

Deployment of VMware® Infrastructure 3 on Dell™ PowerEdge™ Blade Servers

The purpose of this document is to provide best practices for deploying VMware® Infrastructure 3.x on Dell™ PowerEdge™ Blade Servers. The document describes various possible configurations when installing VMware® Infrastructure 3.x along with the advantages and disadvantages for each configuration.



Contents

1. Introduction.....	3
2. Configuration Options.....	3
2.1 Stand-alone and SAN Configurations.....	3
2.2 Configuring the NICs.....	6
2.3 Boot from SAN.....	7
2.4 Embedded Switch vs. Pass-through.....	7
2.5 Virtual Center with PowerEdge 1855/1955.....	8
3. Installing ESX Server Software.....	8
4. Post Installation.....	9
4.1 Setting up VLANs.....	9
4.2 Fault tolerance when using a an Ethernet switch module.....	9
5. References:.....	10

1. Introduction

This document describes various configurations that are possible when installing VMware® Infrastructure 3 on Dell PowerEdge™ 1855 and 1955 Blade Servers. Advantages and disadvantages for each configuration along with the steps to install are detailed.

2. Configuration Options

2.1 Stand-alone and SAN Configurations

Dell has qualified VMware® Infrastructure 3 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 fibre channel switch.

For stand-alone configurations, in addition to one dual port embedded Gigabit Ethernet controller, an Intel or Broadcom dual port Gigabit Ethernet daughter card can be used to provide more network I/O capacity for virtual machines.

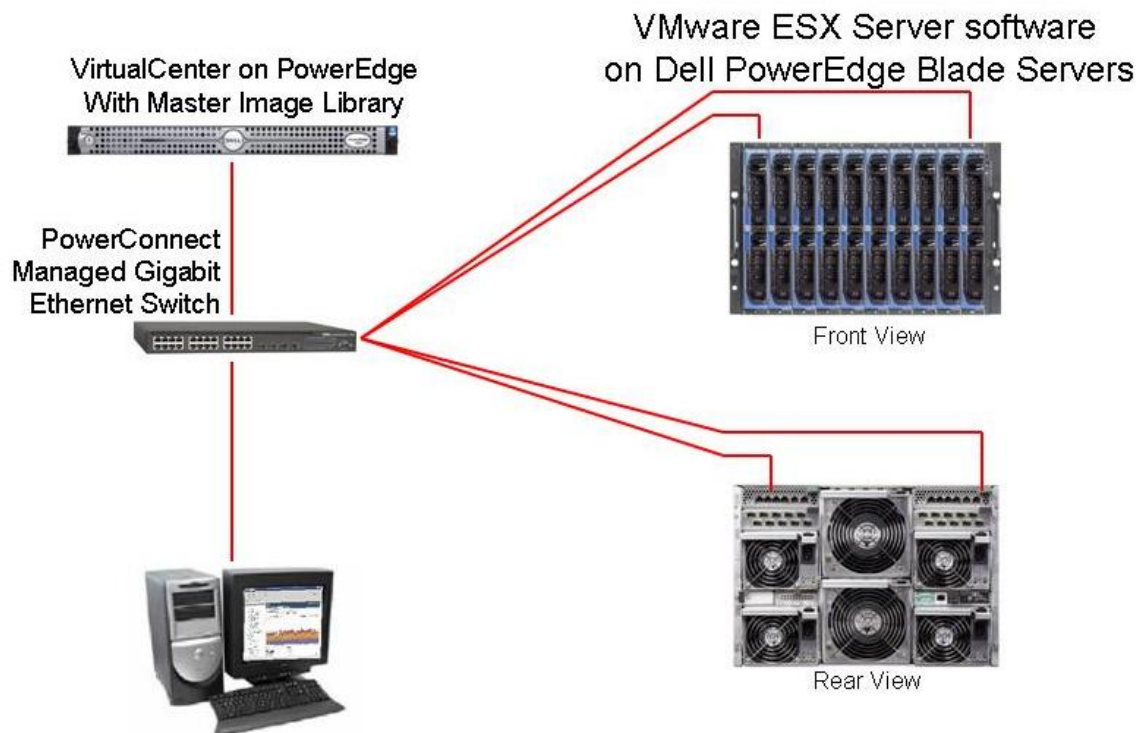


Figure 1: VMware ESX Server Software - Stand-Alone Configuration on PowerEdge Blade Servers

For SAN configurations, a Fibre Channel I/O daughter card and a Fibre Channel switch is required. Direct attached Fibre Channel is not supported. The Fibre Channel switch can

either be an embedded Fibre Channel Switch Module or an external Fibre Channel switch which is connected through a Fibre Channel Pass-Through Module. Table 1 lists Dell supported Ethernet/Fibre Channel daughter cards and switch modules. Table 2 describes differences between the Fibre Channel Switch Module and the Pass-Through Fibre Channel Module.

Table 1: Dell supported peripherals for blade servers

Ethernet Daughter Cards	Intel® Dual Port PCI X Gigabit Ethernet controller Broadcom 5708 Dual Port PCI Express Ethernet controller
Fibre Channel Daughter Cards	Qlogic™ 2342M Qlogic™ QME 2462 Emulex LPe1105
Ethernet Switch Modules	PowerConnect™ 5316M Switch Module Cisco® Catalyst Blade Switch 3030 Module
Fibre Channel Switch Modules	Brocade® SilkWorm® 3014 switch Brocade® SilkWorm® 4016 4Gb switch McDATA® 4314 switch McDATA® 4416 FC switch

Table 2: Fibre Channel switch module vs. Pass-through module

Fibre Channel Switch Module	Pass-Through Fibre Channel Module
<ul style="list-style-type: none"> ▪ Features four uplinks, that are 1/2/4 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 1/2/4 Gb/s uplinks, one for each blade server. ▪ 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 in Figures 2, 3 and 4. Figure 2 shows blade servers connected to Dell/EMC storage via redundant Fibre Channel Pass-Through modules. Pass-Through modules provide one connection for each blade server that is connected to external Fibre Channel switches. Figure 3 illustrates similar configuration but using redundant Fibre Channel Switch Modules available for Dell blade server chassis. Figure 4 shows the configuration when external Fibre Channel switches are used in addition to Fibre Channel switch modules in the chassis.

All these configurations provide protection against failure at a storage processor, switch/pass-through or an HBA level.

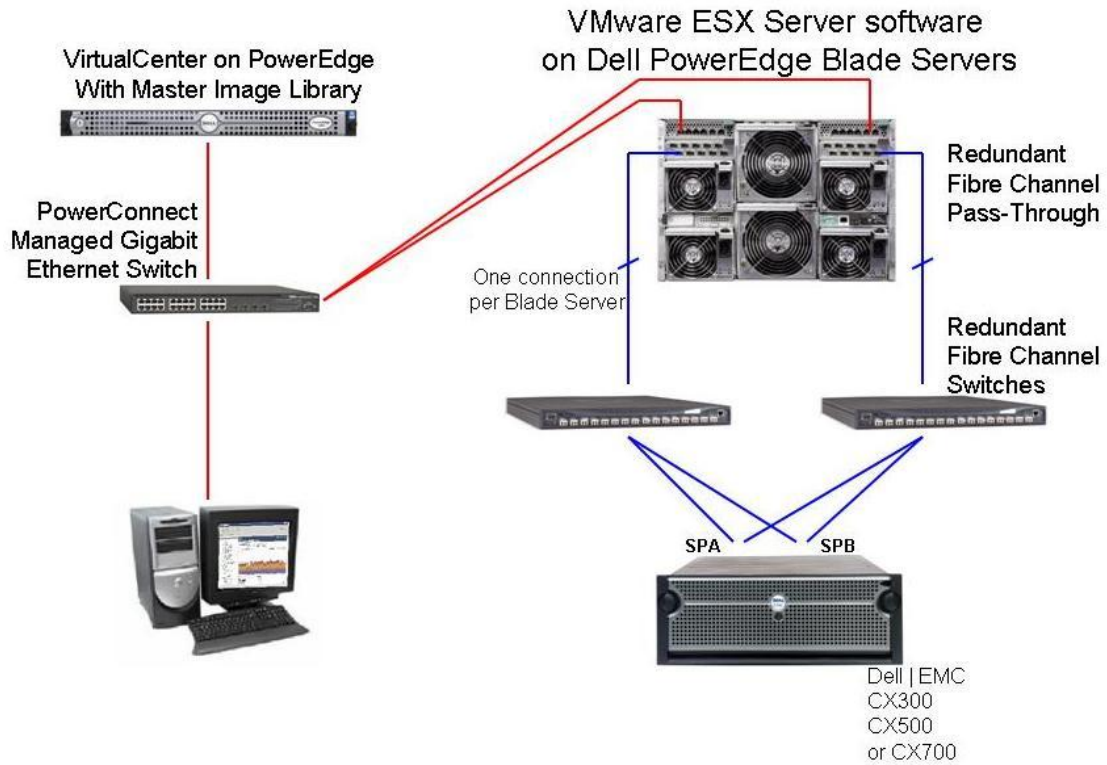


Figure 2: VMware Infrastructure Software - SAN Configuration on PowerEdge Blade Servers with Pass-Through Fibre Channel Modules

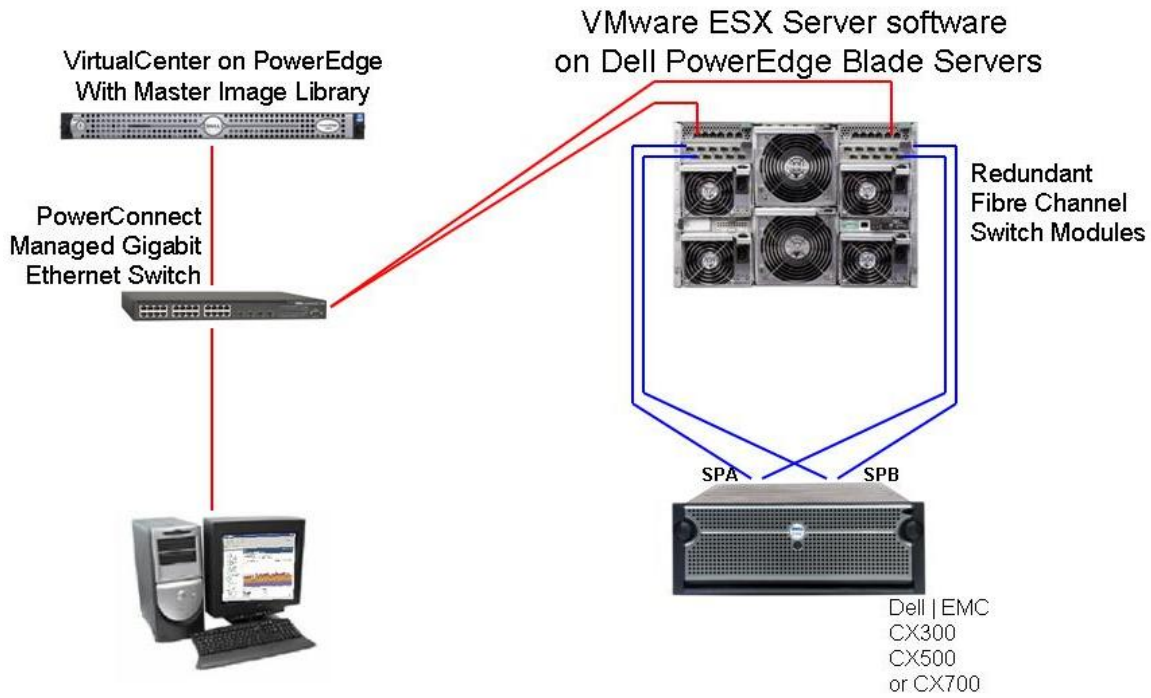


Figure 3: VMware Infrastructure Software - SAN Configuration on PowerEdge Blade Servers with Fibre Channel Switch Modules

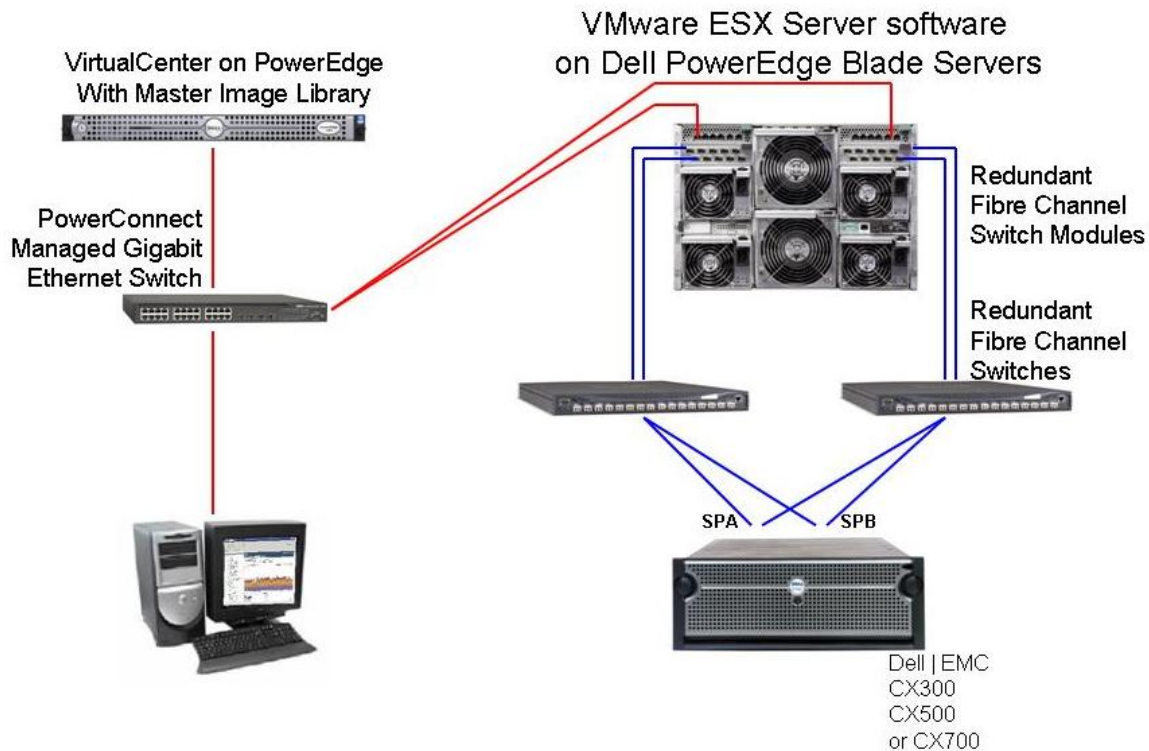


Figure 4: VMware Infrastructure Software - SAN Configuration on PowerEdge Blade Servers with Fibre Channel Switch Modules and external Fibre Channel Switches

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

2.2 Configuring the NICs

In a SAN configuration, each PowerEdge (PE) 1855/1955 blade server configuration contains two embedded Gigabit Ethernet NICs (NIC0 and NIC1) and a dual port Fibre Channel daughter card. Networking can be configured in various ways for traffic isolation or redundancy. In a default installation, Service Console interface is connected to NIC0. A default network for Virtual Machines can be created at installation time that creates a network interface for virtual machines on NIC0. In order to better utilize the network bandwidth and to provide redundancy, the NICs can be configured as shown in Table 2. (this is not required in a stand-alone configuration):

Table 2: NIC Configuration in a SAN environment

Config	NIC 0	NIC 1	Fault Tolerance for VMs	Traffic Isolation?	VM Performance?	Security?
1	Service Console	VMs and VMotion	No	Moderate	Acceptable if VMotion events are infrequent	Moderate - VMotion on production network.
2	Service Console and VMotion	VMs	No	Good	Acceptable	Good -service console on private network.
3	Service Console and VMs	VMotion	No	Moderate	Acceptable if management traffic is infrequent and VMotion events are frequent	Moderate - service console on production network.
4	NIC 0 and NIC1 teamed for Service Console, VM and VMotion traffic		Yes	Moderate and flexible	Acceptable if VMotion events and management traffic are infrequent.	Good - By default VMotion and Service console share the same network. For improved security, VLANs must be used

- ❑ 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.
- ❑ For a NIC team all the team members need to be on the same broadcast domain. The two blade chassis switch modules need to be interconnected either using a direct link or an uplink to an external switch.

2.3 Boot from SAN

ESX Server 3.x Software can be installed and booted from a LUN that resides in a SAN. This technique, known as Boot from SAN (BFS), enables diskless blade servers while consolidating and simplifying storage management. For more information, refer to the *SAN Configuration Guide* at www.vmware.com/support/pubs/vi_pubs.html.

2.4 Embedded Switch vs. Pass-through

Customers can configure a PE 1855/1955 chassis either with an 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. The table below lists the advantages of using either the switch or pass-through configurations.

Table 3: Switch vs. Pass-through module

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

2.5 Virtual Center with PowerEdge 1855/1955

The Dell PowerEdge 1855/1955 chassis can be configured with up to 10 blade servers. One of the blade servers can be installed with VirtualCenter and used as a management console to manage VMs running on other blade servers. For more information on Virtual Center 2.0.1 installation or upgrade, refer to *Installation and Upgrade Guide* at www.vmware.com/support/pubs/vi_pubs.html

3. Installing ESX Server Software

For more detailed step by step instructions refer to *Installation and Upgrade guide* at www.vmware.com/support/pubs/vi_pubs.html and *Deployment Guide - VMware Infrastructure 3.01 for Dell PowerEdge Servers* under *Support Documents* at www.dell.com/vmware.

4. Post Installation

Once the ESX Server software is installed using the above methods, the following post installation steps need to be executed. These steps are required only in a SAN configuration.

4.1 Setting up VLANs

In all four NIC configurations, it is highly recommended to have the VMotion traffic in a separate VLAN. Deploying the VMotion interface to reside in a separate VLAN helps ensure that the VMotion traffic is secure. To set up VLANs in the Virtual Switches of the ESX Server 3.x software, refer to *Server Configuration Guide* at www.vmware.com/support/pubs/vi_pubs.html.

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

4.2 Fault tolerance when using a an Ethernet switch module

This section describes NIC failover when using an Ethernet switch module in Configuration 4, or when using a Gigabit Ethernet daughter card for the stand-alone configuration. The ESX Server can be configured for network redundancy when the embedded Ethernet switch fails or when the external link(s) connected to the embedded switch fails.

❑ Ensure that the Ethernet switch has the latest firmware. For the PowerConnect™ 5316M Switch Module, use firmware version 1.0.0.38 or above.

Network failure detection can be achieved either by using *Link status* or by using *Beacon Probing*. *Link Status* option relies on the connectivity status provided by the network adapter for failover. In the case of PowerConnect Ethernet 5316M switch module, the network adapter will report status as down, when the external switch has failed or when there is no external link (provide the firmware version is 1.0.0.38 or above and the NIC redundancy in the switch is configured appropriately). *Beacon Probing* option sends out and listens for beacon probes on all NICs in the team and uses this information to determine link failure.

For more information on NIC Teaming and Beacon Monitoring refer to the *Networking* section in *Server configuration Guide* at www.vmware.com/support/pubs/vi_pubs.html.

5. References:

1. *Deployment Guide – VMware Infrastructure 3.0.1 for Dell PowerEdge Servers* at www.dell.com/vmware.
2. *Virtual Infrastructure 3 Installation and Upgrade Guide* at www.vmware.com/support/pubs/vi_pubs.html.
3. *Dell PowerConnect 5316M System – User Guide* at <http://support.dell.com>.
4. *VMware ESX Server 3: 802.1Q VLAN Solutions* at <http://www.vmware.com/vmtn/resources/>
5. *VMware Knowledgebase* at <http://kb.vmware.com>

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. © 2006 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.