

Intel IT's Data Center Strategy for Business Transformation

Our refined data center strategy has created new business value in excess of USD 184 million from 2010 to date.

Shesha Krishnapura
Senior Principal Engineer, Intel IT

Shaji Achuthan
Senior Staff Engineer, Intel IT

Bob Barnard
Financial Analyst, Intel IT

Vipul Lal
Senior Principal Engineer, Intel IT

Raju Nallapa
Principal Engineer, Intel IT

Sanjay Rungta
Senior Principal Engineer, Intel IT

Ty Tang
Senior Principal Engineer, Intel IT

Executive Overview

To better meet Intel's business requirements while providing our internal customers with optimal data center infrastructure capabilities and innovative business services, Intel IT has overhauled our data center strategy. Our data center transformation strategy is to run Intel data center services like a factory, effecting change in a disciplined manner and applying breakthrough technologies, solutions, and processes.

We use three key metrics to measure data center transformation success: meet growing customer demand (service-level agreements and quality of service) within constrained spending targets (cost-competitiveness) while optimally increasing the utilization of infrastructure assets (operational efficiency).

Building on previous investments and techniques, our refined data center strategy has created new business value in excess of USD 184 million from 2010 to date. Our key achievements include the following:

- We developed a system software capability called NUMA-Booster, which has delivered USD 55 million in additional server capacity.
- We deployed more than 13,000 Intel® Solid-State Drives as "fast swap" drives, which generated a 27 percent increase in server capacity.
- Four generations of high-performance computing in our Design computing environment created a 30x capacity increase and a 20x quality improvement.
- We adopted new storage capabilities, accelerated storage refresh, and focused

on increasing utilization, generating USD 33 million in cost avoidance.

- We deployed more than 18,000 10 gigabit-per-second network ports, generating more than USD 20 million in cost avoidance.
- An integrated server and network infrastructure provided a 39 percent reduction in hardware across the enterprise.

Over the 2013–2015 time frame, we plan to extend the data center strategy to continue to transform our data center infrastructure. We will do so by using disruptive server, storage, network, and data center facility technologies that can lead to unprecedented quality-of-service levels and total cost of ownership reduction for business applications—all while continuing to improve IT operational efficiency.

Our data center transformation strategy is key for Intel IT to stay competitive, compared to public cloud services. Implementing breakthrough solutions and pursuing aggressive goals are critical factors to success in this transformation.

Contents

Executive Overview..... 1

Background..... 2

 Meeting Compute Environment Challenges..... 2

 Aligning Data Center Investments with Business Needs..... 3

Unique Elements of Our Data Center Strategy 3

 Defining Key Performance Indicators and Goals 4

 Stimulating Bold Innovation through a New Investment Model..... 5

 Implementing a New Unit-Cost Financial Model 5

Results: Building on the Past, Building for the Future..... 6

 Cumulative Results from 2003 to 2013 6

 Results from 2010 to 2013..... 10

Summary of Data Center Best Practices 11

Plans for 2013 through 2015..... 13

Conclusion..... 13

Acronyms..... 14

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

Intel IT operates 64 data centers housing approximately 55,000 servers that underpin the computing needs of more than 104,000 employees.¹ To support the business needs of Intel's critical business functions—Design, Office, Manufacturing, Enterprise, and Services (DOMES)—while operating our data centers as efficiently as possible, Intel IT has engaged in a multi-year evolution of our data center strategy, as outlined in Figure 1.

In the past, we focused our data center investments on improving IT infrastructure as a means to deliver a foundation for the efficient growth of Intel's business. Our primary goal was cost reduction through data center efficiency and infrastructure simplification while reducing energy consumption and our carbon dioxide footprint to improve IT sustainability.

¹ To define "data center," Intel uses IDC's data center size classification: "any room greater than 100 square feet that houses servers and other infrastructure components."

Meeting Compute Environment Challenges

Over the last several years, we have reduced data center energy consumption and greenhouse gas emissions, while at the same time meeting the constantly increasing demand for data center resources. We anticipate these growth rates to continue or even increase further:

- 30 to 40 percent annual growth in compute capacity requirements
- 25 to 35 percent annual growth in storage needs
- 30 to 40 percent annual growth in demand for network capacity

To address these challenges without negatively impacting service delivery, we developed and continue to rely on a number of established industry best practices in all areas of our data center investment portfolio—servers, storage, networking, and facility innovation. Since 2006, these techniques, which are described in detail later in this paper, have enabled us to realize hundreds of millions of U.S. dollars (USD) in cost savings while supporting dramatic growth.

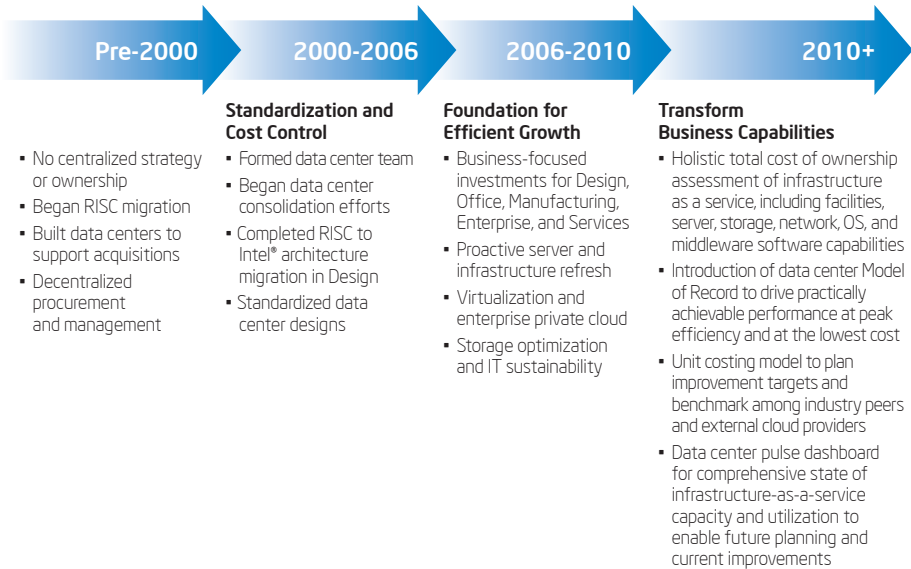


Figure 1. Intel's data center strategy is a continuous improvement process.

Aligning Data Center Investments with Business Needs

We have learned that a one-size-fits-all architecture is not the best approach for our unique business functions. After working closely with business leaders to understand their requirements, we chose to invest in vertically integrated architecture solutions that meet the specific needs of individual business functions.

DESIGN

Design engineers run more than 45 million compute-intensive batch jobs every week. Each job can potentially take several hours to complete. In addition, interactive Design applications are sensitive to high latencies caused by hosting these applications on remote servers. We have used several approaches in our Design computing data centers to provide enough compute capacity and performance to support requirements, including high-performance computing (HPC), grid computing, clustered local workstation computing, using solid-state drives (SSDs) as fast-swap drives, single-socket servers, and a specialized algorithm that increases the performance of the heaviest Design workloads.² Together, these investments enabled Design engineers to run up to 49 percent more jobs on the same compute capacity—which equates to faster design and time to market.

Because Design engineers need to access Design data frequently and quickly, we did not simply choose the least expensive storage method for this environment. Instead, we have invested in clustered and higher performance scale-out, network-attached storage (NAS) in combination with highly scalable parallel storage for our HPC needs. We use storage area networks (SANs) for specific storage needs such as databases.

² Intel uses grid computing for silicon design and tapeout functions. Intel's compute grid represents thousands of interconnected compute servers, accessed through clustering and job scheduling software. Additionally, Intel's tapeout environment uses an HPC approach, which optimizes all key components such as servers, storage, network, OS, applications, and monitoring capabilities cohesively for overall performance, reliability, and throughput benefits. For more information on HPC at Intel, refer to "High-Performance Computing for Silicon Design," Intel Corp., November 2013.

MANUFACTURING

IT systems must be available 24/7 in Intel's Manufacturing environment, so we use dedicated data centers for factories. We have invested heavily over the last few years to develop a robust business continuity plan that keeps factories running even in the case of a catastrophic data center failure. These efforts have paid off, and we have not experienced factory downtime related to data center facilities since 2009.

In our Manufacturing environment, we pursue a methodical, proven infrastructure deployment approach to support high reliability and rapid implementation. This "copy-exact" approach deploys new solutions in a single factory first and, once successfully deployed, we copy that implementation across other factory environments. This approach reduces the time needed to upgrade the infrastructure that supports new process technologies—thereby accelerating time to market for Intel® products. The copy-exact methodology allows for rapid deployment of new platforms and applications throughout the Manufacturing environment, enabling us to meet a 13-week infrastructure deployment goal 95 percent of the time—compared to less than 50 percent without using copy-exact methodology.

OFFICE, ENTERPRISE, AND SERVICES

To improve IT agility and the business velocity of our private enterprise cloud, we have implemented an on-demand self-service model, which has reduced the time to provision servers from three months to three hours. We more than tripled the number of virtualized applications inside the Intel IT Office and Enterprise environments in 2010, from 12 percent to 42 percent, and have achieved virtualization of 75 percent of applications in our Office and Enterprise environments.

In contrast to the Design environment, in the Office, Enterprise, and Services environments we rely primarily on SAN storage, with limited NAS storage for file-based data sharing.

UNIQUE ELEMENTS OF OUR DATA CENTER STRATEGY

Our transformational data center strategy is to run Intel data centers and all underlying infrastructure as if they were factories, with a disciplined approach to change management. By applying breakthrough technologies, solutions, and processes, we can lead the industry and keep up with the accelerating pace of Intel's business.

We have realized hundreds of millions of USD in cost savings since 2006 by proactively refreshing our infrastructure, adopting cloud computing, updating our network, pursuing IT sustainability, and consolidating data centers. In addition, we have supported business growth and capability improvements by deploying unique solutions that benefit Intel's critical business functions—DOMES. We have enhanced our strategy to include several new elements (as detailed in subsequent sections):

- **Key performance indicators (KPIs).** We have implemented three KPIs and have established goals for each of them:
 - Quality of service (QoS), using a tiered approach to service-level agreements (SLAs)
 - Cost efficiency
 - Effective utilization of assets and capacity

Based on improvements each year in technologies, solutions, and processes, we identify the best achievable SLA, the lowest achievable cost, and the highest achievable resource utilization. We call this combination the Model of Record (MOR) for that year. We set investment priorities based on these KPIs to move toward the MOR goal; each year we are getting closer to the MOR while at the same time balancing the three vectors.

- **Investment decision model.** Focusing on the MOR and comparing current data center capabilities to the best

achievable KPIs enables us to prioritize our investment decisions. This approach seeks to remove the conventional improvement mindset, which focuses only on incremental improvements. Instead, we are transforming our capabilities by identifying further groundbreaking innovations—like those already used to implement our private cloud and our highly efficient silicon design computing grid.

- **Unit-costing financial model.** By identifying metrics for improvements in each DOMES area, we can benchmark ourselves and further prioritize our investments.

We believe our new approach to data center costing and investment evaluation, along with a continued focus on meeting business needs, has stimulated a bolder approach to continuous innovation. Our efforts have improved the quality, velocity, and efficiency of Intel IT's business services, creating a sustained competitive advantage for Intel's business. For details, see ["Results: Building on the Past, Building for the Future."](#)

Defining Key Performance Indicators and Goals

The KPIs provide a means to measure the effectiveness of data center investments. Because the service output for each business function is different, we evaluate each business function separately. In our data center investment decisions, we seek to balance and meet all business requirements while optimizing the KPIs.

QUALITY OF SERVICE

We use a tiered approach to SLAs, tailored to each business function's sensitivity to performance, uptime, mean time to repair (MTTR), and cost. Our goal for this KPI is to meet specific performance-to-SLA requirements for defined tiering levels. For example, for our most mission-critical applications, we aim for a higher performance to SLA than for second-tier applications, which are less critical. The end goal and true measure of IT QoS is zero business impact from IT issues.

EFFECTIVE RESOURCE UTILIZATION

Our refined data center strategy represents a dramatic shift in how we view resource utilization. Historically, we measured utilization of IT assets—compute, storage, network, and facilities—by simply determining how busy or loaded an asset was. For example, if a server was working at peak capacity 90 percent of the time, we considered it 90-percent utilized. If 80 percent of available storage was allocated, we considered that 80-percent utilization.


In contrast, we now focus on the actual output of an asset—that is, *effective* utilization. For example, if Intel's design engineers start one million design jobs—thereby keeping the servers very busy—but a third of those jobs terminate before completion because there wasn't enough storage available, that is low effective utilization of compute capacity—only 66 percent. Or, if a customer consumes only 4 GB of a 10-GB storage allocation, the remaining 6 GB is essentially wasted storage—even though it is allocated—and does not represent effective utilization of this asset. Our goal for the effective utilization KPI is to achieve 80-percent effective utilization of all IT assets.

Intel IT Data Center Dashboard

To better monitor and manage our worldwide network of data centers, we developed and deployed an integrated business intelligence (BI) dashboard. This BI tool is modeled on a dashboard used in Intel's Manufacturing environment. By polling more than 192 million data records across our worldwide data center environment, the Intel IT Data Center Dashboard provides a single view of all our data center health metrics—at both the system and component levels. It is capable of delivering on-demand reports to users with a mere 5-second page load.

This dashboard will help us monitor our key performance indicators (KPIs) by highlighting the current state and opportunities for optimization, thereby enabling overall improvements that align with our data center strategy goals.

For example, the dashboard can report on effective utilization of several data center resources, including electronic design automation—meaningful indicator of performance per system (EDA-MIPS); raw and utilized storage capacity; and facilities space, power, and cooling. This data can report statistics by business function or by data center, and can be used to compare KPIs and metrics across several data centers. The figure to the right shows a sample of the dashboard.



The Intel IT Data Center Dashboard provides a holistic view of data center resources to help us track our KPIs and identify opportunities for optimization and improvement.

COST PER SERVICE UNIT

As shown in Table 1, different business functions have a different service unit that we can measure. This unit represents the capacity we enable for our business users.

Our goal for this KPI is to achieve a 10-percent improvement in data center cost efficiency every year. This goal does not necessarily mean we will spend less each year, but rather that we will get more for each dollar we spend. For example, we may spend less for the same number of service units, or we may spend the same amount but get more service output.

Stimulating Bold Innovation through a New Investment Model

Building on a time-tested methodology that has proven successful in Intel's Manufacturing environment over multiple process technology generations, we adopted a new data center investment decision model that compares current data center capabilities to a "best achievable model" that guides us to make investments with the highest impact.

Previously, Intel data center planning teams looked at existing capabilities and funding to establish a plan of record. This plan drove incremental improvements in our existing capabilities; our goal was to minimize total cost of ownership (TCO) and deliver positive return on investment (ROI).

In contrast, the MOR ignores the constraints imposed by what we have today. Instead, it identifies the minimum amount of resources we should *ideally* have to support business objectives—thereby establishing an optimal state with available technology.

By setting a standard of maximum achievable performance, the new model enables us to:

- Determine which investments will have the highest ROI.
- Identify the benefits of using disruptive infrastructure technologies and breakthrough approaches that deliver more optimal data center solutions across all aspects of our infrastructure.

- Make data center location decisions, including identifying potential data centers to consolidate, upgrade, or close.

The new model focuses limited available resources in specific areas for maximum holistic gain.

As shown in Figure 2, because technology is always changing, peak performance also changes—the maximum achievable performance keeps on getting better through innovation. We know that resource constraints make it impossible to ever actually achieve the standard set by the new investment model—although our HPC environment comes very close to that goal. However, the model enables us to identify gaps between where we are and where we'd like to be. We can then identify the biggest gaps in capability to prioritize our budget allocation toward the highest value investments first.

Implementing a New Unit-Cost Financial Model

We evolved our financial model from project- and component-based accounting to a more holistic unit-costing model. For example, we previously used a "break/fix" approach to data center retrofits. We would upgrade a data center facility or a portion of the facility in isolation, looking only at the project costs and the expected ROI of that investment, with no holistic view as to the impact of service unit output. In contrast, today we focus on TCO per service unit—using the entire data center cost stack per unit of service delivered. This cost stack includes all cost elements associated with delivering business services and now considers the worldwide view of all data centers in the assessment of our investments.

As shown in Figure 3, there are six major categories of cost to consider: network, headcount, servers, facilities, OS and management, and storage and backup and recovery. By adding these costs and then dividing by the total number of appropriate service units for the environment, we arrive at a cost per service unit.

Table 1. Service Unit for Each Business Function

Business Function	Service Unit
Design	Cost per electronic design automation-meaningful indicator of performance per system
Office, Enterprise, and Services	Cost per OS instance
Manufacturing	Cost per integrated factory compute environment




Figure 2. Our new data center investment model encourages innovation and provides significant business results.

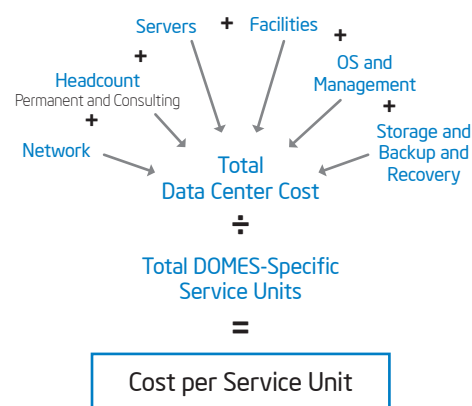


Figure 3. We arrive at a data center unit cost by considering all categories of cost and dividing by the number of units for that environment, such as electronic design automation-meaningful indicator (EDA-MIPS) of performance per system in Design and OS instances in Office, Enterprise, and Services.




Figure 4. Knowing the total unit cost, as well as the individual cost category figures for each business environment, enables us to better choose IT investments that will lower costs the most.




Figure 5. Over the last decade, even as we have met increasing demands for compute and storage resources, we have reduced our data center footprint by 24 percent.

Table 2. Cumulative Data Center Improvements (2003-2013)

Area	Improvement
Data Center-Wide	<ul style="list-style-type: none">Smaller total data center footprintImproved overall storage practicesImproved overall network practices
Design Environment	<ul style="list-style-type: none">More efficient Design compute and storage4th generation of high-performance computingIncreased Design throughput using NUMA-BoosterFaster Design throughput using Intel® Solid-State DrivesDeployment of single-socket servers
Office and Enterprise Environment	<ul style="list-style-type: none">More efficient Office and Enterprise compute and storage

Service-based unit costing enables us to benchmark ourselves and prioritize data center investments. Determining service-based unit costs also allows us to measure and compare the performance of individual data centers to each other, identifying which are underperforming and giving us the tools to decide whether to upgrade or consolidate underperforming data centers.

To show how the new unit-based costing model works, Figure 4 compares Design cost data and Office, Enterprise, and Services cost data. The headcount category accounts for a greater percentage of total cost in Office, Enterprise, and Services than it does in Design; in contrast, servers are more of a cost factor in Design than they are in Office, Enterprise, and Services. Knowing our exact unit cost in each environment, as well as the breakdown of that cost, enables us to develop optimized solutions for each environment that will have the greatest effect on cost efficiency and ROI.

RESULTS: BUILDING ON THE PAST, BUILDING FOR THE FUTURE

This section provides details on some of the improvements and cost savings our data center strategy has enabled over the years. We are building on the success we have already achieved through our data center initiatives over the last decade. Therefore, some of the results shown here are cumulative; others have been achieved over the last three years and are a direct result of our new MOR strategy. Our refined data center strategy will enable us to support the growth of Intel's customers, products, and acquisitions, as well as enhance the quality, velocity, and efficiency of the services we offer to Intel business groups.

Cumulative Results from 2003 to 2013

During the last decade, we have dramatically improved performance and reduced costs for our data centers. Table 2 enumerates several areas of improvement; each of these is discussed in the following subsections.

DATA CENTER-WIDE IMPROVEMENTS

We have improved the performance and cost efficiency of our data centers as a whole.

Smaller Total Data Center Footprint

Figure 5 shows how we have consolidated our data center facilities during the past 10 years. We have reduced the total square footage by 24 percent and reduced the number of sites from 152 to 64. See ["Continued Data Center Consolidation"](#) for a discussion of how the MOR strategy has directly enabled some of these achievements.

Improved Overall Storage Practices

A significant focus on effective utilization in our Design environment has enabled us to improve resource utilization from below 45 percent to more than 50 percent—our goal is to reach 65 percent. This increase in effective utilization has saved Intel USD 5.9 million.

We have applied several storage techniques to enhance storage efficiency and reduce costs:

- Tiered storage.** A five-tier approach to storage has helped us increase effective utilization of storage resources, improve our performance to SLAs, and reduce the TCO for Design storage. The tiers of Design storage are based on performance, capacity, and cost. Tier-1 servers have the highest performance and the least storage capacity. Tier-2 servers offer medium performance but greater storage capacity. Tier-3 servers provide lower performance but emphasize capacity, while Tier-4 and Tier-5 servers have the highest capacity but are used for low-frequency access and read-only archive data.

- **Scale-out storage.** We are making a strategic shift from a fragmented scale-up storage model to a pooled scale-out storage model. Scale-out storage better supports on-demand requests for performance and capacity. In addition, scale-out storage enables transparent data migration capabilities and increases the effective utilization of space freed by the use of efficiency technologies such as deduplication and compression.
- **Storage refresh cycle.** To improve performance and reduce cost, we are implementing a four-year storage refresh cycle. This accelerated refresh cycle enables us to take advantage of servers with better performance and more efficient energy use, thereby reducing both capital and expense costs. For example, a more energy-efficient server can reduce data center power usage; a more powerful server that replaces several older servers can reduce our data center footprint.
- **Data deduplication.** The introduction of new storage to support company growth and our commitment to a four-year refresh are enabling us to use the latest generation of Intel® Xeon® processors. These processors provide us with the processing power to handle data deduplication on our primary storage servers—freeing 1.3 petabytes of capacity, which we are making available for customer use.

Improved Overall Network Practices

To accommodate the increasing demands that data center growth places on Intel's network, Intel IT converted our data center network architecture from multiple 100 megabits per second and 1 gigabit Ethernet (GbE) connections to 10GbE connections. The older, slower connections no longer supported Intel's growing business requirements. The conversion to 10GbE started in 2010; we currently have deployed more than 18,000 10 Gb/s ports. Our new 10GbE data center fabric

design accommodates our current annual 40-percent network capacity growth (see Figure 6). We expect it will also meet increasing network demand in the future.

In addition to increasing the network capacity, we have also increased the effective utilization of network ports over the last three years, from 40 percent to 60 percent (see Figure 7). Higher utilization means we don't have to purchase additional ports to meet network capacity demand growth, saving Intel more than USD 20 million in cost avoidance.

We are also focusing on improving data center stability. In the past, we used a large installation of layer 2-based technology. We are now migrating to a layer 3-based network. This new architecture is enabling us to use all available bandwidth on primary and secondary paths at the same time. Therefore, we are able to use our network capacity more effectively. We are also able to eliminate the spanning-tree protocol within our data centers; this protocol does not scale well for large networks. Because the Internet uses layer 3-based, scalable architecture, using this concept within Intel's data center will make our data center network more scalable and resilient. Also, we are using other technologies such as overlay, multi-chassis link aggregation (MLAG), and tunneling to extend layer 2 across data centers, over the layer-3 topology.

As shown in Figure 8, we tend to adopt higher-speed network technology almost as soon as it is available in the market. We expect to adopt 40GbE technology within some data center environments by the end of 2013, and 100GbE by 2015, to keep pace with network demand.

IMPROVEMENTS IN THE DESIGN ENVIRONMENT

Because silicon chip design represents a significant portion of Intel's business, we have applied our data center strategy to several aspects of Design computing.




Figure 6. Our new 10 gigabit Ethernet (GbE) data center fabric design can accommodate current capacity growth as well as meet increasing network demand in the future.




Figure 7. We have continued to increase effective utilization of our network.




Figure 8. Intel IT adopts higher-speed network technology almost as soon as it is available.




Figure 9. Despite continuing growth in compute and storage demand, our Design data centers are using powerful Intel® technology to meet demand without increasing the number of servers.




Figure 10. Our high-performance computing (HPC) solution, combined with disciplined change management, has steadily increased compute capacity and improved quality of service.

More Efficient Design Compute and Storage

One of the major challenges in our Design environment is that server and storage growth is occurring at a high rate. Compute demand is growing 30 to 40 percent year over year, while storage capacity demand is increasing 30 to 35 percent annually (see Figure 9). As shown in the figure, the power of newer generations of Intel® architecture has enabled us to meet these demands without increasing the number of servers.

We expect server form factor changes to keep the number of servers flat, but the number of cores will continue to increase. We expect to measure data center performance based on number of cores, number of racks, power consumed, and the extent to which we meet the meaningful indicator of performance per system (MIPS) demand.

4th Generation of HPC

Designing Intel® microprocessors is compute intensive. Tapeout is a final step in silicon design, and its computation demand is growing exponentially for each generation of silicon process technology. Intel IT adopted HPC to address this large computational scale and realized significant improvements in computing performance, reliability, and cost.

As shown in Figure 10, our HPC solution has enabled a 30x growth in Design compute capacity from 2005 to 2012. We are now using the 4th generation of our HPC solution and will continue to develop new HPC generations as Intel process technology advances. Figure 10 also shows our commitment to quality. Through a disciplined approach to change management (basically running our data centers as if they are factories), we have reduced the number of compute issues that impact tapeout by 20x.

Increased Design Throughput Using NUMA-Booster

Overall data center optimization includes more than simply looking at server performance and facility efficiency. Application performance and workload optimization can also be contributing factors. We developed a system software capability called NUMA-Booster, which automatically and transparently intercepts our Design workloads and performs workload scheduling better than the default OS scheduling capability. Overall, NUMA-Booster has delivered USD 55 million in additional server capacity. We have achieved the following specific results without any system downtime or end-user impact:

- **Performance.** Our tests showed a 17-percent improvement in design performance (see Figure 11).

Register Transfer Logic (RTL) Performance with NUMA-Booster

Higher is Better




Figure 11. NUMA-Booster has increased Design compute performance by 17 percent.

System with 2x Intel® Xeon® processor X5570, 72 GB DDR3-1333 RAM, 1x 900 GB 10K RPM SAS hard drive, with Linux® 2.6 OS, running Intel silicon design simulation workload, Intel IT measurement.

System with 2x Intel® Xeon® processor X5675, 96 GB DDR3-1333 RAM, 1x 900 GB 10K RPM SAS hard drive, with Linux 2.6 OS, running Intel silicon design simulation workload, Intel IT measurement.

System with 2x Intel® Xeon® processor E5-2670, 128 GB DDR3-1333 RAM, 1x 900 GB 10K RPM SAS hard drive, with Linux 2.6 OS, running Intel silicon design simulation workload, Intel IT measurement.

System with 2x Intel® Xeon® processor E5-2680 v2, 256 GB DDR3-1600 RAM, 1x 900 GB 10K RPM SAS hard drive, with Linux 2.6 OS, running Intel silicon design simulation workload, Intel IT measurement.

- **Data center space and procurement costs.** We have deployed NUMA-Booster on approximately 20,000 servers, thereby reducing the footprint needed to meet demand by 2,200 servers (representing 45 racks of data center space).
- **Carbon footprint.** These 2,200 servers represent a savings of approximately 11.5 million kWh annually, which equals about 6,100 metric tons of CO₂.

We expect to reap even greater results from NUMA-Booster as we retire older servers and deploy newer multicore servers with NUMA-Booster capability.

Increased Design Throughput Using Intel® Solid-State Drives as Fast-Swap Drives

Silicon chip design engineers at Intel face the challenge of integrating more features into ever-shrinking silicon chips, resulting in more complex designs. The increasing design complexity creates large electronic design automation workloads that have considerable memory and compute requirements. We typically run the workloads on servers that need to be configured to meet these requirements in the most cost-effective way.

In Intel IT tests with large silicon design workloads, substituting lower-cost Intel® Solid-State Drives (Intel® SSD) for part of a server's physical memory resulted in a 1.63x performance-normalized cost advantage. Using SSDs as fast-swap drives increased the Design throughput of more than

13,000 servers by 27 percent, representing over USD 12 million in business value.

Deployment of High-Density, High-Frequency Single-Socket Servers

As shown in Table 3, we have determined that a higher frequency, lower core-count, single-socket server based on the Intel® Xeon® processor E3-1200 v3 product family provides better performance, compared to a lower frequency, higher core-count dual-socket server based on the Intel® Xeon® processor E5-2600 v2 product family. Due to this higher performance, these single-socket servers can reduce our application license needs up to 35 percent while completing virtually the same number of workloads. Based on this data, we are deploying single-socket servers in our Design computing environment.

MORE EFFICIENT OFFICE AND ENTERPRISE COMPUTE AND STORAGE

Similar to our Design environment, the compute and storage demands in our Office and Enterprise environment are also growing quickly. Nevertheless, as shown in Figure 12, we continue to meet that demand while steadily reducing the number of physical servers. From 2009 to 2012, we reduced physical servers by 60 percent and achieved an approximate 5.5x increase in the number of virtual machines (VMs). We also greatly increased the ratio of VMs to physical servers—from 2 percent in 2009 to 19 percent in 2012.




Figure 12. A high rate of virtualization combined with Intel® architecture has enabled us to meet growing Office and Enterprise compute and storage demand while significantly decreasing the number of required physical servers.

Table 3. Comparison of Performance between Single-Socket and Dual-Socket Servers

	Rack-Optimized 1U Dual-Socket Servers based on 2x Intel® Xeon® processor E5-2680 v2 (2.8 GHz, 10 cores per socket)	Density-Optimized Single-Socket Servers based on 1x Intel® Xeon® processor E3-1280 v3 (3.6 GHz, 4 cores per socket)
Data Center Space (Standard Rack Configuration)	1 rack	1 rack
Number of Servers per Rack	40	140
Number of CPU Sockets per Rack	80	140
Number of CPU Cores	800 (at 2.8 GHz)	560 (at 3.6 GHz)
Memory per Core	12.8 GB	8 GB
Application Licenses Needed per Rack	800	560
Relative Application Performance per Core for Our Workload	100 percent	135 percent
Application Workload Throughput per Rack	800 (800 × 1.0) 100 percent	756 (560 × 1.35) 96 percent

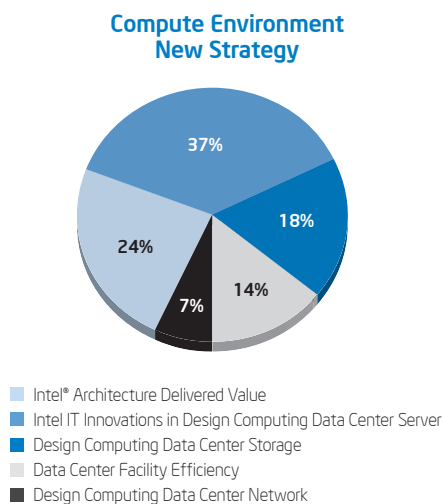


Figure 13. Our new data center strategy has delivered value throughout the enterprise.

Results from 2010 to 2013

As part of our 10-year cumulative efforts to improve data center efficiency and reduce costs, our innovative and transformational data center MOR strategy and execution over the last three years has resulted in over USD 184 million in business value. Here is a summary of the efficiency improvements and cost savings we have achieved in the Design environment from 2010 through the third quarter of 2013:

- **Design computing.** Intel IT innovations in the Design computing data center include the NUMA-Booster solution (17 percent higher performance); Intel SSDs (27 percent higher capacity at lower cost); faster, low-cost single-socket servers (35 percent higher performance); and procurement efficiency. Together, these innovations have resulted in USD 67.71 million in savings—in addition to the gain of USD 44.45 million offered by using newer Intel processor-based systems with similar form factors.
- **Design storage.** We have implemented Design computing data center storage efficiency improvements by adopting new technology capabilities and increasing utilization—generating USD 33.09 million in savings.
- **Design network.** The adoption of a multi-vendor strategy for our Design computing data center network, combined with a focus on reduction of expensive maintenance costs associated with older equipment, generated USD 13.32 million in savings.

Figure 13 summarizes how our new strategy has improved our computing environment.

In addition to the above contributions to the Design-specific environment, our new investment model has enabled us to identify other actionable gaps between the best achievable performance and our current plan.

These actions include the following:

- Continuing to consolidate data centers
- Reducing unit cost for both the Design and Office and Enterprise environments
- Extending our use of blade servers in the Office, Enterprise, and Services environments

CONTINUED DATA CENTER CONSOLIDATION

We used our new investment model to look at the number of data centers we have and the number we should have.

The new investment model identified opportunities to reduce the number of Intel data centers by as much as 35 percent, using techniques such as the following:

- Close, retrofit, or reclassify data centers and improve efficiency.
- Colocate local infrastructure with Design and Manufacturing data centers or provide services from a server closet.
- Manage local infrastructure sites remotely.
- Improve facility power efficiency through strategic investments.

We have targeted 32 inefficient and ineffective data centers, eliminating 61,770 square feet of data center space and converting 23,609 square feet of data center space to low-cost infrastructure rooms, saving Intel USD 25.45 million annually.

REDUCED UNIT COST FOR THE DESIGN ENVIRONMENT

Figure 14 details how our budget has remained relatively flat (left-side of the figure) while unit growth (middle section of the figure) has continued to rise in both the Design and Office and Enterprise environments. Our investment model has enabled us to reduce unit costs in both environments (right-side of the figure)—reducing Design unit cost by 44 percent and Office and Enterprise unit cost by 34 percent.




Figure 14. Our new strategy has enabled us to meet increasing growth and reduce unit cost without increasing our budget.

REDUCED TCO WITH BLADE SERVER TECHNOLOGY

As shown in Figure 15, our new investment model has shown us that moving from rack-mount servers to blade servers can reduce TCO in our cloud computing environment by about 39 percent. This reduction results from reduced port, network, and cable costs. For example, a group of 16 blade servers compared to 16 rack-mount servers requires only 8 Ethernet interfaces instead of 128, and only 4 Fibre Channel interfaces instead of 32. Deploying a newer generation of blade-server technology with converged network fabric within the blade chassis (labeled “Gen-2” in Figure 15) allowed us to reduce the cost even further.

Based on this data, we are actively deploying blade servers to support further virtualization efforts in the Office, Enterprise, and Services environments.

SUMMARY OF DATA CENTER BEST PRACTICES

Over the last decade, we have made many strategic investments and developed solutions to make our data centers more efficient and better serve the needs of Intel’s business. We are now applying our MOR approach across our entire infrastructure stack—compute, storage, networking, and facilities.

Table 4 provides a summary of the best practices we have developed and the business value they have generated.




Figure 15. Blade servers with integrated network/data fabrics have reduced our virtualized environment infrastructure costs by 39 percent.

Table 4. Intel IT Data Center Best Practices

Compute (Servers)	
BEST PRACTICE	BUSINESS VALUE
Regularly refresh servers using the latest generations of Intel® Xeon® processors	<ul style="list-style-type: none"> Virtualization ratios of up to 35:1 Reduced number of servers in Design by 43 percent between 2005 and 2013 Reduced Design environment energy consumption by 10 percent annually since 2008 Greater than a 6x increase in performance between 2005 and 2012
Deploy high-performance computing ⁱ	<ul style="list-style-type: none"> 30x increase in capacity, with a 20x increase in stability Relative throughput scaling has increased to 19.4x from 2006 to 2012 Saved USD 44.72 million net present value (NPV) from 2006 to 2010
Migrate applications from RISC to Intel® architecture ⁱⁱ	<ul style="list-style-type: none"> Enabled significant savings and IT efficiencies Allowed us to realize the benefits of industry-standard OSs and hardware
Adopt virtualization and cloud computing	<ul style="list-style-type: none"> Virtualized more than 60 percent of Office and Enterprise servers, amounting to a reduction of 4,000 servers in our data centers Reduced the time it takes to provision a server from 90 days to 3 hours By implementing a cloud strategy, we have achieved USD 9 million in net savings to date. From 2009 to 2015, we anticipate total program NPV of USD 20 million
Enhance server performance through software optimization	<ul style="list-style-type: none"> Increased Design job throughput up to 49 percent Delivered USD 55 million from 2010 to 2013 from NUMA-Booster USD 12.42 million cost avoidance in 2012 with fast-swap drives based on Intel® Solid-State Drives
Storage	
BEST PRACTICE	BUSINESS VALUE
Refresh and modernize storage using the latest generations of Intel Xeon processors	<ul style="list-style-type: none"> Take advantage of new technology to increase storage capacity, quality, velocity, and efficiency at a lower cost More than twice the I/O throughput than older systems Reduced our data center storage hardware footprint by more than 50 percent in 2011-2012 Reduced backup infrastructure cost due to greater sharing of resources
Right-size storage solutions using a tiered model ⁱⁱⁱ	<ul style="list-style-type: none"> Provide storage resources based on business needs: performance, reliability, capacity, and cost Better management of storage costs while still enabling easy access to necessary data Transition to scale-out storage to reduce operational complexity in tiering data
Continuously monitor and reclaim disk space consumed by aged data	<ul style="list-style-type: none"> More than USD 1 million in capital expenditure avoidance in 2011
Implement thin provisioning and deduplication for storage resources	<ul style="list-style-type: none"> Helps control costs and increase resource utilization without adversely affecting performance Increased storage effective utilization in Design from 46 percent in 2011 to 55 percent in 2012
Network	
BEST PRACTICE	BUSINESS VALUE
Upgrade the data center LAN network architecture to 10 gigabit Ethernet ^{iv}	<ul style="list-style-type: none"> Increased data center network bandwidth by 400 percent over three years, enabling us to respond faster to business needs and accommodate growth Increased the network utilization from 40 percent to 60 percent between 2010 to 2013 Eliminated spanning tree with multi-chassis link aggregation (MLAG) and Layer 3 protocol Reduced network complexity due to fewer network interface cards (NICs) and LAN ports Reduced network cost in our virtualized environment by 18 to 25 percent
Open the data center network to multiple suppliers	<ul style="list-style-type: none"> Generated more than USD 20 million in cost avoidance over three years with new network technology
Facilities	
BEST PRACTICE	BUSINESS VALUE
Increase cooling efficiency	<ul style="list-style-type: none"> Saved close to 16 million kilowatt-hours over 18 months, which is equivalent to reducing our carbon dioxide emissions by 6,800 metric tons
Use a tiered approach to redundancy, availability, and physical hardening	<ul style="list-style-type: none"> Better matching of data center redundancy and availability features to business requirements Reduced wasted power by more than 7 percent by eliminating redundant power distribution systems within a data center
Retrofit and consolidate data centers using a modular design	<ul style="list-style-type: none"> Avoid costly new construction Avoided significant capital expenditures at one data center by not equipping the entire facility with generators Quickly respond to changing data center needs with minimal effort and cost

ⁱ For more information, refer to "High-Performance Computing for Silicon Design."ⁱⁱ For more information, refer to "Migrating Mission-Critical Environments to Intel® Architecture."ⁱⁱⁱ For more information, refer to "Implementing Cloud Storage Metrics to Improve IT Efficiency and Capacity Management."^{iv} For more information, refer to "Upgrading Data Center Network Architecture to 10 Gigabit Ethernet."

PLANS FOR 2013 THROUGH 2015

Our data center strategy is a continuous improvement process—we're always striving to close the gap between current achievements and the best possible scenario.

To that end, we plan to explore the following areas and apply our MOR approach to them.

- **Embrace disruptive servers.** Deploy ultra-dense, power-optimized server nodes to reduce data center space and power consumption for computing needs.
- **Adopt standards-based storage.** Use industry-standard hardware and software for scale-up and scale-out storage, to take advantage of the latest hardware more quickly—enabling higher throughput.
- **Drive network efficiency.** Continue to drive LAN utilization toward 75 percent and pursue top-of-rack architecture to support ultra-high density data center designs. Introduce 40GbE and 100GbE where appropriate and cost effective, in order to meet network capacity demands.
- **Increase facilities efficiency.** Use techniques such as higher ambient temperature for specific data center locations to take advantage of newer equipment specifications, which will reduce cooling needs.

CONCLUSION

To provide a foundation for continuous innovation that will improve the quality, velocity, and efficiency of Intel IT's business services, we have refined our data center strategy, building on the practices established over the last decade. Our refined data center strategy has created new business value in excess of USD 184 million from 2010 to date.

Key achievements include the following:

- We developed a system software capability called NUMA-Booster, which has delivered USD 55 million in additional server capacity.
- We deployed more than 13,000 Intel SSDs as fast-swap drives, which increased server capacity by 27 percent.
- Four generations of HPC in our design computing environment created a 30x capacity increase and a 20x quality improvement.
- We adopted new storage capabilities, accelerated storage refresh, and focused on increasing utilization, generating USD 33 million in cost avoidance.
- We deployed more than 18,000 10 Gb/s network ports, generating more than USD 20 million in cost avoidance.
- An integrated server and network infrastructure provided a 39-percent reduction in hardware across the enterprise.

We have achieved these results by running Intel data centers like a factory, implementing change in a disciplined manner and applying breakthrough technologies, solutions, and processes. Transformational elements of our data center strategy include the following:

- **A focus on three primary KPIs.** These metrics enable us to measure the success of data center transformation: Meet growing customer demand (SLAs and QoS) within constrained spending targets (remaining cost-competitive) while optimally increasing infrastructure asset utilization (asset efficiency).
- **Stimulating bolder innovation by changing our investment model.** Comparing our current capabilities to a "best achievable model" encourages us to strive for innovation that will transform our infrastructure at a faster rate than if we sought only incremental change.
- **New unit-costing financial model.** This model enables us to better assess our data center TCO based on the business capabilities our infrastructure is supporting. The model measures the cost of a unit of service output and enables us to compare investments and make informed trade-off decisions across business functions—thereby maximizing ROI and business value.

The data center transformation strategy, (Figure 16) is a key for Intel IT to stay competitive, compared to public cloud services.

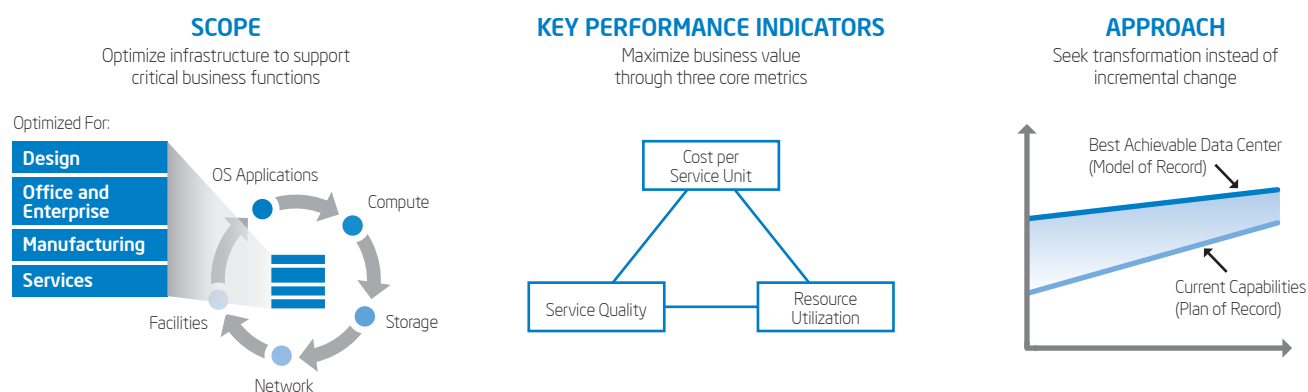


Figure 16. Maximizing the business value of Intel's data center infrastructure requires continued business-driven innovation in the areas of compute, storage, network, and facilities. It also requires us to adopt new metrics and to dramatically shift our investment strategy.



Work smarter

At Insight, we'll help you solve challenges and improve performance with intelligent technology solutions.

Learn more

