Systems Manageability of VMware® ESXi on Dell™ PowerEdge™ Servers

A Dell | VMware Technical White Paper

Viswanathan Balakrishnan

Balasubramanian Chandrasekaran

Sudhir Shetty

Alok Sinha

Charu Chaubal

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.

Table of Contents

Introduction

This article provides a high-level overview of the manageability aspects of the VMware ESXi solution on Dell PowerEdge Servers. It discusses the various facets of manageability and the management tools that an IT Administrator can use to deploy, monitor, inventory, configure, and update an ESXi system.

Virtualization has entered the mainstream and is now critical to building cost-effective, highly-available, and scalable enterprise IT infrastructures, enabling both flexible resource management and automated resource allocation based on strategic enterprise policies. Virtualized infrastructures help simplify IT operations in many ways, from helping to shield software from hardware, to enabling secure resource sharing, to facilitating software deployment and relocation. They also increase business agility by:

- Enabling IT staff to dynamically reallocate resources as needed to avoid planned downtime
- Enhancing the efficiency of application testing and development
- Facilitating rapid, cost-effective disaster recovery

VMware vSphere is the next evolution along this path of innovation. By delivering on the promise of efficiency, control, and choice, vSphere transforms IT infrastructures into self-managed, dynamically-optimized clouds. This giant step forward brings us closer to the vision of computing as a ubiquitously-available, easily-accessible, and reliable utility service—similar to telephone and electrical service. Utilizing a cloud infrastructure built on vSphere, business owners and application owners can deploy new business services under managed service-level agreements (SLAs) without understanding the intricacies of server, storage, and network resources. Well-integrated management capabilities allow IT to control the overall quality of service delivery and to manage external resources as easily as they manage internal resources.

VMware ESXi is the next-generation hypervisor and a core building block for the vSphere platform. This innovative architecture operates independently from general-purpose operating systems, offering improved security, increased reliability, and simplified management. Its compact architecture is designed to be directly integrated into virtualization-optimized and certified server hardware, enabling rapid installation, configuration, and deployment.

Functionally, ESXi is equivalent to VMware ESX; however, the Linux-based service console has been removed, dramatically reducing the footprint to improve security and maintenance. The functionally of the service console is replaced by remote command-line interfaces and adherence to system management standards. In the simplest implementation, ESXi is embedded directly into the firmware of select server models from Dell, allowing the server to boot directly into ESXi.

Because ESXi no longer includes a service console, many of the management activities performed on the ESX platform—for example, configuring user access to the service console and administering management agents running in it—are no longer necessary. Other management tasks previously performed in the service console are now performed in one of several ways:

- Using the vSphere Client to provide a Windows-based graphical user interface for interactive configuration of the platform. The vSphere Client has been enhanced to provide capabilities that were previously available only in the service console.
- Using the vSphere Command Line Interface (vCLI), an interface that enables scripting and command-line-based configuration of the platform from a Linux or Microsoft® Windows® based server, via an encrypted and authenticated communication channel.

- Using tools such as Dell OpenManage™ Server Administrator (OMSA) and Dell Management Console that leverage standard APIs such as WS-MAN to manage, inventory, and monitor ESXi server hardware.

In addition, you can manage ESXi using vCenter, just as you would any ESX system. Distributed virtualization features—such as VMotion and VMware DRS—are designed to work exactly the same on ESXi. This also applies to mixed environments that include ESX and ESXi systems. VCenter presents both types of systems in the vSphere Client user interface—certain features unique to ESXi management appear for hosts equipped with that version.

This article focuses primarily on the hardware management aspects of the Dell/VMware solution.

Architecture Overview

Figure 1 demonstrates the architecture of the Dell/VMware solution ecosystem. The ESXi architecture installed on a PowerEdge server comprises the core virtualization layer, which includes a purpose-built kernel and the processes that run on top of it. The kernel abstracts the underlying hardware from the virtual machines and controls application resource allocation. Some of the processes that run on the kernel are:

- **Direct Console User Interface (DCUI)** – A basic management interface accessed via the local console that provides functionality such as setting the administrative password, configuring networking, and accessing troubleshooting tasks such as viewing logs and restarting management agents.
- **Virtual Machine Monitor (VMM)** – Provides the execution environment for a virtual machine, as well as a helper process known as VMX.
- **Various Agents** – Enables remote manageability from VMware tools such as vSphere Client and VCenter.
- **Management Infrastructure** – VMware bundles components such as the SFCB (Small Footprint CIM Broker) CIMOM and Openwsman, providing the infrastructure for OEM extensions via providers and external accessibility via the WS-MAN protocol.

Dell software components include:

- **OMSA Web Server** is installed on a management station in the enterprise. A single instance of the OMSA Web Server can be used to manage multiple ESXi systems on the network.
- **Dell Provider** plugs into the SFCB CIMOM. It is a lightweight component that interprets requests from the OMSA Web server component and invokes the correct functionality in the Instrumentation stack below it. It also authorizes incoming requests.
- **OMSA Instrumentation** provides the core set of hardware monitoring and management services. It is similar to the services present in other operating system environments that are ported to the ESXi environment. It enables the hardware manageability functions around monitoring, asset inventory, and configuration.
- **Baseboard Management Utilities (BMU)** provides out-of-band (OOB) remote access to the service processor associated with the server. It enables the administrator to run remote IPMI commands against the server to retrieve and clear embedded hardware logs, perform power control operations, identify the chassis, and display sensor and power monitoring information. These utilities are included in the OpenManage media for the management station tools and utilities.

- **Dell Remote Access Controller (DRAC) Console** is a GUI that provides the ability to view system health, sensor information, and embedded hardware logs. Advanced functionality, such as console redirection and virtual media, are also available via this interface.

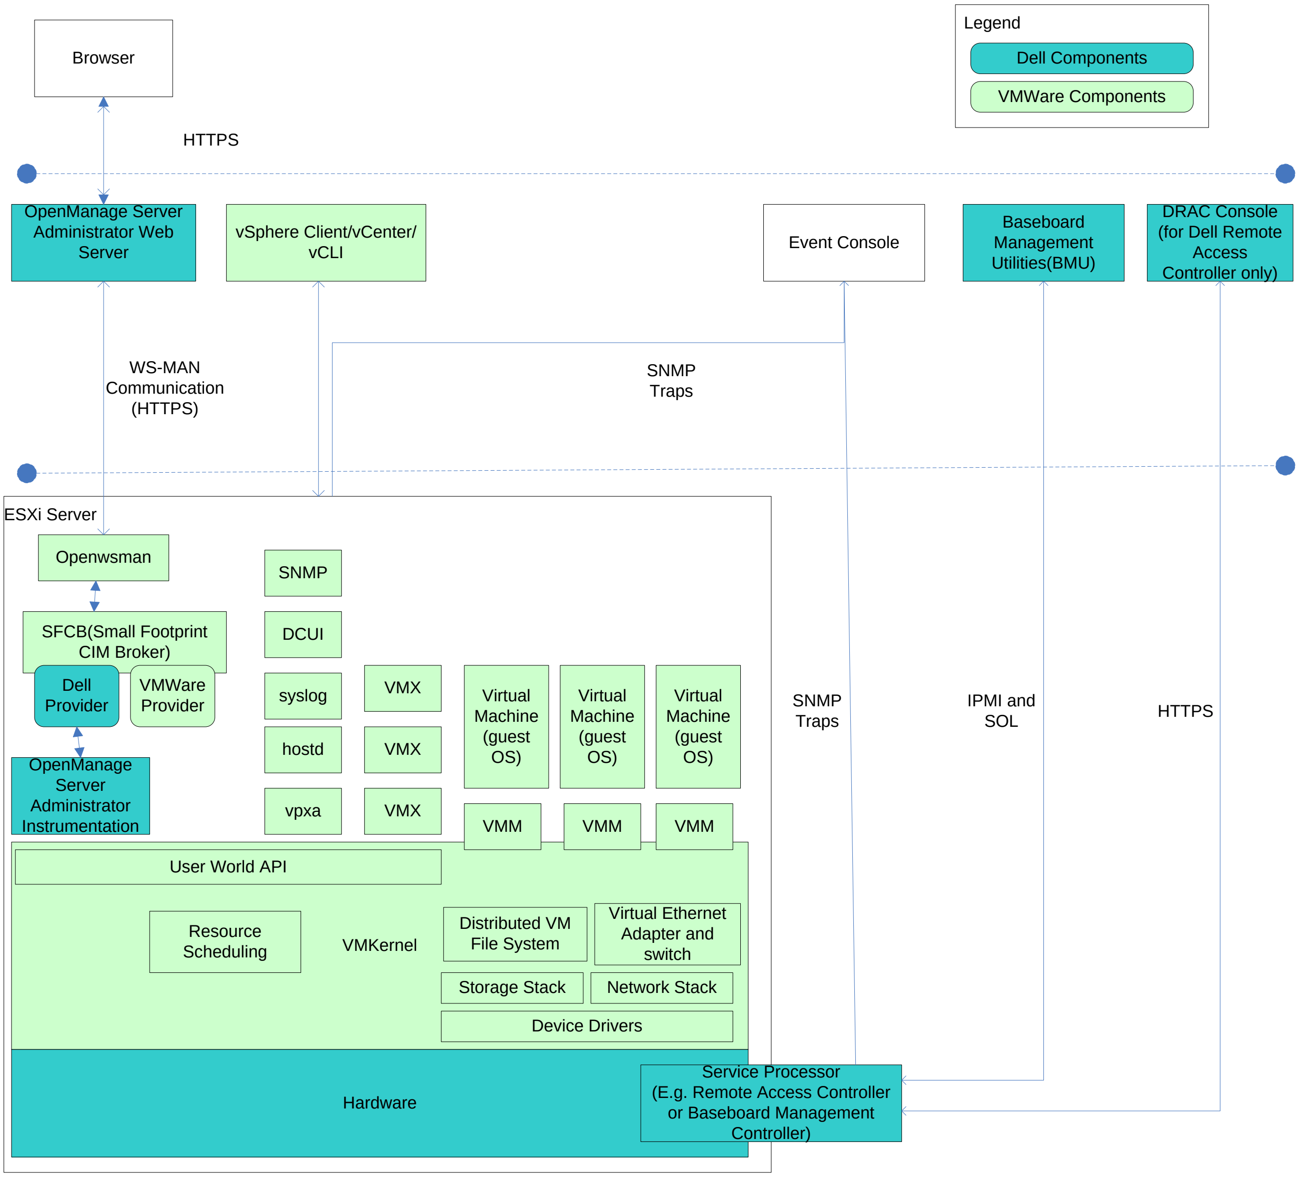
Legend
Dell Components
VMWare Components
Browser
HTTPS
OpenManage Server Administrator Web Server
vSphere Client/vCenter/ vCLI
Event Console
Baseboard Management Utilities(BMU)
DRAC Console (for Dell Remote Access Controller only)
WS-MAN Communication (HTTPS)
SNMP Traps
ESXi Server
Openwsman
SNMP
DCUI
SFCB(Small Footprint CIM Broker)
Dell Provider
VMWare Provider
syslog
hostd
vpxa
VMX
VMX
VMX
Virtual Machine (guest OS)
Virtual Machine (guest OS)
Virtual Machine (guest OS)
OpenManage Server Administrator Instrumentation
VMM
VMM
VMM
User World API
Resource Scheduling
VMKernel
Distributed VM File System
Virtual Ethernet Adapter and switch
Storage Stack
Network Stack
Device Drivers
Hardware
Service Processor (E.g. Remote Access Controller or Baseboard Management Controller)
SNMP Traps
IPMI and SOL
HTTPS

Figure 1: **Overall Software Architecture**

In Figure 1, the event console provides trap reception capabilities, event formatting and display, and the ability to filter and perform actions in response to events. OpenManage IT Assistant and Dell Management Console are examples of event consoles that provide this administrative capability. They are capable of handling Platform Event Traps (PETs) that are generated by the service processor and VMware-issued traps (for example, VM-related events, VM power on and off). They also support SNMP traps issued by OpenManage Server Administrator instrumentation. ESXi can be configured to send SNMP traps to the machine hosting the Event Console using the vSphere Command Line Interface (vCLI).

Systems Manageability

Deployment

You can order PowerEdge servers with the ESXi image pre-installed on internal storage (SD card, USB, or hard drive). This eliminates the need to install a virtualization platform. The ESXi hypervisor automatically starts during boot up, and can acquire network settings via Dynamic Host Configuration Protocol (DHCP). If DHCP is not configured, an administrator can use Direct Console User Interface (DCUI) for basic configuration, making it easier to quickly bring the server online.

Dell software components (OpenManage Server Instrumentation, Dell Provider) are included as part of the OEM-customized ESXi image resident on the boot bank partition of the server's internal storage.

By default, Dell hardware manageability support is disabled. Users can access this functionality by enabling the CIM OEM providers—either via vSphere Client or vCLI—and then rebooting the box. The *Dell OpenManage Server Administrator Installation Guide for VMware ESXi* (available at support.dell.com) contains additional information on enabling this configuration.

Monitoring

The ESXi user interface is very similar to the existing OMSA solution, providing a familiar interface for administrators who have used it to manage ESX or Windows and Linux systems. ESXi enables the user to:

- View server and storage asset data
- View server and storage health information
- View alert and command logs
- Configure hardware (storage, BIOS, etc.)

The user can point a Web browser at the Web server to see the HTML-based user interface. On the login page, an administrator can enter the user credentials and the IP address or host name associated with the ESXi system to initiate remote management of the ESXi system. The software uses the WS-MAN channel to communicate with the ESXi system

The user credentials correspond to local users that have been configured for the ESXi system via vSphere Client. The channel is secured using HTTPS. Dell recommends that the user follow security guidelines around certificate management, as specified in the *ESXi Server Configuration Guide* available at www.vmware.com/support/pubs.

OMSA instrumentation services monitor hardware health. The user interface displays the global status of the main system chassis and the storage sub-system; it also provides the ability to drill down to faulty components. As shown in Figure 2, the user interface shows the health of individual server sub-components (batteries, fans, hardware logs, intrusion, memory, power management, power supplies, and temperature). You can also view storage sub-system health (controllers, batteries, physical disks, virtual disks, connectors, enclosures, enclosure sub-components, Enclosure Management Modules (EMMs), fans, power supplies, temperature sensors, and physical disks) and view details on each of the sub-components via a similar drill-down view.

Administrators can also view key power monitoring data for power consumption, cumulative energy consumed, and system peak power, as shown in Figure 3.

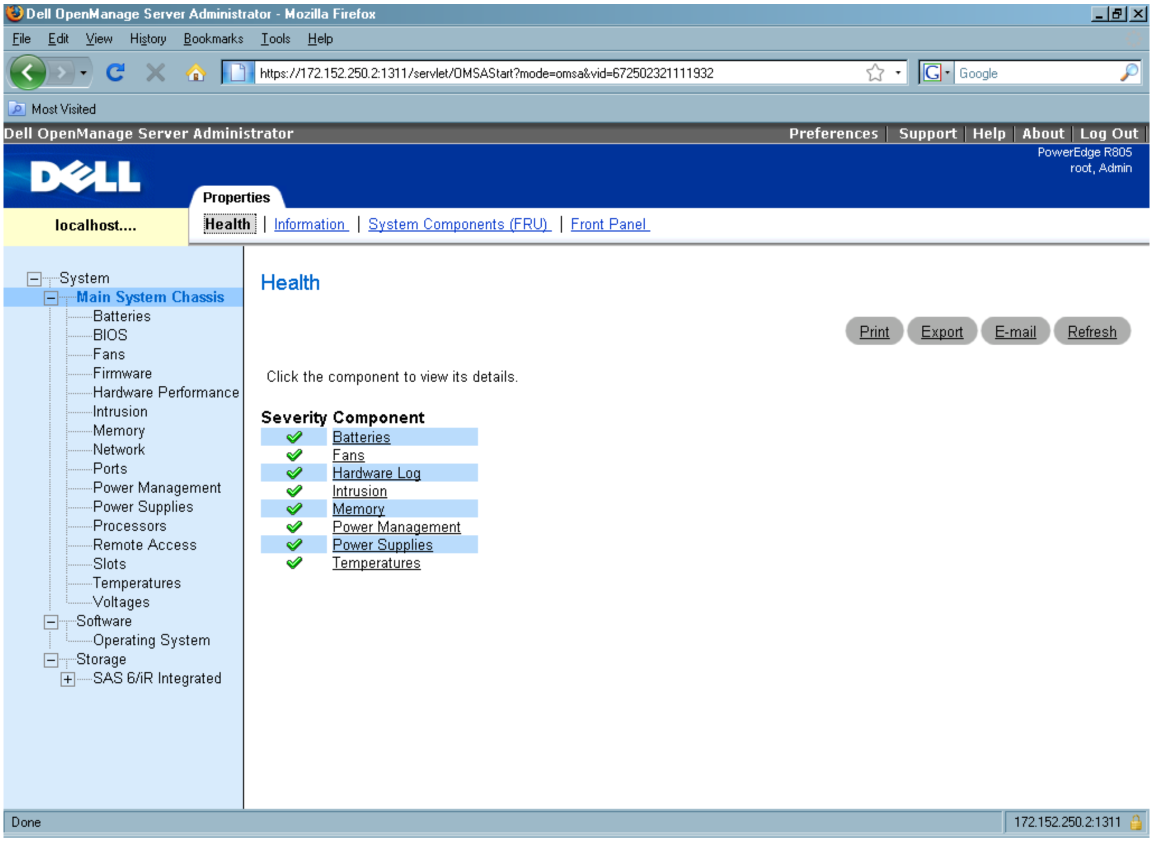

Health

Click the component to view its details.

Severity	Component
✓	Batteries
✓	Fans
✓	Hardware Log
✓	Intrusion
✓	Memory
✓	Power Management
✓	Power Supplies
✓	Temperatures

Figure 2: Health Tab

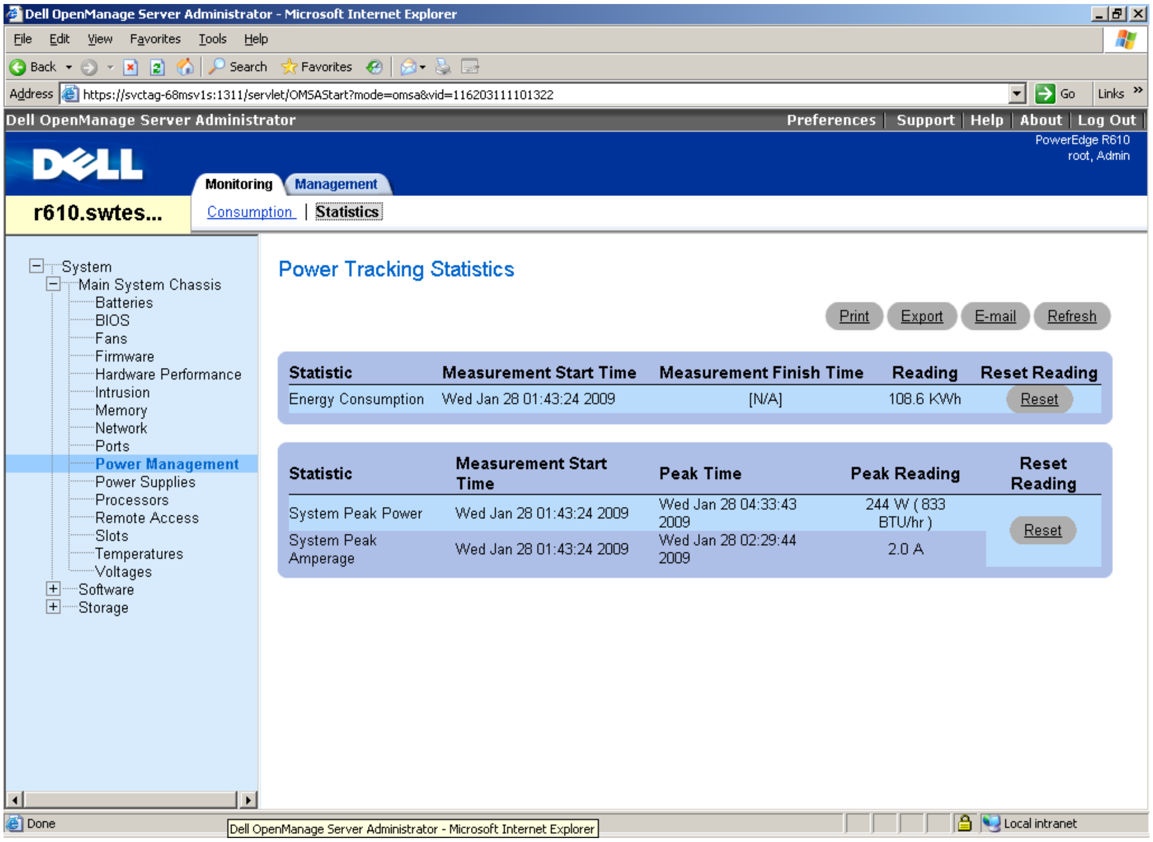

Power Tracking Statistics

Statistic	Measurement Start Time	Measurement Finish Time	Reading	Reset Reading
Energy Consumption	Wed Jan 28 01:43:24 2009	[N/A]	108.6 KWh	Reset

Statistic	Measurement Start Time	Peak Time	Peak Reading	Reset Reading
System Peak Power	Wed Jan 28 01:43:24 2009	Wed Jan 28 04:33:43 2009	244 W (833 BTU/hr)	Reset
System Peak Amperage	Wed Jan 28 01:43:24 2009	Wed Jan 28 02:29:44 2009	2.0 A	

Figure 3: Power Management – Power Tracking Statistics Page

Additionally, OMSA maintains the following logs to provide the user with troubleshooting and diagnostics information:

- **Hardware Log** - Provides historical information about hardware events that are triggered on the system. This data is also viewable through the DRAC user interface and through BMU tools.
- **Alert Log** - Lists events generated by OMSA, including alerts that correspond to the storage sub-system, as shown in Figure 4.
- **Command Log** – Provides an audit trail of each command executed on the server.

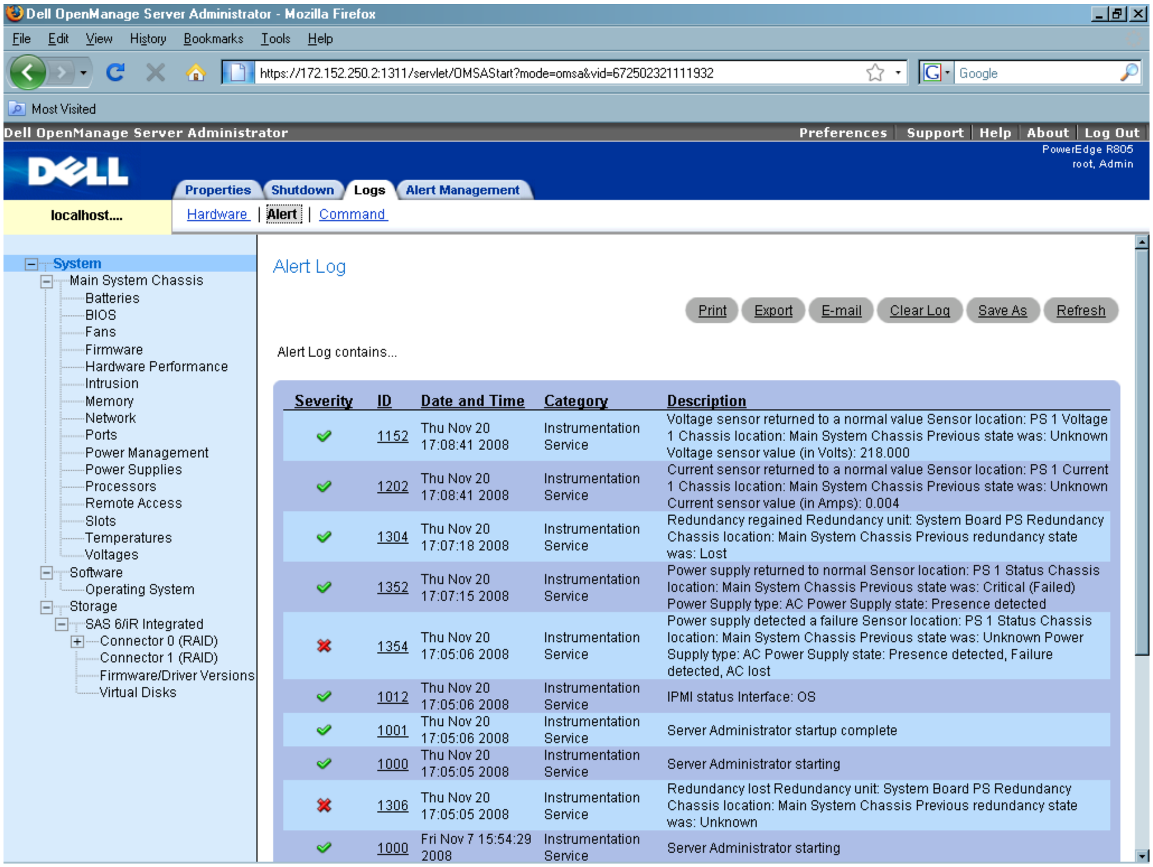

Alert Log

Alert Log contains...

Severity	ID	Date and Time	Category	Description
✓	1152	Thu Nov 20 17:08:41 2008	Instrumentation Service	Voltage sensor returned to a normal value Sensor location: PS 1 Voltage 1 Chassis location: Main System Chassis Previous state was: Unknown Voltage sensor value (in Volts): 218.000
✓	1202	Thu Nov 20 17:08:41 2008	Instrumentation Service	Current sensor returned to a normal value Sensor location: PS 1 Current 1 Chassis location: Main System Chassis Previous state was: Unknown Current sensor value (in Amps): 0.004
✓	1304	Thu Nov 20 17:07:18 2008	Instrumentation Service	Redundancy regained Redundancy unit: System Board PS Redundancy Chassis location: Main System Chassis Previous redundancy state was: Lost
✓	1352	Thu Nov 20 17:07:15 2008	Instrumentation Service	Power supply returned to normal Sensor location: PS 1 Status Chassis location: Main System Chassis Previous state was: Critical (Failed) Power Supply type: AC Power Supply state: Presence detected
✗	1354	Thu Nov 20 17:05:06 2008	Instrumentation Service	Power supply detected a failure Sensor location: PS 1 Status Chassis location: Main System Chassis Previous state was: Unknown Power Supply type: AC Power Supply state: Presence detected, Failure detected, AC lost
✓	1012	Thu Nov 20 17:05:06 2008	Instrumentation Service	IPMI status Interface: OS
✓	1001	Thu Nov 20 17:05:06 2008	Instrumentation Service	Server Administrator startup complete
✓	1000	Thu Nov 20 17:05:05 2008	Instrumentation Service	Server Administrator starting
✗	1306	Thu Nov 20 17:05:05 2008	Instrumentation Service	Redundancy lost Redundancy unit: System Board PS Redundancy Chassis location: Main System Chassis Previous redundancy state was: Unknown
✓	1000	Fri Nov 7 15:54:29 2008	Instrumentation Service	Server Administrator starting

Figure 4: Alert Log

Dell recommends configuring your system to send PET traps and VMware SNMP traps from the ESXi system to an event console in the enterprise, thus notifying you of hardware events that require attention. Typically, event consoles can send notifications via e-mail or SMS, enabling you to log in via the OMSA Web Server or DRAC user interfaces to perform additional troubleshooting and diagnostic functions.

You can also view server hardware health in vSphere Client. vCenter 4 obtains health information from the same source as OMSA. Figure 5 shows what the hardware management tab looks like in vSphere Client. In addition to viewing health status, vCenter allows you to set alarms that alert you to changes in hardware health—several of these alerts are enabled by default when you install vCenter 4.

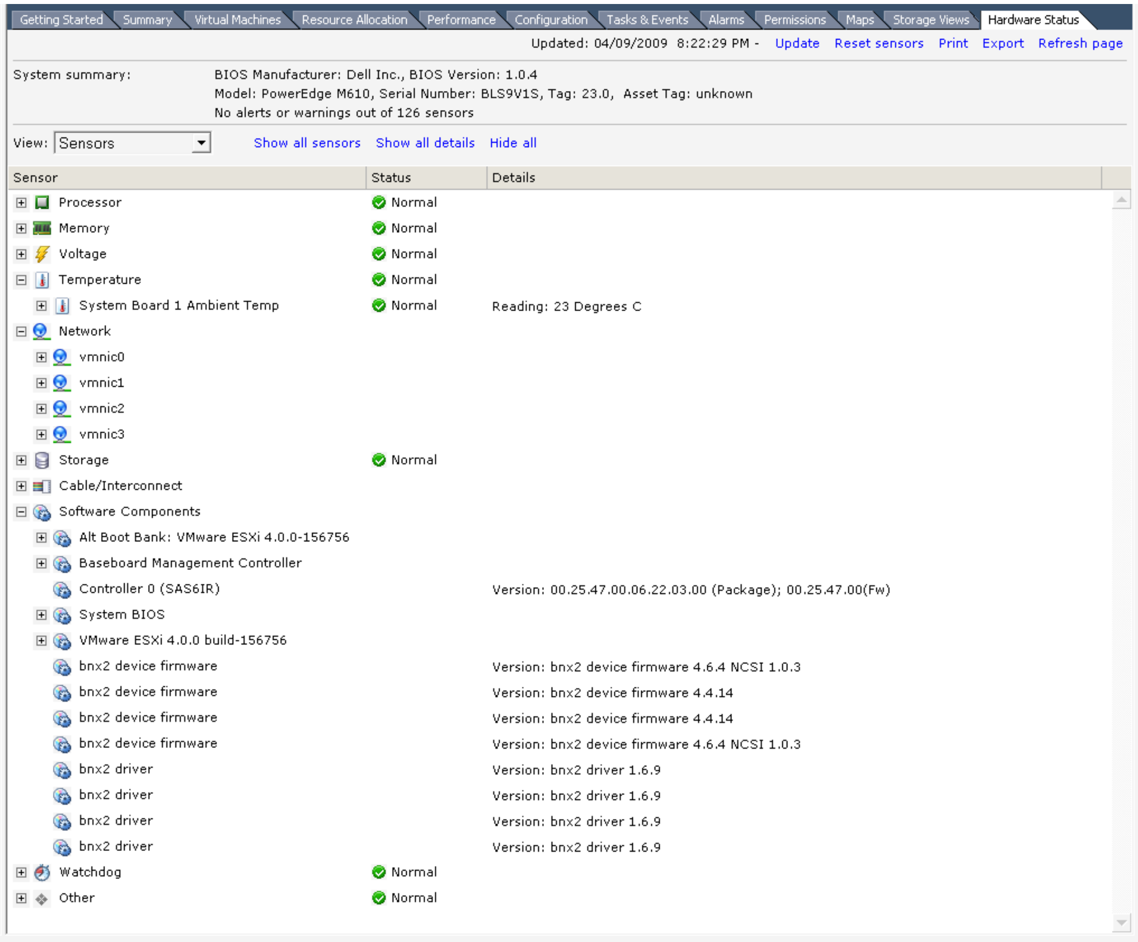

Updated: 04/09/2009 8:22:29 PM -

System summary: BIOS Manufacturer: Dell Inc., BIOS Version: 1.0.4
Model: PowerEdge M610, Serial Number: BLS9V1S, Tag: 23.0, Asset Tag: unknown
No alerts or warnings out of 126 sensors

View: Sensors

Sensor	Status	Details
Processor	Normal	
Memory	Normal	
Voltage	Normal	
Temperature	Normal	
System Board 1 Ambient Temp	Normal	Reading: 23 Degrees C
Network		
vmnic0		
vmnic1		
vmnic2		
vmnic3		
Storage	Normal	
Cable/Interconnect		
Software Components		
Alt Boot Bank: VMware ESXi 4.0.0-156756		
Baseboard Management Controller		
Controller 0 (SAS6IR)		Version: 00.25.47.00.06.22.03.00 (Package); 00.25.47.00(Fw)
System BIOS		
VMware ESXi 4.0.0 build-156756		
bnx2 device firmware		Version: bnx2 device firmware 4.6.4 NCSI 1.0.3
bnx2 device firmware		Version: bnx2 device firmware 4.4.14
bnx2 device firmware		Version: bnx2 device firmware 4.4.14
bnx2 device firmware		Version: bnx2 device firmware 4.6.4 NCSI 1.0.3
bnx2 driver		Version: bnx2 driver 1.6.9
bnx2 driver		Version: bnx2 driver 1.6.9
bnx2 driver		Version: bnx2 driver 1.6.9
bnx2 driver		Version: bnx2 driver 1.6.9
Watchdog	Normal	
Other	Normal	

Figure 5: vSphere Client Health Tab

Asset Inventory

The OMSA user interface also displays asset inventory information about the physical hardware including:

- **Hardware assets** – processor, memory, NIC, power supplies, slots
- **Software** – operating system, BIOS, firmware, systems management agents
- **FRU (Field Replaceable Unit)** – View FRU information as shown in Figure 5; this provides serial number and part number information that is valuable during troubleshooting and diagnostics
- **Storage** – controllers, physical disks, virtual disks, and enclosures
- **Remote Access Information** – remote access devices (BMC, DRAC, or iDRAC) including IPMI version, networking information, and features enabled via the device (LAN Access, Serial-Over-LAN (SOL) access, etc.)

If ESXi is installed on a blade server, information is also provided about the modular chassis enclosure.

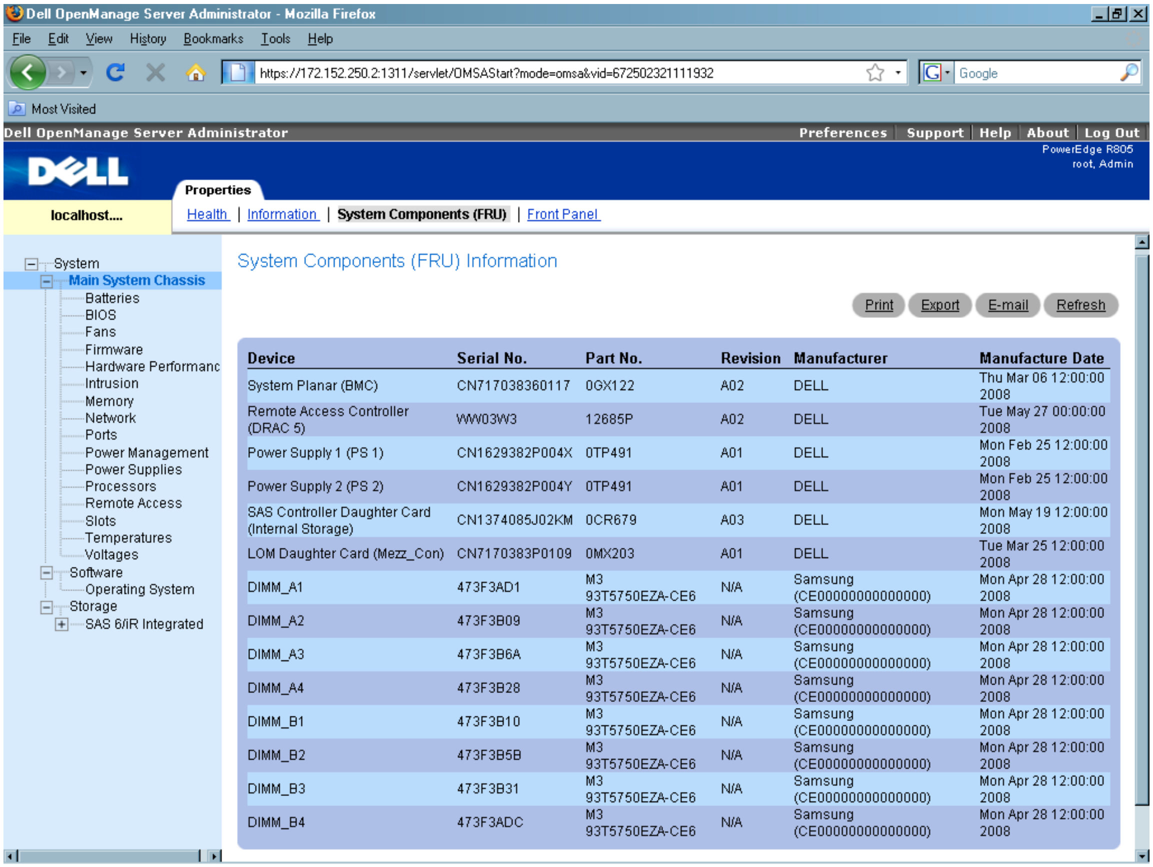

System Components (FRU) Information

Device	Serial No.	Part No.	Revision	Manufacturer	Manufacture Date
System Planar (BMC)	CN717038360117	0GX122	A02	DELL	Thu Mar 06 12:00:00 2008
Remote Access Controller (DRAC 5)	WW03W3	12685P	A02	DELL	Tue May 27 00:00:00 2008
Power Supply 1 (PS 1)	CN1629382P004X	0TP491	A01	DELL	Mon Feb 25 12:00:00 2008
Power Supply 2 (PS 2)	CN1629382P004Y	0TP491	A01	DELL	Mon Feb 25 12:00:00 2008
SAS Controller Daughter Card (Internal Storage)	CN1374085J02KM	0CR679	A03	DELL	Mon May 19 12:00:00 2008
LOM Daughter Card (Mezz_Con)	CN7170383P0109	0MX203	A01	DELL	Tue Mar 25 12:00:00 2008
DIMM_A1	473F3AD1	M3 93T5750EZA-CE6	N/A	Samsung (CE00000000000000)	Mon Apr 28 12:00:00 2008
DIMM_A2	473F3B09	M3 93T5750EZA-CE6	N/A	Samsung (CE00000000000000)	Mon Apr 28 12:00:00 2008
DIMM_A3	473F3B6A	M3 93T5750EZA-CE6	N/A	Samsung (CE00000000000000)	Mon Apr 28 12:00:00 2008
DIMM_A4	473F3B28	M3 93T5750EZA-CE6	N/A	Samsung (CE00000000000000)	Mon Apr 28 12:00:00 2008
DIMM_B1	473F3B10	M3 93T5750EZA-CE6	N/A	Samsung (CE00000000000000)	Mon Apr 28 12:00:00 2008
DIMM_B2	473F3B5B	M3 93T5750EZA-CE6	N/A	Samsung (CE00000000000000)	Mon Apr 28 12:00:00 2008
DIMM_B3	473F3B31	M3 93T5750EZA-CE6	N/A	Samsung (CE00000000000000)	Mon Apr 28 12:00:00 2008
DIMM_B4	473F3ADC	M3 93T5750EZA-CE6	N/A	Samsung (CE00000000000000)	Mon Apr 28 12:00:00 2008

Figure 6: FRU

Configuration

OMSA components allow you to perform hardware configuration via the user interface. Examples of configuration tasks are:

- Enable/disable the system's front panel buttons.
- Configure key BIOS attributes.
- Configure thresholds corresponding to the different sensors and configure the thermal shutdown severity level (warning or critical) to perform an emergency shut-down, if necessary.
- Configure the remote access device (users, networking, Serial-over-LAN (SOL) configuration, and serial port configuration) and configure Platform Event Filters (PEFs) to respond to hardware events that are monitored by the service processor.
- Perform various storage configuration tasks for supported internal and direct-attached storage, including configuring virtual disks, assigning hot spares, and triggering tasks such as rebuilding failed drives (see Figure 7).

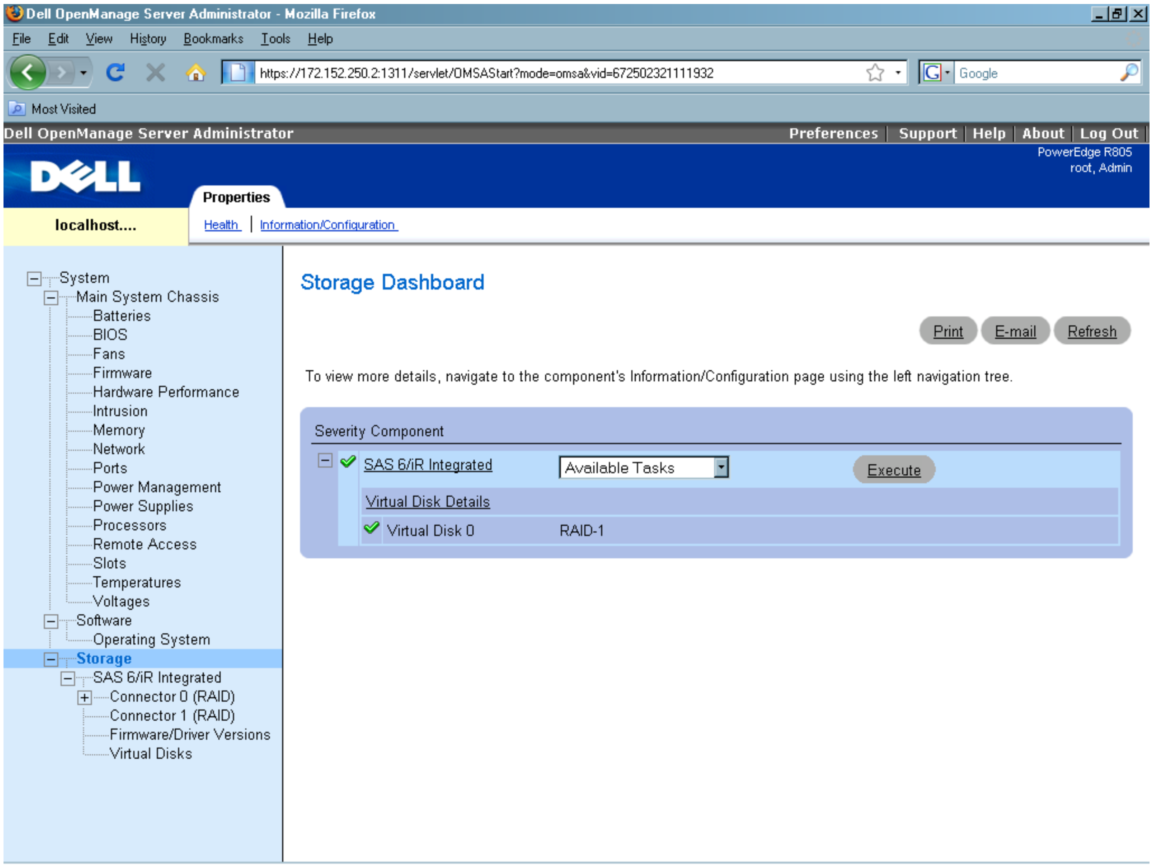

To view more details, navigate to the component's Information/Configuration page using the left navigation tree.

Severity	Component		
✓	SAS 6/iR Integrated	Available Tasks	Execute
	Virtual Disk Details		
✓	Virtual Disk 0	RAID-1	

Figure 7: Storage Configuration

Certain OMSA operations are not permitted—for example, a subset of the power control operations. These operations are documented in the *Dell OpenManage Server Administrator User's Guide* available at www.dell.com/support.

After the initial setup, you can use vSphere Client or vCLI to remotely configure the virtualization layer, including virtual network and storage devices. Alternatively, an advanced capability available within vSphere called Host Profiles allows you to configure an ESXi host to be captured and automatically replicated across similarly-equipped servers.

Software Update

With the elimination of the service console, the number of patches and updates that must be applied to the system is significantly reduced. All updates to ESXi are whole-image replacements, and they include the latest Dell management provider and hardware instrumentation, so there is no need to use a separate tool for those updates.

There are a number of ways to update the ESXi image, with VMware Update Manager being the most important. This tool allows you to download and apply the appropriate update image automatically to the ESXi host, as well as the virtual machines running on it. It is also integrated with VMware's Distributed Resource Scheduler (DRS) to enable zero-downtime patching of ESX/ESXi hosts.

Dell has integrated functionality into third-party consoles to enable deployment of BIOS/firmware updates. However, due to the secure architecture of an ESXi system, deployment of BIOS/firmware bundles to the system through agent software resident in the host operating system is not permitted. To perform firmware updates, an administrator can use the Dell System Build and Update Utility (SBUU), which is bootable media that ships with every PowerEdge server. The update files are available in the Server Update Utility media. Both utilities can be downloaded from Dell's support site at support.dell.com.

Dell recommends migrating virtual machines off of the server before performing BIOS/firmware updates. If VMotion is not configured, power off the virtual machines or suspend them and perform a cold migration.

Summary

VMware ESXi provides virtualization infrastructure that is simple to install and deploy. The manageability functions available through VMware and Dell utilities simplify systems management by providing a centralized administrative console. Dell continues to work with VMware to simplify the systems management experience for virtualization-enabled platforms in future generations of the solution stack, thus enhancing the seamless manageability of an ESXi solution from a centralized remote console.

Viswanathan Balakrishnan is a Software Validation Lead Engineer in *Dell Product Group - Business Software Validation*, specializing in Enterprise and Client Systems Management. He has a Masters in Applied Sciences/Computer Technology from Coimbatore Institute of Technology and an MBA from the University of Madras, India.

Balasubramanian Chandrasekaran is a Systems Engineer in the *Dell Virtualization Solutions Engineering Group*. He has over five years of experience in server virtualization and a Masters in Computer Science from Ohio State University.

Sudhir Shetty is a Software Strategist in the *Dell Systems Management Group*. He evaluates technologies and defines solutions in the areas of client and enterprise systems management and virtualization. He has a Masters in Computer Science from the University of Texas at Austin.

Alok Sinha is a Software Validation Engineer Senior Analyst in *Dell Enterprise Software Validation*. Alok has a Masters in Computer Applications from Indira Gandhi National Open University, and he has worked at Dell for over three years.

Charu Chaubal is a Senior Architect in Technical Marketing at VMware, where he is chartered with enabling customer adoption and driving key partnerships for datacenter virtualization. His areas of expertise include virtualization security, compliance, and infrastructure management. Charu has a Bachelor of Science in Engineering from the University of Pennsylvania and a PhD in Engineering from the University of California at Santa Barbara.

For more information on the tools and utilities mentioned in this article, visit www.support.dell.com and www.vmware.com. Additional information about Dell virtualization solutions is available at www.dell.com/virtualization and www.dell.com/vmware.