

Demonstrating Cisco UCS Manager, Power Assure® and Intel® Data Center Manager to Power Manage Cisco UCS Racks and Blades

A Cisco, Power Assure and Intel Corporation Proof of Concept

Summary

Datacenter power consumption has been increasing over time and this in turn has increased focus on how best to manage and bring this aspect of the IT facility under greater control.

IDC estimates that the energy expense associated with powering and cooling the worldwide server installed base increased 31.2% over the past five years ^a.

There is a need to dynamically control and manage IT resources more closely, to avoid unnecessary data center build outs and operational expense. Below we will briefly describe the overall environment, several use cases for utilizing power management to gain better control and utilization of IT resources and a proof of concept solution jointly developed with Cisco, Power Assure and Intel to show how specific servers can be managed. This is being shown at Intel Beijing Intel Developers Forum (IDF) in April 2011 for the first time.

Market Environment

Datacenter costs are growing rapidly, and over the past few years power and cooling costs, both capital and operating, have emerged as one of the dominant costs in the data center. Estimates from Microsoft indicate that the largest capital expenditure for a new data center is the power and cooling infrastructure.^b

Operational expense related to power is also growing at an increasing rate; IDC estimated that the energy expense associated with the worldwide server installed base represented \$27.5 billion of customer spend in 2010. Operational expense for servers is projected to run at approximately 60% of new server spending, albeit with much higher overall total server performance over time.^c

Worldwide Expense to Power and Cool the Server Installed Base, 2007–2013 (\$M)

	2007	2008	2009	2010	2011	2012	2013
New server spend	59,140	57,050	44,900	44,654	46,123	47,053	48,192
Power and cooling expense	32,003	37,898	29,431	27,524	26,541	28,070	28,002
Share (%)	54.1	66.4	65.5	61.6	57.5	59.7	58.1

Note: See Table 4 for key forecast assumptions.

Source: IDC, 2009

Along with the growth in costs there has also been a greater focus on sustainability and power management. IDC's 2010 Data Center Report predicted that, "sustainability and power management will become critical green IT differentiators across the system infrastructure software landscape"^d and in the Q2'10 Emerson Data Center User Group Survey, adequate

power monitoring in the data center manager was cited as the number one concern, followed by heat density and availability. ^e

An area of focus for the industry has been the low utilization of the servers - in general servers are kept fully operational even when under-loaded. Virtualization technologies have made some inroads into this problem, but without a management infrastructure that can comprehend both the business operational requirements and the power consumption impact of changes to the environment most datacenter operators will simply default to leaving all devices powered on at all times to ensure availability. In fact, EMA research indicates that, "on average, data center servers are kept fully operational 88% of the time and the reason most often cited for this is simply that the business production environment is utilized 24/7." ^f

Businesses must maintain their operations to meet commitments, but this comes at an operational cost, which potentially can be monitored and controlled more precisely. This provides the opportunity to maintain the same (or better) business performance while reducing operational cost. We will explore some of these areas in this paper.

Cisco / Power Assure/ Intel – Proof of Concept implementation

Overview description

This Proof of Concept implementation demonstrates a joint solution from Power Assure, Cisco and Intel to deliver real-time power monitoring and dynamic power saving policies with Cisco Unified Computing System blade and rack servers using the Cisco UCS XML API, Intel Data Center Manager and the Data Center Manageability Interface (DCMI) specification. Power Assure interoperates with Cisco UCS via the XML API and DCMI and provides customers with increased insight and control into their power utilization. The benefits are enhanced business continuity and quality of service delivery, while decreasing power consumption for their data centers.

Power Assure DPM and DPO

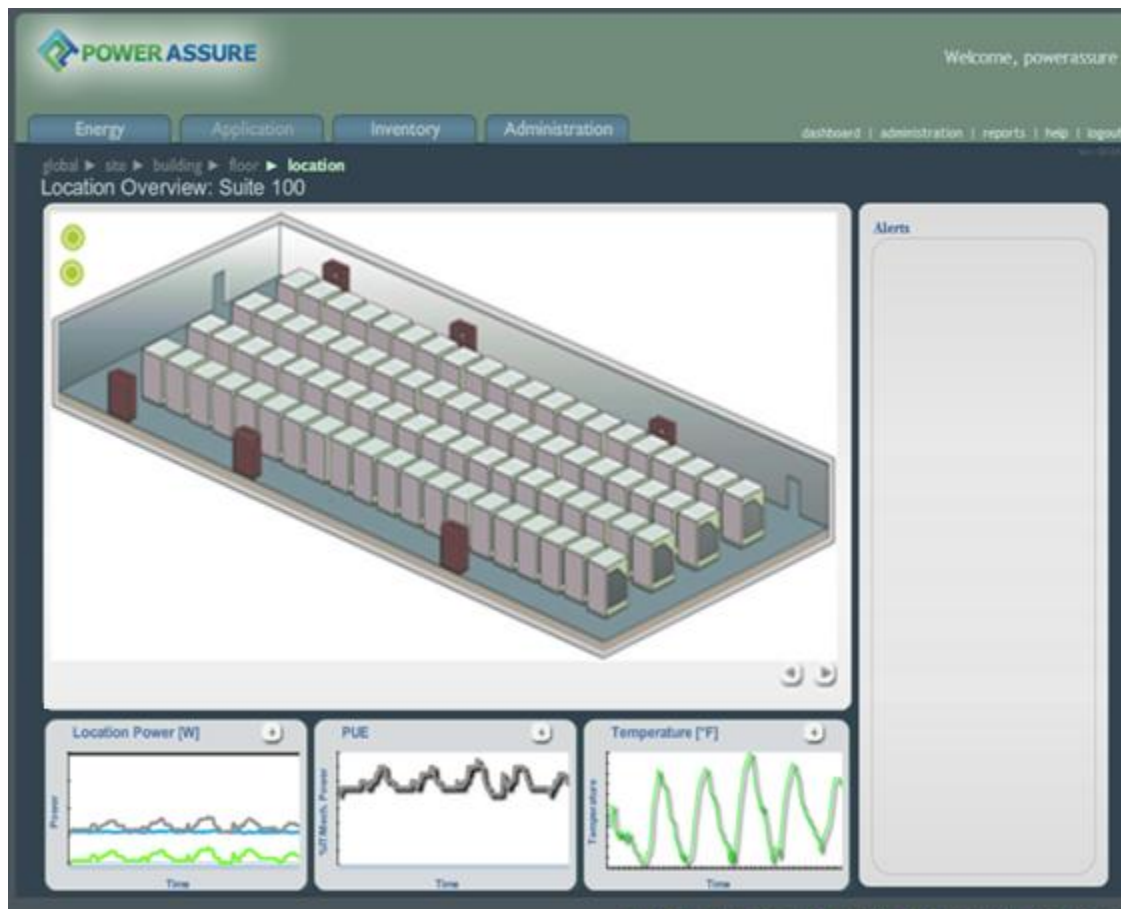
Power Assure software is a Data Center Infrastructure Management (DCIM) solution which provides dynamic power management (DPM) and dynamic power optimization (DPO) for your data centers. The software provides actionable, real-time and historical insight into power utilization and other key performance indicators (KPIs) for the data center, along with the means to optimize energy efficiency while maintaining service levels. Power Assure enables data centers to gain real-time infrastructure insight and intelligence, maximize use of available space and power; extend the useful life of the data center; cut on-going power consumption by an average of 50%; and shift load within and across sites.

The Power Assure DPM software provides insights into all aspects of power consumption in your data center, including utility feeds, generators, power distribution units, power meters, and individual servers. The software is agentless and integrates with virtually any device, sensor, and management system to gather and track relevant information real-time such as energy consumption, temperature, cooling levels, server application load, and available server capacity.

In this demonstration, Power Assure software interoperates with Cisco UCS Manager for the B Series Servers and Intel DCMI for the C Series Servers. The software features a web-based dashboard that allows you to visualize and correlate power, capacity/load, and environmental information from throughout the data center.

The Power Assure DPO software further allows you to adjust capacity and associated resources to reduce power consumption, address power/cooling emergencies to ensure business continuity, and increase efficiency of your data center. For example, DPO enables tight control into the Cisco UCS infrastructure policies to dynamically provision server and application services

such as power capping or shed and shift load. DPO further automates the adjustments through its patent-pending runbook automation to ensure mistake-free operations. ⁹



The Power Assure Dashboard

Cisco UCS

Cisco UCS was architected from the ground up to be a stateless solution for flexible compute environments. It is based around the following key elements:

- An open-standards, Intel based x86 blade-server architecture that embeds a 10GE fabric throughout the server, storage access, network and management infrastructure, providing full 10GE connectivity to each physical server.
- Unified Fabric capabilities, for network, storage access and management, which reduces power consumption by driving cable consolidation and reducing the number of power-consuming components.
- Single point of device management for the entire compute cloud, allowing 3rd party or custom frameworks to manage/monitor all aspects through an XML API. By integrating with this API Power Assure is able to report power information and take advantage of the stateless nature of UCS to migrate and manage workloads such as power cap based on power and cooling metrics.
- A stateless design where all server configurable elements such as firmware revisions, BIOS settings, network and storage identities etc are abstracted into a service profile which can be applied to any suitable server within the environment. This allows for rapid re-purposing and re-provisioning of workloads based on business needs.

This comprehensive set of innovations offers the following tangible benefits in approaching the design of their multi-site data centre:

- Scheduling of workloads on to servers (physical or virtual) within the compute cloud can be based on business policies rather than physical machine/connectivity constraints.
- Network security policies are enforced centrally, preserving SLAs at the logical level, while allowing flexible mobility of virtual workloads between physical resources.
- Overall power consumption is significantly reduced, by greatly reducing the number of components (switch modules, management modules, adapters) requiring power.
- Current investments in SAN infrastructure are preserved, using FCoE to provide lossless bridging of Fibre Channel and Ethernet traffic.
- Operational expenses are reduced by having compute allocation more easily driven through policy-based automation.

Cisco's Unified Computing System (UCS) integrates low-latency Unified Fabric with enterprise-class, x86 based servers, creating an integrated, scalable, multi-chassis platform where all resources participate in a single, unified management domain.

Cisco's UCS enables data center servers to become more dynamic and agile, where the server's personality and workload can be dynamically assigned or migrated to any physical server within the system, regardless of cabling. New computing resources can be deployed "just-in-time". Traditional physical and virtual workloads can be easily migrated between servers through remote management, irrespective of physical connectivity. Hardware resources are completely decoupled from logical workloads.

Cisco's UCS capabilities allow customers to build dynamic data-center clouds that provide on-demand provisioning, access and management of heterogeneous pools of physical resources. Cisco's vision is to provide the best possible, most cost-effective server/network/storage infrastructure, enabling customers to easily scale physical capacity, to be more agile in managing applications and workloads, and to allow customers to align resources based on business policy rather than physical connectivity constraints.

Intel DCM description

Intel® Data Center Manager (Intel® DCM) SDK is a power management solution middleware stack for the data center. It provides accurate, real time power and thermal monitoring as well as precise power management policies all the way from an individual server, logical or physical group of servers designated by IT policy, rack-level to row and control policies and will even provide monitoring of select IT equipment such as PDUs in the data centers. It is a capability that is useful for both IT and facility administrators and allow them to work jointly to reduce the Energy footprint. Intel® Data Center Manager can help enable various use cases:

Intel® DCM can help and address various use cases:

- Power and thermal monitoring allows you to get real-time, accurate power and thermal consumption data, manage data center hot spots, plan and forecast power usage and potentially replacing costly instrumented power strips
- Increasing rack density to maximize server count per rack in a fixed rack power envelope to increase DC utilization
- Power optimizations on specific workload types allows you to provide the optimized power profile per server/rack/floor or workload/application and cut electricity costs
- Business continuity allows to continue or prolong operation during power outage

Data Center Management Console Vendors (ISVs), System Integrators (SIs) and Data Center IT organizations can easily integrate Intel® DCM into their console or command-line applications and provide high value power management features to Large IT organizations, cloud providers and Data centers. DCM is a secure, Web Services interface driven SDK, based on validated implementation of IPMI, DCMI and SNMP protocols.

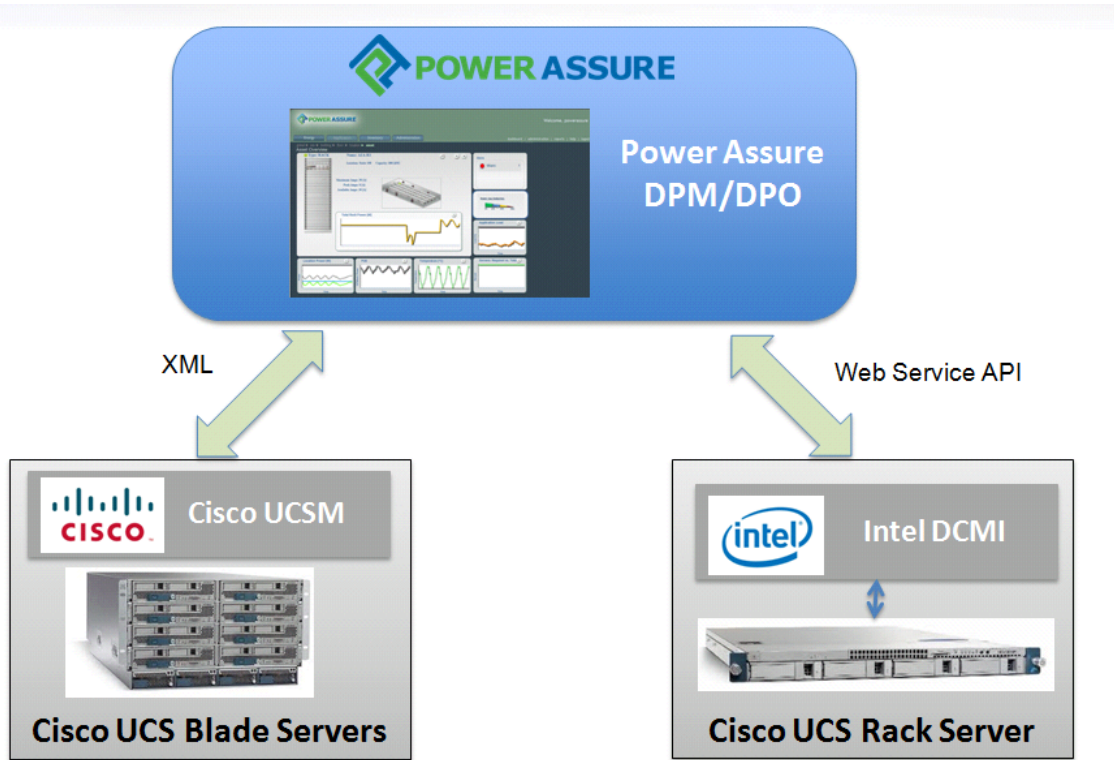
DCMI (Data Center Manageability Interface) Capabilities

DCMIⁱ defines the manageability requirements as a set of standard manageability features and interfaces. By using IPMI 2.0 as the basis, DCMI implementations can take advantage of the existing industry infrastructure of IPMI management controllers, firmware, and expertise. In many cases, the server platform hardware support for DCMI is already present and only the management controller's firmware needs to be changed to support DCMI. The capabilities of DCMI are available to both local and remote management software. These capabilities include:

- Platform Power on/off/reset control - The ability to control platform power and reset independent of host software. This capability is typically used to support remote provisioning and emergency recovery under situations where this cannot be accomplished via local software.
 - Standardized Local and Remote Access Interfaces - DCMI defines standardized interfaces that enable management applications to access the DCMI platform management subsystem locally(also referred to as "in-band" access) or remotely by direct communication with the platform management subsystem over LAN independent of the system's power state or host software (also referred to "out-of-band" access). The platform management subsystem operates on standby power so that it is always available for remote access.
 - Baseline Sensors - DCMI defines a baseline set of required sensors including Inlet Temperature, CPU Temperature, Ambient Temperature and Power Reading.
 - Platform Identification - DCMI defines interfaces for retrieving system ID and asset ID to provide a common mechanism for obtaining positive and unique identification of the managed system.
 - Improved sensor access - The IPMI definition of sensor access requires large number of command transactions and parsing at the management application to discover and/or read the sensors. DCMI defined a small extension command that greatly reduces the steps for accessing the mandatory sensors.
 - Authentication, integrity, and confidentiality -Remote session authentication, integrity, and confidentiality are provided using proven industry hashing and encryption algorithms in conjunction with the RMCP+ session establishment protocol defined by [IPMI 2.0].
 - Platform Event Logging - Sensor 'out-of-range' events and other platform hardware status events are logged to a non-volatile storage area for local or remote retrieval, enabling 'post mortem' and non- real-time analysis of platform events
- Power Assure integrates with Intel DCM for power and thermal monitoring as well as power optimization (such as power cap) to manage and control workloads and reduce power consumption for business continuity, increased server utilization and power savings.

Demonstration Environment

Overview Architecture



Demonstration configuration

<i>Components</i>	
<i>Power Assure</i>	<i>Power Assure DPM and DPO Software delivered as a service</i> <i>Web-based dashboard</i> <i>Power Assure Data Collection Appliance</i>
<i>Cisco</i>	<i>Cisco UCS C210 M2 Rack Mount Server managed via Intel DCMI</i> <i>CISCO UCS 5108 Chassis and B200 Servers managed via UCS XML API</i>
<i>Intel</i>	<i>Intel Data Center Manager 2.0 SDK</i>

Use Case Overview

DCIM solutions, such as Power Assure, measure actual power consumption in data centers down to the granularity of the individual power outlets or servers. This information can be used to improve power efficiency, optimize virtualized infrastructures and make prudent choices during server refresh cycles.

Power and thermal monitoring and event management

Power monitoring is a critical capability that enables workloads to be dynamically monitored and characterized. Real-time power consumption is gathered via the DCMI and UCS XML API interfaces and stored in the Power Assure database. This data can then be used to identify opportunities to increase data center energy efficiency, by measuring the actual power over time related to workload characteristics. Examples would include:

Identifying idle servers

In large server estates there may well be servers that are either unused or that are only used very rarely. Servers that are idle typically will consume 30% to 50% of their maximum power consumption so consolidating lightly or infrequently used servers and/or powering off unused servers can save significant amounts of power and free up power for other systems. Power Assure DPM can monitor server power usages in real-time, identifies servers which are underutilized, and provides the analytic engines to do what-if scenarios of how much power you can save by powering off, power capping, or shed and shift loads. Power Assure DPO can then automate the action and control the Cisco UCS Servers to execute on these power efficiency policies.

Identifying high power usage applications

Applications have widely differing power consumption levels based on their usage profiles and loading. Knowing which applications have high power utilization enables the datacenter operator to more effectively plan where particular servers and services should be located within the datacenter so as to manage the power and cooling load more efficiently. Power Assure DPM can interoperate with Cisco UCS to monitor application power usage, identify high power utilization applications, and allow what-if scenarios on how to shed and shift application loads for the best utilization and energy efficiencies to meet service level requirements.

Power Aware Workload Placement.

As datacenters become more dynamic with the ability to migrate virtualized workloads or UCS Service Profile enabled server re-purposing, understanding the power profiles of the various workloads enables power aware workload placement. This means that high power consumption workloads will not be placed in an area of the datacenter that is already operating under a power or cooling constraint.

Maximizing Data Center Efficiency by Power Capping

Power capping as a technology is not designed to save energy. The goal behind the various power capping implementations is to maximize the use of the power and cooling infrastructure by allowing a data center manager to install more servers in a rack, while ensuring that any power or cooling limits for that rack are maintained.

The Cisco Unified Computing System implements a powerful new concept called a power cap group. A power cap group is a collection of servers or chassis that share a common power allocation or budget, which is then applied in a dynamic manner to servers within the group. This provides the flexibility to power cap based on the service level requirements. ^h Note: This proof of concept did not demonstrate power cap groups, although this capability is part of the Cisco / Power Assure interoperability.

Power Assure is integrated with Cisco UCS to perform power capping. The following screen shot show how a power cap is configured in Cisco UCS and Power Assure dashboard shows the effect of the power cap.



Future directions

Over time, a far deeper and richer set of usage models will provide IT managers with more granular understanding of IT assets under control. For example, more advanced integration will allow Power Assure DPO to implement power capping and throttling based on application loads and service level requirements. For example, mission-critical applications will maintain a higher power cap while lower priority applications may be capped at a lower level or even shut down when an outage occurs, providing mission-critical applications with more power during the outage. Other usages, secondary to power can be potentially deployed for example, cost reduction due to replacement/ displacement of intelligent power strips which has DCMI capabilities to allow better monitoring. Power Assure can then use the historical and trending data to provide analytics and what-if scenarios to analyze potential cost of ownership savings.

Summary

Power Assure integrates with Intel DCMI and Cisco UCS to present visualization of IT, power and thermal monitoring information. Its analytic and optimization engine can then process continuous what-if simulations and automation processes to enable IT managers to make fully-informed decisions on what actions to take in its UCS infrastructure. UCS, together with its DCMI capabilities, has the advanced features to execute these optimization processes.

This helps:

- Maximize data center efficiency and utilization
- Avert disaster for its business continuity plans
- Continue operation under impaired power availability conditions prioritized with multiple service classes
- Extend the life and capacity of their data center, avoiding unnecessary buildouts
- Reduce energy consumption, saving operational costs

References

- a – IDC Datacenter Energy Management: How Rising Costs, High Density, and Virtualization Are Making Energy Management a Requirement for IT Availability April 2010, IDC #223004
- b - C. Belady and G. Balakrishnan, Microsoft, Incenting the Right Behaviors in the Data Center
- c IDC Worldwide Server Energy Expense 2009 – 2013 Forecast – December 2009 #221346,
- d - IDC – Worldwide System Infrastructure Software 2011- Top 10 Predictions – December 2010 #225895
- e - Data Center’s Users Group – Q2’10 – Emerson
- f - EMA - Best Practices in Data Center Power Management – June 2010
- g – Disaster Recovery Journal: April 2011 - Data Center Infrastructure Management: A New Tool for Disaster Recovery Planning and Response
- h Cisco: Power Management in the Cisco Unified Computing System: An Integrated Approach
- i - Data Center Manageability Interface Specification – Intel/ Microsoft August 2010
- j - http://www.att.com/Common/about_us/files/b1414usiness_continuity/BusinessContinuity_2010_Summary.pdf

Contributors/ Authors

Tony Harvey – Cisco – Power and Cooling – Senior product manager
Ivy Lui – Power Assure - Dir. Product Development
Derek Collier – Intel – Data Center Manager – Product Manager
Dror Shenkar – Intel – Data Center Manager – Software Architect

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 the Intel Leap ahead logo, are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Copyright © 2011, Intel Corporation. All rights reserved.