



Data Center Energy Efficiency 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 helps improve cloud data center energy efficiency and operation management.

Intel® Corporation

December 2011

Contents

Executive Summary	2
Intel® Intelligent Power Node Manager and Intel® Data Center Manager	2
POC Test Environments	2
POC Test Cases	3
Power/Thermal Monitoring and Event Management	3
Power Saving/Optimization	3
POC Results	3
Additional Observations:	3
Glossary	4

Executive Summary

Chunghwa Telecom (CHT) in Taiwan is aggressively developing and implementing a cloud business model with the vision of becoming a global ICT leader with its convergence technology. In particular, Infrastructure As A Service (IaaS) is one of the key offerings of CHT cloud businesses. Their business will require a high density compute infrastructure and mandates a high efficiency containerized datacenter.

In this Proof of Concept (POC), CHT and Intel partnered to investigate the use of Intel's power management technologies known as Intel® Intelligent Power Node Manager and Intel® Data Center Manager. The POC was conducted in CHT's existing datacenter with the goal of making CHT's future cloud datacenter more energy efficient, highly available, and with increased server density.

Key findings from the POC include:

- 5%-15% power saving without and with limited performance impacts, respectively;
- Prolonging business continuity time by up to 25% when there is power outage;
- Up to an 83% rack density increase within the same rack-level power envelope when a power management policy is applied
- Successful priority-based allocation of power to workloads under a power-limit scenario;
- Successful alarm triggering and power control when abnormal power or thermal events occur.

Intel® Intelligent Power Node Manager and Intel® Data Center Manager

Intel® Intelligent Power Node Manager provides users with a powerful tool for monitoring and optimization of energy usage, enhancing cooling efficiency, and identifying thermal hot spots in the data center. Intel® Intelligent Power Node Manager is an out-of-band power management policy engine first available with the Intel Xeon 5500 series processor. Intel® Intelligent Power Node Manager works with servers' hardware and software and the operating system's power management features to manipulate the processors' internal power states and dynamically adjust platform power to achieve optimal performance and power.

Intel® Data Center Manager (Intel® DCM) is software technology that provides power and thermal monitoring and management for servers, racks, and the whole data center. It builds on Intel® Intelligent Power Node Manager and customers' existing management consoles to bring platform power efficiency to end users.

When Intel® Intelligent Power Node Manager is paired with Intel® DCM, the combined solution can help more precisely utilize existing power and cooling resources in the entire data centers.

POC Test Environments

The POC test was conducted with a rack of 50 nodes with two Quad-core Intel® Xeon® 5600 series processors each node. The workload was simulated using industry standard benchmark SPECpower_ssj2008. The details of the POC equipments used are listed in Table 1.

Devices	Descriptions
Server Platform	Intel® Xeon® Processor 5600 Series (Quad-Core)
	A container-type rack of servers (50 nodes with two processors per node) with independent fans walls
Server OS	Microsoft® Windows Server® 2008
Tools	Power management tool: Intel® Data Center Manager v2.0
	Workload simulation tool: SPECpower_ssj2008

Table 1 - POC Test Environment

POC Test Cases

The test cases of the POC were as followings:

Power/Thermal Monitoring and Event Management

The 'Power/Thermal Monitoring and Event Management Power Guard' test scenario was designed to record real-time power consumption, record server inlet temperatures, and manage power policies on the servers under test.

Test Case #	Test Case	Descriptions
TC1.1	Real-time power monitoring and power guard	Real-time power monitoring at the server/rack level and power guard to deal with abnormal status
TC1.2	Real-time thermal monitoring and thermal guard	Real-time thermal monitoring of the server inlet temperature and thermal guard to deal with abnormal thermal status

Table 2 - Test Cases of Power/Thermal Monitoring and Event Management

Power Saving/Optimization

The purpose of the 'Power Saving/Optimization' test scenario was to demonstrate energy saving from higher server density and optimization of power consumed by the servers and rack.

Test Case #	Test Case	Descriptions
TC2.1	Power saving and priority based power optimization	Power saving with minimum performance impact and priority-based power optimization by allocating more power to workloads with higher priority under power-limit scenario
TC2.2	Rack density improvement	Rack density improvement through power-workload optimization

Table 3 - Test Cases of Power saving/optimization

POC Results

The POC results were as followings:

- Up to 5% power savings with no performance impact
- Up to 15% power savings with limited performance impacts
- Prolonged the business continuity time by up to 25% when there is power outage;
- Up to 83% increase in rack density with the same rack-level power envelope when aggregated optimal power management policy is applied.
- Demonstrated allocation of power to workloads with higher priority under power-limit scenario
- Successful alarm triggering and power control when abnormal power or thermal events occur.

For CHT's cloud datacenter, the 5%~15% power saving achieved by Intel® Intelligent Power Node Manager and Intel® DCM will save NT\$10,000~ NT\$30,000 annually per rack (assuming a PUE of 1.5, 5KW power consumption per rack, and energy cost of NT\$3/KWh).

Moreover, power-based and thermal-based power management features in Intel® Intelligent Power Node Manager and Intel® DCM further help CHT to optimize power usage and improve service availability under abnormal power or thermal events.

Additional Observations:

Integrating Intel® DCM API with a hypervisor can make server utilization more efficiently without overloading a server, and can help protect content by moving the VM containing the content to safe servers when there are abnormal power or thermal events. Furthermore, with a thermal triggered power policy, the workload can be adjusted, reducing the heat generated by the servers, and resolving hot spots in the data center.

Glossary

BMC	Board Management Controller
Intel® DCM	Intel® Data Center Manager
IaaS	Infrastructure As A Service
POC	Proof of Concept
VM	Virtual Machine
OSPM	Operating System Power Management
PUE	Power Usage Effectiveness

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.

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 © 2011 Intel® Corporation. All rights reserved. Intel®, the Intel® logo Intel®'s Power Management Technologies inside are trademarks of Intel® Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.