

Intel® Cloud Builders Guide: Cloud Design and Deployment on Intel® Platforms

Policy Based Power Management with Dell* and VMware*



Intel® Xeon® Processor 5500 Series

Intel® Xeon® Processor 5600 Series

AUDIENCE AND PURPOSE

This reference architecture is intended for data center administrators and enterprise IT professionals who seek power management solutions to achieve better power efficiency within new or existing datacenters. The techniques and results described can be used as a reference to understand policy based power management solutions implemented with the use of hardware and software components. The reader should be able to develop appropriate power management solutions based on the design options presented.

vmware



Table Of Contents

Executive Summary	4
Introduction	5
Server Power Management	6
Test Bed Blueprint	7
Design Considerations	7
Hardware Description	7
Physical Architecture	8
Installation and Configuration	9
VMware* Components	9
"Plug-In" Components	11
Policy Based Power Management Use Cases	12
Use Case One: Real time Server Power Monitoring and Scheduling	12
Purpose	12
Pre-requisites	13
Steps for Execution	13
Results	13
Use Case Two: Optimize Rack Density/Usage	18
Purpose	18
Pre-requisites	18
Steps for Execution	18
Results	19
Use Case Three: Power Optimized Workloads	19
Purpose	19
Pre-requisites	19
Steps for Execution	19
Results	20
Use Case Four: Disaster Recovery / Business Continuity	21
Purpose	21
Pre-requisites	21
Steps for Execution	21
Results	21
Composite - Policy Based Power Management Use Cases	23
Use Case Five: Datacenter Energy Reduction through Power Aware Support for Multiple Service Classes	25
Purpose	25
Pre-requisites	25
Steps for Execution	26
Results	26
Things to Consider	26
Architectural Considerations	26
Storage	26
Scalability	26
Networking	27
Power Management	27
Hardware	27
Additional Usage Models Under Development	27
Thermal Power Management	27
Development and Usage of "Plug-in"	28
Glossary	28
APPENDIX A: Server Power Management	30
Intel® Power Management Technologies	30
Intel® Intelligent Power Node Manager	30
VMware* Distributed Power Management (DPM)	31
APPENDIX B: Dell* Cloud Solutions	32
Dell* PowerEdge C-Series Servers	32
APPENDIX C: VMware* Power Management Plug-ins	33
Power Management "Plug-in"	33
APPENDIX D: Retrofitting Power Management	35

Executive Summary

To support the evolution to the cloud, Intel works with end users to create an open data center roadmap of usage models that address key IT pain points for more secure, efficient, and simple cloud architectures built on a foundation of transparency. This paper describes a policy power management reference architecture based on Intel, VMware, and Dell solutions. Policy power management refers to usage models aimed at power efficiency.

The goal of policy power management usage models is to optimize productivity per watt in order to reduce total cost of ownership (TCO). Requirements include the capability to monitor and cap power in real-time at server, rack, zone, and data center levels. This means the ability to manage aggregated power consumption within a rack, zone, or data center based on available power and cooling resources. Load migration policies based on power consumption and cooling resources bring further benefits.

In this reference architecture we used Intel® Intelligent Power Node Manager¹ (Intel Node Manager) and, Intel® Data Center Manager (Intel DCM)² to provide power savings through dynamic migration of workloads to fewer servers during periods of low resource utilization and through deactivation of unneeded hosts.

VMware vSphere and Intel Node Manager were integrated with the use of respective software development kits (SDKs). With the use of Intel DCM, the power management “plug-in” (i.e. integration component) extends the ability of cloud and virtualization resource management engines, such as VMware vCenter* Server, to optimize resource distribution. The plug-in used in the tests was developed to highlight one possible usage model of Intel Node Manager. Ideally this power management feature should be directly integrated into the hypervisor or other management applications/plug-ins

that could take into account all the other remaining system parameters beyond just power. We describe the following power management use cases in detail along with experimental results and data.

1. **Real-time Server Power Monitoring** to perform power aware scheduling. Virtual machines can be relocated from power constrained systems to unconstrained systems within the cluster or across different clusters for better system utilization and performance.
2. **Optimization of Rack Density** to maximize available compute resources. In our experiments with dynamic server power capping, we were able to achieve an increase of 30 to 50 percent in server density and maintain the same overall power envelope per rack.³
3. **Power Optimized Workloads** to achieve power efficiency. Workload profiles are built and a maximum performance loss target set. Experiments determine how much capping can be applied before the performance target is hit. The approach is to match actual performance against service level requirements. For workloads that were not processor intensive, we were able to optimize server power consumption by approximately 20 percent without an impact on performance.³ Based on experiments performed in our labs and in advanced technology proofs of concept we discovered that for workloads that were processor intensive, for the same 20 percent power saving, we saw an 18 percent decrease in performance. For a 10 percent power reduction, performance decreased by 14 percent.³
4. **Disaster Recovery/Business Continuity** prioritizes business-critical workloads and distributes them among servers that still have

power in the event of servers going off-line due to power failure. When we substantially reduced the cap for servers running processor-intensive workloads, runtime was increased by approximately 35 percent.³ For servers running I/O-intensive workloads there was minimal performance impact.³

5. **Data Center Energy Reduction** through Power Aware Support for Multiple Service Classes showcases the ability to enforce multiple SLAs across different populations of users with different priority workloads. Workloads that ran over a period of eight hours realized 25 percent less energy consumption.³

The opportunities to reduce energy use through power capping technology alone are limited. For example, the policy may seldom kick in. Policies under dynamic power management take advantage of advanced platform power management technologies. Power capping levels vary over time and become control variables. Selective equipment shutdowns enable reductions in energy consumption, not just power management. The tradeoff for dynamic policies is additional complexity (a closed loop power control system needs to be implemented).

The paradigm of cloud computing brings opportunity for data center efficiency. Policy power management usage models can substantially help to meet power management requirements.

Introduction

To break the trend of the decline in flexibility combined with the increase in costs, a new model for IT services has emerged—cloud computing. Cloud computing is an approach to computing that uses the efficient pooling of an on-demand, self-managed virtual infrastructure, consumed as a service. This approach abstracts out applications and information from the complexity

of underlying infrastructure, so IT can support and enable business value. Cloud computing architectures are built on the foundation of virtualization, and as the customer-proven leader in virtualization, VMware helps to chart the course to cloud computing. In concert with Intel, Dell, and other industry leaders, VMware helps businesses of all sizes migrate to power efficient cloud computing, in order to establish a new era in IT that finally addresses the compounded problems of IT cost and complexity.

At the core of cloud computing is the ability of the underlying compute, network and storage infrastructure to act as an efficient, shared resource pool that is dynamically scalable within one data center or across multiple data centers. With this foundation, critical higher-level capabilities such as policy-based power management, guaranteed quality of service, federation, and data center automation are made possible. Intel, along with leaders in software, works to address these new core innovations in Infrastructure as a Service (IaaS). Intel has initiated a program to rapidly enable enterprises and service providers to identify best practices around design (including reference architectures), deployment, and management. For enterprise IT and cloud service providers who need to utilize their existing data center infrastructure to supply cloud services to their customers, this guide, as part of the Intel® Cloud Builders initiative, provides a comprehensive solution overview that covers technical planning and deployment considerations.

Apart from the compelling operational advantages of a highly scalable and flexible compute infrastructure, both within an organization (private cloud) and externally (public cloud), the potential to enable new services and client devices is vast and in the early adoption stage. Scalability and flexibility in the data center have a direct relationship to increased computing capacity and the

greater need to address potential security threats. In order to realize the vision of cloud computing, new technologies are needed to address security and power efficiency. These will become fundamental to architectures from the microprocessor stage up through the application stack. The focus of this paper is power efficiency and the related usage models.

Based on the Environmental Protection Agency's report to the government, in 2006 data centers in the US consumed about 1.5 percent of the nation's energy and were poised to double this by 2011.³ If storage, network, and computing resources continue to grow at their predicted rate, new power efficient usage models will be required. Higher server utilization, better throughput for network and storage traffic, as well as storage optimized by data type and needs, are a few ways to maximize the existing resources to achieve efficiency.

Companies continue to explore approaches that focus on using existing data center power more efficiently to increase computing capacity, cut power costs, and reduce carbon footprint. Traditionally, organizations have lacked detailed information about actual server power consumption in everyday use. Typically, data center computing capacity has been based on nameplate power, peak server power consumption, or an approximation based on power loads. In practice however, actual power consumption with real data center workloads is much lower than system specifications. This situation results in over-provisioned data center cooling and power capacity, and increased total cost of ownership (TCO). Better understanding and control over server power consumption allows for more efficient use of existing data center facilities. All of this, applied across tens of thousands of servers, can result in considerable savings.

This paper begins with a technical overview of server power management

techniques. We then describe various usage models in detail; we explain their purpose and pre-requisites. We describe the test cases executed and their results, and we include screenshots of the test process. Finally, we describe architectural considerations and additional usage models under development.

Server Power Management

In the past, power consumption used to be an afterthought for server deployment in data centers. Unfortunately, this view persists. For example, in many facilities the utility bill is bundled with the overall building charge which reduces the visibility of the data center cost.

Even though servers have become much more efficient, packaging densities and power have increased much faster. As a result, power and its associated thermal characteristics have become the dominant components of operational costs. Power and thermal challenges in data centers include:

- Increased total operational costs due to increased power and cooling demands
- Physical limitations of cooling and power within individual servers, racks, and data center facilities
- Lack of visibility into actual real-time power consumption of servers and racks
- Complexity of management components and sub-systems from multiple vendors with incompatible interfaces and management applications

These challenges to manage datacenters can be translated into the following requirements:

- Power monitoring and capping capabilities at all levels of the data center (system, rack identification, and data center). What can be done at an individual server level becomes much more compelling once physical or virtual servers are scaled up significantly.
- Aggregation of the power consumed at the rack level and management of

power within a rack group to ensure that the total power does not exceed the power allocated to a rack.

- Higher level aggregation and control at the row or data center level to manage power budget within the average power and cooling resources available.
- Optimization of productivity per watt through management of power at the server, rack, row, and data center levels to optimize TCO.
- Application of standards-based power instrumentation solutions available in all servers to allow management for optimal data center efficiency. Extension of instrumentation to enable load balancing or load migration based on power consumption, and close coupled cooling for the management of pooled power and cooling resources.

Test Bed Blueprint

Intel has worked with Dell to implement a test bed that features Dell's hyperscale-inspired PowerEdge* C servers which are designed specifically for power and space sensitive data centers. The test bed is intended to provide a flexible environment to simulate those aspects of a commercial data center that are relevant to cloud computing usage models.

Design Considerations

Features include:

- Node Manager compliant systems along with advanced configuration and power interface (ACPI) compliant power supply for real-time power monitoring.
- 1 GbE and 10 GbE networks to achieve optimal performance during virtual machine migrations.
- Multiple VLANs to simulate cross site virtual machine migrations.

Physical Architecture

Figure 1 shows the test bed deployment architecture. This shows three different VLANs configured. The VMware vCenter* Server, VMware vSphere Client, and Intel DCM share the same server and reside on

Hardware Description

OEM ESXi* Host	Dell* PowerEdge C1100 ⁴ (2-socket 1U Rack Server)	2-way Intel® Xeon® Processor E5570 @ 2.93 GHz with 12 GB RAM, 250 GB SATA HDD
		Standard baseboard management controller with IPMI 2.0 support DCMI, iKVM, and PXE support Intel® Intelligent Power Node Manager compliant
		Dual hot-plug redundant high-efficiency Power Supplies(650W) – ACPI Compliant
	Software	VMware vSphere Hypervisor* (ESXi 4.1)
OEM ESXi Host	Dell* PowerEdge C2100 ⁴ (2-socket 2U Rack Server)	2-way Intel® Xeon® Processor E5620 @ 2.40 GHz with 12 GB RAM, 500 GB SATA HDD
		Standard baseboard management controller with IPMI 2.0 support DCMI, iKVM, and PXE support Intel® Intelligent Power Node Manager compliant
		Dual hot-plug redundant high-efficiency Power Supplies(750W) – ACPI Compliant
	Software	VMware vSphere Hypervisor (ESXi 4.1)
Management Server	Dell PowerEdge R710 (2-socket 2U Rack Server)	2-way Intel® Xeon® Processor X5670 @2.93 GHz with 16 GB RAM, 500 GB SATA HDD
	Software	Microsoft* W2K8 , IIS, .NET 2.0
		VMware vCenter Server 4.1, VMware vSphere Client 4.1
		VMware vSphere Web Services SDK
		Intel® Data Center Manager 1.5.6, DCM SDK 1.5.6
Data Store NFS Server	Dell PowerEdge R710 (2-socket 2U Rack Server)	2-way Intel® Xeon® Processor X5670 @ 2.93 GHz with 48 GB RAM, 6x300 GB HDD configured as 1.7 TB RAID
	Software	Red Hat* Enterprise Linux* 5.4

Table 1: Hardware Description

one VLAN along with the shared storage. The two Dell systems have VMware vSphere Hypervisor (VMware ESXi 4.1) for a cluster of ESX* nodes and reside on a separate VLAN.

Installation and Configuration

VMware Components

Below are the high level steps required for the installation and configuration of the infrastructure to exercise the platform power management capabilities supported by Intel on the Dell PowerEdge* servers specified above.

The following setup steps assume the reader has a basic understanding of how to install and configure Windows Server* 2008 R2 Enterprise, VMware vCenter Server, and VMware vSphere Client.

- Windows Server 2008 Setup
 - Install Windows Server 2008 on compatible hardware.⁵
 - Configure Web Server Internet information services (IIS) role by choosing WebDAV Publishing, Application Development, Basic Authentication, Windows Authentication, and IIS Management Compatibility services.
 - Ensure IIS is configured to process the ASPX pages as you check on the multipurpose Internet mail extensions (MIME) types supported. If not configured, create a new MIME type for ASPX pages.⁶
- Installation of VMware vSphere Hypervisor hosts
 - Install VMware ESX 4.1 or VMware ESXi 4.1 on the hosts. Both versions of VMware hypervisor support Intel Intelligent Power Node Manager.
- Installation of VMware vCenter Server and VMware vSphere Client
 - Install VMware vCenter Server 4.1 using default settings, except for the following change:

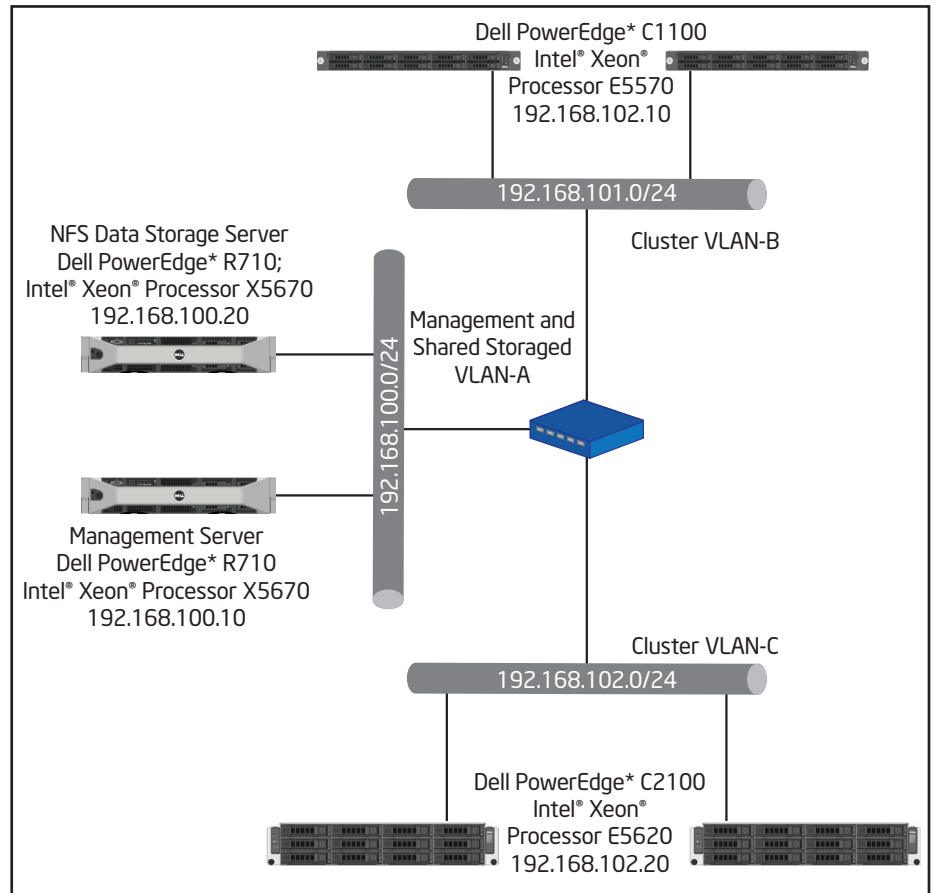


Figure 1: Physical Architecture of Test Bed Setup

- Change the default HTTP port to 81 and HTTPS port to 444 (or other port number of your choice).
Note: This port number needs to be changed since Microsoft IIS runs on the same system, which will use the default HTTP and HTTPS ports.
- Install VMware vSphere Client 4.1 with default values.
- Configuration of VMware vCenter Server
 - Create two different clusters. Hosts on one VLAN configuration will be added to one cluster and hosts on the second VLAN configuration will be added to the second cluster.
- Configure "Manual" distributed resource scheduler (DRS) settings on the clusters and also setup "Enhanced vMotion Compatibility (EVC)" settings based on the ESX hosts that will be added into the clusters. Ex: If your clusters will just contain Intel® Xeon® processor 5500 series and 5600 series servers, you can choose "Intel® Xeon® Core™ i7" as your EVC Mode configuration. This mode will support flex migration of virtual machines between the Intel Xeon processor 5500 series and 5600 series systems.
- Ensure that all required configurations necessary for live VM migrations are completed.⁷

Power Management Plug-In Components

- Installation of the Intel DCM 1.5.6
 - For Dell systems, ensure that Intel DCM 1.5.x or higher is used. Keep the default settings for the installation except for enabling the RMI Port, which is required to make web service calls to DCM.
 - After the installation, ensure that the configuration in Figure 3 is set under the Configuration -> Global options.
 - After the installation is complete, configure Intel® DCM so that it replicates the clusters/hosts hierarchy created in VMware vCenter Server. While you configure the groups and nodes with the reference UI that comes installed by default, ensure the following conditions are met:
 - The group names created in Intel DCM should match the cluster names created in VMware vCenter Server.
 - The node names created in Intel DCM should match the host names in VMware vCenter Server.
 Note: The terminology that Intel DCM and VMware uses for clusters and hosts are different.
- Power Management Plug-in Installation
 - Install the plug-in on the server configured with IIS. Change the default virtual directory where the plug-in needs to be installed.
 - After the successful installation, you need to configure the Web.Config file located in the directory of installation. The following attributes should be changed in the web.config file:
 - URL for VMware vCenter Server
 - URL for Intel DCM installation
 - URL for the plug-in installation

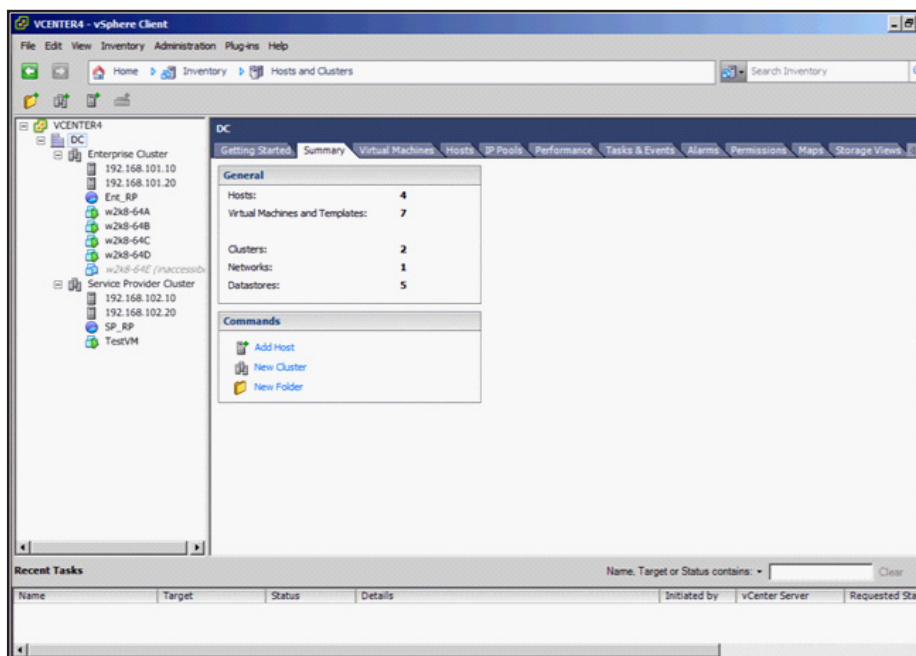


Figure 2: VMware vSphere Client User Interface

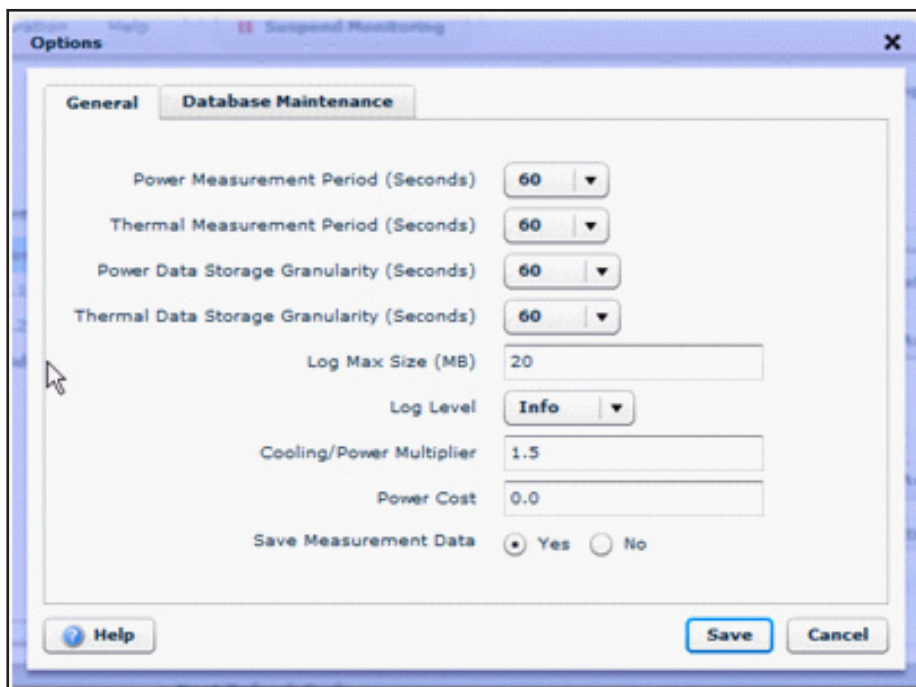


Figure 3: Intel DCM Global Options Configuration

- Register the plug-in with the Managed Object Browser (MOB) tool provided as part of vCenter Server installation. Please refer to <http://communities.vmware.com/docs/DOC-9203> for detailed steps.
- Login to VMware vCenter Server through VMware vSphere Client. Click on the data center name created during the configuration of VMware vCenter Server. A new tab that corresponds to the plug-in installed will appear on the right hand side of the VMware vSphere Client.

Policy Based Power Management Use Cases

Use Case One: Real Time Server Power Monitoring and Scheduling.

Power monitoring is a critical capability that enables us to characterize workloads and identify opportunities to increase data center energy efficiency.

Purpose

Real-time power consumption allows us to perform power aware scheduling. Virtual machines that run the workload can be relocated to optimize and rebalance power margins based on measurements. The virtual machines can be relocated from power constrained systems to unconstrained systems within the cluster or across different clusters for better system utilization and performance.

Server power monitoring helps us find and manage data center hotspots and thereby reduce the chances of hardware failure. The analysis of the historical usage of power consumption data enables us to develop forecast models for data center planning. These forecast models are very accurate as they are based on the real-time data.

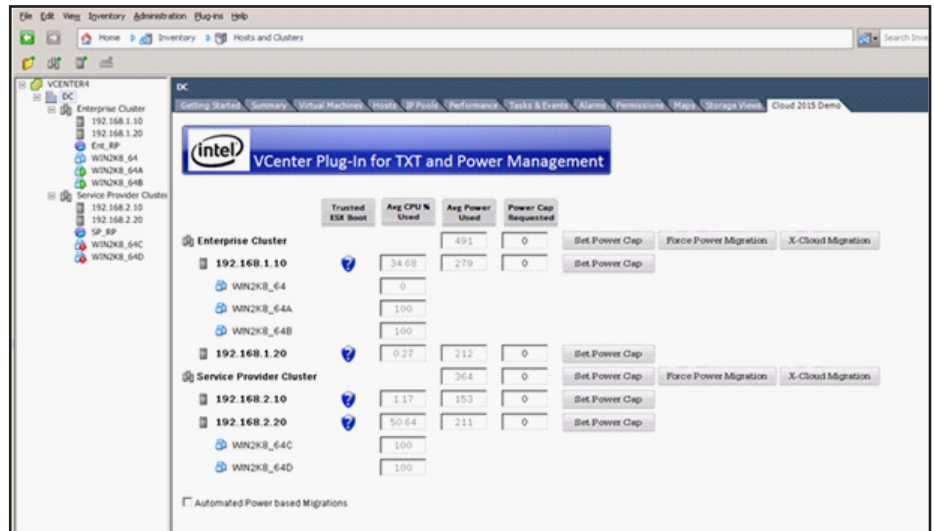


Figure 4: Initial Configuration Without Any Power Constraints

Pre-requisites

- One cluster that consists of ESX hosts with power monitoring capability
- Workloads configured on the virtual machines created on the ESX hosts to drive the host utilization
- Power Management plug-in installed and configured

Steps for Execution

- Use the plug-in UI to set up power caps on the hosts either directly at the cluster level or at the individual server levels. If the power cap is put at the cluster level, DCM will automatically set the power cap of the individual nodes based on historical usage of the nodes.
- Initiate the workload configured on the virtual machines running on the hosts.
- Increase the workload so that one of the hosts on which the workload runs is no longer able to honor the power cap assigned to it because of the increased virtual machine needs. An example scenario involves a three tier application,

wherein the utilization of the DB virtual machines continues to grow as more and more users request data through a cluster of web servers' virtual machines.

- Configure the plug-in to either automatically distribute the workload, or manually click on the migrate button.
- Note: Another scenario can be set up such that there are two clusters, each with at least one ESX host. The ESX host in one of the clusters should be power constrained and the other unconstrained. User can then click on the "X-Cloud Migration" option.

Results

- The highly utilized virtual machine from the power constrained host will be migrated to a low utilization host, which is not power constrained. If all the hosts are power constrained, an appropriate error message is displayed.
- Figure 4 shows the initial configuration without any power constraints wherein the virtual machines receive all the cycles needed for the workload.

- Figure 5 shows that the host “192.168.1.10” is power constrained since it is not able to honor the power policy of 200 watts. There is another host in the same cluster which is not power constrained (“192.168.1.20”).

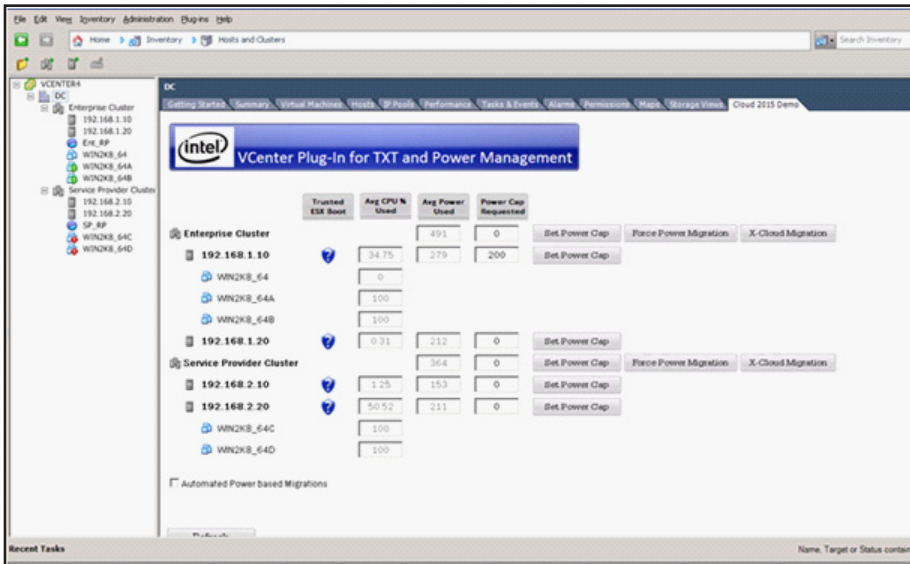


Figure 5: Power Constrained Host Unable to Honor Power Policy of 200 W

- Once the user clicks on the “Force Power Migration” manually or uses the “Automated Power based Migration” option, the highly utilized virtual machine from “192.168.1.10” would be migrated to “192.168.1.20.”

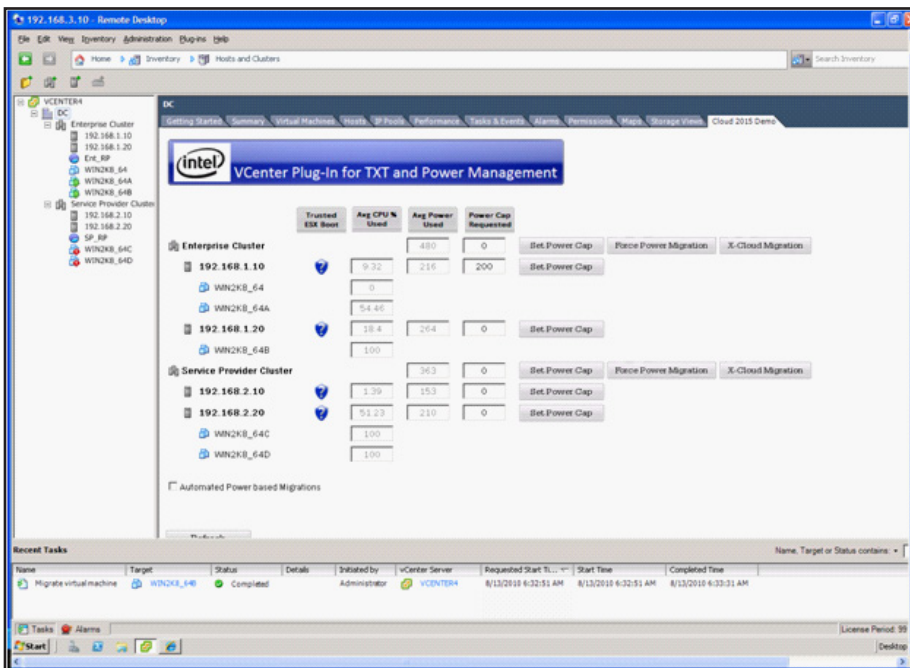


Figure 6: Forced Power Migration

- Now we can set up another scenario to showcase the cross cluster migration. Here we have a power constrained host with a power cap of 190 watts.

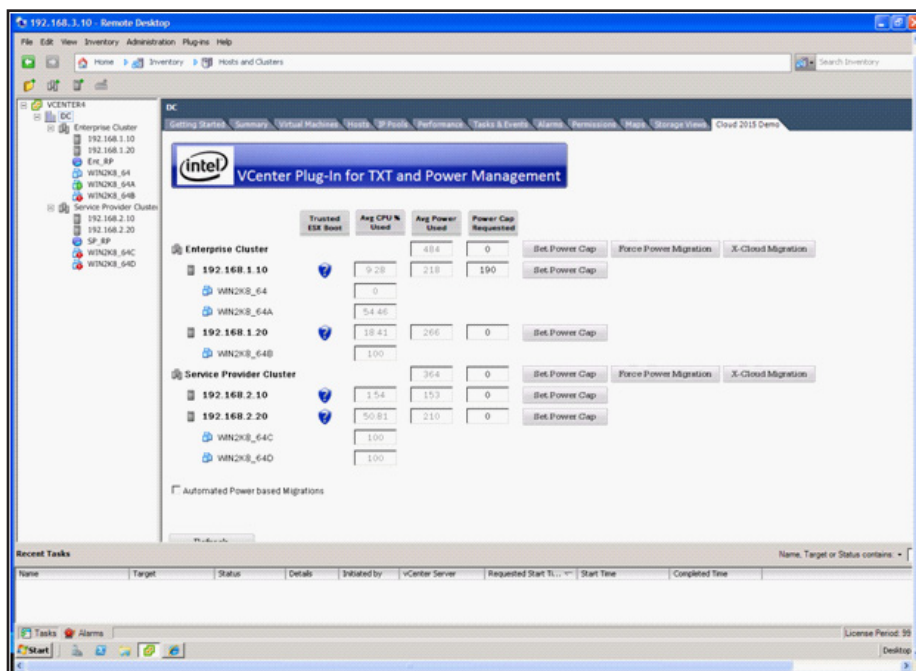


Figure 7: Power Constrained Host

- Figure 8 shows the migration of the high processor utilization virtual machine ("Win2K8_64A") and Figure 9 shows the completed, successful migration.

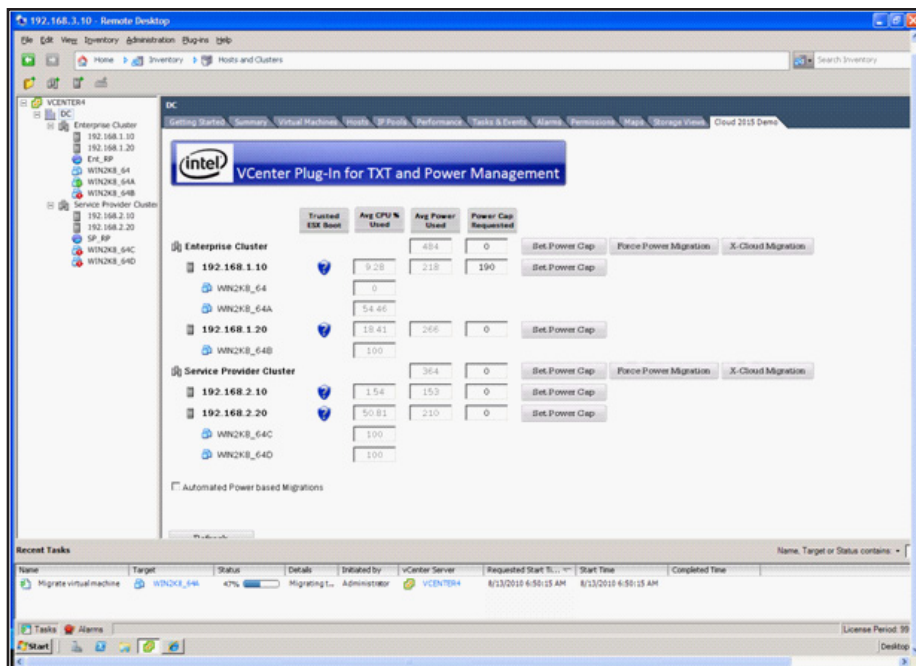


Figure 8: High Utilization Virtual Machine Migration Process

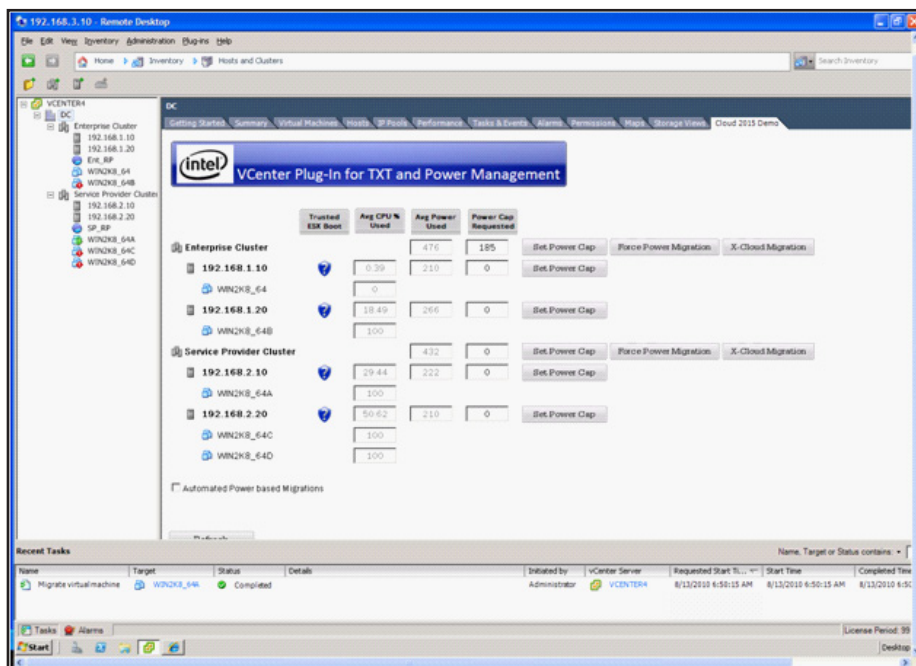


Figure 9: Successful Migration of High Utilization Virtual Machine

Use Case Two: Optimize Rack Density/Usage

Purpose

The collection of real-time power consumption data constitutes an essential capability for power monitoring. Without this data, the best approximation for server power usage comes from the manufacturer's specifications. To use the nameplate numbers as a guidepost requires the allowance of a hefty safety margin. To honor the safety margin in turn leads to data center power over-provisioning and stranded power that needs to be allocated in case it is needed, but is very unlikely to be used. This situation results in over-provisioned data center power, overcooling of IT equipment, and increased TCO.

The availability of power monitoring data allows management by numbers, which tightly matches servers by power consumption to available data center power. The use case is useful in older data centers under-provisioned for power and in host settings with power quotas in effect.

In typical host data centers where the customers are allocated power quotas, the main goal is to optimize the rack utilization so as to place as many servers in a rack as the power limit allows, in order to maximize the microprocessor without interlocked pipeline stages (MIPS) yield. The number of machines will be so large that all machines will likely need to operate under a permanent cap. However, the overall MIPS yield for the collection of machines will be larger than otherwise possible for any combination of machines running uncapped, but whose aggregate power consumption is still subject to the rack power quota.

Pre-requisites

- To understand the SLA requirements of the workload so that the performance gain/loss can be measured against the various power cap values.

Steps for Execution

- Ascertain the power consumption of each server in the rack and also calculate the overall power consumption of the rack.

Note: Assumption here is that the summation of all the nameplate power values of individual servers is equal to the power quota assigned for the rack.

- When you calculate the power consumption, also capture the data related to the SLA of the workload that is running (Ex: response time, utilization etc).
- Use Intel DCM to set the power policy of the rack. Start with lower power cap and increase the value in predefined steps. During each of the steps, capture the SLA related performance data.
- Continue to increase the power cap value to such a point that the performance goes below the SLA required. Capture the previous value in which the performance met the SLA requirements. This would be the optimum power cap value wherein the SLA is still met.

Results

The difference between the total power cap assigned to the rack and the power quota allocated for the rack provides guidance as to how many additional servers with similar power cap settings can be added to the rack without overshooting the power quota allocated. Since we will continue to add additional servers into the rack, the overall performance of the system increases further but stays within the power envelope allocated by the host provider.

In our experiments we have seen an increase of 30 to 50 percent in server density that stays within the same power envelope.² The percentage increase in server density depends on the workload and the SLA requirements.

Please refer to <http://software.intel.com/sites/datacentermanager/whitepaper.php> for real case studies by Intel in concert with external companies.

Use Case Three: Power Optimized Workloads

Purpose

IT organizations (including Intel IT) face significant data center power and cooling challenges. So, companies seek alternative approaches that focus on more efficient use of existing data center power. Power optimization of the workloads is one such approach to achieve power efficiency.

Power optimization requires a table with various workload profiles and a performance loss target not to be exceeded. Developers perform a series of experiments to characterize how much capping can be applied before the performance target is hit. Afterwards, during normal operations, the applications engineer sets power capping targets based on the prior measurements. The system is now said to be “optimized,” because the impact of the application of these caps is now known.

The main benefit of this approach is to match actual Quality of Service (QoS) against service level requirements. Exceeding the SLA generally does not give the provider extra points and indicates unnecessary extra spending. On the other hand, under-delivery on the SLA may result in a noncompliance action by the customer.

Pre-requisites

- IT workload set up on the infrastructure. For this usage model, we used two different types of IT workload. One was a very I/O intensive DB workload and the second one was a high processor workload.

Steps for execution

- Configure the I/O intensive workload on the virtual machines running on the host.
- Run the workload without any power cap and capture the runtime of the workload.
- Now add power cap and gradually increase the power cap value until the runtime starts to increase beyond the baseline value. Note down the power cap value at the point in time there was no runtime impact and beyond which value the runtime started to increase.
- Repeat the above three steps for the processor intensive workload.

Results

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 Node Manager and Intel DCM to throttle back the server processor without an effect on overall performance. As a result, we could reduce server power consumption without risk to service-level agreements (SLAs).

For workloads that were not processor-intensive, we optimized server power consumption by up to approximately 20 percent without an impact on performance, as shown in Figure 10.²




Figure 10: Effects of Capping on Runtime of I/O-intensive Workloads

For workloads that were processor intensive, for the same 20 percent power saving, we saw an 18 percent increase in runtime. Even for a 10 percent power reduction, there was an increase of 14 percent in runtime.²

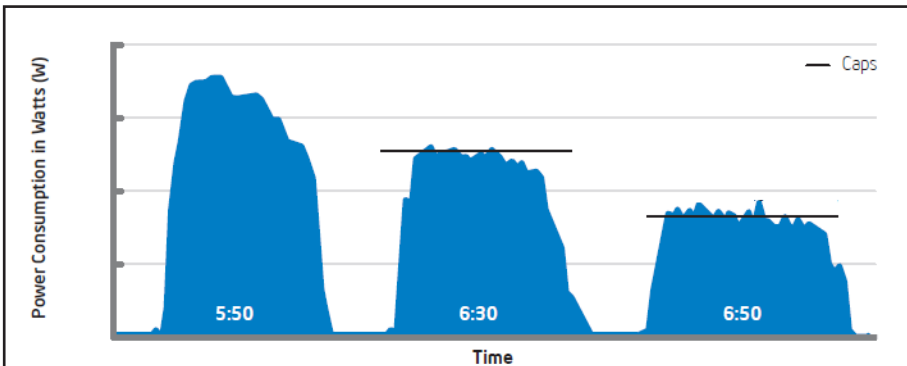


Figure 11: Effects of Capping on Runtime of Processor-intensive Workloads

Use Case Four: Disaster Recovery / Business Continuity

Purpose

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

Pre-requisites

- Both processor intensive and I/O intensive workloads configured.

Steps for Execution

- Simulate partial power outage.
- Based on the power left available, set up power capping to allocate the power that remains among all servers, which will enable all workloads to continue to run.

Results

Since the power capping can be very stringent based on the outage, the performance of the servers will be significantly affected.

Figure 12 shows the difference between the capped and uncapped data for the processor intensive workload. When we substantially reduced the cap for servers running the processor intensive workload, runtime increased by approximately 35 percent.²

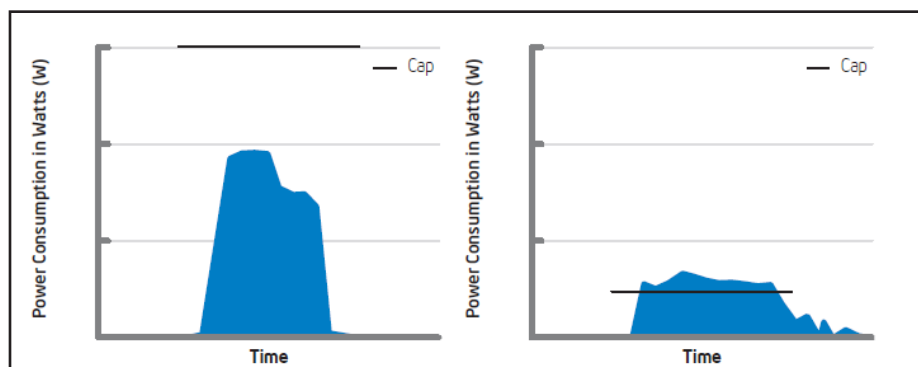


Figure 12: Capped and Uncapped CPU Intensive Workload

Figure 13 shows the data for servers running the I/O-intensive workload. In this scenario there was minimal performance impact.

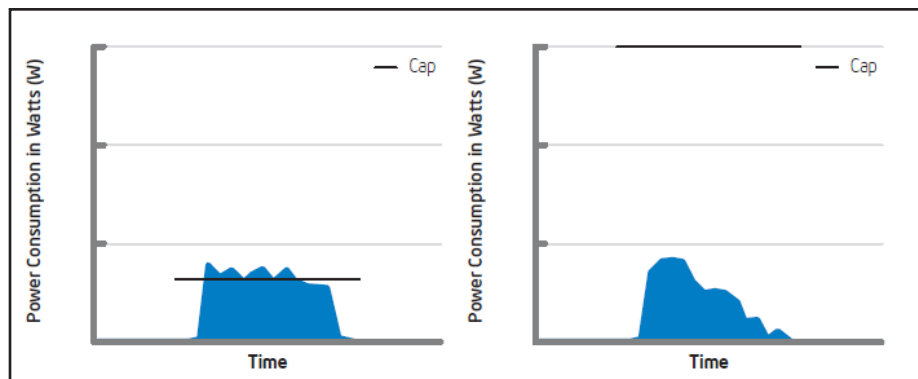


Figure 13: Capped and Uncapped Workload for I/O intensive Workload

Composite - Policy Based Power Management Use Cases

The opportunity to reduce energy usage by power capping alone is limited. For significant energy reduction, sustained power cuts are needed over time. If the policy in effect is capping as a guard rail, the policy seldom kicks in, if at all. Some energy savings are possible under a permanently capped regime, but these are limited by the capping range, or by the need to remove the capping policy to optimize performance yield.

Policies under dynamic power management take advantage of additional degrees of freedom inherent in virtualized cloud data centers as well as the dynamic behaviors supported by advanced platform power management technologies. Power capping levels are allowed to vary over time and become control variables by themselves. Selective equipment shutdowns enable reductions in energy consumption, not just power management. The tradeoff for dynamic policies is additional complexity: if the capping level becomes a control variable, this means a mechanism to exert this control needs to be implemented.

Cloud service workloads may exhibit a more or less predictable pattern, with demand peaks during office hours and deep valleys in the small hours of the morning. In fact, it is not uncommon for demand to vary as much as 10:1 through the day.

Imagine a virtualized cloud workload that takes seven servers to run during peak demand with the 10:1 variance mentioned above.

If the seven servers run 24/7 as is the norm in most data centers today, even if we apply power capping to the lowest possible policy, the best we will do with current technology is to bring power consumption down to 50-60 percent of peak power consumption. This mode of operation is very inefficient during demand valleys considering that the workload demand might be less than 10 percent of peak. This is why most traditional data centers end run at an abysmal 10-20 percent of utilization.

Ideally, if the power consumption per unit of workload demand remained constant, when workload demand drops to 10 percent of peak, so would the power consumption. This concept is known as power proportional computing. There is a bottom for power proportional computing for every known technology. For the present generation of servers, the bottom for an idling server lies at around 50 percent of peak. This means a server that is powered up but doing no work consumes 50 percent of its peak power.

Fortunately, there are additional server states we can exploit under these circumstances. If we know that a server won't be used for a period of time, we can put it to sleep. To be precise, we can put it into ACPI S5 (soft off) or even ACPI S4 (hibernation). A management application can put a server to sleep when not in use and restart it as needed. A sleeping server makes it possible to reduce power consumption by more than 90 percent of peak.

In a common real life analogy, when we leave a room, we turn off the lights. If this is the sensible thing to do, why do we see servers blazing 24/7 in most data centers? This is because most legacy applications will break when the physical server is powered off. However this is no longer true in virtualized environments that allow for the dynamic consolidation

of virtual machines into fewer physical hosts during demand valleys and for their expansion during high demand.

Assume for the moment a workload that takes seven servers to fulfill. At any given time of the day, except for the periods of highest demand, there will be some servers turned off. These servers are said to be "parked." As stated earlier, server parking allows the extension of idle power from 50 percent of peak to 10 percent or less for a pool of servers. This is how we can attain real energy savings.

Power capping is still needed: when demand is lowest, the system may still be over-provisioned with one server running. An application of power capping can further trim down power consumption without undue degradation in QoS. Likewise, since servers are turned on in discrete steps, whenever one is activated the system will likely be over-provisioned. An application of power capping will allow the equalization of supply to demand. Also, the system may support multiple service classes; hence at any given time there may be two or more server sub-pools each allocated to a specific service class with an associated SLA. The total available power is allocated among the different service classes, and those with highest SLA receive the lion's share of available power. The simultaneous application of multiple use cases is called a composite usage.

Use Case Five: Data Center Energy Reduction through Power Aware Support for Multiple Service Classes

Purpose

Consider two service classes for workloads, namely: high and medium priority workloads. The high priority workloads run on unconstrained servers; they can take all the power they need to run as fast as they can. Medium priority workloads are assigned to power capped

servers. These will run more slowly, but they will still run. The customer is charged based on the class of the service chosen.

The main purpose of this usage model is to showcase the ability to enforce multiple SLAs across different populations of users.

Pre-requisites

- Set up a schedule of parked and working servers based on the expected daily cycle demand forecast. An hourly schedule may be sufficient for most workloads.
- Assign power quotas to the active server sub-pools depending on the classes of workloads. These quotas can be set based on the power demand forecast. More precise allocation is possible if the quotas are based on the application's key performance indicators (KPIs).
- Set up a mechanism to tag the workload to a particular service class and to forward the workload to the right set of ESX hosts.

Steps for Execution

- Learn and tune phase
 - Run the application through a few daily cycles with no power management mechanisms to establish the baseline power consumption. This means running the machines 24/7 with no power capping. Note the baseline energy consumption in this operating mode.
 - Establish the allocation schedule for parked and active server sub-pools. Re-run the workload to establish that there is no gross over-allocation or under-allocation. The allocation can be done by time-of-day or in more sophisticated schemes as a control feedback loop that uses KPI monitoring.

- Overlay the power capping schedule to establish the different service classes and perform power consumption curve shaping.
- Re-run the system for a few days to ensure there are no gross mismatches between the power allocation algorithms and workload demand.
- Execution phase
 - Deploy the system previously tuned and monitor the KPIs for a few weeks to ensure there were no corner cases left behind.
 - At this point the system can be released for production.

Results

Workloads run over a period of eight hours used approximately 25 percent less energy.

Things to Consider

Architectural Considerations

Storage

For the sake of cost effectiveness and simplicity, a single network file system (NFS) store (DAS) was used as a shared storage to store virtual machine images. For production deployments, other storage technologies and topologies may need to be used based on the performance, cost, and other factors including storage area networks (SAN) or network attached storage (NAS).

Scalability

Scalability should be considered on two different factors. First, we need to consider the scalability of the underlying infrastructure of VMware vSphere 4 (ESX4.x) that forms the test bed. With VMware vSphere 4.0, VMware has greatly enhanced the scalability of the product. Details on their performance enhancements can be found at http://www.vmware.com/files/pdf/vsphere_performance_wp.pdf.

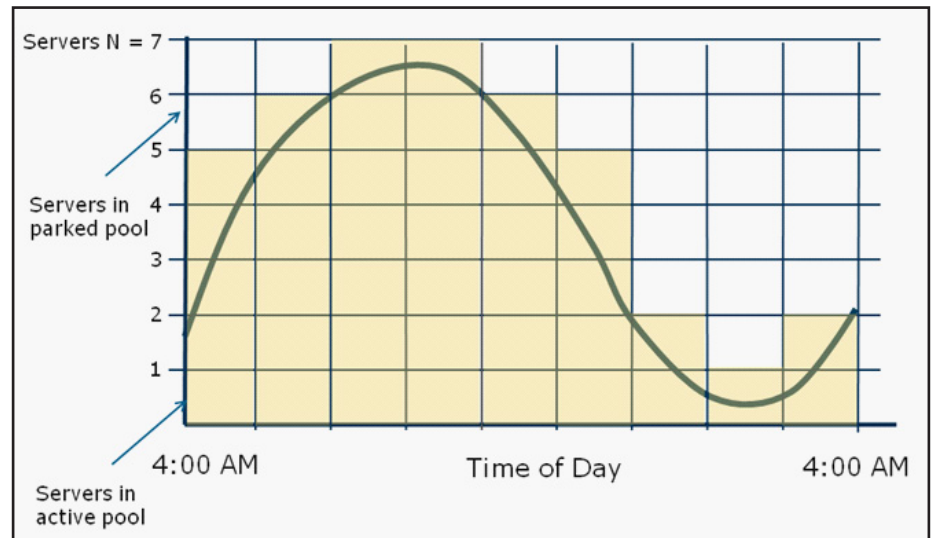


Figure 14: Daily Power Demand Curve and Servers in Active and Passive Pools

Second, on the Intel DCM front, a single stand alone server can manage up to 5000 VMware vSphere Hypervisor hosts.⁸

Networking

We used 1 Gb connections for the infrastructure test bed for both service console/virtual machine networks, and a 100 Mbps link was used to connect to BMC for out of band (OOB) power management. Depending on the customer requirements and usage, production environments might benefit from the use of 10 Gb networks for virtual machine networks.

Power Management

Power management should be considered after careful analysis of the workload performance under various power capping. As mentioned earlier, there are many usage models for which a power management solution would be very beneficial. At the same time, there are scenarios wherein power management may not be the right option. For example, if a highly-sensitive production workload is very processor intensive and the host is already highly utilized, the addition of a power cap would inadvertently affect the performance of the system.

Hardware

A full discussion of processor and overall server performance considerations is beyond the scope of this paper. However, it is important to note that the performance of virtual machines that run in a virtualized platform is heavily influenced by factors of processor architecture and specific feature sets available in the processor. The use of high performance server processors equipped with virtualization and I/O support feature sets, such as Intel Xeon processor 5500/5600 series, which also supports Intel Node Manager is strongly recommended. For more details on Intel® Virtualization technologies please refer to www.intel.com/technology/virtualization/ and download.intel.com/business/resources/briefs/xeon5500/xeon_5500_virtualization.pdf.

Additional Usage Models Under Development

Thermal Power Management

In today's data centers, cooling configurations are determined statically and cooling usage is not matched to the real time demand. Airflow into the rack is optimized for the peak usage and is determined by trial and error method.

With use of the Intel platform that exposes new thermal sensors and events, data center cooling can be optimized dynamically. With the use of real-time data from the group of servers about airflow, server inlet, and exhaust temperatures, optimal rack airflow and temperature can be provided for those servers.

Along with the optimization of cooling and use of the new sensor values, we can avoid platform hot spots when we distribute the workloads to cooler systems. This improves availability of servers because it prevents thermal trips. To determine the platform hot spots, "Thermal Margin" value is used. The higher the value of the thermal margin, the hotter the system is.

Development and Usage of Plug-in

The plug-in discussed in the paper provides the flexibility for Intel to expose its new platform capabilities and the associated usage models. The plug-in concept that is demonstrated can be extended to other hypervisor vendors through use of their respective integration APIs or any other ISV that provides cloud/virtualization management applications.

The plug-in used in the paper was developed to highlight the value of Intel Node Manager for certain usage models. Ideally this feature should be directly integrated into the hypervisor or other management applications/plug-ins to fully benefit from the platform feature, thereby it supports additional usage models and adds value to the end consumer.

The integration with Intel Node Manager can be accomplished through use of the Intel DCM API, which is exposed through a web service and can integrate into any enterprise software management console. Intel DCM in turn uses Intel Node Manager to manage power consumption on each node. Intel provides source code that implements the major usage.

Glossary

Intel® Intelligent Power Node Manager: Node Manager resides on Intel® Xeon® processor 5500 series server (and later) platforms. It provides power and thermal monitoring and policy based power management for an individual server. Capabilities are exposed through standard intelligent platform management interface (IPMI) interface from supported Baseboard Management Controllers (BMC). This requires an instrumented power supply such as PMBus*.

Intel® Data Center Manager: Intel Data Center Manager scales Node Manager functions to racks and groups of servers and enables IT users to benefit from increased rack density, and reduced capital and operational expenses.

VMware® Distributed Power Management (VMware DPM): <http://www.vmware.com/products/drs/> Provides additional power savings beyond server consolidation as it consolidates workloads dynamically even further during periods of low resource utilization.

SDK: Software Development Kit

QoS: Quality of Service

KPI: Key Performance Indicators

SLA: Service Level Agreement

APPENDIX A: Server Power Management

Intel Power Management Technologies

Micro-processors are possibly the most energy intensive components in servers and have traditionally been the focus of power management strategies. Emergent technologies such as solid state drives have the potential to significantly reduce power consumption and in the future, management of memory power consumption may be incorporated.

Intel Intelligent Power Node Manager and Intel DCM are designed to address typical data center power requirements such as described above.

Node Manager is implemented on Intel® server chipsets starting with Intel Xeon processor 5500 series platforms.⁹ Node Manager provides power and thermal monitoring and policy based power management for an individual server and is exposed through a standards based IPMI interface¹⁰ on supported Baseboard Management Controllers (BMCs). Node Manager requires an instrumented power supply that conforms to the PMBus standard.¹¹

Intel DCM SDK 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 to provide high value power management features. Intel DCM capabilities can be found on OEM platforms that support Node Manager. These technologies enable new power management paradigms and minimize workload performance impact.

Intel Intelligent Power Node Manager

Intel Xeon processors regulate power consumption through voltage and clock frequency scaling. Reduction of the clock frequency reduces power consumption, as does lowering voltage. The scale of reduction is accomplished through a series of discrete steps, each with a specific voltage and frequency. The Intel Xeon processor 5500 series can support 13 power steps. These steps are defined under the ACPI¹² standard and are colloquially called P-states. P0 is nominally the normal operating state with no power constraints. P1, P2, and so on aggressively increase the power capped states.

Voltage and frequency scaling also impacts overall system performance, and therefore will constrain applications. The control range is limited to a few tens of watts per individual micro-processor. This may seem insignificant at the individual micro-processor level, however, when applied to thousands or tens of thousands

of micro-processors typically found in a large data center, potential power savings amount to hundreds of kilowatt hours per month.

Node Manager is a chipset extension to the BMC that supports in-band/out-of-band power monitoring and management at the node (server) level. Some of the key features include:

- Real-time power monitoring
- Platform (server) power capping
- Power threshold alerts

Figure 15 Node Manager power management closed control loop shows the Intel NM server power management closed control loop.

VMware* Distributed Power Management (DPM)

Consolidation of physical servers through virtualization has well known cost advantages because it reduces hardware maintenance and power consumption. VMware* Distributed Power Management (DPM)¹³ can provide additional power savings through dynamic consolidation of workloads during periods of low resource utilization. Virtual machines are migrated onto fewer physical servers

and unneeded hosts can be powered off. When workload demands increase, hosts are powered back on and virtual machines are migrated back to them. VMware DPM is an optional feature of VMware Distributed Resource Scheduler (DRS).

The combination of VMware DPM and Intel NM technology can help implement the power management usage models discussed in this paper. Integration of VMware and Intel NM was achieved through a plug-in developed by Intel.

APPENDIX B: Dell* Cloud Solutions

Dell offers solutions and services for customers to build cloud computing data center installations. These solutions and services include:

- New integrated cloud infrastructure solutions comprised of pre-tested, pre-assembled, and fully-supported hardware and software, as well as services that enable public and private cloud builders to easily and quickly deploy and manage cloud infrastructures with confidence.
- New Dell Cloud Partner Program works with leading independent software vendors (ISVs) to offer customers easy-

to-buy and deploy cloud solutions and blueprints optimized for Dell platforms.

- New Dell PowerEdge* C-series servers with hyperscale-inspired design for high-density compute performance with efficient energy consumption and low operational costs.
- New Dell Cloud Services including consulting, deployment and support help customers plan, manage, and reduce complexity when they adopt cloud computing models.

Dell PowerEdge C-Series Servers

Dell has extended its PowerEdge server family with the new C-Series. Designed with inspiration from Dell's DCS business, these new servers are feature-optimized and power-optimized for customers in HPC, Web 2.0, gaming, social networking, energy, SaaS providers, and public and private cloud builders. The new PowerEdge C servers include:

- **PowerEdge C1100:** increased-memory, power-efficient, cluster-optimized compute node server (1U/2S, Up to 144 GB RAM, Intel Xeon processor 5500, 2 x 1 GbE Intel 82576 Kawela cards): Great for power and space sensitive customers requiring maximum memory flexibility.

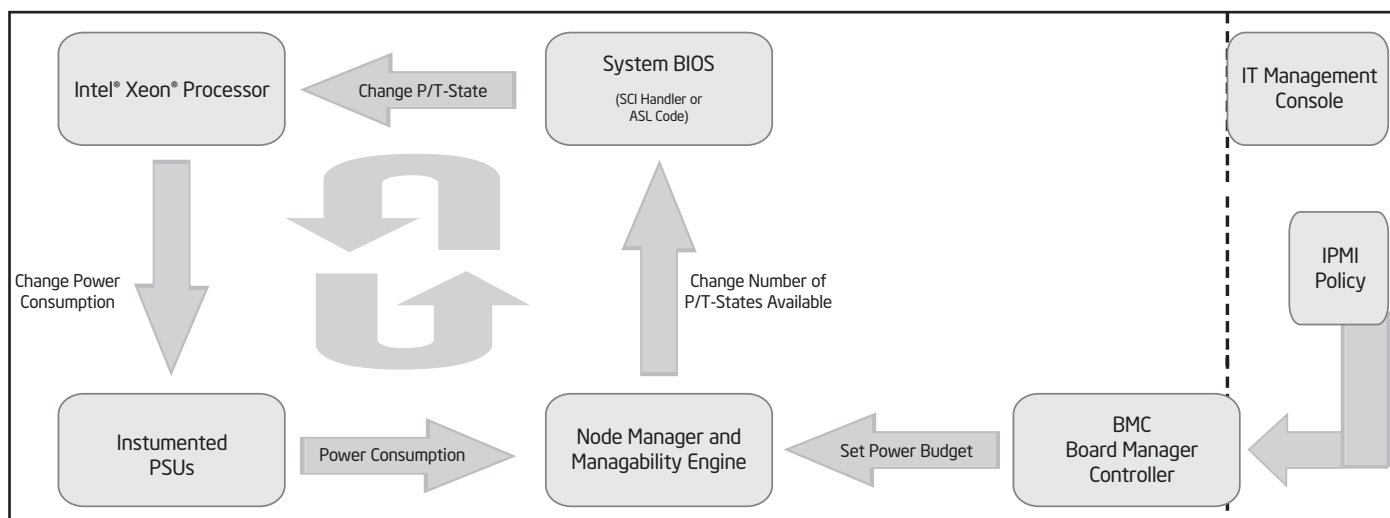


Figure 15: Node Manager Power Management Closed Control Loop

- **PowerEdge C2100:** high performance data analytics, cloud compute platform and cloud storage server (2U/2S, Up to 144 GB RAM, Intel Xeon processor 5500/5600 series, 2 x 1 GbE Intel 82576 Kawa cards): Great for scale-out data center environments where memory and storage density matter most: Hadoop, Map/Reduce, Web analytics, database.
- **PowerEdge C6100:** 4-node cloud and cluster optimized shared infrastructure server (2U/Up to 4 2S server nodes[hot-serviceable], Intel Xeon processor 5500/5600): Great for Hyperscale-inspired building block for high-performance cluster computing (HPCC), Web 2.0 environments and cloud builders where performance is key.

APPENDIX C: VMware Power Management Plug-ins

VMware offers third-party developers and partners the ability to extend the functionality of VMware vCenter Server through VMware vSphere Client Plug-ins¹⁴ with their own product-specific menu selections or toolbar icons that provide access to external, Web-based functionality. Extension Manager is the service interface that supports centralized management of plug-ins (sometimes called extensions) for VMware vCenter Server systems.

Plug-ins have both server and client components. The server component extends VMware vCenter Server functionality. For example, the Storage Monitoring Service provided with VMware vCenter Server uses information from the VMware vCenter Server database to provide reporting level services about storage. Text and graphical user interface (UI) items extend the VMware vSphere Client to support user interaction with the server. For example, reports about storage are available on a storage tab and several context menus in the VMware vSphere Client.

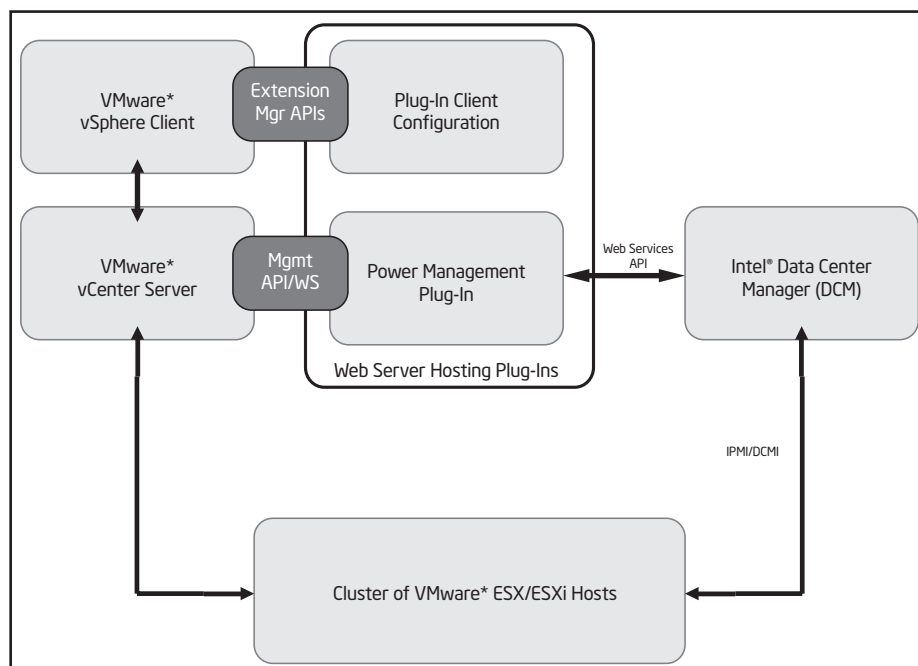


Figure 16: Power Management “Plug-in” Architecture

Power Management Plug-in

In today’s virtualized environments, resource optimization decisions are mainly based on processor and memory utilization which is dependent on the workload. The Power Management plug-in extends the ability of resource management engines for cloud and virtualization environments like VMware vCenter Server to provide a highly optimized resource distribution mechanism that utilizes Intel NM and Intel DCM enabled platform capabilities.

The main concept behind this plug-in is to demonstrate how power consumption of the server can be used as a key parameter for resource distribution and migration decisions that use vSphere Web Services and Intel DCM SDKs. Some of the hypervisor management tools including VMware vCenter Server do not consider power as a parameter in their algorithms for resource distribution. The plug-in demonstrates migrations from a power constrained host to an unconstrained host through the use of VMware vSphere Web

Services and Intel DCM SDK. The plug-in allows the administrator to use its UI to set power policies on the server with Intel DCM SDK instead of with a separate tool for power management. Also the plug-in allows for the administrator to set up power policies either at a server level or at a rack level.

Figure 16 depicts the high level architecture of the plug-in. Components on the left are VMware vCenter Server and VMware vSphere Client. VMware vSphere Client is the front-end for VMware vCenter Server. The middle section is made up of the power management plug-in (web applications) that were developed. The plug-in is registered with VMware vCenter Server through the Extension Manager interface. The right side section of the diagram shows a server with Intel DCM installed. This application is used to manage power consumption of servers with Intel NM enabled. This DCM tool can co-exist on the same server as VMware vCenter Server. For clarity though, they are shown in different boxes. The bottom section

represents the cluster of the ESX hosts that support Intel NM technology. The functionality of the plug-in is exposed through the VMware vCenter Server and Intel DCM APIs.

APPENDIX D: Retrofitting Power Management

The deployment of a power management capability in a data center will almost certainly occur in the context of an existing application. The luxury of defining the power architecture in a green field environment is truly rare. Hence any power management architecture needs to have flexibility to conform to most any pre-existing environment if it is to stand any chance of deployment. A convenient abstraction to integrate a power management capability is that of a pluggable module. This abstraction can take on multiple guises: at a single node level through use of the IPMI or data center manageability interface (DCMI) standard or at the node pool level. Software entities such as Intel DCM SDK or VMware vCenter Server can be used to aggregate the power management capabilities of a pool of nodes into a single logical entity. A specific plug-in is inserted into an application through IPMI or DCMI for single nodes or through DCM or VMware vCenter Server for a group of nodes.

An example of a plug-in with VMware vCenter Server is shown in Appendix C. Next we show an example of a plug-in with DCM. Assume cloud storage application running on a collection of servers in a rack. The architecture is

scalable, which allows an implementation to be extended into multiple racks or rows of racks. Also assume, for simplicity, direct attached storage devices (DASD).

1. The cloud storage application may be implemented with multiple physical racks arranged in a row. We postulate a management application that regulates the power behaviors of the appliances in that particular row, the row power manager. The row power manager may be connected to the building management system (BMS) application, whose role, among others, is to oversee power management across the whole data center.

The row power manager implements a number of policies that get mapped into specific directives to each storage appliance. For instance, if a row is powered by a single 5 KW branch circuit, the row power management imposes a 5 KW power guard rail for each of the storage racks. Without it, an extra branch circuit is required to lead to significant stranded power.

2. In addition to the application of power management-specific technologies, there are emergent technologies that reduce power consumption. A case in point is the replacement of hard drives with solid state drives (SSDs). SSDs consume less than 1 watt at idle and less than 5 watts in general, as opposed to 10 to 15 watts typical of mechanical hard drives. Hence, the use of SSDs will yield a lower Pbaseline.

3. The rack policy engine oversees its power consumption as it monitors the power draw from the power distribution unit (PDU) that feeds the hard drives and the power consumption. It uses the server subsystem as reported by the instance of Intel DCM that regulates the power consumption of the servers.
4. The implementation may provide a monitor-only capability for the storage subsystem, in which case the appliance policy engine needs to meet the power quota for the appliance through regulation of power consumed by the servers in the rack.
5. The rack policy engine in the node control subsystem in the storage application assigns a power target to Intel DCM. This power target can change dynamically based on workload conditions and the policies set at the higher levels.
6. Intel DCM takes the overall power quota for the server subsystem and divides it across the servers in the rack.
7. Node Manager instances set the P-states accordingly to meet the quota imposed by Intel DCM.

The role that Intel DCM plays in component number 5 in Figure 17 is crucial in this reference architecture: it allows power-enabling an application at a fraction of the cost that it would take to implement node control algorithms within the application. In our experience, the

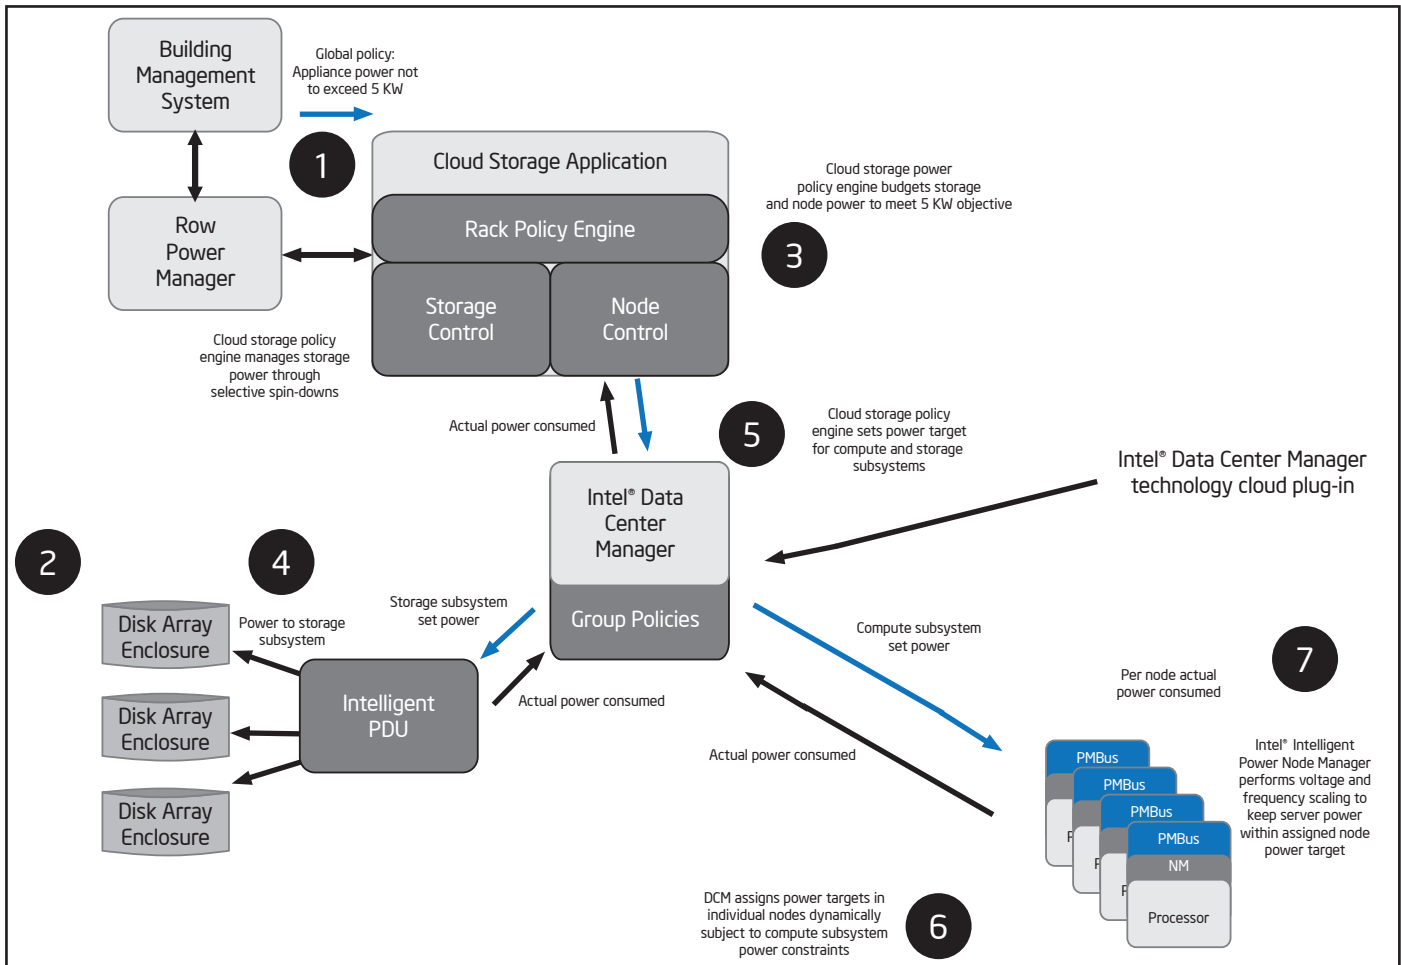


Figure 17: Power Management “Plug-in” Architecture with Data Center Manager

interface of an application to the Intel DCM usually takes less than a week even with the implementation team not previously exposed to the API. All it takes is a few web services calls. Most of the effort goes into the validation of the new capabilities. Because of the small effort involved in the interfacing, Intel DCM, for practical purposes, functions as a plug-in module to quickly add a power management capability to a cloud-

based application, in this case, a cloud storage application. This capability is added without the need to re-architect the original application in any fundamental way.

A final consideration of retrofitting relates to the link up of a power management capability to key performance indicators and service level agreements. Many applications, especially those associated with service oriented industries such

as telecoms, are already wired to monitor key performance indicators (KPI). Given that power management essentially allows an explicit tradeoff between performance and power consumption, power management plug-ins complement these applications nicely: once a plug-in is inserted, it is easy enough to monitor the effect of power management on the KPIs and to ensure the application performance stays within the agreed service level agreement.

Endnotes

1. Intel® Intelligent Power Node Manager, <http://www.intel.com/technology/intelligentpower/index.htm>
2. Intel® Data Center Manager, <http://software.intel.com/sites/datacentermanager/index.php>
3. EPA Report to Congress on Server and Data Center Energy Efficiency, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Report_Exec_Summary_Final.pdf
4. Dell PowerEdge C Series, <http://www.dell.com/content/topics/topic.aspx/global/products/landing/en/poweredge-c-series?c=us&l=en&s=gen&redirect=1>
5. Windows Hardware Compatibility List, <http://www.microsoft.com/whdc/hcl/default.mspx>
6. Creating a new MIME type for ASPX processing, <http://support.microsoft.com/kb/326965>
7. VirtualCenter VMotion Requirements: http://pubs.vmware.com/vi3/resmgmt/wwhelp/wwhimpl/common/html/wwhelp.htm?context=resmgmt&file=vc_create_cluster.7.4.html
8. Intel DCM Scalability, <http://software.intel.com/sites/datacentermanager/datasheet.php>
9. Intel® Xeon® Processor 5500 Series, <http://www.intel.com/itcenter/products/xeon/5500/index.htm>
10. Intelligent Platform Management Interface, <http://www.intel.com/design/servers/ipmi/ipmi.htm>
11. PMBus*, <http://pmbus.org/specs.html>
12. Advanced Configuration & Power Interface, <http://www.acpi.info/>
13. VMware Distributed Power Management Concepts and Use, www.vmware.com/files/pdf/DPM.pdf
14. Getting Started with VC “plug-ins”, http://www.vmware.com/support/developer/vc-sdk/vcplugin/vcplugin_technote_exp.pdf

To learn more about deployment of cloud solutions,
visit www.intel.com/cloudbuilders

Disclaimers

Δ Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

Σ Intel Internal Measurements. Performance results are based on certain tests measured on specific computer systems. Any difference in system hardware, software or configuration will affect actual performance. For more information go to <http://www.intel.com/performance>

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, Go to: http://www.intel.com/performance/resources/benchmark_limitations.htm

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 © 2010 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, Xeon inside, Intel Intelligent Power Node Manager, Intel Data Center Manager, and Intel Virtualization Technology are trademarks of Intel Corporation in the U.S. and other countries.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

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

