



Managing Server Energy Consumption Using IBM PowerExecutive

*By Patsy K. Popa
IBM Systems and Technology Group*

Executive Overview

Some server vendors are content to deliver highly powered, high-density servers into a data center that may be fast approaching the limits of its power and thermal capacity, and then hand off the difficult job of managing the power and thermal needs of those systems to the data center administrator. IBM understands the necessity of managing data center power and cooling resources and costs, and the need to limit the capital costs of upgrading to handle additional equipment.

IBM PowerExecutive™ can help you manage your power and cooling “budget” and stave off those capital improvements for as long as possible.

It is a standard feature built into selected IBM System x™ and BladeCenter® servers that enables you to measure, track and control server power consumption. IBM PowerExecutive is a valuable tool for your data center. It helps you understand how energy is used within the data center, and it enables you acquire the insight required to optimize your servers and their workloads to attain the best performance—while adhering to the power supply and cooling constraints of the data center.

IBM PowerExecutive consists of integrated server hardware and firmware that continuously monitor power consumption, regardless of the operating system or workload used. The IBM PowerExecutive Application Tool, a software extension to IBM Director¹ 5.1x (or later), monitors and displays the collected power measurement data for individual servers. Additionally, it provides analysis of long-term power trends for all servers in the data center with IBM PowerExecutive installed. This assists in the timely planning of costly facilities upgrades so they can be implemented *before* server power and temperature grow beyond the capacity of the data center.

Future releases of IBM PowerExecutive will help enable data center servers to stay within a preset power budget and automatically adjust server power budgets based on changes to the power supply capacity. Such capabilities could help you to push server performance to greater limits within power-constrained data centers. Also IBM plans to create power conservation techniques that could enable you save power when the server system load is low, thus possibly reducing operating costs.

¹ IBM Director is a suite of systems management software that ships standard with most IBM System x and BladeCenter servers.

Contents

- Executive Overview 2**
- Solving the Challenges of Server Density and Power Management 4**
- IBM PowerExecutive and Energy-efficient Components 4**
 - Built-in Metering 4
 - Data Reporting and Trending 5
 - Data Analysis Made Easy 6
 - Event Reporting: Captured and Ready for Review 6
- Future Capabilities 6**
- Conclusion 7**

Solving the Challenges of Server Density and Power Management

Data center performance, space considerations and operating costs are issues that data center administrators confront on a daily basis. The current unpredictability of energy costs, the need to tightly manage data center budgets while striving to increase data center throughput and constraints on capital spending for infrastructure upgrades create ongoing challenges for data center managers and administrators.

Issues limiting full usage of data center capacity include power availability, air conditioning for heat removal, and backup subsystems such as uninterrupted power supplies (UPSs) and generators. Until now, the correct allocation of power and cooling to servers in the data center has been provided using estimated data or by adding an external power meter to the server. Typically, estimation techniques must include large margins to keep server power consumption from exceeding power supply capacity and heat generation from exceeding cooling limits. Some data centers adjust for power and cooling limitations by only partially populating racks and blades in a chassis or by leaving large areas of open space between racks.

Meeting the challenges caused by increased server power consumption and thermal loads requires realtime data about the *actual* power usage and thermal load during data center operation. Administrators need tools that monitor and report average and peak energy demand as well as track power usage and thermal load trends. External power metering and temperature collection is one option for monitoring actual power consumption and heat generation, but it suffers the disadvantages of being costly and bulky. It also adds hardware that complicates the setup of already crowded data centers. Integrating the sensors that measure power and temperature, and then collecting the measured data on a server-by-server basis, is crucial to the practical management of power and heat in the data center.

IBM PowerExecutive and Energy-efficient Components

IBM PowerExecutive is the first product in a rapidly evolving IBM roadmap that focuses on reducing the cost of power in the data center. IBM PowerExecutive consists of hardware and firmware built into IBM servers to assist clients in managing the power and thermal needs of servers in the data center. The IBM PowerExecutive Application Tool includes three components: the PowerExecutive Server, the PowerExecutive Console and the PowerExecutive Database. The PowerExecutive Console provides an intuitive way to access the data while the server and database dependably collect and store measurements and quickly report energy trend data.

IBM PowerExecutive key features include:

- No additional hardware infrastructure to deploy
- Accurate measurement of power on a per-server basis, using sensors fully integrated into the measured system
- Collecting measurement data without affecting the execution of the workload running on the measured system
- Estimating the maximum power consumption of a server, based upon the current hardware configuration, including all installed options
- Managing and exporting power and thermal trend data

Built-in Metering

The IBM PowerExecutive application tool communicates with the management module in BladeCenter chassis and with the Baseboard Management Controllers (BMCs) in rack-mounted IBM servers. It retrieves both temperature and power measurements in realtime from all systems

that support IBM PowerExecutive. For legacy servers and components in a BladeCenter chassis that do not have power-measurement hardware, IBM PowerExecutive uses a static estimate of power.

Figure 1 is a screenshot of IBM PowerExecutive displaying the power consumption of a single IBM BladeCenter HS20 blade server. The tool collects average, minimum, and maximum power values from fine-resolution, sub-second moving averages, so that even short bursts of power usage will be detected and available for planning purposes. PowerExecutive Servers also support industry-standard IPMI commands, so that IBM clients can use their own scripts to monitor server-level power usage.

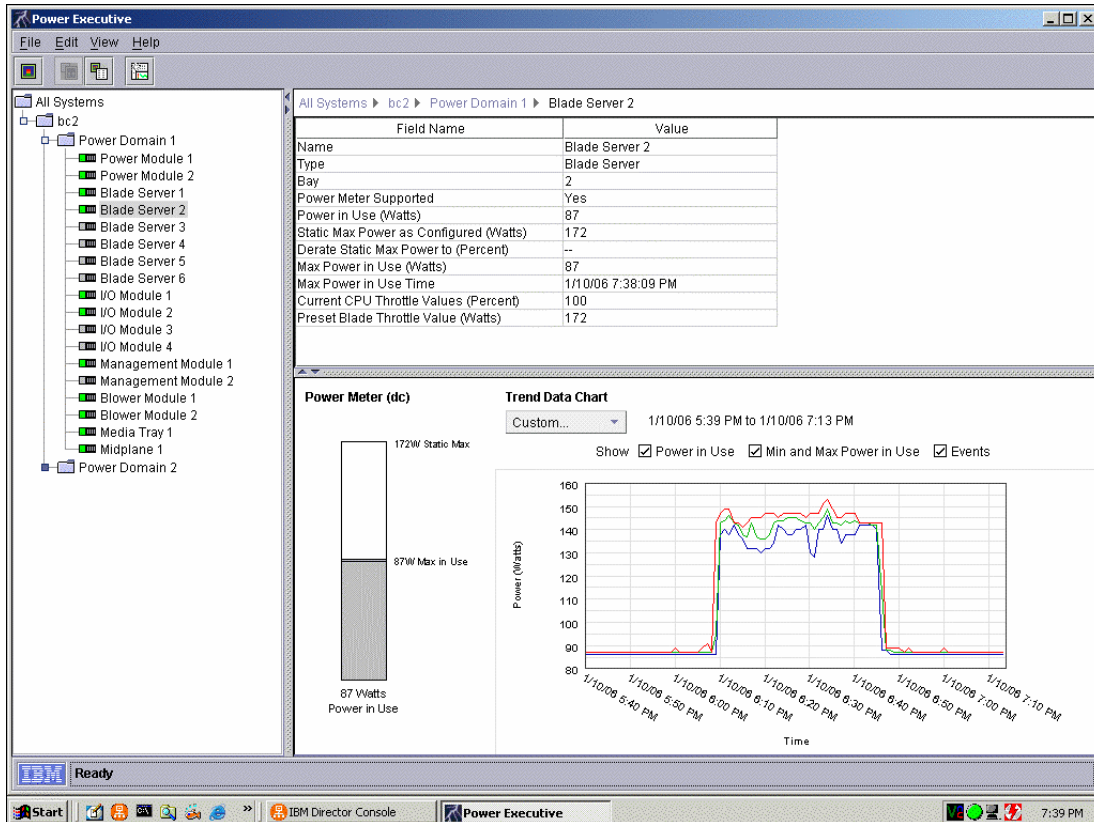


Figure 1. IBM PowerExecutive showing power measurement information for a single BladeCenter server

Data Reporting and Trending

To simplify data management, power usage is displayed hierarchically. Data can be viewed for single servers, for a whole BladeCenter chassis, for a rack and even for groups of racks in a data center. Because all power measurements are tracked inside the PowerExecutive Database, it is easy to track power usage trends over a span of weeks or months. This gives a more complete picture of data center power usage.

Continuous power measurement over long periods exposes the dynamic range of server power consumption. This leads to a better understanding of actual power and cooling utilization in the data center and reveals the margin between maximum server power consumption and power supply capacity. Such measurements are crucial to deciding if more servers can be added to the data center without additional power or thermal upgrades.

IBM PowerExecutive allows IBM clients to measure and track power consumption trends for specific IBM System x servers and BladeCenter blade servers that use Intel™, IBM PowerPC® and AMD processors. The *initial* list of supported IBM servers and BladeCenter chassis includes:

- BladeCenter HS20 server (machine type 8843)
- BladeCenter JS21 server (8844) — support planned
- AMD LS20 for IBM BladeCenter server (8850)
- All BladeCenter, BladeCenter H and BladeCenter T chassis with installed Advanced Management Modules
- System x server models x3550 and x3650

Data Analysis Made Easy

IBM PowerExecutive supports exporting power data to an external resource, such as a spreadsheet or proprietary energy analysis program. This information can then be used to analyze power consumption and energy needs for the data center as a whole, and even to analyze data center power usage compared to overall building or site energy usage. Accurate power and energy measurements and the analysis of such measurements are crucial for data center planning.

Event Reporting: Captured and Ready for Review

IBM PowerExecutive reports events that can affect power consumption. These reports provide the user with a quick and a reliable method for determining the cause of changes in the power consumption pattern. Among the events reported are:

- Blade server insertion
- Blade server removal
- Chassis going online
- Chassis going offline
- Changes in acoustic mode settings
- Changes to the power allocation policy

Future Capabilities

The data center is a dynamic, changing environment. Unexpected changes in power capacity can be caused by workload spikes, server redeployments and power subsystem failures. IBM intends, in future releases of IBM PowerExecutive, to enable servers to control their own power consumption to fit within a changing power budget. This unique ability to finely adapt power consumption would allow servers to run closer to the power and thermal limits of the data center without a loss in performance. It also enables servers to gracefully survive unexpected reductions in power and cooling subsystems and run longer on backup power. Additionally, you would be able to limit server power consumption to a predetermined budget by the system administrator. Idle servers will use power-saving modes making their power usage allocation available to other, busier servers. Thus, power can be directed and redirected to where it can best be used in the data center. This will make IBM PowerExecutive a unique tool for controlling rising utility costs and managing capital costs and, potentially, avoiding or postponing expensive upgrades.

Additional capabilities are planned that support the general philosophy of providing a toolset that lets IBM servers adapt to a changing environment of workloads, cooling and power distribution, rather than burdening the client with adapting the data center environment to the server.

Conclusion

IBM PowerExecutive is available today on both BladeCenter and rack-mounted IBM servers. It delivers timely, accurate power-consumption information for both individual servers and entire data centers. Trending and analysis tools give IBM clients a clear view of power supply and demand within the data center. The ability to plan and predict power consumption reduces the infrastructure required for redundancy, allows more servers to be installed on smaller power feeds, and thus helps lower overall data center capital and operation costs. Ultimately, this helps IBM clients optimize existing computing resources by using them at levels close to the cooling and power limits of the data center and thus improving workload throughput performance. Future enhancements to IBM PowerExecutive would include the ability to precisely control the maximum power consumption of a server, giving IBM clients flexibility in deploying servers in power-constrained environments.

To download a copy of IBM PowerExecutive, visit our Web site at http://ibm.com/servers/eserver/xseries/systems_management/ibm_director.



For More Information

IBM System x and xSeries Servers	ibm.com/systems/x
IBM BladeCenter Server and options	ibm.com/systems/bladecenter
IBM Rack Configurator	ibm.com/pc/us/eserver/xseries/library/configtools.html
IBM ServerProven® Program	ibm.com/pc/us/compat
IBM Technical Support	ibm.com/server/support

Legal Information

© IBM Corporation 2006
IBM Systems and Technology Group
Dept. U2SA
3039 Cornwallis Road
Research Triangle Park, NC 27709

Produced in the USA

May 2006

All rights reserved.

For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, the e-business logo, BladeCenter, PowerExecutive, PowerPC, ServerProven, System x and xSeries are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see <http://ibm.com/legal/copytrade.shtml>.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows and the Windows logo are trademarks or registered trademarks of Microsoft Corporation.

Other company, product and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.