

Reference Architecture for Dell™ VES Self-Service Creator and VMware® vSphere® 4

Solutions for Large Environments

Virtualization Solutions Engineering

Ryan Weldon and Tom Harrington



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.

© 2010 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, the *DELL* badge, *EqualLogic*, *PowerEdge*, and *PowerConnect* are trademarks of Dell Inc. *VMware*, *vMotion*, *vCenter*, and *vSphere* are registered trademarks or trademarks of VMWare, Inc. in the United States or other countries. *Microsoft*, *Active Directory*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *SUSE* is a trademark of Novell Inc. in the United States and other countries. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the U.S. and 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.

Table of Contents

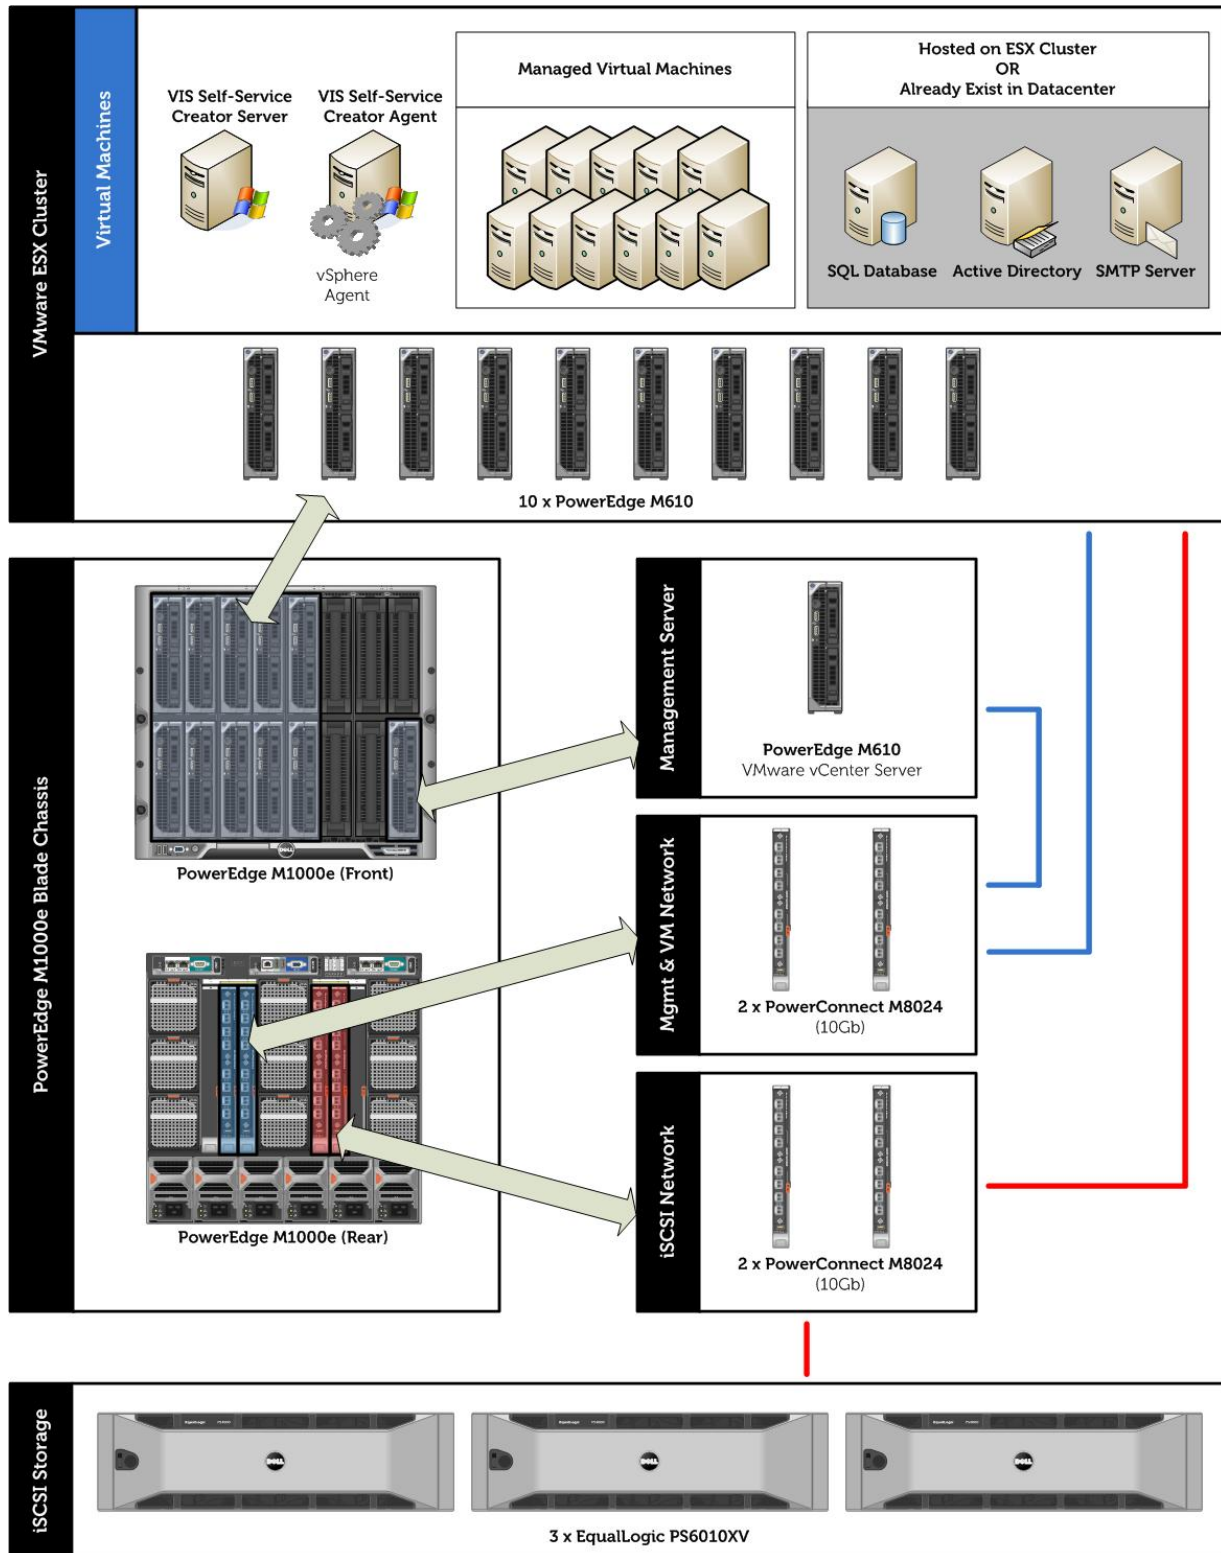
Overview	2
Figure 1: Reference Architecture Overview	3
Design Principles	4
Solution Capabilities	4
VIS Self-Service Creator	4
vSphere.....	5
EqualLogic PS6010XV iSCSI Storage Array	5
PowerConnect M8024 Switches.....	5
Solution Components Details	6
VIS Self-Service Creator	6
vSphere.....	7
Networking	8
iSCSI Storage.....	8
Additional Resources.....	9

Overview

This white paper provides a sample implementation using VIS Self-Service Creator and vSphere 4. The sample implementation provided is designed for large configurations that target a from-the-ground up solution, or are looking to provide private cloud capabilities for an existing environment. This white paper is intended to provide an overview of the requirements to support a VIS Self-Service Creator implementation; however, it is not intended to serve as an implementation guide. This white paper does not discuss mapping your IT and business policies, procedures, and standards into the VIS Self-Service Creator software which is critical for a successful implementation. Dell Global Services consultants provide both planning and implementation support for VIS Self-Service Creator and virtualization environments. For additional details on Dell Global Services offerings, refer to <http://www.dell.com/services>.

Additional details on the underlying hardware and hypervisor configuration utilized to support this solution are available in a separate Dell white paper, "Reference Architecture for VMware vSphere 4 in a 10 Gigabit iSCSI Environment with Dell PowerEdge Blade Servers", available at <http://www.dell.com/virtualization/businessready>. The recommendations for VIS Self-Service Creator provided in this paper are targeted for environments that have the entire managed infrastructure within a single site, and have less than 1500 virtual machines that are required to be managed. If the targeted implementation does not meet these parameters, then modifications may need to be made to support your requirements. Additional reference architectures are available at <http://www.dell.com/vis> that may meet your requirements. Figure 1 on the following page provides an overview of a standard implementation.

Figure 1: Reference Architecture Overview



Design Principles

The following principles were utilized during the design of this reference architecture:

- **Optimal Configuration for Large Environments:** The reference architecture solution components are specifically designed for the requirements of large environments where the managed virtual environment contains up to 1500 virtual machines within a single data center.
- **Optimal Server Configuration for Virtualization:** The recommended architecture is designed with an optimal hardware configuration to support virtualized workloads.
- **Redundancy:** Redundancy is incorporated in every aspect of the solution, including networking and storage to avoid any single-point-of-failure.
- **Isolated, Redundant, and High-Performance Network Architecture:** The proposed 10 Gb Ethernet network architecture is designed to support isolation of the various traffic types required in a virtualized environment. It is designed to have no single point-of-failure, and have optimal performance through NIC teaming and load balancing.
- **Scalability:** Scaling the reference architecture is possible by changing the server model (such as PowerEdge™ M910s), increasing the number of PowerEdge M610 servers and/or EqualLogic™ iSCSI storage arrays; in addition, adding additional memory or increasing processor speeds on the PowerEdge servers will provide support for additional virtual machines on the ESX cluster. Depending on the scalability level required, stand-alone (outside the blade chassis) switches, such as the PowerConnect™ 8024F, may need to be added to the environment. Scaling beyond 1500 virtual machines and multiple sites/data centers is supported by VIS Self-Service Creator; however, that configuration is outside the scope of this document.
- **Integration with an Existing Environment:** VIS Self-Service Creator can manage previously deployed virtual machines on a variety of supported hypervisors/deployment solutions (see the *VIS Self-Service Creator Support Matrix* available at <http://support.dell.com/manuals>). From a networking perspective, the PowerConnect switches may be uplinked into the existing network infrastructure, and it is assumed that a network infrastructure is in place to support the out-of-band hardware management network.

Solution Capabilities

VIS Self-Service Creator

- **Self-Service Provisioning:** End users utilize a Web interface to request that VMs be provisioned and manage previously provisioned VMs. Access controls are in place to support separate views per user and per group.
- **Control VM Sprawl:** Supports administrator-defined lease periods for VMs with automatic reclamation of resources at lease expiration. Administrators may request or force reclamation of VMs prior to the defined expiration date, and end users may return unused VMs prior to the end of the lease.
- **Automate IT Governance Policies:** Supports the automation of numerous IT governance policies including complex organizational approval workflows, VM placement enforcement, and administrator defined archival periods.
- **Integrate with and manage existing virtualized infrastructure:** Supports management of previously deployed virtualization solutions and import of their associated VMs.

- **IT Cost Allocation:** Supports both allocation-based and consumption-based chargeback models, with cost transparency throughout the product. Enables users and administrators to understand the cost of IT services consumed, even if departments aren't actually charged for the resources.

vSphere

- **Advanced virtualization features:** Features such as VM High Availability and vMotion® decrease both planned and unplanned downtime, while providing administrators the flexibility to perform hardware operations without impacting end users. Features such as Distributed Resource Scheduler (DRS) and Distributed Power Management (DPM) ensure optimal resource utilization of ESX cluster servers.
- **Template-based cloning:** Through vCenter® Server, VMs can be converted into templates that VIS Self-Service Creator can clone to quickly deploy VMs.
- **Increase resource utilization:** Consolidation of multiple servers onto a small collection of ESX servers provides a significant increase in resource utilization and reduces space, power, and cooling requirements.

EqualLogic PS6010XV iSCSI Storage Array

- **Redundancy:** Available with two controllers, power supplies, and support for multiple RAID types, the EqualLogic PS6010XV array is designed for mission critical applications.
- **Thin provisioning:** Thin provisioning is a forward planning tool for storage allocation that provides all of the storage an application/VM might need upfront. When the iSCSI volume is created, a small percentage of the total capacity is physically allocated on the array. The allocated capacity for the volume grows real-time as the storage is requested.
- **Advanced performance monitoring:** Array performance monitoring is supported using SNMP, and through EqualLogic SAN Headquarters (SAN HQ) software. SAN HQ provides detailed utilization metrics along with historical data to provide actionable information to administrators.
- **Business continuity:** EqualLogic includes Auto-Snapshot Manager/VMware Edition (ASM/VE) that provides integration with vSphere for backup and restore capabilities through array-based VM snapshots.

PowerConnect M8024 Switches

- **Built for performance:** The PowerConnect M8024 supports up to 24 ports of wire-speed, 10 Gigabit Ethernet supporting 240 Gbps of switching capacity, and 357 Mpps forwarding rate. The PowerConnect M8024 provides ample performance for server aggregation and iSCSI storage.
- **Link aggregation:** Link aggregation groups (LAGs) allow multiple individual links to be combined and act as a single link utilizing the aggregate bandwidth of the links in the group. The reference architecture implements LAGs for the inter-switch links to ensure adequate bandwidth for cross-switch traffic. LAGs may also be utilized for connecting to any existing network infrastructure.

- Uplink to existing environment: In addition to the 16 internal 10Gb ports, the M8024 delivers flexible external I/O connectivity choices, supporting up to two 10Gb Base-T, CX-4 and/or SFP+ uplink modules.

Solution Components Details

VIS Self-Service Creator

The required components for a successful deployment include a VIS Self-Service Creator server, VIS Self-Service Creator agent(s), Microsoft® SQL Server, Microsoft Active Directory®, and a SMTP (mail) server.

- VIS Self-Service Creator server

The VIS Self-Service Creator server hosts the VIS Self-Service Creator software and Web interface that administrators utilize to configure the solution and perform VM lifecycle management activities; end users utilize the Web interface as well to request and manage their virtual machines. The VIS Self-Service Creator server interacts with the VIS Self-Service Creator Agents, SQL database, and SMTP server.

The VM that hosts VIS Self-Service Creator is required to run a supported version of Microsoft Windows Server® and should meet the following minimum requirements:

- vCPU: 2
- RAM: 4 GB
- Disk: 50 GB

Note: It is recommended to configure this virtual machine as a high-availability VM within vCenter.

If existing infrastructure is available to support the VIS Self-Service Creator server, it is not required to host the VM on the ESX cluster defined in this reference architecture. If this option is implemented, then it is recommended to dedicate a system (physical or virtual) to VIS Self-Service Creator.

- VIS Self-Service Creator Agent(s)

VIS Self-Service Creator agents are responsible for interacting with the virtualization platform and the hosts on which VIS Self-Service Creator provisions machines. In this reference architecture, the VIS Self-Service Creator agent for vCenter Server is installed on a dedicated virtual machine. If the existing infrastructure has other vCenter instances, or other supported provisioning methods, then the VIS Self-Service Creator server requirements and placement of the agents may be impacted; refer to the Installation Guide for more details.

The VIS Self-Service Creator agent host is required to run a supported version of Microsoft Windows Server and should meet the following minimum requirements:

- vCPU: 2
- RAM: 2 GB

- Disk: 20 GB

Note: It is recommended to configure this virtual machine as a high-availability VM within vCenter.

- Microsoft SQL Server

VIS Self-Service Creator uses a Microsoft SQL Server database to maintain information about the systems it manages and its own elements and policies. If a SQL Server already exists in the data center, then that instance can be used by VIS Self-Service Creator. Otherwise, the SQL instance can be hosted in a VM on the ESX cluster defined in this reference architecture.

The SQL Server requirements for a dedicated VIS Self-Service Creator instance are the following:

- vCPU: 2
- RAM: 2 GB
- Disk: 40 GB

Note: Configure no more than 75% of the total physical RAM for the SQL Server to use.

- Microsoft Active Directory

Access to the VIS Self-Service Creator console is based on each user's Windows credentials. VIS Self-Service Creator recognizes valid users of the Active Directory (AD) domain it is installed on, as well as valid local VIS Self-Service Creator server users. VIS Self-Service Creator authorizes each user for various activities based on the roles within VIS Self-Service Creator. If AD already exists in the data center, then that instance may be utilized by VIS Self-Service Creator; otherwise, AD can be hosted in a VM on the ESX cluster defined in this reference architecture.

- SMTP Server

VIS Self-Service Creator supports automatic emails to inform the administrators and users about required actions (e.g. approve a virtual machine build request) and VM build status (e.g. VM build complete and ready for use). If an SMTP server already exists in the data center, then that instance can be used by VIS Self-Service Creator; otherwise, the SMTP server can be hosted in a VM on the ESX cluster defined in this reference architecture.

vSphere

VIS Self-Service Creator supports several deployment methods, such as ESX template cloning, WIM Imaging, Linux Kickstart, and SUSE™ AutoYaST; these features require ESX/ESXi servers and a vCenter instance that manages those servers.

- ESX Cluster

The ESX cluster contains ten PowerEdge M610s running vSphere Enterprise Plus. The cluster can host both the VIS Self-Service Creator components, as well as the VMs deployed by VIS Self-Service Creator. The following configuration is the minimum recommended for each PowerEdge M610:

Component	Details
CPU	2 x Intel® Xeon® E5520, 2.26Ghz, 8M cache
Memory	48 GB
Ethernet Ports	4 x 10Gb

The four ports are provided by two add-in mezzanine network adapters (each with two network ports) that support redundancy for the virtual machine, VMKernel, service console, and iSCSI traffic. Details on mapping the network adapter to a virtual switch are available in the white paper, "Reference Architecture for VMware vSphere 4 in a 10 Gigabit iSCSI Environment with Dell PowerEdge Blade Servers".

- vCenter Server

The VIS Self-Service Creator agent communicates with the vCenter Server to perform various actions such as VM provisioning, control VM status (e.g. power on/off, reboot), and retrieve VM utilization/usage attributes. vCenter Server Standard Edition is used to provide support for large scale ESX clusters; other editions are supported, however the features and scalability goals identified in this paper may not be supported. The following configuration is the minimum recommended for the PowerEdge M610 management server:

Component	Details
CPU	1 x Intel Xeon E5520, 2.26Ghz, 8M Cache
Memory	8 GB
Ethernet Ports	4 x 10Gb

Networking

Four PowerConnect M8024 switches are used in this reference architecture to provide physical segregation of iSCSI and non-iSCSI traffic, and redundancy for each traffic type. Two switches are configured to support iSCSI traffic, while the other two support management, VM, and VMKernel traffic. Traffic isolation is accomplished for management, VM, and VMKernel traffic through the use of VLANs. Details on link aggregation, VLAN settings, and inter-switch links are available in the previously mentioned white paper, "Reference Architecture for VMware vSphere 4 in a 10 Gigabit iSCSI Environment with Dell PowerEdge Blade Servers".

iSCSI Storage

The EqualLogic PS6010XV array consists of two 10Gb Ethernet network ports per controller for data (iSCSI), and one port for out-of-band management; at publication, the array can be configured with capacities ranging from 4.8 to 9.6 Terabytes. If additional capacity is required, or the performance requirements necessitate an upgrade, then an additional array or arrays can be added. These arrays can be managed as a single unit, and data can be automatically distributed to the arrays to increase performance.

Additional Resources

- VIS Self-Service Creator Release Notes, <https://partners.dynamicops.com/DellVISCreator>
- *VIS Self-Service Creator User Guide*, <https://partners.dynamicops.com/DellVISCreator>
- *VIS Self-Service Creator Operating Guide*, <https://partners.dynamicops.com/DellVISCreator>
- *Reference Architecture for VMware vSphere 4 in a 10 Gigabit iSCSI Environment with Dell PowerEdge Blade Servers*, <http://www.dell.com/virtualization/businessready>
- Additional VIS Self-Service Creator Reference Architectures, <http://www.dell.com/vis>
- EqualLogic White Papers & Tech Reports, <http://www.equallogic.com/resourcecenter/documentcenter.aspx>