

Best Practices Guide: Network Convergence with Emulex LP21000 CNA & VMware® ESX Server



How to deploy Converged Networking with VMware
ESX Server 3.5 Using Emulex FCoE Technology



Table of Contents

Introduction	3
Drivers of Network Convergence	3
Benefits of FCoE Enabled Network Convergence	4
FCoE Technology Overview	6
Introduction to Emulex LP21000 CNA	7
Converged Network Deployment models	7
Converged Networking with FCoE in an ESX Environment	8
CNA Instantiation Model in ESX 3.5	9
Configuring the LP21000 as a Storage Adapter.....	10
Configuring the LP21000 for Multipathing in ESX	12
Configuring the LP21000 as a Network Adapter.....	14
Configuring the LP21000 for NIC Teaming in ESX	15
Scaling to Enterprise Manageability	16
Conclusion	18

Introduction

Fibre Channel over Ethernet (FCoE) technology enables convergence of storage and LAN networks in enterprise data centers. The key benefits of the technology include cost savings and improved operational efficiencies in deploying and maintaining enterprise data centers. But the technology is new and deployment configurations are less well understood compared to proven Fibre Channel solutions.

Server virtualization using VMware ESX 3.5 enables consolidation of server resources by allowing multiple virtual machines to exist on a single physical server. This level of consolidation increases the demand for I/O bandwidth and thus increases the need for multiple I/O Adapters. FCoE-based converged networking addresses the requirement for increased bandwidth while significantly simplifying the I/O infrastructure.

This best practices guide, targeted at server, storage and network administrators, is intended to present readers with information about network convergence using FCoE technology and also to provide an overview about how to extract the maximum benefits from network convergence using FCoE technology in an ESX Server environment.

Drivers of Network Convergence

Enterprises rely on their computing infrastructure to provide a broad array of services. As enterprises continue to scale their computing infrastructure to meet the demand for these services, they inherently face an infrastructure sprawl fueled by the proliferation of servers, networks and storage. This infrastructure growth typically results in increasing costs, real estate, cooling requirements, management complexity and thereby reducing the overall efficiency of operations.

In order to overcome these limitations, the industry is moving towards empowering data centers with technologies that enable consolidation. Server virtualization and blade server technologies are helping address the consolidation of physical servers and rising rack densities in the data center. A related trend associated with server virtualization and blade server deployments is the expanded use of higher I/O bandwidth and adoption of external storage which enable consolidation and mobility.

Since the late 1990s data centers have maintained two sets of networks – a Fibre Channel Storage Area Network for storage I/O traffic (SAN) and a local area network (LAN) for data networking traffic. With the growing implementation of SANs, accelerated by the adoption of blade servers and server virtualization, the penetration of SANs is expected to grow much higher.

While IT managers can continue to maintain two separate networks, in reality this increases the overall operation costs by multiplying the number of adapters, cables, and switch ports required to connect every server directly with supporting LANs and SANs. Further, the resulting network sprawl hinders the flexibility and manageability of the data center.

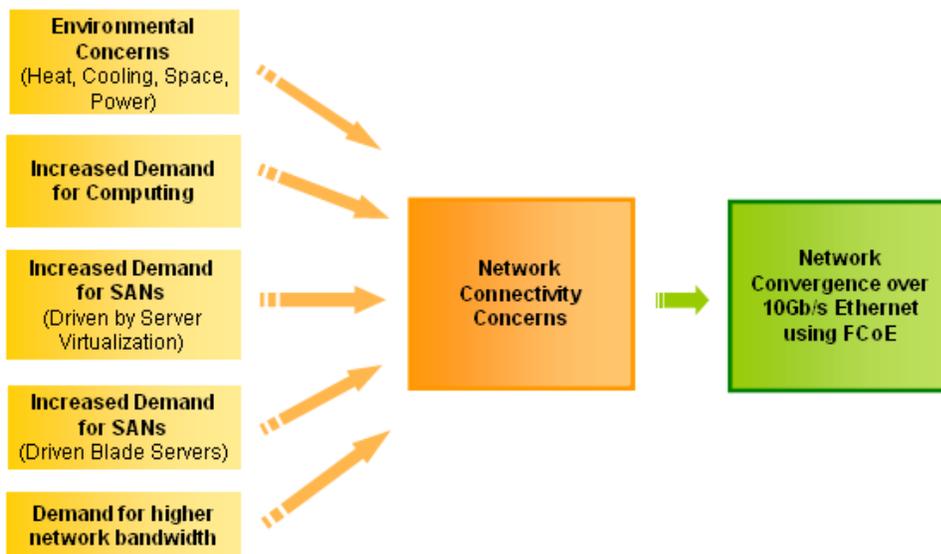


Figure 1: Key factors driving network convergence.

Network convergence, enabled by FCoE, helps address the network infrastructure sprawl while fully complementing the server consolidation efforts and improving the efficiency of overall operations of the enterprise data center.

Benefits of FCoE Enabled Network Convergence

An FCoE enabled converged network provides the following key benefits to IT organizations:

- **Lower total cost of ownership through infrastructure simplification**

A converged network based on FCoE lowers the number of cables, switch ports and adapters required to maintain both SAN and LAN connectivity. The reduction in the number of adapters facilitates the use of smaller servers and the expanded offload to the adapter enables higher CPU efficiency which directly affects power and cooling costs.

- **Consolidation while protecting existing investments**

FCoE lets organizations phase the roll out of converged networks. FCoE-based network convergence drives consolidation in new deployments without affecting the existing server and storage infrastructure or the processes required to manage and support existing applications. For example, the use of a common management framework across Fibre Channel and FCoE connected servers protects existing investments in management tools and processes and thus lowers the long-term operating cost of the data center.

- **Increased IT efficiency and business agility**

A converged network streamlines and eliminates repeated administrative tasks such as server and network provisioning with a “wire once” deployment model. The converged network also improves business agility letting data centers dynamically and rapidly respond to requests for new or expanded services, new servers and new configurations. The converged network fully complements server virtualization in addressing the on-demand requirements of next generation data center where applications and the infrastructure are provisioned on the fly.

- **Enables Seamless Extension of Fibre Channel SANs in the Data Center**

Fibre Channel is the predominant storage protocol deployed by enterprise data centers. With the adoption of blade servers and server virtualization there is an increased demand for Fibre Channel SANs in these environments. FCoE addresses this requirement by leveraging 10Gb/s Ethernet and extending proven Fibre Channel SAN benefits to these environments. The use of light-weight encapsulation for FCoE also facilitates FCoE gateways that are less compute intensive and thus ensures high levels of storage networking performance.

- **Maintains alignment with existing operational models and administrative domains**

FCoE based converged networks gives IT architects the ability to design a data center architecture that aligns with existing organization structure and operational models. This minimizes the need to unify or significantly overhaul operational procedures used by storage and networking IT staff.

FCoE Technology Overview

Fibre Channel over Ethernet (FCoE) is the industry standard being developed to drive network convergence in the enterprise data center. The FCoE technology leverages lossless Ethernet infrastructure to carry native Fibre Channel traffic over Ethernet.

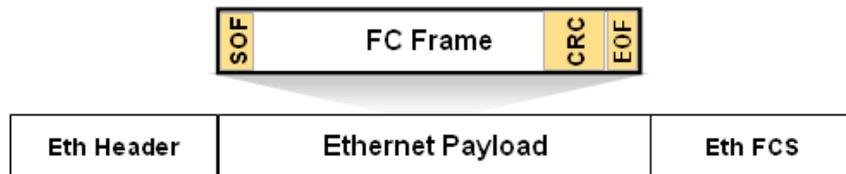


Figure 2: FCoE Encapsulates complete Fibre Channel frames on to an Ethernet frames

Leveraging 10Gb/s Ethernet to carry both the data networking and storage networking traffic enables convergence in the data center networks and thus reduces the number of cables, switch ports and adapters which in turn lowers the overall power and cooling requirements and the total cost of ownership.

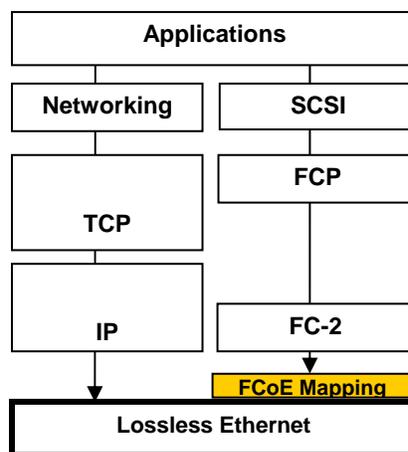


Figure 3: Networking and Fibre Channel stacks leverage a common Enhanced Ethernet infrastructure to achieve convergence in the data

By continuing to maintain Fibre Channel as the upper layer protocol, the FCoE technology also fully employs existing investments in Fibre Channel infrastructure.

Introduction to Emulex LP21000 CNA

Emulex is one of the leading proponents and developers of Fibre Channel over Ethernet (FCoE) technology. The Emulex LP21000 Converged Network Adapters (CNA) product family is designed to support high performance storage and networking applications in enterprise data centers. The CNAs simplifies network deployment by converging networking and storage connectivity while providing enhanced performance for virtualized server environments.



Figure 4: Emulex LP21000 Converged Network Adapter

Servers supporting **PCI Express (x8 or higher)** will be able to utilize the full potential of Emulex LP21000's 2x10Gb/s Ethernet interfaces, balancing the storage and networking I/O requirements. The priority flow control prevents degradation of traffic delivery and enhanced transmission selection enables partitioning and guaranteed service levels for the different traffic types.

The CNA fully leverages the proven Emulex LightPulse® architecture and thus extends the reliability, performance and manageability characteristics of the proven Fibre Channel host bus adapters to the converged network adapter.

Converged Network Deployment models

Emulex is closely partnering with CEE switch manufactures to develop the next generation of FCoE based converged network infrastructure. The combination of lossless Ethernet switches and Emulex CNAs will help facilitate the adoption of FCoE and pave the way for 10Gb/s Ethernet as the converged network fabric in the enterprise data center.

The lossless Ethernet switch, typically provisioned with native Fibre Channel ports and an FCoE to Fibre Channel Forwarder, provides all the services of a Fibre Channel switch. The Switch can either be connected directly to a Fibre Channel Storage array directly using the F-port functionality or the solution can be connected to another Fibre Channel switch or a Fibre Channel Director (using E-port or NPIV mode) which in turn are connected to the rest of the SAN.

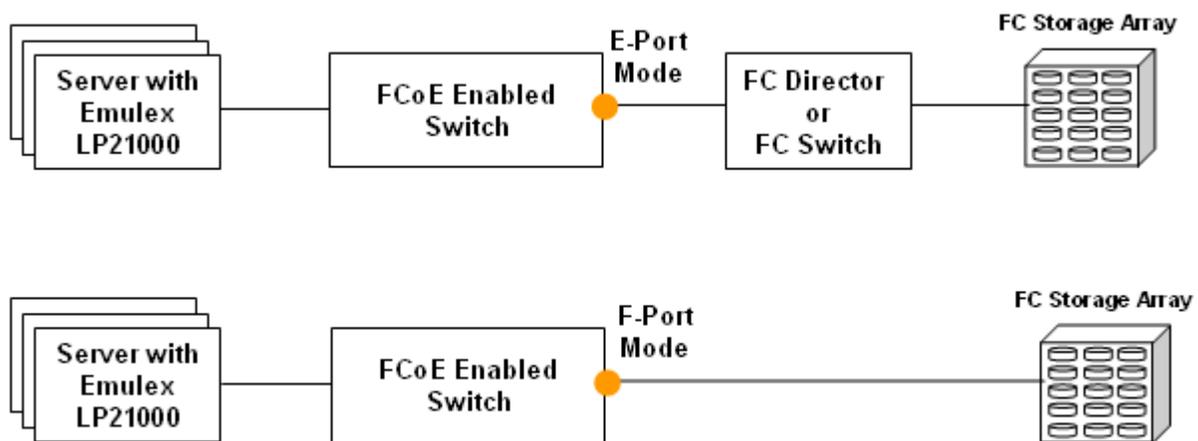


Figure 5: Converged network deployment and access to SAN

Converged Networking with FCoE in an ESX Environment

Server virtualization is one of the key drivers of Fibre Channel over Ethernet deployment. FCoE enabled network convergence significantly reduces the number of cables, switch ports and adapters as show in the picture below –

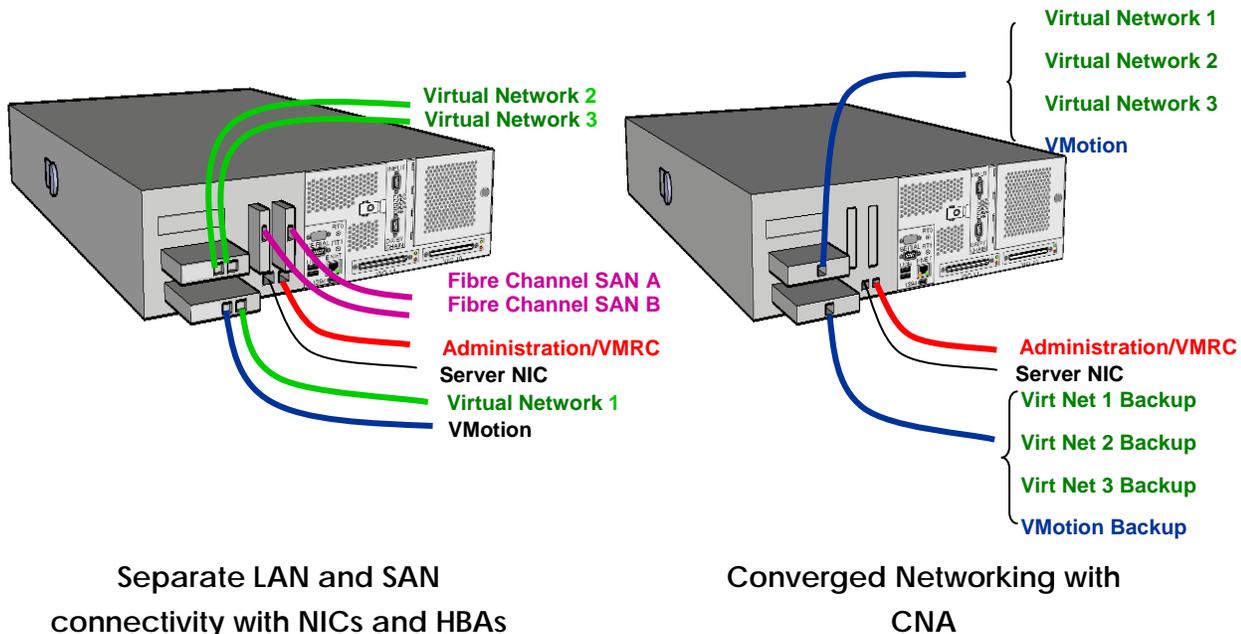


Figure 6: Effects of Convergence using a CNA in a Rack-mount Server running ESX

The following sections of this guide provide a detailed description of how to deploy and configure an Emulex LP21000 in a VMware ESX Server environment.

CNA Instantiation Model in ESX 3.5

Although the CNA is a single physical entity, to an operating system the adapter is represented as a HBA (Host Bus Adapter) and as a NIC (Network Interface Card). 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.

Extending this concept to VMware ESX, the Emulex LP21000 CNA is recognized by ESX as both a HBA (storage adapter) and a NIC (network adapter). The ESX server in turn provides the networking and storage networking access for the virtual machines hosted on that ESX server. The configuration and manageability of the ESX server is performed by VMware Virtual Center (typically residing on a separate management server).

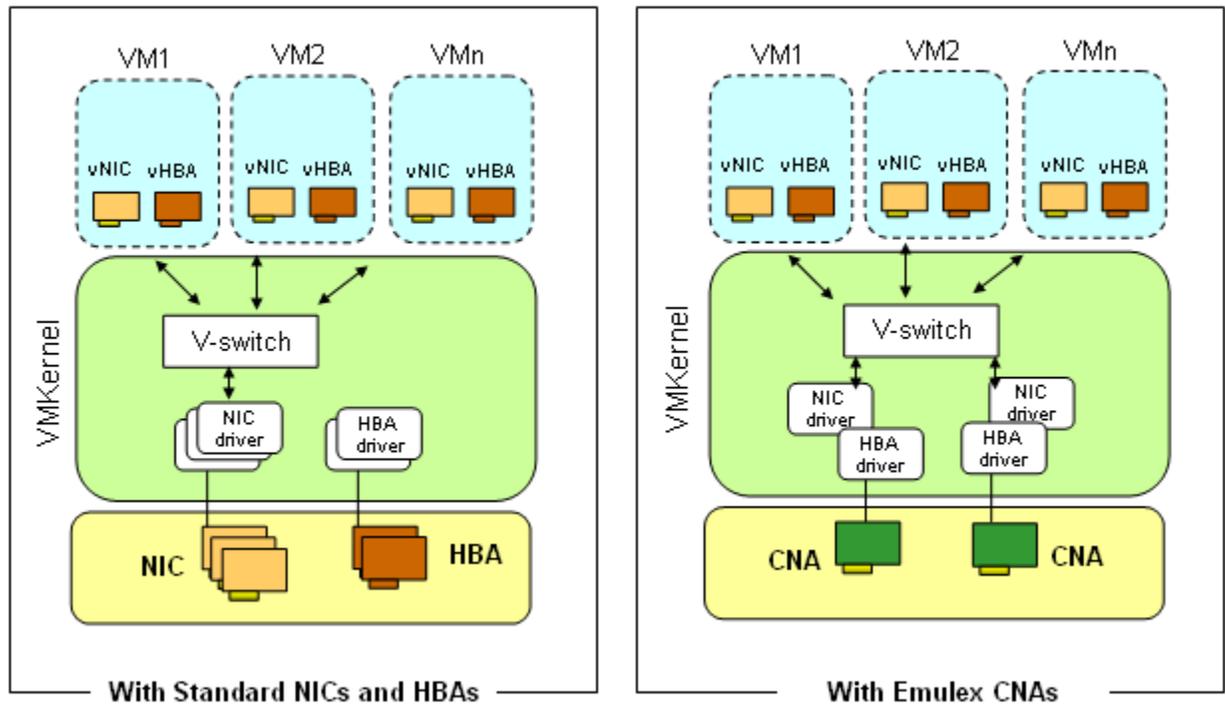


Figure 7: The hypervisor/VMKernel views the CNA as an HBA and a NIC

When an Emulex LP21000 CNA is installed on an ESX server, the adapter is visible in Virtual Center Configuration Tab in the storage adapter view as well as the network adapter view. When deploying multiple CNAs for failover and load balancing, both Fibre Channel multipathing and NIC teaming must be configured on both the CNAs.

Configuring the LP21000 as a Storage Adapter

Upon installation the LP21000 appears (see Figure 8) as a storage adapter similar to that of a HBA. The LP21000s are each associated with a standard 48-bit Fibre Channel WWN (World Wide Name), and in the *details* area for the selected LP21000, all LUNs visible to that adapter are displayed. The LUN masking and zoning operations required for LUN visibility, are carried out similar to configuring the Fibre Channel HBA.

By selecting *storage* from within the *configuration* tab (see Figure 9), the Add Storage Wizard can be used to create the data stores and select either Virtual Machine File System (VMFS), or Raw device Mapping (RDM) access method (for more information refer to the

VMware ESX 3.5 configuration guide). Both access methods can be used in parallel by the same ESX server using the same physical HBA or CNA.

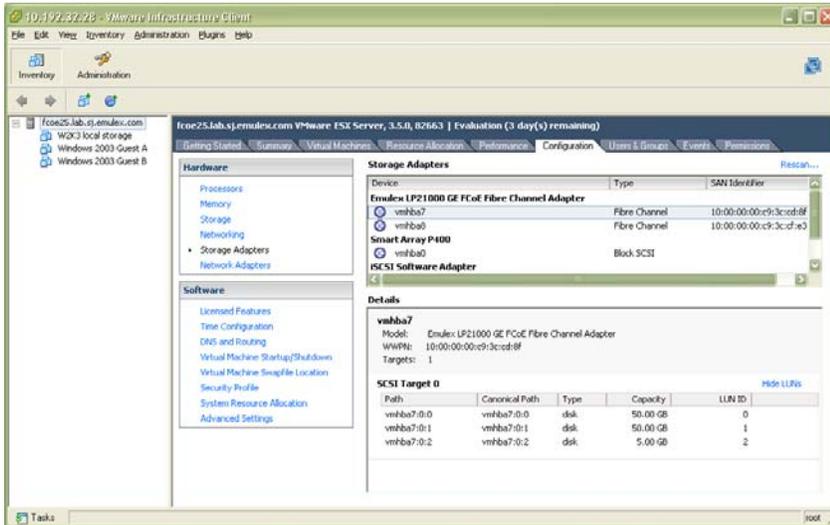


Figure 8: Two LP21000 CNAs appear as fibre channel storage adapters

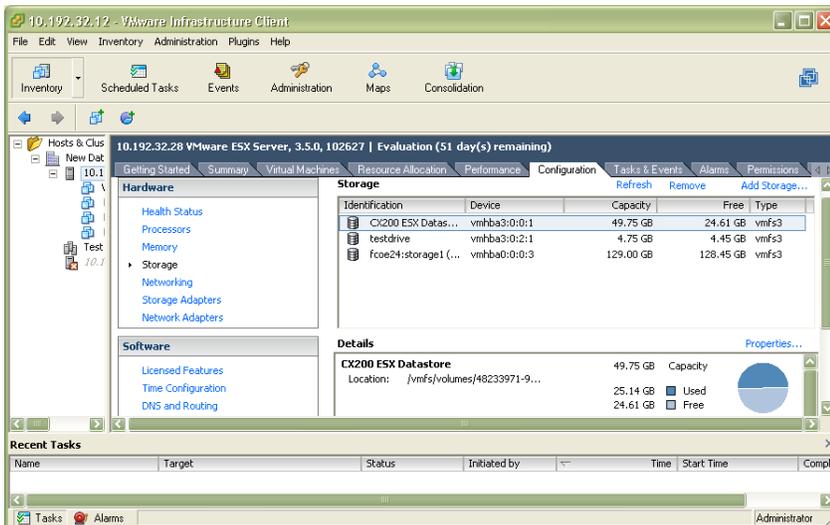


Figure 9: Use the “Add Storage Wizard” to create the datastore

An option for IT organization is to leverage Emulex’s Virtual HBA technology, based on the ANSI T11 **N_Port ID Virtualization (NPIV)** standard. This virtualization technology allows a single physical CNA or HBA port to function as multiple logical ports, each with its own separate identity. NPIV is currently supported in ESX Server 3.5 and up, and currently requires use of Raw Device Mapping (RDM) rather than VMFS. Using NPIV, each virtual machine can attach to its own Virtual port (or VPort) which consists of the combination of a distinct Worldwide Node Name (WWNN) and up to four Worldwide Port Names (WWPNs), as shown in Figure 10. Again, the procedures for using this technology are the same as

those for Fibre Channel HBAs, including full support for virtual machine and virtual port migration by VMotion. For best practices on leveraging NPIV and RDM in ESX Environments see the companion best practices guide “Best Practices Guide: Emulex Virtual HBA® Solutions and VMware® ESX 3.5” available at the following location:

<http://www.emulex.com/white/hba/emulex-vmware-npiv-best-practice.pdf>

Configuring the LP21000 for Multipathing in ESX

Multipathing is common technology used by data centers to protect against adapter, switch, or cable failure. Enabling multipathing with the CNA typically means providing link redundancy either in the form of dual-port configurations or a configuration with two single port CNAs to support failover. Similar to HBA multipath configuration, each of the CNA ports used for multipathing must have visibility to the same LUNs including proper physical connectivity, LUN masking, zoning, and any other array configuration. When each LP21000 has the required visibility to the same LUNs and an ESX data store has been appropriately created on those LUNs, then ESX will provide the path switching or the multipathing to that data store, providing a redundant path for VMs to their corresponding storage store. No special drivers are required, since the multipathing functionality is part of the value-add of ESX.

A single dual-port CNA can provide port redundancy with each port connecting to a different CEE switch

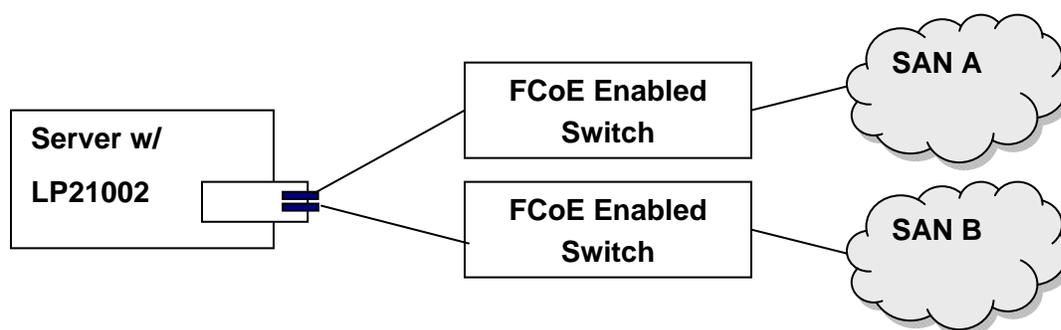


Figure 10: Port level redundancy using a single Dual-Ported CNA

Though the above setup provides multiple data paths to the storage and provides active-standby links for networking, it does not eliminate the CNA as the single point of failure. This situation is avoided by the next configuration. Added redundancy can be provided with

two CNAs serving one host. Installing two CNAs in one host is the **recommended** configuration for high availability converged networking.

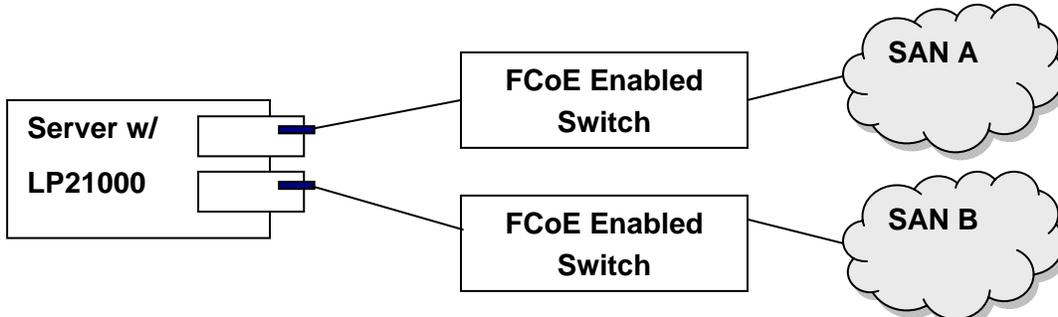


Figure 11: CNA level redundancy using two Single-Ported CNAs

In this configuration, even in the extreme case of a CNA failure, the links for storage and networking can be immediately failed over without any interruption to applications.

Within ESX, multipathing can be configured from the configuration tab selecting storage, followed by selecting properties for the appropriate data store. From the data store properties windows, selecting Manage Paths will reveal the multiple paths to the data store and the policy for multipathing access (fixed, last active, or round robin with the latter being experimental) see Figure 12. Once the multipathing configuration has been completed for a set of CNAs on an ESX server, redundancy configuration for the storage functionality of the CNAs will be complete. However, the network side of the CNA still requires configuration (NIC teaming).

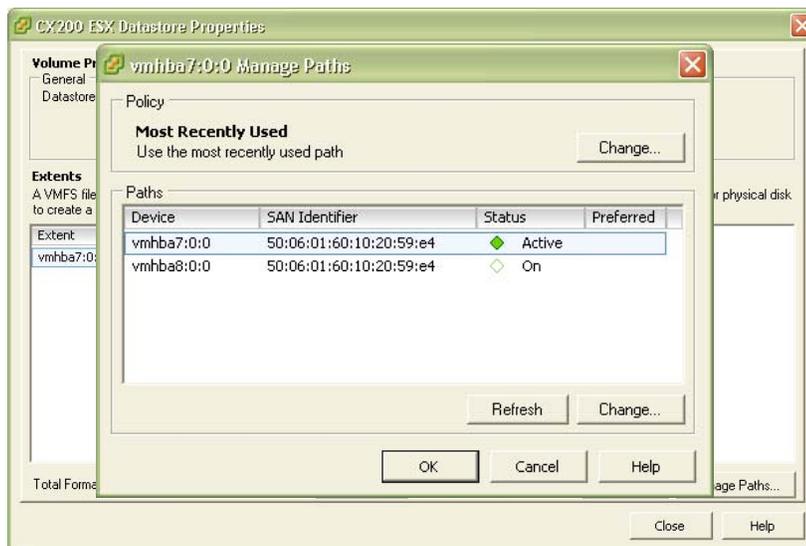


Figure 12: CNA initiated multipathing accessing a data store

Configuring the LP21000 as a Network Adapter

Upon installation, the LP21000 appears as a 10 Gigabit NIC in the Configuration Tab's network adapter view, as shown in Figure 13 where two LP21000s appears as two 10 Gigabit network adapters (also shown are two 100 Mbps NICs which are built into the server motherboard.).

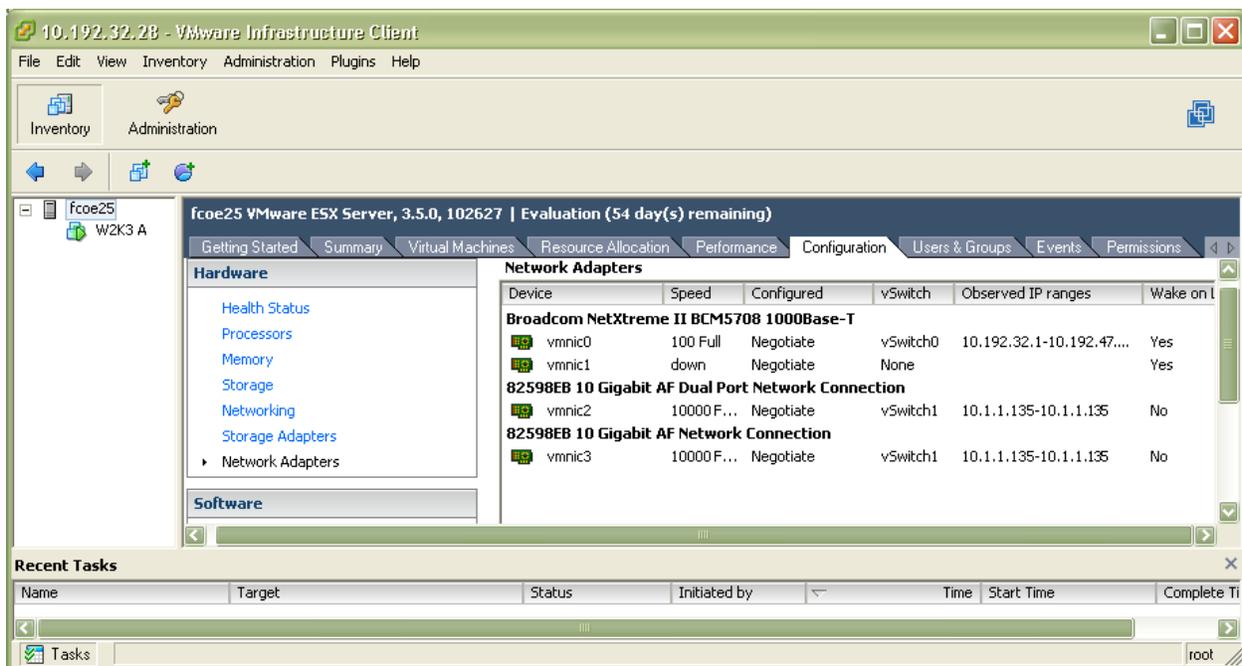


Figure 13: Two LP21000 CNAs appear as two 10 Gigabit network adapters

It is recommended that the Service Console traffic is configured on a separate physical network, using the 100 Mbps links. This provides the enhanced security for the service console as well as enables CNA link bandwidth to be shared by multiple virtual machines. Virtual machines reach the external network through the virtual network (which comprises of one or more vSwitches) which in turn are linked with the physical adapters. In order to create a virtual network and assign it a virtual machine, begin by selecting Networking from the Configuration tab, select Add Networking to start the Add Networking Wizard. Select *create a virtual switch* and select the network adapters (which are representations of CNAs) from the list to map them as uplinks to the virtual switch. Since multiple virtual machines with variety of traffic types can now access the shared 10Gb/s link, it is also recommended that VLANs be used to segregate traffic from different virtual machines. The VLAN information can be configured as part of Port Group properties under Connection Settings screen of the Add Networking Wizard.

Configuring the LP21000 for NIC Teaming in ESX

NIC teaming occurs when multiple uplink adapters are associated with a single vSwitch to form a team. A team can either share the load of traffic between physical and virtual networks among some or all of its members or provide passive failover in the event of a hardware failure or a network outage. Figure 14 shows the configuration of two vSwitches within an ESX server. The Service Console traffic is assigned to a vSwitch that in turn is connected to 100 Mbps NIC. The second vSwitch provides networking for the VMkernel service and for a virtual machine. On the uplink, the vSwitch is connected to two LP21000 CNAs which provides redundancy through NIC teaming.

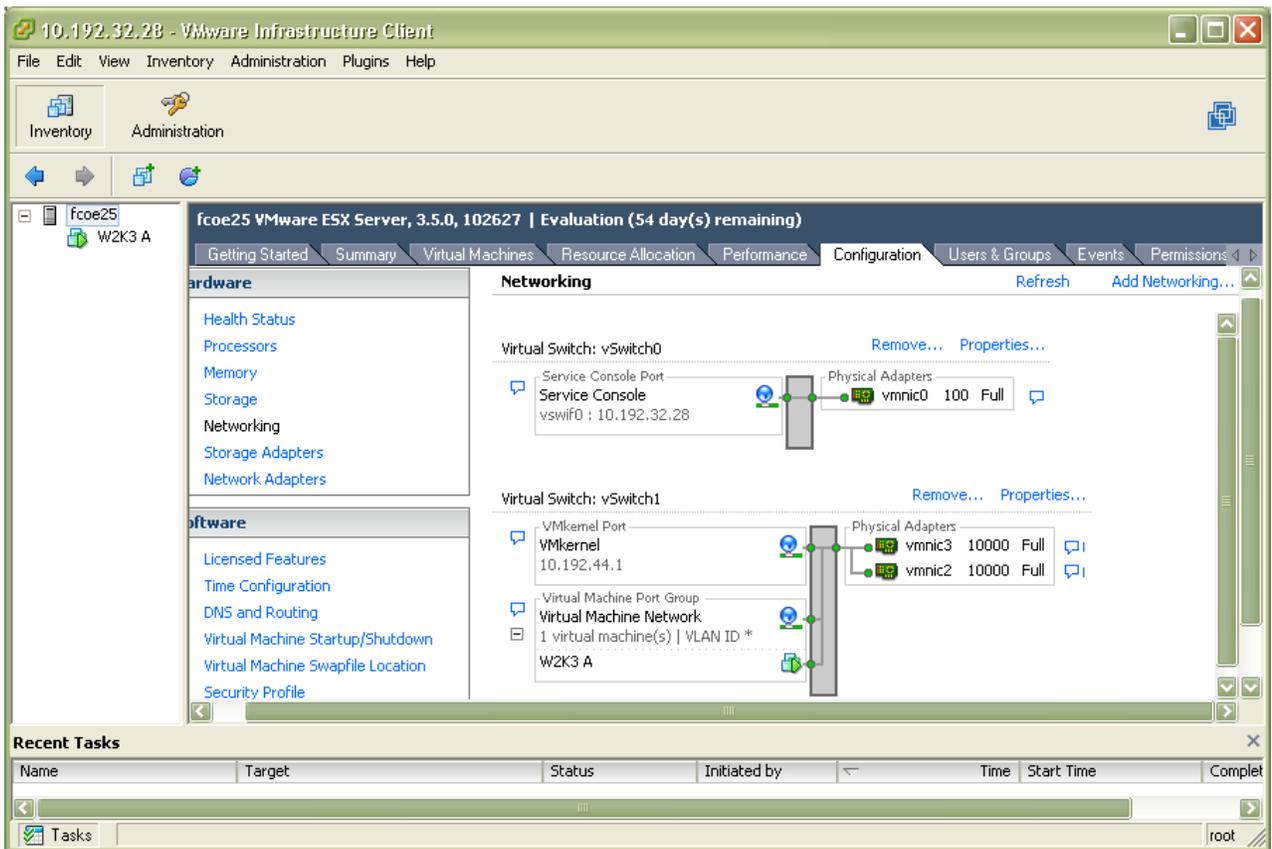


Figure 14: A network configuration where a vSwitch is uplinked with two CNAs for NIC teaming

Scaling to Enterprise Manageability

As FCoE-based convergence makes its way into the enterprise data centers, one of the key requirements is the ability to manage both the FCoE and Fibre Channel SAN-based infrastructure without increasing the overall management complexity. Emulex provides software applications that streamline deployment, simplify management and virtualization as well as provide multipathing capabilities in mission critical environments. Emulex HBAnyware 4.0 is a unified management platform for Emulex LightPulse Fibre Channel HBAs and FCoE CNAs. It provides a centralized console to manage Emulex HBAs and CNAs on local and remote hosts. Some new features of HBAnyware 4.0 include -

Unified management of Emulex Fibre Channel HBAs and CNAs

- Virtual Port Monitoring and Configuration
- Online Boot from SAN Configuration
- On-line WWN management
- Flexible, customizable reporting capability

The availability of graphical user interface (GUI) and command line interface (CLI) options provide greater flexibility with HBAnyware. The application is available for a broad set of operating systems including Linux, Windows, Solaris and VMware ESX platforms and supports cross-platform SANs. Some of the highly valued capabilities of HBAnyware that make it more attractive for enterprise scalability and high availability are firmware management, with the ability to update or “flash” firmware without having to bring down the CNA or HBA, and the ability to flash a set of adapters on remote servers with just a few mouse clicks. Such capabilities enhance SAN availability and is one of the reasons Emulex adapters have been known to run continuously for years.

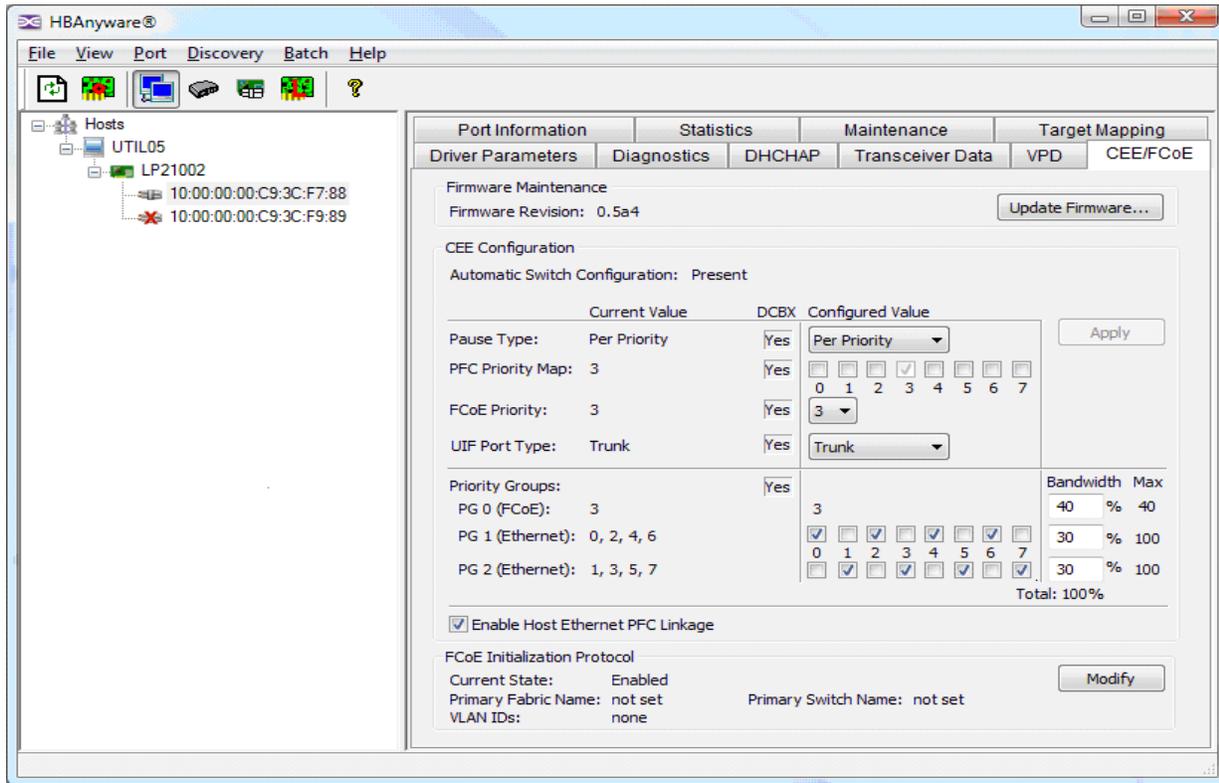


Figure 15: HBAAnyware 4.0 common Management Application for HBAs and CNAs that can be local or remote

Of additional value for data center manageability is the common driver architecture where the same driver can be used with all generations of Emulex HBAs and CNAs. This capability reduces driver update errors and simplifies administration.

Overall the manageability value of the LP21000 CNA as illustrated with HBAAnyware 4.0 and the Emulex common driver architecture enables existing data centers to seamlessly extend their Fibre Channel management processes to FCoE attached servers, thereby facilitating efficiency in a multiprotocol, multi-operating system enterprise SAN, while lowering costs and protecting data center investment. The Emulex HBAAnyware application also provides functions for assigning bandwidth for individual traffic types by apportioning bandwidth assignments for each priority group as shown in the above figure.

Conclusion

Server virtualization has been one of the key drivers of converged network in the data center. A converged network fully complements server virtualization and enables the roll out of on-demand services where applications and network services are provisioned on the fly. Fibre Channel over Ethernet technology enables network convergence in enterprise data centers and provides benefits such as infrastructure simplification, investment protection and lowered total cost of ownership. The key advantage with FCoE is that it extends these benefits without disrupting existing architectures and operational models.

Emulex has partnered with VMware in enabling FCoE based network convergence for the ESX 3.5 Server environments. This document provided an outline for how to deploy this emerging technology in an ESX environment.

VMware, Inc. 3401 Hillview Drive Palo Alto CA 94304 USA Tel 650-427-5000

www.vmware.com

Emulex Corp. 3333 Susan Street Costa Mesa CA 92626 USA Tel 714-662-5600

www.emulex.com

© Emulex Corporation. All Rights Reserved. No part of this document may be reproduced by any means or translated to any electronic medium without prior written consent of Emulex.

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

VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

