ENHANCED POWER MONITORING FOR DELL POWEREDGE SERVERS

By Santosh Bhadri G. V.
Ramakrishna M. S.
Surendra Bhat

The enhanced power monitoring and management features available in supported Dell™ PowerEdge™ servers enable IT administrators to easily track and manage energy use through tools such as the Dell OpenManage™ suite—ultimately helping optimize server power consumption and control operational costs in their data centers.

Related Categories:

Dell OpenManage
Dell PowerEdge servers
Green IT
Power and cooling

Visit DELL.COM/PowerSolutions for the complete category index.

The rapid expansion of enterprise data centers and the advent of high-performance servers have driven up energy requirements and their related costs for many organizations. Optimizing energy consumption has therefore become a top priority for IT managers. Dell PowerEdge servers equipped with enhanced power monitoring and management features—including PowerEdge models 1950 III, 2950 III, 2970, R805, R900, and R905—along with tools such as the Dell OpenManage suite and the Baseboard Management Controller (BMC) Management Utility (BMU) can help enterprises take a first step toward streamlining and controlling energy use.

KEY POWER MONITORING FEATURES

The enhanced power monitoring features on Dell PowerEdge servers use Power Management Bus (PMBus) technology. PMBus is a two-wire serial bus protocol based on the System Management Bus (SMBus) protocol, which itself is an enhanced derivative of the popular Inter-Integrated Circuit (I^2C) bus protocol designed to provide additional functionality for power control applications. PMBus is supported by several power supply manufacturers.[1]

Servers supporting these features require hardware and firmware with PMBus support; in the current implementation, the power supply units (PSUs) and BMC must be PMBus compliant. The firmware on the PSU measures the instantaneous current (in amps), power (in watts), and input voltage (in volts) at regular intervals, and the BMC on the server planar accesses these values from the PSU through the PMBus, also at regular intervals. Configurations using noncompliant or mixed-mode PSUs do not support power monitoring, and the Power Monitoring link does not appear in the left pane of the Dell OpenManage Server Administrator (OMSA) graphical user interface (GUI) for systems using these configurations.

In addition, the BMC can initiate administrator-configured actions when the system power consumption exceeds a defined threshold value. These actions can include throttling the processors, sending alerts, powering down the system, and so on using Platform Event Filters (PEFs). Figure 1 illustrates the high-level interface between the BMC and PSUs using the PMBus, a typical configuration on PowerEdge servers.

POWER MONITORING USING DELL OPENMANAGE SERVER ADMINISTRATOR

OMSA is a server configuration and monitoring tool for Dell PowerEdge servers. OMSA 5.3 and later support power monitoring and configuration on

[1] For more information on PMBus, visit www.pmbus.org.

supported Dell PowerEdge servers with power monitoring capabilities, enabling reporting of power consumption data such as energy use and peak power consumption. The OMSA instrumentation layer gathers power management data from the BMC and populates it in OMSA, where administrators can access it through either the OMSA GUI or the OMSA command-line interface (CLI). Administrators can also use OMSA to configure power management options such as PEFs and Simple Network Management Protocol (SNMP) traps.

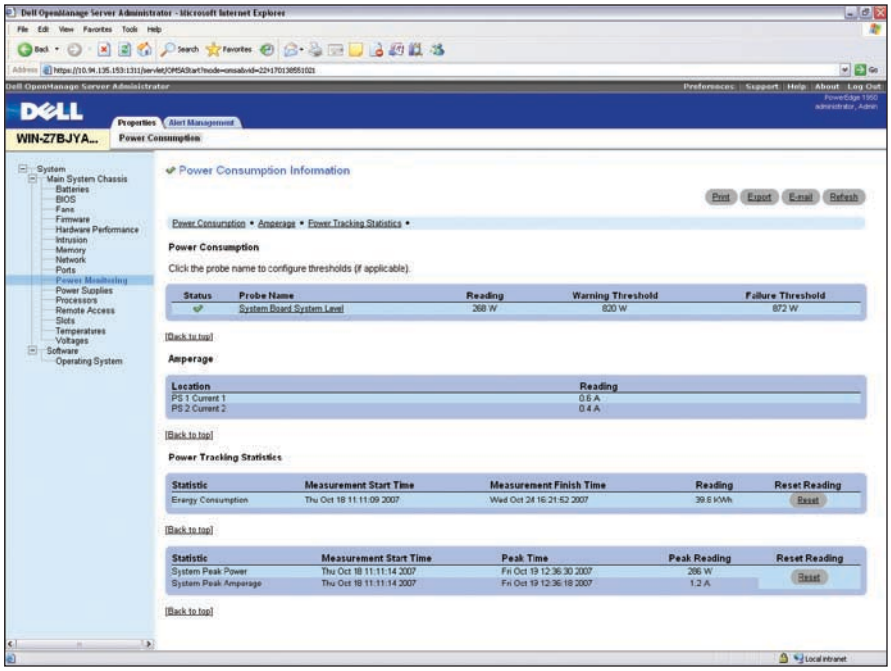

Figure 2. *Power Consumption Information screen in Dell OpenManage Server Administrator*

Viewing power consumption data

Administrators can view power consumption data in the OMSA GUI by selecting System > Main System Chassis > Power Monitoring in the left pane (see Figure 2). The Power Consumption section shows the system's instantaneous power consumption (in watts), shown as the System Board System Level power consumption. This instantaneous value is the average value of readings over a one-minute period. Administrators can set the warning threshold for this component.

The Amperage section shows the instantaneous current (in amps) of the current sensors on each of the PSUs, shown as "PS 1 Current 1" and "PS 2 Current 2." Like the power consumption value, this instantaneous value is an average value of readings recorded over a one-minute period, as retrieved from the BMC.

The Power Tracking Statistics section shows several values related to power consumption. The Energy Consumption statistic is the cumulative power consumed by the system (in kilowatt-hours), and includes the start and finish time for this cumulative reading. The start time is defined as the time when the system first received power or when the reading was last reset, while the finish time is typically the current time. OMSA administrators can reset the cumulative system power reading using the OMSA GUI or CLI. This reading also resets following a complete loss of power or a firmware flash, but not when the system is powered down or power cycled. Reset actions are recorded in the OMSA command log.

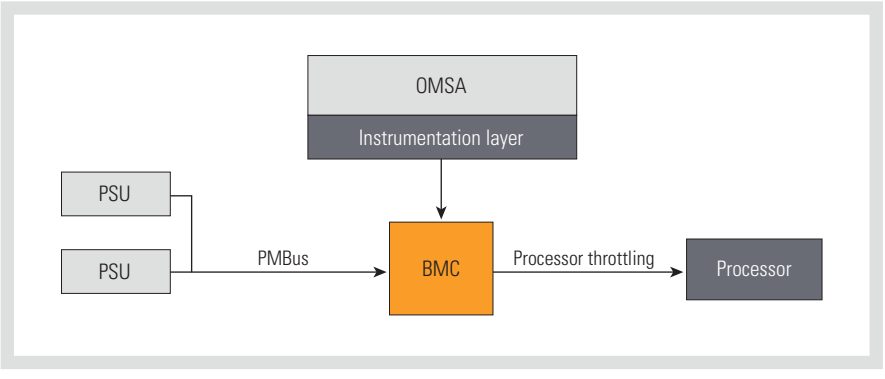

Figure 1. *Power monitoring architecture on Dell PowerEdge servers*

The System Peak Power statistic is the peak power consumed by the system (in watts), along with the measurement start time and the time of peak power use since the start time. The System Peak Amperage statistic is the peak current reading (in amps), along with the measurement start time and the time of peak current since the start time. Administrators can reset both the peak power and peak amperage readings using the OMSA GUI or CLI.

Power consumption details are available for privileged user levels such as User, Power User, and Administrator in a read-only view through the OMSA instrumentation. To provide access control and help increase security, certain options are available only for certain user levels, as shown in Figure 3.

Administrators should note that modular blade server systems have power supply sensors and voltage probes only for the enclosure as a whole, not for individual server blades.

Using the command-line interface

Power monitoring options available in the OMSA GUI are also available through the CLI. For example, administrators can use

	User	Power User	Administrator
View power monitoring data	✔	✔	✔
Configure power monitoring thresholds		✔	✔
Reset power monitoring readings			✔
View alert management data	✔	✔	✔
Set alert actions		✔	✔
Set PEF actions			✔
Enable SNMP traps		✔	✔

Figure 3. *Feature access for different user levels in Dell OpenManage Server Administrator*

the `omreport chassis pwrmonitoring` command to view power tracking statistics, and the `omconfig chassis pwrmonitoring` command along with the appropriate sub-command to configure thresholds or reset readings.

PowerEdge servers that support power monitoring now include a set of power supply voltage probes that monitor the instantaneous input voltage for each PSU. Administrators can view these readings by selecting System > Main System Chassis > Voltages in the left pane of the OMSA GUI or by using the `omreport chassis volts` command in the CLI.

OMSA can also determine whether individual PSUs support power monitoring. Administrators can view this information by selecting System > Main System Chassis > Power Supplies in the left pane of the OMSA GUI, then clicking the Properties tab. For systems that support power monitoring, the Individual Power Supply Elements table includes a Power Monitoring Capable column. Administrators can use the `omreport chassis pwrsupplies` command to display this data.

Configuring Platform Event Filters

PEFs are a BMC feature that administrators can configure to take appropriate action when a warning or failure is detected for a particular system component. These actions occur at the system board level and are independent of the server OS. PowerEdge servers that support power monitoring now include two additional PEFs: System Power Probe Warning and System Power Probe Failure.

Administrators can view, enable and disable, and configure PEFs in the OMSA GUI by selecting the System link in the left pane, then selecting Alert Management > Platform Events. The PEF list for supported servers also adds a Power Reduction column that displays the current setting for the System Power Probe Warning and System Power Probe Failure PEFs (see Figure 4). When enabled, these PEFs allow the system to take appropriate action to reduce power consumption if the system crosses the warning or critical threshold for power consumption. Administrators can also configure Platform Event Traps using PEFs for these power monitoring probes.

Administrators should keep in mind that using PEFs to reduce power consumption places the processor in a degraded state that can reduce overall system performance. When the system is in a reduced-power mode and power consumption remains below the warning threshold, the system can begin to regain its performance. Administrators can view the processor status by selecting System > Main System Chassis > Hardware Performance in the left pane of the OMSA GUI (see Figure 5). OMSA also records the events for this action in the Embedded Server Management (ESM) and Alert logs.

Configuring SNMP traps

Administrators can use OMSA to configure and generate SNMP traps for events related to power monitoring. Administrators with appropriate privileges can configure alerts for the System Power Probe Warning and System Power Probe Failure traps in the OMSA GUI by selecting the System link in the left pane, then selecting Alert Management > SNMP Traps.

Figure 4. *Platform Event Filters in Dell OpenManage Server Administrator*

POWER MONITORING USING THE BMC MANAGEMENT UTILITY

Administrators can access and configure power monitoring information directly on the BMC using the BMU, which connects to remote BMCs using an out-of-band interface. Administrators can use the following commands to view power monitoring and sensor information (see Figure 6):

```
ipmish -ip ipaddress -u
   username -p password
   powermonitor
ipmish -ip ipaddress -u
   username -p password
   sensor
```

The commands `ipmish -help powermonitor` and `ipmish -help sensor` display additional information on using these commands. The BMU also supports configuring thresholds, alert actions, and PEFs, but these operations are limited to the Administrator user level.

POWER MONITORING USING DELL OPENMANAGE IT ASSISTANT

Dell OpenManage IT Assistant is a remote, network-based, many-to-one management console for Dell PowerEdge servers. IT Assistant 8.1 and later support power management for individual servers or groups of servers on a network. This power management feature is designed to serve as an extension of the IT Assistant performance monitoring feature, helping administrators collect, store, and display instantaneous values of power consumed (in watts), power drawn by each power supply (in amps), and total energy use (in kilowatt-hours) at defined intervals. Power monitoring requires that OMSA 5.3 or later be installed on the managed servers.

Creating power monitoring tasks

Administrators can create power monitoring tasks in IT Assistant by selecting Manage > Performance and Power Monitoring, then specifying the IP address, start and end time, attributes, sampling interval, and security credentials. For the attributes, they can select appropriate performance and power monitoring counters: Processor Utilization Time and Available Memory for performance, and Aggregate Energy, Aggregate Power, Amperage per Power Supply, Energy Consumption, Peak Amperage, and Power Consumption for power monitoring. Administrators can configure warning and critical threshold values for the Aggregate Power, Amperage per Power Supply, and Power Consumption counters, which trigger an appropriate alert when the system exceeds the defined thresholds.

When creating a task, choosing an appropriate polling interval can be

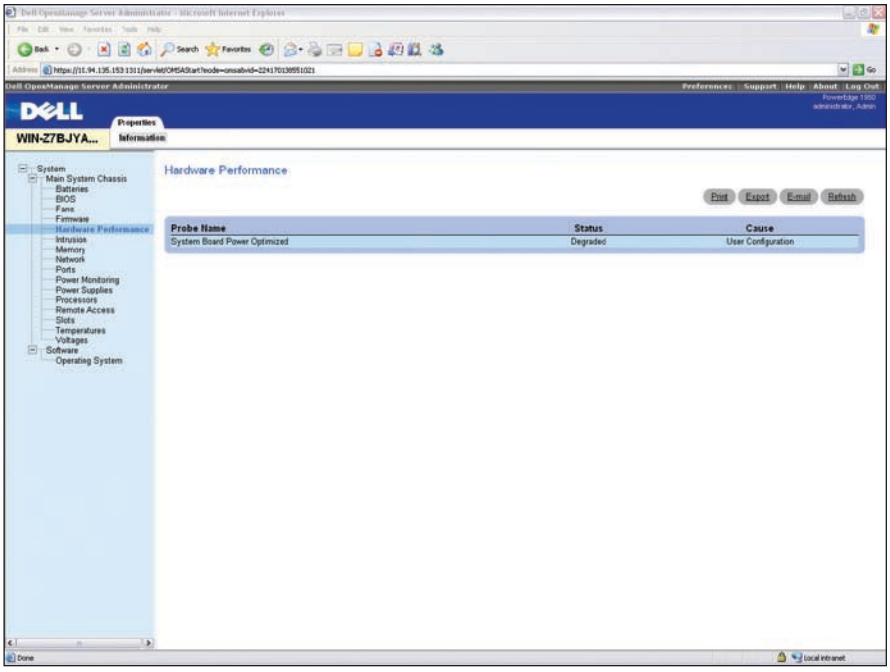

Figure 5. *Hardware Performance status for a Dell PowerEdge server in a reduced-power mode*

```
Power Tracking Statistics
Statistic                : Energy Consumption
Measurement Start Time   : Wed Oct 24 16:45:36 2007
Measurement Finish Time  : Sat Nov 03 02:31:10 2007
Reading                  : 58.3 kWh

Statistic                : System Peak Power
Measurement Start Time   : Wed Oct 24 16:45:40 2007
Peak Time                : Tue Oct 30 22:10:44 2007
Peak Reading             : 307 W

Statistic                : System Peak Amperage
Measurement Start Time   : Wed Oct 24 16:45:40 2007
Peak Time                : Tue Oct 30 22:10:34 2007
Peak Reading             : 1.3 A
```

Figure 6. *Example power monitoring output from the BMC Management Utility*

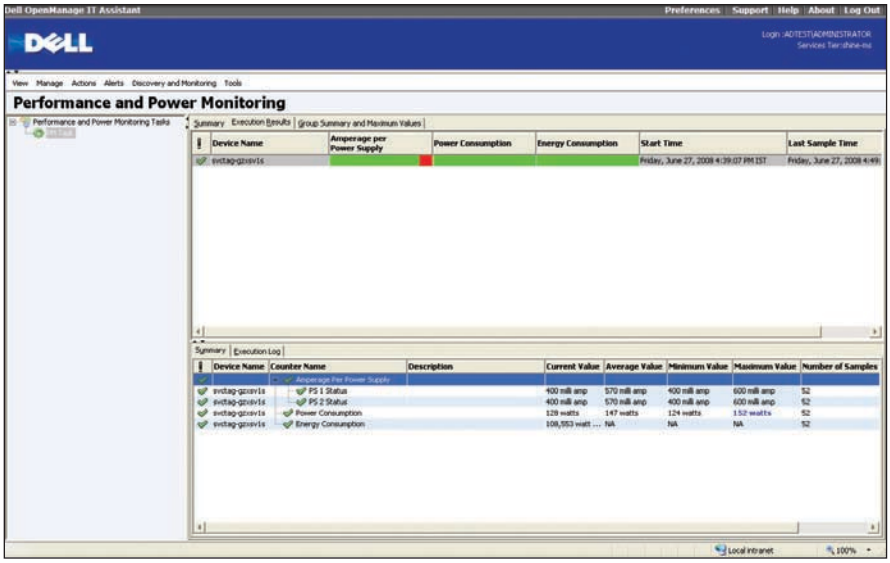

Figure 7. *Execution Results tab in Dell OpenManage IT Assistant*

critical to generating useful data. A polling interval that is too large can prevent IT Assistant from adequately capturing variations in performance and power consumption, which may result in inaccurate instantaneous power consumption charts. Combining a small sampling interval with a large number of counters and devices may cause excessive resource utilization on the managed servers and network.

Analyzing task results

After initiating a task, administrators can click the Summary tab to see an overview of the task and selected attributes, which are regularly updated with the most recent sample time. The Execution Results tab displays the current status of each server, with the colors of the bars determined by the percentage of time for which a particular attribute had a critical (red), warning (yellow), or normal (green) status (see Figure 7). After selecting a server, administrators can view detailed information on the counters at the bottom of the page. They can also view these readings in a chart by right-clicking on a system and selecting "View chart."

IT Assistant also collects the maximum power consumption and current values for each system during each polling cycle, and regularly updates the peak values by comparing the existing values with the current readings. Administrators can view the latest available data in the Group Summary and Maximum Values tab on the Performance and Power Monitoring page, and can view a chart of this data by right-clicking on a particular counter and selecting "View chart."

After running a task on an individual server or a group of servers, administrators can view the collected data on the Performance and Power Monitoring tab for each system or on the Execution Results tab for each task. Figure 8 shows the results of a task run on a single server for the Processor Utilization Time, Power Consumption, and Available Memory counters. Figure 9 shows a processor utilization chart from the same server.

Analyzing results such as those shown in Figures 8 and 9 can help administrators understand overall performance and power consumption in their data center—including identifying underutilized and overutilized systems, determining the appropriate relocation of hardware or applications to different systems based on resource availability and system load, tracking system load during peak and off-peak hours, identifying ways to balance loads between systems, and evaluating scalability based on current usage and available resources.

OPTIMIZED ENERGY USE FOR DELL POWEREDGE SERVERS

Taking advantage of the enhanced power management features in supported Dell PowerEdge servers can help enterprises optimize their data centers in multiple ways. For example, because nearly all of the power consumed by a server is

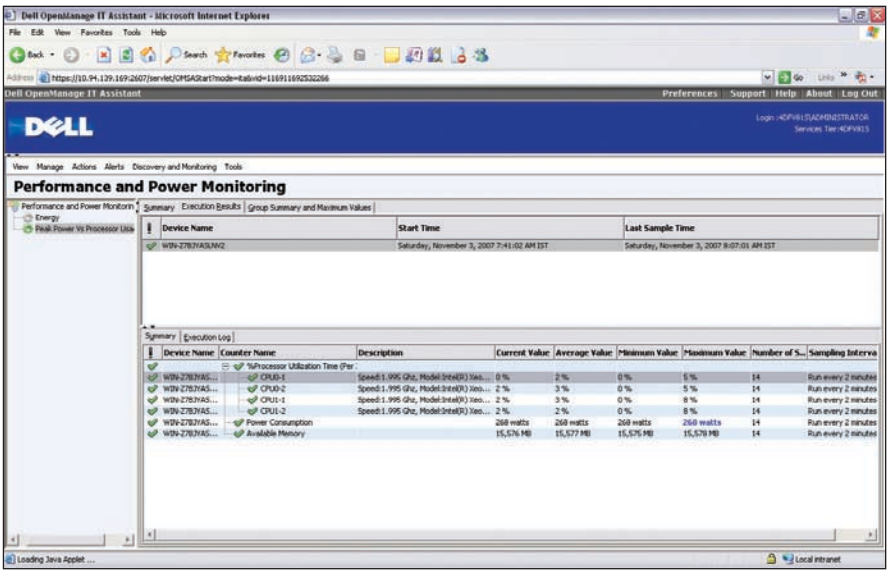

Figure 8. *Results of a performance and power monitoring task run on a single server in Dell OpenManage IT Assistant*

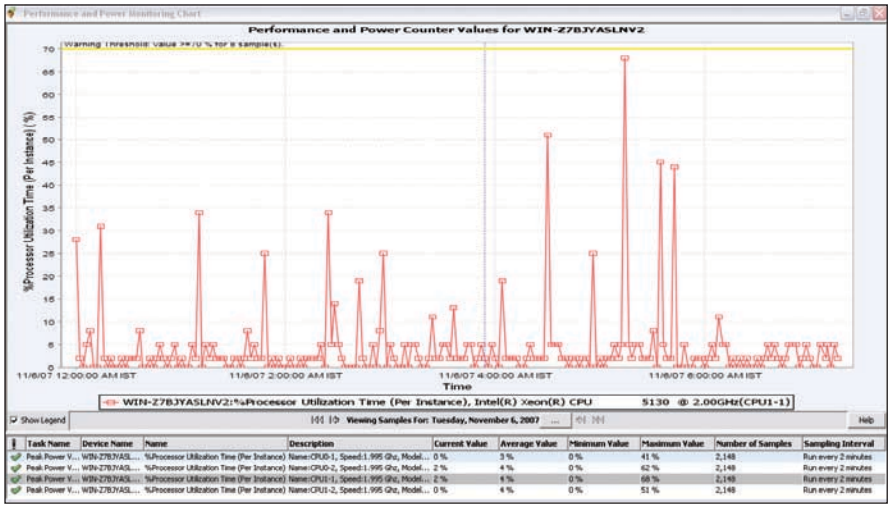

Figure 9. *Processor utilization chart from a performance and power monitoring task in Dell OpenManage IT Assistant*

dissipated as heat, administrators can use these features to help improve their data center designs to avoid hot spots. Optimizing power consumption in conjunction with monitoring enables administrators to run their servers at defined power thresholds to help meet efficiency requirements, which in turn can help increase the efficiency of cooling systems. These features can also help identify and implement appropriate load balancing based on overall system load and the needs and capabilities of specific systems. By combining these features with simplified management tools such as the Dell OpenManage suite, administrators can take the first step toward a truly energy-efficient, cost-effective IT environment.

Santosh Bhadri G. V. is a Product Group engineer senior analyst in the Sustaining Business organization at the Dell India R&D Center. His interests include computer architecture and storage technologies, and he has a bachelor's degree in Electronics and Communication from B.M.S. College of Engineering.

Ramakrishna M. S. is a Product Group adviser in the Enterprise Server Test Group at the Dell India R&D Center. His interests include electronics and music, and he has a bachelor's degree in Information Technology and a diploma in Electronics from the University of Mysore.

Surendra Bhat is a Product Group test engineer senior analyst in the Enterprise Server Test Group at the Dell India R&D Center. His interests include networking and storage technologies, and he has a bachelor's degree in Electronics Engineering from the University of Mumbai.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell OpenManage:
DELL.COM/OpenManage