WHITE PAPER
Data Center Efficiency and Optimization
Data Center Management

The Efficient Data Center

Using data center instrumentation and server refresh to optimize compute performance in constrained environments.

"Power and cooling, rather than floor space, are becoming the largest limitation in the data center."

– Kathleen Broderick
IDC Analyst
November 2010 whitepaper (IDC 225616)

David Jenkins
Server Technology Marketing Manager,
Server Platforms Group

Introduction

Data center managers will face a lot of challenges over the next few years. In many facilities, three pressing problems need to be solved in the immediate future: limited power, increasing cooling demands, and space constraints.

The growing urgency of power and cooling problems is apparent to anyone who has been operating a data center over the last 10 years. Power and cooling are becoming as expensive as the IT equipment itself. In fact, power and cooling-driven utility bills now represent 25 percent to 40 percent of the operational expenditures in most data centers.

What about space? Consider that corporate data is expected to increase by 650 percent in just five years, from 2009 to 2015.[1] The electronic containers for that data—files, images, packets, and tag contents—are growing 50 percent faster than the data itself.[2] 2.5 billion users will connect to the Internet over the next five years,[1] with more than 10 billion devices.[2] This will require eight times the amount of storage capacity, 16 times the network capacity, and more than 20 times the current compute capacity by 2015.[3]

Figures like these may seem overwhelming, but they cannot be ignored, especially given the sobering cost of building new data centers. Construction costs are projected to exceed $1,000 per square foot or $40,000 per rack.[4]

Obviously it's time to look for new efficiencies within the data center to help stave off the power limitations, cooling issues, and high cost of adding new space. With 81 percent of IT managers saying they will exceed capacity for power or space in the next five years,[5] there's no time to waste. Additional sources of data and points of control are needed to get the most out of every watt consumed within the entire facility.

Table of Contents

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

Optimizing Power Consumption Based on Actual Work

Power efficiency requires looking at all of the power consumed in the data center, not just the power consumed by IT. The goal must be to optimize power consumption at all levels for the work done. Power Usage Effectiveness (PUE), a metric developed by The Green Grid, lets you quantify overall efficiency, evaluate the effects of changes, and monitor performance.

PUE is calculated by dividing Total Facility Power by IT Equipment Power. Total Facility Power is defined as the power measured at the utility meter that is used in the data center. It includes the power used in the infrastructure for lighting and cooling the building, the power used by IT equipment, and the power lost in the distribution/conversion system. IT Equipment Power, on the other hand, is the power used to manage, process, store, or route data within the data center.

The ideal PUE score is 1.0, but essentially impossible to achieve. Today's data centers are more likely to have PUE values in the range of 2.0 to 2.5, where 2.0 means one watt of electricity is lost to power overhead for every watt used to fuel computing equipment. How much room is there for improvement? Consider that Google, which has built some of the most efficient data centers in the industry, achieved a PUE score of 1.12 on one of its most efficient data centers.[6]

Once a PUE is computed for a data center, ways to improve PUE can be evaluated and selected based on their potential for improving efficiency and delivering a solid ROI through efficiency and performance gains.

For IT equipment, an extremely important measurement to obtain is the actual power being consumed to do the work. In most data centers, this is different than what might be calculated for IT Equipment Power because the power allotted to each server in a rack and to the rack as a whole is often based on the power rating of each server's power supply. Most data centers build to the maximum specification of each server's power supply. But consider that in many cases servers with a 500-watt power supply might only consume around 250 watts while running their workload. In fact, that might be true for hundreds of servers in a facility. This suggests that facilities may be over-provisioning IT equipment and paying for capacity that is not performing real work.

The challenge here is determining how much power each server really requires to handle its workload. The goal in data center efficiency should be to optimize power for the work being delivered, whether it is IT power, power conversion, or non-IT power such as lighting and building cooling.

Instrumenting Servers to Better Understand Their Power Needs

Data center designers have often had to rely on inadequate power usage data in deploying servers and racks. Lacking actual power usage information for each server based on its intended load, they stay on the conservative side and populate server racks based on available server power information, using the declared power rating from the server nameplate. This nearly always leads to an overstatement of a server's actual power need, resulting in underutilization of the budgeted rack power and providing more cooling than is necessary for the actual loads.

What's needed is better power consumption instrumentation at every level of the data center, from the processor to the platform to the entire data center. Fortunately, servers that provide instrumentation measuring actual power consumption in real-time and as a time series are now available. Even better, technology is available that enables IT managers to cap platform power to keep it within a targeted power budget.

The Power of Power Capping

For many data centers, one important advantage of a server refresh is the opportunity to obtain the necessary system instrumentation and power capping tools to limit system power to what's required for the actual workload. Being able to control the power the processor uses is the first level of instrumentation, and is essential to data center productivity and efficiency.

Intel® Xeon® processor-based platforms, for instance, are available with an innovative solution known as Intel® Intelligent Power Technology that scales power to workload demand. An embedded feature of this technology, Intel® Intelligent Power Node Manager, enables dynamic power management of IT equipment. Used with operating system tools, Intel Intelligent Power Node Manager allows data center managers to set a power budget (through third-party management consoles) for a server. Enabled consoles can aggregate Intel Intelligent Power Node Manager-equipped servers to set budgets for a rack, a row of servers, or even entire data centers, for up to 40 percent denser deployments.[7,+]

Intel Intelligent Power Node Manager is an out-of-band (OOB) power management policy engine that enables regulation of power consumption (power capping) by adjusting processor and memory subsystem power. Intel Intelligent Power Node Manager works with the BIOS and operating system power management (OSPM) to regulate and dynamically adjust platform power while delivering maximum performance and power for a single node within that power envelope.

Several server manufacturers also offer server power capping solutions that enable IT managers to set power consumption thresholds. Intel Intelligent Power Node Manager's advantage is that it's a vendor-agnostic, closed-loop system that interacts with the data center management console and automatically controls the server to enforce system policies. Intel Intelligent Power Node Manager uses the processor, memory, and power supply instrumentation as the controls to report system power consumption and cap system power to whatever policy is defined by the management console.

"High computing density is the name of the game for us because it drives profitability," says Don Nalezyty, director of IT enterprise architecture at Oracle.

The entire solution makes it possible to move power from one portion of the data center to another dynamically throughout the day—depending on where the need for power or cooling is the most urgent.

Through third-party management consoles, Intel Intelligent Power Node Manager reports actual power consumption and allows IT managers to react accordingly. For instance, a console could establish a policy that continuously monitors virtualized server racks to determine if they are approaching their aggregate power capacity. If a rack is near capacity, the console could reduce rack power draw by migrating workloads to another virtualized server rack that has power headroom to spare. The result maximizes performance without exceeding power constraints. A second use case is close-coupled cooling—where the power consumed by the servers is reported to the HVAC system, which then scales cooling output to the current heat load demands generated by the IT equipment.

A third use case is increasing rack density. An IT manager might evaluate a situation for a few weeks and notice one server never goes above 200 watts. In such a circumstance, the IT manager could safely lower the budget from 500 watts to 300 watts without restricting the performance capability of the workload. By doing so, the IT manager will free up 200 watts of stranded power. Freeing up such stranded power in several servers in a rack can enable a data center to deploy more servers into a rack and remain within the rack's power budget.

To summarize, server power capping can be used to:

- **Increase rack density:** By maintaining server budget policies set according to real workload by a data center management system, power capping enables stranded (unused) power to be rescued, allowing more servers to be added to the rack.
- **Maximize work during power and thermal spikes:** By dynamically capping power to shed load, data center managers can keep critical jobs running during power and thermal events.
- **Reduce power consumption:** Having actual server power reported on instrumented systems enables cooling capacity to be more closely linked to actual rack temperatures.
- **Facilitate power-based load balancing:** Power consumption is becoming part of the load balancing equation in virtual environments. Power capping enables data centers to dynamically manage power- and thermal-constrained systems and racks.

Dynamically Managing Power and Cooling in the Data Center

Coupling solutions like Intel Intelligent Power Node Manager with certain management consoles enables data center managers to dynamically balance power at the server, rack, and data center level. Intel, for instance, offers a plug-in—Intel® Data Center Manager—that enables a management console to aggregate and control large numbers of Intel Intelligent Power Node Manager-enabled systems. This makes it possible to deliver the processing needs of a business within a set power budget and manage energy bills in a predictable, consistent manner. The entire solution makes it possible to move power from one portion of the data center to another dynamically throughout the day—depending on where the need for power or cooling is the most urgent.

Power and thermal monitoring capabilities can be especially helpful in fine-tuning modern cooling solutions such as chimney cabinets, cool aisle containment, liquid-cooled cabinets, hybrid solutions, or the trend toward warmer data centers. Just as fine-tuning power saves money, so does tailoring the cooling to the precise needs of each server, rack, cabinet, and row. Here instrumentation can be a real aid in providing real-time temperatures and power consumption related to actual loads and providing guidance on power capping to achieve cooler operating temperatures without revving up the cooling system.

The ability to monitor power usage on the server and rack level is also a powerful tool in assessing the actual savings obtained at the component level by selecting innovations such as more energy-efficient power supply units and voltage regulation modules. It can also help determine the cost-effectiveness of something like switching to solid-state drives—a move anticipated to reduce hard-drive consumption fivefold by eliminating the energy required to drive moving parts and by generating less heat.

The ROI of Server Refresh

In any data center looking to improve efficiency, servers are a great place to start, particularly if current servers lack power instrumentation and capping capabilities. As the productivity engines of the data center and the target of much of the power and cooling, servers offer some of the greatest opportunities for improvement.

The performance increases of the Intel® Xeon® processor 5600 series are a good example, delivering up to 1.6 times the performance for enterprise computing than the previous generation Intel® Xeon® processor 5500 series.[8,+] The real story here is not just raw performance, though. It's intelligent performance. The Intel Xeon processor 5600 series uses Intel Intelligent Power Technology to automatically scale server performance to workload needs while requiring up to 30 percent lower power compared to the previous generation.[9,+]

Taking advantage of performance achievements enables data centers to realize fast ROI on server refreshes. A good example comes from a study Intel did comparing servers five years apart in generation.

In this study, the performance delivered by 15 racks of the five-year-old equipment could be refreshed with one rack of servers based on Intel Xeon processor 5600 series, while also using up to 94 percent less space and consuming up to 95 percent less energy.[10,+]

The cost savings realized from replacing 500 five-year-old servers with 30 Intel Xeon processor-based servers—taking into consideration utility costs, software, and warranty—are estimated to be more than $100,000 per month, which would pay for the 30 new servers in approximately two months.[10,+]

By selecting such energy-efficient architecture, IT acquires room to grow and increases IT performance in the freed-up space. This demonstrates how a refresh can address both cost and performance.

The cost savings from replacing 500 five-year-old servers with 30 Intel Xeon processor-based servers ... are estimated to be more than $100,000 per month ...

Refresh to Drive Efficiency

Every data center is unique, but nearly every data center faces a constraint, whether it's power, cooling, or space. And nearly every data center needs to increase its processing capabilities to handle the ever-increasing data, transaction processing, and other demands required by the modern enterprise. For almost every data center using servers older than three years old, refreshing with new servers is an important first step. This step alone can alleviate all three of the above constraints and deliver fast ROI. What's more, the instrumentation available on these new servers at both the component and platform level can be used to get the most performance out of every dollar, kilowatt, and square foot that is in the data center.

Here are some recommendations for quickly improving data center efficiency and getting a handle on power management through instrumentation.

At the system level

- Determine what's in the data center.
- Refresh inefficient legacy machines.
- Deploy in their place proven efficient systems (based on industry metrics like SPECpower*) that provide advanced power management features and tools.
- Take advantage of energy-saving component choices such as energy-efficient power supplies. Such components quickly pay for their initial extra cost.

At the platform level

- Choose server form factors that meet your data center density requirements.
- Consolidate workloads to drive up system utilization.
- Use system instrumentation to perfect load balancing and load shedding strategies.

At the facility level

- Develop a top-down plan.
- Set a PUE target goal.
- Choose the right equipment, power, and cooling options to help achieve that goal.

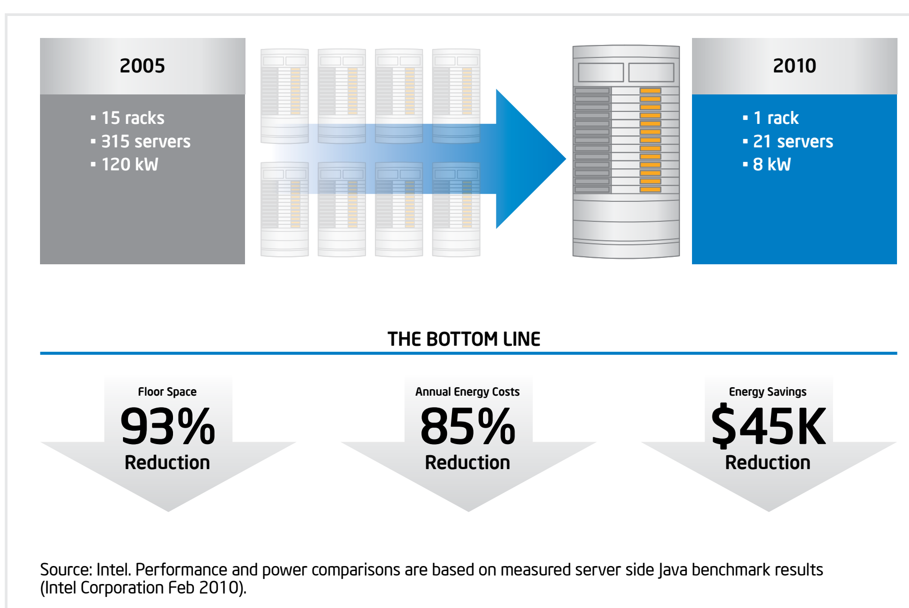

Figure 1. Example of the ROI of a server refresh.[10,+]

Additional Resources

Intel® Xeon® processor E3 family product brief
www.intel.com/Assets/PDF/prodbrief/325197.pdf

Intel® Xeon® processor 5600 series product brief
www.intel.com/Assets/en_US/PDF/prodbrief/323501.pdf

Intel® Xeon® processor E7 family product brief
www.intel.com/Assets/PDF/prodbrief/325213.pdf

Intel® Intelligent Power Node Manager and other power-saving technologies:
Power capping animation:
communities.intel.com/community/openportit/server/blog/2009/05/11/intel-xeon-processor-5500-instrumentation-enables-power-capping

"A Dynamic Approach to Power Budgeting." *Intel® Software Insight Magazine*, August, 2008.
softwarecommunity.intel.com/articles/eng/3802.htm

"Baidu Proof of Concept White Paper."
communities.intel.com/docs/DOC-1492

For additional information go to:
www.intel.com/content/www/us/en/servers/server-products.html

* Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark* and MobileMark,* are measured using specific computer systems, components, software, operations, and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. 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, visit Intel Performance Benchmark Limitations: http://www.intel.com/performance/resources/benchmark_limitations.htm

1 Gartner Group, "Hot Trends and Innovations in Data Centers."

2 "The Diverse and Exploding Digital Universe," IDC, March 2008.

3 Source: Intel Market Projections, 2009-2010. 8x Network: 800 terabytes per second of IP traffic estimated on internal Intel analysis "Network Supply/Demand 2010-2020" forecast; 16x Storage: 60 exabytes of data stored from Barclays.

4 "Datacenter Trends Survey," IDC, 2007.

5 Green Tech World, TMC 2007.

6 "Efficient UPS Aids Google's Extreme PUE," Data Center Knowledge, April 1, 2009.

7 Dynamic Power Optimization for Higher Server Density Racks – A Baidu Case Study with Intel® Dynamic Power Technology (Intel, 2008). For more details, see http://software.intel.com/sites/datacentermanager/intel_node_manager_v2e.pdf

8 Source: Internal Intel measurements for Xeon® X5680 vs. Xeon® X5570 on BlackScholes*. Up to 1.6x performance compared to Xeon 5500 series claim supported by a CPU-intensive benchmark (Blackscholes). Intel internal measurement. (Feb. 25, 2010)

- Configuration details: Blackscholes.*
- Baseline Configuration and Score on Benchmark: Intel pre-production system with two Intel® Xeon® processor X5570 (2.93 GHz, 8 MB last-level cache, 6.4 GT/sec QPI), 24GB memory (6x4GB DDR3-1333), 4 x 150GB 10K RPM SATA RAID0 for scratch, Red Hat* EL 5 Update 4 64-bit OS. Source: Intel internal testing as of February 2010. SunGard v3.0 source code compiled with Intel v11.0 compiler. Elapsed time to run benchmark: 18.74 seconds.
- New Configuration and Score on Benchmark: Intel pre-production system with two Intel® Xeon® processor X5680 (3.33 GHz, 12 MB last level cache, 6.4 GT/sec QPI), 24GB memory (6x4GB DDR3-1333), 4 x 150GB 10K RPM SATA RAID0 for scratch, Red Hat* EL 5 Update 4 64-bit OS. Source: Intel internal testing as of February 2010. SunGard v3.0 source code compiled with Intel v11.0 compiler. Elapsed time to run benchmark: 11.51 seconds.

9 Source: Fujitsu Performance measurements comparing Xeon L5650 vs. X5570 SKUs using SPECint_rate_base2006. See http://docs.ts.fujitsu.com/dl.aspx?id=0140b19d-56e3-4b24-a01e-26b8a80cfe53.

10 Source: Intel measurements as of February 2010. 5 month ROI claim estimated based on comparison between 2S Single Core Intel® Xeon® 3.80 with 2M L2 Cache and 2S Intel® Xeon® X5680 based servers. Calculation includes analysis based on performance, power, cooling, electricity rates, operating system annual license costs and estimated server costs. This assumes 8kW racks, $0.10 per kWh, cooling costs are 2x the server power consumption costs, operating system license cost of $900/year per server, per server cost of $7200 based on estimated list prices and estimated server utilization rates. All dollar figures are approximate. Performance and power comparisons are based on measured server side java benchmark results (Intel Corporation Feb 2010). Platform power was measured during the steady state window of the benchmark run and at idle. Performance gain compared to baseline was 15x.

Baseline platform: Intel server platform with two 64-bit Intel Xeon Processor 3.80 Ghz with 2M L2 Cache, 800 FSB, 8x1GB DDR2-400 memory, 1 hard drive, 1 power supply, Microsoft* Windows* Server 2003 Ent. SP1, Oracle* JRockit* build P27.4.0-windows-x86_64 run with 2 JVM instances.

New platform: Intel server platform with two Intel® Xeon® Processor X5680 (12M Cache, 3.33 GHz, 6.40 GT/s Intel® QPI), 24 GB memory (6x4GB DDR3-1333), 1 SATA 10krpm 150 GB hard drive, 1 800w power supply, Microsoft Windows Server 2008 64 bit SP2, Oracle* JRockit* build P28.0.0-29 run with 4 JVM instances.

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

Printed in USA 1212/DEJ/HBD/PDF ♻ Please Recycle 322706-003US