

Emulex 10GbE Virtual Fabric Adapter II for IBM BladeCenter

IBM Redbooks Product Guide

The Emulex 10 GbE Virtual Fabric Adapter II and Emulex 10 GbE Virtual Fabric Adapter Advanced II are enhancements to the existing IBM® BladeCenter® Virtual Fabric portfolio. These adapters allow IBM clients to simplify their I/O infrastructure by reducing the number of switches needed inside the chassis. The Emulex 10GbE Virtual Fabric Adapter II is a dual-port 10 Gb Ethernet card that supports 1 Gbps or 10 Gbps traffic, or up to eight virtual NIC (vNIC) devices.

The virtual NICs are configured to meet the client's mix of network connectivity and throughput demands for today's complex server application environments. Each physical 10 Gbps port can be divided into four virtual ports with bandwidth allocation in 100 Mbps increments to the maximum 10 Gbps per physical port. The Emulex 10GbE Virtual Fabric Adapter II Advanced adds FCoE and iSCSI hardware initiator functionality.

The Emulex 10 GbE Virtual Fabric Adapter II supports the IBM Virtual Fabric 10Gb Switch Module in IBM Virtual Fabric Mode (vNIC1). The card supports Switch Independent Mode (vNIC2) with other IBM BladeCenter 10 GbE switch offerings from Cisco and Brocade, and the 10Gb Ethernet Pass-Thru Module for IBM BladeCenter. Figure 1 shows the adapter.



Figure 1. Emulex 10GbE Virtual Fabric Adapter II for IBM BladeCenter

Did you know?

This adapter is based on industry-standard PCIe architecture and offers the flexibility to operate as a Virtual NIC Fabric Adapter or as a dual-port 10Gb or 1Gb Ethernet device. Because this adapter supports up to eight virtual NICs on a single dual-port Ethernet adapter, you see benefits in cost, power/cooling, and data center footprint by deploying less hardware. The adapter enables you to configure up to 12 NIC devices on a single blade server: eight virtual NICs with the Emulex 10 GbE Virtual Fabric Adapter II plus the two onboard NICs plus two from a CIOv expansion card.

Part number information

Table 1. Ordering part number and feature code

Description	Part number	Feature code
Emulex 10GbE Virtual Fabric Adapter II - IBM BladeCenter	90Y3550	A1XG
Emulex 10GbE Virtual Fabric Adapter II Advanced - IBM BladeCenter	90Y3566	AIXH
Emulex 10GbE Virtual Fabric Adapter II - IBM BladeCenter	00Y3266*	A3NV
Emulex 10GbE Virtual Fabric Adapter Advanced II - IBM BladeCenter	00Y3264*	A3NW
Emulex 10GbE Virtual Fabric Advanced Upgrade for IBM BladeCenter	49Y4265	2436

* These options are refreshed offerings of the existing VFA II adapters. The new adapters support pNIC, vNIC1, and vNIC2 out-of-box, while the existing adapters require the firmware upgrade to support vNIC2 functionality (pNIC and vNIC1 is supported out-of-box).

The adapter part number includes the following items:

- One Emulex 10GbE Virtual Fabric Adapter II or Emulex 10GbE Virtual Fabric Adapter II Advanced
- Documentation CD
- Important Notices flyer

The standard Emulex 10GbE Virtual Fabric Adapter II can be upgraded to the same features as the Emulex 10GbE Virtual Fabric Adapter II Advanced with the addition of the Emulex 10GbE Virtual Fabric Advanced Upgrade option, part number 49Y4265.

Features

Emulex Virtual Fabric Adapter II and Emulex Virtual Fabric Adapter Advanced II have these features:

- Connection to either 1 Gb or 10 Gb data center infrastructure (1 Gb and 10 Gb auto-negotiation)
- PCI Express 2.0 x8 host interface
- CFFh form factor, can be combined with a CIOv expansion card on the same blade server
- It operates either as an 8-port virtual NIC or as a 2-port 1/10 Gb Ethernet adapter:
 - Virtual port bandwidth allocation in 100 Mbps increments
 - Supports up to eight virtual ports
- Wake On LAN support
- Full-duplex (FDX) capability
- Bus-mastering support
- Direct memory access (DMA) support
- Preboot Execution Environment (PXE) support
- IPv4/IPv6 TCP, UDP checksum offload
 - Large send offload (LSO)
 - Large receive offload
 - Receive side scaling (RSS)
 - IPv4 TCP Chimney Offload
 - TCP Segmentation Offload
- VLAN insertion and extraction

- Jumbo frames up to 9000 bytes
- Load balancing and failover support including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), teaming support, and IEEE 802.3ad
- Enhanced Ethernet (draft)
 - Enhanced Transmission Selection (ETS) (P802.1Qaz)
 - Priority-based Flow Control (PFC) (P802.1Qbb)
 - Data Center Bridging Capabilities eXchange Protocol, CIN-DCBX and CEE-DCBX (P802.1Qaz)
- The Emulex Virtual Fabric Adapter Advanced II provides FCoE and iSCSI hardware initiator support.
- Supports Serial over LAN (SoL) and concurrent KVM (cKVM)
- Supports two vNIC modes: IBM Virtual Fabric Mode and Switch Independent Mode.

Modes of operation

The original Emulex 10GbE Virtual Fabric Adapter can be configured in pNIC or vNIC mode. The new Emulex 10GbE Virtual Fabric Adapter II now offers two vNIC modes of operation in addition to pNIC mode:

- IBM Virtual Fabric Mode, which is the same as vNIC mode on original Emulex 10GbE Virtual Fabric Adapter. This mode only works in conjunction with IBM Virtual Fabric 10Gb Switch Module installed in I/O module bays 7 and 9 in the chassis. In this mode, the Emulex adapter communicated with the switch module to obtain vNIC parameters (using DCBX). Also, a special tag within each data packet is added and later removed by the NIC and switch for each vNIC group to maintain separation of the virtual channels.

In vNIC mode, each physical port is divided into four virtual ports providing a total of eight virtual NICs per adapter. The default bandwidth for each vNIC is 2.5 Gbps. Bandwidth for each vNIC can be configured at the IBM switch from 100 Mbps to 10 Gbps, up to a total of 10 Gb per physical port. The vNICs can also be configured to have 0 bandwidth if you must allocate the available bandwidth to fewer than eight vNICs. In IBM Virtual Fabric Mode, you can change the bandwidth allocations through the IBM switch user interfaces without requiring a reboot of the server.

When storage protocols are enabled on the Emulex 10GbE Virtual Fabric Adapter Advanced II, six ports are Ethernet, and two ports are either iSCSI or FCoE

- Switch Independent Mode, where the adapter works with the following switches:
 - IBM Virtual Fabric 10Gb Switch Module
 - Cisco Nexus 4001I Switch Module
 - Brocade Converged 10GbE Switch Module
 - 10Gb Ethernet Pass-Thru Module and a top-of-rack (TOR) switch

These switches must be installed in I/O module bays 7 and 9 in the chassis. Beyond the way that the adapter communicates with the switch, Switch Independent Mode offers the same capabilities as IBM Virtual Fabric Mode in terms of the number of vNICs and the bandwidth each can be configured to have. Switch Independent Mode extends the existing customer VLANs to the virtual NIC interfaces. The IEEE 802.1Q VLAN tag is essential to the separation of the vNIC groups by the NIC adapter or driver and the switch. The VLAN tags are added to the packet by the applications or drivers at each end station rather than by the switch.

- In pNIC mode, the expansion card can operate as a standard 10 Gbps or 1 Gbps 2-port Ethernet expansion card. When in pNIC mode, the expansion card functions with any available 10 Gb switch or 10 Gb pass-thru module installed in I/O module bays 7 and 9.

In pNIC mode, the Advanced adapter or the Advanced upgrade operates in traditional Converged Network Adapter (CNA) mode with two Ethernet ports and two storage ports (iSCSI or FCoE) available to the operating system..

Supported servers

The Emulex 10GbE Virtual Fabric Adapter II and Emulex 10GbE Virtual Fabric Adapter II Advanced are supported in the IBM BladeCenter servers listed in Table 2.

Table 2. Supported BladeCenter servers

Part number	Feature code	Product description	HS12 (8028)	HS22 (7870)	HS22V (7871)	HS23 (7875)	HS23E (8038)	HX5 (7873)
90Y3550	A1XG	Emulex 10GbE Virtual Fabric Adapter II	N	Y	Y	N	Y	Y
90Y3566	A1XH	Emulex 10GbE Virtual Fabric Adapter Advanced II	N	Y	Y	N	Y	Y
00Y3266	A3NV	Emulex 10GbE Virtual Fabric Adapter II	N	Y	Y	N	Y	Y
00Y3264	A3NW	Emulex 10GbE Virtual Fabric Adapter Advanced II	N	Y	Y	N	Y	Y
49Y4265	2436	Emulex 10GbE Virtual Fabric Advanced Upgrade	N	Y	Y	N	Y	Y

Figure 2 shows where the CFFh card is installed in a BladeCenter server.




Figure 2. Location on the BladeCenter server planar where the CFFh card is installed

Supported BladeCenter chassis and I/O modules

IBM BladeCenter chassis support is based on the blade server type in which the expansion card is installed. Consult the IBM ServerProven® website to see in which chassis each blade server type is supported: <http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/>

Table 3 lists the supported chassis and I/O module combinations that the Emulex 10GbE Virtual Fabric Adapter II supports, in either standard physical port mode (pNIC) or one of the two virtual NIC (vNIC) modes, IBM Virtual Fabric Mode or Switch Independent Mode.

Table 3. I/O modules supported with the Emulex 10GbE Virtual Fabric Adapter II

Description	Part number	BladeCenter S	BladeCenter E	BladeCenter H	BladeCenter T	BladeCenter HT	MSIM	MSIM-HT
IBM Virtual Fabric (vNIC1) Mode								
IBM Virtual Fabric 10Gb Switch Module	46C7191	N	N	Y	N	Y	N	N
BNT 6-port 10 Gb High Speed Switch	39Y9267	N	N	N	N	N	N	N
10Gb Ethernet Pass-Thru Module	46M6181	N	N	N	N	N	N	N
Cisco Nexus 4001I Switch Module	46M6071	N	N	N	N	N	N	N
Brocade Converged 10GbE Switch Module	69Y1909	N	N	N	N	N	N	N
Switch Independent (vNIC2) Mode								
IBM Virtual Fabric 10Gb Switch Module	46C7191	N	N	Y	N	Y	N	N
BNT 6-port 10 Gb High Speed Switch	39Y9267	N	N	N	N	N	N	N
10Gb Ethernet Pass-Thru Module	46M6181	N	N	Y	N	Y	N	N
Cisco Nexus 4001I Switch Module	46M6071	N	N	Y	N	Y	N	N
Brocade Converged 10GbE Switch Module	69Y1909	N	N	Y	N	Y	N	N
pNIC mode								
IBM Virtual Fabric 10Gb Switch Module	46C7191	N	N	Y	N	Y	N	N
BNT 6-port 10 Gb High Speed Switch	39Y9267	N	N	Y	N	Y	N	N
10Gb Ethernet Pass-Thru Module	46M6181	N	N	Y	N	Y	N	N
Cisco Nexus 4001I Switch Module	46M6071	N	N	Y	N	Y	N	N
Brocade Converged 10GbE Switch Module	69Y1909	N	N	Y	N	Y	N	N

In BladeCenter H, the ports of CFFh cards are routed through the midplane to I/O bays 7, 8, 9, and 10, as shown in Figure 3. The BladeCenter HT is similar in that the CFFh cards are also routed through the midplane to I/O bays 7, 8, 9, and 10.

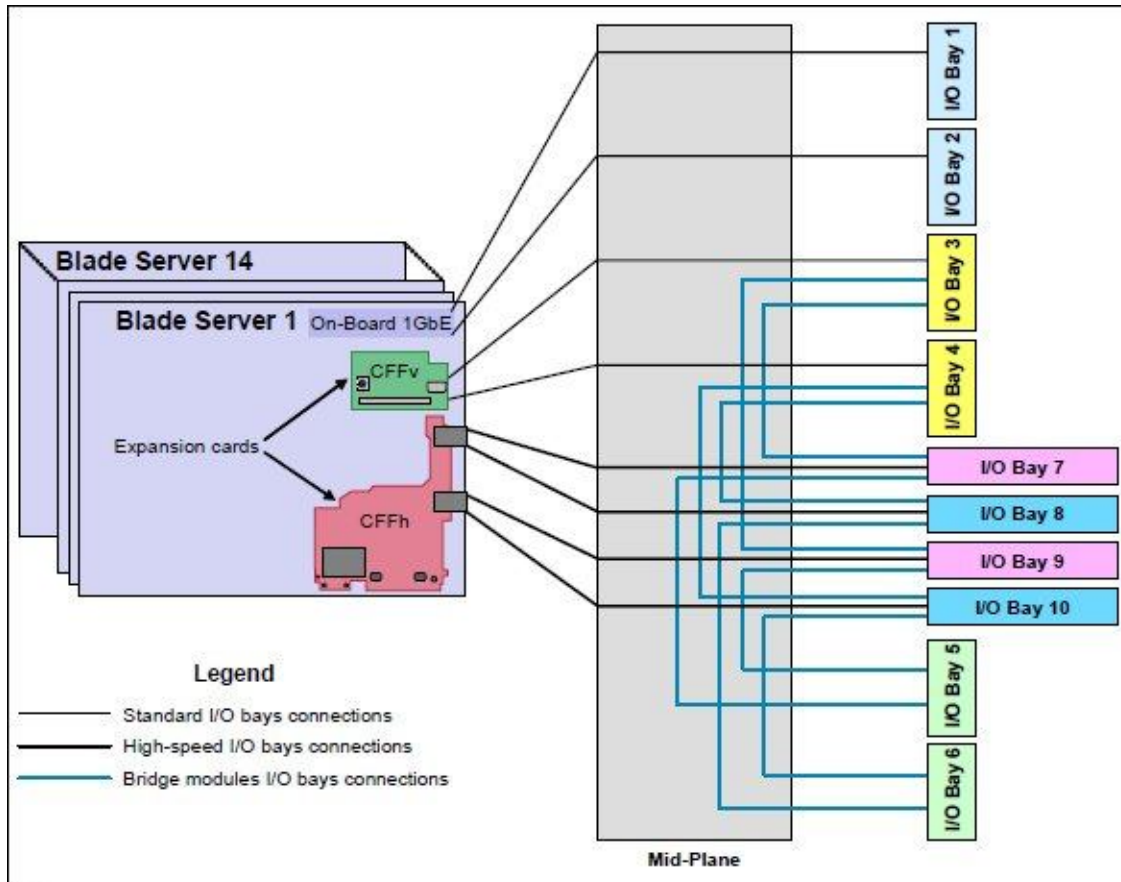


Figure 3. IBM BladeCenter H I/O topology showing the I/O paths from CFFh expansion cards

The Emulex 10GbE Virtual Fabric Adapter II and Emulex 10GbE Virtual Fabric Adapter II Advanced require that two I/O modules be installed in bays 7 and 9 of the BladeCenter H and HT chassis (Table 4).

Table 4. Locations of I/O modules required to connect to the expansion card

Expansion card	I/O bay 7	I/O bay 8	I/O bay 9	I/O bay 10
Emulex 10GbE Virtual Fabric Adapter II	Supported I/O module	Not used	Supported I/O module	Not used
Emulex 10GbE Virtual Fabric Adapter II Advanced	Supported I/O module	Not used	Supported I/O module	Not used

MAC addresses are assigned to the expansion card. In Virtual Fabric Adapter mode (vNIC mode), each of the eight virtual ports is assigned a MAC address. In 2-port 10 Gb Ethernet mode (pNIC mode), each of the two physical ports is assigned a MAC addresses.

Popular configurations

This section shows a configuration that uses the Emulex 10GbE Virtual Fabric Adapter II.

Figure 4 shows the use of the Emulex 10GbE Virtual Fabric Adapter II to route either two physical ports or eight virtual ports to the IBM Virtual Fabric 10Gb Switch Modules installed in the BladeCenter chassis. Two IBM Virtual Fabric 10Gb Switch Modules are installed in bay 7 and bay 9 of the BladeCenter H chassis. All connections between the controller, card, and the switch modules are internal to the chassis. No cabling is needed.




Figure 4. Configuration using the Emulex 10GbE Virtual Fabric Adapter II

Depending on how you configure the switch and expansion card, you can get either of these combinations:

- Eight virtual Ethernet ports (either vNIC mode). Each virtual port can be configured to operate at 100 Mbps to 10 Gbps, up to a total maximum of 10 Gb per physical port. The default bandwidth is 2.5 Gb per virtual port.
- Two 1 Gb or 10 Gb Ethernet ports (pNIC mode).

The components used in this configuration are listed in Table 5.

Table 5. Components used when connecting the Emulex adapter to two IBM switch modules

Diagram reference	Part number/ machine type	Description	Quantity
1	7870	IBM BladeCenter HS22	1 to 14
2	90Y3550	Emulex 10GbE Virtual Fabric Adapter II	1 per server
3	8852 or 8740/8750	BladeCenter H or BladeCenter HT	1
4	46C7191	IBM Virtual Fabric 10Gb Switch Module	2
5	44W4408	IBM 10GBase-SR SFP+ Transceiver	Up to 20*

* To communicate outside the chassis, you must have either one SFP+ transceiver or SFP+ direct-attached cable (DAC) for each of the external 10 Gb ports on the IBM 10Gb High Speed Switch Module that you wish to enable.

Operating environment

This adapter complies with temperature rating requirements for BladeCenter H and BladeCenter HT chassis and is supported in the following environment:

- Temperature - Operating:
 - 10 to 52° C (50 to 125° F) at an altitude of 0 to 914 m (0 to 3,000 ft)
 - 10 to 49° C (50 to 120° F) at an altitude of 0 to 3,000 m (0 to 10,000 ft)
- Temperature - Non-operating:
 - 40 to 65°C (-40 to 149° F) at an altitude of 0 to 12,000 m (0 to 39,370 ft)
- Humidity
 - Operating: 8 to 80%, noncondensing
 - Non-operating: 5 to 80%, noncondensing

Operating system support

The Emulex 10GbE Virtual Fabric Adapter II supports the following operating systems:

- Microsoft Windows Server 2003/2003 R2, Datacenter Edition
- Microsoft Windows Server 2003/2003 R2, Datacenter x64 Edition
- Microsoft Windows Server 2003/2003 R2, Enterprise Edition
- Microsoft Windows Server 2003/2003 R2, Enterprise x64 Edition
- Microsoft Windows Server 2003/2003 R2, Standard Edition
- Microsoft Windows Server 2003/2003 R2, Standard x64 Edition
- Microsoft Windows Server 2008 R2
- Microsoft Windows Server 2008, Datacenter x64 Edition
- Microsoft Windows Server 2008, Datacenter x86 Edition
- Microsoft Windows Server 2008, Enterprise x64 Edition
- Microsoft Windows Server 2008, Enterprise x86 Edition
- Microsoft Windows Server 2008, Standard x64 Edition
- Microsoft Windows Server 2008, Standard x86 Edition
- Microsoft Windows Server 2008, Web x64 Edition
- Microsoft Windows Server 2008, Web x86 Edition
- Red Hat Enterprise Linux 5 Server Edition
- Red Hat Enterprise Linux 5 Server x64 Edition
- Red Hat Enterprise Linux 6 Server Edition
- Red Hat Enterprise Linux 6 Server x64 Edition
- SUSE LINUX Enterprise Real Time 10 AMD64/EM64T
- SUSE LINUX Enterprise Server 10 for AMD64/EM64T
- SUSE LINUX Enterprise Server 10 for x86
- SUSE LINUX Enterprise Server 11 for AMD64/EM64T
- SUSE LINUX Enterprise Server 11 for x86
- Solaris 10 Operating System
- VMware ESX 3.5
- VMware ESX 4.0
- VMware ESX 4.1
- VMware ESXi 4.1
- VMware vSphere 5

Support for operating systems is based on the combination of the expansion card and the blade server on which it is installed. See the IBM ServerProven Web site for the latest information about the specific versions and service packs supported. Select the blade server, and then select the expansion card to see the supported operating systems. <http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/>

Related publications

For more information refer to the following resources:

- *Emulex 10GbE Virtual Fabric Adapter II for IBM BladeCenter Installation and User's Guide*
<http://www.ibm.com/support>
- BladeCenter Ethernet home page
<http://ibm.com/systems/bladecenter/hardware/openfabric/ethernet.html>
- IBM RETAIN® tip H197476: vNIC support - Emulex 10 Gigabit Ethernet Virtual Fabric Adapters
<http://ibm.com/support/entry/portal/docdisplay?Indocid=MIGR-5085659>
- Emulex product page for the 10GbE Virtual Fabric Adapter
<http://www.emulex.com/solutions/ibm-bladecenter-virtual-fabric/emulex-virtual-fabric-adapter.html>
- *IBM BladeCenter Interoperability Guide*
<http://www.redbooks.ibm.com/big>
- IBM Redbooks® publication *IBM BladeCenter Products and Technology*, SG24-7523
<http://www.redbooks.ibm.com/abstracts/sg247523.html>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2011. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document was created or updated on June 28, 2013.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbook@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips0828.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>. The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter®
IBM®
Redbooks®
Redbooks (logo)®
RETAIN®
ServerProven®

The following terms are trademarks of other companies:

BNT, and Server Mobility are trademarks or registered trademarks of Blade Network Technologies, Inc., an IBM Company.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.