IT@Intel White Paper

Intel IT
IT Best Practices
Data Centers and IT Innovation
December 2011

Intel IT's Data Center Strategy for Business Transformation

We believe our new approach to data center costing and investment evaluation will stimulate a bolder approach to continuous innovation and will improve the quality, velocity, and efficiency of Intel IT's business services, creating a sustained competitive advantage for Intel's business.

Shesha Krishnapura
Senior Principal Engineer, Intel IT

Vipul Lal
Senior Principal Engineer, Intel IT

Raju Nallapa
Principal Engineer, Intel IT

Isaac Priestley
Strategic Financial Analyst, Intel IT

John Pereira
Director, Data Centers and Hosting, Intel IT

Sanjay Rungta
Senior Principal Engineer, Intel IT

Ananth Sankaranarayanan
Technical Program Manager, Intel IT

Executive Overview

To better meet Intel's business requirements while providing our internal customers with optimal data centers and innovative business services, Intel IT is updating our data center strategy. Building on previous investments and techniques, our refined strategy will continue to create new business value while improving data center efficiency.

We have realized hundreds of millions of U.S. dollars 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—Design, Office, Manufacturing, Enterprise, and Services (DOMES).

Going forward, we have enhanced our strategy to include several new elements:

- Changing our investment decision model, which now compares current data center capabilities to a "best achievable model." This model seeks to remove the conventional improvement mindset, which only focuses on incremental improvements. Instead, the model will help us transform our capabilities by identifying further groundbreaking innovations—like those already used to implement our private cloud and our highly efficient silicon design computing grid.
- Implementing holistic key performance indicators and associated goals for cost per service unit, quality of service, and effective resource utilization.
- Using a new unit-costing financial model that enables us to benchmark ourselves and prioritize our investments.

We believe our new approach to data center costing and investment evaluation will stimulate a bolder approach to continuous innovation and will improve the quality, velocity, and efficiency of Intel IT's business services, creating a sustained competitive advantage for Intel's business.

Contents

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 90 data centers housing approximately 75,000 servers that underpin the computing needs of more than 90,000 employees.[1] To support the business needs of Intel's critical business functions—Design, Office, Manufacturing, Enterprise, and Services (DOMES)—while operating our world-class 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 carbon dioxide footprint to improve IT sustainability.

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

- 30 to 45 percent annual growth in compute capacity requirements
- 35 to 40 percent annual growth in storage needs
- 35 to 40 percent annual growth in network bandwidth[2]

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.

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.

[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."

[2] Internet bandwidth as a representation of overall Intel network traffic.

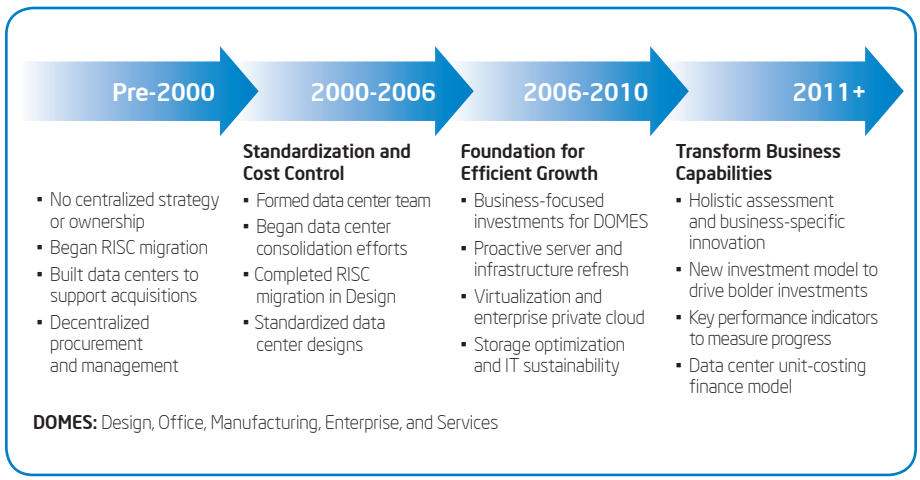

Figure 1. Intel's data center strategy is a continuous improvement process.

DESIGN

Design engineers run 20 to 30 million compute-intensive batch design 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, and a specialized algorithm that increases the performance of the heaviest Design workloads.[3] Together, these investments enabled Design engineers to run 25 percent more jobs without adding more 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 network-attached storage (NAS), along with parallel storage—which is highly scalable in performance—for our HPC needs. We use storage area networks (SANs) for specific storage needs such as databases.

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. That number is greater than 60 percent today, and we remain on track to virtualize 75 percent of the 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.

[3] 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 a high-performance computing (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 2009.

REFINING OUR STRATEGY

As the pace of Intel's business accelerates, we must continually refine our data center strategy to deliver world-class capabilities in a cost-effective manner. Our refined strategy includes a new investment model that helps us determine which investments will have the greatest business value, key performance indicators that help us measure the success of our investments, and a new unit-costing model that helps us better understand the true cost of providing IT services to each business function.

These new elements of our data center strategy, along with a continued focus on meeting business needs, will help build on the success we have already achieved through our data center initiatives over the last decade. These successes include significant data center consolidation and dramatic IT cost efficiencies. The refined strategy will enable us to support the future growth of Intel's customers, products, and acquisitions, as well as enhance the quality, velocity, and efficiency of the services IT offers to Intel business groups.

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 (POR). 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).

Enhancing Design Job Throughput through Software Optimization

We developed a system software capability called NUMA-Booster, which performs more optimal scheduling of our workloads compared to default OS scheduling on the latest multi-core non-uniform memory access (NUMA) systems. Using the NUMA-Booster algorithm increased job throughput of our most critical Design workloads by 17 percent and generated USD 20.4 million net present value (NPV) in savings. In addition, the algorithm earned a ranking in IDG's InfoWorld Green 15 Awards.

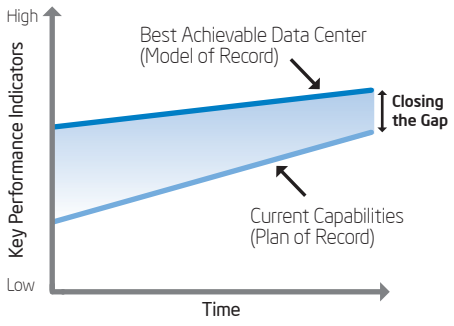

Figure 2. Our new data center investment model encourages innovation and provides significant business results.

In contrast, the new investment model, called model of record (MOR) internally, 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.

Using 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.

APPLYING THE NEW INVESTMENT MODEL

We have begun applying our new investment model to identify actionable gaps between the best achievable performance and our current plan. The model has identified three areas of investment with potential for high ROI for one or more business functions:

- Further data center consolidation.
- Reducing unit-cost for the Design environment.
- Extending our use of blade servers in the Office, Enterprise, and Services environments.

Data Center Consolidation

We used our new investment model to look closely at how many data centers we have currently and how many we should have ideally.

The new investment model identified opportunities to reduce Intel's data centers by as much as 35 percent. Since we identified this gap, we are developing a plan with positive ROI to close the gap through efforts such as:

- Closing, retrofitting, or reclassifying data centers and improving inefficiencies.
- Co-locating local infrastructure with Design and Manufacturing data centers, or providing services from a server closet.
- Managing local infrastructure sites remotely.
- Improving facility power efficiency through strategic investments.

Reduced Unit-Cost for the Design Environment

Our investment model has shown that we can improve Design unit cost by 29 percent through continuing investments in four-year server refresh, storage optimization, and deployment of a customized software algorithm to optimize grid performance. We are basing this projection on the 17-percent design job throughput improvement enabled by the NUMA-Booster algorithm (see sidebar), newer servers that offer more meaningful indicator of performance per system (MIPS) for the same cost compared to the prior year, storage optimizations that have reduced costs by USD 1.8 million, and improvements in utilization.

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 system and component levels. It is capable of delivering on-demand reports to users with a mere five-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.

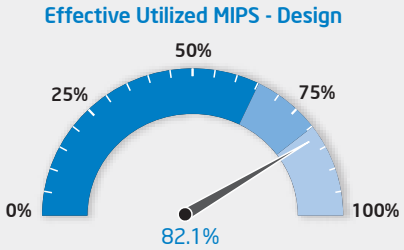

The Intel IT Data Center Dashboard provides a holistic view of data center resources to help us track our key performance indicators (KPIs) and identify opportunities for optimization and improvement.

Reduced TCO from Blade Server Technology

Our new investment model has shown us that moving from rack-mount servers to blades can reduce TCO in our cloud computing environment by about 29 percent through reduced port, network, and cable costs. For example, a group of 16 blade servers compared to 16 rack-mount servers requires only eight Ethernet interfaces instead of 128, and only four Fibre Channel interfaces instead of 32.

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

Defining Key Performance Indicators and Goals

The second major change we have made to our data center strategy is to define three new key performance indicators (KPIs) that enable us to measure the effectiveness of data center investments: cost per service unit, quality of service, and effective utilization. Because the service output for each business function is different, we do not evaluate all business functions in the same way. Our data center investment decisions seek to balance and meet all business requirements while optimizing the KPIs.

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.

QUALITY OF SERVICE

We use a tiered approach to service-level agreements (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 at a higher performance-to-SLA than for second-tier applications, which are less critical. The end goal and true measure of IT quality of service is zero business impacts from IT.

Table 1. Service Unit for Each Business Function

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

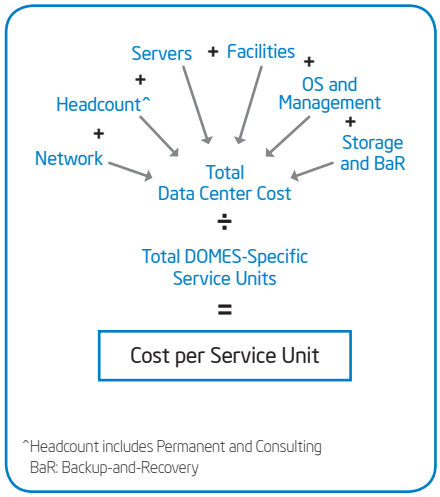

Figure 3. We arrive at a data center unit cost by considering all categories of cost and dividing by the number of units for that environment, such as electronic design automation-meaningful indicator of performance per system (EDA-MIPS) in Design and OS instances in Office, Enterprise, and Services.

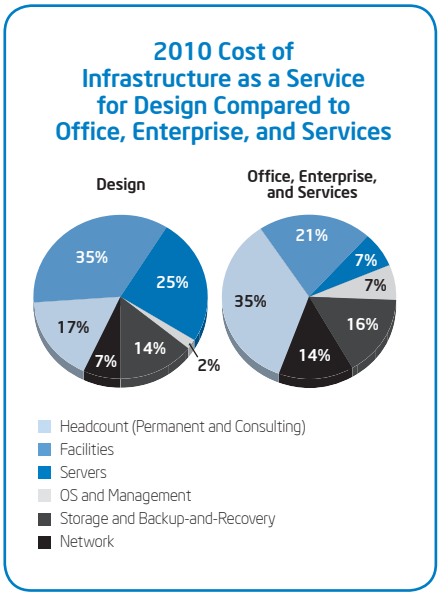

Figure 4. Knowing the total unit cost, as well the individual cost category figures for each business environment, enables us to better choose IT investments that will have the greatest impact on lowering costs.

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.

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, head-count, facilities, servers, OS and management, and storage and backup and recovery (BaR). 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.

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—giving us the tools to make decisions on 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.

DATA CENTER BEST PRACTICES

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.

These investments are spread across our entire infrastructure stack—compute, storage, networking, and facilities. These best practices, and the business value they have generated, are described in Table 2.

Table 2. Intel IT Data Center Best Practices

Compute (Servers)	
Best Practice	**Business Value**
Regularly refresh servers using the latest generations of Intel® Xeon® processor.	▪ Virtualization ratios of up to 20:1. ▪ Reduced number of servers in Design by 40 percent between 2005 and 2010. ▪ Reduced Design environment energy consumption by 10 percent annually since 2008. ▪ Greater than 5x increase in performance between 2005 and 2010.
Deploy high-performance computing.[i]	▪ 3x throughput improvement for tapeout jobs. ▪ 10x improvement in compute environment reliability. ▪ Saved USD 44.72 million net present value (NPV) from 2006 to 2010.
Migrate applications from RISC to Intel® architecture.[ii]	▪ Enabled significant savings and IT efficiencies. ▪ Allowed us to realize the benefits of industry-standard OSs and hardware.
Adopt virtualization and cloud computing.	▪ 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 three 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.	▪ Increased Design job throughput by 17 percent. ▪ Projected USD 20.4 million NPV from 2010 to 2013.
Storage	
Best Practice	**Business Value**
Refresh and modernize storage using the latest generations of Intel Xeon processor.	▪ 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. ▪ Will reduce 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]	▪ Provide storage resources based on business needs: performance, reliability, capacity, and cost. ▪ Better manage storage costs while still enabling easy access to necessary data.
Continuously monitor and reclaim disk space consumed by aged data.	▪ More than a USD 1 million in capital expenditure avoidance in 2011.
Implement thin provisioning for storage resources.	▪ Help control costs and increase resource utilization without adversely affecting performance.
Network	
Best Practice	**Business Value**
Upgrade the data center LAN network architecture to 10 gigabit Ethernet.[iv]	▪ Respond faster to business needs. ▪ Accommodate current growth. ▪ Meet increasing network demand in the future. ▪ 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.	▪ Lower component costs.
Facilities	
Best Practice	**Business Value**
Increase cooling efficiency.	▪ 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.	▪ 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.	▪ 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."

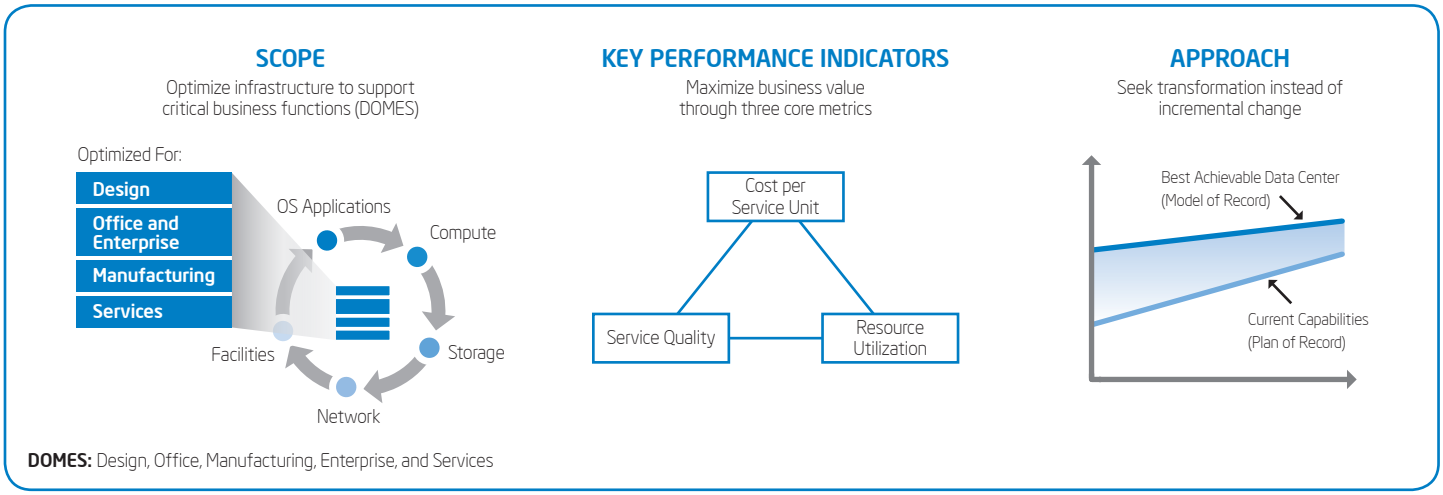

Figure 5. Maximizing the business value of Intel's data center infrastructure requires continued business-driven innovation in the areas of compute, storage, network, and facilities; the adoption of new metrics; and a dramatic shift in investment strategy.

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.

New elements of our data center strategy include:

- **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 only sought incremental change.
- **A focus on three primary KPIs.** Our goals include achieving an annual 10-percent decrease in cost per service unit, meeting business-specific performance-to-SLA goals at defined tiering levels, and striving to achieve an 80-percent effective utilization of our data center assets.
- **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.

Figure 5 summarizes our refined data center strategy, which has the overall goal of operating a world-class data center infrastructure to deliver a competitive advantage for Intel's business.

Combining the new elements of our data center strategy with the best practices of the past will enable us to continue to build world-class data centers to support the growing and changing needs Intel's business in a cost-effective manner.

For more information on Intel IT best practices, visit www.intel.com/it.

CONTRIBUTORS

David Aires
Bradley Ellison
Bill Guyon

ACRONYMS

BaR	backup and recovery
BI	business intelligence
DOMES	Design, Office, Manufacturing, Enterprise, and Services
EDA-MIPS	electronic design automation-meaningful indicator of performance per system
HPC	high-performance computing
KPI	key performance indicator
MIPS	meaningful indicator of performance per system
MOR	model of record
MTTR	mean time to repair
NAS	network-attached storage
NIC	network interface card
NUMA	non-uniform memory access
NPV	net present value
POR	plan of record
ROI	return on investment
SAN	storage area network
SLA	service-level agreement
TCO	total cost of ownership
USD	U.S. dollar

Printed in USA Please Recycle 1211/ABC/KC/PDF 325588-001US