


Preserving Performance While Saving Power Using Intel® Intelligent Power Node Manager and Intel® Data Center Manager

Deploying Intel® Intelligent Power Node Manager and Intel® Data Center Manager with a proper power policy can decrease server and data center power usage without impacting performance



BMW GROUP

David Helman
BMW Group

Matthew Kelly
BMW Group

Martin Guttman
Intel Corporation

Additional Contributors

Bill Carter
Enrique Castro-Leon
Thomas Kellerer
Dror Shenkar
Intel Corporation

EXECUTIVE SUMMARY

Intel Intelligent Power Node Manager and Intel® Data Center Manager (Intel® DCM) are platform-software-based features that reside on the Intel® Xeon® processor 5500 series, providing power, thermal monitoring, and policy-based power management for individual servers, racks of servers, and/or data centers.

To determine how much Intel Intelligent Power Node Manager and Intel DCM can decrease power consumption in a data center *without* impacting performance, BMW Group and Intel conducted a technical proof of concept (PoC) and series of tests. This paper describes the procedures used and the results from testing Intel Intelligent Power Node Manager and Intel DCM with data from a live BMW Group workload.

In the BMW Group PoC, the functional capabilities of Intel Xeon processor 5500 series-based systems with Intel Intelligent Power Node Manager and Intel DCM made it possible to lower server power consumption by 18 percent and increase server efficiency by approximately 19 percent. In designing data center infrastructure to effectively use the capabilities of Intel Intelligent Power Node Manager and Intel DCM, the results show greater overall server power savings, increased server rack utilization, reduced IT infrastructure costs, and lower overall data center power consumption.

Business Challenge: Decreasing Data Center Power Costs Without Compromising Performance

Key goals for BMW Group are to minimize the company's impact on the environment, reduce data center power consumption, and increase density as required at the system, rack, and data center levels—all without compromising the performance of business solutions and applications. Data centers can use Intel

Intelligent Power Node Manager in conjunction with their operating system to control the operating frequency (i.e., P-states) of a processor by adjusting the processor's clock speed on the fly. This frequency adjustment results in dynamic changes to the power consumption of a server, rack of servers, or data center. This capability helps companies like BMW Group reach their business goals and also minimize their impact on the environment.

Table of Contents

Executive Summary1
Business Challenge: Decreasing Data Center Power Costs Without Compromising Performance1
Intel Intelligent Power Node Manager2
Intel Data Center Manager3
Test System Configuration4
PoC Testing Approach and Use Cases5
PoC Results7
Conclusions7

To prove the value of Intel Intelligent Power Node Manager, Intel and BMW Group set up a series of experiments using Intel® Server Boards with integrated Intel Intelligent Power Node Manager and Intel DCM to simulate the real workload of a BMW Group server. The goal was to generate power policies with Intel Intelligent Power Node Manager that will decrease BMW Group’s power usage without affecting the system performance.

The rising cost of power remains the highest priority for today’s data centers. Virtualization has helped to reduce power consumption by eliminating under-used servers, but virtualization by itself cannot solve the increasingly large power demands of data centers.

In computing, there are conflicting tradeoffs between processing speed and cost. In data centers, there are conflicting tradeoffs between power usage and performance. Service level agreements (SLAs) are paramount in modern data centers, since the cost of violating them—even for a few hours—can have a significant impact on business. Because of this, power consumption policies have often given way to performance concerns.

IT organizations facing power and cooling challenges can benefit from:

- **Increased rack density** within space, power, and cooling constraints through fine-grained power control.
- **Reduced capital costs**, a benefit of right-sizing power and cooling infrastructure based on actual power consumption trends.
- **Reduced operations costs**, a benefit of eliminating worst-case headroom buildup during rack power provisioning.

BMW GROUP

As a global player, the BMW Group is represented through its premium products—the BMW*, MINI*, and Rolls-Royce* brands—in more than 140 countries the world over. A flexible network of 23 production plants in 12 countries ensures that all customers receive exactly the car they ordered, tailored to their specific wishes and preferences.

The BMW Group’s production network operates hand-in-hand with the Group’s international research and development centers as well as with external partners. This provides the option to introduce new products into the marketplace quickly and efficiently, adjusting production without delays due to specific customer wishes and requirements.

Along with its automotive concerns, the BMW Group’s activities include developing, producing, and marketing motorcycles as well as comprehensive financial services for private and business customers.

(Power is provisioned to servers based on workload SLAs.)

Intel Intelligent Power Node Manager
 Intel Intelligent Power Node Manager is a power management engine embedded in Intel® server chipsets. Processors carry the capability to regulate their power consumption by manipulating the processors’ internal power states. Intel Intelligent Power

Node Manager works with servers' hardware and software and the operating system power management (OSPM) to perform this manipulation and dynamically adjust platform power to achieve maximum performance and power for a single node (Figure 1).

Intel Intelligent Power Node Manager features include:

- **Dynamic power monitoring:** Measures actual power consumption of a server platform. Intel Intelligent Power Node Manager gathers information from PM-BUS instrumented power supplies, provides real-time power consumption data in single-time increments or as a time series, and reports through the IPMI interface.
- **Platform power capping:** Sets platform power to a targeted power budget while maintaining maximum performance for the given power level. Intel Intelligent Power Node Manager receives power policy from

an external management console through the IPMI interface and maintains power at the targeted level by dynamically adjusting CPU P-states.

- **Power threshold alerting:** Intel Intelligent Power Node Manager monitors actual platform power and compares it to the target power budget. Intel Intelligent Power Node Manager directs the operating system or the hypervisor to change to a target P-state. The firmware ensures that the individual server meets the assigned power consumption target. When the target power budget cannot be maintained, Intel Intelligent Power Node Manager alerts the management console. When the management console is used, the console can change power allocations when this alert is received.

Intel Data Center Manager

Intel DCM is an SDK that provides power and thermal monitoring and management for servers, racks, and groups of servers in data

centers. Management console vendors (ISVs) and system integrators (SIs) can integrate Intel DCM into their console or command-line applications and provide high-value power management features to IT organizations.

Intel DCM:

- **Has a built-in, policy-based intelligent heuristics engine** that can maintain group power capping while dynamically adapting to changing server loads and minimizing the performance impact of workloads.
- **Uses Intel Intelligent Power Node Manager** for node power and thermal management.
- **Is designed as an SDK** to integrate into existing management software products through a Web service application programming interface (WSDL API).

Table 1 shows Intel DCM's features and benefits.

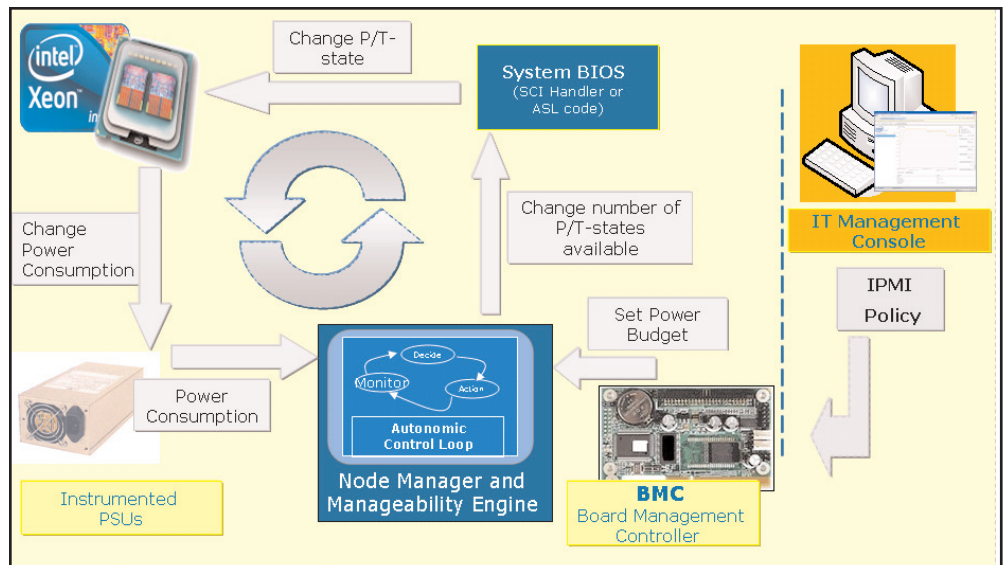


Figure 1. System and platform architecture for Intel Intelligent Power Node Manager and Intel DCM

Test System Configuration

For the PoC, the engineering team used CPU utilization data from BMW Group to represent a typical day’s workload from servers in the production grid. The utilization data, from IBM System x3850* enterprise servers with 32 GB of memory each, was averaged hourly.

These servers experienced similar loads every day, making them good candidates for the PoC and test experiment. The test servers used in the PoC were Intel® Server Boards S5520UR, each with two quad-core Intel Xeon processors X5560 series with 12 GB of memory, running Windows* Server 2008. The testing used software solutions and utilities including workload software, system, CPU, and performance utilities including a

sample DCM GUI (Figure 2). This sample graphical implementation, based on Adobe Flex* technology, demonstrates how a management console can integrate DCM SDK using its Web service APIs and use it to aggregate Intel Intelligent Power Node Manager to manage groups of servers. The GUI provides broad capabilities and functionality to set server and rack policies and monitor, track, and collect power consumption data.

Table 1. Intel DCM Features and Benefits

Flexible data center hierarchy support	Supports management simultaneously at all levels of data center hierarchy
Power and thermal data aggregation and trending	Monitors node power and inlet temperature data in real time Aggregates power and inlet temperature data Stores trend data for up to one year
Custom alerts	Provides alerts based on custom power and thermal events Automates event handling by defining policy based actions to change power threshold
Intelligent group power capping	Supports for multiple policies, depending on user’s power threshold target or goal to minimize power consumption Maintains group power capping while dynamically adapting to changing server loads Accepts SLA priority as policy directive Automatically manages rack and group power consumption and safeguards from sudden power spikes
Configurable policies	Schedules power capping policies by time of day and manage to the specifics of your data center environment. This allows IT organizations to provision power to workloads dynamically based on need rather than one-time average or worst-case setting with low administration overhead Maintains multiple policies simultaneously
Agentless out-of-band management	Does not require installation of any software agents on managed nodes
Easy integration and co-existence	Exposes Web Services Description Language (WSDL) APIs; this enables easy integration into existing management software, either commercial or home-grown Optionally resides on an independent management server or co-exists with ISV product on the same server
Scalability support	Manages up to 1,000 managed nodes
Security compliance	Ensures security by complying with IPMI 2.0 BMC authentication, integrity, and confidentiality code
Flexible partner innovation opportunities	Fosters partner innovation by enabling custom filtered reports of current or historical data for any hierarchical entity (e.g., at node, rack, and row or group level) and for any time period
Reference graphical user interface	Provides a simple and functional interface that implements all leading usage models Source code sharing enables integration into other applications and customization based on need

PoC Testing Approach and Use Cases

The PoC testing used the high-level approach outlined in Figure 3.

Pre-Deployment Phase

The pre-deployment phase focused on gathering, processing, and analyzing BMW Group system production server performance data. System CPU utilization data was collected for a full 24-hour work day. (Figure 4 shows the average hourly CPU utilization data.) This baseline data was an integral part of the PoC and testing efforts. Because the daily usage was cyclical, with periods of low CPU utilization during the evening, the team decided that one possible scenario for saving power was to create an after-hours policy. After analyzing the data, the team also identified

two other policies. The Intel Intelligent Power Node Manager usage model and power policies tested were:

1. **Informal (flat) power policy:** Select a policy that is uniformly applied across a 24-hour period.
2. **After-hours policy:** Select a policy for off-operational hours such as weekends.
3. **Planned capping policy:** Select numerous policies across the week and day.

After analyzing the potential use cases to address broader business needs for the PoC and testing, the testing applied a planned capping policy. This would make it possible to apply numerous optimum power policies throughout the day.

Power Discovery Phase

The power discovery phase is a critical effort that helps to determine and select the optimum server power settings that will not impact the performance of the systems and solutions.

Power Savings Evaluation Phase

The power savings evaluation phase is designed to execute given workloads, gather systems performance and power data, and analyze potential server and data center savings as gained from applying power policies. This was done without impacting the overall system performance.

For details on the applied approach, test methodology, and data analysis to determine the appropriate server power settings, see Figure 5.

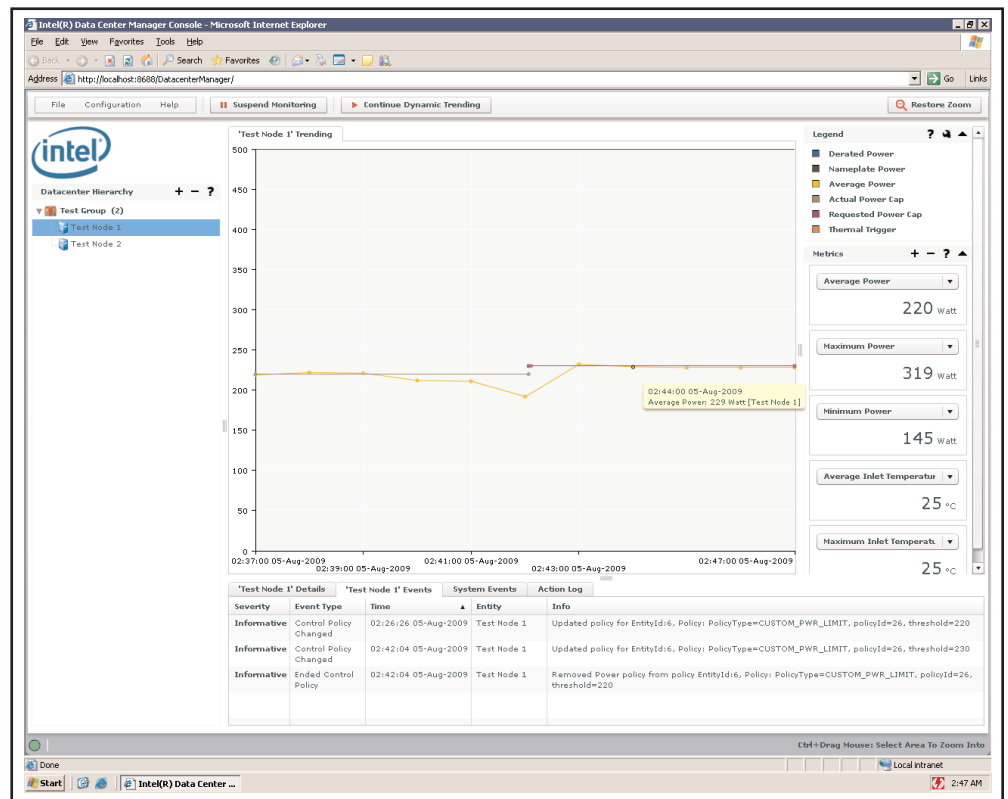


Figure 2. Intel DCM demonstration software graphical interface for the PoC

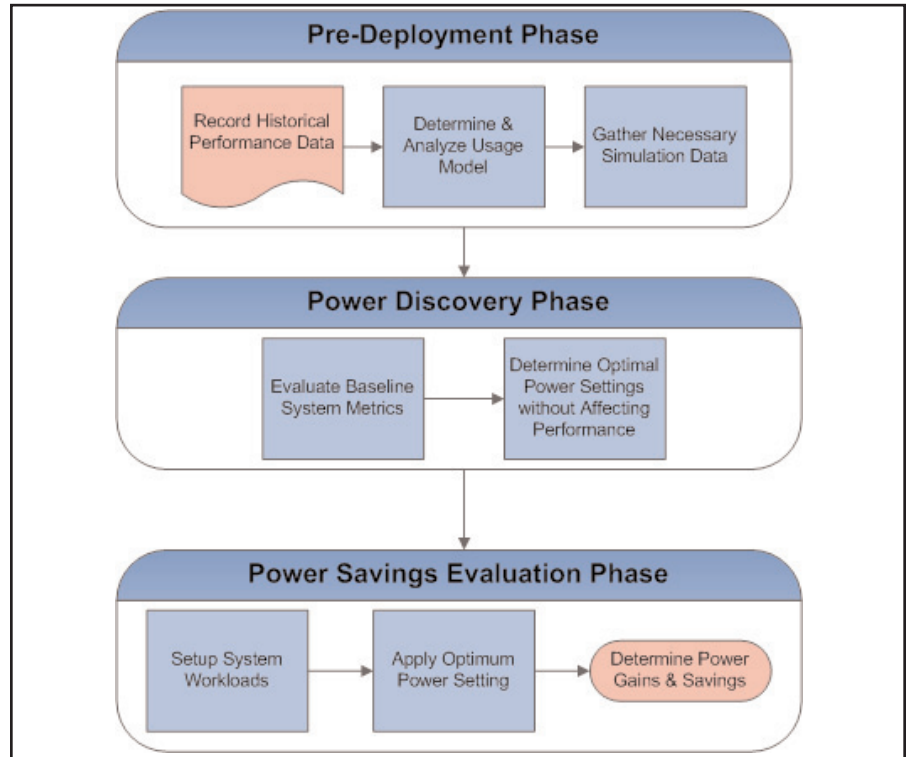


Figure 3. High-level test approach and methodology

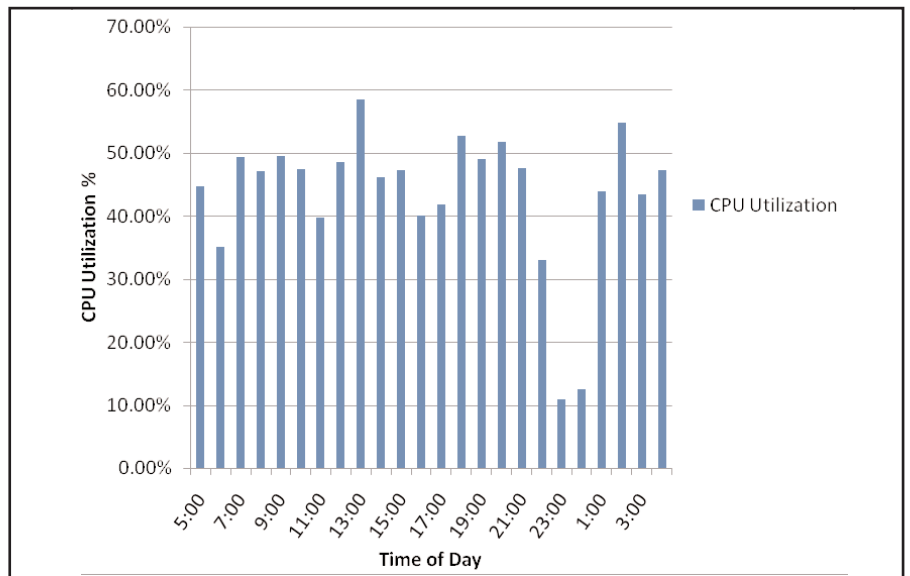


Figure 4. Collected average CPU utilization over 24 hours

The team used the processes outlined in Figure 3 and 5 to determine the optimum power policy setting that matched the performance requirements of the historical performance of the servers. Using this systematic process made it possible to arrive at optimal power settings without impacting the system performance.

To determine and gather the existing BMW Group system workloads and CPU system utilization in the PoC, the test used system monitoring and performance tools. The data was collected over a period of 24 hours and then used to calculate the hourly average CPU utilization (Figure 4).

The testing applied CPU utilization and workloads to test and gather server system performance data including:

- **Measuring server power (watts) consumption.** The base server power consumption can be used as a guide to set new power policies and to calculate potential power and data center savings with Intel Intelligent Power Node Manager and Intel DCM without impacting system performance.
- **Determining the base number of executed transactions.** The base number of transactions at specific CPU utilization is critical information since it was used to assure that system performance was not impacted due to a lower power policy, thus not impacting the overall system performance.

The next step was to vary the system power until the lowest power usage was determined that would still allow the number of instructions to be processed. The testing:

- **Applied** BMW Group CPU utilization workloads

- **Lowered** server power settings and used the Intel DCM capabilities (Figure 2)
- **Simultaneously recorded** the base number of instructions the system executed

If the decreased power setting resulted in an insufficient number of instructions executed, the team increased the power setting and then repeated the process and tested again until the resulting number of instructions processed was greater than or equal to the base number of instructions. The resulting power setting represented the ideal power conditions required to execute the same number of instructions.

When all conditions were met, the power usage result yielded an optimum power policy setting without impacting the system performance.

Once the lowest acceptable power setting was determined, the team applied the server power policies with given CPU utilization and transactions and performed and executed the system tests.

To determine the overall power saving and effect of power policies, the team:

- **Measured** the new system performance and server power consumption
- **Calculated** the server and data center power gains and cost savings

PoC Results

The goal of the PoC was to determine and apply the maximum acceptable power capping policy without impacting system performance. For each CPU's workloads and utilization, the testing modified power policy settings to a point where there would be no impact on system performance. The PoC clearly shows the business benefits of using the functional capabilities of Intel Intelligent Power Node Manager and Intel DCM.

The PoC test results (Figure 6) show the percentage of maximum power savings versus CPU utilization that will not affect system performance. During the test, simulated workloads operating with an optimum power policy show optimal power savings with 45 percent CPU utilization. Note that it is possible to increase server power savings, but applying lower power policy settings would decrease the overall system performance.

Figure 7 shows the amount of server power savings versus CPU utilization without affecting system performance. The greatest savings can be achieved between CPU utilization of 40 and 80 percent.

Figure 8 shows the amount of server power savings versus CPU utilization as well as the potential data center power savings for data centers with power usage effectiveness (PUE) of 2.0.

Table 2 shows the test data collected running through the procedure shown in Figure 5 using various load test values. The respective power savings are the result of applying lower power policies without negatively impacting performance.

Table 3 shows the approximate power savings shown in the tests and modeled for a data center with 1,000 servers. The data center power savings were calculated using a data center PUE of 2.0.

Conclusions

Testing showed that using Intel Intelligent Power Node Manager and Intel DCM enabled BMW Group to address its power and data center business challenges. The tests enabled BMW to decrease power usage between four and 18 percent, with a wide variety of utilization, without affecting system performance. The tests showed these savings were possible at all levels of CPU utilization.

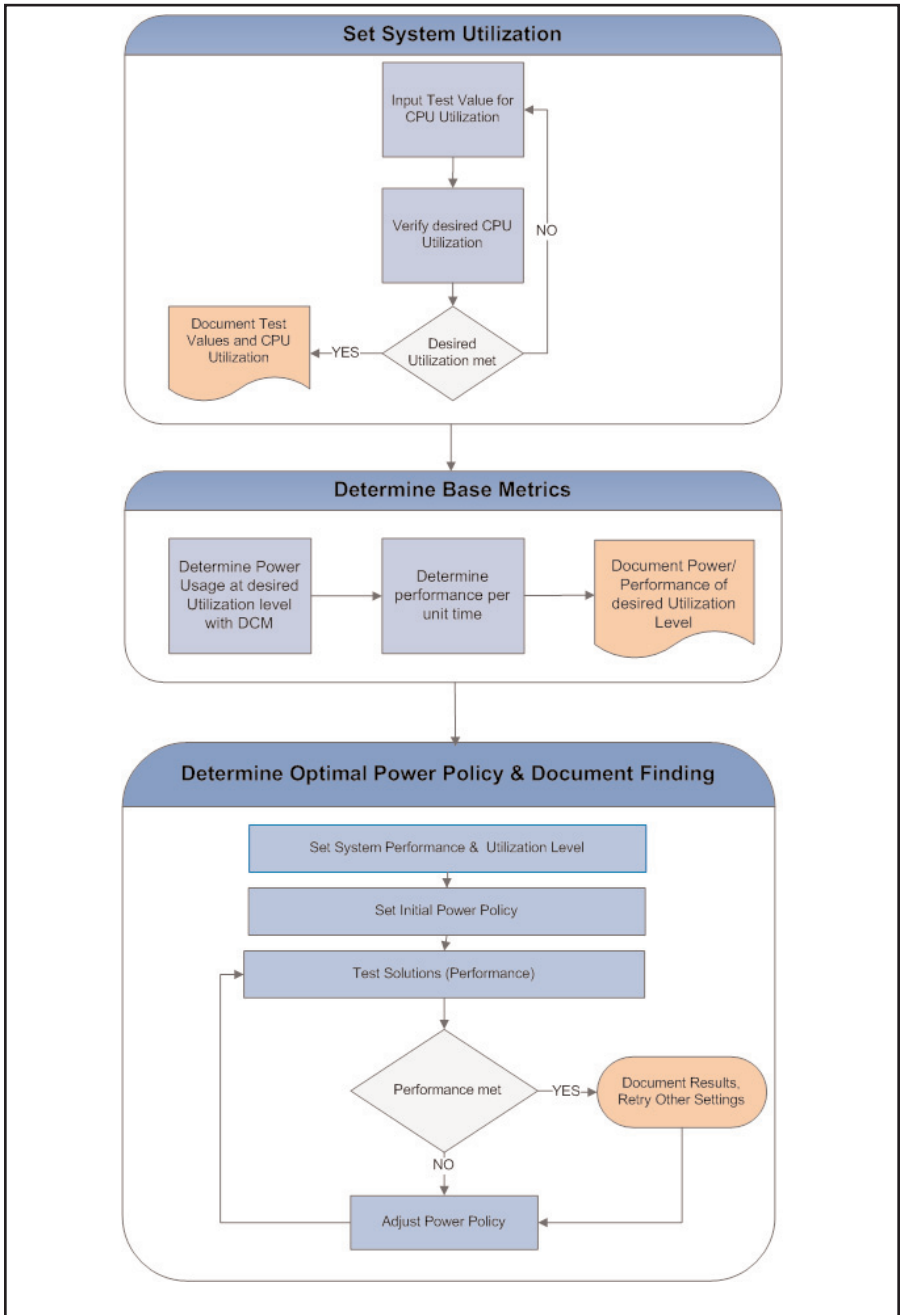


Figure 5. Detailed PoC and test methodology

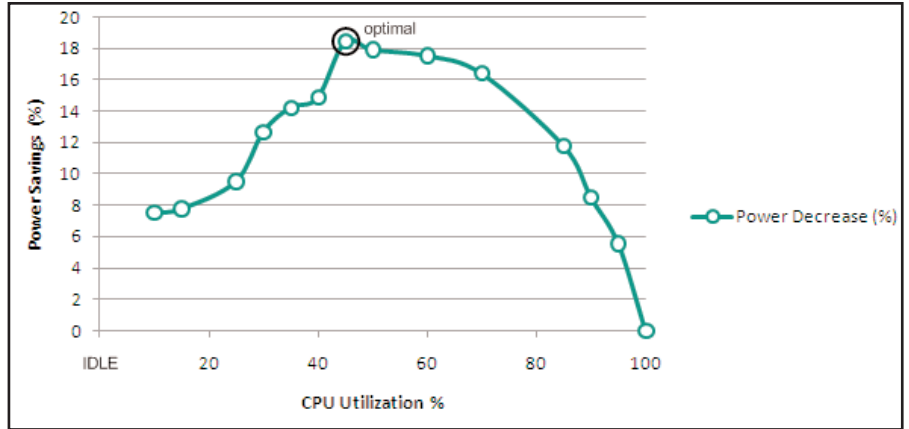


Figure 6. Maximum power savings versus utilization without affecting performance

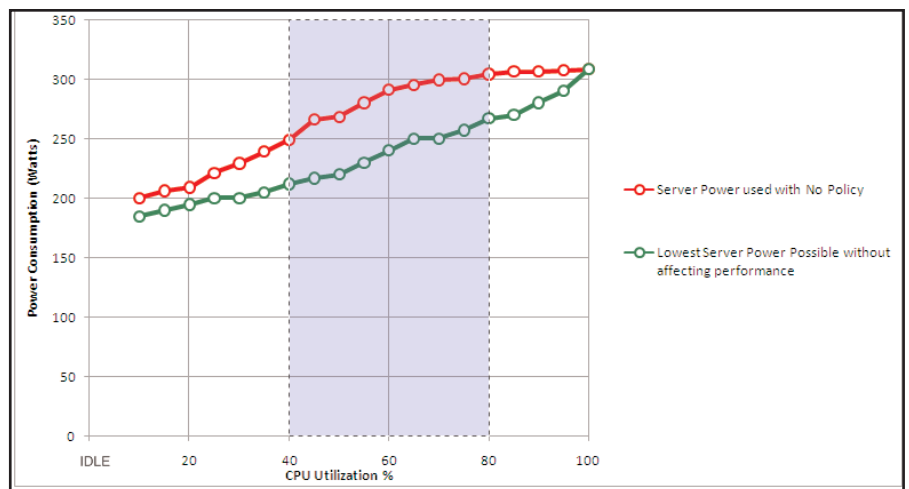


Figure 7. Power savings versus CPU utilization without a power policy and with a power policy

The optimal savings were achieved with CPU utilization of approximately 45 percent. The potential savings would decrease if the CPU were used below approximately 40 percent or above 70 percent. One option to effectively maximize power savings is to use virtualization solutions with Intel Intelligent Power Node Manager and Intel DCM. If system performance is not a required business element, Intel Intelligent Power Node Manager and Intel DCM can still be used with a lower-than-optimal power policy to achieve power savings.

Applying optimum power policies to approximately 1,000 servers would result in savings of server power and also decrease in data center power consumption:

- **Server power savings** of approximately 18 percent KWh
- **A data center with a PUE of 2.0** would yield approximate power savings of 32 percent KWh
- **Server efficiency** increase of approximately 19.5 percent

Intel Intelligent Power Node Manager and Intel DCM can provide numerous benefits besides direct cost savings. For instance, some data centers leave empty spaces in their server racks due to thermal, electrical, or power constraints or requirements. One option is to use the power savings gained from using Intel Intelligent Power Node Manager and Intel DCM to increase the overall number of servers or server density in a given rack without increasing power utilization within a given power envelope. This would postpone construction of new infra-




Figure 8. Projected CPU and data center power savings versus CPU utilization

Table 2. Test Data Using Various Load Test Values

CPU Utilization (%)	Ops./Min., No Policy (Billions)	Load Test Value	Power Used, No Policy (W)	Lowest Power w/out Affecting Perf. (W)	Ops./Min. w/Policy (Billions)	Power Decrease (W)
10	94	23	200	185	120	7.50
15	122	30	206	190	146	7.77
25	176	35	221	200	211	9.50
30	194	44	229	200	211	12.66
35	232	48	239	205	238	14.23
40	261	53	249	212	288	14.86
45	317	62	266	217	319	18.42
50	329	66	268	220	332	17.91
60	391	80	291	240	402	17.53
70	423	83	299	250	423	16.39
85	466	90	306	270	477	11.76
90	489	97	306	280	490	8.50
95	493	99	307	290	504	5.54
100	511	100	308	308	511	0.00

Table 3. Approximate Power Savings for a Data Center with 1,000 Servers

Projected Savings	No Power Policy	Uniform Power Policy	After-Hours Policy	Planned Power Policy
Server Power Usage/Hr. (KWh)	259	240	237	217
Power Usage/Day (KWh)	6,225	5,760	5,711	5,209
Power Usage/Week (KWh)	43,575	40,320	39,977	36,463
Power Usage/Month (KWh)	186,750	172,800	171,330	156,270
Power Usage/Yr. (KWh)	2,272,125	2,102,400	2,084,515	1,901,285
Server Efficiency increase (%)	0.00	8.07	9.00	19.50
Savings Gained vs. No Power Policy (KWh/Yr.)	--	169,725	187,610	370,840
Data Center Power Savings (PUE=2.0) (KWh/Yr.)	--	339,450	375,220	741,680

Table 4. Comparative Potential Data Center Cost Savings for 1,000 Servers

Total Cost for Power (US\$)	No Power Policy	Uniform Power Policy	After-Hours Policy	Planned Power Policy
Total Cost Per Day	454	420	416	380
Total Cost Per Week	3,180	2,943	2,918	2,661
Total Cost Per Month	13,632	12,614	12,507	11,407
Total Cost Per Year	165,865	153,475	152,169	138,793
Server Savings Per Year	--	12,389	13,695	27,071
Data Center Cost Savings Per Year (PUE=2.0)	--	24,779	27,391	54,142

structure. For example, Baidu Inc., a large Internet search portal, was able to use Intel Intelligent Power Node Manager and Intel DCM in its testing to demonstrate how to increase server rack density between 20 and approximately 40 percent without increasing the overall rack power.

Both companies and suppliers are increasingly focusing on green computing and addressing computer efficiency, data center power efficiency, and performance per watt. The PoC results show the benefits of adapting and using Intel platforms with Intel Intelligent Power Node Manager and Intel DCM. IT organizations facing power and cooling challenges can benefit from intelligently applying power policies and utilizing Intel platforms with Intel Intelligent Power Node Manager and Intel DCM to:

- **Address server data center power and cost savings** without impacting system performance.
- **Increase rack density** within space, power, and cooling constraints through fine-grained power control.
- **Reduce capital costs** by right-sizing power and cooling infrastructure based on actual power consumption trends.
- **Reduce operations costs** by eliminating worst-case headroom buildup during rack power provisioning. Power is provisioned to servers based on workload SLAs.

While there are limits to applying power policies without impacting servers' performance, decreasing system power while keeping performance level the same will increase the overall system efficiencies.

For BMW Group, the benefits of Intel Intelligent Power Node Manager and Intel DCM were reduced server and data center infrastructure power consumption for increased server rack server density—which directly translates into financial gains.

For more information on Intel Xeon processor 5500 series, visit www.intel.com/xeon or www.intel.com/references.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

Copyright © 2009 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others. Printed in USA SS/MG/0909 Please Recycle

