

Intel IT Data Center Solutions: Strategies to Improve Efficiency

Our long-term data center planning process will help us realize an estimated USD 1 billion in savings while enabling the agility to respond faster to business needs.

Executive Overview

Over time, Intel IT has evolved our strategies to optimize our data center infrastructure to respond faster to business needs while enhancing the services and value IT brings to the business. Our new data center strategies shift the emphasis away from reducing the number of physical data center facilities to, instead, focusing on approaches that leverage the full potential of our data centers worldwide. This helps increase business value across the entire data center infrastructure. We expect our methods to achieve a nominal cost savings of about USD 1 billion by 2014.¹

Our approach centers on three strategies: optimization, utilization, and strategic investments. The key elements of these strategies include:

- Accelerating server refresh to take advantage of performance and power efficiency improvements.
- Consolidating and virtualizing our server resources.
- Upgrading facilities to improve facilities capability, utilization, and energy efficiency.
- Adding capacity in a modular, scalable way.
- Locating new data centers in aggregation sites whenever possible.
- Eliminating data centers whenever feasible and financially viable.

▪ Optimizing WAN configurations to substitute for localized data center capacity and to support remote data center services.

▪ Switching to a parallel storage solution in the high-performance computing (HPC) area.

Our strategies are already delivering results. For example, from 2006 to 2008, the increased use of optimization tools, such as virtualization and server refresh, reduced data center facilities capital investment by 65 percent.

Our long-term planning process will help us continue to drive broader efficiencies across the data center environment as well as respond faster to business needs—which in turn enhances the services and value that IT brings to the business.

¹ Net present value of USD 550 to 650 million.

Matt Beckert
Strategic Financial Analyst, Intel IT

Bradley A. Ellison
Senior Data Center Architect, Intel IT

Shesha Krishnapura
Senior Principal Engineer, Intel IT

Sanjay Rungta
Principal Engineer, Intel IT

Arlene Sanchez
Project Manager, Intel IT

Ananth Sankaranarayanan
Technical Program Manager, Intel IT

CONTENTS

Executive Overview	1
Background	2
Our Initiative	2
Data Center Strategy	3
Evolving our Strategy	3
Process and Communication	4
Planning for the Long Term	4
Facilities Design Strategy	5
Compute Servers	6
Optimizing Performance and Total Cost of Ownership with Accelerated Server Refresh	6
Virtualizing Server Resources	7
Data Center Virtualization	8
Compute Strategy Results	8
Network Environment	9
WAN Optimization: Managing Bandwidth	9
Storage	10
Optimizing Storage	10
Parallel Storage Milestone	10
Future Alignment	11
Next Steps	11
Cloud Computing	11
IT Sustainability	11
Solid-State Drives	11
An On-going Initiative	12
For More Information	12
Acronyms	12

IT@INTEL

IT@Intel is a resource that enables IT professionals, managers, and executives to engage with peers in the Intel IT organization—and with thousands of other industry IT leaders—so you can gain insights into the tools, methods, strategies, and best practices that are proving most successful in addressing today's tough IT challenges. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

Data centers are at the heart of Intel's massive worldwide computing environment. With almost 443,000 square feet distributed over 97 data centers, these facilities house approximately 100,000 servers and carry an average of 148 million e-mail messages and 2,183 terabytes (TB) of WAN traffic each month.²

Business needs are growing rapidly, increasing the demands on our data center resources. We serve four main computing application types (referred to as DOME)—Design (sometimes referred to as engineering computing), Office, Manufacturing, and Enterprise. Each has its own set of characteristics and requirements:

- **Design computing.** Supporting Intel's chip design community, this group has most of the servers—about 70 percent.
- **Office computing.** This general-purpose group, which has about 10 percent of the servers, comprises typical IT and customer services.
- **Manufacturing computing.** With about 10 percent of the servers, this group supports fabrication, assembly, and testing manufacturing.
- **Enterprise computing.** Enterprise computing supports e-Business and enterprise resource planning, and comprises about 10 percent of the servers.

Of the 97 data centers currently in use, most of the office and enterprise computing is housed in three primary data centers.

Responding to the specific needs of each of these groups requires us to continually explore

innovative ways to utilize the full potential of our data centers, while effectively dealing with changing business conditions.

Our approach to managing our data centers reflects this ongoing process.

Our Initiative

Our data center initiative, which began in 2000, is a process of continuous and ongoing improvement (see Figure 1).

- **1995 to 2000.** During this time, we had no overall picture of compute demands, the number of data centers, rate of growth, or the costs to the company. We also did not have a holistic strategy in place or any defined data center standards. Data centers of various designs were added as components of a variety of business acquisitions. Each business unit had its own servers and built its own data centers; the total number of data centers increased as did costs.
- **2000 to 2004.** We created a group whose responsibility was to plan and manage data centers. Part of that endeavor was to complete an inventory to assess the number and makeup of data centers throughout the company. Similar to the undertakings by other large companies at the time, Intel began to focus on establishing a process to guide data center investments.
- **2005 to 2007.** We launched a data center efficiency program with an eight-year goal in mind to reduce the number of data centers and thus achieve greater efficiency by eliminating facilities overhead. The program achieved significant results, consolidating roughly 150 data centers to fewer than 100. However, we learned that data center consolidation wasn't appropriate for all data centers. Each manufacturing facility, for example,

² Intel IT data from May 2009. These numbers are dynamic and will change over time.

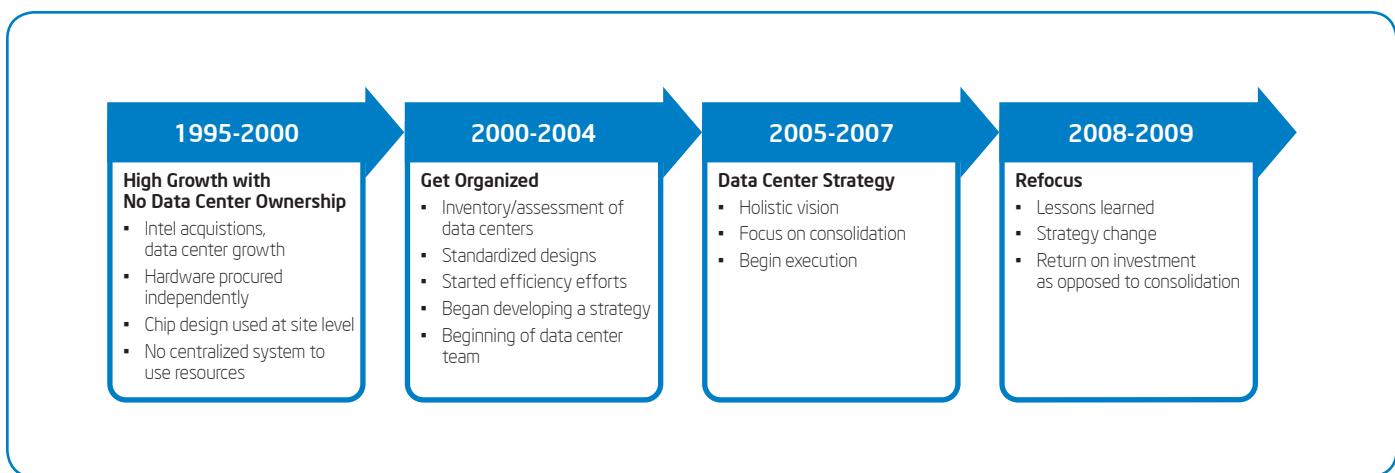


Figure 1. Intel's data center initiative is a continuous and ongoing improvement process.

requires a local data center as a fail-safe for the production process, reducing the likelihood of an interruption in the 24x7 manufacturing process. Within the engineering computing environment, our interactive workload required a data center nearby to eliminate the latency induced by the distance from the design data.

▪ **2008 to 2009.** Although our data center efficiency program helped to jump-start the process, we realized that instead of an eight-year plan, we needed a continuous and ongoing improvement process that included optimizing all data center facilities and IT infrastructure, not just the ones appropriate for consolidation. This strategy had to not only change the mindset within Intel IT, but also establish an internal, clearly understood set of guidelines that could be shared with our internal business associates in their own business terms.

DATA CENTER STRATEGY

Over the last two years, we've evolved our data center strategy from an initiative that focused on data center consolidation to a broad-based approach that focuses on optimizing business value across the installed base of data centers.

Our new approach emphasizes long-range planning and business-based, financially guided decisions on how to respond to demand. We look at how to improve and fully utilize our existing data centers before adding new capacity. When growth is necessary, we increase financial efficiency by using modular and scalable building blocks, locating new data centers in aggregation sites when possible. Retrofits and new construction make use of the latest designs to improve energy and materials efficiency, and reduce the carbon footprint.

Evolving our Strategy

We are no longer focusing on the one-time goal of reducing our data centers to a specific number. Instead, we are focused on achieving optimal business efficiency through the predictable use of all IT and facilities assets. This long-term approach emphasizes three broad strategies:

- **Optimization.** Before we consider consolidating or constructing new facilities, we analyze the potential benefits of upgrading our existing data centers by optimizing the various components. This may include refreshing the servers and the storage and networking capabilities, or the facilities themselves depending on which achieves the highest return on investment.
- **Utilization.** We focus on increasing the available capacity of our existing data centers. For example, virtualization allows us to install multiple applications on a single

We learned that data center consolidation wasn't appropriate for all data centers. Each manufacturing facility, for example, requires a local data center as a fail-safe for the production process, reducing the likelihood of an interruption in the 24x7 manufacturing process. Within the engineering computing environment, our interactive workload required a data center nearby to eliminate the latency induced by the distance from the design data.

server, increasing utilization by a factor of five. In turn, this helps to maximize the infrastructure utilization of our existing data centers—saving on space, floor loading, power, and cooling—and minimize our capital investments. We also look at opportunities to pool batch servers into virtualized data centers, driving utilization of those servers to approximately 80 percent or above. Improved infrastructure utilization also means greater energy efficiency, as fewer watts are consumed by idle equipment.

- **Data center investment.** When demand grows beyond existing capacity, we invest in low-cost modularized facilities that can quickly come online. Wherever possible, we locate these facilities in centralized data center aggregation points. For interactive design and manufacturing sites, we continue to invest in local data center capabilities.

We are using these strategies to increase efficiency while reducing the cost of key physical assets such as data centers, servers, and storage, which account for a significant proportion of our operating costs. For each of these areas we have different sets of goals, measurements for success, and appropriate actions to cumulatively drive affordable computing services.

For example, one of the major goals of our optimization strategy is to reduce the power usage effectiveness (PUE) ratios of our data centers and thus increase the balance of watts available to support compute demands without increasing our carbon footprint. In addition, a central goal of our compute server strategy is to improve server performance, which can be achieved by purchasing higher performing and more energy-efficient servers and refreshing our servers on a four-year cadence. This leverages power and performance improvements made available by Moore's Law. Both strategies support the overall goals of data center efficiency.

We can achieve optimal efficiency by analyzing technology, business, and demand requirements and applying the most appropriate combination of strategy and activities. This may include a combination of:

- Accelerating server refresh.
- Consolidating and virtualizing servers.
- Upgrading facilities to improve their capability, utilization, and energy efficiency.
- Adding capacity in a modular, scalable way.
- Locating new data centers in aggregation sites whenever possible.
- Optimizing WAN configurations to substitute for local data center capacity.
- Optimizing storage.

We have expanded our focus to include changes to workflow models, tiered services models, cost transparency, business continuity, disaster recovery, and virtualization.

Process and Communication

With so many financial and technical interdependencies, we have found that organizing and executing our processes are critical to making progress in data center efficiency. The development and execution of our strategy requires coordination among business units, IT, corporate finance, facilities engineering, planning, and senior management.

It is critical that we look at costs holistically rather than from each group's isolated perspective. In addition, we will execute the following:

- **Architecture.** The blueprint for how to design and build, what to build, where to build, and how to integrate the many interdependencies into the decision making process.
- **Engineering.** The design and assembly of new technologies into the building blocks that deliver IT infrastructure solutions.

- **Operations.** The consistent and standardized operational processes that contribute to running our data centers like factories.
- **Finance.** The cash implications and consistent unit cost analysis to capture the cost of producing a cycle of compute or megabyte of content.
- **Communications.** The means of aligning and coordinating the activities that support our strategy through the Internet and other communications vectors.
- **Planning.** The integration of business requirements, technology feasibility, and financial analysis to deliver cost-effective services.

Planning for the Long Term

To anticipate compute and storage demands, our planning team collects both short-range and long-range demand requirements from each business unit—our customers—once a quarter. We then review and analyze this demand to determine if our data center facilities can meet the requirements. This drives the preparation of a supply response. If we cannot meet the demand requirements, we identify the most cost-effective approach by performing a financial and technical analysis comparing alternatives. In this way, we gain a holistic view of overall business requirements, which helps us forecast over two to three years to strategically plan for additional capacity.

This long-term view allows us to make smart investments in our data centers that pay for themselves. For example, by leveraging application virtualization and replacing four-year-old servers with new high-performing multi-core systems, we've reduced spending in data center facilities capital investment by 65 percent in 2008 compared to similar spending in 2006 while increasing performance by 2.5x (see Figure 2).

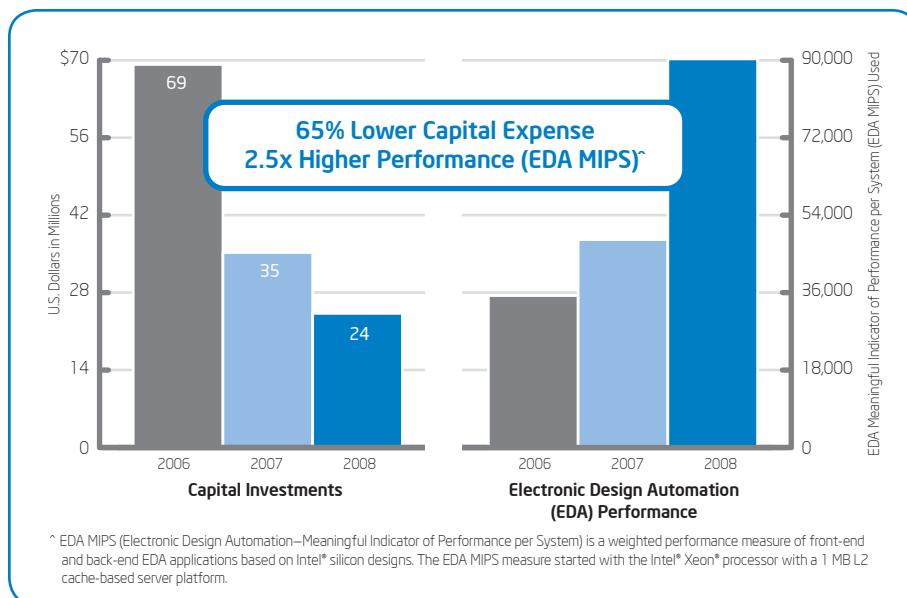


Figure 2. We reduced data center infrastructure capital investments by 65 percent from 2006 to 2008 while increasing performance 2.5x in the same time period.

FACILITIES DESIGN STRATEGY

As compute demands continue to increase, we search for ways to reduce costs and run our data center facilities and IT infrastructure more efficiently.

In designing an efficient data center, we take into account power, cooling, space, and floor load requirements. An imbalance of any one of these components will constrain the data center's capacity. We tackle this challenge two ways: (1) We use modular designs that employ the latest technologies to allow these characteristics to be addressed incrementally, and (2) we build higher density data centers, where applicable, rather than spreading out server load. Increasing the density of the computing power contained in each data center is generally a much more cost-effective and energy-efficient way to design the data center. Scaling up to add capacity reduces overhead for cooling, building material, cabling, and other costs.

We are also optimizing specific facilities to eliminate waste in cooling and are achieving measurable results (see sidebar story). Intel's goal is to reduce energy usage by

Increasing Data Center Power Efficiency

We are taking a variety of innovative approaches to improving data center facilities power usage and lowering power usage effectiveness (PUE). Here are a few examples:

Wet-side economizers. Our newest data centers employ decoupled wet-side economizers to provide free cooling as much as 40 percent of the year.

Air economizers. In a 10-month study, we used an air economizer to cool production servers using 100 percent outside air at temperatures up to 90 degrees. The results showed an estimated power savings of 67 percent using the economizer 91 percent of the time, as well as a lower rate of failure than with the standard air-conditioned control environment. Based on our study, the use of air economizers could potentially reduce annual operating costs by up to USD 2.87 million for a 10-megawatt data center.

Containment. Air leakage in many of our data centers caused cooling loss that resulted in sub-optimal capacity utilization. Open spaces in the racks allowed the hot air to flow through the cabinet and mix with the cooling air, reducing capacity. We addressed this problem in two upgrades. Our first-level upgrade re-routed the wires and cables out of the airflow path. We then blocked empty slots to prevent the intermixing of hot and cold air. This simple bypass management solution doubled our cooling capacity from the typical 2 to 4 kilowatts (kW) to 4 to 8 kW per rack data center. In our second-level upgrade, we introduced additional containment and isolated the hot air from the cool air using "chimney cabinets." This innovative approach delivered even more impressive results, doubling the cooling capacity again to 8 to 14 kW, helping us drive toward lower PUE.

White papers on these and other Intel data center innovations and approaches can be found on www.intel.com/it.

five percent in 2009 and continue to reduce energy usage every year thereafter.

Another goal is to have metering and monitoring capabilities in our data centers so we can properly measure our energy consumption. Preliminary results from a sampling of our data centers indicate an annual average PUE ranging from 1.2 in our newest high-density data centers to greater than 2.0 in our older facilities. Our expanded knowledge about our worldwide energy consumption will allow us to use best practices to optimize and improve our energy consumption over the next few years.

Looking ahead, additional cost savings and efficiency will come as we adopt new technology and modularized approaches to our data centers, allowing us to deploy additional highly efficient capacity as needed. We continue to run various proofs of concept (PoCs) to test the feasibility of innovations such as enterprise-level instrumentation and metering, containerization, air-side economization, and high ambient temperature operation in our data centers. This ongoing effort to challenge legacy designs and technologies offers the potential for breakthroughs that may radically change the way we design, build, and operate the data centers of the future.

COMPUTE SERVERS

In the past two years, we have made major changes in our compute server strategy to increase efficiency while reducing cost and energy consumption.

These changes include accelerating server refresh based on the latest high-performing Intel® architecture-based servers, virtualizing our server resources, and utilizing data center virtualization.

Optimizing Performance and Total Cost of Ownership with Accelerated Server Refresh

The demand for server compute capacity at Intel is increasing at a rapid rate as the complexity of design computing grows (see Figure 3). Because more than 70 percent of our server purchases are for design computing, it's critical to achieve the optimal performance for our workloads while increasing utilization and reliability.

In the past, we focused on maximizing the useful life of servers. Like many organizations, we depreciate the value of a server over four years. However, we typically continued to use servers even after they were fully depreciated,

rather than replace them. While well-intended, this practice actually increased the total cost of ownership (TCO) because older machines run less efficiently and require more maintenance. Maintenance beyond the initial four years is also typically more expensive.

We have changed our approach to one that focuses on optimizing our server purchases based on higher performing systems and refreshing our servers once they are fully depreciated. Server refresh enables us to accelerate the shift to more power-efficient and cost-effective servers based on multi-core Intel® Xeon® processors. This translates into a decrease in both the data center footprint and power requirements.

The following examples help to illustrate how using the highest performing processors and refreshing servers every four years has increased efficiency in our data centers.

- We found that the performance of a server equipped with the latest quad-core Intel® Xeon® processor 5500 series was 13x higher than a four-year-old server based on a single-core processor. This means that 13 four-year-old servers in the data center could be replaced with just one new server (see Figure 4).

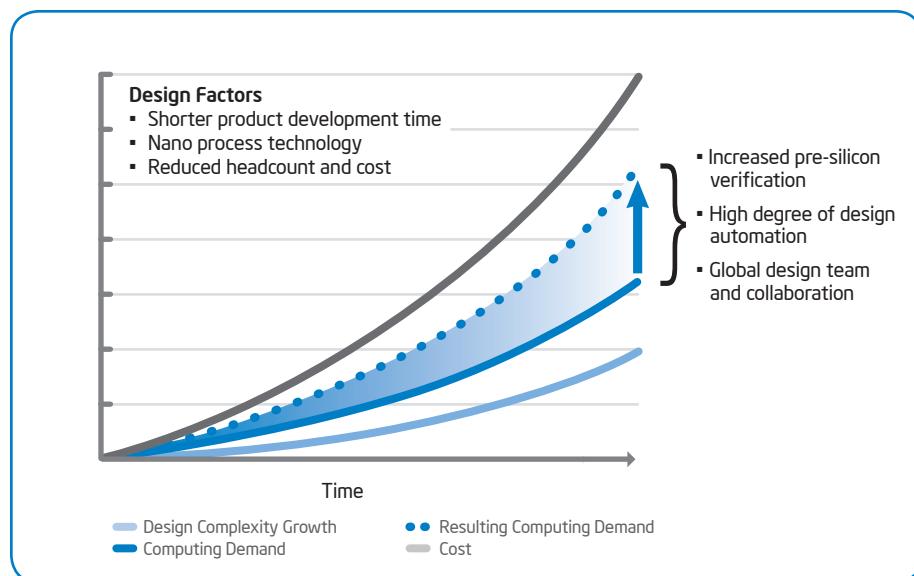


Figure 3. Increased design complexity and business requirements lead to exponential compute capacity growth and increased costs.

- Our studies also found that using one server with a quad-core Intel® Xeon® processor X5570 today can take the place of 10 servers based on the single-core 64-bit Intel® Xeon® processor with 2 MB L2 cache purchased in 2005, offering equal design application throughput (see Figure 5).³

Our studies also show that the high-end Intel Xeon processors are delivering the maximum

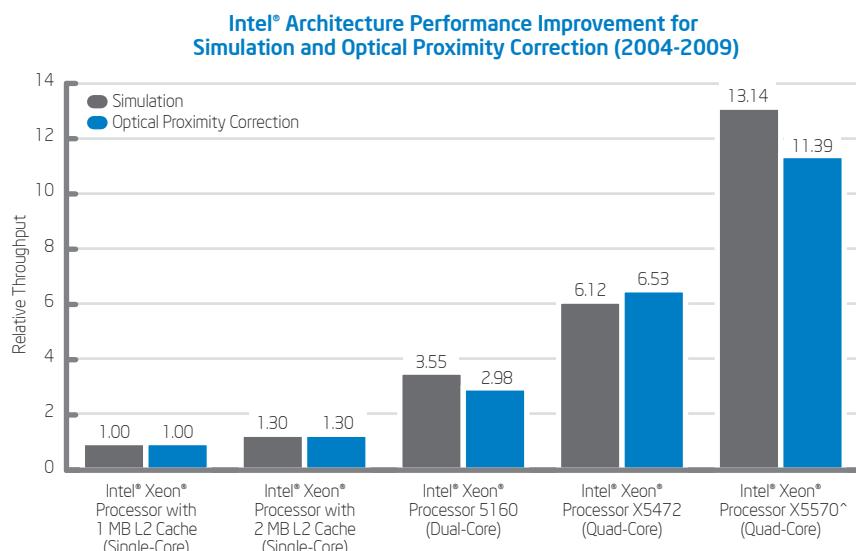
3 Intel internal measurements May 2007, November 2007, and February 2009.

value to Intel IT. Our analysis showed that four-year server TCO is dominated by the costs associated with software, such as application, database, OS, and middleware, which is 3.8x the cost of the hardware platform. The server platform, and hence the processor, accounts for only a small percentage of overall cost, as Figure 6 shows. As a result, the cost difference between platforms with low-end and high-end processors has little impact on overall server TCO. Because high-end processors

substantially increase performance, they deliver better value to Intel IT.

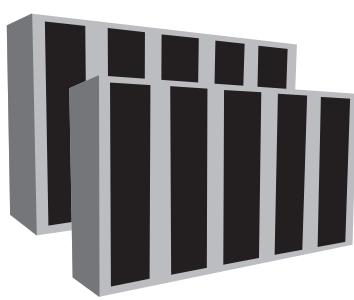
Virtualizing Server Resources

Over the last two years, we have focused our efforts primarily on optimizing efficiency and performance in both our design and manufacturing environments. These data centers make up most of our server and network environments and provide critical differentiated capabilities for Intel's business strategy.



^ For the throughput tests on the Intel Xeon processor X5570, Intel® Hyper-Threading Technology (Intel® HT Technology) was enabled for the simulation application and disabled for the Optical Proximity Correction application. For more details, see "Improving EDA Batch Application Performance," Intel Corporation, March 2009.

Figure 4. Servers using the quad-core Intel® Xeon® processor 5500 series attain the highest throughput for simulation and optical proximity correction.



2005

- 10 Racks
- 840 Servers
- 4 GB RAM per Server
- 400 Square Feet (Sq. Ft.)
- 225 Kilowatts (kW)
- Intel® Xeon® Processor (3.8 GHz) with 2 MB L2 Cache



2009

- 1 Rack
- 84 Servers
- 48-96 GB RAM per Server
- 40 Sq. Ft.
- 30 kW
- Intel® Xeon® Processor X5570 (2.93 GHz)

Same Performance/Throughput Using 90% Less Data Center Space and 86% Less Power

Figure 5. One rack of servers based on Intel® Xeon® processor 5500 series can take the place of 10 racks of servers based on Intel® Xeon processor-based servers while offering equal design application throughput. Standardizing on high-end processors has helped us attain a 90 percent reduction in data center space and an 86 percent reduction in power for equal electronic design application performance throughput.

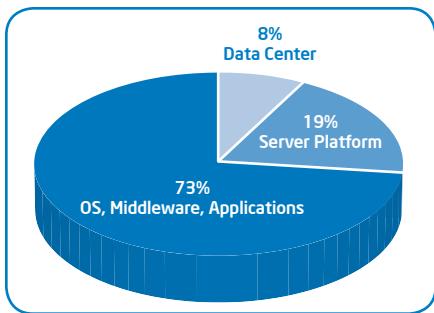


Figure 6. The hardware platform accounts for a small proportion of server total cost of ownership (TCO).

As we move into 2010, we are planning to accelerate our infrastructure efficiency efforts in the office and enterprise environments through virtualization. Currently, we have virtualized 10 to 20 percent of the servers in these data centers, which has improved server utilization, further reduced physical hardware cost and energy consumption, and increased ROI. Furthermore, we are able to provision new applications within days rather than weeks, allowing us to respond more quickly to rapidly changing business needs while reducing overall costs.

Our testing with the new Intel Xeon processor 5500 series and past experience with virtualization deployments demonstrates that Intel Xeon processor 5500 series can support approximately twice the number of virtual machines as previous generation two-

socket servers for the same TCO, and can achieve consolidation ratios of up to 20:1.

As part of our four-year server refresh strategy, we plan to aggressively deploy the Intel Xeon processor 5500 series in our office and enterprise environments to further accelerate virtualization adoption, with a goal of virtualizing 70 to 80 percent of these environments over the next several years.

Data Center Virtualization

Traditionally, Intel semiconductor design groups have relied mostly on servers in local design computing data centers. To obtain the growing computing power needed for their projects, design groups typically added enough local capacity to meet peak demand. The peaks and valleys of use meant that in between the demand peaks, many of these local servers were underutilized. Worldwide, this added up to significant underused capacity.

We have addressed this issue through data center virtualization, which lets design teams use idle servers at remote data centers. Applying Intel's global computing resources to individual projects has helped to increase overall server utilization and smooth fluctuations in demand, reducing the need to add local compute capacity and cutting costs. Figure 7 illustrates how data center virtualization has helped us maximize the use of existing server capacity.

Our goal is to make using remote resources highly automated and transparent. Design teams should be able to use computing resources anywhere, without being concerned about their physical location. The bulk of the batch simulation design activities that account for nearly 70 percent of the available compute cycles in a silicon design project can be served through this approach. However, interactive design and backend design activities require local compute capacity.

Data center virtualization is an essential element of our broader data center efficiency program because eliminating site dependencies enables us to consolidate data centers and boost utilization. We have located batch data processing in three major aggregated sites. Batch processing software has enabled engineers at geographically dispersed design sites to access the new large pools of server resources, which helps to drive global compute capacity utilization—targeted to run at 80 to 85 percent—to meet peak project demands (see Figure 8).

Compute Strategy Results

Our focus on server refresh has delivered significant results, reducing the footprint of the data center and increasing performance at lower power requirements. In 2008, we removed 20,000 servers based on single-core processors and replaced them with

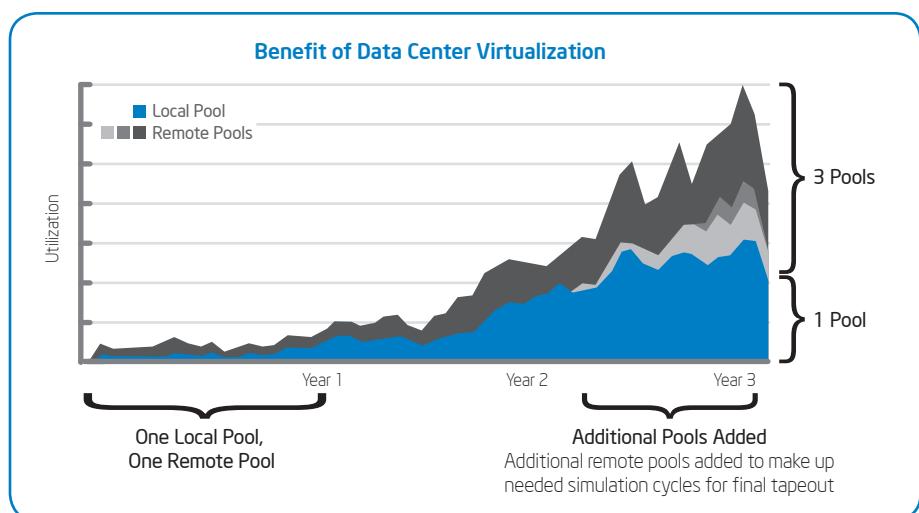


Figure 7. Data center virtualization helps to increase overall server utilization, ultimately reducing costs.

4,000 servers based on the quad-core Intel Xeon processor 5400 series. These server refresh activities enabled a savings of USD 45 million for 2008:

- USD 5 million in operating cost savings
- USD 40 million in reduced, delayed, or eliminated data center capacity

In addition, because of the power savings from the newer servers, Intel was able to take advantage of an energy conservation program in the state of Oregon, resulting in a USD 250,000 incremental benefit for a five-percent utility run rate consumption reduction.

Through data center virtualization, we have boosted utilization rates for our batch design computing data centers from 55 percent in 2006 to nearly 80 percent in 2008 (as shown in Figure 8). We have developed tools to make this possible, including a scheduler that assigns jobs to server pools spread across multiple data centers and software that automatically provisions OSs to as many as 100 servers in 30 minutes. Approximately 50 percent of all batch jobs now run remotely.

While increasing utilization, we also improved throughput time for many projects by making additional compute capacity available in remote data centers. This enabled design groups to meet or improve on key milestone dates for their projects.

Significantly, our increased use of optimization tools, such as virtualization and server refresh, has resulted in a 65 percent reduction in the data center facility capital investments over the last three years.

NETWORK ENVIRONMENT

We continue to improve our data centers by increasing standardization in our network configurations. Standardization includes physical interconnects, cabling, and physical and logical network topologies.

Instead of the customized, special-purpose data center approach we used in the past to meet the needs of individual business units, we now design modular multi-tenant data centers that can serve a larger community of users across the DOME environment. This type of design is efficient and cost effective because we use the same network standards and reference designs, service offerings, and access layer infrastructure in every data center.

We expect significant savings in LAN costs resulting from improved computing efficiency. For example, we expect that server consolidation will decrease the number of network ports, reducing cost because we require fewer network components. This in turn reduces the cable infrastructure cost.

WAN Optimization: Managing Bandwidth

As we shift computing resources from local data centers to aggregation sites around the world, we increased our WAN investment to provide enough bandwidth to maintain service-level agreements for many applications.

Table 1 gives examples of WAN latency between Intel data center sites in three different regions.

We also researched and implemented WAN optimization technologies that feature caching, compression, and Transmission Control Protocol (TCP) optimization to further minimize the investment required. In lab and production tests, we found that these tools were extremely effective in reducing the load on our WAN, substantially accelerating data transfers and reducing transmitted file sizes by as much as 99 percent⁴ for some applications.

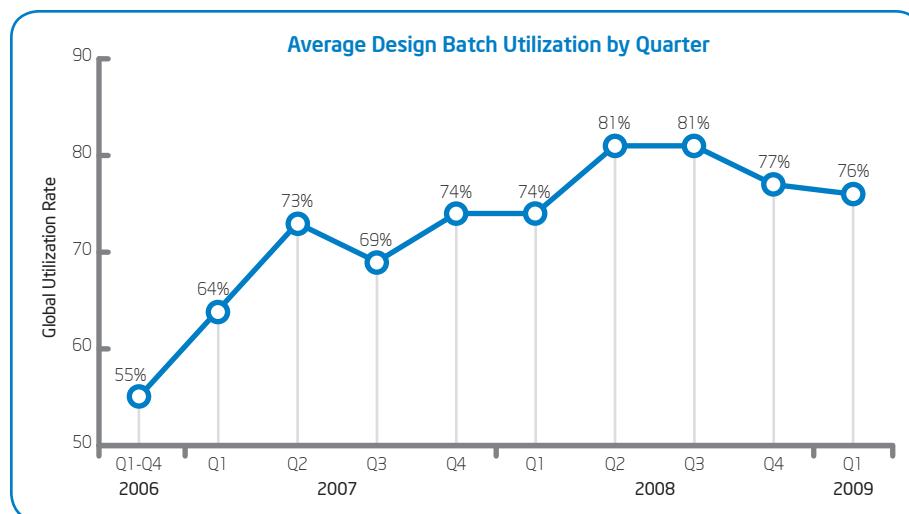
Optimizing our WAN network has enabled us to consolidate data centers at our smaller facilities, supporting the elimination of 45 data centers between 2006 and 2008.

Table 1. WAN Latency between Data Sites

Region	Delay
Within America Region	20-35 ms
Asia to/from America Region	200 ms
Europe to/from America Region	200 ms

4 "Optimizing WAN Performance for the Global Enterprise," Intel Corporation, May 2006.

Figure 8. Locating batch data processing in three major sites has significantly increased server utilization.



STORAGE

Today, we manage about 18 petabytes (PB) of primary and backup storage in our design computing, office, and enterprise environments. We anticipate that a variety of factors will stimulate future growth, including the increasing complexity of silicon designs, the growth of enterprise transactions, cross-site data sharing and collaboration, regulatory and legal compliance, and the ongoing need for retention.

Our present storage landscape is mapped to multiple computing areas: design, office, manufacturing, and enterprise. We choose storage and backup and recovery solutions based on the application use models for these respective areas, (see Table 2).

- Design computing.** Our silicon design computing uses the Linux* OS environment and primarily relies on network attached storage (NAS) for file-based data sharing. In addition to NAS, we use parallel storage for our high-performance computing (HPC) needs. We have more than 8 PB of NAS storage capacity and 1 PB of parallel storage in our design computing environment. We use slightly less than 1 PB of storage area network (SAN) storage in design computing, primarily to serve database and post-silicon validation needs.
- Office, enterprise, and manufacturing.** We rely primarily on SAN storage for block-level data sharing, with more than 8 PB of capacity. Limited NAS storage is used for file-based data sharing. For both NAS and

SAN, storage is served in a three-level tier model (Tiers 1, 2, and 3) based on required performance, reliability, availability, and cost of various solutions offered in respective areas.

- Backup, archive, and restore.** Backup, archive, and restore are major operations used in data management. We use both disk and tapes for our backups. Tapes are used for archive functions to facilitate long-term offsite data storing for disaster recovery. The tapes remain offline, which saves significant energy and offers a cost-effective solution. Our disk-based backups serve specific needs whenever faster backup and recovery are required. Our virtual tape library serves the disk-to-disk backup for faster backup and recovery needs, especially in the office and enterprise computing areas.

Optimizing Storage

We have focused our overall storage strategy on optimizing metrics for reliability, availability, performance, and scalability. Achieving results in these areas will in turn increase efficiency and reduce capital and operational costs.

Although our storage strategy encompasses office and enterprise data centers, our emphasis at this time is optimizing storage in the area of design computing.

In design computing, we identified storage performance and scalability as the major bottlenecks: The storage solution we had used historically could serve only 400 distributed clients and had a 400-GB volume size limit.

In fact, we needed to support up to at least 4,000 clients and up to 3 TB in volume.

Remaining with our previous storage solution meant adding 10 additional storage server racks and possibly even more over time for 4,000 clients, which would increase our data center footprint as well as power and cooling costs. This was not an option. In addition, replicating large design data across multiple storage servers to work around scalability limitations affected the productivity of the design engineers.

Taking a different approach, we decided to research parallel storage solutions that would not only satisfy our current storage needs, but also easily scale to future demands. The storage solution also needed to deliver higher performance at a significantly lower cost of ownership.

Parallel Storage Milestone

The deployment of our parallel storage solution is a milestone: It's the first time such a solution has been deployed in production for specific use in the semiconductor industry as well as in an IT organization.

Since we replaced our previous storage solution with a parallel storage solution, we have saved USD 22.68 million over three years (2006 to 2008) and have achieved significant results in a number of areas:

- Scalability.** We were able to replace every 10 conventional storage servers with one parallel server. This translates into 10:1 consolidation and huge costs savings in both space and energy usage.

Table 2. Intel IT's Storage Landscape

	Design Computing	Office Infrastructure Computing	Manufacturing Computing	Enterprise Computing
Network Attached Storage (NAS)	High- and medium-performing servers	Medium-performing servers	Medium-performing servers	Medium-performing servers
Storage Area Network (SAN)	Tier 2	Tier 2	Tier 1 and Tier 2	Tier 1, Tier 2, and Tier 3
Backup and Recovery	Tape library, disk-to-disk	Tape library, disk-to-disk	Tape library, disk-to-disk	Tape library, disk-to-disk

- Performance.** For select applications, we saw a performance improvement of greater than 300 percent from the parallel storage server compared to the previous storage solution.
- Volume size.** The volume size increased by a factor of 16, from 400 GB to 6,400 GB. We achieved the ability to support 6,400 GB without compromising our backup, archive, and restore service levels.

Table 3 summarizes the improvements we've gained over the previous solution.

Future Alignment

The HPC system we chose is customized for our design needs and aligns with Intel's process technology roadmap. Looking ahead, this storage solution will continue in alignment as Intel integrates improvements in process technology—moving from 45nm to 32nm technology and eventually to 22nm technology.

NEXT STEPS

Through our ongoing involvement in industry research and technology development groups, we continue to have active dialogues and evaluate emerging technologies that have the potential to reduce our energy consumption and improve our data center efficiency even further.

Here are some examples of our current areas of exploration.

Cloud Computing

Our initial cloud strategy, developed in mid-2008, incorporates three models: software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS). Intel has been using SaaS from third parties for several years and has experienced success in application domains that are not competitive differentiators for the business and which have standardized workflows. Intel has not used external PaaS and IaaS services to any significant degree.

As a part of our overall cloud computing strategy, we are pursuing a two-pronged approach: improving the internal hosting environment—making it more cloud-like and improving agility, availability, efficiency, and cost-effectiveness—as well as exploring external clouds. In the first half of 2009 we completed multiple PoCs to study external PaaS and IaaS offerings, which provided us additional data to help us refine our future plans. Based on the results from these PoCs, we have concluded that at this time the outstanding concerns pertaining to security, compliance, TCO, and lack of standards limit applicability of these external clouds to specialized scenarios in the near term. However, the pace of innovation in the cloud space continues at a high rate, and we continue to monitor the rapid developments of new cloud computing technologies and services that may result in broader adoption of these services in the Intel computing strategy at a later date.

IT Sustainability

Adding to our ongoing focus on sustainability, in 2009 Intel IT created a sustainability program

office to establish a CO₂ usage baseline and address consumption and waste. So far, our efforts have given us a clear understanding of our direct IT carbon footprint.

The program has also established a roadmap for energy reductions that aligns with the corporate goal of a five percent reduction in 2009. We are currently involved in many active projects—ranging from application end-of-life and server refresh to power metering and monitoring innovations and proof-of-concept testing. Collectively, we anticipate these projects will save or avoid more than five to 10 percent of our CO₂ baseline. The team also recognizes successes in these areas, which is helping to bring awareness of our accomplishment across the organization and embedding sustainability principles into our decision making as well as our day-to-day business tasks.

Solid-State Drives

We're investigating the use of solid-state drive (SSD) technology for enterprise and design computing. We have found that SSDs on a notebook client can increase a user's productivity by reducing the time it takes to start up the laptop and load applications. We are currently conducting a PoC to determine the benefit of SSDs in the server space. SSDs may be particularly useful in design computing where large amounts of memory are required, helping to achieve higher performance. In the enterprise environment, we are studying the use of SSD technology for virtual machine boot and swap.

Table 3. Storage and Backup Improvements

Area	Category	Previous Solution	Parallel Storage Solution
Storage and Backup (I/O throughput, scalability, and performance)	Theoretical I/O operations per second	512	5,210
	Practical I/O operations per second	Unknown	3,500+ MBps
	Volume size	400 GB	6,400 GB
	Single-stream performance	70 MBps	160 MBps

AN ON-GOING INITIATIVE

With our strategy in place to guide future data center development, we know that we are not pursuing a one-time goal. As business conditions continue to change and we investigate innovative technologies, our data center strategy must be dynamic enough to integrate these new factors into our long-term vision.

To achieve this vision requires us to balance resources, using the full capacity of our existing data centers. It also means knowing 18 to 20 months ahead of time when a data center will reach maximum capacity, triggering a decision to add a new one.

We have also realized the importance of communication, planning, and coordination among all of stakeholders. As with many IT organizations, understanding the big-picture view and gaining buy-in from all parties requires input from business units, corporate finance, facilities engineering, and senior management.

Through our long-term planning approach that emphasizes three strategies—optimization, utilization, and strategic investments—we will continue to drive broader efficiencies across the data center environment. Focusing our strategies on the key areas of facilities, compute servers, networking, and storage will help us respond more quickly to business needs and in turn enhance the services and value that IT brings to the business.

For more straight talk on current topics from Intel's IT leaders, visit www.intel.com/it.

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, reference www.intel.com/performance/resources/benchmark_limitations.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

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

FOR MORE INFORMATION

For an in-depth look at the data center innovations and approaches discussed in this paper, download additional IT@Intel publications at www.intel.com/it.

- "IT@Intel Information Technology 2008 Performance Report"
- "A Quantitative Approach to ERP Server Platform Sizing"
- "Building a Long-Term Strategy for IT Sustainability"
- "Developing an Enterprise Cloud Computing Strategy"
- "Faster Chip Design with Intel® Xeon® Processor 5500 Series"
- "Implementing Virtualization in a Global Business-Computing Environment"
- "Improving EDA Batch Application Performance"
- "Increasing Data Center Efficiency through Metering and Monitoring Power Usage"
- "Realizing Data Center Savings with an Accelerated Server Refresh Strategy"
- "Reducing Data Center Cost with an Air Economizer"
- "Reducing Data Center Energy Consumption with Wet Side Economizers"
- "Reducing Storage Growth and Costs: A Comprehensive Approach to Storage Optimization"

- "Selecting Server Processors to Reduce Total Cost"
- "Solid-State Drives in the Enterprise: A Proof of Concept"
- "Virtualization with the Intel® Xeon® Processor 5500 Series: A Proof of Concept"

ACRONYMS

DOME	design, office, manufacturing, and enterprise
HPC	high-performance computing
IaaS	infrastructure as a service
kW	kilowatt
NAS	network area storage
PaaS	platform as a service
PB	petabytes
PoC	proof of concept
PUE	power usage effectiveness
SaaS	software as a service
SAN	storage area network
SSD	solid-state drive
TB	terabytes
TCO	total cost of ownership

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 Xeon are trademarks of Intel Corporation in the U.S. and other countries.

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

Copyright © 2009 Intel Corporation. All rights reserved.

Printed in USA
0909/KAR/KC/PDF

Please Recycle
321370-001US

