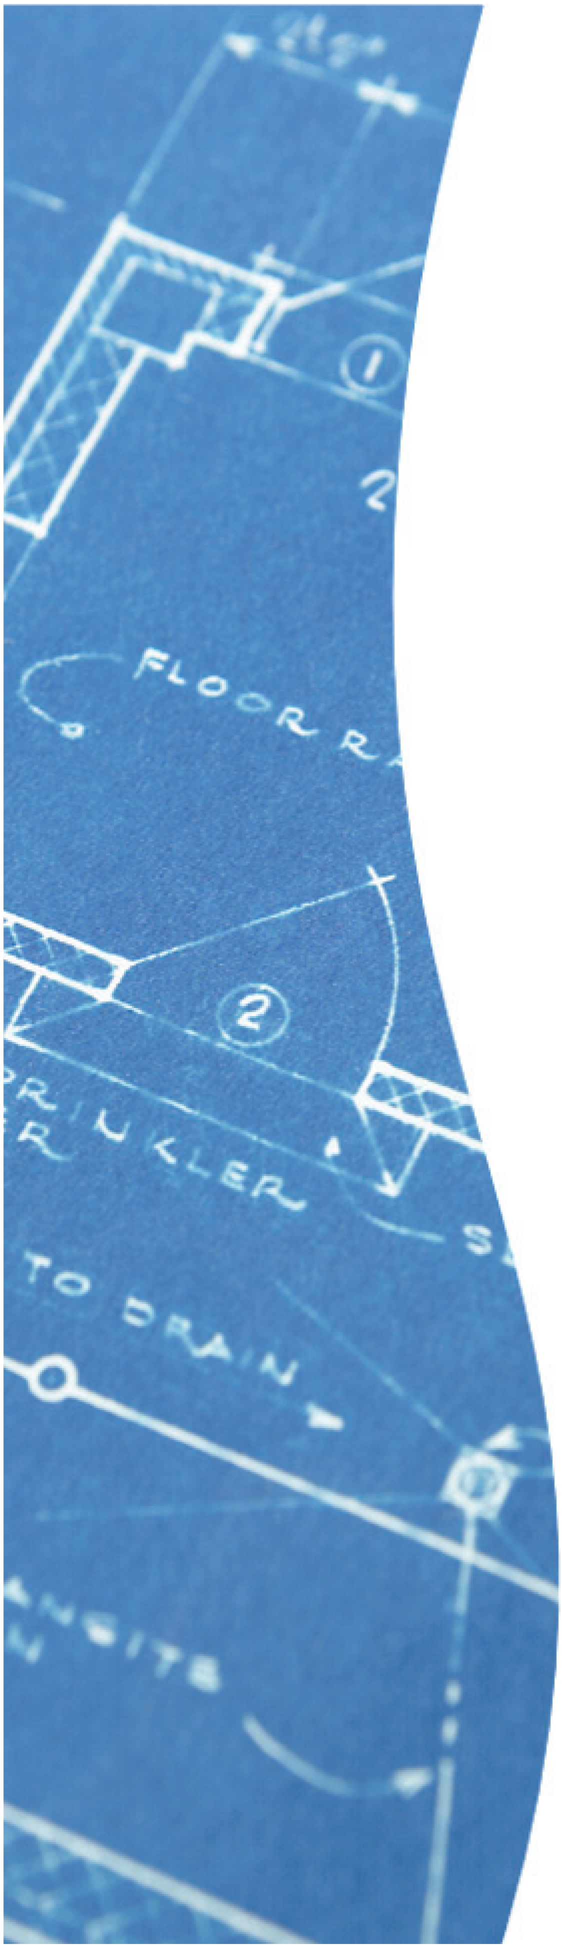

SUN'S APPROACH TO **INTELLIGENT POWER MONITORING**

Subodh Bapat
Kenny Gross
Phil Morris
Vidhya Srinivasan
Marina Thiry

Sun BluePrints™ Online

Part No 820-5826-10
Revision 1.0, 8/5/08

Table of Contents

Chapter 1
Introduction

Never before in the history of computing has so much attention been paid to the sheer cost of running systems. Over the last fifteen years, information technology (IT) promised to catapult businesses across geographical boundaries, accelerate operations to Web speed, and entice consumers to move from passive point-and-click browsing to active participation in civic, social, and consumer Net events. To capitalize on this potential, datacenters emerged around the globe—with huge server farms enabling enterprises and online communities to thrive and grow. As these promises were realized datacenters flourished, expanding to capacity and consuming more and more power per square foot. All of this success has come at a price.

Today many enterprises are faced with a lack of leeway in the datacenter. Facilities are physically constrained and there is no room left to expand. In many regions, utility companies are at their limit—there is not a single Watt more to add to the datacenter bill of fare. Moreover, the cost of energy continues to skyrocket. Over half of every dollar spent on a new server goes toward power and cooling. Indeed, the cost of keeping a volume server powered over its service lifetime often exceeds its purchase price. Equally formidable, social and government pressures are increasingly holding companies accountable for the carbon dioxide (CO_2) and pollution emitted by the generation of power that companies use. Consequently, businesses are now focused on reducing their power bill, economizing power consumption throughout the datacenter, and managing and reporting on their carbon footprint as a measure of eco-responsibility.

In response to this need, Sun is pioneering a number of advanced technologies and services that can help companies to more effectively measure, analyze, manage, and report on power consumption. This Sun BluePrints™ article introduces the key principles driving innovation in this area, examines Sun's approach to intelligent power monitoring, and highlights other power management technologies that are in development across Sun's product portfolio.

Chapter 2
Sun's Approach to Intelligent Power Monitoring

Whether looking to lower energy costs, optimize power usage for capacity expansion, or both, effectively managing datacenter power consumption requires understanding where and at what rate power is being used, and the efficiency levels across datacenter assets. The more accurate the available data, the more efficiently companies can manage the current state and reliably forecast resources for the future. In addition, understanding the relationship between the types of power-related data is essential to taking steps to economize. Finally, the insight gathered from these tasks can help answer the nagging question of whether green IT investments are being effectively leveraged to return a real energy savings.

Accurate measurement and monitoring are fundamental to gaining control of power usage. However, not every measurement technique is sufficient. Several common approaches fall short of providing adequate insight. For example:

- System name plate power levels typically are based on worst-case scenarios, such as running the server with maximum theoretical demand levels for every internal component in extreme physical environments. Power requirements that are identified for extreme cases are often grossly in excess of what is suitable for most environments.
- Power calculators use look-up tables that may not account for a particular configuration or rely on users to estimate workloads. While power calculators are likely to produce a more accurate assessment than the system name plate power, they provide a static estimation and cannot account for dynamic variations in workloads for CPUs, memory, disk drives, and fan speeds.
- Handheld power meters provide accurate readings—for that particular moment in time. Traditional datacenter assessments require personnel to take readings from each system. While accurate, these readings are limited in use and require a significant investment in time and labor. While non-handheld external power meters, such as clamp-on and networked external power meters, reduce the recording labor, installation is cumbersome and expensive.
- Typically hardware power distribution units (PDUs) are installed throughout the datacenter with multiple servers plugged into each PDU. Most metered PDUs allow the aggregate power of all equipment plugged into the PDU to be monitored. However, few PDUs are metered at the socket level, and most do not provide data on the power consumption of individual servers. Furthermore, even socket-level metering PDUs cannot distinguish between servers, switches, and storage devices in terms of the loads in the rack that draw power.

Table 2-1. Advantages of dynamic power monitoring for servers

		Advantages				
		Realistic Power Draw	Continuous Power Draw Ratings	Ease of Installation	Server-Level Power Information	Correlation of Power with System and Environment Parameters
Power Monitoring Methods	Nameplate Rating	X	X	N/A	√	X
	Online Power Calculators	X	X	N/A	√	X
	Handheld Power Meters	√	X	X	√	X
	Networked External Power Meters	√	√	X	√	X
	Rack-Level Power Monitoring (Metered PDUs)	√	√	X	X	X
	Server-Level Dynamic Power Monitoring	√	√	√	√	√

Unfortunately, static power readings provide limited value for troubleshooting power anomalies, predicting failures, and administering unexpected changes with regard to optimizing power and other resources. In contrast, dynamic power monitoring overcomes the limitations of static power readings, providing an accurate chronological record of actual power usage. However, the approach to dynamic power monitoring also impacts the benefits to be gained. Economizing power usage requires a dynamic power monitoring solution that takes into account server-level power usage and other power-related data, such as load utilization and temperature, to provide the insight necessary to confidently assess the power consumption baseline of the datacenter and manage resources more effectively.

Sun understands the criticality of optimizing every unit of energy in the datacenter. With accurate real-time power measurement, even routine tasks such as provisioning and decommissioning servers can be done in a way that yields an optimal return on power—which can help reduce energy costs or let power be distributed to another system. Since system power draw approximations obtained from measuring power at the rack or uninterruptible power supply (UPS) level include overhead, maximizing energy savings starts with dynamically assessing power consumption at the system level. In addition, the insight gained from system-level power monitoring can help identify and validate more dramatic opportunities to economize power through server refresh, consolidation, or virtualization efforts.

To help this effort, Sun designed a smart energy service called the Sun Intelligent Power Monitoring Service. Built on innovative Sun patented technologies, the service makes it possible to discern more accurate power utilization measurements—faster and at lower cost to organizations.

- *Continuous System Telemetry Harness (CSTH)*
  The CSTH systematically monitors and records a rich array of system hardware physics and performance metrics—system throughput, CPU and memory loads, distributed temperatures in the system, voltage, current, fan speed, and other operational and physical profiles—to help identify problems before they occur. It provides a comprehensive data set of system parameters that characterize the operational state, power and thermal efficiency, and the electronic health of every server under surveillance. In addition, the CSTH aids the proactive detection of subtle developing anomalies before failure occurs, provides quantitative telemetry signatures that foster quick and accurate root cause analysis of problems that impact power efficiency or availability goals, and offers continuous system operability, efficiency, and health metrics.

- *Inferential form of Real-Time Power Harness (iRTPH)*
  The iRTPH monitors server and blade total power utilization—even those containing power supplies without voltage or current sensors. It uses advanced pattern recognition from other correlated telemetry metrics to infer power draw, and returns highly accurate measurements. For example, monitoring a platform with hardware sensor readings resulted in a margin of error of +/- 18 Watts. Using the iRTPH, the same platform had a margin of error of only +/- 2.5 Watts—a 7x improvement in accuracy. Because it does not rely on current or voltage sensors, the iRTPH affords the added advantage of covering both legacy platforms and those from other vendors, many of which lack these sensors.

The combination of CSTH and iRTPH delivers powerful technology differentiation that fosters a high degree of power optimization yielding better business results. These innovations support substantial energy savings by providing accurate, real-time 3D temperature profiles throughout server interiors. The conventional approach is to estimate CPU temperatures inside servers based on external ambient temperature, computational fluid dynamics modeling, worst-case assumptions about load levels, and extreme environmental conditions, such as high altitude or humidity. This conservative approach results in significant over-provisioning of cooling for most servers.

In comparison, CSTH and iRTPH provide continuous temperature readings at many interior points throughout the server. For example, a 4U Sun server simultaneously samples 60 internal temperatures throughout the regions containing CPUs, FBDIMMs, ASICs, hard disks, and power supplies. The Sun Intelligent Power Monitoring Service is the first in the industry to intelligently distribute internal cooling to precisely match the desired temperatures of system components. It minimizes direct system power and cooling requirements, including fan speeds and acoustics, without impacting reliability.

Chapter 3
Sun Intelligent Power Monitoring Service

The Sun Intelligent Power Monitoring (IPM) Service is a network service for effectively managing datacenter power consumption by making more efficient use of power draw, cooling requirements, and load utilization for servers and storage systems. A Web-based interface provides access to monitoring, reporting, and forecasting capabilities that provide direct visibility into system power consumption and other power-related data. With this information, organizations can efficiently manage datacenter power bills and resources, effectively plan for capacity expansion, and meet emerging regulatory requirements for energy reporting. Furthermore, the underlying technologies of the service foster a high degree of power optimization, yielding better business results. Table 3-1 lists the key features and benefits of the Sun IPM Service.

Table 3-1. Key features and benefits of the Sun IPM Service

Feature	Function	Benefit
• View power consumption, temperature, and load utilization by system or group	• Measure and analyze energy consumption of heterogeneous server and storage devices by identifying systems with maximum and minimum values, highest deltas, and faults	• Understand the power consumption profile of the datacenter to establish a baseline and then determine how to optimize
• Generate trend reports using charts or by exporting raw data to comma separated value (CSV), Microsoft Excel, or XML format	• Correlate power draw and temperature utilization patterns for one or more systems, and accurately forecast energy and cooling requirements based on actual power draw	• More effectively plan for capacity expansion or unexpected changes and take advantage of energy rebates
• Compare readings between two or more systems	• Measure energy consumption by time period, such as hh:mm, day, week, month, quarter, and year	• Know whether eco-efficiency investments return a real energy savings
• Customize groups of servers to view and report by rack, location, system type, system name, or cost center	• Organize and sort systems logically for administration tasks	• Understand power profiles by group, support chargeback billing by group, and help ensure the datacenter is being used efficiently
• Receive alerts and escalations for power spikes, thermal anomalies, and other power-related irregularities	• Provide proactive anomaly detection	• Proactively detect faults and determine the root cause of power and thermal events before problems occur
• Use power consumption data to determine optimal placement of virtual machine (VM) workloads on physical resources	• Provide intelligent workload provisioning	• Help ensure that server assets are being used at optimum efficiency

Technical Design and Architecture

The following sections highlight the technical approach used to implement and deliver the Sun Intelligent Power Monitoring Service, and provide a detailed view of the logical and core components in its technical design and implementation.

Primary Principles Behind the Design Approach

- *Ease of adoption and procurement*—The Sun IPM Service is a remote network service that is accessible via a Sun portal. It provides a secure, cost-effective way for organizations to monitor and manage services—without upfront investments in infrastructure or software licensing. Using a Software as a Service (SaaS) model, organizations can provide remote monitoring and management services within a matter of days, rather than the weeks or months required with typical off-the-shelf solutions or those requiring extensive hardware modifications to enable system-level power and temperature monitoring.
- *Data privacy and security*—The security strategy is designed with multiple layers of encryption on the customer datacenter and Sun hosted infrastructure to keep power consumption, thermal flux, energy efficiency, and other power-related data private.
- *Non-intrusive*—The Sun IPM Service software that resides on the host management system in the datacenter is architected to be agent-less. As a result, the monitored system is not burdened with a continuously running process that might cause concern about resource consumption.
- *Flexibility*—The Sun IPM Service can operate independently, or the software can be integrated with system management software, such as Sun™ xVM Ops Center, IBM Tivoli, or HP OpenView, simplifying deployment.

How It Works

Figure 3-1 illustrates how the Sun IPM Service works.

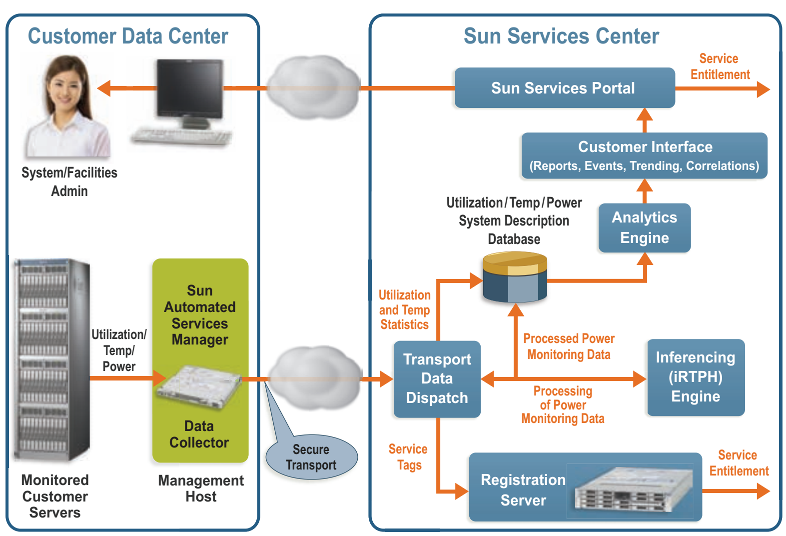

Figure 3-1. How the Sun IPM Service works

Core Components

Staying competitive in today's business climate means monitoring and managing the energy that powers the datacenter to help ensure optimal performance and economies of scale. With hundreds or thousands of IT systems consuming power, automating this process is essential. The system infrastructure that comprises the Sun Intelligent Power Monitoring Service is built to provide data security in each core component as explained below.

Customer Datacenter Infrastructure

The Sun Automated Services Manager and the IPM Data Collector comprise the software that supports the customer datacenter infrastructure.

Sun Automated Services Manager

Based on the security inherent in the Solaris™ 10 Operating System (OS), the Sun Automated Service Manager is responsible for receiving telemetry time-series signatures from monitored assets and delivering them to the Sun IPM Service core infrastructure, as well as acting as a gatekeeper for incoming remote management sessions. A flexible deployment model lets multiple instances be deployed to meet the needs of a single organization to address size or geographic diversity concerns. Upon initialization, the system must be registered with the Sun IPM Service core infrastructure, and a private/public encryption key exchange performed. These 1024-bit Rivest, Shamir, and Adelman (RSA) keys are used for signing all future messages into or out of the Sun Automated Service Manager.

The Sun Automated Service Manager can be deployed in a number of ways within an organization to meet security compliance requirements—within the end-customer's demilitarized zone (DMZ), within a trusted network, or even outside all of a customer's perimeter firewalls. This flexibility allows telemetry reception within the datacenter environment to comply with the organization's security policies and other governance requirements.

IPM Data Collector

The IPM Data Collector collects one or more data sources associated with each monitored system. For this power monitoring service, the data source is comprised of environmental data, such as temperature, voltage, and current. A variety of data sources are available. For example, Sun Fire™ 6800 servers include system controllers that can be connected to via `telnet` in order to obtain environmental data from various sensors throughout the system. Many Sun x64 servers incorporate Intelligent Platform Management Interface (IPMI) compliant Integrated Lights Out Managers (ILOMs) that allow the `ipmitool` utility to gain access to similar data. Other Sun systems have various Lights Out Managers that provide similar services. In addition, a system shell login is also a data source.

Sun Hosted Infrastructure

The IPM Data Processing Service and IPM Portal Tier give organizations insight into the power consumption profile of their datacenter.

IPM Data Processing Service

The IPM Data Processing Service collects parametric telemetry data from the target customer’s system, assembles it as XML, and securely transports the information to Sun’s endpoint. The parametric data endpoint resides within Sun's datacenter for security, processing, and temporary storage within the transport queue. Using authentication, the transported data objects are dispatched from Sun's transport endpoint, formatted, submitted, and processed by Sun's patented iRTPH software processor. The iRTPH algorithm is designed to process and analyze the raw parametric data, producing extremely accurate power values that reflect the same level of accuracy one would obtain with an expensive attached power meter reference device. Power, individual CPU utilization, and target system internal component temperature signals are stored in an Oracle database for later retrieval for accurate trend data reporting by the customer.

IPM Portal Tier

A Web portal provides the primary human interface for the Sun IPM Service. The IPM Portal Tier provides the power consumption profile for systems at a group, individual system, or sub-system level. The connection points are secured via the 128-bit HTTPS protocol, and all access requests are authenticated in real time against Sun Online℠ service accounts for access. If the credentials presented are authenticated, the request is tested against the authorization models currently in the system to verify the user or system is authorized to perform the submitted request.

Getting Started

The software is quick and easy to install. The following steps explain how to install the software package and activate the service on systems.

1. Enable the Sun Intelligent Power Monitoring Service in a datacenter by downloading the software package from *http://power.sun.com*. Log in using a Sun Online account and accept the evaluation license agreement.
2. Install, configure, and activate the service on the system. The installation script is interactive, walking users through the service installation and configuration process. Once installed, systems that need to be monitored are identified and the service is activated. A README file is provided to clarify or answer questions as part of the package.
3. View reports on the Sun Network Services portal. Once the Sun Intelligent Power Monitoring Service is activated, organizations can monitor power consumption for systems in the Sun Network Services Portal. Users can login with a Sun Online account and view power consumption reports for the systems that are enabled.

Interpreting Reports

Several reports are provided to help users view and analyze information.

Power Consumption Dashboard

The service supports the creation of a dashboard that represents status at a group level. Up to five groups can be configured for viewing. Status information is based on Sun recommended thresholds for each system type. Figure 3-2 shows a typical dashboard.

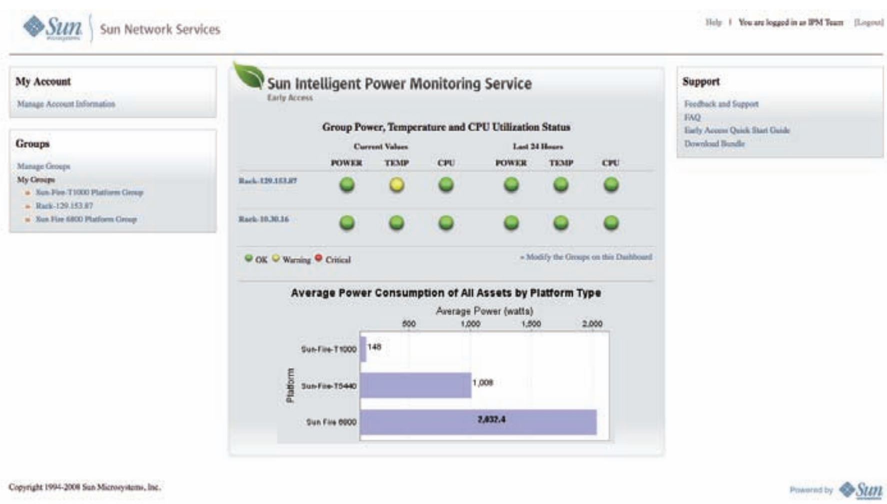

Figure 3-2. An example Sun IPM Service dashboard viewed through the Sun Network Services portal

Creating Groups

Groups can be created based on organizational needs (Figure 3-3). Pre-defined groups are created by the service based on subnet (three octets) and system type.

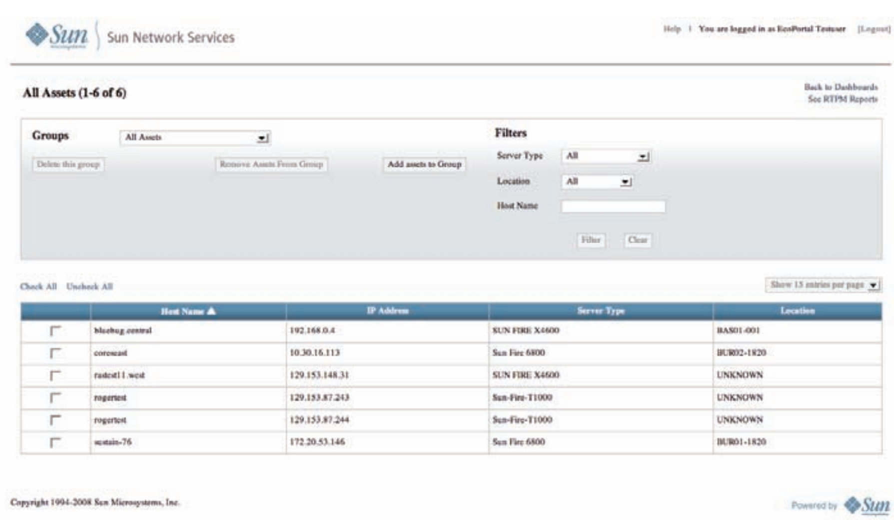

Figure 3-3. Creating groups in the Sun IPM Service

Reports

Major reporting components include:

- *Power consumption, load utilization, and temperature graphs*
  These graphs can be viewed at a group, system, or system component level for certain types of systems. Group-level reports help customers with capacity planning activities, such as taking systems off a rack based on load requirements. In addition, companies can use this information for chargeback billing to departments for power consumed.

- *Data points*
  The default view of reports shows the average reading for temperature, power, and load. Users also can view other statistics in the graph, including maximum, minimum, range, and median readings.

- *Trend data*
  Trend reporting is available for any duration specified. In addition, the current reading is available and displayed alongside the trend reports.

- *Power, temperature, and load correlation*
  Viewing the graphs together can help with analysis. A spike in power might have a corresponding spike in load utilization. This can help identify potential cooling requirements in the datacenter.

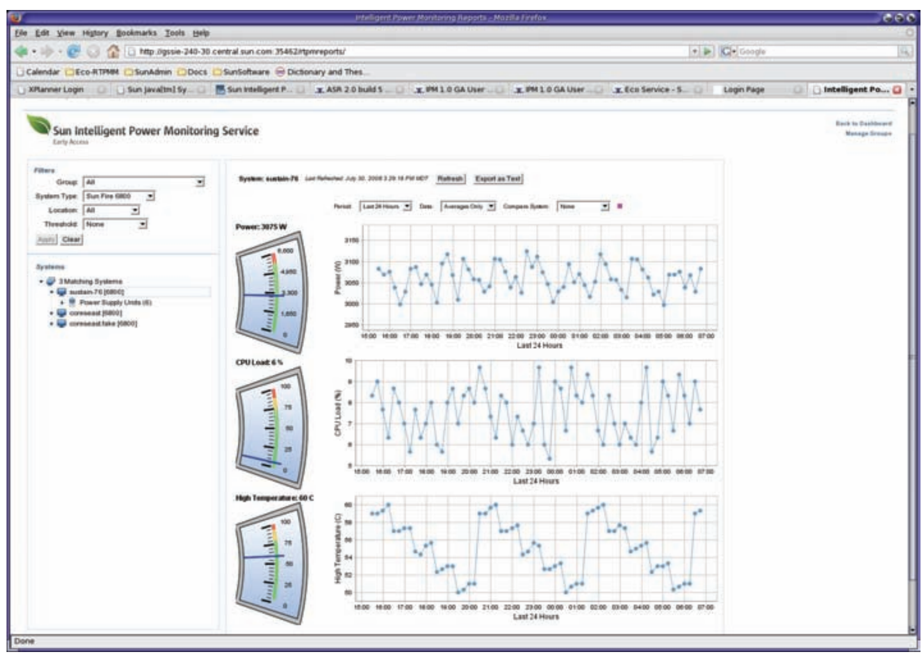

Figure 3-4. Power, temperature, and load trend report

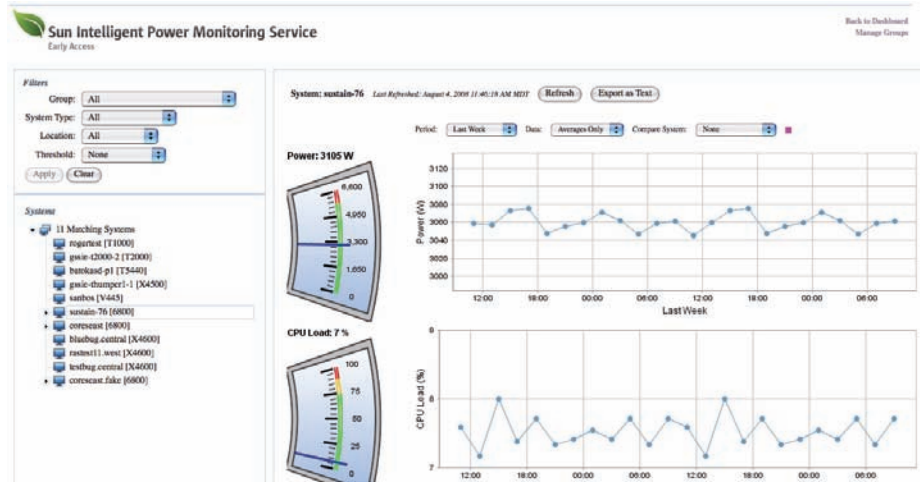

Figure 3-5. Understanding the relationship between power and load utilization helps support workload adjustments

Chapter 4
Use Case Scenarios

The following real-world scenarios exemplify the value of the Sun Intelligent Power Monitoring Service. Figure 4-1 illustrates how information is used to economize power consumption in the datacenter.

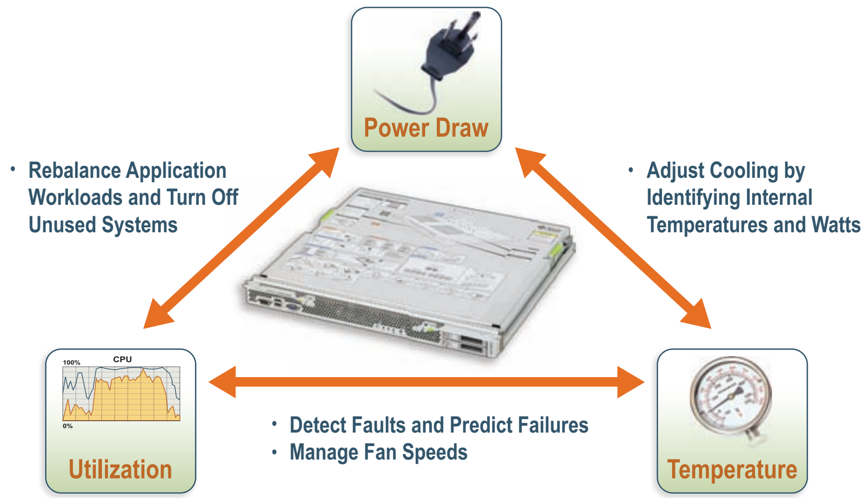

Figure 4-1. Understanding the relationship between power draw, temperature, and load utilization helps measure, analyze, and manage power consumption at optimum levels

Example 1—Reducing the Power Bill

The CIO of a large manufacturing enterprise is tasked with reducing the datacenter electricity bill by 25 percent. The datacenter is comprised of hundreds of heterogeneous systems, most of which are older models. Without power monitoring capabilities, and just a big power bill available for perusal, the CIO has no idea where to even start.

Using the Sun Intelligent Power Monitoring Service the CIO obtained an accurate, seasonal power consumption profile of the datacenter and learned where power was being consumed, the rates of consumption, and the efficiency levels across datacenter assets. By establishing a baseline profile, groups of servers identified as inefficient were chosen for server refresh and consolidation efforts, allowing many older servers to be replaced with fewer newer servers. Commensurate savings were also realized in the cooling system due to reduced air conditioning load. In addition, the service helped the CIO to employ a chargeback billing solution to further defer operating costs.

By using the Sun Intelligent Power Monitoring Service and Sun Professional Services assessment and optimization services, the company was able to:

- Reduce the datacenter electricity bill by over 30 percent
- Optimize power consumption based on peak trends and server utilization
- Economize power usage according to seasonal utilization levels

Example 2 — Going Green

A large financial institution set a goal to become the greenest bank in the world. Many tactical projects are in place to reduce energy consumption and purchase clean energy. Circuit-level power monitoring is available, however the company has no idea at the system level about how to identify and reconcile the worst power offenders. In addition, three corporate datacenters are located in large urban areas where electricity and land prices are extremely high. Yet the organization cannot leave the area due to applications with massive latency issues.

Economizing power usage requires a dynamic power monitoring solution that takes into account server-level power draw and other power-related data to provide the insight necessary to manage resources more effectively. With the Sun Intelligent Power Monitoring Service, the company was able to measure and compare energy consumption between systems, and groups of systems, and identify and reconcile the worst power offenders. With this insight, the bank was able to simply turn off approximately 10 percent of datacenter servers that remained powered on even though business units stopped using the applications running on them.

Next, the bank identified servers that consumed significant power but had relatively low utilization. These physical servers were replaced with virtual servers and consolidated onto fewer systems supporting virtualization. In addition, the company was able to correlate measured latency on each system with measured power draw, thereby allowing the organization to standardize on systems with the least latency per power draw.

By using the Sun Intelligent Power Monitoring Service the company was able to:

- Validate that each physical server runs with high utilization levels and support the claim of being the greenest bank
- Achieve an average ratio of four virtual servers on each physical server in the datacenter
- Make more effective decisions on how to replace leased systems
- Significantly reduce power consumption by identifying and replacing power hungry systems with newer, energy-efficient servers

Example 3 — Economizing Under Constraints

A power company told a U.S. telecommunications provider that it cannot supply them with more power, as the company is the largest energy consumer in the region. The telecommunications provider turned to Sun to help analyze current power consumption and efficiency levels, and determine how to make optimal use of the available power — without disrupting the customer base and still support growth. The existing datacenter includes thousands of Sun systems and approximately four times as many platforms from other vendors.

Sun understands the importance of optimizing every unit of power in the datacenter. With the Sun Intelligent Power Monitoring Service, the telecommunications provider was able to accurately characterize the power draw and utilization of all servers. Since traffic such as voice calls and Short Message Service (SMS) text messages has predictable utilization patterns over each 24-hour period, the service helped identify periods when server capacity is at its highest and lowest. By adjusting server capacity during periods of low payload traffic, and temporarily repurposing spare server capacity each night to non-real-time tasks such as bill generation, the telecommunications provider was able to use their limited power more efficiently. As a result, the company was able to:

- Free 2 megawatts of power for use for business growth
- Save capital expenditures by dramatically increasing the life of the datacenter
- Save operating expenditures (OPEX) that would have been spent on a co-location facility for additional space

Chapter 5

Other Power Monitoring and Management Technologies from Sun

Principles

Sun's eco datacenter vision foresees the emergence of ultra-intelligent datacenters that are self-aware, fully automated, and capable of dynamically monitoring and fine-tuning energy flows in a manner that is calibrated to computational demands. Sun has articulated three key principles of energy-efficient datacenters.

- *The principle of totality*
  An energy-efficient datacenter has every component appropriately harnessed to deliver the most energy-efficient operation possible. The corollary to this principle is that datacenter energy efficiency is not isolated to a chip, server, hypervisor, operating system, application, network, storage, PDU, UPS, air conditioning, or chiller. All need to be involved.
- *The principle of agility*
  An energy-efficient datacenter must be responsive to the dynamics of all its inputs. A datacenter must be able to dynamically adjust its energy consumption in response to:
  - Fluctuating demand inputs, such as network traffic and compute workloads
  - Fluctuating supply inputs, such as changing grid power availability, utility brownouts, and electricity price differential signals
  - Changing IT requirements, such as corporate policies around total OPEX and carbon footprint targets
- *The principle of proportionality*
  An energy-efficient datacenter calibrates its energy consumption in proportion to the computational work it is expected to deliver. The corollary to this principle is that it is not just individual servers or storage devices that must calibrate their power draw to current compute workloads. The entire datacenter must do so. For example, a datacenter containing only a Web farm is not fully utilized when there is no external HTTP traffic coming into the datacenter, and therefore must scale back its energy usage to a minimum base level.

Power Management Innovations

Sun continues to work on a number of advanced technologies to help companies efficiently utilize datacenter infrastructure. Sun innovations driving these principles include:

- *Energy-proportional computing*
  A server must be able to calibrate its power consumption to the computational load it is currently experiencing, thereby only using energy proportional to the computational work it performs. Many of today's servers burn nearly the same amount of energy whether they are 10 percent or 90 percent utilized. In contrast, an energy-proportional server recognizes changes in its utilization level and adjusts its energy consumption appropriately.
- *Comprehensive energy management*
  Server energy efficiency requires all layers in the stack to be involved to run the server at an optimal energy level. Sun's approach to power monitoring is to deliver energy-efficient, innovative technology at every level—not just in individual components or systems.
- *Energy management by policy*
  Sun continues to work on server technologies designed to let organizations express specific server energy management policies applied to individual systems, racks, aisles, or datacenters. Energy policies include power capping, scheduled or opportunistic reduction of computational capacity in response to fluctuating workloads, transition to quiescent states, and more.

Sun Projects and Technologies

Sun intends to provide energy-efficiency innovation in multiple projects and technologies across its product portfolio.

- *Power policy modes*
  Sun continues to work on designing platforms that let administrators and datacenter operators express specific power directives on systems, including Watts ceilings, various sleep states, and self-calibration of power draw in response to changing utilization, temperature, and supply constraint conditions. System firmware APIs that can be exercised for this purpose by administrators and high-level datacenter management software continue to be developed.
- *Microprocessors*
  Sun continues to invest in multiple microprocessor power management capabilities, including disablement and power throttling of individual cores, voltage scaling, and frequency scaling.

- *System devices*
  Sun continues to make systems smarter—capable of calibrating their power draw or shutting down in response to utilization levels. Power-managed devices include memory DIMMs, disks, I/O devices, and fans.
- *Intelligent fan control*
  Today most system fans are over designed and sized for worst-case ambient temperature, humidity, and altitude conditions. As a result, a significant portion of system power is consumed by system fans. Sun continues to work on intelligent fan control algorithms designed to permit the scaling of fan power consumed for airflow in proportion to system temperature as well as ambient conditions, such as altitude, humidity, acoustics, smoke, particulates, and other environmental parameters.
- *Power supplies*
  Sun's power supplies are on a trajectory to exceed industry standards set by consortia such as Climate Savers and 80plus.org for the relevant class of systems. In addition, Sun continues to work on power supplies with unique technological innovation that maintain optimum power supply efficiency at low load levels and under fluctuating utilization conditions.
- *Virtualization*
  Virtualization management is an important element of datacenter power management. Sun xVM software provides several key capabilities, including letting virtual machines read resource utilization levels in a physical system, enforcing resource utilization limits, and migrating virtual machines to run the workload in the most efficient location in the datacenter.
- *Operating systems*
  The OpenSolaris™ project continues to work on power management, including the ability to reduce unnecessary process wake-ups through clock interrupts, power-aware dispatching of workloads, power managing microprocessors according to policy, and exposing APIs for applications to run as efficiently as possible.
- *Systems management*
  Sun's systems management activities include furthering the ability to power manage individual servers and racks. At the datacenter level, Sun xVM Ops Center software continues to be refined to support the migration of virtualized workloads and locate virtual machines on servers that are most appropriate to run them in accordance with availability, reliability, and energy efficiency policies.

- *Professional services*
  Sun’s datacenter professional services include energy efficiency considerations in every aspect of service delivery, from simple virtualization, consolidation, migration, and server refresh projects, to large-scale datacenter design, layout, and optimization efforts. Sun's professional services competencies include system, network, and storage architectures, as well as datacenter space, power, and cooling architectures.

Sun's comprehensive approach to datacenter energy efficiency encompasses semiconductor devices, system devices, firmware, operating systems, and systems management software. While all these technologies help Sun servers to consume less energy and behave responsibly as good datacenter citizens, power monitoring technology provides deeper insight to hone power management effectiveness and demonstrates the power savings achieved. The combination of energy-efficient products and intelligent power monitoring can help companies to economize energy usage throughout the datacenter, and tie into IT chargeback billing solutions, compliance reporting tools, and carbon footprint assessment techniques.

Chapter 6
For More Information

Many organizations can benefit from more efficient use of energy in the datacenter. The Sun Intelligent Power Monitoring Service provides an innovative approach to economizing energy consumption while preserving infrastructure service-level expectations. To learn more, *visit http://www.sun.com/service/power*. Be sure to investigate the try-and-buy program and experience some of the capabilities free of charge for 60 days. In addition, Sun sales representatives are available to answer questions about how to use intelligent power monitoring to reduce energy costs while optimizing datacenter investments.

About the Authors

Subodh Bapat is a Distinguished Engineer and the Vice President for Energy Efficiency at Sun Microsystems. He is responsible for coordinating Sun's systems-level energy efficiency technologies cross-divisionally across Sun's product portfolio, from the microprocessor level to the datacenter level. Subodh previously served as VP and Chief Technology Officer of Sun's $4B Volume Systems division, and led a number of acquisitions for the company. Subodh is a published author, frequent industry speaker, and a holder of multiple patents.

Kenny Gross is a Distinguished Engineer at Sun Microsystems, and is team leader for the System Dynamics Characterization and Control team. Kenny specializes in advanced pattern recognition, continuous system telemetry, and dynamical system characterization for improving the reliability, manageability, and energy efficiency of enterprise computing systems. Kenny has 202 U.S. patents issued and pending, 169 scientific publications, and was awarded a 1998 R&D 100 Award for one of the top 100 technological innovations of that year, for an advanced statistical pattern recognition technique (MSET) that is now being used for a variety of applications to optimize energy efficiency and reliability for enterprise computer servers.

Phil Morris is a Technical Director and Chief Technologist at Sun Microsystems and has won the President's Award for his contributions to the Sun/SAP Competency Centers and Sun Professional Services SAP Practice. Phil's current specialty is around eco datacenters, including facilities planning, architecture design, cooling and airflow analysis, and maximizing power efficiencies utilizing consolidation, virtualization, and overall power savings and reduction of the carbon footprint. He has been an acknowledged leader within the company for over eight years.

Sri “Vidhya” Srinivasan is a Principal and Engineering Manager at Sun Microsystems, and is responsible for the delivery of innovative network services. She works closely with services marketing and services readiness groups to rapidly prototype and deliver service solutions that are of value to customers. Since joining Sun in 2000 as an engineer, Vidhya has played a variety of different roles including program management and Business Architect. Prior to Sun, Vidhya was a consultant for the Interlink Group in Denver, Colorado.

Marina Thiry, a Principal with Services Marketing at Sun Microsystems, is primarily responsible for building revenue opportunities by identifying, evaluating, and creating innovative services. She works closely with customers and technologists to create targeted business solutions from broadly defined technologies, then develops and implements business plans to enter new markets. Previously, Marina was the Vice President of Content at Quaartz, an Internet start-up acquired in 2000 by a NASDAQ company. Prior to that, she spent 16 years producing television and radio programs in Los Angeles and interactive television and multimedia content in Silicon Valley.

References

Sun Intelligent Power Monitoring Service:
`http://sun.com/service/power`

Sun Servers:
`http://sun.com/servers`

Sun Services:
`http://sun.com/service`

Sun xVM Software:
`http://sun.com/xvm`

Ordering Sun Documents

The SunDocs[SM] program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

Accessing Sun Documentation Online

The `docs.sun.com` Web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is
`http://docs.sun.com/`

To reference Sun BluePrints Online articles, visit the Sun BluePrints Online Web site at:
`http://www.sun.com/blueprints/online.html`

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN (9786) **Web** sun.com

Printed in USA 08/08