

Data Center Energy Efficiency with Intel® Power Management Technologies

- Ability to monitor and cap server power consumption
- Up to 20 percent reduction in server power use with no impact on runtime
- Potential for increased computing capacity, power savings, and business continuity

Intel IT conducted a technology evaluation of Intel® Intelligent Power Node Manager and Intel® Data Center Manager (Intel® DCM). Our goals were to assess the potential of these Intel® power management technologies to increase data center energy efficiency, and to validate potential usage models.

We conducted our evaluation in a test environment representing a virtualized data center, using servers based on Intel® Xeon® processor 5500 series.

We successfully used Intel Intelligent Power Node Manager and Intel DCM to monitor and cap power consumption across individual servers and groups of servers. For workloads that were not processor-intensive, we optimized server power consumption by up to approximately 20 percent without impacting performance, as shown in Figure 1.

Power monitoring is a critical capability that enables us to characterize workloads and identify opportunities to increase data center energy efficiency. Our evaluation showed that Intel power management technologies can address key data center power and cooling challenges, helping to increase computing capacity, reduce power consumption, and maintain business continuity.




Figure 1. We used capping to reduce server power consumption by approximately 20 percent when running an I/O-intensive workload, without impacting runtime.

Business Challenge

Intel IT, like other organizations, faces significant data center power and cooling challenges. Rapid growth in demand drives a continual need for more computing resources. This is straining the limits of data center power and cooling capacity. At the same time, power and cooling costs are becoming an increasingly important component of total cost of ownership (TCO).

Ways to accommodate the increasing demand include building new data centers or adding power and cooling capacity to existing data centers. However, both of these options are extremely expensive and take a long time to complete.

Because of this, we are increasingly applying alternative approaches that focus on using existing data center power more efficiently in order to increase computing capacity, cut power costs, and reduce Intel's carbon footprint.

Traditionally, because IT organizations have lacked detailed information about actual server power consumption in everyday use, data center computing capacity has been based on nameplate power, peak server power consumption, or derated power. In practice, server power consumption with real data center workloads is nearly always lower than this. This situation results in overprovisioned data center power, overcooling of IT equipment, and increased TCO.

If we can better understand and control server power consumption, we can more efficiently use existing data center power, resulting in benefits such as increased capacity and reduced power costs. Applied across tens of thousands of servers, this could result in considerable savings.

To this end, we conducted an evaluation of two new Intel power management technologies: Intel Intelligent Power Node Manager and Intel

Data Center Manager. These technologies allow us to monitor and measure server power consumption across the data center, and cap this power consumption at targeted levels.

INTEL® INTELLIGENT POWER NODE MANAGER

Intel Intelligent Power Node Manager is an out-of-band power management policy engine available with the Intel Xeon processor 5500 series. It enables regulation of individual server power consumption (power capping) through modulation of the processor's performance (P) and throttle (T) states. It measures and reports actual server power consumption, and provides for power capping to a targeted power budget specified by an enabled console manager. If this budget cannot be maintained, Intel Intelligent Power Node Manager can send alerts to the management console.

INTEL® DATA CENTER MANAGER

Intel DCM is a software development kit and reference user interface that provides power and thermal monitoring and management for servers, racks, and groups of servers in data centers. Management console providers and system integrators can use it to build management consoles that employ Intel Intelligent Power Node Manager to monitor and cap the power consumption and temperature of individual servers or groups of servers. Intel DCM can aggregate data from systems enabled with Intel Intelligent Power Node Manager. It lets administrators adjust caps based on server utilization, business conditions, and power consumption trends.

Data Center Use Cases

We identified several potential data center use cases that take advantage of Intel Intelligent

Power Node Manager and Intel DCM to increase computing capacity, save power, and improve business continuity.

CAPACITY MANAGEMENT

By measuring server power consumption, we can more accurately determine data center capacity based on the power that is actually required. This enables us to exploit existing unused power capacity to increase computing capacity and rack density in a power-constrained data center without the need for expensive data center construction or additional power circuits and cooling units.

POWER SAVINGS

For workloads that are not constrained by processor performance—such as I/O-intensive and memory-intensive workloads—we may be able to use Intel Intelligent Power Node Manager and Intel DCM to throttle back the server processor without affecting overall performance. As a result, we could reduce server power consumption without risk to service-level agreements (SLAs).

BUSINESS CONTINUITY

Once we have employed server power monitoring to characterize our application workloads, we can use capping to help maintain business continuity in the event of a power outage. We can use capping to distribute the remaining power among servers and prioritize business-critical workloads.


Technology Evaluation

In our technology evaluation, we used Intel Intelligent Power Node Manager and Intel DCM to monitor and cap power consumption of Intel Xeon processor 5500 series-based servers in a test configuration that modeled a virtualized IT environment. Our goals were to:

- Test the features of Intel Intelligent Power Node Manager and Intel DCM
- Explore the value of these technologies for our key data center use cases.

TEST ENVIRONMENT

Our test environment is shown in Figure 2. The environment included four servers based on Intel Xeon processor X5550, each running a hypervisor and one or more virtualized test workloads. Three of these servers each supported three test workloads in separate virtual machines (VMs); the fourth server ran the test



VM Virtual Machine

Figure 2. Test environment.

automation framework plus one test workload. This created a total of 10 test workloads, all of which executed simultaneously in our tests.

We used a synthetic workload that generated an adjustable mix of computations and database I/O operations. We tuned the workload to make it more compute-intensive or I/O-intensive, depending on the test. This approach enabled us to create predictable, consistent test loads and to observe how capping affected performance with different types of workload.

FEATURE TESTING

We successfully tested the capabilities of Intel Intelligent Power Node Manager and Intel DCM to monitor power consumption and temperature of individual servers and groups of servers, to set caps, and to retain server power utilization data over time.

TEST CASES

Our test cases investigated the use of capping for capacity management, power savings, and business continuity.

Capacity Management and Power Savings

We measured the effect of capping on runtimes with I/O-intensive and CPU-intensive workloads. In each test, we simultaneously ran 10 identical I/O-intensive or CPU-intensive workloads dispersed across the VMs on our test servers. We first ran the workloads without capping power consumption. We measured power consumption throughout the test and recorded job completion time, defined as the time required for power consumption to return to idle level across the environment.

We then repeated the test using Intel DCM policies and Intel Intelligent Node Manager settings to cap power consumption. In successive repetitions, we progressively reduced power consumption by 100 watts (w) at a time.

With the I/O-intensive workload, we were able to reduce server power consumption by approximately 20 percent, or 200 W, without performance impact, as shown in Figure 1. Each reduction caused Intel Intelligent Power Node Manager to dynamically change the processor P states and T states to help ensure the target server power consumption was reached. Because the workload was not CPU-constrained, this did not affect overall performance until the reduction in power consumption exceeded 20 percent.

When we conducted similar tests with a CPU-intensive workload, we observed a different effect: Reducing power consumption by 20 percent resulted in an approximately 18 percent increase in runtime, shown in Figure 3.

Business Continuity

Our test simulated an outage causing a partial loss of data center power. We used capping to allocate the remaining power among all servers, enabling all workloads to keep running. We prioritized the server workloads by capping power consumption at different levels.

We used two different workloads, as shown in Figure 4: one that was CPU-intensive (left) and one with more limited resource requirements (right). When we substantially reduced the cap for servers running the CPU-intensive workload (top) runtime was increased by approximately 35 percent; in contrast, capping the servers running the job with limited performance requirements (bottom) resulted in a minimal performance impact. This suggests that capping can be used to avoid downtime and maintain performance of high-priority workloads in the event of a partial outage.




Figure 3. Effects of capping on runtime of CPU-intensive workloads. Reducing power consumption by 20 percent resulted in an 18 percent increase in runtime.




Figure 4. Using capping to help maintain business continuity.

Analysis

Our results suggest that we can use capping to reduce server power consumption without impacting SLAs, enabling us to increase data center computing capacity and reduce power costs. With server workloads that are not CPU-intensive, we reduced server power consumption by up to about 20 percent without affecting performance.

Our evaluation also showed that once we have characterized and comprehended the application environment, it is possible to use Intel Intelligent Power Node Manager and Intel DCM for business continuity, using capping to avoid downtime and maintain performance of high-priority workloads.

ADDITIONAL POTENTIAL BENEFITS

In addition to the tested use cases, we have identified several additional potential benefits for future investigation.

Monitoring and managing server power consumption allows us to understand and control actual power utilization across a data center room. This enables better utilization and sizing of facility power equipment such as uninterruptible power supply (UPS), reducing overall data center TCO.

Intel Intelligent Power Node Manager and Intel DCM also provide server-level thermal

monitoring with alarming. This could allow us to safely increase data center operating temperatures, letting us increase compute capacity without adding cooling equipment.

Conclusion

Our evaluation indicated that Intel power management technologies could enable us to expand data center capacity, reduce power consumption, help ensure business continuity, and improve facilities management.

Capping may be particularly useful in reducing the power consumption of servers running workloads that are not CPU-constrained, such as I/O-intensive or memory-intensive workloads.

Server monitoring is a critical capability enabling us to characterize and control our application environment. Once we have comprehended the environment, we can evaluate the tradeoffs between performance and power consumption, and determine where to use capping to realize the benefits of increased data center energy efficiency.

**Learn more about Intel
IT's best practices at
www.intel.com/IT.**

AUTHORS

Alan Ross
Principal Engineer, Intel IT

Bob Stoddard
Manageability Architect, Intel IT

CONTRIBUTORS

Chris Black
Todd Christ
Madhu Datla
Tarun Viswanathan

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, reference www.intel.com/performance/resources/benchmark_limitations.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR

ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

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.

Copyright © 2010 Intel Corporation. All rights reserved.

Printed in USA
0210/KAR/KC/PDF

 Please Recycle
322023-001US

