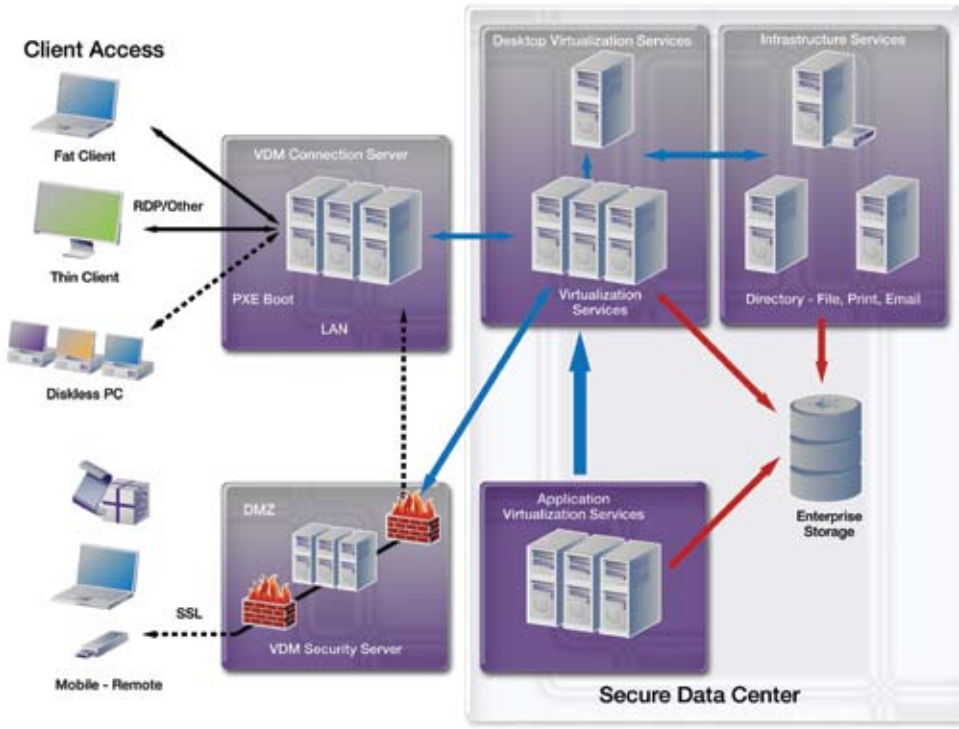

Figure 38: Network convergence simplifies virtual desktop architecture

Regulatory Compliance

In regulated businesses such as financial and health care, enterprises are under constant pressure to ensure that their internal networks, applications and related data are always available and highly secure. Some regulations also require active monitoring and reporting of network events.

Figure 39: With converged networks, fewer nodes are required for compliance monitoring and managing.

Internal business initiatives and external regulations are constantly adding to these challenges and are testing the capabilities of status quo networks. Although IT managers could continue to deploy multiple networks and ensure compliance, the process gets tedious with the changing dynamics of SAN expansion driven by virtual servers and blade servers. A simplified approach to networking provides competitive advantages in the face of new business initiatives and helps meet regulatory compliance obligations.

9 Network Convergence Deployment Phases

The deployment of converged networks in data centers is expected to happen in multiple phases over the next few years. The key deployment phases are outlined in this section of the guide.

Phase-1: Pilot Deployment Phase

This initial phase began in mid-2008 with the general availability of FCoE products. The bulk of customer installations during this phase are expected to be pilot deployments, primarily due to lack of OEM and operating system qualifications, and standards maturity. Concurrent with this early availability, the FCoE standard is being completed in INCITS and the necessary standards for the prerequisite Ethernet enhancements are being completed in the Ethernet Alliance.

These early market deployments are primarily focused on server I/O consolidation. With this focus, server administrators will be working through the transition from managing distinct NICs and HBAs to managing Converged Network Adapters. With the deployment of unified top-of-rack switches, role-based administration will support the management of SAN resources by the storage administrator, and LAN resources by the LAN administrator.

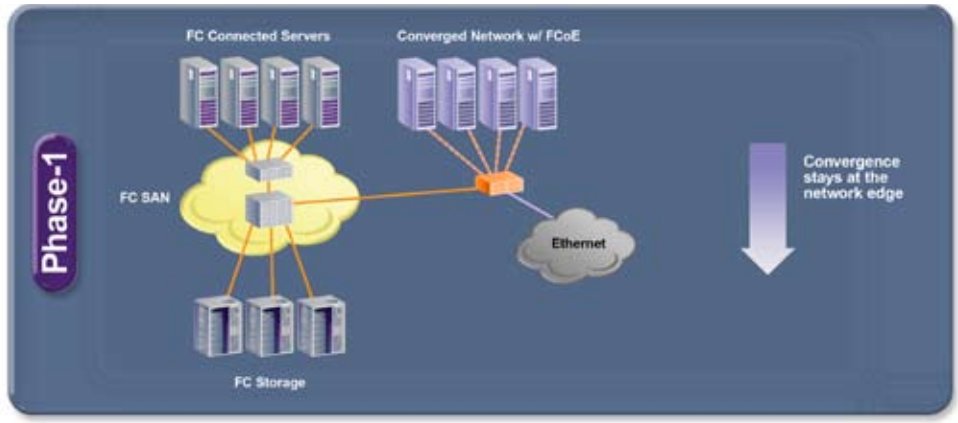

Figure 40: Phase-1 of converged network deployment

Phase-2: Volume Deployment Phase

This phase is expected for new data center deployments in late 2009 and early 2010. Organizational dynamics will likely keep deployments focused on server I/O consolidation. Blade server deployments will begin as single-chip CNA mezzanine cards become available. Low cost interconnect will continue to be based on SFP+ Direct Attach and top-of-rack or middle-of-row switching architectures. As third generation 10GBASE-T transceivers become available, this will enable end-of-row and other more centralized designs where low-cost cabling beyond 7 meters is enabled.

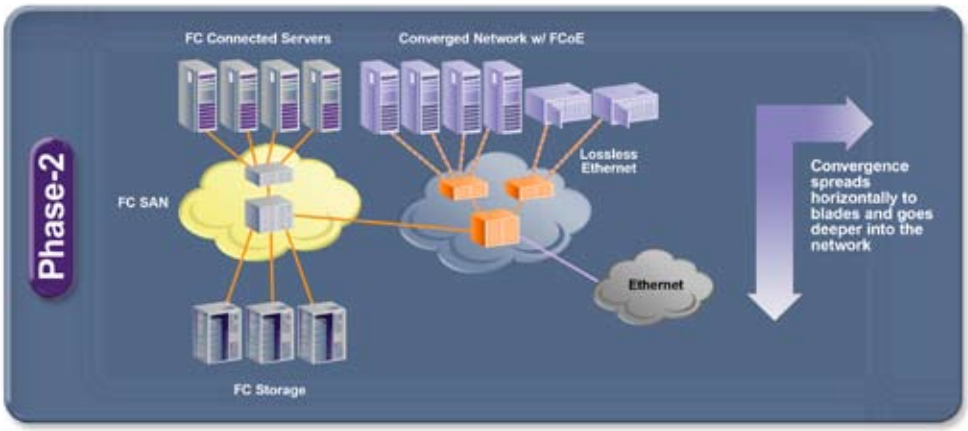

Figure 41: Phase-2 of converged network deployment

Phase-3: End-to-End Convergence Phase

This phase is expected for new data center deployments in 2011 and beyond. As organizations evolve, the processes will be developed to optimize the management of unified fabric, such as consolidated change control, fault isolation of virtual Fibre Channel resources and traffic engineering. A unified fabric also requires additional protocol support, including end-to-end congestion management being developed in the IEEE 802.1Qau project, and upgrades to spanning tree being developed in the IETF with the TRILL project.

In this phase, storage devices natively running FCoE are expected to make their presence. Organizational dynamics will likely keep deployments focused on server I/O consolidation.

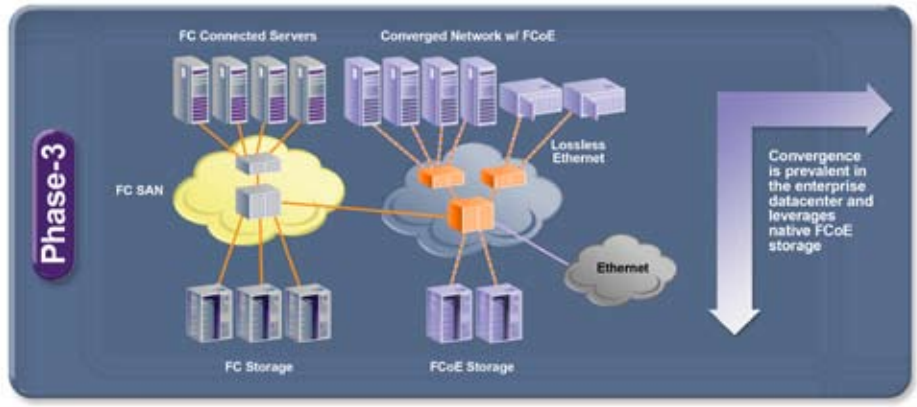

Figure 42: Phase-3 of converged network deployment

Appendix A: Introduction to Emulex LP21000

Emulex is one of the leading proponents and developers of FCoE technology. The Emulex LP2100 series CNAs are designed to support high performance storage and networking applications in enterprise data centers.

Product Description

The Emulex LP21000 is a multi-protocol PCI-Express 1.0a Converged Network Adapter (CNA) designed for enhanced 10Gb/s Ethernet fabrics performing both traditional networking interface card (NIC) functionality and Fibre Channel HBA functionality using FCoE technology.

The Emulex LP21000 FCoE CNA is comprised of four components, the 10GB/s network interface controller, Fibre Channel I/O controller, FCoE Controller and the PCI express switch. The Fibre Channel controller leverages eight generations of field-proven Emulex LightPulse® technology to deliver seamless interoperability and exceptional reliability in data center environments.

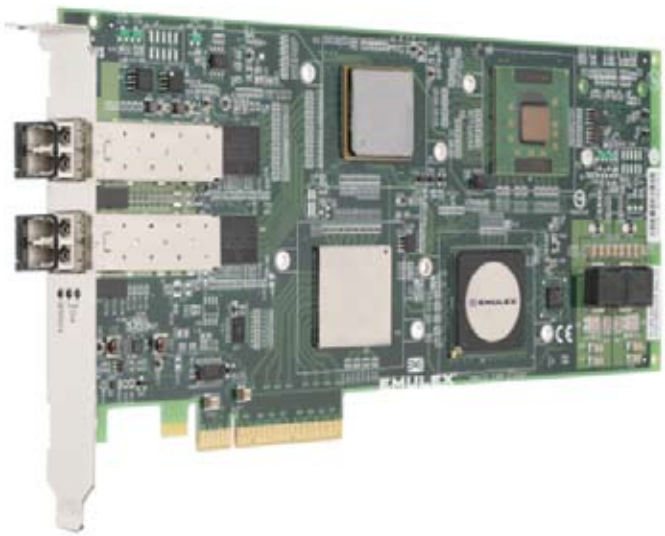

Figure A1: Emulex LP21002 CNA

Key Features

- Single and dual port 10Gb/s Ethernet configurations with optical and copper interfaces:
  - LP21000-M for single port – Optics
  - LP21002-M for dual port – Optics
  - LP21000-C for single port – Copper
  - LP21002-C for dual port – Copper
- Single 8-lane standard Gen 1, PCIe 1.0a interface shared by the Fibre Channel and MAC controllers through an IDT 3-port PCIe switch.
- Standard PCIe 1.0a form factor PCBA; full height, 9.1" long.
- On-board flash memory: 8MB for Fibre Channel controller program storage, 1MB for FCoE program storage, and 128Kb for PXE Ethernet boot.
- On-board configuration EEPROM for the Ethernet MAC and FCoE bridge.
- Common Fibre Channel driver for use on LightPulse CNAs and HBAs.
- Interoperates with existing Fibre Channel SAN infrastructures - switches, arrays, SRM tools (including Emulex utilities).
- Unified Ethernet to FC SAN connectivity provided by the FCoE switch.
- Support for MSI-X, jumbo frames (16K bytes), VLAN tagging (802.1Q, PER priority pause / priority flow control, and advance packet filtering.
- No host operating system changes required - i.e. NIC and HBA functionality (including device management, utilities, etc.) transparent to the host operating system.

Unique Product Advantages

- Eight generations of Emulex Fibre Channel controller technology and eight generations of Fibre Channel controller innovations, and proven LightPulse drivers.
- Interoperability with a wide range of FC storage and switching equipment.
- Use of Emulex HBAnyware management tool across both CNAs and HBAs.
- High performance and low CPU utilization by offloading significant portions of FC processing to the CNA.

- Boot from FC SAN.
- Supported with a wide range of operating systems.

Functional Architecture

The LP21000 appears as two PCI devices to the host — a network adapter and Fibre Channel adapter. Host networking and storage drivers communicate with the appropriate PCI function in the LP21000. The I/O operations are routed to the appropriate controller for processing. Networking transactions are delivered to the lossless NIC, and forwarded on to the unified fabric.

For Fibre Channel transactions, the Fibre Channel frames are sent to the FCoE encapsulation engine, and then transferred to the lossless NIC for delivery to the unified fabric. Received traffic is handled much the same way. Incoming traffic is processed by the lossless NIC, then filtered based on FCoE ETHERTYPE (type is specified in the Ethernet header). Non-FCoE traffic is delivered to the Ethernet NIC, and decapsulated FCoE traffic is then forwarded to the Fibre Channel IOC for further processing. The Fibre Channel blocks are then relayed to the appropriate host device driver.

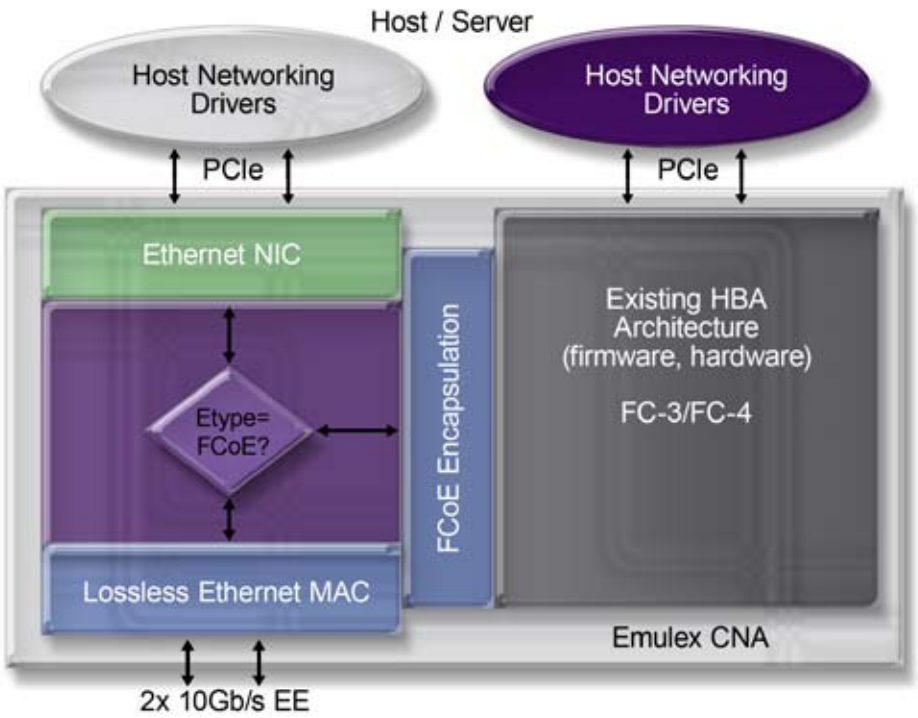

Figure A2: Emulex LP21000 functional architecture

Device Manager View in the Operating System

Since the Emulex LP21000 offerings provide Fibre Channel and NIC connectivity, Window Device Manager will show 10Gb NIC and 4Gb/s Fibre Channel devices in the device tree. This level of transparency ensures minimum impact management processes.

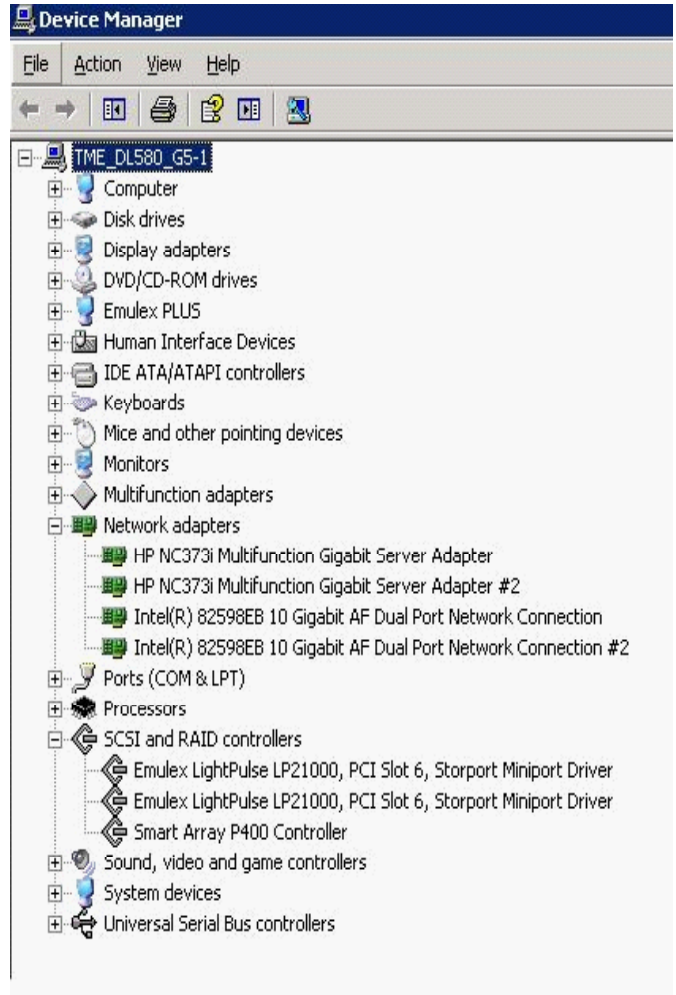

Figure A3: View of LP21002 in Windows Device Manager

Software Configuration Parameters

The enhanced Ethernet specific parameters for the LP21000 offerings can be configured with Emulex's HBAnyware management application.

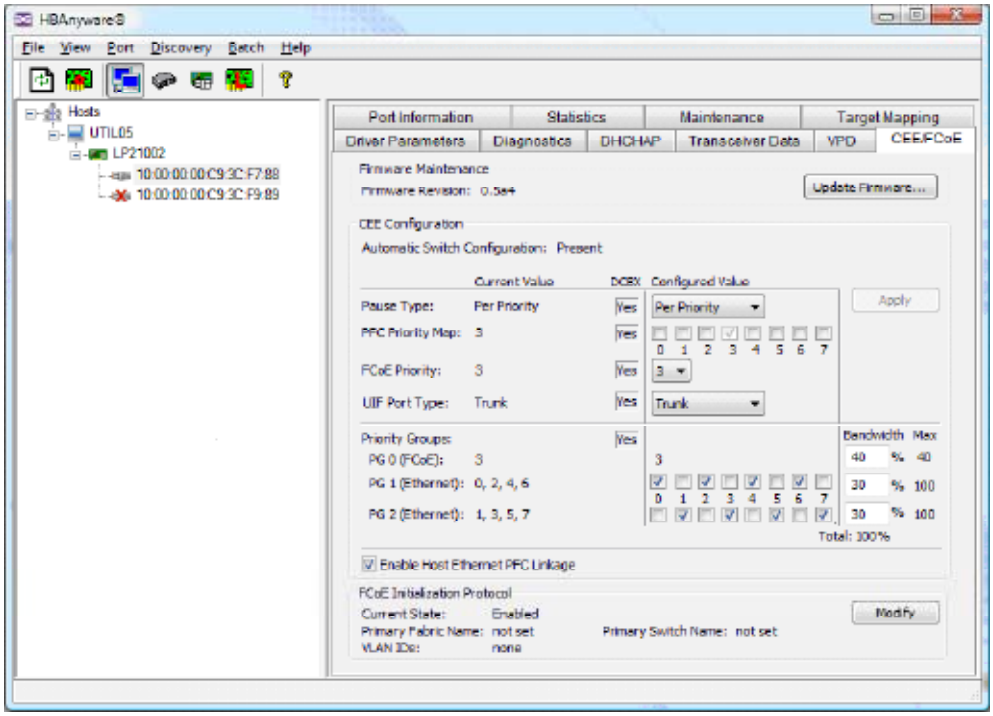

Figure A4: HBAnyware configures the traffic priorities and per Priority flow control.

Automatic Switch Configuration

This is a non-configurable field which displays whether the Automatic Switch Configuration feature is present on the attached CEE/FCoE switch.

Values displayed by this parameter include:

- Present
- Not Present

Pause Type

This is a configurable parameter. The drop-down menu may be used to select the Ethernet flow control type- standard PAUSE flow control or per priority-based PAUSE flow control. With per priority-based flow control means the Ethernet network will be seen as 8 virtual lanes (or priorities) of traffic rather than one.

Possible values:

- Standard
- Per Priority

FCoE Priority

This is a drop-down menu of the available values for the FCoE priority parameter. The FCoE traffic from the CNA will be transmitted with the assigned priority number.

Possible values: 0, 1, 2, 3, 4, 5, 6, 7

PFC Priority Map

This is a series of eight checkboxes to select values for various flow control priorities. The priority assigned for FCoE is selected by default. The other priorities can be individually selected.

Possible values: 0, 1, 2, 3, 4, 5, 6, 7

Priority Groups

LP21000 supports three priority groups (PGs) and bandwidth can be configured for these PGs. Ethernet enables an end point to define eight traffic priorities (this information is captured using 3-bits in the 802.1Q tag also known as VLAN Tag).

Priority groups is a new concept introduced in lossless Ethernet which enables a group of priorities to be bound together for bandwidth allocation. For example, priorities 0, 1 and 2 may be assigned to PG-0; priority-3 may be assigned to PG-1; and priorities 4,5,6,7 may be assigned to PG-15. This enables simplification of bandwidth assignments.

The priority groups are used to carry the three different traffic types typically seen in the data center networks:

- Low latency traffic (for clustering)
- Lossless traffic (for storage)
- Best-effort (for standard networking traffic)

Appendix B: FAQ

Can I replace FCIP with FCoE?
FCIP was primarily developed to provide wide area connectivity between distant Fibre Channel SANs using the routing capabilities of IP. Designed for the data center, FCoE is not routable and cannot replace FCIP.

When will FCoE be widely deployed?
The industry consensus is that FCoE volume deployment could begin sometime in 2010. Enterprises are expected to have their initial deployments analyzed and qualified during the interim timeframe.

What is the status of the standards behind FCoE and Ethernet?
The proposals for standardization of Ethernet enhancements to achieve convergence in access networks were complete as of July 2008 and FCoE standardization was ratified in June 2009. This enables enterprises to move forward with initial FCoE deployments in the edge/access networks. To enable convergence in the core network, additional Data Center Bridging (DCB) standards are required, which are likely to be ratified during the second half of 2009. The network convergence solutions that are deployment-ready today have factored in these staggered standardization timelines and are designed to support seamless field upgrades to incorporate the new DCB enhancements as they are standardized.

Emulex CNAs will be compliant with ratified standards, and any final standards changes will be accommodated in the CNA firmware updates.

Is SFP+ Twin-ax copper a standard?
The SFF Committee defined the electrical parameters for a pluggable interface supporting 8.5 and 10 Gigabit per second data rates. These SFP+ modules are typically used to support optical interfaces; however, the SFF-8431 specification also standardizes the requirements for Passive Direct Attach SFP+ Cable.

How does SFP+ direct attach copper compare with 10GBASE-T?
10GBASE-T will support much longer distances than SFP+ direct attach copper, but has limitations due to high power requirements and high latency. Customers can choose SFP+ direct attach copper for low cost, low power cabling within a rack, and optical connectivity for other topologies.

Why not use 10GBASE-T copper cabling for convergence?
The 10GBASE-T solutions could be an attractive option for relatively longer distance cable deployments. For the connectivity from the back of the server to the top of the

rack switch, the SFP+ direct attach copper solutions provide significantly low latency and low power consumption.

How are the Fibre Channel SAN management practices such as LUN masking and zoning preserved?

FCoE requires no changes to the Fibre Channel stack. Concepts such as WWN and fabric login are carried over to FCoE. With the WWN as the foundation, the concept of LUN masking (at the storage array) and zoning (at the switch) are fully preserved.

Appendix C: Related Resources

Related Standards

IEEE 802.3
http://standards.ieee.org/getieee802/802.3.html

IEEE 802.1 activities
http://www.ieee802.org/1/files/public/docs2007/new-cm-barrass-pause-proposal.pdf
http://www.ieee802.org/1/files/public/docs2007/new-cm-pelissier-enabling-block-storage-0705-v01.pdf
http://www.ieee802.org/1/files/public/docs2007/au-ko-fabric-convergence-0507.pdf
http://www.ieee802.org/1/pages/802.1au.html
http://www.ieee802.org/1/files/public/docs2008/az-wadekar-dcbcxp-overview-rev0.2.pdf

FCoE
http://www.t11.org/
http://www.fibrechannel.org/OVERVIEW/FCIA_SNW_FCoE_WP_Final.pdf
http://www.t11.org/ftp/t11/pub/fc/bb-5/08-569v1.pdf

TRILL (Transparent Interconnection of Lots of Links)
http://www.ietf.org/html.charters/trill-charter.html

Emulex Resources

Network Convergence Solution Site
- Resources from Emulex and its partners for deploying converged network in the data center
  http://www.emulex.com/solutions/convergence.html

White Papers
- Best Practices Guide: Network Convergence with Emulex LP21000 CNA & VMware® ESX Server
  http://www.emulex.com/artifacts/96fe478b-d7ed-43dd-ae39-3e8478ce69bd/vmware-esx-fcoe.pdf
- Cisco Nexus 5000 and Emulex LP21000 CNA – An FCoE Solution Overview

http://www.emulex.com/artifacts/7eea7e69-106a-42d3-a4f7-12aba7632169/nexus5000_lp21000.pdf

- Convergenomics – Network Convergence Cost Savings Calculator
  http://www.emulex.com/files/tools/FCoE-calc-rev5a.html

World Headquarters 3333 Susan Street, Costa Mesa, CA 92626 +1 714 662 5600
Wokingham, UK +44 (0) 118 977 2929 | **Munich, Germany** +49 (0) 89 97007 177
Paris, France +33 (0) 158 580 022 | **Beijing, China** +86 10 68499547