

DELL

Dell™ | Citrix XenServer

Dell Edition

Reference Architecture

VIRTUALIZATION SOLUTIONS ENGINEERING

September 2008



THIS WHITE 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.

Dell, the *DELL* logo, *EqualLogic*, *OpenManage*, *PowerVault*, and *PowerEdge* are trademarks of Dell Inc.; *Xen*, *XenServer*, *XenCenter* and *XenMotion* are either registered trademarks or trademarks of Citrix Inc.; *Intel* and *Xeon* are a registered trademarks of Intel Corporation in the U.S. and other countries; *AMD* is a registered trademark and *Opteron* is a trademark of Advanced Micro Devices Inc.; *Microsoft*, *Windows*, *Windows Vista*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

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 disclaims proprietary interest in the marks and names of others.

©Copyright 2008 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Information in this document is subject to change without notice.

Table of Contents

- EXECUTIVE SUMMARY 5
- INTRODUCTION AND SCOPE 5
- XenServer Dell Edition Introduction 6
 - Citrix XenServer Dell Edition Features 6
 - Factory installed from Dell 6
 - Integrated systems management and monitoring 6
 - XenServer Local Console 6
 - Optimized footprint and controlled environment 6
 - Improved reliability 6
 - Improved XenServer updates 6
 - Pre-certified and supported configurations 6
- Dell | Citrix XenServer Reference Architecture 7
 - Small Configuration 9
 - Medium Configuration 10
- Hardware Considerations for XenServer Environments 11
- References 13

EXECUTIVE SUMMARY

This Dell™ Virtualization Reference Architecture (RA) white paper describes validated architectures for Dell|Citrix XenServer™ Dell Edition 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 document 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.

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, 1-to-many management, high availability and enterprise-class server and storage products. The goals of the document are as follows:

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

The scope of this document is limited to Dell | Citrix XenServer Dell Edition virtualization solutions only. It defines 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 document 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/citrix or www.dell.com/virtualization for the latest updated versions.

The document is arranged as follows: Section 2 includes the introduction and details the scope of the document; Section 3 describes Dell | Citrix XenServer Dell Edition feature, Section 4 describes two reference configurations , Section 4 describes general sizing considerations and Section 5 contains the reference links.

XenServer Dell Edition Introduction

With the 64-bit open-source Xen hypervisor at its core, Citrix® XenServer Dell Edition™ is a powerful virtualization solution that enables efficient resource consolidation, utilization, dynamic provisioning, and integrated systems management. XenServer Dell Edition has a small foot print and is optimized to run from an internal flash storage in Dell PowerEdge™ servers. Dell and Citrix have partnered to bring pre-qualified and virtualization-ready platforms for today's dynamic and growing data centers.

Citrix XenServer Dell Edition Features

Factory installed from Dell

Citrix XenServer Dell Edition is factory installed from Dell on select PowerEdge servers. This reduces the installation and deployment time required and simplifies the deployment process to get your infrastructure ready to use. With minimal configuration, the XenServer host is available to run virtual machines.

Integrated systems management and monitoring

Citrix XenServer Dell Edition comes pre-installed with Dell OpenManage™ Server Administrator. This enables systems management right out of the box without any additional need to install an agent on the host.

XenServer Local Console

The XenServer Dell Edition includes a new XenServer Local Console user interface to enable local administration of the host. XenServer Local Console enables users to configure and view host specific properties such as management network configuration, local storage for virtual machines, etc. XenCenter, the standard Microsoft® Windows® management console for the XenServer, is also available.

Optimized footprint and controlled environment

XenServer Dell Edition is optimized for a smaller disk footprint and writes to flash storage. The majority of the XenServer Dell edition file system is Read Only and thus provides a tighter control over the XenServer operating environment. The XenServer Dell Edition host agent software has been significantly optimized to minimize the number of write cycles to the flash storage. XenServer writes to flash storage only when something important has changed and must be recorded. Write minimization helps improve the life of the flash storage.

Improved reliability

Running XenServer Dell Edition on a server's internal flash storage provides improved reliability over running on traditional hard disks. Flash storage does not have moving parts and is more reliable than a hard disk.

Improved XenServer updates

To improve reliability of software upgrades, the XenServer Dell Edition image contains a primary and secondary copy of the XenServer file system. Any time an update is applied, the secondary copy gets updated leaving the primary copy in a known good state. The secondary copy becomes the primary after a successful upgrade. An update to the XenServer host can be applied using the XenServer local console or XenCenter.

Pre-certified and supported configurations

Citrix XenServer Dell Edition is certified and fully supported by Dell for select server and storage configurations. Both Dell and Citrix have worked closely to provide the highest quality product and user experience

Dell | Citrix XenServer Reference Architecture

The Dell | Citrix XenServer 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 Citrix XenServer infrastructure is built. Specific XenServer features should be chosen based on the datacenter needs. Once chosen, these key XenServer features will drive hardware requirements. Therefore, the features are as important to the design and decision process as the traditional QoS deliverables.

Two sample configurations: **Small** and **Medium** are listed here. These configurations are generic deployments based on solution size and capability and are independent of the target application. The 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.

The following table describes sample components for each configuration. Citrix XenServer features are listed below and enabled with recommended Dell server and storage. These are examples of baseline architectures that can be bundled together into a datacenter to solve customer problems.

Table 1: Hardware and Software features of the Baseline Configuration

	Small	Medium
Description	Simple entry-level configuration	Designed for small & medium business
Server Configuration		
Servers	PE2900 III, PE1950 III PE R805	PE M600 PE M605 PE M805 PE M905 PE 2950 III PE R805 PE R900 PE R905
CPU	2 Sockets / Dual-Core	2 Sockets / Quad-Core

Minimum NIC ports	Two (two LAN on Motherboard ports)	Four to six (Two to four LAN on Motherboard ports plus additional add-in if necessary)
DRAC	Optional	Yes
Local Storage	5x 73GB / 5x146GB	2x 73GB
Internal storage controller	PERC 6/i or SAS6/iR (RAID 1,5)	PERC 6/i or SAS6/iR (RAID 1,5)
Storage Configuration		
Storage Fabric	Local or external SAS/SATA storage	Network-based iSCSI
Storage Array	Local Storage, MD1000, MD1120 MD3000	MD3000i, PS5000E, PS5000X, PS 5000XV
Backup Server	N/A	PE2900
Backup Software	N/A	Backup Exec 11D
Backup Device	N/A	PowerVault TL4000
Software Configuration		
	Citrix XenServer Dell Express or Enterprise Edition	Citrix XenServer Dell Enterprise Edition
XenMotion live migration	Yes*	Yes

High Availability	Yes*	Yes
Resource QoS Controls	Yes*	Yes
Multi-server management	Yes*	Yes
Resource pools	Yes*	Yes
Shared IP-based storage	Yes*	Yes
VLAN confirmation	Yes*	Yes
Dell EqualLogic Storage Adapter	Yes	Yes
Management		
XenCenter	Yes	Yes
ITA	Optional	Yes
OpenManage Administrator	Yes	Yes

* Feature is available only in Citrix XenServer Dell Edition

Small Configuration

The Small configuration is designed to be a simple entry-level configuration. Each server is running Citrix XenServer Dell Express or Enterprise Edition. The storage requirements are expected to be low so local storage may be suitable for some deployments. But direct attach options are also available with the MD1000, MD1120, or MD3000.

Figure 1 below is a sample instance of a small configuration and does not reflect all possible small configurations. It provides a high level view of a small configuration; but it also highlights that some key XenServer features are not configurable within the construct of a small configuration. Therefore, there is a need to abstract higher order configurations such as the Medium and Large configurations. In this example, Citrix XenServer Dell Edition is deployed on two PowerEdge 2900 III servers that are attached to a PowerVault MD3000 array. Using MD3000 array, the two XenServer hosts can share a storage LUN, thus enabling live migration or XenMotion of virtual machine across two hosts. PowerVault MD3000 provides an ideal environment for small use where host scalability is not a big concern. PowerVault MD3000 provides up to 15 TB of raw disk capacity using 15 SAS or SATA drives. Up to two PowerVault

MD1000 disk enclosures, each enclosure providing a maximum of 15 TB of raw storage can be daisy chained to PowerVault MD3000 to increase raw storage capacity to a total of 45 TB. Using SAS HBAs on XenServer hosts and dual controllers on MD3000 provides redundancy against an HBA, controller or SAS interconnect failure.

As a best practice it is recommended to segregate XenServer host management traffic from other management traffic such as storage management. Note that in XenServer, the management network interface is used for XenMotion or live migration traffic. Traffic segregation can be achieved by using separate physical switches or using separate VLANs. As shown in the figure, high availability and load balancing for virtual machine traffic can be achieved by bonding two NICs on each XenServer host.




Figure 1: A sample small configuration for Citrix XenServer Dell Edition

Medium Configuration

The Medium configuration is targeted towards small and medium business customers to enable both host and storage scalability that is more than what is offered by the small configuration. Each PowerEdge server running Citrix XenServer Dell Enterprise Edition enables customers to use advanced features such as resource pools, integrated HA, XenMotion, and resource QoS.

Figure 2 below is a sample instance of a medium configuration and does not reflect all possible medium configurations. Citrix XenServer Dell Edition runs on Dell PowerEdge R805 servers, configured in a resource pool. All XenServer hosts are connected via an Ethernet fabric and share storage on the Dell EqualLogic PS5000 array(s). Two NICs on each host are configured as a bond to provide high availability for iSCSI storage traffic. The iSCSI traffic is isolated from other network traffic using separate physical Ethernet switches. Two NICs on each host are bonded to provide high availability and load balancing for virtual machine traffic. Two NICs on each host are bonded to provide high availability and load balancing for the host management and XenMotion traffic. This configuration illustrates traffic segregation using

separate physical switches; however isolation can also be achieved using VLANs. XenCenter can be installed on any workstation or server machine that has access to the host management network.




Figure 2: A sample medium configuration for Citrix XenServer Dell Edition

Up to 16 hosts can be configured as part of a single resource pool. Being part of a resource pool allows XenServer hosts to share iSCSI storage on PS Series arrays. Shared storage is a requirement for live migration or XenMotion of virtual machine, providing uninterrupted VM migration from one physical host to another. Using shared storage also allows users to configure the XenServer pool for High Availability, a new feature available in XenServer 5.0. Depending on the pool configuration and restart priority of VMs, in case of a host failure, XenServer HA feature restarts virtual machines on remaining healthy hosts in the pool.

XenServer 5.0 also introduces a direct storage adapter for Dell EqualLogic PS Series arrays. Using direct storage adapter for PS Series, storage management tasks for virtual machines can be achieved directly using XenCenter interface. The storage adapter also enables advanced PS series features such as cloning to instantaneously deploy new virtual machines, virtual machines snapshots, thin provisioning, etc.

Hardware Considerations for XenServer Environments

This section provides general scaling considerations for hardware when designing a virtualized infrastructure. The guidance focuses around the four major hardware subsystems: (1) CPU; (2) memory; (3) disk; and (4) network. These four subsystems are important to understand in the general sense because not planning for them can lead to poor quality of service (QoS) in the datacenter.

The CPU subsystem is a key component in delivering acceptable QoS in the datacenter. The decision hinges on choosing the right number of processors and processor type. In addition, choosing between

dual-core processors and quad-core processors involve many factors such as cost and performance. Ultimately, the server needs enough physical CPU resources to meet the needs of the virtualized infrastructure and the near term growth. In short, improving 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.

When sizing CPU requirements for XenServer, it is fair to assume that one processor core will be used by the XenServer platform while the remaining available for use by the virtual machines. The control domain is allocated one CPU and hence uses one of the available cores in the host. XenServer Dell Edition supports up to 8 virtual CPUs per VM, however the guest OS may support a different number. It is recommended to configure a virtual machine is one virtual CPU and add more virtual CPUs if necessary.

The memory subsystem is another important component in delivering QoS in the datacenter. The main decision is to choose 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 will ensure enough capacity to sustain and grow the datacenter.

In XenServer, out of the total physical memory available on the host, a VM is allocated the full chunk of configured memory. Hence when planned memory requirements for XenServer host, aggregate memory requirements for all VMs should be taken into consideration. Memory consumed by XenServer platform (Xen hypervisor and control domain) varies with the total physical memory available in the server. The Xen hypervisor uses 128MB of memory regardless of the physical memory on the server. The memory used by the control domain is always at least 200MB, and is never more than 752MB; within that range it is scaled as a linear function of total host RAM. For hosts with up to 3.5GB of physical RAM, the control domain usage remains at 200MB; on a 5GB host the control domain will use 228MB; on a 16GB host the control domain consumes 454MB; and on hosts with 32GB or more the control domain consumes 752MB. Therefore, the total memory used by XenServer platform ranges anywhere between 328MB to 880MB.

The disk subsystem needs to provide enough storage capacity while delivering QoS to the datacenter through acceptable IO 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. The second design criteria involve performance sizing the disk subsystem and providing enough disk spindles and bandwidth to handle the IO need. To provide better service to I/O hungry virtual machines, XenServer Dell Enterprise Edition allows one to specify virtual disk I/O priority for virtual disk on LVM based storage repositories.

The network subsystem needs to deliver QoS on IO requests, connect the datacenter, and enable key hypervisor features. The decision points involve choosing enough NIC ports to handle the VM data traffic, management traffic, and storage traffic (if any) and building a solid fabric infrastructure to meet current needs and allow for scalability as the datacenter grows.

XenServer Dell Edition supports up to 6 active physical network interfaces (or up to 6 pairs of bonded network interfaces) per XenServer host and up to 7 virtual network interfaces per VM. XenServer 5.0 introduces active/active NIC bonding that enables load balancing and failover for virtual machine traffic and failover for other types of network traffic. For finer grain control over network bandwidth Quality of Service, I/O priority can be set (XenServer Dell Enterprise Edition only) for a virtual machine's network interface to limit the transfer rate to desired kilobytes per second. For IP storage, as a best practice it is recommended to use a highly available network infrastructure: host NIC ports configured as a bond, redundant switches and storage array with redundant network ports/controllers. A dedicated private network for IP storage is recommended to isolate storage traffic from virtual machine and host management traffic.

References

- Citrix XenServer Dell Edition Solution Guide: <http://support.dell.com/support/edocs/software/Citrix/>
- Dell | Citrix Solutions: www.dell.com/citrix
- Dell Virtualization Solutions: www.dell.com/virtualization
- XenServer Dell Edition Product Downloads: www.citrix.com/xenserver/dell
- Dell PowerVault Storage: www.dell.com/powervault
- Dell EqualLogic Storage: www.dell.com/psseries

Trademarks used in this text: *Dell*, the *DELL* logo, *EqualLogic*, *OpenManage*, *PowerVault*, and *PowerEdge* are trademarks of Dell Inc.; *Xen*, *XenServer*, *XenCenter* and *XenMotion* are either registered trademarks or trademarks of Citrix Inc.; *Intel* and *Xeon* are a registered trademarks of Intel Corporation in the U.S. and other countries; *AMD* is a registered trademark and *Opteron* is a trademark of Advanced Micro Devices Inc.; *Microsoft*, *Windows*, *Windows Vista*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

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.