

# **BUSINESS READY CONFIGURATIONS FOR A VIRTUAL READY INFRASTRUCTURE**

**A Solution Guide**

**FOR DELL™ POWEREDGE™ BLADE SERVER, DELL  
EQUALLOGIC™ STORAGE, AND DELL ADVANCED  
INFRASTRUCTURE MANAGER™**

**Dell End-to-End Solutions Engineering**

**<http://www.dell.com/scalent/businessready>**



Information in this document is subject to change without notice.

© Copyright 2010 Dell Inc. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

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.

*Intel* is a registered trademark of Intel Corporation; *Microsoft* and *Windows* are registered trademarks, and *Hyper-V* is a trademark of Microsoft Corporation in the United States and/or other jurisdictions. 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.

# Contents

- 1 Introduction .....5
- 2 Audience and Scope.....5
- 3 Overview .....5
  - 3.1 Hardware.....6
  - 3.2 Software .....8
  - 3.3 Why Dell blades, Dell EqualLogic, Dell PowerConnect and Scalent?.....8
  - 3.4 Roadmap for rolling out a Virtual Ready Infrastructure .....9
- 4 Solution Specification.....10
  - 4.1 How to order.....12
  - 4.2 Dell Global Services .....12
- 5 Reference Architecture .....12
  - 5.1 Design Principles.....12
  - 5.2 Starting Configuration.....12
  - 5.3 Redundant Configuration .....13
- 6 Network Architecture.....14
  - 6.1 System Control Network.....15
  - 6.2 Other VLANs .....16
  - 6.3 Best Practices .....16
  - 6.4 Configuring Dell PowerConnect Ethernet switches.....17
- 7 Storage Configuration .....17
  - 7.1 Dell EqualLogic PS6000X .....17
  - 7.2 Volume Size Considerations .....17
  - 7.3 Array RAID Configuration.....17
- 8 Management.....18
  - 8.1 Scalent Management .....18
  - 8.2 Systems Management.....18
  - 8.3 Storage Management.....18
- 9 Deploying and Configuring the Solution.....19
  - 9.1 Phase 1 .....19
    - 9.1.1 Install Scalent Controller.....19
  - 9.2 Phase 2 .....23
    - 9.2.2 Migrate OS images to iSCSI target.....23

9.3 Phase 3 .....	30
9.4 Phase 4 .....	33
9.5 Troubleshooting.....	35
9.5.1 Best Practices.....	35
9.5.2 Known Issues .....	35
9.5.3 Basic Troubleshooting .....	35
10 References.....	36
11 Appendix.....	37
A Scalent Console Views .....	37
B Configuring Dell PowerConnect Ethernet switches .....	39

## **I Introduction**

The Business Ready Configurations for Dell™ PowerEdge™ blade servers, Dell PowerConnect™ switches, Dell EqualLogic™ SAN, and Dell Advanced Infrastructure Manager™ (Dell AIM) provides a detailed reference architecture for deploying and utilizing Dell AIM on Dell blades, switches and iSCSI storage environments. Based on extensive engineering work in architectural design and certification, customers can quickly and confidently deploy these proven architectures into production environments, thereby helping to eliminate much of the costly and time consuming trial and error often encountered during complex infrastructure design and implementation. The solution is designed to provide a completely managed infrastructure with the ability to heal itself in the event of failure. The solution includes the network architectures, switch configurations, storage configurations, and best practices necessary for deploying and configuring the solution. The solution will help customers achieve the full benefits of Dell AIM, Dell PowerEdge blade servers, Dell PowerConnect network switches, and Dell EqualLogic PS Series arrays.

## **2 Audience and Scope**

The intended audiences for this white paper are IT administrators, IT managers, and channel partners who are planning to deploy Dell AIM using Dell blade servers, Dell PowerConnect switches, and Dell EqualLogic SAN. This white paper provides an overview of the recommended servers, switches, storage, software, and services. It can be used to plan and procure the required components to set up a virtualization infrastructure.

This white paper provides reference architecture and best practices for deploying and configuring Dell's 11<sup>th</sup> generation blade server – M610, Dell EqualLogic PS6000 Series arrays, and Dell AIM. The solution uses iSCSI as the storage environment. Fibre Channel is not discussed in this white paper. The solution is designed to be standalone, with the blade enclosure(s) and storage arrays connecting to switches mounted in the rack. Based on customer requirements, further customization of the recommended architecture may be required.

## **3 Overview**

This section provides a high-level product overview of the Dell AIM, Dell PowerEdge blade servers, and Dell EqualLogic PS Series arrays communication through a Dell PowerConnect managed switch environment.

## 3.1 Hardware



### Dell Advanced Infrastructure Manager

- Rapid provisioning of server images
- Automatic Server Failover
- Server identity management for network and storage connectivity

### Dell PowerEdge Blade Servers

- Energy efficient PowerEdge M1000e enclosure
- M610 blade server
  - Intel “Nehalem” processor
  - Intel QuickPath memory technology
  - Lifecycle Controller
- CMC and iKVM for enclosure management



### Dell EqualLogic PS6000 series iSCSI Arrays

- (4) Gigabit Network ports per controller
- 2 GB cache per controller
- Intelligent and Automated management tools
- Self Managing Array
- Seamless scalability
- Top-tier data protection software



### Dell PowerConnect 6000 series Managed Layer 3 Gigabit Ethernet Switches

- 24 or 48 ports per switch
- Advanced layer 3 capabilities

Figure 1 – Overview of Dell Blade Server, Dell EqualLogic SAN and Dell AIM

### 3.1.1 Dell PowerEdge Blade Server

**Blade Modular Enclosure:** The Dell PowerEdge M1000e is a high-density, energy-efficient blade chassis that supports up to sixteen half-height blade servers, or eight full-height blade servers, and six I/O modules for the blade servers. A high-speed passive mid-plane connects the server modules to the I/O modules, management, and power in the rear of the enclosure. The enclosure includes a flip-out LCD screen for local configuration. The enclosure also includes six hot-pluggable and redundant power supplies and nine hot-pluggable N+1 redundant fan modules.

**Blade Servers:** The Dell PowerEdge M1000e supports the 11<sup>th</sup> generation Dell PowerEdge M610 blade servers based on the new Intel® Nehalem processors. The Dell PowerEdge M610 supports 12 DIMM slots. The blade servers use Intel QuickPath technology to provide a high-speed link to the memory modules. The 11<sup>th</sup> generation blade servers come with the next generation system management tool:

Unified Server Configurator (USC). This helps customers reduce operating costs by simplifying deployment and management. USC supports diagnostics, firmware updates, and hardware configuration.

**I/O Modules:** The enclosure provides three redundant fabrics using six I/O modules. The modules can be populated with Dell PowerConnect™ switches.

**Management:** The Dell PowerEdge M1000e has integrated management through a redundant Chassis Management Controller (CMC) module for enclosure management and integrated keyboard, video, and mouse (iKVM) modules.

For more information on Dell blade servers, see <http://www.dell.com/blades>.

### 3.1.2 Dell EqualLogic PS6000 Series iSCSI Arrays

The Dell EqualLogic PS6000 is the latest iSCSI storage device from Dell. Its features include four network ports per controller, faster processors, 2 GB cache per controller, support for RAID 6, increased drive capacity, and a new monitoring application, SAN HQ, at no additional cost.

In addition to the new features described above, Dell EqualLogic SAN devices provide the following capabilities:

**Reliability:** Dell EqualLogic PS6000 Series arrays have hot-swappable redundant components, a choice of RAID types, and hot-spare disks. They also include the Auto-Stat Disk Monitoring System (ADMS) which proactively scans disk drives in the background to help detect media anomalies and correct them.

**Scalability:** As each array is added to the storage group, the storage capacity and performance, in terms of both bandwidth and IOPS, are increased. This increased capacity can be utilized without downtime. EqualLogic PS Series arrays in a SAN work together to automatically manage data, load balance across all resources, and expand to meet growing storage needs.

**Self-Managing Arrays:** The arrays offer many self-managing features such as automatic load balancing and storage tiering. A single storage pool can have different models that offer a range of capacity and performance parameters. In addition, different arrays in a storage pool can be configured with different RAID levels, and volumes will automatically be migrated between the RAID levels based on performance data and usage patterns. All data and volume movement can be performed online with zero downtime.

**Top-Tier Data Protection Software:** Advanced data protection features such as Auto Replication and Auto-Snapshot Manager come standard. Role based administration and a full suite of data protection, availability and recovery tools to insure data integrity. Instant Volume Restore, Multi-Path I/O, Multi-Volume, Writeable Snapshots, Volume Cloning, and Volume Consistency Sets are all included.

For more information on Dell EqualLogic storage, see <http://www.dell.com/equallogic>.

### 3.1.3 Dell PowerConnect 6224 and 6248 Managed Layer 3 Gigabit Ethernet Switches

The PowerConnect 6224 AND 6248 are two of Dell's most advanced switch offerings with advanced core switching capabilities. These Gigabit Ethernet Layer 3 switches are stackable and offer optional 10 Gigabit Ethernet uplinks and support the latest version of the Internet Protocol—IPv6—enabling broader worldwide scalability.

Dell PowerConnect switches provide the following capabilities:

**Advanced Layer 3 Capabilities:** The PowerConnect 6224 supports advanced Layer 3 routing and multicast protocols to help reduce congestion and manage traffic in the network. Frequently-used LAN routing protocols such as RIPv1/v2, OSPFv2/v3, VRRP, IGMP, DVMRP, PIM and LLDP-MED, are also supported.

**Optional 10 Gigabit Ethernet:** The PowerConnect 6224 switch supports up to four 10 Gigabit fiber & two 10GBase-T copper Ethernet uplinks for connectivity directly to 10GE servers, routers, enterprise backbones, and data centers.

For more information on Dell PowerConnect switches, see <http://www.dell.com/networking>.

## 3.2 Software

Dell AIM software environment consists of the following components:

- **The Controller** is the Scalent software running on a dedicated server that manages the Dell AIM environment. The controller can be setup in an active/passive pair for redundancy. The controller does not reside in the data path.
- **Personas** are fully equipped boot images that run on servers the Controller manages. Personas are created from existing servers and can be locally booted or booted from a central storage repository like FC-SAN, iSCSI or NAS. Booting from a central storage offers the most flexibility and the greatest return on investment.
- **The Console** is a Flash-based web graphical user interface used to monitor and configure the Scalent environment from anywhere on the network. Administrative roles can be used to control access to the consoles capabilities.
- **The Scalent SDK** includes a command-line interface (CLI), a simulator, and APIs.

For more information on Dell's Advanced Infrastructure Manager Software, see <http://www.dell.com/scalent>

## 3.3 Why Dell blades, Dell EqualLogic, Dell PowerConnect and Scalent?

Dell AIM on Dell blade servers, PowerConnect switches, and Dell EqualLogic SAN provides a compelling infrastructure that meets the demands of today's data centers. The combined solution offers customers optimized resource utilization, simplified management, seamless scalability, and energy efficiency with no compromise in performance all in a compact package.

### 3.3.1 Simplified Management

One of the key benefits of Scalent on Dell blades and storage is simplified management. Scalent software integrates seamlessly with Dell's blade enclosure and PowerConnect switches providing customers with a simplified systems management interface. Integration and management through the CMC and iDRAC connections allows Scalent access to create and manage virtual network resources, power control, and storage access. The Dell EqualLogic PS Series arrays provide a single-pane management view of the entire storage pool in the SAN. The Dell EqualLogic PS6000 Series arrays intelligently balance workloads across the available arrays with no manual intervention.

### 3.3.2 Seamless Scalability

Dell servers and storage provide a highly scalable solution to meet the growing IT demands of today's data centers. The Dell M1000e enclosure and M610 blade servers with Scalent allow for the creation of server pools which provide automatic persona failover to provide maximum uptime. Scalent can manage any number of similar servers and manage the personas, network connections, and storage access on each.

### 3.3.3 Optimized Resource Utilization

Scalent's ability to rapidly re-deploy server Personas across pools of like servers provides data center managers the ability to change server roles based on demand or schedule.

### 3.3.4 Energy Efficiency

The Dell PowerEdge M1000e enclosure takes advantage of thermal design efficiencies such as ultra-efficient power supplies and dynamic power-efficient fans with optimized airflow design to efficiently cool the chassis and enable better performance in a lower power envelope.

## 3.4 Roadmap for rolling out a Virtual Ready Infrastructure

Any transformation in the data center infrastructure must be rolled out in a planned and careful manner. The responsibility for data center management, at least in large organizations, is distributed among multiple groups, typically along the lines of the 4 functional units – server, storage, networks, and operating systems/applications. Dell AIM provides a methodology for gradual transformation of a data center from static to dynamic. This method is designed to make the transformation more manageable, particularly in environments where existing operational policies must be respected – IP address administration, network access restrictions, storage security, service level requirements, etc.

This methodology begins with the server department. Ultimately, the requirements for a dynamic data center are driven by this department. Dell provides simple and nonintrusive mechanisms for deploying a Dell AIM environment and building an inventory of servers. Centralized booting is mandatory for creating a dynamic data center. Dell AIM includes tools and methodologies for migrating existing workloads to central storage and gradually transitioning from local booting to booting from central storage.

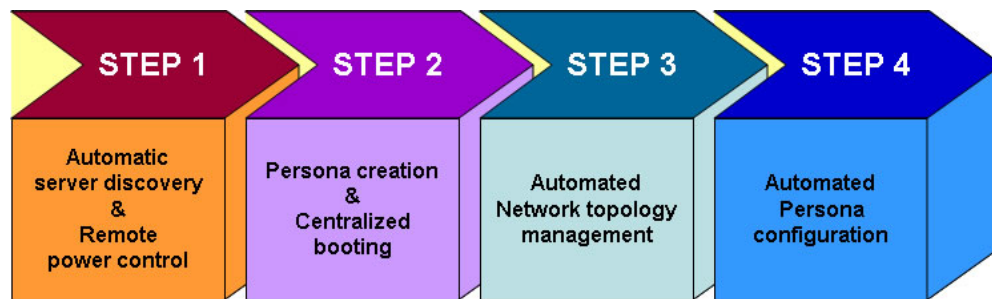


Figure 2 – Methodology for implementing Dell AIM

The next step in this methodology requires coordination across departments, particularly between the server and network departments. Dell AIM supports multiple modes of network management, such as read-only, fully managed, shared vs. dedicated control, that allow the network administrators more control over the pace of and the extent to which network changes are automated as a result of dynamic server repurposing.

The final step, automating server image configuration and integration with other management software in the ecosystem through Dell AIM abstraction and APIs, completes the transformation.

## **4 Solution Specification**

The following table provides the specification of the two business ready configurations. The remainder of the document discusses reference architectures for these configurations, how to set them up, and best practices for the deployment.

	Starter Configuration	Redundant Configuration
<b>Solution Summary</b>		
<b>Solution ID</b>	908256.1	908271.1
<b>Blade Chassis</b>	(1) M1000e with (2) Dell M6220 I/O modules	(2) M1000e with (4) Dell M6220 I/O modules
<b>Managed Blade</b>	(15) M610 with 6GB memory in each blade	(30) M610 with 6GB memory in each blade
<b>Storage Device</b>	(1) PS6000X	(2) PS6000X
<b>Management Blade</b>	(1) M610 with 6GB memory	(2) M610 with 6GB memory
<b>Chassis Configuration</b>		
<b>I/O module for A1</b>	Dell PowerConnect M6220	Dell PowerConnect M6220 per chassis
<b>I/O module for A2</b>	Dell PowerConnect M6220	Dell PowerConnect M6220 per chassis
<b>Management</b>	(2) Redundant Chassis Management Controllers (CMC)	(4) Redundant Chassis Management Controllers (CMC)
<b>KVM</b>	Integrated Avocent keyboard, video, and mouse (iKVM) switch	(2) Integrated Avocent keyboard, video, and mouse (iKVM) switch
<b>Managed Blade Configuration</b>		
<b>Blade Server Model</b>	M610	M610
<b>Processor</b>	(2) Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache	(2) Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache
<b>Memory</b>	6GB Memory (6x1GB), DDR3	6GB Memory (6x1GB), DDR3
<b>Local Storage and Controller</b>	SAS6/IR with (2) x 73GB 10K RPM SAS Hard Drives	SAS6/IR with (2) x 73GB 10K RPM SAS Hard Drives
<b>Storage Configuration</b>		
<b>Storage Device</b>	(1) PS6000X	(2) PS6000X
<b>Drives</b>	8 x 450GB, 10K SAS	8 x 450GB, 10K SAS
<b>Storage Capacity</b>	3.6 Terabyte capacity	3.6 Terabyte per device
<b>Management Blade Configuration</b>		
<b>Blade Server Model</b>	M610	(2) x M610
<b>Processor</b>	(2) Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache	(2) Intel Xeon (Nehalem) E5520, 2.26Ghz, 8M Cache
<b>Memory</b>	6GB Memory (6x1GB), DDR3	6GB Memory (6x1GB), DDR3
<b>Local Storage and Controller</b>	SAS6/IR with (2) x 73GB 10K RPM SAS Hard Drives	SAS6/IR with (2) x 73GB 10K RPM SAS Hard Drives
<b>Rack Configuration</b>		
<b>Rack</b>	Dell 2420 24U Rack Enclosure	Dell 4220 42U Rack Enclosure
<b>Top of Rack switch</b>	Dell PowerConnect 6224	(2) Dell PowerConnect 6224
<b>Software and Services</b>		
<b>Additional Software</b>	Scalent Activation Key - 15 managed (2-Socket) server nodes	(2) Scalent Activation Key - 15 managed (2-Socket) server nodes
<b>Services</b>	<ul style="list-style-type: none"> <li>• 3 Year ProSupport for IT and Mission Critical 4HR 7x24 Onsite Pack</li> <li>• EqualLogics PS Array Installation</li> <li>• M1000e Blade Chassis On-site installation service</li> </ul>	<ul style="list-style-type: none"> <li>• 3 Year ProSupport for IT and Mission Critical 4HR 7x24 Onsite Pack</li> <li>• EqualLogics PS Array Installation</li> <li>• M1000e Blade Chassis On-site installation service</li> </ul>

Table 1 – Solution Specification

## 4.1 How to order

To order the solution discussed in this paper, contact your Dell Sales Representative with the *Solution ID* provided above. Customers can order either of the two pre-configured solutions. Customization can be made to the configurations to fit specific needs. Based on the customization, some of the best practices discussed in this white paper may not apply.

The solution can also be directly ordered online with a few clicks using this link to Dell Business Ready Configurations for a Virtual Ready Infrastructure: <http://www.dell.com/scalent/businessready>. Customization can also be made to the configurations using the Dell online ordering tool.

## 4.2 Dell Global Services

Today's financial environment dictates that IT reduces its budget while still providing ever increasing infrastructure services. To assist with this goal, Dell Global Services can also be engaged to order the solution. Dell Global Services helps customers find the "better path" to a virtual ready infrastructure to reduce total cost of ownership, speed time to ROI, increase agility, and reclaim IT resources.

## 5 Reference Architecture

This section describes the reference architecture for the solution.

### 5.1 Design Principles

The following principles were used during the design of this architecture.

- **Redundancy with no single point of failure:** Redundancy is incorporated in every aspect of the solution, including networking and storage.
- **Scalability:** The solution is highly scalable and is based on the architecture of the M1000e chassis and the EqualLogic PS6000 Series arrays. Guidelines to deploy additional blade servers and storage to increase the compute and storage capacity respectively are included.
- **Isolated, redundant and high-performance network design:** The network is designed to support isolation of various networks through the use of VLANs. It is designed to have no single point of failure and have optimal performance when used in conjunction with Dell AIM.

### 5.2 Starting Configuration

This solution consists of a Dell M1000e chassis populated with M610 blade servers. One blade is used for management (Scalent Controller) and the rest run iSCSI booted personas managed by the controller. A PowerConnect 6224 switch serves as a top of rack switch, and connects the EqualLogic PS6000X array to the network. Figure 3 illustrates a high-level reference architecture for the solution based on the starting configuration described in the Specification section.

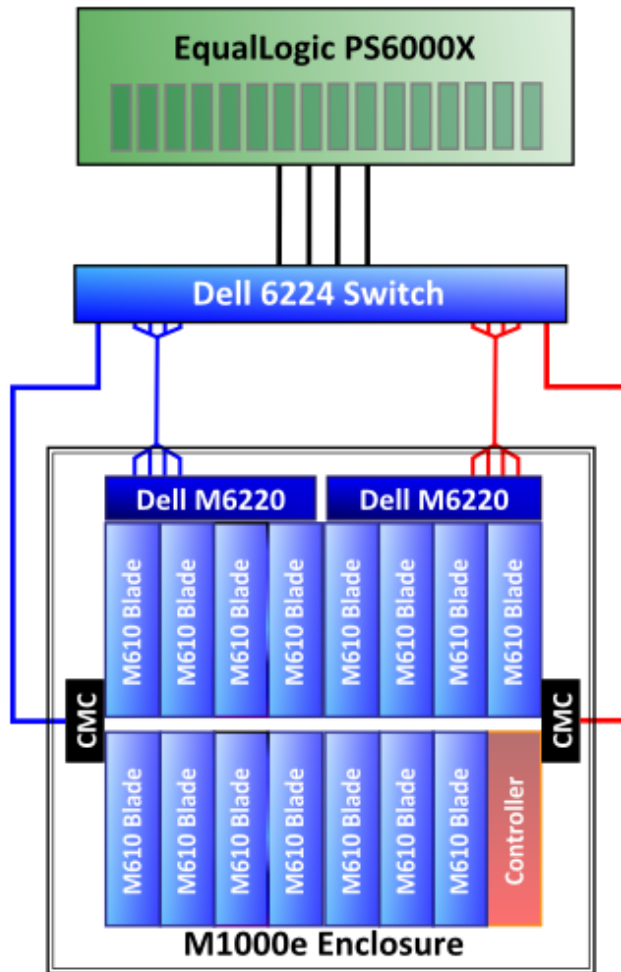


Figure 3 – Starting Configuration

### 5.3 Redundant Configuration

This solution consists of two Dell M1000e chassis populated with M610 blade servers. One blade in each chassis is used for management (resilient Scalent Controllers) and the rest run iSCSI booted personas managed by the controllers. Two PowerConnect 6224 switches serve as top of rack switches, and connect the two EqualLogic PS6000X arrays to the network. The EqualLogic arrays are configured in a storage group, allowing the servers to access them from a single group address. Figure 4 illustrates a high-level reference architecture for the solution based on the redundant configuration described in the Specification section.

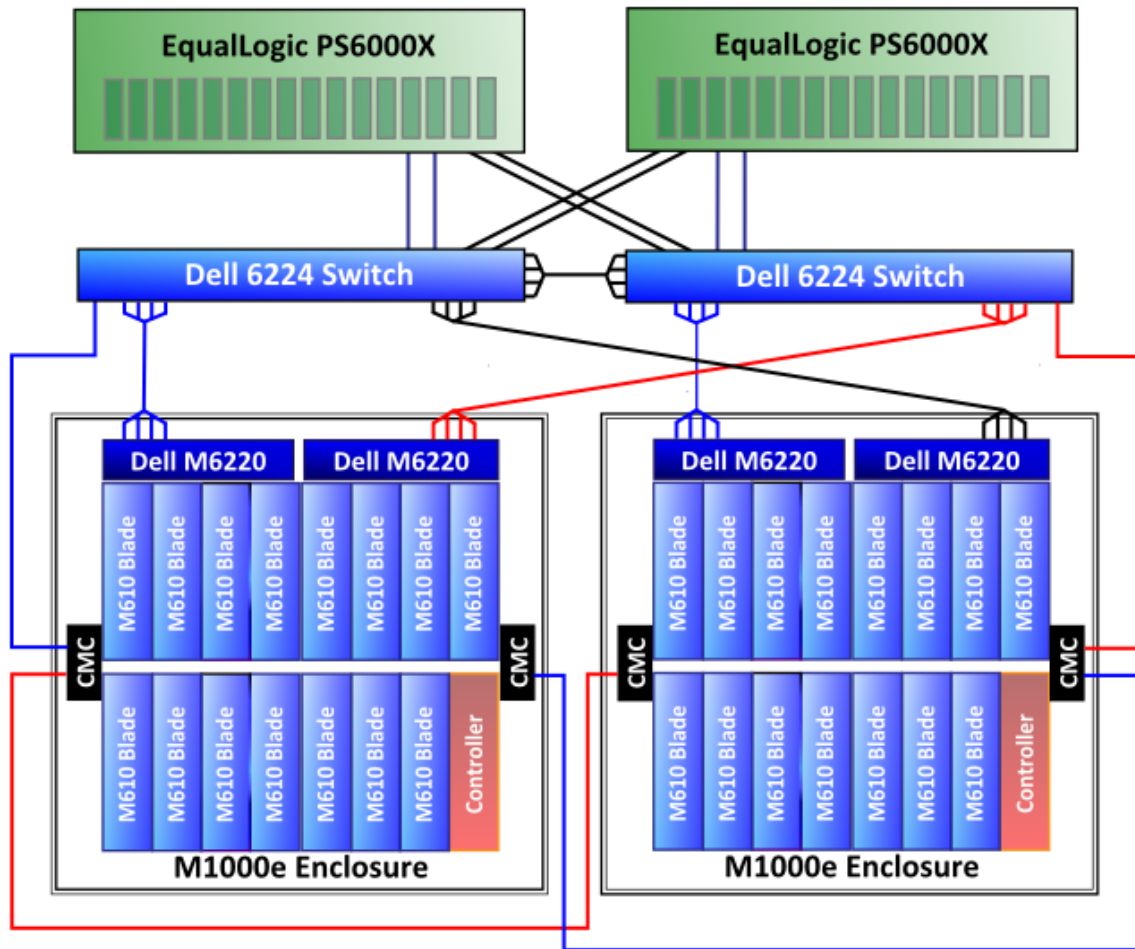


Figure 4 – Redundant Configuration

## 6 Network Architecture

Design Principles for Network Architecture: The following design principles are used to develop the network architecture to support the above traffic types:

- Redundancy: Blade chassis have redundant I/O modules. Redundancy of the network adapters is achieved through VNICs on the VLANs and redundant managed switch modules.
- Simplified management through Virtual Networking: Switches servicing the same traffic type are combined into logical fabrics using VLANs managed through the Scalent software.
- iSCSI SAN physical isolation: The iSCSI SAN should be logically separated from the LAN to reduce any network congestion.

- Optimal performance: Link aggregation and iSCSI load balancing is used to achieve the highest throughput possible.

## Physical Network

The solution network consists of a blade chassis connected to interconnect switches. The Dell blade chassis can support three separate fabrics; this solution will use only the first fabric. Each blade chassis will contain two Dell M6220 PowerConnect switch modules. Each switch module will be connected to an interconnect switch, a Dell PowerConnect 6224, via an aggregated link of four ports. Each switch module will also have two connections from each iSCSI storage controller to each switch.

In the case of the *Redundant* configuration, the interconnect switches will also be connected to each other via an aggregated link of four ports. The *Redundant* configuration also provides redundancy via the second switch and redundant cable pathways.

## Virtual Network

Virtual networks behave in much the same way in a Scalent-managed environment as conventional networks in a conventional data center, except that they're built out of virtual LANs (VLANs) running over physical NICs, switches, and cables, instead of those physical elements themselves.

Because they are virtual, Scalent virtual networks are easier to create and manage than conventional networks: simply drag as many vSwitches, vNICs, and other resources as needed into the workspace in Virtual Networks view, draw cables between them to connect them, experiment with different topologies, and configure all their properties in the Console. There is no need to configure physical NICs on individual machines, configure physical switches, or move physical cables.

Now virtual LANs can be created that segregate network traffic, without the constraints of physical LANs. Also, Controller-assigned networks simplify how IP address ranges are allocated and managed.

## 6.1 System Control Network

The Scalent Controller uses the System Control Network (SCN) to discover new servers and their capabilities, to communicate status and configuration changes between itself and personas and VMRacks, to connect servers with the network storage devices that contain the images from which personas and VMRacks boot, and to manage many other aspects of how personas are configured, including how they are connected to virtual networks.

For communication from the Controller to elements that aren't on the System Control Network (SCN), such as persona storage servers, chassis and server management modules, switch management interfaces, and so on, the Controller must have a route to those networks, typically via the default gateway for the Controller real IP addresses.

In a typical production deployment, the Controller is connected to a network that is independent of the Scalent environment.

Because the Controller responds to DHCP requests from elements on the SCN in the Scalent environment, the Controller (or at least its connection via the SCN Services interface) should reside in its own broadcast domain to prevent responses to DHCP requests from servers it doesn't manage.

## 6.2 Other VLANs

Following are some examples of additional VLANs that can be created in a virtual ready infrastructure.

### 6.2.1 iSCSI VLAN

This solution isolates iSCSI traffic to a dedicated iSCSI VLAN. This VLAN will provide access to the Dell EqualLogic PS6000 storage that will house the images of network-booted personas.

### 6.2.2 Lights-out management VLAN

The vRack (blade chassis) servers' Chassis Management Controllers (CMC) will have addresses on these networks and are connected to them.

The router transmits traffic between this network and the Controller's Controller Services address or to the Controller's SNMP Trap Collector address if a separate address was configured for it.

### 6.2.3 Active Directory VLAN

A VLAN will be created that will allow for a separate broadcast domain for Active Directory services such as DHCP and DNS. A separate network insures that client DHCP requests are not answered by the Controller.

### 6.2.4 Controller VLAN

A VLAN dedicated network to which the Controllers are connected. This network is in its own broadcast domain, so it doesn't interfere with DHCP servers for elements that aren't in the Scalent environment. The Controllers' real IP addresses (set in their operating systems) and the Controller Services virtual IP address are on this network.

## 6.3 Best Practices

The System Control Network (SCN) is VLAN 4004. This VLAN contains the SCN discovery, DHCP and PXE services. Both of the Controllers, all switches, and all managed servers should use this network.

In both configurations, the PowerConnect 6224 switches should be configured as routers to enable communication across all VLANs. In the redundant configuration, setting up Virtual Router Redundancy Protocol (VRRP) will allow routing to continue even if one of the two interconnect switches fail.

When configuring Scalent's switch permissions, all ports not being controlled by Scalent should be set as "Unmanaged". Any port not marked as "Unmanaged" will be shut down automatically if it is not configured in the console.

Any VLAN manually created in the Dell AIM environment (such as the private VLAN used between resilient controllers) that are to remain active once Scalent takes full control of switches must have an external network with the same VLAN created in the Scalent console.

### 6.3.1 Routing

Routing between all of the VLANs and the SCN must be properly configured in order for the Scalent controllers to monitor and manage the VLANs and their associated resources. The Scalent controller must have access to all networks that contain manageable resources.

## 6.4 Configuring Dell PowerConnect Ethernet switches

The network switches in the setup need to be configured for the Scalent Controller to manage them. This is done by connecting to the switch via the CMC for M6220 switches, or through to serial console on 6224 switches. The System Control Network (SCN) needs to be added to the switches, and aggregated links configured for maximum throughput and resiliency. This configuration is covered in greater depth in Appendix B.

## 7 Storage Configuration

### 7.1 Dell EqualLogic PS6000X

The storage configurations for this solution are based on the configuration selected.

**Starting Configuration:** This configuration uses a single Dell EqualLogic PS6000X array. A single storage group will be created and assigned an IP address on the iSCSI VLAN. The storage contained in the array will be used to store the Dell AIM database and all iSCSI booted personas. If more storage is needed, another array can be installed and subsequently added to the original storage group. Once the additional array is added to the group, it will be available to the environment.

**Redundant Configuration:** The redundant configuration uses two Dell EqualLogic PS6000X arrays. As with the starting configuration, a storage group is created on the first array and an IP address is assigned on the iSCSI VLAN. Once the array and group are configured, the second array can be added to the storage group, which will make the combined storage available to the environment. Additional arrays can be added to provide more storage capacity if desired.

PS Series arrays provide dynamic load balancing. As the workload changes, volume data and network I/O are load balanced within and across the arrays in a group with no impact on applications and no user intervention. The more disk spindles available the better the overall performance will be, because all active disks per array are always utilized and data is spread across all disks within the RAID sets per 3U enclosure.

All iSCSI traffic travels over standard switched networks and the Dell AIM solution is designed to manage these networks. This makes iSCSI storage a particularly compelling storage solution in this environment.

For more information on Dell EqualLogic storage, see: <http://www.dell.com/equallogic>.

### 7.2 Volume Size Considerations

Volume sizes depend on the customer environment and the type of workload. Remember to size the volume appropriately to allow space for all software and updates required for the persona's intended function. Thin-provisioning can be utilized to allow volumes to grow on demand when additional storage is necessary, which will increase storage utilization efficiency.

### 7.3 Array RAID Configuration

The RAID type chosen is highly dependent on the workload in the environment. The Dell EqualLogic PS6000 array supports four RAID types – RAID 5, RAID 6, RAID 10, and RAID 50. RAID 10 will generally provide maximum throughput and performance at the expense of storage capacity. RAID 50 will provide more usable storage, but has less performance than RAID 10 in random I/O situations and requires

additional overhead in the case of a drive failure. RAID 5 provides the most storage capacity at the expense of slightly lower performance and availability.

## 8 Management

### 8.1 Scalent Management

The Scalent software solution provides unified management of infrastructure resources in a data center. The data center abstraction provided by Scalent effectively unifies the management of a data center infrastructure that is potentially heterogeneous in each and every operating department. Data center infrastructure is constantly undergoing change. The abstraction provided by Scalent insulates higher order management software from underlying changes in infrastructure. Access to the abstraction is provided conveniently through a scriptable command line interface and web services. Although this document is mainly focused on managing a homogeneous environment comprised of Dell PowerEdge blade servers and Dell PowerConnect switches, the true value of Scalent lies in its ability to manage heterogeneous infrastructure. These include heterogeneous servers (regardless of form factor, make and model), network switches, and central storage repositories. Scalent also provides a unified management of physical and virtual resources. The virtual resources include hypervisors such as VMware ESX, Microsoft® Hyper-V™, and Red Hat Xen, as well as virtual machines that are deployed on these hypervisors.

### 8.2 Systems Management

The following tools and software can be used to manage the hardware.

#### Dell OpenManage

Dell OpenManage Server Administrator (OMSA) can be installed on Dell AIM personas and used to manage the host hardware. For more information on Dell OpenManage and its capabilities, see [www.dell.com/openmanage](http://www.dell.com/openmanage).

#### Deployment and change management using Unified Server Configurator (USC)

The Dell PowerEdge M610 blade servers come with USC. This helps customers reduce operating costs by simplifying deployment and management. Key features include: diagnostics, self-update (UEFI, Driver Pack update), firmware updates (BIOS, NIC FW, RAID Controllers), and hardware configuration.

#### Out-of-band CMC and iDRAC

CMC provides a single, secure interface to manage the inventory, configuration, monitoring, and alerting for chassis components (iKVM, CMC), I/O modules, servers, and iDRAC. It also provides excellent real-time power management, monitoring, and alerting capabilities. The Dell chassis also provides the users with system-level power limiting, slot-based prioritization, and dynamic power engagement functionalities. iDRAC on each server provides the flexibility to remotely manage the server through console redirection and virtual CD-ROM/DVD/floppy/flash capabilities.

### 8.3 Storage Management

Dell EqualLogic provides a rich set of management features that is available at no additional cost and come with exceptionally easy-to-use management tools. For more information on Dell EqualLogic features and management capabilities, see [www.dell.com/equallogic](http://www.dell.com/equallogic).

## 9 Deploying and Configuring the Solution

### 9.1 Phase 1

#### 9.1.1 Install Scalent Controller

##### **Pre-Requisites**

The Scalent Controller can be installed in a dedicated blade server with the following minimum configuration:

- 2 CPUs (two sockets).
- 2 GB of RAM.
- 40 GB hard disk.
- An optical drive to install Red Hat Linux and Scalent software on the server

One of the NICs in the server set aside to function as the Controller will be used for connecting to the Scalent environment, hosting the Console and SDK connections, receiving SNMP traps, connecting to managed devices, and connecting to storage arrays.

To install the Scalent Controller, you will need:

- A licensed copy of the 32-bit Red Hat Enterprise Linux 5 Update 2 operating system
- Scalent Linux ISO distribution

##### **Installation checklist**

It is advised to go over the Scalent Pre-Installation Checklist included with the “Deployment Planning and Scalent Controller Installation Guide”. When installing the Scalent Controller software, one of the two Controller installation options must be selected: basic or advanced.

A basic installation prompts for the minimum information required to install the Controller. Any additional configuration changes can be made by connecting to the Scalent Console and using the Console’s graphical user interface. This is the best choice for a standalone deployment.

This information is needed before beginning the installation process:

1. Controller IP address – This is the IP address that is assigned to the primary NIC of the server and is configured typically during the process of OS installation.
2. System ID – More than one Scalent environment can be installed in the same data center, or in data centers that can communicate with each other. To ensure that the MAC addresses and other configurations the Controller creates are unique, a unique number from 0 to 31 must be assigned to each Scalent environment when installed.

3. Controller Services Virtual IP Address – This is the IP address that is configured to host the three Controller services: Scalent Console access (from GUI, CLI and SDK), communication with elements in the Scalent environment, and receiving SNMP traps sent by network switches and other devices. In a basic installation, a single IP address is configured to host all three services. This is a virtual IP address and gets instantiated when the Controller starts and typically binds to the same NIC on which the Controller's real IP address is configured.
4. System Control Network (SCN) and Scalent DHCP – This is a private network used by the Controller to communicate with the managed entities in the Scalent environment. The Scalent DHCP network is used temporarily during the process of discovering new servers. The range of IPs for both networks can be modified post installation. However, it is preferable to have this information entered during the installation process.

**Note:** Dell FlexAddress should not be used in conjunction with the Dell AIM solution.

### **Installation Process**

1. Install the Red Hat 5 Update 2 operating system on the server set aside to function as the Scalent Controller. Most default settings should suffice with the exception of SELINUX which should be disabled. This can be done during the installation or later. Detailed instructions can be found in the “Deployment Planning and Controller Installation Guide”. It is recommended that ‘iptables’ is enabled prior to installing the controller software.
2. Install and Configure Scalent Controller Software using the steps listed below.
  - a. Mount the Scalent Linux media and install the controller RPMs (packages) by executing the “installController.sh” utility found in the root directory of the image. This utility will also install any missing OS packages needed for the functioning of Scalent Controller. This process will typically take less than 5 minutes.
  - b. Configure the Scalent Controller by executing the “/opt/scalent/bin/setupController.sh” utility. This utility will bring up a wizard that will walk through the configuration process. The questions are self explanatory and pick the “Basic Installation”. Once the information that was gathered in the previous step is entered, the utility configures the installation and will start the Scalent Controller after confirmation.

Note that in a “Basic” installation, very few inputs are needed during the configuration process. The installer will assume default values for all the networks that are managed by the Scalent controller. These have been listed below.

**Infrastructure VLAN range**– The range of VLANs that the controller uses to create the infrastructure that supports the Scalent environment including the System Control Network (SCN). The default range is VLAN 4002 to VLAN 4089. The default SCN VLAN is 4004.

**Scalent Assigned VLAN range** – The range of VLANs available for the controller to assign to virtual networks that are created. The default range is VLAN 3746 to VLAN 4001.

**Customer Assigned VLAN range** – The range of VLANs that can be assigned to virtual networks created in the environment. The default range is VLAN 2 to VLAN 3745.

VLANs that are not in one of these ranges are not managed by the Controller.

## **Verify Installation**

After the installation process is completed, it can be verified by checking if the Scalent service is running. The Scalent Controller can be easily accessed through the Scalent Console. Open a web browser on any system that has connectivity to the network where the Scalent Controller is installed. Enter the Scalent Controller Services Virtual IP address in the URL. The default username and password to login to the Scalent console are both 'admin'. On the initial log in, the environment will have no elements under management. The Scalent console provides multiple views of the environment. The most commonly used views are the physical network, virtual network and catalog views.

**Physical Network View** – This view shows all the physical elements of the environment. These include physical network switches and their ports, chassis, blades, servers, external switch ports, and physical cabling.

**Virtual Network View** – This view enables creation of a network topology comprising of server images (or Personas). The Personas are interconnected by means of virtual switches (VLANs) and virtual cables.

**Catalog View** – The Scalent controller manages a variety of elements in the environment. These include servers, racks, personas, switches, and so on. The catalog view groups all the like elements together and presents them either in the form of a list or icons.

Other information entered during the configuration process (SCN and DHCP address range, System ID, etc) can be seen in the **environment view** of the Scalent Console.

Sample screenshots of the various views provided by the Scalent console have been included in Appendix A.

### **9.1.2 Installing a Resilient Scalent Controller (Optional)**

In order to protect against the failure of a single Controller, Scalent provides an option to install and configure a secondary controller. This can be done during the initial install or at a later stage.

#### **Pre-Requisites**

To deploy a pair of resilient controllers, you will need:

1. Oracle Cluster File System 2 (OCFS2) installation on both the servers that will function as Scalent Controllers. The software packages are included in the Scalent distribution.
2. Shared Database location – The resilient controllers need to have access to a common shared data store that will function as the Scalent Configuration repository. The shared data store is a 10GB data LUN on the Dell EqualLogic iSCSI SAN. Details about configuring the shared storage can be obtained from the "Deployment Planning and Scalent Controller Installation Guide".
3. OCFS2 Heartbeat Network - This is a dedicated network between the Controllers for the OCFS2 lock manager to exchange keep-alive traffic. In a blade environment, this is accomplished by assigning a unique VLAN to the switch ports connected to the NICs configured for the private network.

## **Installation Process**

Installing the secondary controller is very similar to the process earlier described for installing a single controller.

1. Install the Red Hat 5 Update 2 Operating System on the both of the Controllers.
2. Verify that both of the Controllers can access the shared LUN and can communicate on the dedicated heartbeat network. If the hosts are not stored in DNS, the host names have to be manually added to the “/etc/hosts” file on both the hosts.
3. Install the OCFS2 packages from the Scalent software distribution using the command “installcontroller.sh resilient=yes” on the primary node.
4. Configure the Controller using the “/opt/scalent/bin/setupcontroller.sh” utility and during the process enter the OCFS2 related information.
5. Repeat steps 3 and 4 on the secondary Controller.

### **9.1.3 Accessing Scalent Controller**

In addition to the web console, the Scalent Controller can be accessed using the CLI. This can be done in 2 ways:

1. On the Scalent Controller server – Login to the Scalent Controller using a secure session and execute the command “/opt/scalent/bin/sdk”. This will bring up the Scalent CLI prompt. Providing an optional argument “tty=true” will enable tab completion for the CLI session.
2. Any Windows Client – The Scalent SDK client can be installed and invoked on any desktop/server running Windows XP or later. Typically the executables are stored in the C:\Scalent\SDK\bin folder and the CLI is started by executing the voeCli.bat script. Once the CLI starts up, the following 2 commands need to be executed:

```
OPEN
```

```
SET HOST=<IP Address of Scalent Controller>
```

### **9.1.4 Discovering Servers**

Scalent Controller can discover running servers in the data center without having to install any agent software. Server discovery tools for each of the supported operating systems are included in the Scalent software distribution.

1. Discovering servers running Windows: Login to the server as Administrator and configure the Windows firewall to allow ICMP requests. This can be done from the GUI or by issuing the following command:

```
>netsh firewall set icmpsetting 8 enable
```

Execute the batch file discovery\discover.bat from the Scalent software distribution for Windows. There is no need to copy the utility to the server.

2. Discovering servers running Linux: Login to the server as root user and run the shell script “discover.sh” from the root directory of the Scalent software distribution for Linux. Yet again, the software does not have to be copied or installed on the server.

Discovering Bare Metal Servers: For servers that are devoid of any operating system, the discovery process is different. After the server is powered on, the BIOS needs to be configured and PXE booting on one or both of the NICs should be enabled. When the server starts booting, the Scalent Controller will respond to the PXE boot requests and boot a discovery image on the server. The discovery image will inspect the system configuration and the information is stored in the Scalent configuration repository.

**Note:** Dell FlexAddress should not be used in conjunction with the Dell AIM solution.

## 9.2 Phase 2

### 9.2.1 Server Power Control

During the server discovery process, the Scalent controller will also discover the information needed to access and login to the Integrated Dell™ Remote Access Controller (iDRAC) on the servers. In the case of blades it is the information related to the IPMI interface. After the completion of server discovery process, this information can be accessed from the Scalent console for each of the discovered servers. At this stage, the servers can be powered on/off from the Scalent Console regardless of whether they are disk booted or net booted.

In the Scalent Console physical view, point to the server running Windows that was discovered using the server discovery utility. The lights out management can be entered by selecting the ‘management settings’ associated with that server. After entering that information, the server can be powered off by clicking on the “Stop Persona” link. Servers running Windows OS need to be powered off before proceeding to the next step. ***Before shutting down the server, there is one important step that is essential for the image to be able to boot over iSCSI. Log into the Microsoft Windows® server and ensure that the Microsoft iSCSI initiator is installed and enabled for booting. This is handled differently depending on whether a Windows 2003 image or a Windows 2008 image is being migrated.***

Windows 2008 – By default a Windows 2008 installation includes the Microsoft iSCSI boot initiator. This service must be set to “automatic”.

Windows 2003 – If the initiator is not already installed, download the initiator from Microsoft for the appropriate architecture – 32 or 64 bit. While installing the iSCSI initiator, the boot over iSCSI option should be checked and the appropriate NIC on the server should be selected.

### 9.2.2 Migrate OS images to iSCSI target

#### Storage Configuration

Before proceeding with the migration process, a volume of appropriate size will need to be created in the storage array and setup the appropriate access control to the volume using the Dell EqualLogic Group Manager management tool. Refer the product documentation for detailed instructions.

## Scalent Migration Utility

The Scalent software includes migration utilities for Windows and Linux. The primary function of these utilities is to copy boot and data images from local disk to a central storage repository. After the migration, the utility also carries out offline editing to enable the image to boot from central storage on the source server as well as other compatible servers. In this case, we will focus the discussion on migrating to the Dell EqualLogic iSCSI storage array.

### Windows edition

The migration utility is one of the DVDs in the media kit. Alternatively, download the bootable ISO which can be burned on a CD or copied to a bootable USB drive. The server to be imaged needs to be shut down and booted to the migration utility. The migration utility which is based on Microsoft Windows Preinstallation Environment (Windows PE) presents a menu driven interface. Execute the following steps to carry out the migration:

1. Create a volume in the EqualLogic storage and configure the access so that this server can have visibility to the volume. This can be done in multiple ways. Refer to the EqualLogic documentation for the procedure.
2. From the Advanced Menu in the migration utility, establish connectivity to the iSCSI target and verify visibility to the volume.
3. Go back to the Main menu and copy the boot image by selecting the local OS disk as source and the iSCSI volume as destination.

The migration process is very fast and efficient and typically averages 2 GB/min. Once the migration completes, the server can be shutdown.

### Linux edition

The Linux migration utility is included in the Linux ISO as a standalone executable called "copy\_persona.sh". This utility is a general-purpose tool for copying personas from one kind of storage device to another. It can also be used to partition a device before copying a persona onto it. This section covers the simple case of copying a Linux image running on a local disk to an iSCSI volume, using the default partitioning settings. Execute the following steps to carry out the migration of a Red Hat 5 image:

1. Create a volume in the EqualLogic storage and configure the access so that this server can have visibility to the volume. This can be done in multiple ways. Refer to the EqualLogic documentation for the procedure.
2. If necessary, install the iSCSI initiator utilities on the source image. The package to be installed is `iscsi-initiator-utils-6.2.0.742-0.5.el5.i386.rpm`. It can be found on the Red Hat Linux installation DVD.
3. On the source server, edit the `/etc/iscsi/initiatorname.iscsi` file and add the initiator name planned for this migration. After saving the file, restart the iSCSI daemon.

4. Force the iSCSI process to discover the iSCSI target server. The command in this case is

```
# iscsiadm -m discovery -type st -portal target_ip:target_port
```

For example:

```
# iscsiadm -m discovery --type st --portal 10.30.2.5:3260  
10.30.2.5:3260,1000 iqn.2001-05.com.equallogic:0-8a0906-0cd0f7b04-  
0dd9fef9f324a958-sql-server
```

The target name returned in the previous command needs to be provided as input to the next command.

```
# iscsiadm -m node --target targetname --portal targetip:targetport --login
```

For example:

```
# iscsiadm -m node --target iqn.2001-05.com.equallogic:0-8a0906-  
0cd0f7b04-0dd9fef9f324a958-sql-server --portal 10.30.2.5:3260 --login
```

```
Login session [iface: default, target: iqn.2001-05.com.equallogic:0-  
8a0906-0cd0f7b04-0dd9fef9f324a958-sql-server, portal: 10.30.2.5,3260]
```

5. View the available iSCSI devices by executing the 'fdisk -l' command and initiate the migration by executing the "copy\_persona.sh" utility

When an image is copied from local disk to an iSCSI volume, there needs to be a differentiation between how the iSCSI volume is mounted now and how it will be mounted when it becomes the boot device. For example, a disk-booted Linux image typically boots from /dev/sda and mounts the iSCSI volume as the next available device /dev/sdb. However, when booting the same image from the iSCSI target in the future, it will mount its boot device as /dev/sda. The current and future iSCSI volume paths must be specified when using copy\_persona.sh.

In the following command:

- o Use the -d argument to specify the iSCSI volume onto which the image will be copied, which was determined when the 'fdisk -l' command was executed.
- o Use the -r argument to specify the device the new persona will boot from after it is created. Typically this is /dev/sda.
- o Use the -t argument to specify the size of the swap disk needed, in megabytes (3000, or 3 gigabytes, is a typical value).
- o The argument -n n is appended to specify that the iSCSI-booted image will not have a Dell AIM agent installed. Agent installation and functions provided by the agent is discussed in a later section.

```
# /opt/scalent/bin/copy_persona.sh -d /dev/sdb -r /dev/sda -t swapsize -n n
```

After checking that there's enough room on the iSCSI volume, copy\_persona.sh prompts for confirmation that the partition on the destination iSCSI volume should be deleted (answer yes). Then the script asks if

Local Volume Manager (LVM) on the destination device should be used, presents a default partitioning scheme for the iSCSI volume and asks if it should be used. The `copy_persona.sh` utility creates partitions on the iSCSI volume, copies the disk-booted persona onto the iSCSI volume, and reports success when it's done.

### 9.2.3 iSCSI Bootable personas

In a Scalent environment, a bootable image is often referred to as a 'Persona'. The Persona definition includes the OS image as well as other attributes that allows the image to be retargeted across different servers. In this section, we will see how the images that were migrated in the previous section can be booted up on the source server itself.

#### Adding an iSCSI booted Windows Persona

1. From the Scalent Console, delete the previously discovered server that was running Windows OS on the local disk.
2. The next step is to add the disk image that was migrated to the iSCSI target, as a retargetable Windows Persona in the Scalent environment. This can be done either using the Scalent Console or the CLI.
3. Adding an iSCSI-booted Windows Persona using the Scalent console:
  - a. Switch to the "Catalog" view in the Scalent console.
  - b. In the left side bar, select Personas and click on "New Persona" seen at the bottom of the sidebar.

Please enter the following information in order to add a new persona to the system.

ID	<input type="text"/>	
Name	<input type="text"/>	
OS Family	linux	▼
OS Architecture	x86_32	▼
Boot Type	iSCSI-Booted [iscsi]	▼
Server Type	<input type="text"/>	Image ID <input type="text"/>
Target IP Address	<input type="text"/>	<input type="checkbox"/> Template
Target Name	<input type="text"/>	
Initiator Name	<input type="text"/>	
Device	<input type="text"/>	
File System Type	<input type="text"/>	

Figure 5 – Adding a new iSCSI booted persona

- c. In the add a new persona dialog box, enter the following mandatory information
    - i. Name – Although this is optional, it is a good practice to assign a meaningful name to all the Personas in the Scalent environment.
    - ii. OS Family – Select Windows
    - iii. OS Architecture – x86\_32 or x86\_64 depending on the Windows build that was installed on the source server
    - iv. Boot Type - Select iSCSI
    - v. Server Type – Enter “EqualLogic” (without quotes)
    - vi. Target IP Address – Enter the IP address of the EqualLogic IP SAN
    - vii. Target Name – This should be obtained from the EqualLogic storage’s group manager console. This will be something like iqn.2001-05.com.equallogic:0-8a0906-0cd0f7b04-0dd9fef9f324a958-sql-server
    - viii. Initiator Name – This is the name that was configured in the EqualLogic Storage that will have access to the target. This will be a string like iqn.1991-05.com.microsoft:sql-server.
    - ix. Uncheck the “Agent Exists” check box.
  - d. Click OK to add the Persona and save the changes. A new dormant Persona will be added to the catalog view in the Scalent console.
4. Alternatively, the Scalent CLI can be used to add the Persona by issuing a command similar to the one shown below.

```
add persona osFamily=windows osArch=x86_64  
bootType=iscsi image="iscsiwbi|EqualLogic|10.30.2.5|3260| iqn.2001-  
05.com.equallogic:0-8a0906-0cd0f7b04-0dd9fef9f324a958-sql-server | iqn.1991-  
05.com.microsoft:sql-server |iscsiwbi" name="Win-iSCSI-Test "
```

The parameters in the CLI command are the same as shown above in step 3.

### Booting the newly created iSCSI booted Windows Persona

In order to boot the Persona created in the previous step, the source server must be re-discovered as a bare metal server. This can be done by rebooting the server and changing the boot order in the BIOS so that it boots over the network. The server will be discovered and be visible in the Scalent console. The lightning bolt symbol on the server is a representation that the server is configured to boot from the network. If the IPMI credentials are correctly entered, the server will be automatically powered off. The LED on the console shows the power state.

The next step is to associate the persona with the server. This can be done by making an explicit assignment from the Scalent console. On the physical view, select the server and select the ‘server assignment’ option from the drop down list on the top right. Select the “Use only for persona” option and

select the persona that was created in the previous step. Save the configuration and click on the “Start Persona” link. The server will be powered on and the Windows image will be booted over iSCSI on that server. This can be monitored by looking at the console of the server. After the persona is up and running, the state change is reflected on the Scalent console and the color of the persona will change to “green”.

### Adding an iSCSI booted Linux Persona

1. From the Scalent Console, delete the previously discovered server that was running Linux OS on the local disk.
2. Add the disk image that was migrated to the iSCSI target as a retargetable Linux Persona in the Scalent environment. This can be done either using the Scalent Console or the CLI.
3. Add an iSCSI-booted Linux Persona using the Scalent console:
  - a. Switch to the “Catalog” view in the Scalent console.
  - b. In the left side bar, select Personas and click on ‘New Persona’ seen at the bottom of the sidebar.
  - c. In the add a new persona dialog box, enter the following mandatory information
    - i. Name – Although this is optional, it is a good practice to assign a meaningful name to all the Personas in the Scalent environment.
    - ii. OS Family – Select linux
    - iii. OS Architecture – x86\_32 or x86\_64 depending on the Windows build that was installed on the source server
    - iv. Boot Type - Select iSCSI
    - v. Server Type – Enter “EqualLogic” (without quotes)
    - vi. Target IP Address – Enter the IP address of the EqualLogic IP SAN
    - vii. Target Name – This should be obtained from the EqualLogic storage’s group manager console. This will be something like iqn.2001-05.com.equallogic:0-8a0906-0cd0f7b04-0dd9fef9f324a958-sql-server
    - viii. Initiator Name – This is the name that was configured in the EqualLogic Storage that will have access to the target. This will be a string like iqn.1991-05.com.microsoft:sql-server.
    - ix. Device – This is the boot device of the new iSCSI volume, where the image was previously copied. Eg: /dev/sda1
    - x. Kernel – This must match the Kernel of the source image. Check ‘uname -r’ on the source server to determine the value.
    - xi. Uncheck the “Agent Exists” check box.

4. Alternatively, the Scalent CLI can be used to add the Persona by issuing a command similar to the one shown below.

```
add persona osFamily=linux osArch=x86_64 bootType=iscsi
image="iscsiwbi|EqualLogic|10.30.2.5|3260|iqn.2001-05.com.equallogic:0-8a0906-0cd0f7b04-
0dd9fef9f324a958-apache|iqn.1991-05.com.apache |/dev/sda1|ext3|iscsiwbi" kernel="vmlinuz-2.6.18-
53.el5" name="Lin-iSCSI-Test"
```

### Create a RAM disk and matching Kernel for iSCSI-booting

In the case of Linux, there is one more step that is required before proceeding to iSCSI boot the copied image. This is the creation of RAM disk on the source and copying it over to the Scalent controller. Execute the following steps on the source server to accomplish this.

1. Install the Persona software (present in the Linux ISO of the Scalent software distribution) on the source server.

```
#installPersona.sh
```

**Note:** The Scalent service does not need to be started

```
#service scalent stop
```

```
#chkconfig scalent off
```

2. Run the utility to build the RAM disk.

```
#/opt/scalent/bin/mkscalentrd
```

The utility looks for the sources it needs to build the most inclusive RAM disk and reports on what it discovers. Then it builds the RAM disk and reports where it stores it in the /tmp directory on the server, and offers a suggestion for how to copy the RAM disk and kernel to the appropriate directory on the Scalent Controller.

3. Copy the RAM disk and kernel to the appropriate directory on the Scalent controller.

```
# scp /tmp/initrd-2.6.9-34.ELsmp.img.gz <controller_ip>:/var/opt/scalent/tftpboot/ramdisk
```

```
# scp /boot/vmlinuz-2.6.9-34.ELsmp <controllerip>:/var/opt/scalent/tftpboot/kernel
```

The example above shows the copy command for a 32-bit image. The locations will vary for a 64-bit image. Follow the instructions presented by the mkscalentrd utility for exact commands required to copy the 64-bit RAM disk and kernel to the Controller.

### Booting the newly created iSCSI booted Linux Persona

In order to boot the Persona created in the previous step, the source server needs to be re-discovered as a bare metal server. Follow the same steps as before when the server was rediscovered and the persona was assigned to the server. The Linux persona will be assigned to the selected server in this case. Once the server powers on, the booting process can be observed from the server console.

## **Retargeting iSCSI Booted Personas**

The ability to decouple a server image from its underlying server is gained by creating a persona. The persona allows server image mobility. The process of bringing up a persona on a different server is called retargeting. Personas can be retargeted across disparate servers.<sup>1</sup> In order to retarget a running persona, execute the steps as shown below on the Scalent console:

1. Stop a running persona by clicking on the Stop Persona link from either the physical or catalog view. Every configuration change has to be confirmed by saving the configuration.
2. After a few minutes, the persona state changes and the server is powered off.
3. In the physical view, select the server that was running the persona. From the “Server” drop down list (on the top right), first select ‘server assignment’ and then select “Use for any persona or VM Rack”. After the configuration is saved, the persona is disassociated from the server.
4. Pick a different bare metal server that was discovered previously.
5. Assign the same persona that was stopped to this server. Follow the same steps as before for assigning and starting the persona.
6. After a few minutes, the new server will be up and running the same persona.

The same set of tasks can also be executed using Scalent CLI commands.

## **9.3 Phase 3**

### **9.3.1 Discover Network Switches**

In the previous sections, we have seen how Scalent can be used to discover servers and the associated server images. In the next few sections, we are going to explore how Scalent can be used to discover network switches and how these switches can be managed from within a Scalent environment. In general, a Scalent environment will have 2 classes of switches: chassis switches and top-of-the-rack switches. The top-of-the-rack switches can further be categorized as vRack switches and Interconnect switches. This configuration includes chassis switches and interconnect switches. Discussion about the vRack switches are beyond the scope of this document.

#### **Adding the Chassis switches in Read-Only mode**

From the Scalent Console’s physical view, drag a New Chassis from the resources in the left side bar. An alert prompts for entry of basic information about the chassis switches, including the IP address and authentication settings that have been configured for each switch. Select “Read-Only Switch” for the Switch Management Mode. At a later stage, we will convert the switch to be a “Fully Managed” switch. As the name implies, the Scalent controller only reads information about the switch and its ports. Enter the SNMP Community Name to add a switch in “Read-Only” mode. When the switch is in “Read-Only” mode, the switch ports have to be manually configured to carry the appropriate VLANs for the role each port plays in the environment.

---

<sup>1</sup> iSCSI booted Windows personas can only be retargeted across like servers

The alert also prompts for the entry of the authentication settings that were configured for the chassis' management module.

Please enter the following information in order to add a new chassis to the system.

Chassis Type

Management Mode

Management IP Address

SNMP Community Name

Confirm SNMP Community Name

Username

Password

Confirm Password

Management Module

Management IP Address

Username

Password

Confirm Password

First Switch

Additional Switches

Channel	Management IP Address
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>

Figure 6 – Adding a new chassis to the Dell AIM console

After the configuration is saved, the blade servers that were discovered previously and shown in the default container will automatically “jump” into the blade chassis that was added . At this stage the Scalent controller makes the association between the blade servers and their position in the physical network topology.

**Switch Port Roles**

A switch port in the Scalent environment can have one of the following roles. These roles can be viewed by pointing to a chassis switch in the physical view and selecting “Port Role Settings” from the “Chassis Switch” drop-down list on the top right corner in the physical view of Scalent console.

Server (S) ports connect a chassis to a blade server. Server ports are not shown in the Console Physical Network and other views to reduce clutter. They can be viewed in each switch’s Port Role Settings sidebar. On chassis switches, the server ports are fixed (they can be only server or unmanaged ports), because they are physically connected to the backplane into which the chassis blade servers plug. Server ports are members of the System Control Network (SCN)

External (E) ports connect to networks outside the Scalent environment. Any switch port that is not configured as a server or unmanaged port is configured as an external port by default. External ports are displayed at the side of chassis in the Console Physical and other views, where cables can be connected or disconnected from other switch ports or from external networks. If an external port is cabled on one

switch to the external port on another switch in the Scalent environment, both ports are changed to interconnect ports.

Interconnect(I) ports connect managed switches within the Scalent environment. An interconnect port is created in the Console Physical Network view by cabling an external port on one switch in the Scalent environment to an external port on another switch in the environment, which converts both ports to interconnect ports. If the cable is disconnected from an interconnect port, it becomes an external port again.

*Interconnect ports are members of all VLANs; their native VLAN is the sink VLAN, which is VLAN 4005 (unless a different sink VLAN is specified during Controller installation).*

Unmanaged (U) ports are ports that Scalent does not manage; they must be configured with the switch's own management tools. Unmanaged ports are not shown in the workspace. A cable or server must be deleted if it is connected to a port before a switch port's port role can be changed to unmanaged.

### **9.3.2 Converting the 'Read-Only' switches to be 'Fully Managed'**

When a switch management mode is "Fully Managed", Scalent Controller will be able to make changes to the switch (using a combination of SNMP and telnet). To convert a previously discovered switch, point to the switch from the physical view of the Scalent console and select "Management Settings" from the Chassis Switch dropdown list. Change the Management Mode field to "Fully Managed" and enter the switch username and password when prompted. Once this information is entered and the configuration is saved, Scalent controller gains the ability to make changes to the switch configuration. This enables automated network configuration which is one of the pre-requisites for deploying a virtual ready infrastructure.

### **9.3.3 Creating Basic Network Topologies**

In this section, we will explore some basic capabilities that are enabled by having the Scalent controller manage the network switches in the environment. We will also be introducing new Scalent terms as we go along. Virtual networks in a Scalent-managed environment behave as conventional networks in a data center, except that they're built out of virtual LANs (VLANs) running over the physical NICs, switches, and cables, instead of those physical elements themselves. Because they are virtual, Scalent virtual networks are easier to create and manage than conventional networks. In order to build network topologies in a Scalent environment, switch to the 'Virtual' view in the Scalent console.

vSwitch - A vSwitch(Virtual Switch) is a virtual equivalent of a physical switch. In other words, a vSwitch is the Scalent term used to refer to a VLAN. vSwitches can be added by simple drag-and-drop operations in the Scalent console. Drag a vSwitch icon from the left sidebar on to the workspace. On the right sidebar, the VLAN ID can be entered if desired. If left blank, Scalent controller will assign the next available VLAN ID from the block of VLANs in the "Controller-Assigned" VLAN range. Once the configuration is saved, a new VLAN is created in the environment. This can be viewed by logging into the chassis switches and accessing the switch configuration.

vNIC – A vNIC is a Virtual NIC (Network Interface Card). The vNICs are configured on the Persona and in the most basic configuration each vNIC will correspond to each physical NIC on the server that is running the Persona. A vNIC can be created by dragging a vNIC on to a Persona in the 'Virtual'

view of the Scalent console. After the vNIC is created (or while creating a vNIC), the IP address for the interface can also be assigned. Scalent provides multiple options to set IP addresses.

A simple network topology can be created by executing the following steps:

1. Choose two running Personas.
2. Drag a vNIC to each of the Personas and assign them static IP addresses in the same subnet. For example, 192.168.0.10 and 192.168.0.20
3. Drag a vSwitch to the workspace.
4. Draw cables from both the vNICs to the vSwitch.
5. Save the configuration.

The topology can be tested in multiple ways. First, ping both IP address from each server. Next, check the switch configuration and examine the VLAN membership of the switch ports connected to the servers that are running the Personas. The VLAN ID of the vSwitch that was created on both the ports will be visible. Deleting one of the cables and saving the configuration will automatically update the switch configuration.

## 9.4 Phase 4

### 9.4.1 Scalent Agent Installation

The Dell AIM solution includes an optional agent component that can be installed on the Personas to carry out some advanced functions. These functions include Advanced Health Monitoring, OS side Virtual Networking, and Workload Configuration and Management. In this section, we will see how the agent software can be installed and how the various functions provided by the agent can be enabled.

#### Installing Scalent agent software for Windows

Log into the server running the persona as the Administrator. Mount the Scalent software distribution for Windows and run the executable – setup\_persona.exe present in the root directory. When asked to confirm the Scalent agent software install on the server, click OK.

If a sufficiently recent version of the driver is not available, the installer will prompt for an updated NIC driver. The setup application installs the Scalent agent software in the C:\Scalent\persona directory. It then configures and starts the persona agent service. If firewall software is configured on the server, the persona disables it: The persona can't operate correctly with a local firewall running. If any services (such as Apache or IIS) are installed on the persona, they should be configured to start *after* the persona agent service.

#### Installing Scalent agent software for Linux

Log in to the server running the persona as the root user. Mount the Scalent software distribution for Linux and execute the following command from the root directory of the distribution. The script installs the agent software.

```
# ./installPersona.sh
```

When the installation is complete, start the Scalent service using the following command:

```
# service scalent start
```

### **9.4.2 Persona Network Management**

The Scalent agent provides OS side virtual networking capabilities. In the previous section, we briefly discussed how vNICs (Virtual NICs) can be created on Personas. Without the agent, the number of vNICs is limited by the number of physical NICs in the server. Installing the agent overcomes the limitation imposed by the number of physical NICs. Multiple vNICs can be instantiated on a single physical NIC.<sup>2</sup>

In order to create virtual NICs, ensure that the “Networking Enabled” persona property under the persona drop-down list is checked. This can be done either from the catalog or virtual view of the Scalent console. Persona properties can be changed only when the persona is in a dormant state. If the persona is running, it needs to be stopped first. However, vNICs can be added to personas, regardless of whether they are running or dormant.

### **9.4.3 Workload Management**

The Scalent agent also provides workload configuration and management capabilities. The workload configuration function enables defining the startup and shutdown of application related services. These are referred to as persona extensions. The extensions are external scripts that are automatically executed when there is a transition in the persona state. Let us take an example of a persona running a web server workload. During the process of booting up, a persona extension can be executed to automatically update the load balancer routing table. Details regarding workload management capabilities and persona extensions are beyond the scope of this document.

---

<sup>2</sup> This feature of creating Virtual NICs is currently not available for Windows 2008 personas

## 9.5 Troubleshooting

### 9.5.1 Best Practices

**Note:** Dell FlexAddress should not be used in conjunction with the Dell AIM solution.

Please refer to the “Operations Topics” book within the documentation set for additional best practices.

### 9.5.2 Known Issues

#### Link Aggregation with PowerConnect 6224

**Description:** When using Dell PowerConnect 6224 switches running firmware version 2.2.0.3 with Dell AIM 3.1, there is a possibility of VLAN information on that switch being lost when a specific set of conditions are met.

**Reproduction:**

- Link Aggregation enabled
- Port Channel not cabled in Scalent Console (port is shutdown)
- Switch is reloaded

**Solution:** This issue can be corrected by recreating the Link Aggregation on the appropriate ports on the switch.

**Avoidance:** This issue can be avoided by making sure all necessary ports are configured and cabled in the Scalent console at all times.

### 9.5.3 Basic Troubleshooting

The Scalent support portal includes an extensive knowledgebase which can be accessed for a variety of commonly encountered issues as well as troubleshooting tips.

## 10 References

### **Integrating Blade Solutions with EqualLogic SANs**

[http://www.dell.com/downloads/global/partnerdirect/apj/Integrating\\_Blades\\_to\\_EqualLogic\\_SAN.pdf](http://www.dell.com/downloads/global/partnerdirect/apj/Integrating_Blades_to_EqualLogic_SAN.pdf)

### **Deploying Pools and Tiered Storage in a PS Series SAN**

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5239>

### **Deploying Thin Provisioning in a PS Series SAN**

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5245>

### **Network Connection and Performance Guidelines**

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5311>

### **Network Connection Guidelines**

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5331>

### **PS Series Array Network Performance Guidelines**

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229>

### **Dell PowerConnect Switches**

<http://www.dell.com/powerconnect>

### **Dell PowerConnect M6220 Product Page**

[http://www.dell.com/us/en/enterprise/networking/pwcnt\\_6220/pd.aspx?refid=pwcnt\\_6220&s=biz&cs=555](http://www.dell.com/us/en/enterprise/networking/pwcnt_6220/pd.aspx?refid=pwcnt_6220&s=biz&cs=555)

### **Dell PowerConnect 6224 Product Page**

[http://www.dell.com/us/en/enterprise/networking/pwcnt\\_6224/pd.aspx?refid=pwcnt\\_6224&s=biz&cs=555](http://www.dell.com/us/en/enterprise/networking/pwcnt_6224/pd.aspx?refid=pwcnt_6224&s=biz&cs=555)

### **Dell PowerConnect 62xx Command Line Interface Guide**

[http://support.dell.com/support/edocs/network/pc62xx/en/CLI/PDF/cli\\_en.zip](http://support.dell.com/support/edocs/network/pc62xx/en/CLI/PDF/cli_en.zip)

### **Dell PowerConnect 62xx Command Line Interface Addendum**

[http://support.dell.com/support/edocs/network/pc62xx/en/CLI\\_Add/cli\\_add.zip](http://support.dell.com/support/edocs/network/pc62xx/en/CLI_Add/cli_add.zip)

### **Dell PowerConnect 62xx User's Guide**

[http://support.dell.com/support/edocs/network/pc62xx/en/UG/PDF/ug\\_en.zip](http://support.dell.com/support/edocs/network/pc62xx/en/UG/PDF/ug_en.zip)

### **Dell PowerConnect 62xx User's Guide Addendum**

[http://support.dell.com/support/edocs/network/pc62xx/en/UG\\_Add/PDF/ug\\_add.zip](http://support.dell.com/support/edocs/network/pc62xx/en/UG_Add/PDF/ug_add.zip)

### **Scalent Support Portal**

<http://www.scalent.com/support>

### **Scalent Documentation**

Can be found on media shipped with solution or at <http://www.scalent.com/support>

# 11 Appendix

## A Scalent Console Views

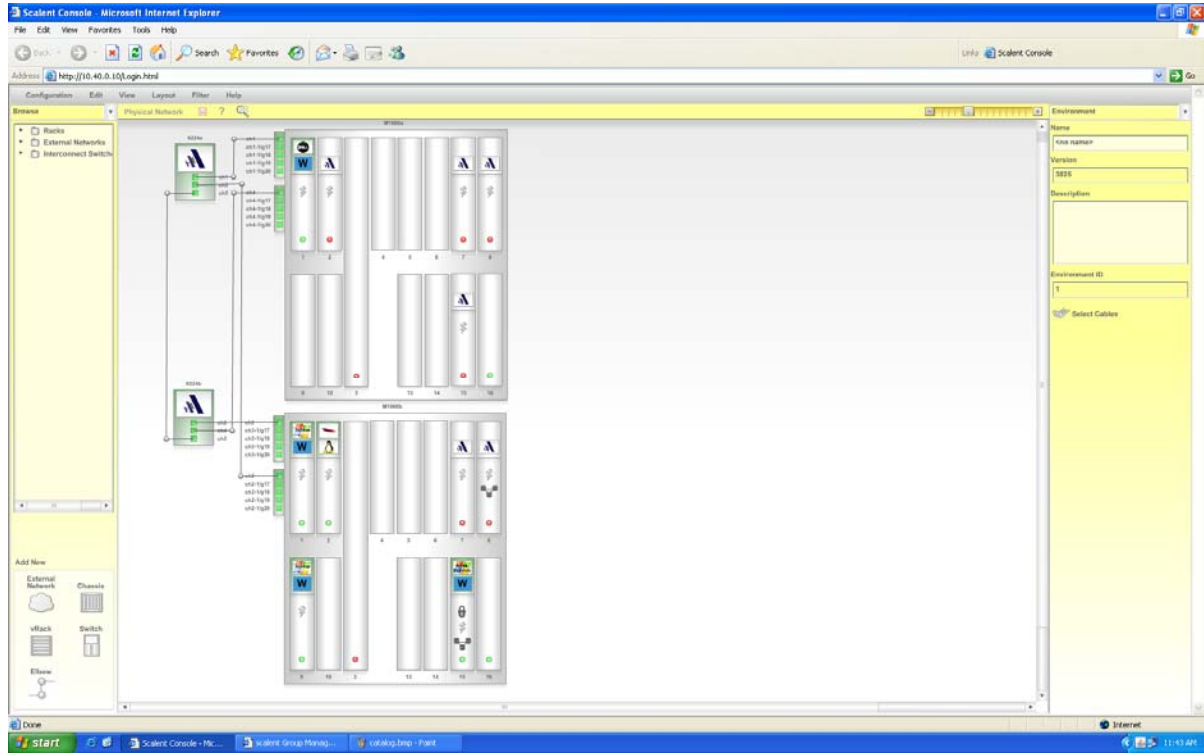


Figure 7 – Physical View

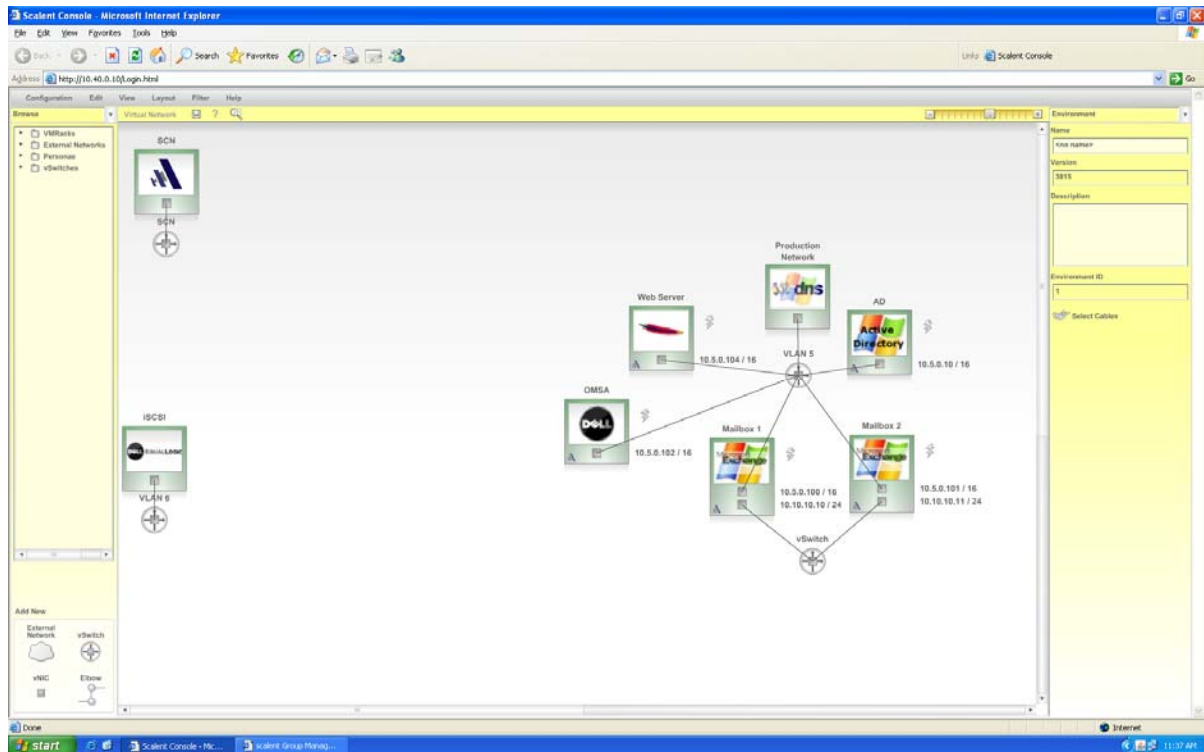


Figure 8 – Virtual View

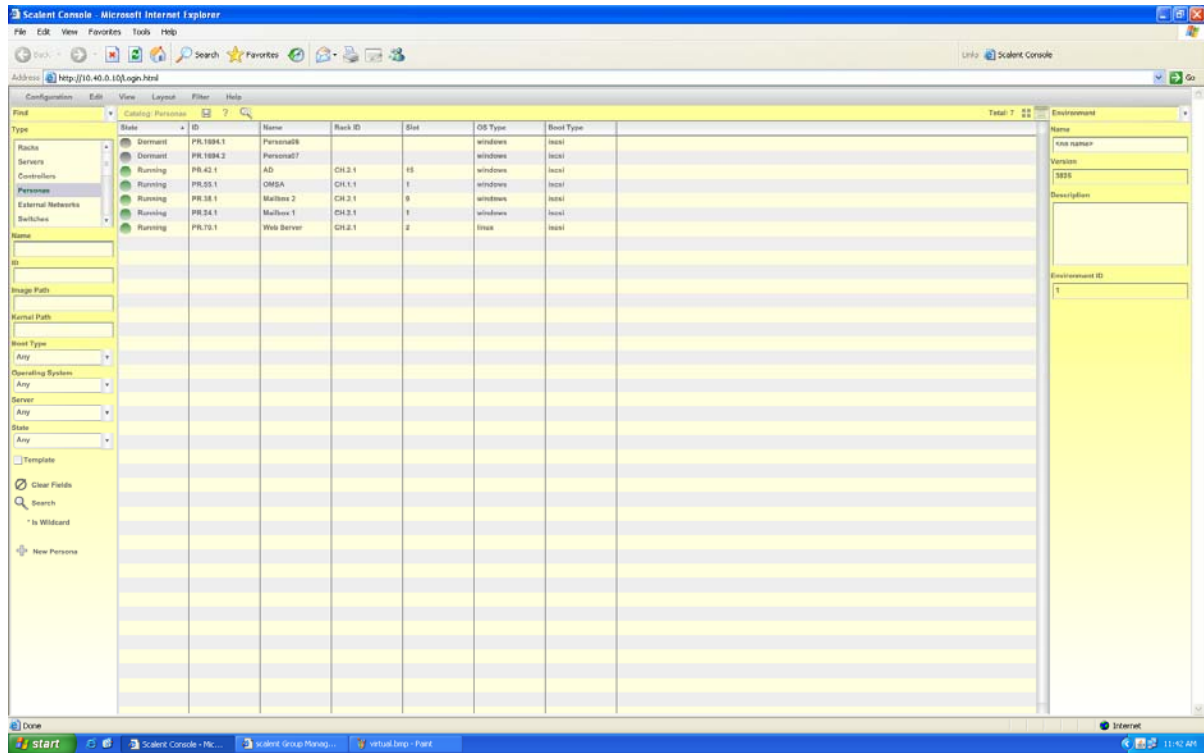


Figure 9 – Catalog View

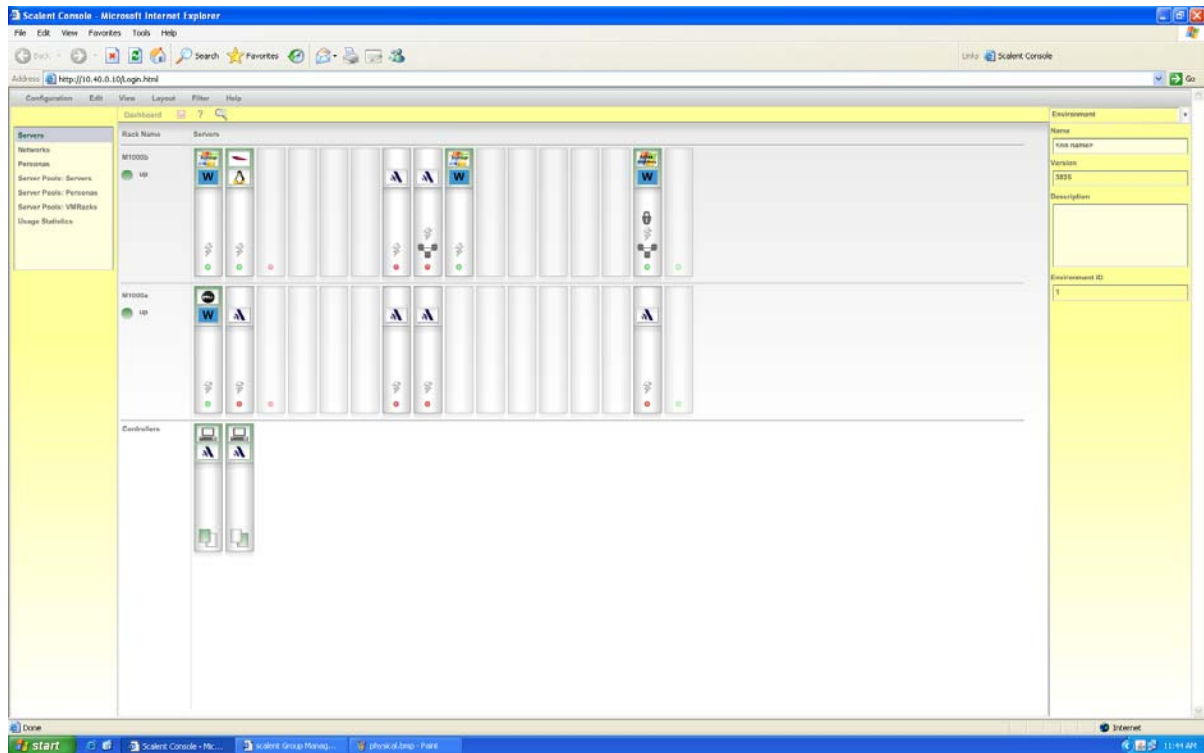


Figure 10 – Dashboard View

## B Configuring Dell PowerConnect Ethernet switches

The PowerConnect M6220 switch has 20 gigabit Ethernet ports, g1 through g20. Ports g1 through g16 connect to blade servers inserted in the chassis' 16 blade server slots. Port g17 through g20 correspond to the switch's 4 external ports, for connecting to other switches, or to devices outside the Scalent environment. By default, Scalent reserves the first external port (g17) for managing the switch. In this configuration, we will configure ports g17 through g20 in an aggregated link for improved bandwidth and resiliency. Follow the steps below to configure the chassis switch module.

The PowerConnect 6224 switch needs to be configured similarly, but the ports are not assigned to a particular function. The ports connecting the switches together need a similar aggregated link configured, and the SCN (VLAN 4004) needs to be added to all interconnect ports.

1. Disconnect all cables from the switch's external ports
2. Establish a telnet connection to the CMC and log in as the 'root' user.
3. Connect to the switch module from the CMC. In the following command, replace x with the number of the switch to be connected to:

```
$ connect switch-x
connect: acquiring remote port.
Connected to remote port.
Escape character is '^\'.
```

The switch presents its prompt, typically `console>`

4. Check that a supported firmware version is running. If not, install a supported version of firmware, as described in the documentation that came with the chassis.
5. Enter 'enable' mode to change switch configuration.

```
console> en
console#
```

6. Reset switch to factory default settings.

```
console# delete startup-config
Delete startup-config? [Y/N](N): y
Startup file was deleted
console# reload
You haven't saved your changes. Are you sure you want to continue
(y/n) [n] ? y
This command will reset the whole system and disconnect your
current
session. Do you want to continue (y/n) [n] ? y
```

7. Reconnect to the switch and when prompted, do not run the setup wizard.
8. Enter Configuration mode.

```
console> en
console#
```

```
console# configure
console(config)#
```

9. Add an account for Scalent controller to manage the switch.

```
console(config)# username username password password level 15
```

10. Configure the SCN (System Control Network) on the switch.

```
console(config)# vlan database
console(config-vlan)# vlan 4004
console(config-vlan)# exit
console(config)# interface vlan 4004
console(config-if-vlan4004)# name "SystemControlNetwork"
console(config-if-vlan4004)# exit
```

11. Configure VLAN mode for all interfaces.

```
console(config)# interface range ethernet all
console(config-if)# switchport mode general
console(config-if)# exit
```

12. Configure an IP address for the switch. For a very simple configuration, where the Controller is connected to a single managed switch and all devices are on the same network, the values could be those of the System Control Network (SCN), which is VLAN 4004 by default.

```
console(config)# ip address ipAddress netmask
console(config)# ip default-gateway defaultGateway
console(config)# vlan database
console(config-vlan)# vlan vlanx
console(config-vlan)# exit
console(config)# ip address vlan vlanx
```

13. Configure how the switch connects to the other networks in the data center. The following example commands configure switch ports 1/g17 through 1/g20 to carry traffic for two networks, the SCN (VLAN 4004) and a single data center network (vlanx). A similar aggregated link will need to be configured on the switch connected to these ports.

```
console(config)# interface range ethernet 1/g17-1/g20
console(config-if)# channel-group 1 mode auto
console(config-if)# exit
console(config)# interface port-channel 1
console(config-if-ch1)# switchport mode general
console(config-if-ch1)# switchport general pvid vlan vlanx
console(config-if-ch1)# switchport general allowed vlan add 4004 untagged
```

```
console(config-if-ch1)# switchport general allowed vlan add vlanx untagged
console(config-if-ch1)# exit
```

14. Configure the port on the switch that corresponds to the chassis slot the controller blade server will connect to. In the following commands, replace controllerPort with the switch port that Controller will connect to.

```
console(config)# interface ethernet controllerPort
console(config-if-port)# switchport general pvid 4004
console(config-if-port)# switchport general allowed vlan add 4004
console(config-if-port)# exit
console(config)# exit
```

15. Save the switch configuration.

```
console# copy running-config startup-config
```

16. Restart the switch.

```
console# reload
```

Connect the switch's external ports (g17-g20) to the data center network.