DELL

Dell™ | VMware ESX/ESXi 3.5 Virtualization Reference Architecture

Dell | Virtualization Solutions Engineering

September 2008

Table of Contents

EXECUTIVE SUMMARY

This *Dell™ | VMware ESX/ESXi 3.5 Virtualization Reference Architecture* white paper describes validated architectures for Dell | VMware™ virtual infrastructures. The Virtualization RA is intended to provide customers with an example basis to consider, evaluate, and select the Dell virtualization solution components that best fit their requirements. This white paper is not intended to be an exhaustive study of specific architectures for every environment. For a complete evaluation of your datacenter's virtualization requirements, please contact your Dell Sales representative.

SECTION 1: INTRODUCTION AND SCOPE

Customers are faced with many choices when planning a virtualization deployment. Server, storage, and software configurations can vary depending on workload and feature requirements. This paper proposes several sample virtualization architectures based on Dell's server and storage products. Architectures are divided into three main categories: Small, Medium, and Large. These categories are based on the features and capacity of each product, as well as the complexity of the overall architecture. The Small architecture is designed with simplicity in mind and focuses on providing basic functionality with minimal hardware. It does not enable advanced virtualization features. The Medium configuration is designed to meet the production requirements of small and medium businesses by providing economical SAN solutions, enhanced virtualization features, and backup options. Finally, the Large configuration is a full-feature virtualization solution using redundancy, one-to-many management, high availability and enterprise-class server and storage products.

The goals of the white paper are as follows:

- Introduce customers to the Dell | VMware sample reference architectures.
- Clarify, aid, and simplify the decision making process for customers.

The scope of this white paper is limited to Dell | VMware virtualization solutions architectures using Dell servers, storage, and networking components. Smaller architectures using free products are possible, but are not described here. Custom configurations for specific customer engagements may have third party hardware or software and hence are out of scope.

This white paper is a starting point in the design of a virtualization solution. Dell provides several resources for the complete design and deployment of virtualization solutions. For further information please visit **www.dell.com/vmware** or **www.dell.com/virtualization** for the latest versions of the virtual solutions.

This white paper is arranged as follows:

- "SECTION 1: INTRODUCTION AND SCOPE" includes the introduction and details the scope of this white paper

- "SECTION 2: VMware ESX/ESXi 3.5 INTRODUCTION" introduces the VMware infrastructure including the ESX Server software and VirtualCenter.
- "SECTION 3: Dell | VMware ESX/ESXi 3.5 REFERENCE ARCHITECTURE" describes the Dell | VMware reference architecture including sample Small, Medium, and Large configurations.
- "SECTION 4: HARDWARE CONSIDERATIONS FOR ESX ENVIRONMENTS" describes general sizing considerations.
- "SECTION 5: ADDITIONAL INFORMATION" provides resources for additional information.

SECTION 2: VMware ESX/ESXi 3.5 INTRODUCTION

The VMware infrastructure consists of the ESX Server software and VirtualCenter. ESX Server enables multiple operating systems to run simultaneously in virtual environments on a single physical server. These virtual environments, referred to as virtual machines, can be created and managed faster than conventional systems. The virtual machines do not interact directly with the physical hardware. Each of these virtual machines runs in a resource-isolated, secure environment and can include Microsoft® Windows®, Linux, Novell® NetWare®, and Sun Solaris operating systems with their associated applications.

VirtualCenter is a management application that monitors and manages virtual machines and ESX Server hosts. With VirtualCenter, users can create, start, stop, and migrate virtual machines across the entire physical data center. Advanced features such as VMware Distributed Resource Scheduling (DRS) and VMware High Availability (HA) are available with VirtualCenter 2.x.

SECTION 3: Dell | VMware ESX/ESXi 3.5 REFERENCE ARCHITECTURE

The Dell | VMware reference architecture represents a possible sample solution that provides a customer with options to consider when meeting their datacenter requirements. In Dell's reference architecture, Dell server and storage are the foundation on which VMware ESX infrastructure is built. Specific ESX features should be chosen based on the datacenter needs. Once chosen, these key ESX features drive hardware requirements. Therefore, the features are as important to the design and decision process as the traditional Quality of Service (QoS) deliverables.

Three sample architectures, Small, Medium and Large, are listed in this section. These sample configurations are generic deployments based on solution size and capability and are independent of the target application. These recommendations focus on hardware capacity and software features; they do not attempt to prescribe a particular server model or chipset, as these preferences may vary based on customer requirements.

Table 1 describes sample components for each configuration. VMware ESX features are listed along with the recommended Dell server and storage. These are examples of baseline architectures that can be bundled together into a datacenter to solve specific customer problems.

Table 1: Hardware and Software Features of the Baseline Configuration

	Small	Medium	Large
Description	Simple entry-level configuration	Designed for small & medium business	Large enterprise environments
Server Configuration			
Servers	PE 1950 III PE 2950 III PE 2970 PE 2900 III	PE M600 PE M605 PE M805 PE M905 PE 2950 III PE 2970 PE 2900 III PE R805	PE M600 PE M605 PE M805 PE M905 PE 2950 III PE 2970 PE 2900 III PE R805 PE R900 PE R905
CPU	2 Sockets Dual/Quad Core	2 Sockets Quad-Core	2 Sockets Quad-Core or 4 Sockets Dual/Quad-Core
Minimum Recommended Network Interface Controllers (NICs)	Two LAN on Motherboard+ One dual port add-in NIC	Two to Four LAN on Motherboard + Two dual port add-in NIC controllers	Two to Four LAN on Motherboard + Two dual port add-in NIC controllers
DRAC	Optional	Yes	Yes

Local Storage	5x 73GB/146GB (2x 73GB for PE1950 III) if hosting virtual machines (VMs) on local storage or 2x 73GB if hosting VMs on a Serial Attached SCSI (SAS) disk expansion enclosure	2x 73GB	2x 73GB
Controller **Redundant Array of Independent Disks (RAID)**	PERC 6/i, SAS 6/iR RAID 1 and 5	PERC 6/i, SAS 6/iR RAID 1 and 5	PERC 6/i, SAS 6/iR RAID 1 and 5
Storage Configuration			
Storage Fabric	Local or Disk Expansion Enclosure (SAS) storage	Network-based Internet Small Computer Systems Interface (iSCSI)	Fiber Channel or iSCSI Storage
Storage Array	Local Storage, MD1000, MD1120 and MD3000	MD3000i, AX4-5I, PS5000E, PS5000X, PS 5000XV, PS5500	PS5000E, PS5000X, PS 5000XV, PS5500, CX4-120, CX4-240, CX4-480, CX4-960
Backup			
Backup Server	N/A	PE2900	PE2900
Backup Software	N/A	Backup Exec 11D™	Backup Exec 11D
Backup Device	N/A	PowerVault™ TL4000	PowerVault ML6000
Software Configuration			

VMware Infrastructure Edition	Foundation	Enterprise	Enterprise
VMware VMotion™	N/A	Yes	Yes
VMware Distributed Resource Scheduling (DRS)	N/A	Yes	Yes
VMware High Availability (HA)	N/A	Yes	Yes
VMware Storage VMotion	N/A	Yes	Yes
Management			
VMware VirtualCenter (VC)	Yes	Yes	Yes
ITA	N/A	Yes	Yes

Small Configuration

The Small configuration is an example of a simple entry-level configuration. Customers can get a feel for how their applications perform under virtualization. The Small configuration is shown in Figure 1 with one server running ESX Server Foundation Edition. Eight GB is the minimum memory recommended in this configuration. The storage requirements are expected to be low and local storage may be suitable in some instances. Direct attached disk expansion enclosures can be attached if larger storage capacity is required. The available options for direct attached storage are the PowerVault MD1000, PowerVault, MD1120, and PowerVault MD3000.

Figure 1 illustrates an example Small configuration that demonstrates feature simplicity and easy deployment, but does not reflect all possible Small configurations. This example provides a high level view of a Small configuration; but it also highlights that some key ESX features are not configurable within the limits of a Small configuration. Implementing advanced features of the Dell | VMware reference architecture requires a Medium or Large configuration. In the example Small configuration, ESX Server is deployed on a Dell PowerEdge™ 2900 III that is attached to a Dell PowerVault™ MD1120, which offers disk capacity and redundancy.

Figure 1: Sample Small Configuration Demonstrating Feature Simplicity and Easy Deployment

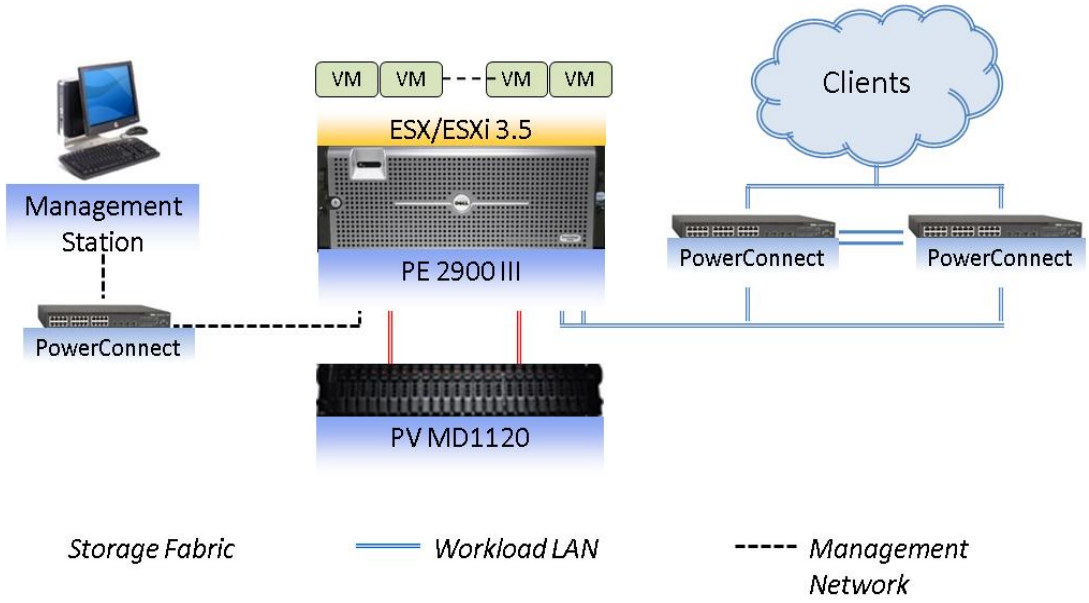

Medium Configuration

The Medium configuration enables features that are designed to meet the needs of small and medium business customers. The Medium configuration is a representation of an ESX feature oriented infrastructure. Host servers run ESX/ESXi Enterprise Server and are managed using the VirtualCenter Server. Each server is attached to shared storage, which enables additional ESX features such as VMotion, High Availability, Distributed Resource Scheduling, and Storage VMotion. The Medium configuration offers uptime, maintenance scheduling, resource scheduling, one-to-many manageability options and backup and recovery options.

Figure 2 illustrates an example Medium configuration that demonstrates these additional ESX features but does not reflect all possible Medium configurations. Servers are running on Dell PowerEdge R805. They are connected using an Ethernet fabric and share storage on the PowerVault MD3000i. Virtual Center Server runs on a Dell PowerEdge 1950 III to configure and manage the MD3000i and to manage the ESX Servers.

Figure 2: Sample Medium Configuration Demonstrating Additional ESX Features Such As High Availability, Distributed Resource Scheduling, VMotion, and Storage VMotion

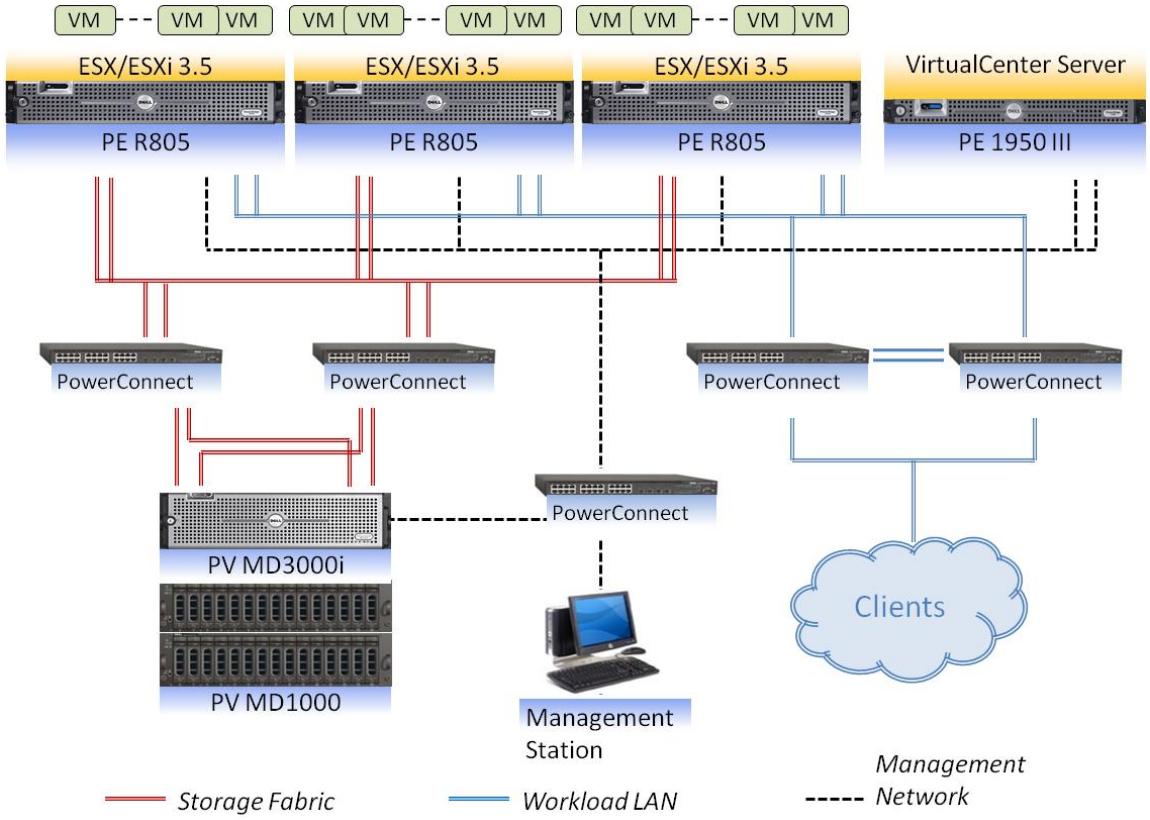

Large Configuration

The Large configuration is a fully featured virtualization solution targeted towards large enterprise needs. Each server is running ESX Enterprise Server. The key features enabled are high availability, distributed resource scheduling, VMotion, Storage VMotion, manageability, and operability. The focus is on performance, scalability, space, capacity, and power efficiency.

Figure 3 illustrates an example Large configuration that demonstrates a fully-featured virtualization solution but does not reflect all possible Large configurations. In this example, the ESX servers are running on Dell PowerEdge M600 blade servers, which focus on optimizing space, performance, and power consumption. The storage backend is the Dell EqualLogic™ PS5000 to provide host expandability, performance, and storage manageability. VirtualCenter Server runs on a Dell PowerEdge 2950 III attached to a PowerVault MD1120, which offers RAID storage and redundancy to the VC Server.

Figure 3: Sample Large Configuration Demonstrating a Fully Featured Virtualization Solution

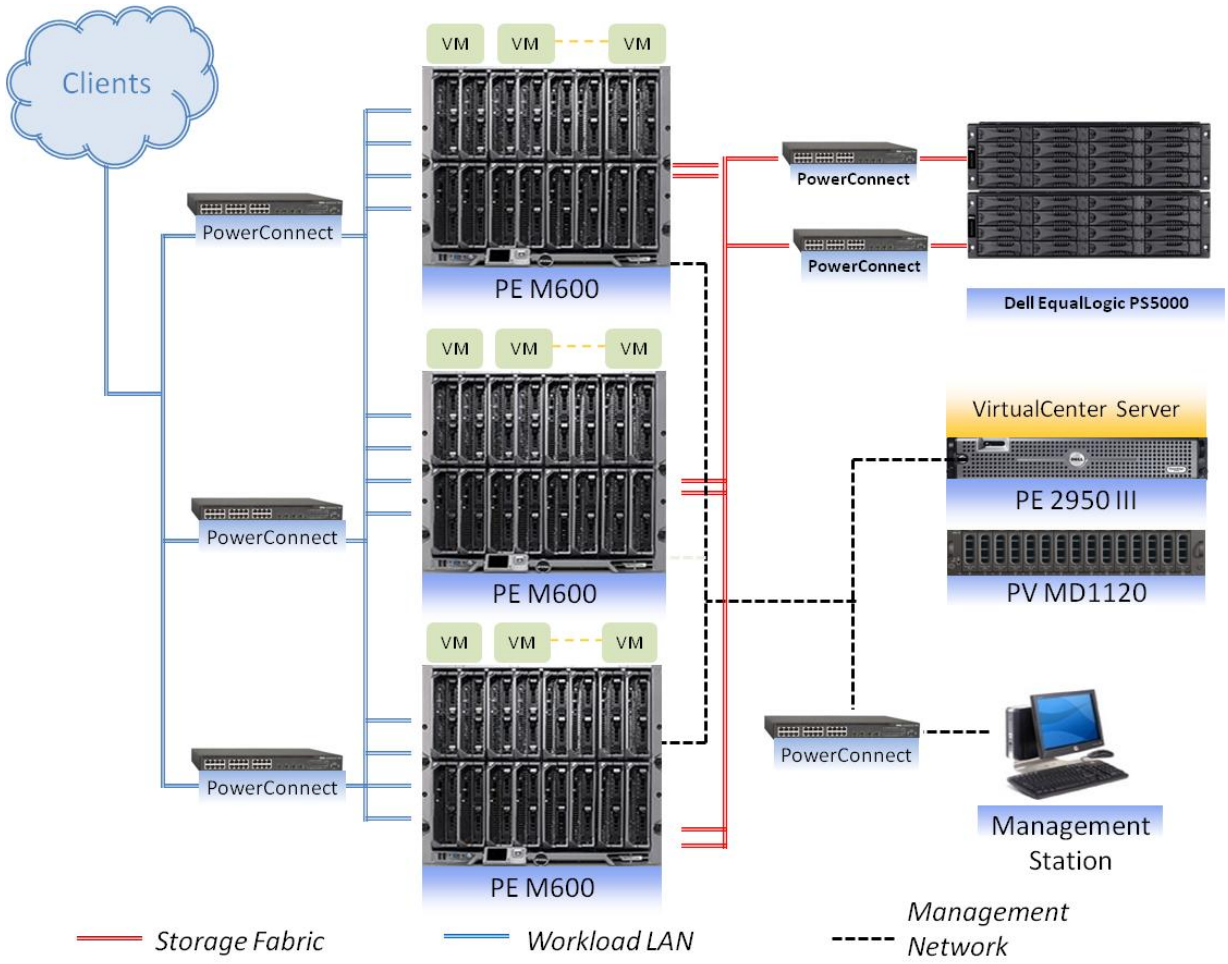

Disaster Recovery for Large Configurations

Disaster recovery for large configurations can be achieved through VMware Site Recovery Manager (SRM). SRM makes disaster recovery rapid, reliable and manageable so that you can meet recovery objectives. It eliminates complex manual recovery steps and enables a non-disruptive and automated way to test the recovery plan. Dell EqualLogic and Dell EMC® storage (with replication software) supports SRM. For more information on Site Recovery Manager, visit: **http://www.vmware.com/products/srm/** .

SECTION 4: HARDWARE CONSIDERATIONS FOR ESX ENVIRONMENTS

This section provides general scaling considerations for hardware when designing a virtualized infrastructure.

This discussion focuses on the four major hardware subsystems:

- CPU
- Memory
- Disk
- Network

It is important to understand these four subsystems because failure to plan for them properly can lead to poor quality of service (QoS) in the datacenter.

CPU Subsystem

The CPU subsystem is a key component in delivering acceptable QoS in the datacenter. The key decision points are choosing the right number of processors and processor types. In addition, choosing between dual-core processors and quad-core processors involves many factors such as cost and performance.

Ultimately, the server needs enough physical CPU resources to meet the needs of the virtualized infrastructure for the near term and future growth. In short, improved CPU performance can be accomplished by choosing servers with more processor capacity, picking processors with higher frequency bins or upgrading to quad-core processors from dual-core processors.

For ESX/ESXi, the virtual machine CPU requirements are distributed across all physical cores. The hypervisor consumes some CPU resources to manage the virtual machines and their requests. Each VM can support up to 4 virtual CPUs. It is always good to start with 1 virtual CPU and increase the number of virtual CPUs if needed.

Memory Subsystem

The memory subsystem is another important component in delivering QoS in the datacenter. The main decision is choosing the right amount of memory to meet the needs of the datacenter and balance it against the cost of memory and the cost of adding additional systems. Next, the memory need is driven by the number of virtual machines, the amount of memory needed by each virtual machine, the amount of memory needed by the hypervisor to manage those VMs, and the memory cushion for migration considerations. Since memory tends to be the first resource to be completely utilized, careful planning is required to ensure enough capacity to sustain and grow the datacenter.

In ESX/ESXi, the aggregate memory requirement for all VMs should be taken into consideration. Even with memory sharing and some of the memory optimizations that ESX/ESXi makes available, it is much safer and more stable to give a VM the full allocation of its configured memory. The memory consumed by the hypervisor varies based on the number of VMs and the memory allocated to each VM. The

memory usage starts at 384MB and adds an additional 32MB per VM. In a heavily loaded scenario, up to 1GB RAM could be used by the hypervisor.

Disk Subsystem

The disk subsystem needs to provide enough storage capacity while delivering QoS to the datacenter with acceptable I/O request response times. In addition, the number of hosts connecting to the disk subsystem and the fabric type need to be factored in order to pick the correct Dell storage enclosure. The capacity portion is a simple calculation based on RAID level choice and capacity needs for virtual machines and applications. An additional design criterion involves performance sizing the disk subsystem and providing enough disk spindles and bandwidth to handle the I/O need.

Network Subsystem

The network subsystem needs to deliver QoS on I/O requests, connect the datacenter, and enable key hypervisor features. The decision points involve choosing enough NIC ports to handle the data traffic, management traffic, and hypervisor-specific traffic and building a solid fabric infrastructure to meet current needs and allow for scalability as the datacenter grows. Bandwidth and performance can be increased by adding hardware to the infrastructure. The general recommendation is to have dedicated NIC for management (two if VMware HA is required), a dedicated NIC for VMotion, two or more dedicated NIC for redundant iSCSI traffic and two or more redundant NIC for virtual machines.

SECTION 5: ADDITIONAL INFORMATION

Please check **www.dell.com/vmware** or **www.dell.com/virtualization** for the latest updated versions of Dell whitepapers and solutions guides. From **dell.com/vmware**, click on **Support Documents** under the **Resources** tab and select the relevant support document. Information on the VMware Site Recovery Manager is located at **http://www.vmware.com/products/srm/** .

For additional Dell | VMware resources, please visit **http://www.delltechcenter.com/page/VMware .**

All Dell hardware-specific manuals, white papers, and other technical documentation can be found on **www.support.dell.com**. Once here, click **Technical Support** in the **Browse by Category** link box, and select your hardware platform.