

Deployment of VMware ESX 2.5.1 Server Software on Dell PowerEdge Blade Servers

The purpose of this document is to provide best practices for deploying VMware ESX 2.5.1 Server Software on Dell PowerEdge Blade Servers. The document describes various possible configurations when installing ESX Server Software along with the advantages and disadvantages for each configuration.



1. Introduction

This document describes various configurations that are possible when installing VMware® ESX 2.5.1 Server Software on Dell™ PowerEdge™ 1855 Blade Servers. Advantages and disadvantages for each configuration along with the steps to install are detailed.

2. Configurations Options

2.1 Stand-alone and SAN Configurations

Dell has qualified VMware ESX Server 2.5.1 software in two basic configurations. Blade servers can be used in a “stand-alone” mode (figure 1), or they can be connected to a shared SAN through a switch.

❑ PE 1855 Dual-Port Gigabit Ethernet daughter card is not supported with ESX configurations in this release.

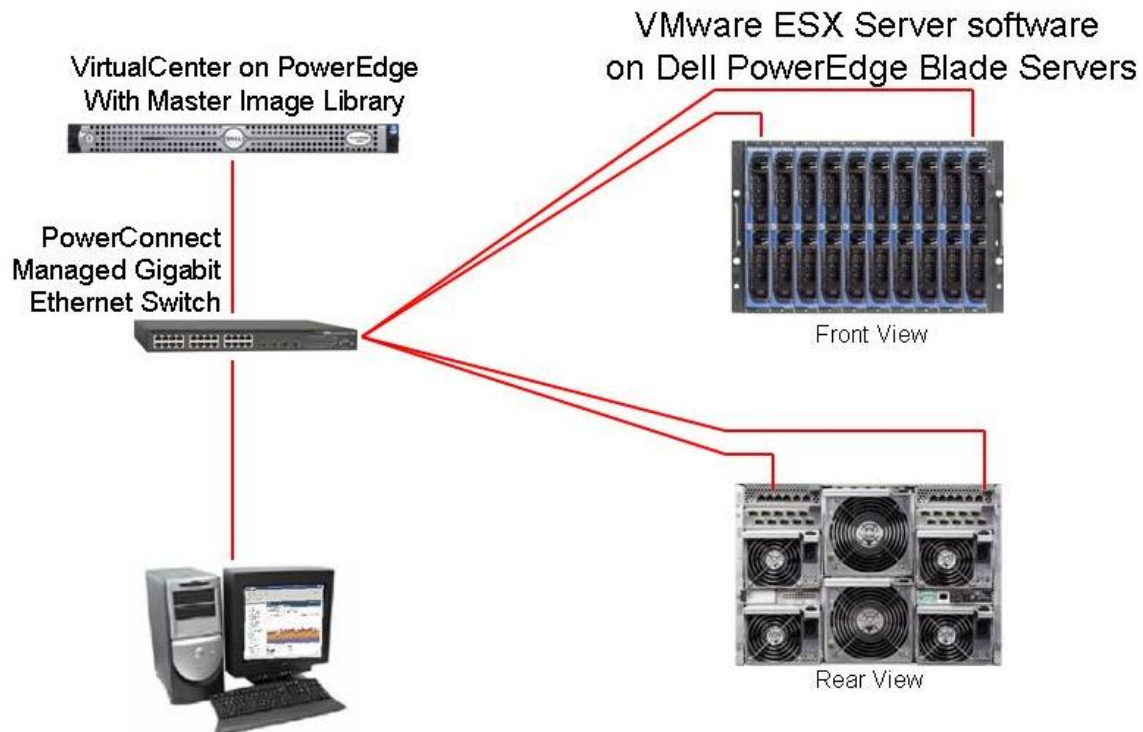


Figure 1: VMware ESX Server Software - Stand-Alone Configuration on PowerEdge Blade Servers

For SAN configurations, a Fibre Channel switch is required. Direct attached Fibre Channel is not supported. The Fibre Channel switch can either be the embedded Brocade SilkWorm 3014 Fibre Channel Switch Module or an external Fibre Channel switch which

is connected through a Pass-Through Fibre Channel Module. The differences between the Brocade SilkWorm 3014 Fibre Channel Switch Module and the Pass-Through Fibre Channel Module are listed below.

Table 2: Fibre Channel switch module Vs Pass-through module

Brocade SilkWorm 3014 Fibre Channel Switch Module	Pass-Through Fibre Channel Module
<ul style="list-style-type: none"> ▪ Features four uplinks, that are 1Gb/s and 2 Gb/s auto-sensing ▪ Oversubscribed 10 to 4 (ten links to the blade server and four external up-links) ▪ Saves on rack space and cabling 	<ul style="list-style-type: none"> ▪ Features ten 2 Gb/s uplinks. ▪ Provides maximum Fibre Channel bandwidth and the ability to seamlessly utilize existing Fibre Channel investments ▪ Requires additional external switches, rack space and cabling

The different qualified SAN configurations are shown below:

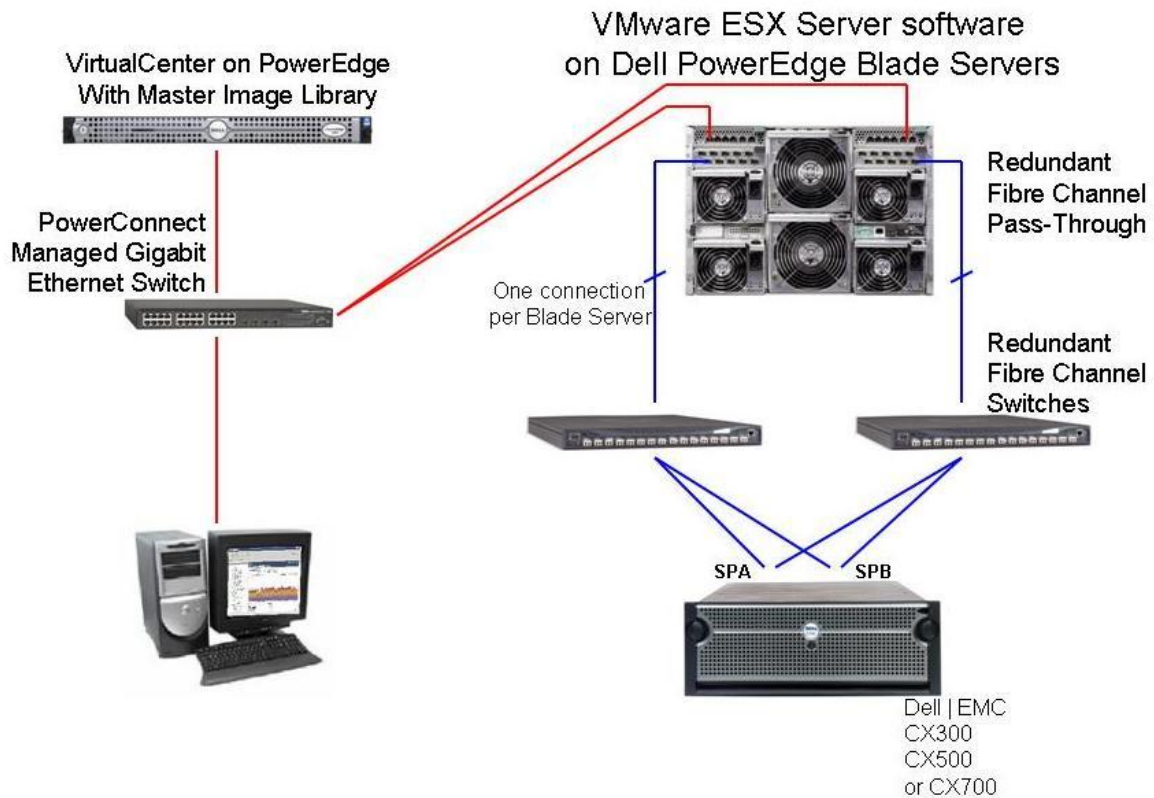


Figure 2: VMware ESX Server Software - SAN Configuration on PowerEdge Blade Servers with Pass-Through Fibre Channel Modules

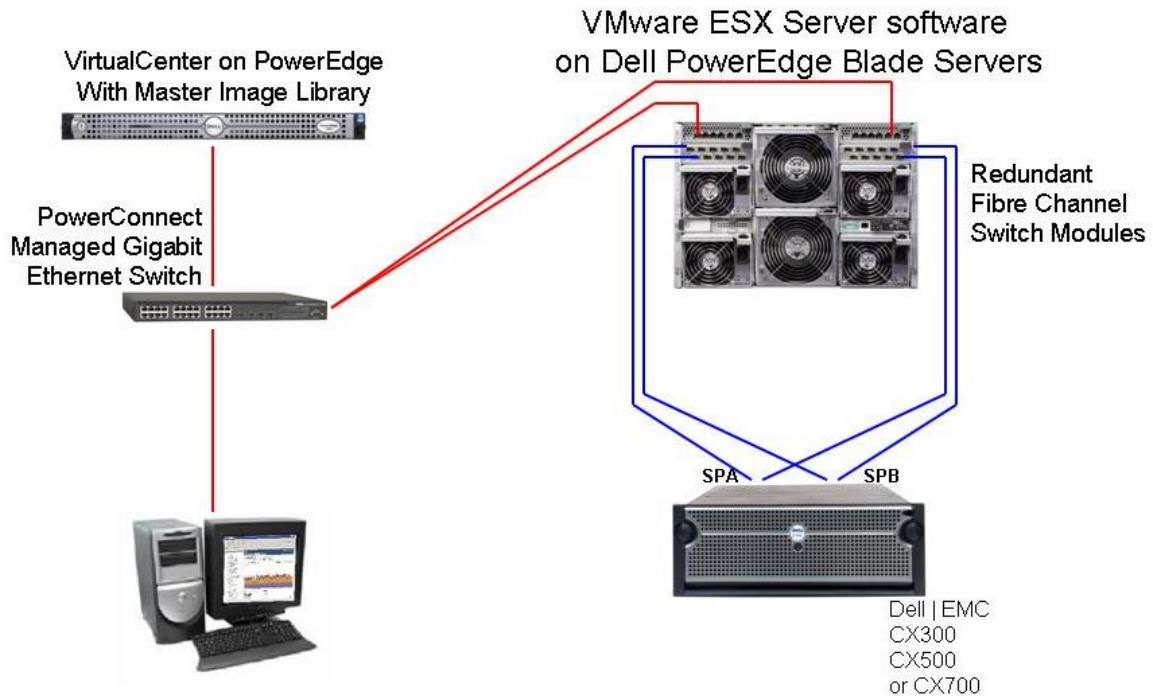


Figure 3: VMware ESX Server Software - SAN Configuration on PowerEdge Blade Servers with Fibre Channel Switch Modules




Figure 4: VMware ESX Server Software - SAN Configuration on PowerEdge Blade Servers with Fibre Channel Switch Modules and external Fibre Channel Switches

■ A minimum of two Inter Switch Links (ISL) between two directly communicating switches is recommended.

2.2 Configuring the NICs

Each PowerEdge (PE) 1855 blade server comes with two embedded Gigabit¹ Ethernet NICs (NIC0 and NIC1). The NICs have to be either dedicated to the Service Console, dedicated to the Virtual Machines or shared between the Service Console and Virtual Machines. In a default installation, NIC0 is dedicated to the Service Console and NIC1 is dedicated to the Virtual Machine.

In order to better utilize the network bandwidth and to provide redundancy, the NICs can be configured in the following ways:

Table 2: NIC Configuration

Config	NIC 0	NIC 1	Fault Tolerance for VMs	Installation	Traffic Isolation?	VM Performance?	Security?
1	Service Console	VMs and VMotion	No	Standard	Moderate	Acceptable if VMotion events are infrequent	Moderate - VMotion on production network.
2	Service Console and VMotion (shared)	VMs	No	Command line post-Installation steps	Good	Acceptable if management traffic is infrequent	Good - service console on private network.
3	Service Console and VMs (shared)	VMotion	No	Command line post-Installation steps	Moderate	Acceptable if management traffic is infrequent and VMotion events are frequent	Moderate - service console on production network.
4	NIC 0 shared between Service Console and VMs. NIC 0 and NIC 1 are teamed and used by VMs and VMotion		Yes	Command line post-Installation steps	Moderate and flexible	Acceptable if VMotion events are infrequent.	Poor – VMotion and Service console on production network.

Configurations 2, 3 and 4 require sharing of NIC 0 between the service console and Virtual Machines.

■ Service console traffic is encrypted by default.

■ VMotion traffic is not encrypted. In order to help ensure effective and secure VMotion, it is highly recommended to have the VMotion NIC on a separate VLAN or separate physical NIC.

2.3 Boot from SAN

ESX 2.5.1 Server Software can be installed and booted from a LUN in a SAN. This enables diskless blade servers and isolates and simplifies storage management. For more information, refer to the *ESX 2.5 Server SAN Configuration Guide* at www.vmware.com.

2.4 Embedded Switch vs Pass-through with PE 1855

Customers can configure a PE 1855 chassis either with a PowerConnect™ 5316M Ethernet switch module or with a pass-through Ethernet module. There are no special installation steps required for VMware ESX Server software to use the pass-through module. Steps to configure the ESX Server software and the switch module to attain fault tolerance along with VLAN setup are described in the Post Installation Section. The table below lists the advantages of using either the switch or pass-through configurations.

Table 3: Switch vs Pass-through module

PowerConnect 5316M Ethernet Switch Module	Pass-Through Ethernet Module
<ul style="list-style-type: none"> ▪ Features six 10/100/1000BaseT Ethernet uplinks, providing full Layer 2 functionality ▪ Oversubscribed 10 to 6 (ten links to the blade server and six external up-links) ▪ VMotion traffic between blade server in the same chassis is enclosed within the switch ▪ Saves on rack space and cabling 	<ul style="list-style-type: none"> ▪ Features ten 1000BaseT Ethernet uplinks. ▪ Provides maximum network bandwidth and the ability to seamlessly utilize existing network investments ▪ VMotion traffic affects external switches connected to the chassis ▪ Requires additional external switches, rack space and cabling

2.5 Virtual Center with PE 1855

The PE 1855 chassis can be configured with up to 10 blade servers. One of the blade servers can be installed with Virtual Center and used as a management console to manage virtual machines in other blade servers.

3. Installing ESX Server Software

This section gives an overview of different ways that VMware ESX 2.5.1 Server software can be installed on the PE 1855. For more detailed step by step instructions refer to *ESX 2.5 Server Installation Guide* at http://www.vmware.com/pdf/esx/esx25_install.pdf and *Deployment Guide - VMware ESX Server 2.5.1 Software for Dell PowerEdge Servers* in the Resource Center at www.dell.com/vmware.

3.1 Using CD-ROM

The simplest way to install ESX Server software is to use a CD-ROM drive. A USB CD-ROM drive is required for installing ESX Server software on a PE 1855. Even though installing via a CD-ROM is straightforward, it requires manual intervention and can be time consuming for installation on multiple blade servers. Typically, the first installation is done through a CD-ROM and the subsequent installations use a scripted installation process as described in the next section.

3.2 Using Scripted Installation

VMware installation can be scripted. This unattended installation is useful for deploying ESX Server software on multiple blade servers. The following are the steps for installing ESX using Scripted Installation.

- Configure one of the ESX Servers to support Scripted Installation. This server will be the deployment server housing the contents of ESX installation CD-ROM.
- Using the ESX Management User Interface (MUI), create a boot floppy (or a kick-start file)
- Boot the target server using the boot floppy (or kick-start file deployed through PXE)

For a step-by-step instruction on how to install ESX Server software using scripted installation, refer to *Remote and Scripted Installation* section in *ESX 2.5 Server Installation Guide*.

Note that if a static IP address is configured into the boot-floppy image then a unique boot floppy must be created for each blade server. If the boot-floppy is configured for DHCP address, then the same image can be used on multiple blade servers. It is highly recommended to use a static IP address for the ESX Server software.

4. Post Installation

Once the ESX Server software is installed using the above methods, the following post installation steps must be executed.

4.1 Configuring the NICs

Table 1 lists four configurations for NICs when installing ESX Server software. This section provides instructions for enabling the four different configurations.

4.1.1 Configuration 1:

In this configuration, NIC 0 is dedicated to the service console and NIC 1 is dedicated to the VMkernel. NIC 1 is used both by VMs and VMotion.

This is the default configuration and no special steps are required to enable configuration 1.

4.1.2 Configuration 2:

In this configuration, NIC 0 is shared between the service console and the VMkernel. NIC 1 is dedicated to the VMkernel. NIC 0 is used by VMotion, in addition to the service console. NIC 1 is used by VMs.

The steps to enable this configuration are:

1. Log in with root level privileges to the service console and execute the command `vmkpcidivv -i`. This is an interactive command to allocate devices between the service console and virtual machines. Configure NIC 0 to be shared between service console and the Virtual Machines.
2. Reboot the server
3. Create two virtual switches: one for Virtual Machines (NIC 0), and one for VMotion (NIC 1), using the MUI.

❑ To perform VMotion, the switch names should be identical across all blade servers.

4.1.3 Configuration 3:

In this configuration, NIC 0 is shared between the service console and the VMkernel. NIC 1 is dedicated to the VMkernel. NIC 0 is used by VMs, in addition to the service console. NIC 1 is used only for VMotion.

The steps to enable this configuration are similar to the steps for configuration 2 and are as follows:

1. Log in with root level privileges to the service console and execute the command `vmkpcidivv -i`. This is an interactive command to allocate devices between the service console and virtual machines. Configure NIC 0 to be shared between the service console and Virtual Machines.
2. Reboot the server
3. Create two virtual switches: one for Virtual Machines (NIC 1), and one for VMotion (NIC 0), using the MUI.

■ To perform VMotion, the switch names should be identical across all blade servers.

4.1.4 Configuration 4:

In this configuration, NIC 0 is shared between the service console and the VMkernel. NIC 1 is dedicated to VMkernel. NIC 0 and NIC 1 are teamed by the VMkernel to provide a bond. This bonding provides for NIC redundancy and load balancing.

The steps to enable this configuration are:

4. Log in with root level privileges to the service console and execute the command `vmkpcidivv -i`. This is an interactive command to allocate devices between the service console and virtual machines. Configure NIC 0 to be shared between service console and Virtual Machines.
5. Create a NIC team by adding the following lines to the end of the file `/etc/vmware/hwconfig`:

```
nicteam.vmnick0.team = "bond0"  
nicteam.vmnick1.team = "bond0"
```

6. One can make NIC 1 as the primary NIC for the bond, so that NIC 1 will be used for the VMs and VMotion, while NIC 0 will be used by the service console. This is optional. To make `vmnick0` (which is the typical name given by ESX Server software for the NIC 1 described in this document) as the primary NIC, add the following line to `/etc/vmware/hwconfig`:

```
nicteam.bond0.home_link = "vmnick0"
```

■ Use `vmkpcidivv -i` to find out the exact name of the NIC, which needs to be made as primary NIC.

7. Reboot the server.

4.2 Setting up VLANs

In all four NIC configurations, it is highly recommended to have the VMotion NIC in a separate VLAN. Having the VMotion NIC in a separate VLAN helps ensure that the VMotion traffic is secure. To set up VLANs in the Virtual Switches of the ESX Server software, refer to *ESX 2.5 Server Administration Guide*.

To set up VLANs in a PowerConnect 5316M switch module, refer to the *Dell PowerConnect 5316M System – User Guide* at <http://support.dell.com>.

4.3 Setting up fault tolerance when using a PowerConnect 5316M switch module

To enable NIC failover when using a PowerConnect 5316M switch in [configuration 4](#), click on the advanced configuration in the Options tab of the ESX Server Management User Interface (MUI) and change the following parameter value to 1:

```
Net.ZeroSpeedLinkDown = 1
```

5. References:

1. *Deployment Guide - VMware ESX Server 2.5.1 Software for Dell PowerEdge Servers* at www.dell.com/vmware.
2. *ESX 2.5 Server Installation Guide* at www.vmware.com.
3. *ESX 2.5 Server Administration Guide* at www.vmware.com.
4. *Dell PowerConnect 5316M System – User Guide* at <http://support.dell.com>.

¹This term does not connote an actual operating speed of 1 Gb/sec. For high speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

THIS PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND. INFORMATION IN THIS DOCUMENT IS SUBJECT TO CHANGE WITHOUT NOTICE.

Dell, the DELL logo, PowerEdge, and PowerConnect are trademarks of Dell Inc. VMware is a registered trademark of VMware, Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own. © 2005 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden. For more information, contact Dell Inc.