



Addressing Power and Thermal Challenges in the Datacenter

Solutions for Optimizing Density, Availability, and Cost

Power and cooling are becoming critical issues in today's datacenters. This paper:

- Offers practical recommendations for addressing these issues today.
- Introduces a new Intel® technology that can reduce average platform power and cooling demands by up to 25 percent.
- Discusses the accelerated move toward multi-core Intel processor architectures, along with other power-related innovations at the processor, rack, and datacenter levels.

Table of Contents

Executive Summary	2
Today's Challenge	3
Better Power Efficiency at the Microprocessor Level	4
Silicon Innovation	4
Architectural Advances	4
Dynamic Power Management	5
Rack and Datacenter Issues	6
De-Rating Strategies for Rack Optimization	6
Moving Toward Integrated Datacenter Management	8
Best Practice Recommendations	8
System-level Optimizations	8
Rack-level Optimizations	9
Datacenter-level Optimizations	9
The Bigger Picture	10
Conclusion	11
Resources	11
Intel® Solution Services	11
IT@Intel: Discover How Intel Does IT	11
Other Datacenter Management Resources	11
Sidebars	
<i>The Power of Parallelism</i>	4
<i>How Much Can You Save with DBS?</i>	6
<i>Support for DBS in Intel® Processors</i>	8
<i>Support for Datacenter Optimization</i>	10

Executive Summary

Rising utility rates and escalating compute requirements are creating new power and thermal challenges for datacenter managers. Today's high-density rack and blade servers bring these issues into especially sharp focus. Since these architectures are inherently more scalable, adaptable, and manageable than traditional platforms, they deliver much-needed relief in complex, crowded datacenters. Yet they also introduce power and thermal loads that are substantially higher than those of most currently deployed systems. In some cases, they may even push the cooling infrastructures of older facilities beyond their design limits.

For the most part, these are not insurmountable challenges. As will be described in this paper, the implementation of current datacenter best practices for power and cooling management can improve overall efficiency, and may enable much higher densities without disruptive change. Of course, power-related challenges will increase in the future, as total datacenter compute power continues to climb, and IT organizations take advantage of more powerful and higher density architectures. The answer is not to reverse this trend, which offers incredible benefits in performance and total cost of ownership (TCO); but rather to develop new and better ways to contain and manage power and thermal requirements.

A long-term solution will require broad industry innovation and collaboration. At the component level, Intel is working toward a number of fundamental microprocessor enhancements to address these issues. The most important is a shift to multi-core architectures that incorporate two or more processing units on a single chip. This will deliver major performance improvements while helping to contain power usage.

For more immediate relief, the latest Intel® Xeon™ processor offers a technology called Demand Based Switching (DBS), which reduces average system power consumption and cooling costs by as much as 25 percent, with minimal performance penalty.¹ This can potentially save millions of dollars in utility bills for large companies with multiple datacenters. It may also enable savings in cooling infrastructure costs, by reducing total requirements for new equipment.

This white paper discusses these and other developments in power and thermal control and management. It also offers practical recommendations that can be implemented now to optimize power and cooling strategies in the datacenter, and to accommodate higher density server solutions without major infrastructure overhauls.

¹Based on preliminary Intel power testing. Actual power savings will vary based on system configurations and workloads.

Today's Challenge

Several critical factors have converged to create today's power and thermal challenges. The first is the dramatic increase in the number of servers deployed in most facilities. Worldwide, total server numbers have increased approximately 150 times in less than a decade. During that same period, average compute power per server has increased by a factor of 10, and smaller form factors have packed that power into much smaller volumes.

The magnitude of these changes was not anticipated when most of today's active datacenters, infrastructure equipment and operational practices were designed. Most existing facilities are designed to accommodate only 40-75 watts per square-foot, about a quarter of what would be ideal for today's needs. They also employ cooling technologies and strategies that have not changed appreciably in over 30 years.² In many cases, IT and facilities managers are not even sure if their facilities will operate as expected, since they are only now being tested at their design limits. For all these reasons, power and

cooling can be a limiting factor in deploying the most effective computing solutions in a given environment.

Power and cooling are also becoming more significant contributors to TCO in today's datacenters. This is made more apparent by recent budget constraints and rising utility rates. Though these conditions may ease, the general trend is likely to continue. Businesses are deploying compute power at ever-increasing rates, so total datacenter power consumption will undoubtedly continue to rise, despite dramatic improvements in performance-per-watt in current and future processors and servers. This means that power and thermal management will most likely remain important issues for IT and facilities managers into the foreseeable future.

A solution to this challenge will require industry innovation on an ongoing basis, and at all levels—from silicon, to racks, to datacenters—to pave the way for continuing, cost-effective IT growth (Figure 1). It will also require a move toward more rigorous best practices for datacenter power and cooling management.

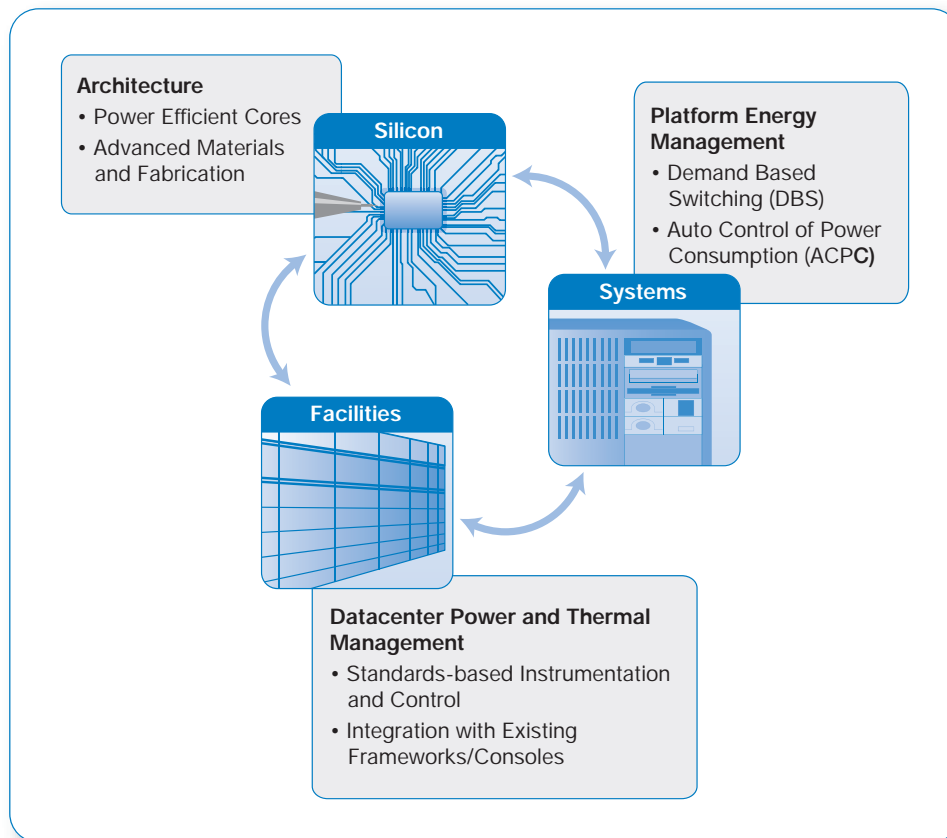


Figure 1. Intel is working on multiple fronts to address today's power and thermal challenges, from silicon innovations that will dramatically increase component efficiency, to system- and datacenter-level initiatives that involve broad industry collaboration.

²For more information, see *Essential Cooling System Requirements for Next Generation Datacenters, White Paper #5*, by American Power Conversion. Available at www.apc.com.

The Power of Parallelism

Parallelism is the ability to process multiple software instructions, threads or jobs simultaneously. It can dramatically improve overall performance, without the power consumption issues associated with very high processor frequencies. Intel® architecture supports multiple options for enhancing parallelism based on specific workload requirements. At the highest level, these options include large symmetric multiprocessing systems (SMPs) with up to 128 processors. These platforms deliver exceptional parallelism per platform for multiple, simultaneous software stacks, or for complex, threaded applications. Alternatively, workloads that can be load-balanced across multiple systems can be run on clustered architectures, enabling massive parallelism using affordable platforms. As a real-world example, Google uses this strategy to cost-effectively deliver exceptional response times for a very high volume of search engine queries.^a

Other Intel technologies can be used to complement both clustering and large, SMP architectures. Depending on workload characteristics, one or more of the options below can drive substantially higher parallelism, with corresponding performance and power benefits.

- **Intel® Hyper-Threading Technology (HT)**—With HT, a single processor (or processor core) can simultaneously process two software threads. This can improve performance by up to 30%.^b
- **Multi-Core Processor Architectures**—Discussed in this paper.
- **Explicitly Parallel Instruction Computing (EPIC)**—EPIC was designed from the ground up to maximize parallelism for high-end, 64-bit applications, not only at the hardware level, but also through compilers that optimize software code for parallel processing. With the introduction this year of the Intel Itanium® 2 processor with 9M cache, this architecture will have enabled Intel to maintain the same thermal envelope across three generations of processors, while more than doubling performance. The first dual core Itanium processor (available in 2005) is expected to increase performance by about 50%, while consuming less power than its single core predecessor.

^a Google also reports a more than 30% performance boost when taking advantage of Intel Hyper-Threading technology (HT), and expects substantial benefits from multi-core processors. For more information, including a discussion of power and thermal issues, see Web Search for a Planet: The Google Cluster Architecture, available at: http://www.acroterion.ca/Knowledge_Base/Web_Search_For_A_Planet.html

^b Hyper-Threading Technology requires a computer system with an Intel Pentium® 4 or Xeon™ processor supporting HT Technology and an HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See <http://www.intel.com/info/hyperthreading/> for more information including details on which processors support HT Technology.

Better Power Efficiency at the Microprocessor Level

Intel is working on several fronts to reduce microprocessor power consumption, from fundamental transistor fabrication to overall chip architecture. A few of the most important developments are described below.

Silicon Innovation

Current and ongoing Intel advances in materials, basic transistor fabrication, and circuit design will all contribute to improved performance and power efficiency in future microprocessors. As one example, a recent technological breakthrough achieved record-setting arithmetic logic unit (ALU) performance. According to Steve Pawlowski, an Intel Fellow in the Corporate Technology Group and the director of Intel's Microprocessor Technology Lab: "These technologies enable a huge performance improvement while cutting the energy requirement in half. They reduce active power leakage four times over the best conventional arithmetic circuits available." Since ALUs are both power-hungry and performance-critical, this is an important advance. Intel is currently integrating this new ALU design into multiple, existing processor designs.

A more fundamental and far-reaching breakthrough is the successful development by Intel researchers of record-setting, high-performance transistors using a new high-K material.³ When used in conjunction with specialized metals, also developed by Intel, this material decreases gate leakage current by a factor of up to 100, drastically reducing one of the industry's most significant contributors to unproductive power consumption. High-k/metal gate materials are on track for inclusion in Intel's 45nm process which is scheduled for high volume production in 2007.

Another Intel innovation is "strained silicon," which is already in production on Intel's 90nm process generation. This technology can be used either to improve transistor performance or to decrease power consumption. With enhanced ALUs, high-K transistors, strained silicon and a variety of other silicon and circuit innovations, Intel engineers are confident they can contain power consumption through the next decade, while continuing to provide dramatic increases in processor performance.

Architectural Advances

"Essentially, the design paradigm has shifted at Intel, and all the resources we have are dedicated to multi-core [processors]."

– Paul Otellini, President and Chief Operating Officer, Intel Corporation⁴

³The high-K material is used for the transistor "gate" dielectric, and corresponding new metal materials have been developed for the gate, itself. For more information, visit the Intel Web site at: <http://www.intel.com/pressroom/archive/releases/20031105tech.htm>

⁴Intel shoots for dual cores, wireless profits, John G. Spooner, CNET News.com, May 13, 2004.

Intel is now focusing virtually all of its processor development resources on multi-core architectures that will help to enable dramatic performance increases within existing or better power and thermal envelopes⁵. Multiple cores improve parallelism, which is the ability to process multiple instructions and multiple software threads per clock cycle. They also help to avoid performance stalls, since the processor can continue working, even if one core is stuck waiting for data or instructions.

One analyst has likened this approach to designing automobiles with multiple cylinders.⁶ Instead of a single, very large cylinder running at extremely high speeds (and correspondingly high temperatures), each individual cylinder can be designed for optimal power efficiency. Combined, they deliver far better performance than the single cylinder, with substantially lower total power consumption.

In addition to power efficiency, multi-core architectures will provide greater flexibility for shaping processor functionality to match the diverse needs of enterprise applications. As the move toward multi-core solutions gains momentum, Intel will take advantage of its leadership in silicon design and manufacturing to integrate more and better performing cores in each processor, which will continue to drive both price/performance and performance/watt.

Of course, these advances will take place within the context of other platform enhancements. Examples include enhanced processor support for virtualization and Intel® Extended Memory 64 Technology (Intel® EM64T)⁷, which enables a processor to run both 32-bit and 64-bit operating systems and applications. Though such advances do not focus explicitly on power or thermal issues, they enable more efficient use of computing resources in diverse environments, which increases the business value of all datacenter infrastructure, including power and cooling systems.

Dynamic Power Management

The latest Intel Xeon™ processor supports Demand Based Switching (DBS), a capability that can automatically reduce average system power consumption by up to 25 percent in a typical datacenter environment.⁸ The principle is relatively simple, and is based on Enhanced Intel® SpeedStep® Technology, which Intel has used with great success to enable improved battery life in laptops and handheld devices.

Traditionally, a microprocessor operates only at a single frequency and voltage, regardless of its workload. It is therefore always “on” and always consuming full power. Processors with Enhanced Intel® SpeedStep Technology are designed to run at multiple frequency/voltage settings (Figure 2).

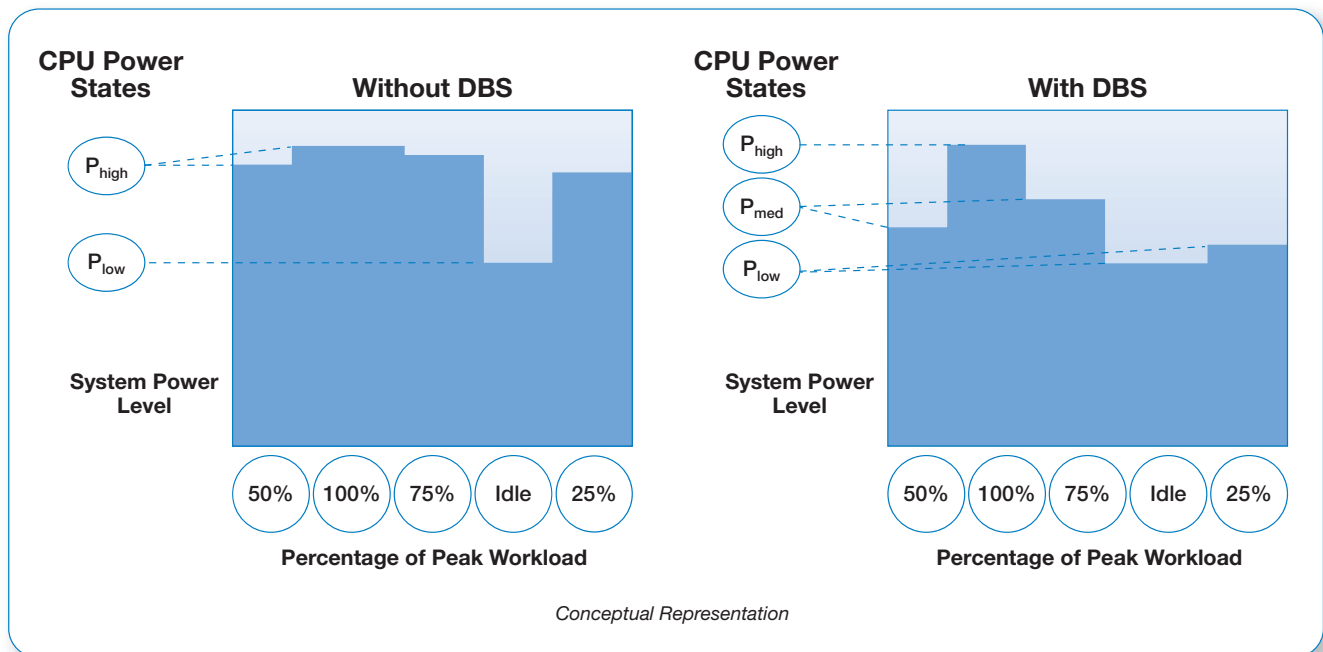


Figure 2. Without DBS, a processor is always set in its highest power state. With DBS, a processor has multiple power states to choose from, so it can consume less power when processing less intensive workloads. This can reduce average system power consumption by up to 25%.⁸

⁵ For more information, see *Growth at the Platform Level*, by Paul Otellini, President and Chief Operating Officer, Intel Corporation, available at: http://www.intel.com/intel/finance/presentations/PDF_Files/sam_2004_PSO.pdf

⁶ Nathan Brookwood, as reported in *Intel decides two cores are better than one*, by Tom Krazit, IDG News Service, Boston Bureau, May 10, 2004, available at: <http://www.itworld.com/Comp/1982/040510inteltwocores/>

⁷ Intel® Extended Memory 64 Technology (Intel EM64T) requires a computer system with a processor, chipset, BIOS, OS, device drivers and applications enabled for Intel EM64T. Processor will not operate (including 32-bit operation) without an Intel EM64T-enabled BIOS. Performance will vary depending on your hardware and software configurations. Intel EM64T-enabled OS, BIOS, device drivers and applications may not be available. Check with your vendor for more information.

⁸ Based on preliminary Intel power testing. Actual power savings will vary based on system configurations and workloads.

In conjunction with an operating system (OS) and BIOS that support DBS, the processor will automatically operate at the lowest setting that is consistent with optimal application performance.⁹ The OS monitors processor utilization multiple times per second and downshifts to a lower frequency and voltage as appropriate. Power usage is therefore automatically tailored to match server workloads, which substantially reduces waste with minimal impact on peak performance capabilities.¹⁰

Actual power savings from DBS will depend on several factors, including application workloads, the number and definition of power states, and the policies that determine when and how these power states are changed. DBS typically delivers minor power savings during peak workloads, with substantial savings during intermediate loads and halt states, which account for the majority of server time in most environments. Power savings of 25 percent, as cited above, is a reasonable expectation based on typical datacenter conditions. As processor frequencies increase in future releases, operating systems that support DBS will automatically take advantage of additional power states as they become available, to further optimize power efficiency.

It must be emphasized that DBS, by itself, does not significantly decrease *peak* power consumption, so it should not be used in de-rating platforms to determine maximum rack densities. However, it *can* significantly reduce utility costs (see the sidebar, “How much can you save with DBS?”).

Enabling a processor to perform reliably in multiple power states is not a trivial enhancement. It introduces major design challenges that must be addressed in every aspect of processor functioning, from the switching characteristics of individual transistors to signal timing across the entire device. Intel validation processes include extensive testing to ensure reliable and optimized performance in each power state.

Rack and Datacenter Issues

As utility costs increase and higher density platforms become available, it becomes more and more worthwhile to invest in technologies and operational strategies to improve power and thermal control at the rack and datacenter levels. This section discusses current strategies and future developments at these levels.

De-Rating Strategies for Rack Optimization

For many years, most IT organizations based power requirements for a rack on the “nameplate” specifications of installed systems. This is a very conservative approach, since nameplate values are based on fully configured systems, which are

How Much Can You Save with DBS?

Less power consumption means lower utility bills for power and cooling, but by how much? Many factors influence potential savings, including the number of servers that support DBS, as well as their configurations and workloads. Systems operating 24x7 near peak utilization will experience little benefit, but such systems are rare. Typical servers with lower or more sporadic utilization will see much greater savings.

In a collaborative study with a global financial institution, Intel engineers estimated direct savings of \$250K to \$300K annually for 876 servers in a single datacenter.^a These benefits were based on a detailed analysis of existing workloads, and would require DBS-capable processors and operating systems on all servers.

^a Results are based on unpublished data. The study was performed in 2004.

relatively rare in most environments. Recently, it has become common to “de-rate” systems, which simply means to reduce the nameplate power rating by some percentage based on estimates of actual power and cooling needs.

De-rating may deliver cost benefits via increased efficiency and higher rack densities. However, depending on the strategy, de-rating can increase the risk of downtime due to overheating or a tripped breaker.

In general, there are five basic options (Table 1 on page 7):

- 1. No De-Rating**—This is the simplest and most conservative option, though it may entail substantial waste of power, cooling and rack space.
- 2. De-Rating Based on Actual System Configurations**—With this approach, maximum power consumption per rack is based on actual system configurations (P_{config}). For typically configured systems, this strategy generally enables about 50% more systems per rack. For minimally configured systems, it may allow up to 3 times as many systems per rack. Some server vendors currently offer software utilities that will automatically calculate P_{config} .

Since system power consumption should not exceed P_{config} , there is little additional risk in this strategy, if it is managed appropriately. However, IT staff must be aware of the implications, and re-evaluate P_{config} and associated rack densities when systems are replaced or upgraded.

⁹ OS and BIOS support for DBS is expected sometime in 2H'04 for workstation and server platforms in both Windows and Linux environments. Check with your vendor for details.

¹⁰ Even at peak software workloads, a platform will often experience reduced CPU utilization and halt states at millisecond time scales, and will therefore downshift as appropriate to lower power states. This may minimally impact peak performance capabilities. It will also reduce peak power consumption by comparable or greater percentages.

Table 1. De-Rating Options for Optimizing Rack Density

Option	Definition	Admin Costs	Pros/Cons
No De-Rating	Load rack based on the full nameplate power of each system.	Lowest	Minimizes risk, but rack density may be much lower than necessary.
P_{config}	Calculate the maximum power load based on the actual configurations of the systems in the rack.	High	Risk remains low, and rack density may be substantially improved (as much as triple the density for a rack with minimally configured systems).
PMSI (power supply monitoring)	Track actual power consumption using a console manager or electronic nameplate.	Med	With sufficient data collection, it may be possible to fully populate racks, while managing risk per business requirements.
PMSI plus P_{config}	Track actual consumption versus P _{config} and delete systems as necessary.	Med	Simplifies density optimization and provides continuous monitoring of actual consumption.
ACPC	Same as PMSI with P _{config} , with additional assurance that system power will not exceed the specified setting.	Med (low when automated)	Enables high rack density with relatively low risk of downtime; some (controllable) risk of performance throttling.

3. De-Rating Based on Power Supply Monitoring (PMSI)—

By using power supplies with monitoring capabilities, IT can track actual power consumption over time, and establish rack densities based on these measurements. Since real-world variables, such as application ratios and CPU utilization, lead to power consumption levels lower than P_{config}, this can lead to significant additional density gains. Of course, these variables are subject to change, so care must be taken to monitor trends and to provide appropriate power and thermal headroom to balance risks and benefits.

4. De-Rating Based on P_{config} and PMSI—

Using both strategies together, IT staff can track actual power consumption versus P_{config}, and delete systems as necessary to balance density and risk more effectively.

5. Automatic Control of Power Consumption (ACPC)—

This is an option Intel is developing and will make available to server vendors. ACPC enables IT administrators to set limits on system power consumption by monitoring power supplies and controlling microprocessor power states. If a system attempts to draw power beyond specified limits,

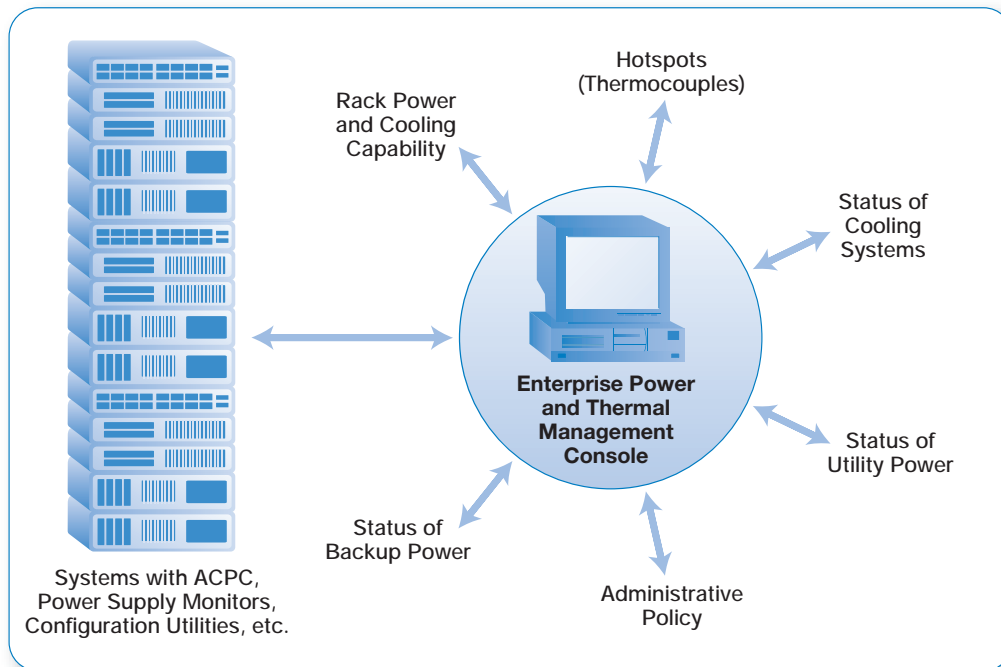


Figure 3. Intel is working to develop a standards-based framework that will support integrated and automated power and thermal management throughout the datacenter, enabling improved efficiency, higher rack densities, and reduced TCO.

the frequencies of microprocessors that support multiple power states can be automatically down-shifted to reduce power consumption without shutting down systems.

In essence, ACPC will allow IT staff to allocate appropriate amounts of power and cooling to systems in order to optimize rack density, while managing the risk of downtime due to electrical or cooling issues. ACPC can also enable facility managers to reduce platform power consumption during a power or thermal crisis without shutting down the system. ACPC will be designed to integrate readily with standard management consoles and frameworks.

Moving Toward Integrated Datacenter Management

With the strategies outlined in the previous section, there is clearly a tradeoff between power and cooling efficiency and administrative overhead. Automated P_{config} utilities from system vendors and automated power supply monitoring can reduce those costs. However, to reap full benefit from these strategies, monitoring and control capabilities must be fully automated to reduce the human costs of implementation.

Intel is currently working on a framework that will make this possible. The goal is to enable IT and facilities managers to automate and optimize power and cooling management across the datacenter, and to monitor and control all relevant variables at the component, system, rack, and datacenter level (Figure 3). The framework will be designed to work seamlessly with standards-based management tools and infrastructure solutions. Of course, this will require broad collaboration among component, system, and infrastructure vendors. Ultimately, it should provide the ability to balance workloads, system utilization, and power and cooling delivery to optimize performance versus costs in a very comprehensive and granular way.

Best Practice Recommendations

The following recommendations come from Intel® Solution Services, a professional services organization within Intel, which consults with clients worldwide on enterprise datacenter management. In applying these recommendations, please be aware that specific cooling challenges can be complex, and many organizations will benefit from consulting with outside specialists.

System-level Optimizations

1. Take Advantage of DBS for New Server Deployments— DBS is currently supported on the latest Intel Xeon processor with Intel EM64T and will be increasingly integrated into next-generation Intel server processors. OS and BIOS support for DBS is expected sometime in 2H'04 for workstation and server platforms in both Windows and Linux environments. Check with your vendor for details.

Support for DBS in Intel® Processors

DBS is one of many new features introduced in the latest Intel® Xeon™ processor. Other features in the processor and supported platforms include:

- Higher frequency for faster performance
- Intel® Extended Memory 64 Technology (Intel® EM64T), which extends memory addressability and enables the processor to run both 32-bit and 64-bit operating systems and applications
- PCI Express for faster and more scalable I/O
- An 800MHz system bus for faster on-board communications
- DDR2 memory for higher capacity and performance with reduced power consumption
- C1E (enhanced halt state) to reduce power consumption in idle systems

Current plans call for DBS to be supported in the Intel® Xeon™ processor MP and the Intel Itanium® processor families in 2005. In selecting processors, be aware that the performance advantages of the Intel Xeon processor MP in multiprocessor configurations, and the even greater advantages of the Intel Itanium 2 processor, may enable fewer systems to handle the same workload, thereby reducing total power consumption. Intel also offers a low power (62W) DP Itanium processor. Support for DBS should therefore be considered as one of many factors in addressing performance, cost and power requirements.

- 2. Check With Your Server Vendor About Other Power Management Capabilities—**Power configuration (P_{config}) utilities, ACPC and other tools may significantly extend the benefits of DBS, and also help IT staff increase rack densities without undue risk of downtime due to electrical or thermal issues.
- 3. Move to Newer Operating Systems for All Platforms—**Microsoft Windows Server 2003 makes use of "halt states" to reduce power consumption for idle servers. This is beneficial even on older platforms. However, DBS-enabled systems have lower-power halt states that add substantially to the total power savings.

Rack-level Optimizations

Most system, rack-, and room-level cooling issues are created due to insufficient airflow or inadvertent mixing of hot and cold air. The need for sufficient airflow is obvious, but is often overlooked by IT personnel who are focused on other concerns. The mixing of hot and cold air is a more subtle issue, but equally problematic, since it can dramatically reduce the efficiency of a cooling system, and may also impact airflow.

1. Understand Airflow Requirements for Specific

Equipment—There are four basic airflow scenarios: front-to-back, side-to-side, bottom-to-top, and top-to-bottom. Understanding the requirements for specific equipment will enable an efficient rack-level design and cooling strategy.

2. Standardize on Racks Designed for High-Density

Environments—Standardizing on an appropriate rack design makes it much easier to establish and enforce effective power and thermal policies. Avoid shallow racks to ensure in-rack cabling does not obstruct airflow. Consider racks that support retrofit fan or cooling units (but verify the benefits of these add-on units). The ability to supplement cooling for individual racks may be important to accommodate high-density systems without compromising room-wide efficiency.

3. Arrange Racks in Rows to Establish Hot and Cold

Aisles—Racks should be aligned front-to-front along cold aisles, and back-to-back along hot aisles. Within each row, racks should be tightly abutted. For this strategy to be effective, cold air must be delivered to cold aisles and hot air extracted from hot aisles. Work to eliminate hot air remixing, which will cause short cycling of the cooling system.

4. Use Blanking Panels

—Blanking panels improve airflow through the rack, minimize air loss and help to prevent exhaust air recirculation.

5. Ensure Adequate Airflow to Individual Racks and

Systems—Clearly define power and cooling requirements at the room, row, and cabinet level. Ensure sufficient airflow to racks based on system-level inlet air temperature and airflow requirements, and use thermal and aerodynamic analysis tools to model and design your cooling solutions. Insufficient airflow will often result in hotter systems and turbulence that decrease cooling efficiency. For example, if a rack requires more cold air than the room provides, its fans will pull in a mix of hot and cold air. This will result in reheating of the room, hotter systems, unhealthy airflows, and a substantial reduction in the efficiency of the cooling system.

6. Explore the Benefits of Blade Servers

—Blade architectures may reduce total power consumption (per unit of compute power) and deliver substantial TCO benefits through reduced cabling, easier provisioning and improved

modularity. However, they may increase power and cooling *density*. It is therefore important to look at total costs, risks and benefits within your particular physical and operational environment.¹¹

Datacenter-level Optimizations

1. Understand Datacenter Airflow—The locations of cooling systems and ductwork are obviously critical, but so are the locations of racks, cable trays, firewalls and other infrastructure elements. Blank off any floor opening that allows access air to escape the plenum. Software tools are now available that greatly simplify airflow and thermal analysis. Consider consulting with facilities cooling specialists for complex implementations.

2. Optimize Room Temperature Settings—Consider increasing the Delta T of your cooling system to more closely match IT equipment specifications. This may allow you to reduce total airflow, while meeting the same cooling capacity and reducing operational costs. (As an example, Intel IT has found it beneficial to lower supply air temperatures to between 55 and 65 degrees Fahrenheit while increasing Delta T values to 26 degrees Fahrenheit.)

3. Pay Attention to Infrastructure Efficiency—It is generally worthwhile to spend more for infrastructure components that run efficiently at anticipated loads. Power loss in uninterruptible power supplies, power distribution units, cooling systems, etc., just add to the thermal load.

4. Perform Regular Power and Thermal Audits—New systems, upgrades, and room changes can have unintended consequences, so it is important to monitor airflow, temperature, and other environmental factors on a regular basis.

5. Avoid Over-Design—Right-sizing power and cooling infrastructure is one of the most effective ways to reduce capital and operational costs in the datacenter. Work to understand lifecycle requirements and size infrastructure accordingly. Track vendor innovations, and, whenever possible, move toward more modular, flexible, and standardized solutions that improve agility and scalability.

6. Establish Policies and Educate Personnel—Best practices for power and thermal management must become an integral component of datacenter operations. Everything from temperature and humidity settings to new system and cable deployments should follow well-understood guidelines that optimize cooling efficiency and minimize airflow obstructions and hot/cold air mixing.

7. For New Datacenters, Establish a Master Plan Based on Usage—Different usage models require different layouts and capacities to enable optimized cooling solutions (Table 2 on page 10).

¹¹ A recent report by Gartner (*Electrical Requirements for Blade Servers*, April 24, 2004) compares power consumption between blade and traditional rack architectures. It cites approximately 50% power savings for single-processor blades, 30% for dual-processor blades, and minimal if any savings for quad-processor blades. For a more comprehensive look at the issues, see the IT@Intel white paper, *Blade Server Benefits—Evaluating new server technology in an IT environment*, available on the Intel Web site, at: http://www.intel.com/business/bss/infrastructure/enterprise/blade_server.pdf. This paper documents a proof-of-concept study by the Intel IT organization, which was performed to assess the potential benefits of blade servers in Intel's own IT environment.

Table 2. Understanding Power Requirements for New Datacenter Designs

	Datacenter Focus	
	Technical Design	Business Transactions
Typical Configuration	60% High Density 30% Low Density 10% Network Devices	30% High Density 50% Low Density 20% Network Devices
Typical Rack Power	High Density (recent past): 3kW-10kW High Density (current): 10kW-14kW Blade Racks (future): 15kW-20kW	Rack load steady at 3kW-10kW (High density systems can be spread out to dissipate heat)

Support for Datacenter Optimization

Intel® Solution Services

As businesses grow their IT infrastructure, they must occasionally make critical decisions about buying, building, optimizing, or consolidating datacenter facilities—decisions that can impact business and IT options for many years. Effective solutions must take into account not only changing business and IT requirements, but also technology trends that may substantially alter needs and capabilities over the life of the facility or equipment.

Intel Solution Services, the worldwide professional services organization within Intel, offers expert assistance based on formal, best known methods for optimizing the use of current facilities, projecting future requirements and developing long-term strategic plans that balance business opportunity against TCO. Intel Solution Services has worked successfully with many clients to develop detailed business, technical, and financial analyses that drive better and faster decisions on crucial datacenter issues.

For more information, visit:

www.intel.com/go/intelsolutionservices

The Bigger Picture

Power and thermal management is just one of the many challenges facing IT and facilities managers today. In a typical datacenter, 70-80% of IT costs are directed at maintaining existing infrastructure and applications.¹² More efficient solutions are imperative to reduce these costs and free up resources for new deployments and upgrades that contribute more directly to business value (Figure 4).¹³

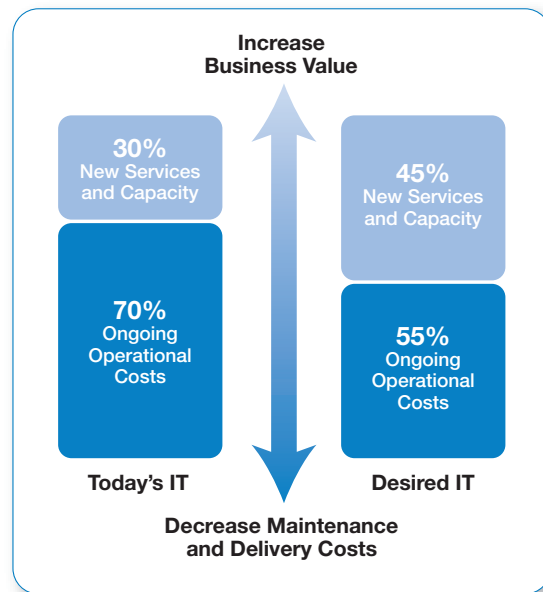


Figure 4. Most businesses devote 70-80% of their IT budget to maintenance. It is imperative to reduce these costs and free up resources for new deployments and upgrades that contribute more directly to business value.

¹² Based on a quote by Kevin Rollins, President and COO, Dell Corporation, as reported in *Dell and Sun Offer Different Visions*, InformationWeek.com, by Larry Greenemeier, September 17, 2003.

¹³ For recommendations on reducing TCO for datacenter management, see the Intel white paper, *Optimizing Enterprise IT—Essential Strategies for Reducing Operating Costs and Ensuring Business Value*, available on the Intel Web site at: <http://www.intel.com/business/bss/infrastructure/managing/optimizing.pdf>.

The industry is driving toward a solution to this challenge, and most vendors share a common vision for tomorrow's datacenter:

- **Self-managing Infrastructure Components**—Vendors are building increasing levels of manageability and automation into platforms, operating systems, and applications. The goal is to provide solutions that can configure, monitor, optimize, heal, and protect themselves.
- **Modular Datacenters that Are More Agile and Scalable**—Ideally, tomorrow's datacenter will be managed as a single system from a single console. Virtual compute, storage, and network resources will be dynamically allocated based on business and application requirements.

Of course, management complexity will not simply disappear, but must be modeled and managed in silicon and software. These developments will therefore depend on increasingly powerful server platforms—and increasing compute density. Though it can be assumed that performance per watt will continue to rise rapidly, it can also be assumed that total datacenter power consumption will also increase. It will therefore be essential to manage power and cooling resources within the context of total costs and performance.

Conclusion

Rising utility rates, aging facilities, and high-density server solutions are introducing new power and thermal challenges for IT and facilities managers. The first step in addressing these challenges is to employ datacenter best practices, as described in this paper. In many cases, this will enable substantially better rack utilization, while improving overall power and thermal efficiency.

The second step is to reduce processor-based power and cooling requirements in new server deployments via DBS. DBS can reduce average CPU power consumption by up to 25%, with minimal performance penalty.¹⁴ It is supported today in the latest Intel Xeon processor, which also includes Intel EM64T and other significant enhancements.

Finally, stay informed about evolving options for managing and optimizing datacenter power and cooling efficiency. Intel is currently working to develop a standards-based framework that it will make available to the industry. The goal is to put more manageable and power-efficient building blocks in the hands of IT and facilities managers, along with the management tools they need to optimize price/performance for their entire facility.

Meanwhile, Intel will continue to focus on advances in silicon technology and microprocessor architecture to drive higher performance with better power efficiency and lower total cost. Next-generation multi-core processors, in particular, will help to eliminate many current power and thermal barriers in processor and platform design. Intel will also continue to work with other leaders in the industry to develop power and thermal solutions that improve efficiency at every level of the datacenter.

Resources

Intel Solution Services

Intel Corporation's worldwide professional services organization helps enterprise companies capitalize on the full value of Intel architecture through consulting focused on architecture transitions. Backed by the world's largest silicon manufacturing company and one of the largest e-Business corporations in the world, Intel Solution Services uses its expertise in Intel architecture and next-generation technologies, as well as its relationships with key industry alliances, to design cost-effective, leading-edge solutions that help deliver superior business results.

For more information about Intel Solution Services, visit: www.intel.com/products/services/intelsolutionservices

IT@Intel: Discover How Intel Does IT

Intel's IT organization is a critical component in every aspect of Intel's business. It runs a global infrastructure that supports Intel's diverse workforce in more than 250 locations around the world. It also provides products and services to Intel's external customers, including help desk services to 80,000 individual IT users and information solutions to Intel Business Groups. The IT@Intel program offers a variety of white papers and other resources providing useful insights into critical IT issues.

For more information about Intel IT, visit: <http://www.intel.com/it/>

Other Datacenter Management Resources

- APC: <http://www.apc.com>
- AFCOM: <http://www.afcom.com>
- 7x24 Exchange: <http://www.7x24exchange.org/>
- Microsoft:
<http://www.microsoft.com/management/default.mspx>
<http://www.microsoft.com/technet/itsolutions/techguide/msm/default.mspx>
- Unisys: http://www.unisys.com/products/data_center/
- IBM Intel-based solutions: <http://www.pc.ibm.com/us/eserver/xseries/>
- HP Mission Critical Support: <http://www.hp.com/hps/mission/>
- Gartner: <http://www.gartner.com/>
- Forrester: <http://www.forrester.com>
- META Group: <http://www.metagroup.com>

¹⁴ Based on preliminary Intel power testing. Actual power savings will vary based on system configurations and workloads.



Copyright ©2004 Intel Corporation. All rights reserved.

Intel and the Intel logo, Intel Itanium, Intel Xeon and Intel SpeedStep 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. Information regarding third-party products is provided solely for educational purposes. Intel is not responsible for the performance or support of third-party products and does not make any representations or warranties whatsoever regarding quality, reliability, functionality, or compatibility of these devices or products.

302739-001
