

Deployment Guide: Network Convergence with Emulex[®] OneConnect[™] FCoE Technology and VMware[®] Server Virtualization

How to Deploy Converged Networking with
VMware ESX[™] Server and ESXi Using Emulex OneConnect UCNAs

WHITE PAPER

Table of Contents

Introduction	3
Drivers of Network Convergence	3
The Data Center Networking Challenge	3
Data Center Bridging: The Foundation for Network Convergence	3
Benefits of Network Convergence and FCoE	3
FCoE Technology Overview	3
Introduction to Emulex OneConnect™ FCoE CNA	4
Converged Network Deployment Models	4
Converged Networking with FCoE in an ESX Server Environment	5
CNA instantiation model in ESX Server	6
Configuring the OneConnect UCNA as a storage adapter	6
Installing the OneConnect UCNA for VMware platforms	6
Verifying installation	7
Emulex CIM provider for ESXi	7
Configuring the OneConnect UCNA for multipathing in ESX Server	8
Configuring the OneConnect UCNA for NPIV	9
Configuring the OneConnect UCNA as a network adapter	10
Configuring the OneConnect UCNA for NIC teaming in ESX Server	10
Scaling to enterprise manageability	11
Emulex OneCommand Manager application diagnostics	11
Conclusion	12

Introduction

Fibre Channel over Ethernet (FCoE) is a standard developed by INCITS T11 that fully leverages the enhanced features of 10 Gigabit Ethernet (10GbE) for I/O consolidation in the data center. The Emulex OneConnect Universal Converged Network Adapter (UCNA) is the first adapter to provide protocol offloads for TCP/IP, FCoE and iSCSI with a single product platform. It allows companies to wire once, creating a seamless upgrade path to solidify the benefits of network convergence.

Cloud computing is emerging as the next cost-saving technology for data centers. IT departments are building on their VMware server virtualization deployments and implementing hybrid-cloud solutions to increase operating efficiencies and improve resource utilization. In order to satisfy performance demands with a highly-efficient infrastructure, cloud providers are also moving to 10GbE converged networks.

Emulex has a long and successful history of working closely with VMware and other partners to leverage new technologies that evolve in converged network environments. This deployment guide presents information about network convergence and also provides an overview of how to realize the maximum benefits when deploying OneConnect FCoE technology in a converged network supporting a VMware ESX Server environment.

Network Convergence

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

- Multiple network fabrics, each dedicated to a specific type of traffic
- High numbers of adapters and switch ports
- Complicated cabling infrastructure
- Complex management of switch and adapter firmware and associated service contracts
- How to maximize server virtualization return on investment (ROI)

Data Center Bridging and FCoE

The IEEE Data Center Bridging (DCB) standard enables multiple traffic types over a single link, or network convergence, with FCoE and Ethernet using the same cable. With DCB, data centers can achieve the following:

- **Consolidate**—Replace multiple 1GbE links with a high-bandwidth 10GbE infrastructure, significantly reducing the number of switch ports, adapter ports and cables.
- **Converge**—Unify networking and storage traffic to a single network, eliminating the need for dedicated networks for each traffic type.
- **Virtualize**—Virtualize the underlying physical network infrastructure and provide service level guarantees for each type of traffic. In addition, the solution must be capable of responding to dynamic changes in network services based on the business demands of data center applications.

FCoE Technology Overview

The FCoE protocol specification provides for direct encapsulation of a complete Fibre Channel frame (including checksum and framing bits) within an Ethernet frame and avoids the overhead of any intermediate protocols (see Figure 1).




Figure 1: FCoE encapsulation in Ethernet

This lightweight encapsulation ensures that FCoE-capable Ethernet switches are less compute-intensive, providing high performance that is typical of a Fibre Channel network. By retaining Fibre Channel as the upper layer protocol, FCoE fully supports critical Fibre Channel capabilities, including fabric login, zoning and Logical Unit Number (LUN) masking.

Benefits of Network Convergence and FCoE

The combination of network convergence and FCoE provide the following key benefits to IT organizations:

- **Infrastructure simplification**—Network and storage I/O is converged over a single 10GbE infrastructure that supports SAN and local area network (LAN) connectivity with common adapters, switch ports and cables.
- **Protection for existing Fibre Channel investments**—FCoE is designed to fully leverage existing Fibre Channel investments. Deployments begin at the edge of the data center and storage traffic is routed by a Fibre Channel forwarding switch into the existing Fibre Channel SAN. Full interoperability with installed Fibre Channel switches and arrays means no need to “rip and replace”. Even more critical, there’s no requirement to move and potentially disrupt access to data. A common management framework across Fibre Channel and FCoE connected servers protects existing investments in management tools and processes and lowers long-term operating costs.

- **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.

Emulex OneConnect FCoE CNA

Emulex OneConnect FCoE CNAs are single-chip 10GbE server adapters that support FCoE, while leveraging mature Fibre Channel management tools and processes (see Figure 2). The OCE10102-F and OCE11102-F offload Fibre Channel protocol processing from the CPU, providing high performance storage connectivity and optimized CPU utilization. The additional processing power retained by CPUs provides more CPU bandwidth for enterprise applications and enables the deployment of more virtual machines (VMs) per virtualized server.

Emulex FCoE CNAs leverage ten generations of advanced, field-proven LightPulse Fibre Channel HBA technology, making them ideal for server I/O consolidation.

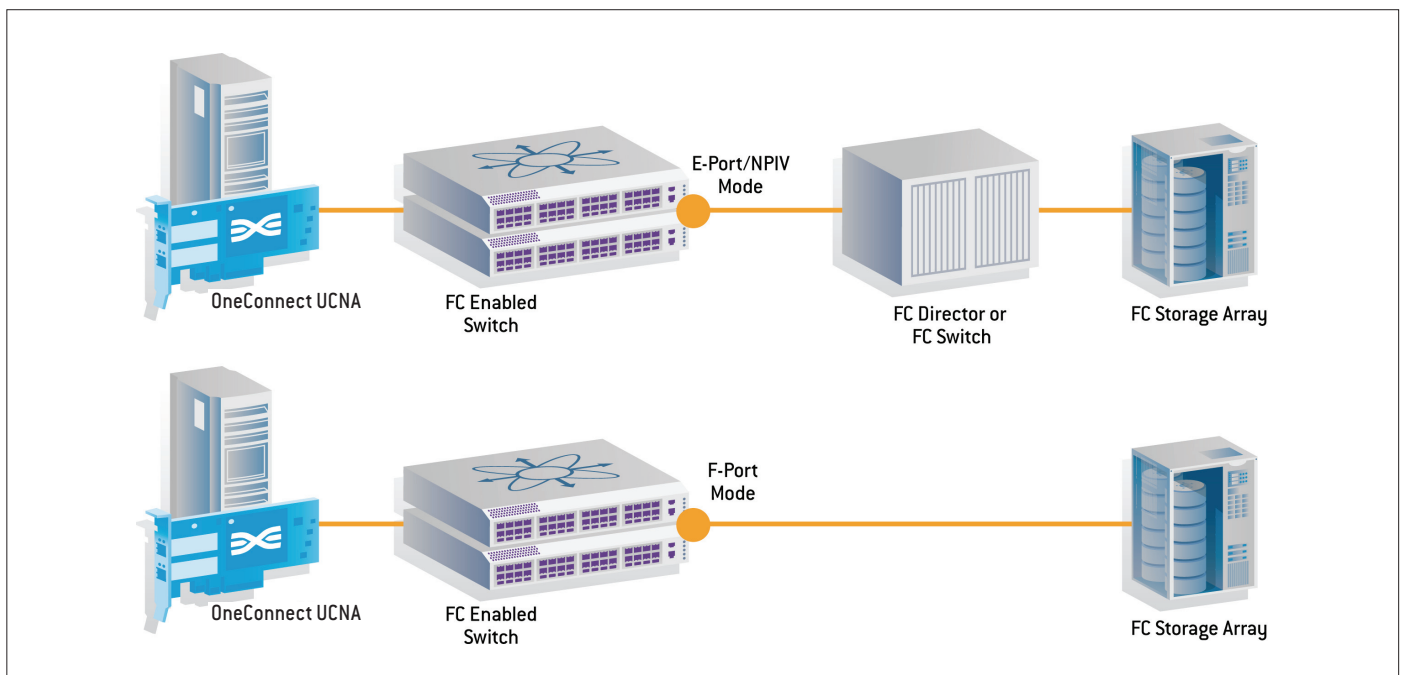


Figure 2: Converged network deployment and access to SAN

Converged Networking with FCoE in an ESX Server Environment

Server virtualization is one of the key drivers of FCoE deployment. FCoE-enabled network convergence significantly reduces the number of cables, switch ports and adapters, as shown in Figure 3, for an ESX Server environment.

The following sections provide a detailed description of how to deploy and configure an Emulex OneConnect OCe1102-F or OCe10102-F FCoE CNA in a VMware ESX Server environment. Prior to the deployment of virtualized environments using VMware ESX platforms and Emulex adapters, the following areas should be considered:

- Access to networked storage through a Fibre Channel or FCoE switch
- Multipathing schemes for failover

- Disaster recovery requirements to minimize potential disruption, exposure and data loss
- Using Virtual Machine File System (VMFS) or Raw Device Mapping (RDM)

VMFS is the easiest and fastest data store to deploy. VMFS is a virtual, clustered file system that allows multiple VMware ESX servers to access the same data store and associated LUN, which can be expanded with dynamic LUN extension. VMFS requires minimal administrative overhead and provides acceptable performance for the majority of VMs.

RDM maps a specific LUN to a single VM by using a mapping file inside VMFS that acts as a proxy for a raw device, allowing direct block-level access from the VM to the LUN (see Figure 4). The mapping file is presented to the management software as an ordinary disk file available for the usual file system operations. Emulex's NPIV-enabled vPort capabilities are only supported with RDM.

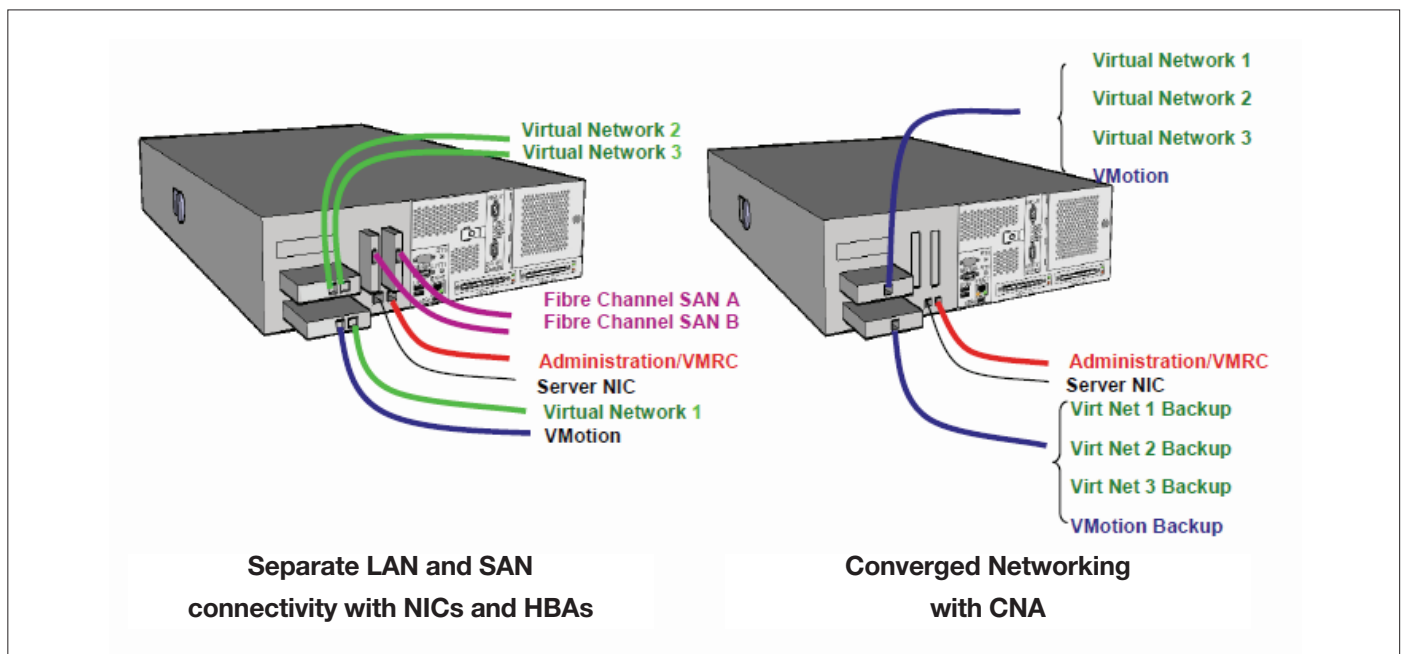


Figure 3: Port reductions using a OneConnect UCNA in a rack-mount server running ESX Server

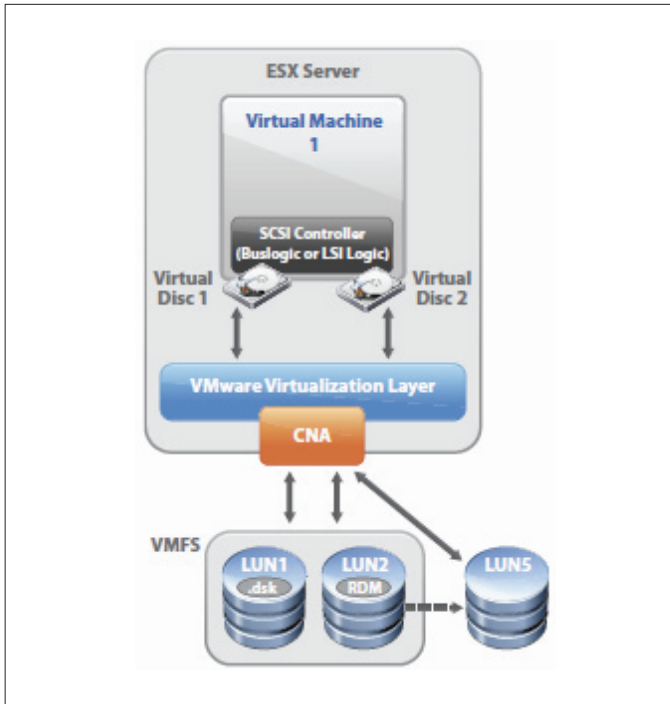


Figure 4: Storage connectivity with CNA in ESX environment

CNA instantiation model in ESX Server

Although the CNA is a single physical entity, to an operating system (OS), each CNA port is represented as an HBA and a NIC (network interface card), as shown in Figure 5. This level of transparency to the OS ensures that the CNA configuration and management practices are performed in the same manner as would be done with separate NIC and HBA devices.

Extending this concept to VMware ESX Server, the Emulex OneConnect UCNA is recognized by ESX Server as both an HBA and a NIC. ESX Server, in turn, provides the networking and storage networking access for VMs hosted on it. Configuration and management of the ESX Server is performed by VMware vCenter™ Server (typically residing on a separate management server).

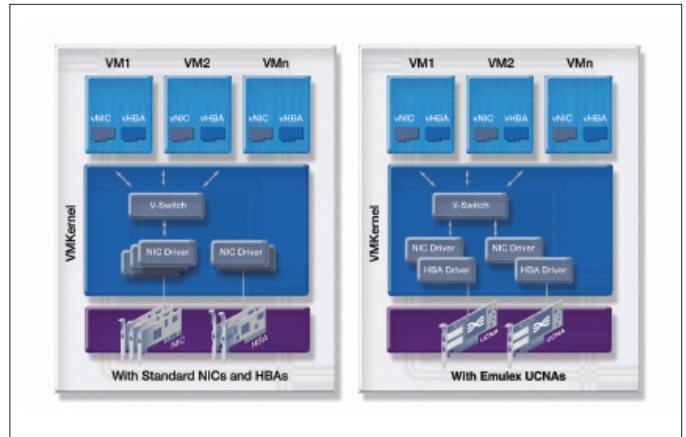


Figure 5: The ESX Server hypervisor views the CNA as an HBA and a NIC

When an Emulex OneConnect UCNA is installed on an ESX server, the adapter is visible in the vCenter Configuration Tab in the storage adapter view and the network adapter view. When deploying multiple UCNAs for failover and load balancing, both Fibre Channel multipathing and NIC teaming would be configured.

Configuring the OneConnect UCNA as a storage adapter

OneConnect UCNAs present a storage adapter to the hypervisor, which works just like an HBA (see Figure 6). CNA ports have a standard 48-bit Fibre Channel Worldwide Port Name (WWPN) and all LUNs visible to that adapter are displayed in vCenter in the details area for the selected CNA. LUN masking and zoning operations are the same for CNAs and HBAs.

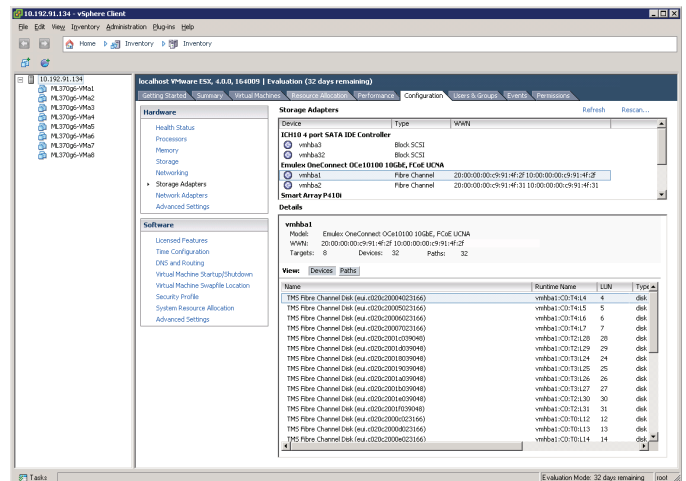


Figure 6: Two OneConnect OCE1000 FCoE CNAs appear as Fibre Channel storage adapters

By selecting storage from within the configuration tab (see Figure 7), the “Add Storage Wizard” can be used to create the data store and select the VMFS or RDM access method (for more information, refer to the VMware ESX configuration guide). Both access methods can be used in parallel by the same ESX server using the same HBA or CNA.

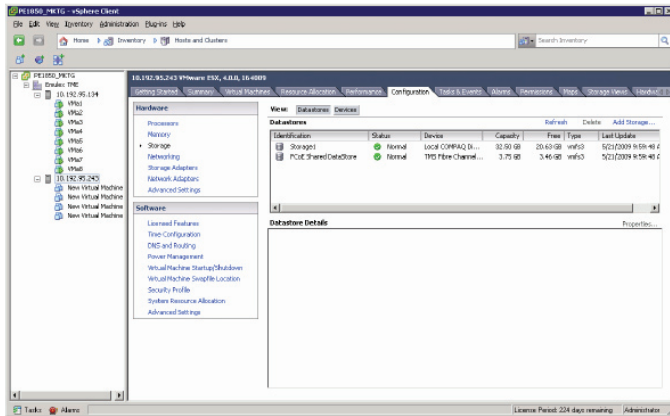


Figure 7: Use the “Add Storage Wizard” to create the data store

Installing OneConnect UCNAs for VMware vSphere

There are three software components for installation of OneConnect UCNAs:

1. Driver (inbox or download via Emulex.com)
2. Core Kit (download from Emulex.com)
3. OneCommand Manager application (download from Emulex.com)

Drivers for Emulex adapters are included in the VMware ESX Server distribution and driver updates are available for download via Emulex.com.

After you install the driver, download the Core Kit from Emulex.com and install it in the VMware Service Console. The Core Kit must be installed on all servers that will be managed by the Emulex OneCommand Manager application. Once installed, data center administrators can use OneCommand Manager to manage Emulex adapters on remote VMware servers.

The last step in setting up the solution is installation of the Emulex OneCommand Manager application on the server running VMware vCenter. With OneCommand Manager installed, IT administrators can centrally manage Emulex adapters on ESX Server, improving overall management efficiency within the data center.

Verifying installation

To ensure proper operation, you should verify the installation process using vCenter, which confirms adapters are installed and working properly and connected targets and storage resources can be accessed.

OneCommand Manager for VMware

Emulex also offers the OneCommand Manager for VMware Plug-in enabling comprehensive control of Emulex HBAs and CNAs from the VMware vCenter management console.

Configuring the OneConnect UCNA for multipathing in ESX Server

Multipathing technology is used to protect against adapter, switch or cable failure. Similar to HBA multipath setup, each CNA port used for multipathing must have visibility to the same LUNs, including proper physical connectivity, LUN masking, zoning and any other array configuration. When each OneConnect UCNA port has the required visibility to the same LUNs and an ESX server datastore has been appropriately created on those LUNs, ESX Server will enable path switching or multipathing, providing a redundant path for VMs to their corresponding storage store. Special drivers are not required since the multipathing functionality is part of ESX Server distribution.

A dual-port CNA can provide port redundancy with each port connecting to a different switch port (see Figure 8). Though a dual-port UCNA provides multiple data paths to the storage and provides active standby links for networking, it does not eliminate the UCNA as a single point of failure. This situation is avoided by adding redundancy with two UCNAs installed on the host. Two UCNAs per host is the recommended configuration for high availability converged networking. In the rare case of an adapter failure, the links for storage and networking can be immediately failed-over with no interruption to applications.

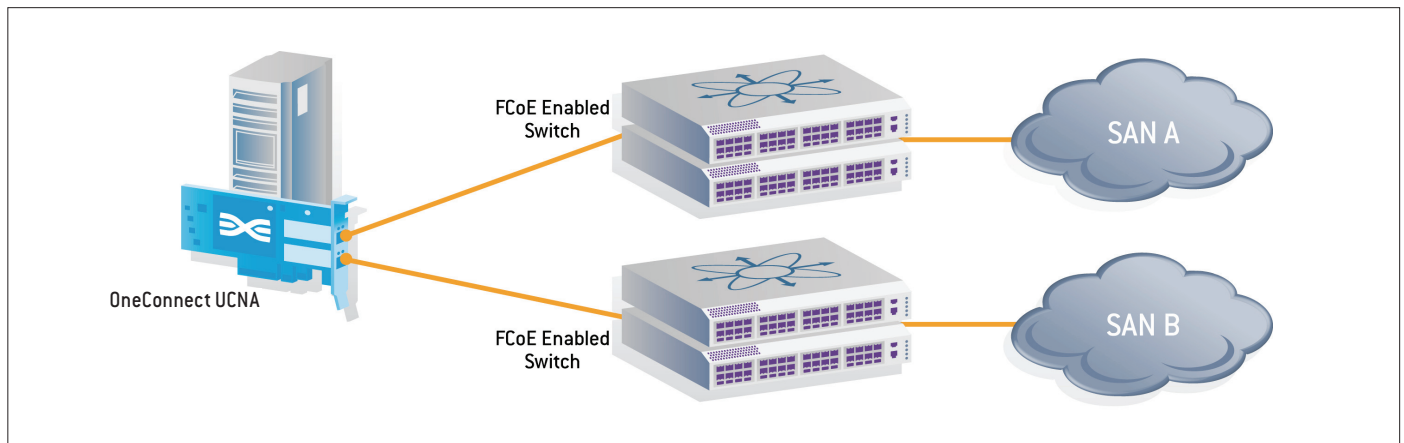


Figure 8: Multipathing with dual-port configuration

Within ESX Server, multipathing can be configured from the configuration tab by following these steps:

- Select “Storage”, followed by properties for the appropriate datastore.
- From the datastore properties window, select “Manage Paths” to reveal the multiple paths to the data store and the policy for multipathing access (fixed, last active, or round robin - see Figure 9).
- Once the multipathing configuration has been completed for a set of CNAs, redundancy configuration for the storage functionality of the CNAs will be complete. However, redundancy for the network function of the CNA still requires configuration (NIC teaming, described later).

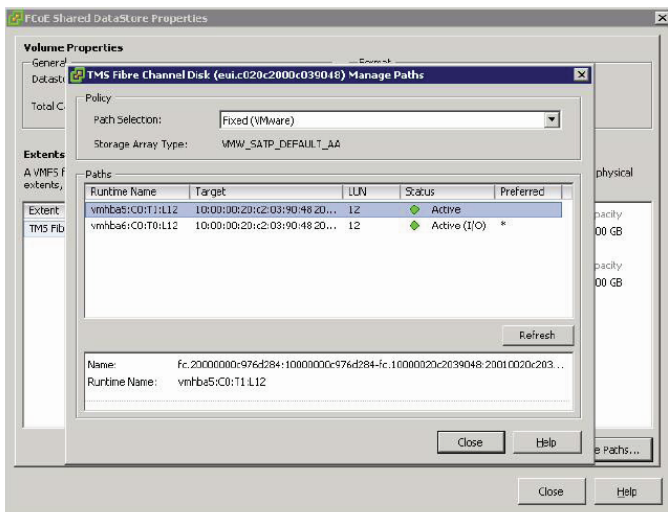


Figure 9: CNA initiated multipathing accessing a data store

Configuring the OneConnect UCNA for NPIV

A VMFS datastore provides a simple model to allocate storage space for VMs without exposing them to the complexity of the underlying storage. A Fibre Channel SAN-based datastore is the most common deployment as it enables VMotion, ESX Server boot from SAN and support for high availability clusters. By default, VMs are associated with Vdisks, which also store VM definitions.

Each VM is stored as a set of files in its own directory in the datastore. A Virtual Machine Disk (VMDK) is a file that resides in a datastore associated with an ESX server. A datastore will reside on a VMFS volume for file-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 VMDK can be manipulated (i.e., copied, moved, backed up) just like any other file. For guest OSs that support hot-adds of new disks, a new VMDK can be added without having to shut down the VM. For more secure configurations, you can leverage NPIV capabilities within Emulex adapters to allocate a specific LUN to a specific VM and limit the access to that LUN using SAN zoning based on an NPIV-enabled vPort.

To enable NPIV, use the following steps:

- Select the “Custom” setting on the Create New Virtual machine dialog box
- Select “Raw Device Mapping”, which is required for NPIV
- Select the “Edit virtual machine setting before completion”
- Choose the “Options” tab and select “Fibre Channel NPIV”
- Select “Generate new World Wide Names” to create an NPIV-enabled vPort

Configuring the OneConnect UCNA as a network adapter

Upon installation, the OneConnect UCNA appears as a 10GbE NIC in the Configuration Tab's network adapter view, as shown in Figure 10, where the dual-port CNA appears as two 10GbE network adapters (also shown are NICs which are built into the server motherboard).

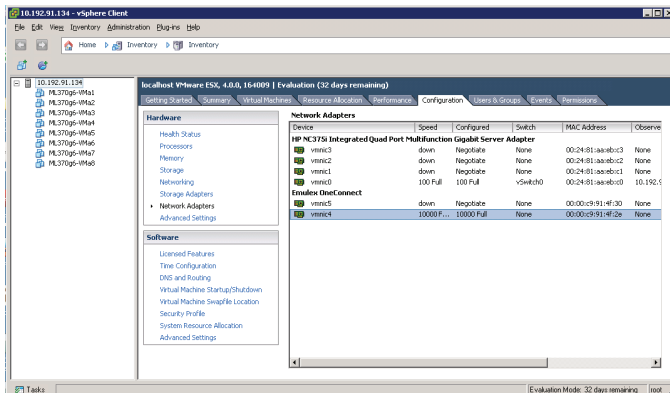


Figure 10: Two OneConnect UCNAs appear as two 10GbE network adapters

For ESX Server, it is recommended that the service console traffic be on a separate physical network, which could be a slower NIC port on the server motherboard. This provides the enhanced security for the service console, as well as preserving high-speed CNA bandwidth for VM I/O. VMs reach the external network through the virtual network (which is comprised of one or more virtual switches [vSwitches]) which, in turn, is linked with the physical adapters.

To create a virtual network and assign it to a VM, use the following steps:

- Select "Networking" from the Configuration tab
- Select "Add Networking" to start the Add Networking Wizard
- Select "Create a virtual switch"
- Select the CNAs from the list to map them as uplinks to the vSwitch

Since multiple VMs with a variety of traffic types can now access the shared 10Gb/s link, it is also recommended that Virtual LANs (VLANs) be used to segregate traffic from different VMs. The VLAN information can be configured as part of Port Group properties under the Connection Settings screen of the Add Networking Wizard.

Configuring the OneConnect UCNA for NIC teaming in ESX Server

NIC teaming occurs when multiple uplink adapters are associated with a single vSwitch to form a team. A team can share the load of traffic between physical and virtual networks or provide passive failover in the event of a hardware failure or network outage. Figure 11 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 a 100 Mbps NIC. The second vSwitch provides networking for the VMkernel service and for a VM. On the uplink, the vSwitch is connected to two OneConnect UCNAs which provides redundancy through NIC teaming.

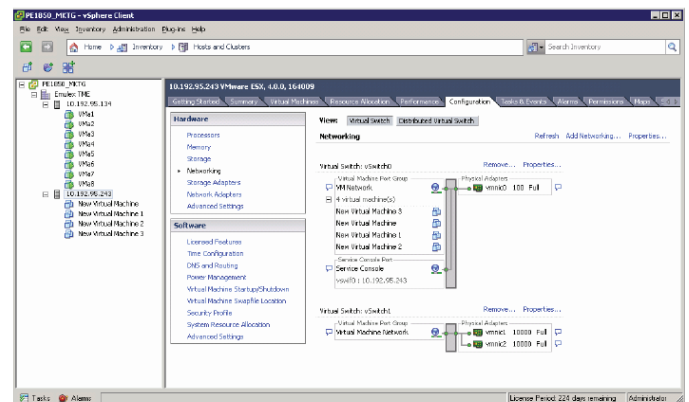


Figure 11: A network configuration where a vSwitch is uplinked with two CNA ports for NIC teaming

Scaling to enterprise manageability

One of the key requirements for network convergence is the ability to manage both the FCoE and Fibre Channel SAN 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.

The Emulex OneCommand Manager application is a unified management platform for Emulex LightPulse Fibre Channel HBAs and OneConnect UCNAs. It provides a centralized console to manage Emulex adapters on local and remote hosts. Updates can be performed on multiple servers at one time. Some important features of OneCommand Manager include:

- Unified management of Emulex Fibre Channel HBAs and OneConnect UCNAs
- vPort monitoring and configuration
- Online boot from SAN configuration
- Online WWN management
- Flexible, customizable reporting

The OneCommand Manager application provides graphical user interface (GUI) and command line interface (CLI) options and is available for a broad set of OS platforms including Linux, Windows, Solaris and VMware. Some of the highly valued capabilities of OneCommand Manager for enterprise scalability and high availability are firmware management, with the ability to update or “flash” firmware without having to bring down the adapter, and the ability to flash a set of adapters on remote servers with just a few mouse clicks.

Of additional value for data center management is the common driver architecture where the same driver can be used with all generations of Emulex HBAs and UCNAs. This reduces driver update errors and simplifies administration. Data center administrators can fine tune adapter driver settings to optimize performance of storage resources. OneCommand Manager also provides functions for assigning bandwidth for individual traffic types by apportioning bandwidth assignments for each priority group, as shown in Figure 12.

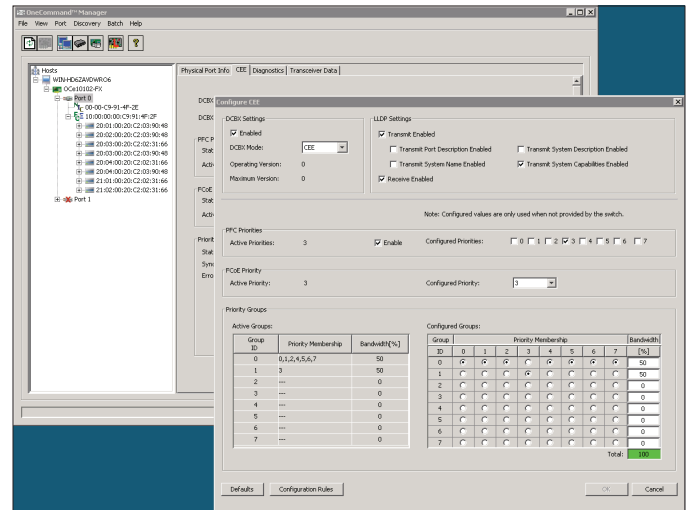


Figure 12: OneCommand Manager for HBAs and CNAs that can be local or remote

The benefits provided by the OneCommand Manager application and the Emulex common driver architecture enable existing data centers to seamlessly extend their Fibre Channel management processes to FCoE attached servers, facilitating efficiency in a multi-protocol, multi-OS enterprise SAN.

Emulex OneCommand Manager application diagnostics

When trouble occurs in a SAN, the first challenge is pinpointing the cause of the problem. The OneCommand Manager application offers a wide set of diagnostic resources that enable IT administrators to quickly and efficiently identify and address SAN issues.

Advanced diagnostics tools include a variety of loop-back and end-to-end tests that can be performed with the OneCommand Manager application to validate connectivity. One example is adapter beaconing which allows quick visual identification of a specific adapter within server racks. OneCommand Manager also supports a Power-On Self Test (POST) while the adapter is active. This enables a complete reset of the adapter without costly system downtime.

There is also a diagnostic log file command, which creates a log file for any Emulex adapter that is being managed. This log retrieves and records all firmware and driver status information which can be saved for reference or emailed to the Emulex support group for remote troubleshooting. An abbreviated log file is also created for quick review by system administrators.

Conclusion

Server virtualization is one of the key drivers of converged networks 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 dynamically. FCoE technology enables network convergence in enterprise data centers and provides benefits such as infrastructure simplification, investment protection and lowered TCO. FCoE extends these benefits without disrupting existing architectures and operating models.

Emulex OneConnect UCNAs are a family of single-chip, high-performance 10GbE multi-function adapters that provide server connectivity for network and storage traffic. OneConnect supports FCoE with lossless Ethernet, seamlessly converging storage and network traffic while leveraging mature Fibre Channel management tools and processes to maximize virtualization ROI.

Together Emulex and VMware provide intelligent solutions that ease the deployment and management of converged networks for virtual environments. Look to these companies to continue developing converged networking solutions for virtualized deployments and other leading edge technologies.

