



# Power and Cooling Innovations in Dell PowerEdge Servers

This technical white paper describes the Dell PowerEdge Energy Smart Architecture and the new and enhanced features designed into Dell PowerEdge 13th generation servers



Dell Enterprise Systems Engineering  
March 2016

John Jenne  
Craig Klein

# Revisions

Date	Description
March 2016	Initial release

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. Copyright © 2016 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



## Table of contents

Executive summary.....	4
Introduction .....	4
Dell’s Energy Smart Architecture .....	5
Design .....	6
Measure.....	9
Control.....	10
Report.....	12
Additional Information .....	15
About the Authors.....	15

## Tables

Table 1 Design core tenets - Dell Energy Smart technologies.....	5
--	---

## Figures

Figure 1 Dell PowerEdge SPECPower evolution .....	4
Figure 2 PowerEdge PSU portfolio.....	7
Figure 3 Power supply power loss across system loads.....	8
Figure 4 PowerEdge power monitoring capabilities .....	9
Figure 5 System-level power monitoring accuracy .....	9
Figure 6 Power capping example .....	10
Figure 7 EPR example with redundant (1+1) PSU configuration.....	11
Figure 8 PSU hot spare—redundant PSU power dissipation.....	12
Figure 9 iDRAC8 Power Monitoring GUI .....	13
Figure 10 OpenManage Power Center power monitoring GUI .....	14



# Executive summary

Dell PowerEdge servers continue to evolve to deliver industry-leading value by reducing power consumption and providing intelligent power management features at the server and data center level. In this white paper, you will learn about enhancements to Dell's Energy Smart Architecture that are integrated into the latest Dell PowerEdge 13th generation servers.

## Introduction

Dell's latest generation of PowerEdge servers build upon the Dell Energy Smart Architecture (DESA) introduced in previous generations of PowerEdge servers. DESA took the approach of including high power efficiency and intelligent power management as part of the platform's base features, enabling Dell customers to compute more while consuming less.

DESA features offer high efficiency, intelligent power management capabilities, increased ROI, and new or improved data center usage models. Examples include:

- Power supply right sizing
- IDLE power efficiency
- Circuit breaker power capping
- High accuracy power monitoring

Dell PowerEdge servers continue to evolve by significantly reducing power consumption with each new generation as illustrated in Figure 1.

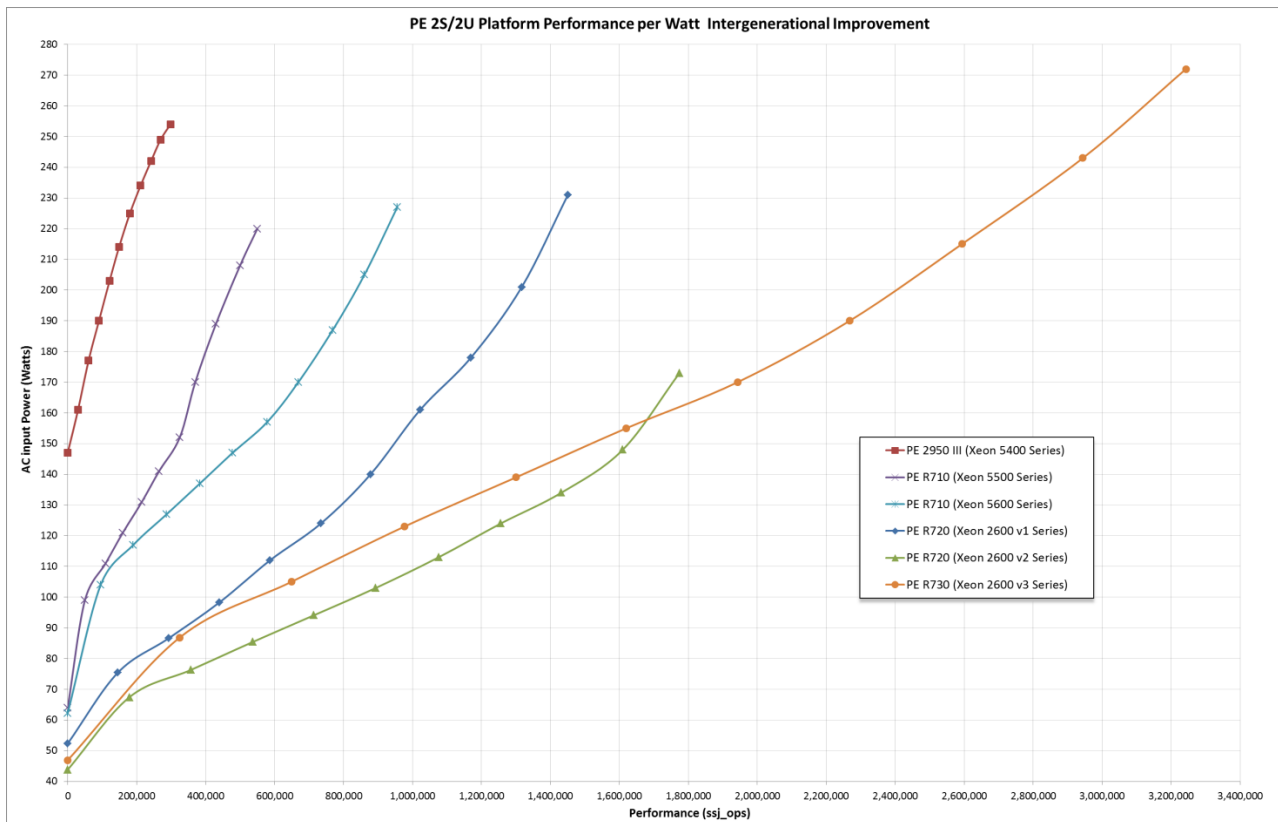


Figure 1 Dell PowerEdge SPECpower evolution



While this white paper focuses on highlighting the DESA features in Dell PowerEdge servers, visit the Dell Power and Cooling Technologies page<sup>1</sup> for additional information on values, strategy, implementation, components, and best practices.

## Dell's Energy Smart Architecture

Dell's Energy Smart Architecture is comprised of four core tenets: design, measure, control, and report. New and enhanced Dell Energy Smart technologies have been developed for each of the four core tenets. The majority of these technologies are leveraged across the Dell PowerEdge 13th generation server portfolio (including rack, tower, and blade servers) to provide consistency and compatibility. Table 1 highlights some of the notable enhanced technologies found in Dell PowerEdge 13<sup>th</sup> generation servers.

Table 1 Design core tenets - Dell Energy Smart technologies

Design		
<b>Power supply</b> <ul style="list-style-type: none"> <li>• Common form factor</li> <li>• Extended Power Range for CPU transients</li> <li>• High efficiency main and auxiliary power rails</li> <li>• Right-sized for the system</li> <li>• Hot Spare</li> <li>• Smart Power Factor Correction</li> </ul>	<b>Voltage regulator (VR)</b> <ul style="list-style-type: none"> <li>• High efficiency</li> <li>• Switching</li> <li>• Phase shedding</li> <li>• Configuration-based tuning</li> </ul>	<b>Components</b> <ul style="list-style-type: none"> <li>• Low voltage processors</li> <li>• Low voltage DDR4 DIMMs</li> <li>• SSDs</li> <li>• Software RAID option</li> <li>• HDDs with standby</li> </ul>
<b>Thermal</b> <ul style="list-style-type: none"> <li>• High-airflow chassis</li> <li>• Optimized heat sinks</li> <li>• Independent fan control</li> <li>• Pulse width modulation fans</li> <li>• Low power fans</li> </ul>	<b>BIOS</b> <ul style="list-style-type: none"> <li>• Processor P, C, T-states</li> <li>• DDR4 frequency selection</li> <li>• DDR4 CKE and self refresh</li> <li>• Turbo Boost disable</li> </ul>	<b>Board design</b> <ul style="list-style-type: none"> <li>• Low-loss layout</li> <li>• Low-loss materials</li> <li>• Low-loss connectors</li> </ul>
Measure		
<b>Temperature</b> <ul style="list-style-type: none"> <li>• Ambient, exhaust</li> <li>• Components: processor, chipset, DIMMs, hard drives,</li> <li>• PCIe adapters, power supplies</li> </ul>	<b>Power</b> <ul style="list-style-type: none"> <li>• Amperage, BTUs, voltage, and watts</li> <li>• High accuracy</li> <li>• Components: power supplies, processor, memory, storage, I/O, fans</li> </ul>	<b>Performance</b> <ul style="list-style-type: none"> <li>• Processor utilization</li> <li>• Memory throughput</li> <li>• I/O throughput</li> </ul>

<sup>1</sup> Dell Power and Cooling Technologies: <http://www.dell.com/en-us/work/learn/power-and-cooling-technologies>



Control		
BIOS <ul style="list-style-type: none"> <li>Enhanced Active Power Controller</li> <li>(OS independent process power management)</li> <li>OS enabled processor power management</li> <li>System Profiles               <ul style="list-style-type: none"> <li>Performance</li> <li>Performance per Watt</li> <li>Dense configuration</li> </ul> </li> <li>Processor core disable</li> <li>Host LOM port disable</li> <li>External and internal USB port disable</li> <li>Unused component disable</li> </ul>	Firmware <ul style="list-style-type: none"> <li>Data center electrical and thermal power capping</li> <li>Non-linear fan curve</li> <li>Adaptive thermal algorithm</li> <li>Closed loop thermal monitoring and throttling</li> <li>Dynamic DIMM 2x refresh</li> <li>System thermal power capping</li> <li>Management LOM port disable</li> <li>Memory VR phase shedding</li> </ul>	<ul style="list-style-type: none"> <li>Power inventory/budget</li> <li>Staggered power on</li> <li>Remote power On/Off</li> <li>Right-size PSU protection</li> <li>Extended Power Range</li> <li>Physical Disk Power Management</li> <li>IDLE memory control</li> </ul>
	Dell OpenManage™ <ul style="list-style-type: none"> <li>Remote power On/Off</li> </ul>	Dell OpenManage Power Center <ul style="list-style-type: none"> <li>Group level (rack, aisle, data center) power capping</li> </ul>
Report		
Dell iDRAC8 <ul style="list-style-type: none"> <li>Node level power reporting: averages, peaks, real-time</li> <li>Node level reporting mechanisms include tables and graphs</li> <li>Node level alarms/alerts, including power cap not maintained</li> </ul>	Dell OpenManage Power Center <ul style="list-style-type: none"> <li>Provides group level (rack, aisle, data center) reporting</li> <li>Group level power reporting: averages, peaks, real-time</li> <li>Group level reporting mechanisms include tables and graphs</li> <li>Group level alarms/alerts, including power cap not-maintained</li> </ul>	

## Design

With the 12<sup>th</sup> generation of Dell PowerEdge servers, Dell introduced a power supply (PSU) portfolio with a common form factor. This strategy has been carried over into the next generation, with new and enhanced designs for Dell PowerEdge 13<sup>th</sup> generation servers. The common PSU form factor lets you help you select from a rich set of PSU options for capacity, input (AC or DC), or efficiency.

Figure 2 highlights the PSU options that are supported by the majority of the new 13<sup>th</sup> generation PowerEdge server portfolio. Some value platforms also support a low-cost cabled PSU option.



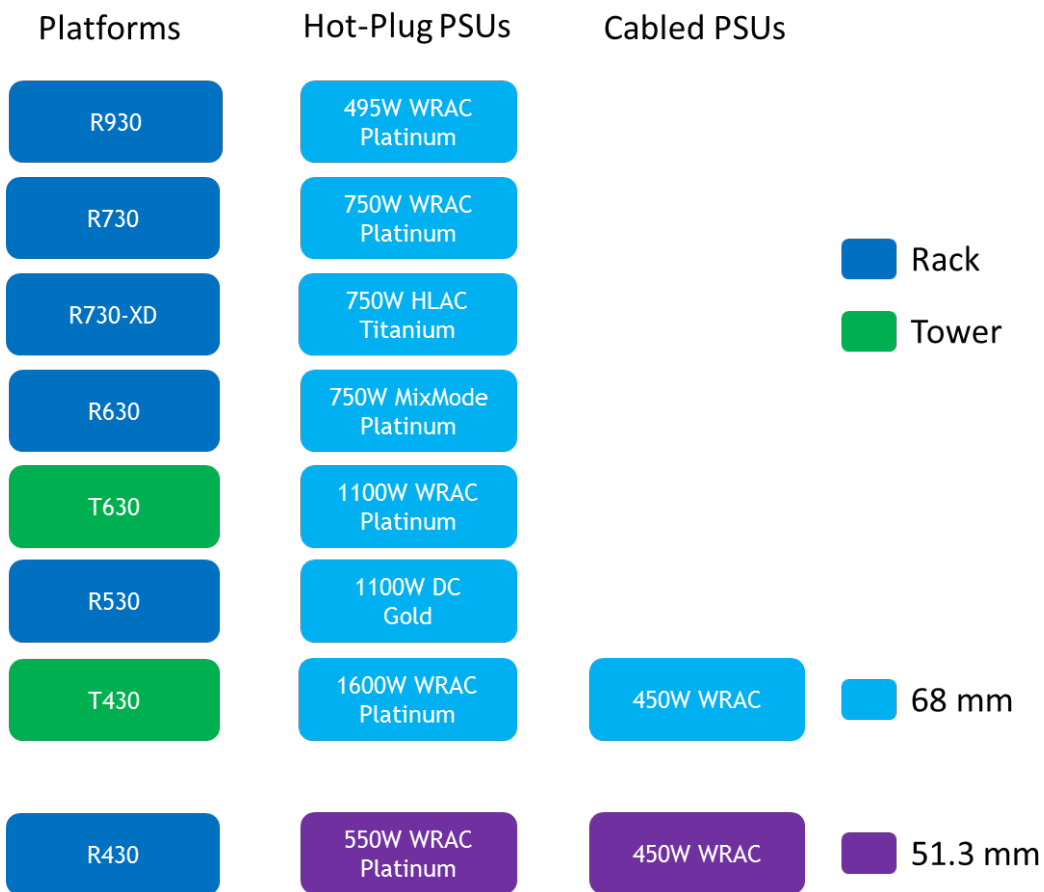


Figure 2 PowerEdge PSU portfolio



The PowerEdge PSUs meet 80 PLUS® efficiency requirements for Platinum level with an Energy Smart option to upgrade to Titanium level. Dell PowerEdge servers are equipped with industry-leading PSU efficiencies that meet 80 PLUS’s most stringent Titanium efficiency requirements.

Figure 3 highlights the reduced power dissipation benefits of right sizing and high efficiency of PowerEdge 86mm PSUs.




Figure 3 Power supply power loss across system loads

Power delivery efficiency is also improved through voltage regulator (VR) enhancements such as high efficiency design, intelligent phase shedding, and ability to tune operation for a given server configuration.

PowerEdge servers support a software RAID option for storage configurations with four SATA HDDs or less. While hardware RAID support improves performance and robustness, there is a power penalty. Depending on your requirements, software RAID provides an excellent RAID alternative to reduce platform power consumption.





## Measure

PowerEdge servers integrate a custom version of Intel® Node Manager firmware, the product of a joint Dell and Intel development effort. Node Manager provides rich power monitoring capabilities, including system-level (PSU) and subsystem-level (processor, memory, I/O, storage, and fan) monitoring. Node Manager also provides sampling rates up to ten samples per second to improve accuracy of averaging and to enable faster response to extreme power transients.




Figure 4 PowerEdge power monitoring capabilities

Power monitoring accuracy on PowerEdge servers exceeds the EPA's ENERGY STAR® certification requirements by 87%, and rivals many external high-cost power meters.




Figure 5 System-level power monitoring accuracy

Thermal monitoring capabilities include exhaust, power control unit (PCU), memory and additional component temperatures such as hard drives and PCIe adapters. The additional thermal monitoring provides pinpoint temperature information to the thermal algorithm for finely tuned control of the system thermals.



# Control

Enhanced power control features designed into PowerEdge servers target power-capping, power efficiency, IDLE power, and deployment simplification.

In addition to rich power monitoring capabilities, Dell's custom version of Intel's Node Manager provides sophisticated power capping capabilities to support various data center and platform usage models. Fast and intelligent, the Node Manager power-capping solution limits power utilization below the user-defined limit within one second, while optimizing performance within the power constraints. With PowerEdge servers, you can use power capping to limit server power consumption due to data center or rack electrical or thermal limitations. Examples of other usage models include:

- Control peak power usage during a defined time period to control billing rates
- Provisioning power to workloads based on an SLA that does not require fully configured capability
- Power cap due to power utilities mandating brownout energy consumption reduction
- Limiting power due to extreme outside temperature that limits the data center's cooling capability
- Power cap to protect the data center against power and thermal excursions (such as HVAC failure)

Dell's OpenManage Power Center provides a one-to-many power management console that can be used to set group-level (rack, aisle, data center) power cap policies. This console provides a simple interface for managing the power consumption of PowerEdge servers in the data center.

The iDRAC8 Enterprise license is required for power capping through the iDRAC8 GUI or OpenManage Power Center.

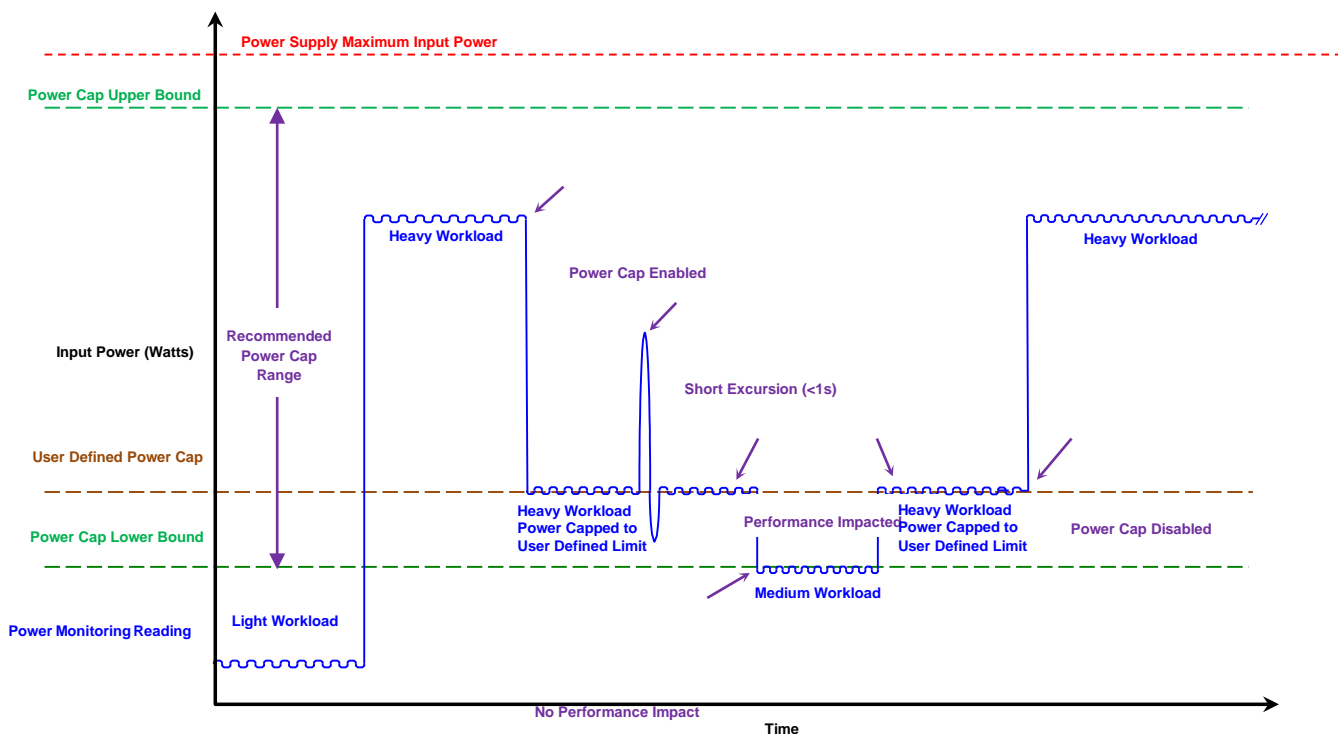


Figure 6 Power capping example



Dell’s PSU right-sizing strategy to improve power efficiency and reduce costs greatly benefits from a PowerEdge server feature called Extended Power Range (EPR) which temporarily extends the output capability of the power supply. Under typical workloads, power consumption is well below the output capability of the PSU, but if the workload spikes power beyond the PSU capability, ultra-fast hardware protection power cap policies are triggered to reduce power consumption. Extended Power Range allows richer system configurations to be supported within the power constraints of a smaller PSU.




Figure 7 EPR example with redundant (1+1) PSU configuration

The PSU feature Hot Spare improves the operating efficiency of redundant PSU configurations for light to typical workloads by putting one of the PSUs in sleep state. At higher loads, the sleeping PSU autonomously wakes to optimize operating efficiency. Visit the Dell Power and Cooling Technologies page<sup>2</sup> to view a video demo of this feature. Figure 8 highlights the hot spare reduction in PSU power dissipation.

<sup>2</sup> Dell Power and Cooling Technologies: <http://www.dell.com/en-us/work/learn/power-and-cooling-technologies>






Figure 8 PSU hot spare—redundant PSU power dissipation

Dell’s Active Power Controller (DAPC), which provides OS-independent processor power management, has been fine-tuned for next-generation Intel processors. DAPC intelligently manages Intel’s Turbo Boost feature for a near-linear performance-per-watt response, compared to non-Turbo Boost levels. DAPC allows workloads to utilize Turbo Boost when needed, providing an optimal balance between performance and power consumption.

Other control features target IDLE power. DIMM IDLE power is minimized by supporting CKE power down and self-refresh features. Storage IDLE power is minimized by the Physical Disk Power Management feature; Dell hardware RAID controllers will spin down IDLE disks. These IDLE power controls complement other power delivery efficiencies at light loads and existing IDLE power controls (for example, P-states or C-states).

System Profiles provide “easy buttons” for configuring a server for performance, performance-per-watt, or dense (conservative) configurations that adjust various BIOS performance, power, and RAS settings. The Dense configuration focuses on prioritizing RAS over performance and power for dense memory configurations.

Available Thermal Profiles complement the System Profiles by providing options for Performance or Performance-per-Watt thermal control algorithms. These thermal controls let you set a maximum exhaust temperature or increase system fan speeds by a fixed offset for additional cooling to I/O slots.

## Report

iDRAC8’s power reporting in the iDRAC webserver GUI includes a power cap alert function that alerts you when a user-defined power cap cannot be maintained. iDRAC8 also allows you to specify a power cap limit below the recommended range, and lets you know when the platform cannot maintain the specified limit.



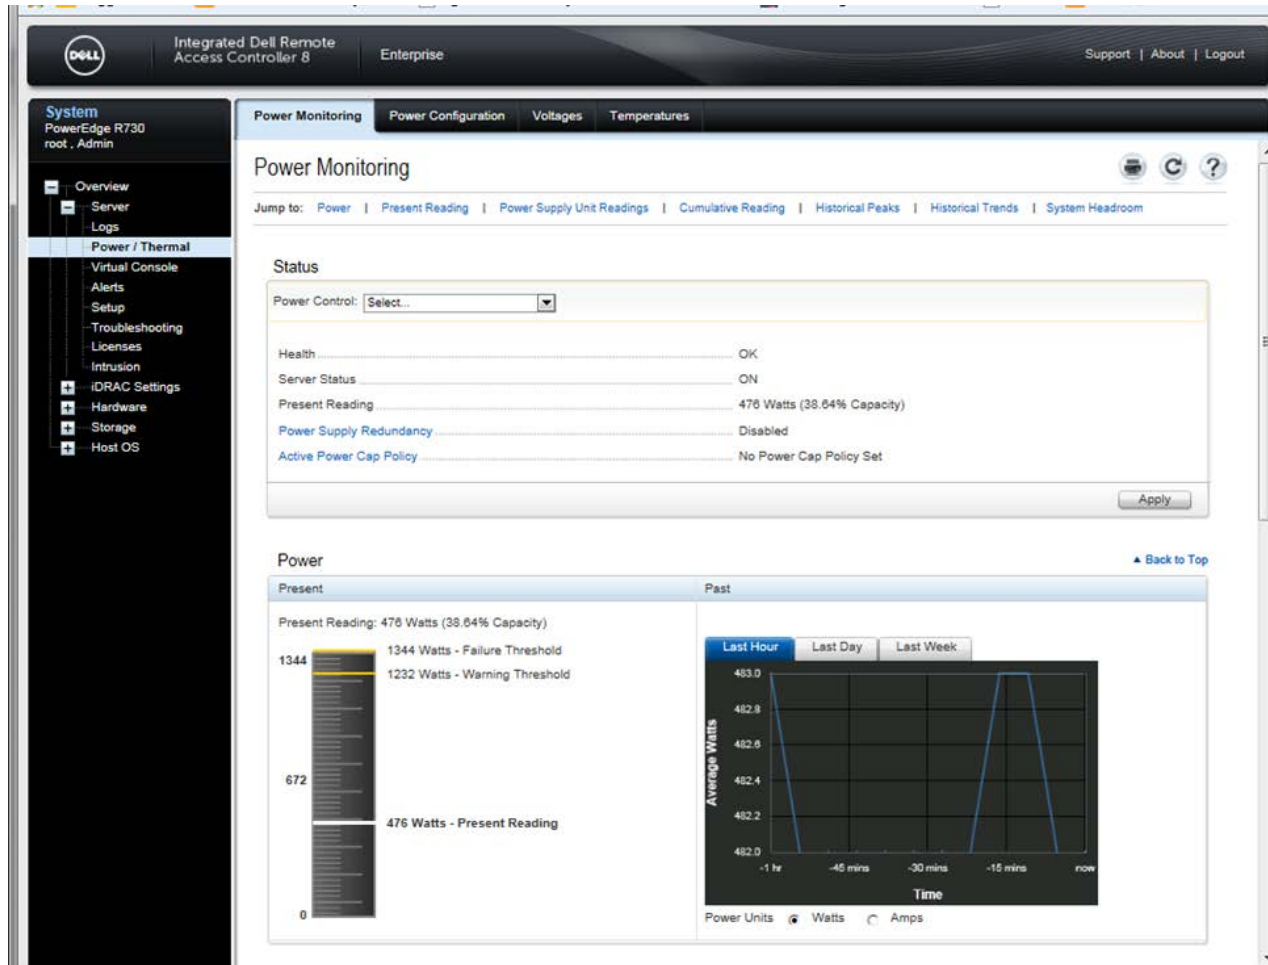


Figure 9 iDRAC8 Power Monitoring GUI

OpenManage Power Center provides powerful group-level power reporting capabilities at the rack, aisle and data center levels. This one-to-many power management console provides a clean, easy-to-use interface for users to track power consumption throughout the data center. Like iDRAC8, Power Center provides an alert function when a group-level power cap cannot be maintained. iDRAC8 Express provides power monitoring capabilities for the iDRAC8 GUI and OpenManage Power Center. An iDRAC8 Enterprise license is required for power capping in the iDRAC8 GUI or in OpenManage Power Center.





Figure 10 OpenManage Power Center power monitoring GUI



## Additional Information

For additional details and information, see the following documents and resources available on Dell.com:

- "Power Consumption Reduction: Hot Spare"
- "Power Consumption Reduction: High Efficiency Power Supplies"
- "Increasing Energy Efficiency through Modular Infrastructure"
- [Dell.com/PowerandCooling](http://Dell.com/PowerandCooling)
- [Dell.com/PowerCenter](http://Dell.com/PowerCenter)
- [Dell.com/FreshAir](http://Dell.com/FreshAir)

## About the Authors

John Jenne is a Distinguished Engineer in the Dell Enterprise Product Group focused on Node Power across the Dell Server portfolio. John has previously held positions with Compaq, MaXXan Systems, and Newisys. He has a B.S. in Computer Engineering from Clemson University, an M.S.E.E. from the University of Houston, and currently holds 16 patents.

Craig Klein is a Principle Engineer in the Dell Enterprise Product Group working on Node Power across the Dell Server portfolio. Craig has previously held positions with IBM, and Jabil Circuit. He has a B.S. in Electrical Engineering from Texas A&M University, and currently holds 148 patents.

