

Maximizing BladeCenter I/O Capabilities for VMware Environments with Dedicated I/O for Ethernet and Fibre Channel

AT A GLANCE

An IBM BladeCenter[®] blade server is a flexible, robust server that is excellent for server consolidation and virtualization. Each blade server has the optional capability of accommodating one or more I/O expansion cards to provide multiple Fibre Channel and/or Ethernet connections to the BladeCenter switch.

The most flexible I/O configuration of an IBM BladeCenter blade server is two Fibre Channel Ports and six Ethernet Ports (2x6). Virtualization solutions, such as VMware, are maximized with dedicated Ethernet ports, while access to the Storage Area Network (SAN) is achieved through the Fibre Channel ports. Future migrations to technologies such as 10Gb/s Ethernet or Fibre Channel over Ethernet (FCoE) are also made easier with the 2x6 I/O configuration since Ethernet and Fibre Channel are connected with dedicated expansion cards. Instead of having I/O topologies (Ethernet & Fibre Channel) on one card, dedicated I/O expansion cards simplify administration and allow for easier management. This Technical Brief describes this preferred 2x6 I/O configuration and the many benefits associated with it.

PRODUCTS NECESSARY

The 2x6 I/O configuration is accomplished with the following two products:

- **Emulex LightPulse Dual-channel 4Gb/s Fibre Channel Expansion Card**—(CFFv) (IBM part number 43W6859). Specifically designed for the IBM Blade Servers.
- **IBM 2/4 Port Ethernet Expansion Card for IBM BladeCenter**—(CFFh) (IBM part number 44W4479). For the BladeCenter S, this expansion card enables two additional Ethernet ports while on the BladeCenter H, this expansion card enables four additional Ethernet ports and this will be the configuration used for this tech brief.

Problem & Solution

The following can be said regarding customers that use IBM blade servers:

- ▶ Applications are generating increasing amounts of data that is seen as a critical resource.
- ▶ This critical data is best accessed through Fibre Channel based storage area networks (SANs) as they have the performance, bandwidth headroom and low CPU overhead that enterprise data centers require.
- ▶ Server consolidation is being used to save energy, save costs and save data center space.
- ▶ Server virtualization on IBM blade servers are being used for this server consolidation.
- ▶ Highly utilized server virtualization (such as VMware) on IBM blade servers perform best with two Fibre Channel and six Ethernet ports.
- ▶ With the release of the IBM 2/4 Port Ethernet expansion card, IBM blade servers can be configured with the ideal configuration (2x6 I/O configuration) for highly utilized VMware applications of 2 Fibre Channel Ports for backend (storage I/O) and 6 Ethernet Ports for front-end (network I/O).

“As trends like multicore and virtualization continue to emerge, clients are looking to get more out of their I/O. The most commonly requested blade I/O topology today is multiple ports of Ethernet (4 or 6 per blade) and Fibre Channel concurrently. Until now the only way to achieve this I/O port count was through the use of combo cards that had Ethernet and Fibre Channel ASICs on the same card. With the introduction of IBM’s 2/4 Port Ethernet expansion card, users can achieve up to 6 Ethernet ports and 2 Fibre Channel ports to each blade—utilizing all currently available CFFv and future channel I/O virtualization style I/O adapter cards. This simple topology makes it easier for clients to get the exact I/O topology they require with extreme flexibility and no wasted ports.”

—**Scott Tease**, worldwide product marketing manager for IBM’s BladeCenter products

2x6 Blade Server I/O Configuration Benefits

The benefits for this 2x6 I/O configuration for an IBM blade server (2 Fibre Channel & 6 Ethernet Ports) is as follows:

- ▶ Server virtualization & VMware implementations highly recommend a high number of dedicated Ethernet ports for services and virtual machines.
- ▶ Dedicated I/O (Fibre Channel & Ethernet separated) allow for easier management and more flexibility in transitioning to new technology.

Server Virtualization & VMware Implementations

Data centers have green computing initiatives to lower energy consumption and save on power costs. Server virtualization and blade servers are key technologies being deployed to reduce capital, management and operating costs. The Fibre Channel expansion card that links servers to storage networks is a critical link in the flow of data. Emulex performance, reliability and virtualization technology empower data centers to more effectively deploy server virtualization, follow SAN best practices, and leverage popular storage management applications. Leveraging blade servers and server virtualization technologies yields more efficient use of assets and a noticeable reduction in TCO. Server virtualization is a key requirement of data centers. Thus, the 2x6 I/O configuration is ideal for virtualization solutions, such as VMware.

Table 1 shows the different I/O configurations (2x2, 2x4 & 2x6) that are available for an IBM blade server, the necessary expansion cards, and when that configuration should be used in a VMware environment.

Table 1—Fibre Channel and Ethernet I/O Configurations

Fibre Channel x Ethernet		
	Expansion Cards Used	Suggested Purpose
2x2	<ul style="list-style-type: none"> · Ethernet with LOM · Fibre Channel with Emulex CFFv Expansion Card 	Small VMware and Limited Applications
2x4	<ul style="list-style-type: none"> · Ethernet with LOM · Fibre Channel and Ethernet on same I/O Card (CFFh) 	Medium Applications with No Future Growth Plans
2x6	<ul style="list-style-type: none"> · Ethernet with LOM · Ethernet with IBM CFFh Expansion Card · Fibre Channel with Emulex CFFv Expansion Card 	(Current or Future) Large VMware, Large Applications or Heavy Network I/O

Table 1 highlights that the 2x6 configuration is for (current or future) large VMware, large application or heavy network I/O blade servers.

Exploring a VMware solution in detail, VMware best practices suggest dedicated Ethernet ports for the following services:

- ▶ Service Console
- ▶ VMkernel
- ▶ Virtual Machines (VM)

Each of these services has different demands on availability, performance and security and thus benefit greatly from dedicated Ethernet ports.

Table 2 shows how a 2x6 I/O configuration should be implemented in a VMware environment.

Table 2—VMware Implementation with IBM and Emulex Expansion Cards

Interface	Location	Service
Ethernet 0	Gigabit Ethernet, onboard	Service Console
Ethernet 1	Gigabit Ethernet, onboard	VMkernel
Ethernet 2	IBM Ethernet Expansion Card (CFFh)	Virtual Machines
Ethernet 3	IBM Ethernet Expansion Card (CFFh)	Virtual Machines
Ethernet 4	IBM Ethernet Expansion Card (CFFh)	Virtual Machines
Ethernet 5	IBM Ethernet Expansion Card (CFFh)	Virtual Machines
Fibre Channel 0	Emulex Fibre Channel Expansion Card (CFFv)	Datastore (VMFS/RDM on SAN)
Fibre Channel 1	Emulex Fibre Channel Expansion Card (CFFv)	Datastore (VMFS/RDM on SAN)

In Table 2, the Service Console and the VMkernel functions have dedicated interfaces which are the two onboard Ethernet ports. The virtual machines are mapped to the four ports on the IBM Ethernet expansion card. Datastores (VMFS) are connected through the two Fibre Channel ports.

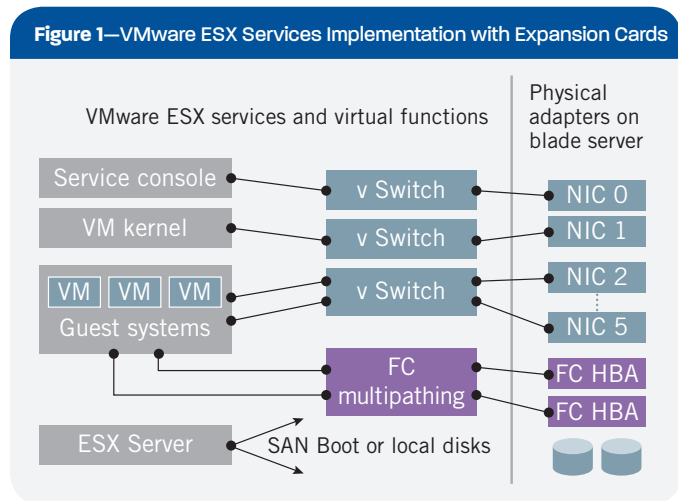
The Service Console is vital for system administration through which the Virtual Infrastructure Client (VI Client) remotely accesses the ESX server. Network performance requirement is very low whereas security might be of prime concern. However, if hardware service consoles with VLANs are implemented on a network and not through a dedicated physical network, this practice can also be applied to the ESX service console.

All configuration tasks for ESX Server are performed through the Service Console including configuring storage, controlling aspects of VM behavior and setting up virtual switches or virtual networks. Because the Service Console is a point of control for ESX Server, safeguarding it from misuse is crucial. A dedicated Ethernet port for the Service Console provides the highest security possible.

Moving a virtual machine from one host to another is called migration. Migrating a VM that is powered on is called VMotion. It is important to have VMotion on a separate Ethernet port because when migration with VMotion occurs, the contents of the guest operating system's memory are transmitted over the network. A dedicated Ethernet port guarantees a separation of VMotion traffic from other administrative traffic or even VMs.

Virtual Machines usually require a great deal of network throughput and low latency in order to increase overall guest system performance. The use of dedicated Ethernet ports is thus recommended. If VMs are for different customers (or departments inside of an organization) then VMs for a particular customer should be on a dedicated Ethernet port.

Figure 1 displays how ESX services and virtual functions relate to the physical adapter cards on blade server.



Availability is always a key concern for all three service categories but external factors, (i.e., availability of the physical network infrastructure in general) should be taken into account. VMware ESX uses virtual switches. Ethernet ports can be assigned to the virtual switches.

High Ethernet demand is generally driven by application requirements and not necessarily VMware requirements.

Dedicated I/O (Fibre Channel & Ethernet)

The IBM Ethernet expansion card allows customers to achieve the network I/O bandwidth necessary for high demand user and device connectivity. Cabling is easier to understand with four Ethernet ports in the Multi-Switch Interconnect Module (MSIM) slots rather than two in the MSIM and two in standard bays. According to best practices, the two onboard Ethernet ports (LOM) would be used for management such as service console and VMkernel. Also according to best practices, the four remaining Ethernet ports all go to the MSIM slots, which would be used for the virtual machines (VMs). This would allow administrators to understand and trace cables more easily and thus being easier to administer. Combining Ethernet and Fibre Channel on one expansion card (shared I/O configuration) complicates management.

Transition to New Technology

Separating I/O (Fibre Channel and Ethernet) topologies on different expansion cards supports IT transition strategies to newer technologies. For example, dedicated I/O allows organizations to upgrade to 10Gb/s Ethernet or 8Gb/s Fibre Channel more cost-effectively than a “rip and replace” process that is associated with shared I/O on one card. That is, when organizations want to upgrade to 10Gb/s Ethernet or 8Gb/s Fibre Channel independent of each other, they can just replace one card instead of purchasing two cards. As organizations evaluate FCoE, those who implemented the 2x6 I/O configuration have the comfort of knowing that migration to FCoE will not incur unnecessary costs. It is not best practices to implement two new technologies at the same time. Therefore, with I/O topologies on dedicated I/O expansion cards, new technologies can be implemented at different times.

Background Information (CFFv, CFFh, Bays)

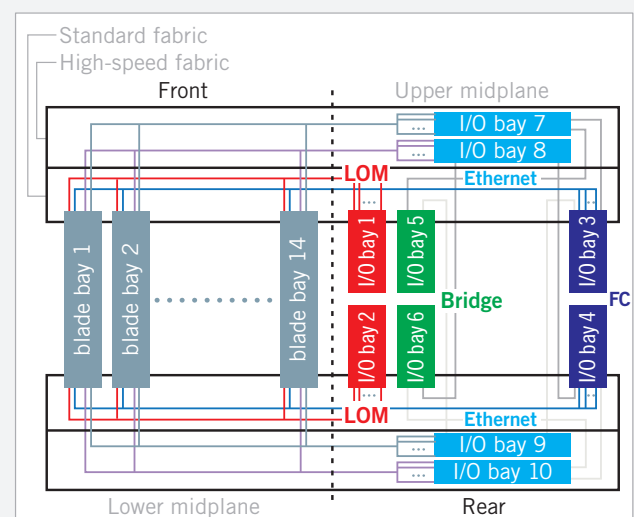
An IBM BladeCenter blade server connects I/O expansion cards through either a Compact Form Factor Horizontal/High Speed (CFFh) method or a Compact Form Factor Vertical/Standard (CFFv) method. The CFFv uses the PCI-X connector on the blade server. The CFFh attaches to the PCI Express connector of the blade server.

The IBM Ethernet Expansion Card for BladeCenter is a CFFh Ethernet card with four Ethernet ports. This expansion card complements the two Ethernet ports contained on the motherboard (LOM).

The Emulex Fibre Channel expansion card for BladeCenter is a CFFv card which, when combined with the IBM CFFh 2/4 Ethernet Port Expansion Card, delivers a flexible and powerful 2x6 I/O configuration where there are two Fibre Channel ports to access the SAN and six Ethernet ports for high performance applications. With only one location for each CFFv and CFFh card, the 2x6 I/O configuration delivers the maximum number of Fibre Channel and Ethernet ports available today.

IBM BladeCenter H chassis delivers high performance through a total of ten hot-swap I/O bays shown in Figure 2.

Figure 2—IBM BladeCenter H Connection Paths



The I/O bays are separated into three groups which are:

- ▶ **Four standard bays** (bays 1 thru 4)
- ▶ **Two bridge bays** (bays 5 and 6)
- ▶ **Four high-speed bays** (bays 7 thru 10)

The bays are used as follows:

- ▶ **Bays 1 and 2** support standard Ethernet-compatible I/O modules only. These are routed internally to the onboard Ethernet controllers on the blade servers.
- ▶ **Bays 3 and 4** can be used either for standard switch or pass-through modules (such as Fibre Channel connectivity) or for bridge modules. These are routed internally to the PCI-X connector on the blade servers.
- ▶ **Bays 5 and 6** are dedicated for bridge modules only and do not directly connect to the blade bays. Bridge modules provide links to the I/O bays 7 thru 10 and can be used as additional outputs for I/O modules in those bays. In case I/O bays 3 and 4 are used for bridge modules, they are not directly connected to the blade servers, and bay 3 provides redundancy for bay 5, and bay 4 provides redundancy for bay 6.
- ▶ **Bays 7 through 10** are used for high-speed switch modules (Ethernet). These are routed internally to the PCI Express connector on blade servers that have it.

The Emulex Fibre Channel and the IBM Ethernet I/O expansion cards allow connectivity with both standard and high-speed switch module bays at the same time. The CFFv (vertical/standard) expansion card connects to the blade server PCI-X bus to provide I/O connections to standard switch I/O bays 3 and 4. The CFFh (horizontal/high speed) expansion card connects to the blade server PCI-Express bus in order to provide I/O connections to the high-speed switch modules in I/O bays 7 through 10. Each blade server bay has a total of eight dedicated connection paths to the I/O modules (six for Ethernet and two for Fibre Channel).

The IBM CFFh Ethernet Expansion Card is used in conjunction with the Multi-Switch Interconnect Module (MSIM) (IBM Part No. 39Y9314), and requires in the MSIM a supported Ethernet switch module installed in the left bay for the Ethernet connections. These are the high speed bays (bays 7 thru 10). Each MSIM can house two switch modules. IBM BladeCenter H chassis can support two MSIMs.

Emulex Fibre Channel Expansion Card Benefits

The Emulex Fibre Channel expansion card provides reliable, high performance, low CPU overhead 4Gb/s connectivity, enabling high-availability access to scalable storage in order to meet the needs of the most demanding applications and environments.

Following are key benefits of using the Emulex Fibre Channel expansion card:

1 Compatibility and Investment Protection—Like IBM, Emulex delivers confidence that today's purchases will be valuable in the years to come due to:

- Fibre Channel expansion card has full driver compatibility with:
 - All Emulex Host Bus Adapters (HBAs) for standard and blade servers
 - Forward compatibility with Emulex Converged Network Adapters (CNAs)
- One driver on each operating system supports all generations of Emulex adapters
- Adaptive firmware-based architecture that simplifies upgrades and provides long term investment protection. **There are no server re-boots required during firmware upgrades.**

2 Reduced Complexity—The Emulex Fibre Channel expansion card is easy to install and manage (minimizing administration costs) due to:

- Emulex's common driver model reduces management complexities and streamlines installation and upgrades.
- Centralized management with Emulex's HBAnyware management application enables management of all Emulex Fibre Channel expansion cards, HBAs & CNAs across the SAN. HBAnyware incorporates driver-based technology for automated discovery and control. Emulex Fibre Channel expansion card includes time- and cost-saving functionality, such as firmware upgrade, driver reconfiguration and persistent bindings. HBAnyware allows Emulex adapters to be managed from any console within the SAN.

3 Superior Quality and Reliability—Emulex Fibre Channel expansion card has high reliability gained from the experience of eight generations of industry-leading Emulex host bus adapters (HBAs).

4 Fast Diagnosis and Recovery—Comprehensive diagnostic functions, coupled with detailed event logging and tracing, provide for fast and efficient SAN troubleshooting.

5 Seamless Integration into IBM Director—Adherence to the Storage Management Initiative - Specification (SMI-S) standard and common HBA API enable upward integration into IBM Director, an integrated suite of system management tools.

Due to these benefits, the Emulex Fibre Channel expansion card is an excellent choice for data centers that use IBM blade servers.

Links

Detailed VMware planning and configuration information can be found in the “Server Configuration Guide” from VMware:

http://www.vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_3_server_config.pdf

IBM Hardware Announcement for 2/4 Port Ethernet Expansion Card (CFFh) for IBM BladeCenter provides multiple interfaces with blades:

http://www-01.ibm.com/common/ssi/rep_ca/1/897/ENUS108-601/ENUS108-601.PDF

IBM HBA Interoperability Search Tool:

<http://www-03.ibm.com/servers/storage/support/config/hba/index.wss>

Configurator Tools:

<http://www-03.ibm.com/servers/eserver/xseries/library/config.tools.html>

Emulex information on HBAs for IBM:

<http://www.emulex.com/ibm/index.jsp>

Summary

Emulex and IBM together deliver a 2x6 I/O configuration consisting of two Fibre Channel ports and six Ethernet ports that is excellent for highly utilized virtualized IBM blade servers.

This 2x6 I/O configuration maximizes the flexibility required for data center operations, particularly for server virtualization, especially with VMware which recommends 6 Ethernet ports.

The separation between the I/O topologies (Fibre Channel and Ethernet) allows IT management to transition independently to FCoE or 10Gb/s Ethernet without requiring to implement both at the same time.

This separation of Fibre Channel and Ethernet onto different expansion cards also allows for easier administration.

For high performance IBM blade servers, the combination of IBM and Emulex Expansion Cards delivers the powerful solution that data centers require.

This document refers to various companies and products by their trade names. In most, if not all cases, their respective companies claim these designations as trademarks or registered trademarks. This information is provided for reference only. Although this information is believed to be accurate and reliable at the time of publication, Emulex assumes no responsibility for errors or omissions. Emulex reserves the right to make changes or corrections without notice. This report is the property of Emulex and may not be duplicated without permission from the Company.

09-179 · 8/08



www.emulex.com

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