White Paper

September 2025

IT@Intel Data Center Strategy Leading Intel's Business Transformation

As we continue to apply breakthrough technologies and solutions while evolving our processes, we enable the acceleration of Intel's business

Intel IT Authors

Shesha Krishnapura
Intel Fellow and Intel IT CTO

Shaji Kootaal Achuthan
Senior Staff Engineer

Murty Ayyalasomayajula
Senior Staff Engineer

Vipul Lal
Senior Principal Engineer

Raju Nallapa
Senior Principal Engineer

Sanjay Rungta
Senior Principal Engineer

Ty Tang
Senior Principal Engineer

Himayun Zia
Technical Program Manager

Executive Summary

Intel IT runs Intel data center services like a factory, accomplishing change in a disciplined manner and applying breakthrough technologies, solutions, and processes. This enables us to optimally meet Intel's business requirements while providing our internal customers with effective data center infrastructure capabilities and innovative business services. Building on previous investments and techniques, our data center strategy has generated savings exceeding USD 11.41 billion from 2010 to 2024.

We are constantly enhancing our data center strategy to continue our data center transformation. Using disruptive server, storage, network, infrastructure software and data center facility technologies can lead to unprecedented quality-of-service levels and reduction in total cost of ownership for all applications. They also enable us to continue to improve IT operational efficiency and be environmentally responsible.

$11.41 Billion

2010-2024 Cumulative Savings

44% Savings with a **Disaggregated Server Design Refresh** compared to a full-acquisition

400% Increase in **Data Transfer Rates** between sites through international WAN links

1-Day Deployment using our **Process Transformation** for new physical server development

631x Increase in our **HPC Environment** with 322x improvement in quality

Contents

Background

Intel IT operates 53 data center modules at 15 data center sites. These sites have a total capacity of 133 megawatts, housing more than 464,000 servers that underpin the computing needs of more than 100,000 employees.[1] To support the business needs of Intel's critical business functions—Design, Office, Manufacturing and Enterprise (DOME)—while operating our data centers as efficiently as possible, Intel IT has engaged in a multiyear evolution of our data center strategy, as outlined below.

[1] Number of employees as of January 1, 2025. Number of data centers and servers as of August 2025. 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."

Intel IT Data Center Strategy Evolution

Intel's Data Center Strategy is a Continuous Improvement Process

Pre-2000

Ad-hoc/Unstructured Growth

- No centralized strategy or ownership
- Built data centers to support acquisitions
- Decentralized procurement and management
- RISC migration to Intel® architecture begins

2000-2006

Standardization and Cost Control

- Formed data center team
- Completed RISC to Intel Architecture migration in Design
- Standardized data center designs
- Began data center consolidation efforts

2006-2010

Foundation for Efficient Growth

- Business-focused investments for DOME
- Proactive server and infrastructure refresh
- Virtualization and enterprise private cloud
- Storage optimization and IT sustainability

2010-2013

Transform Business Capabilities

- TCO assessment of Infrastructure as a Service
- Introduction of data center MOR
- Unit-cost model to plan improvement targets and benchmark
- Pulse dashboard for comprehensive state of Infrastructure-as-a-Service capacity and utilization

2013-2023

Focus on Resource and Energy Efficiency

- Breakthrough disaggregated server architecture innovation
- Centralized batch computing capacity in two mega-hubs
- Combined high-frequency servers and optimal workloads for platform pairings
- Centralized management of servers and resource
- Converted older wafer fabrication facilities into data centers
- Custom rack design to optimize space, compute, and power density
- Environmental sustainability—either free-air cooling or evaporative cooling-tower water to condition the data centers
- State-of-the-art electrical density and distribution system

2024 and Beyond

Intelligent Data Center Operations

- Real-time control to optimally utilize data center power, cooling, and density while achieving maximum performance from Intel® Xeon® processor-based servers
- Implemented machine learning for data center workload management, including resource prediction for incoming jobs, smart job placement during scheduling, dynamic runtime resource adjustment, idle job detection, and removal
- Job profile-based auto-tuning for shared storage IOPS reduction and alerting of potential workload performance bottlenecks

2010-2024: $11.41 Billion Cumulative Cost Savings

Meeting Compute Environment Challenges

In the past, we focused our data center investments on improving IT infrastructure 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 CO2 footprint to improve IT sustainability.

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

- 30 to 40% in compute capacity requirements
- 35 to 50% in storage needs
- 30 to 40% in demand for network capacity

We needed to address these challenges without negatively impacting service delivery. We developed and continue to rely on many established industry best practices in all areas of our data center investment portfolio. These areas include servers, storage, networking, and facility innovation. Since 2010, these techniques, described in detail later, have enabled us to realize USD 11.41 billion in cost savings while supporting significant growth.

Disaggregated server at scale

Breakthrough Disaggregated Server Architecture

By decoupling the CPU/DRAM and NIC/Drives modules from other server components, we can independently refresh servers' CPU and memory without replacing other server components. This results in faster technology adoption, which in turn puts new technology at our Design engineers' fingertips.

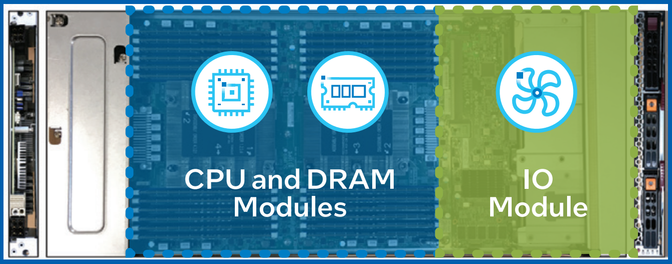

Learn More:

- Disaggregated Servers Drive Data Center Efficiency and Innovation Paper
- Disaggregated Servers Blog
- Mission - Green Computing Video

Aligning Data Center Investments with Business Needs

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

Design

Design engineers run more than 308 million compute-intensive batch jobs every week. Each job can take a few seconds to several days 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. These approaches include high-performance computing (HPC), grid computing and clustered local workstation computing.[2] We used SSDs as fast local data cache drives, single-socket servers, and a specialized algorithm that increases the performance of the heaviest Design workloads. Together, these investments enable Design engineers to run up to 49% more jobs on the same compute capacity. This 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 in combination with caching on local storage and automatic block tiering to on-premises low-cost object storage for our HPC needs. We use storage area networks 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 info on HPC at Intel, refer to "High-Performance Computing for Silicon Design," Intel Corp., December 2015.

Manufacturing

IT systems must be available 24/7 in Intel's Manufacturing environment, so we use dedicated data centers co-located with the factories for Manufacturing. We have invested heavily over the last few years to develop a robust business continuity plan. Our plan keeps factories running even in the case of a catastrophic data center failure.

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 us to quickly deploy new platforms and applications throughout the Manufacturing environment. This helps us meet a 13-week infrastructure deployment goal 95% of the time—compared to less than 50% without using copy-exact methodology.

Office and Enterprise

To improve IT agility and the business velocity of our private enterprise cloud, we have implemented an on-demand self-service model. This model has reduced the time to provision servers from three months to on-demand provisioning. We have achieved a mature level of virtualization in our Office and Enterprise computing environment and have deployed containers technology to further improve the agility in managing infrastructure and application; software development and testing; and scalable services deliveries.

In contrast to the Design environment, in the Office and Enterprise environments we rely primarily on a storage area network, with limited network-attached storage for file-based data sharing.

Defining a Model of Record

Our transformational data center strategy involves running Intel data centers and underlying infrastructure as if they were factories, with a disciplined approach to change management. Applying breakthrough technologies, solutions and processes in an effective controlled manner can help us be an industry leader and to keep up with the accelerating pace of Intel's business.

Based on improvements each year in technologies, solutions, and processes, we use three key performance indicators (KPIs) to define a model of record (MOR) for the year. These KPIs—which are discussed in more detail in subsequent sections—include the following: best achievable quality of service (QoS) and service-level agreements (SLAs); lowest achievable unit cost; and highest achievable resource utilization.

We set investment priorities based on the KPIs to move toward the MOR goal. As shown below, each year we get closer to the MOR while at the same time balancing the KPIs.

We use five primary tactics to achieve our MOR goals:

- Embrace disruptive servers
- Adopt tiered storage
- Increase facilities efficiency
- Drive network efficiency
- Improve operational efficiency

More information is provided about each of these tactics in subsequent sections.

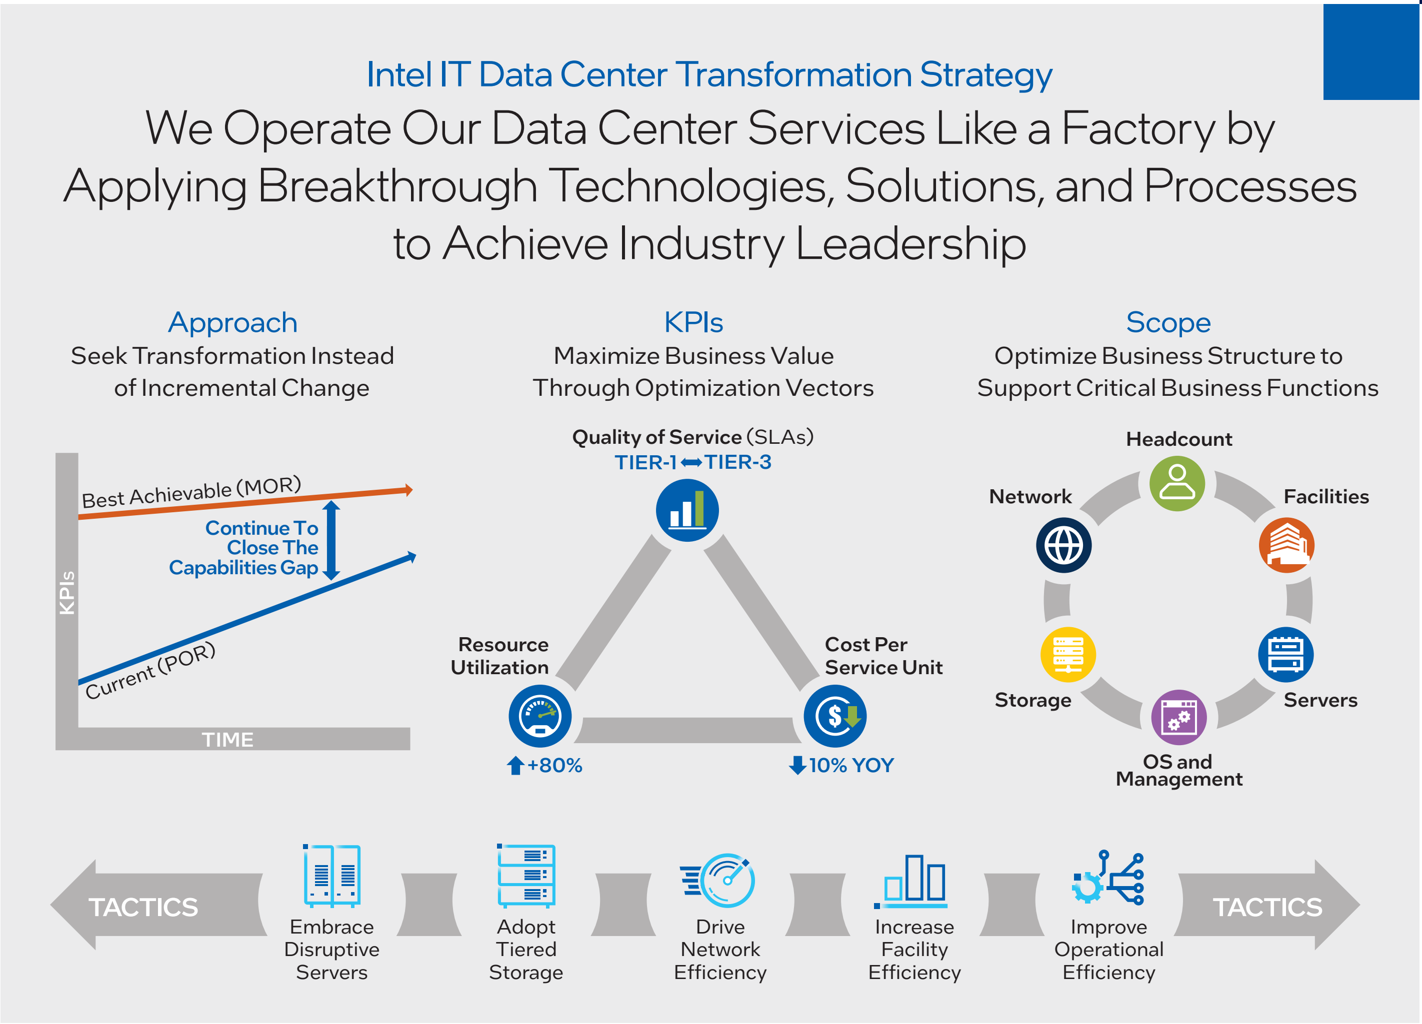

Approach

Seek transformation instead of incremental change

Achieving Economic Value

Our new data center investment model encourages innovation and provides significant business results. We have realized substantial cost savings since 2006 by proactively refreshing our infrastructure. For example, Intel® Xeon® processor-based servers have contributed significant economic value. During this time, we have delivered substantially higher computational throughput as measured by a practical electronic design automation (EDA) workload. Further cost savings result from adopting cloud computing-like technologies, 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—DOME.

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."

Disaggregated server at scale

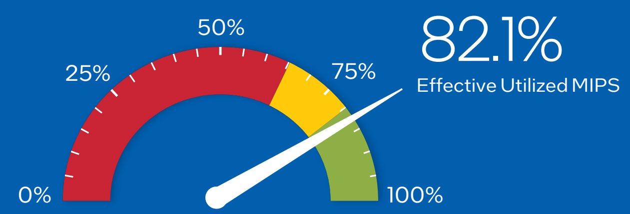

Intel IT Data Center Dashboard

To better monitor and manage our worldwide network of data centers, we developed and deployed an integrated Intel IT Data Center Dashboard. This dashboard is modeled on a dashboard used in Intel's Manufacturing environment.

This dashboard helps us monitor our KPIs by highlighting the current state and opportunities for optimization. We can thereby achieve 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 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 above shows a sample of the design environment dashboard.

KPIs

Maximize business value through optimization vectors

Defining KPIs 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 them 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 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.

Cost per Service Unit

Different business functions have a different service unit that we can measure. This unit represents the capacity we enable for our business users. The service unit for each business function is:

- **Design:** Cost per EDA-MIPS
- **Office and Enterprise:** Cost per OS instance
- **Manufacturing:** Cost per integrated factory compute environment

Our goal for this KPI is to achieve a 10% improvement in data center cost efficiency every year. This goal does not necessarily mean we will spend less each year, but 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.

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% of the time, we considered it 90% utilized. If 80% of available storage was allocated, we considered that 80% utilization.

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

Stimulating Bold Innovation through a New Investment Model

Our efforts are based 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." This model 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.

Because technology is always changing, peak performance also changes—the maximum achievable performance keeps improving through innovation. We know that resource constraints make it difficult to achieve the standard set by the new investment model. However, our HPC environment comes very close to that goal. The model enables us to identify gaps between where we are and where we would like to be. We can then identify the biggest gaps in capability to prioritize our budget allocation toward the highest value investments first.

Scope

Optimize business structure to support critical business functions

Implementing a New Unit-Cost Financial Model

We evolved our financial model from project- and component-based accounting to a more holistic unit-cost 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 return on that investment. We had 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.

Figure 1 shows the six major categories of cost to consider: headcount, facilities, servers, OS and manageability, storage and backup/recovery and network. By adding these costs and then dividing them by the total number of appropriate service units for the environment, we arrive at a cost per service unit.

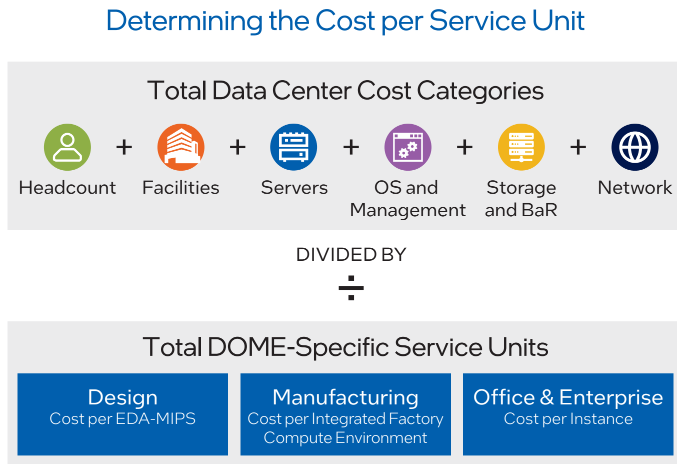

Figure 1. We arrive at a data center unit cost by considering all categories of cost and dividing by the number of units for that environment. Unit examples include EDA-MIPS in Design and OS instances in Office and Enterprise.

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. This comparison helps us identify which data centers are not performing optimally and decide whether to upgrade or consolidate them.

To show how the new unit-based costing model works, Figure 2 compares Design cost data and Office and Enterprise cost data. The network category shows a larger total cost in Office and Enterprise compared to Design. In contrast, servers are more of a cost factor in Design than they are in Office and Enterprise. 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.

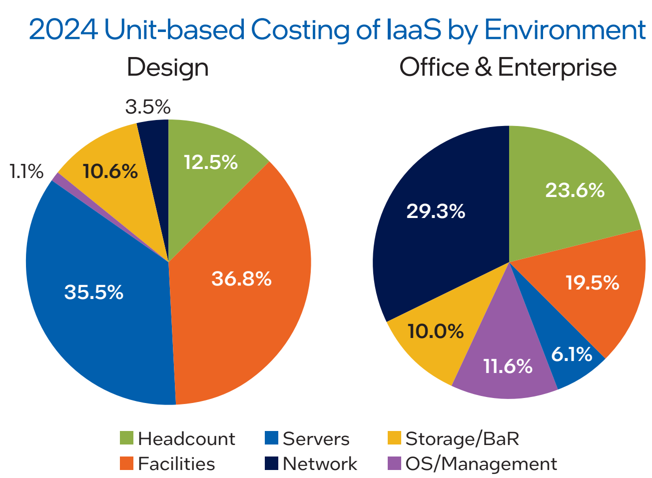

Figure 2. Knowing total unit costs and individual cost category figures for each business environment, we can better choose IT investments that lower costs the most.

Results

Building on the Past, Building for the Future

This section details some of the improvements and cost savings our data center strategy has enabled over the years, using our five primary tactics of embracing disruptive servers, adopting tiered storage, increasing facilities efficiency, driving network efficiency, and improving operational efficiency. We are building on previous successes. Therefore, some of the results shown here are cumulative; others have been achieved over the last few years as a direct result of our MOR strategy. Our refined data center strategy enables us to support the growth of Intel's customers, products, and acquisitions. It also helps to enhance the quality, velocity, and efficiency of the services we offer to Intel business groups.

We have dramatically improved performance and reduced costs for our data centers:

- **Data Center-wide**
  - Smaller total data center footprint
  - Improved overall storage and network practices
  - Implemented design zones to achieve highly resilient scaling
  - Increased data center facilities efficiency
  - Global street-to-server audit helps prioritize investments
- **Design Environment**
  - Deployed disaggregated servers
  - More efficient Design compute and storage
  - Increase Design throughput and performance using NUMA-Booster, smart jobs placement and scheduling, shared storage IOPS reduction, and idle jobs detection
  - Faster Design throughput using SSDs
- **Office and Enterprise Environment**
  - More efficient Office and Enterprise compute and storage

Disaggregated Server Architecture

The First Major Server Innovation Since the Introduction of Blade Servers In 2005

Intel IT has developed a disaggregated server architecture. The architecture separates the CPU/DRAM module and the NIC/Drives module on the motherboard. Redesigning the server to be modular enables us to upgrade the CPU/DRAM module while retaining the other components that are not ready for end-of-life. These include fans, power supplies, cables, network switches, drives, add-on module/accelerator, and chassis. The disaggregated server architecture is characterized by a CPU/DRAM complex or module and a NIC/Drives module.

These modules can be refreshed independently of each other and of the rest of the server components.

We have found that the disaggregated design offers the following benefits:

- No need to replace perfectly good components.
- No need to reinstall the OS.
- Cuts refresh costs by a minimum of 44%.
- Reduces technician time spent on refresh by 77%.
- Decreases refresh materials' ship weight by 82%.

1-Socket Disaggregated Server Example

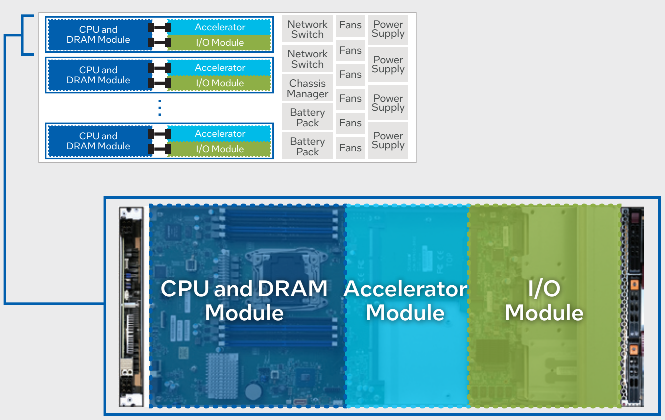

2-Socket Disaggregated Server Example

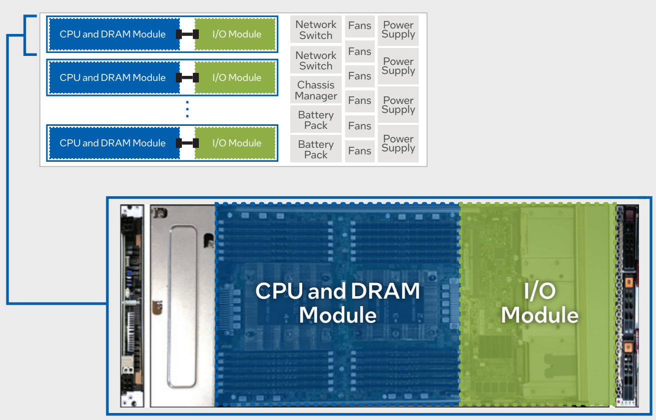

Disaggregated Server Innovation Reduces TCO and TCE

One of our leading tactics to achieve our MOR goals is to adopt disruptive server technology. To this end, we are deploying disaggregated servers throughout our data centers. It makes little sense to replace an entire light fixture when all that is needed is a more energy-efficient or powerful light bulb. Likewise, replacing an entire server is not necessary when all that is needed is a more advanced CPU and DRAM.

Our disaggregated server architecture has the potential to dramatically change how data centers around the world perform server refreshes. It will lead to significant refresh savings (see Figure 3) and the opportunity to quickly take advantage of the latest compute technology. This technology is already being used in Intel's data centers in Santa Clara, California. These data centers have the world's best power usage effectiveness (PUE) rating of 1.06.

The ability to spend less time and money on refreshing servers means Intel IT can afford to refresh faster, bringing the most advanced Intel Xeon processor-based technology into Intel's data centers. We are excited about the resulting opportunities to boost data center efficiency and more effectively power Intel's silicon design jobs. We have deployed more than 395,000 disaggregated servers so far, based on multiple generations of Intel Xeon processors.

In addition to the TCO benefits of 44% lower refresh cost over a full acquisition (rip-and-replace) refresh, reduced provisioning time of 77% and reduced shipping costs, disaggregated servers have total cost to environment (TCE) benefits of 82% reduction in material shipping weight and significantly reduced e-waste.

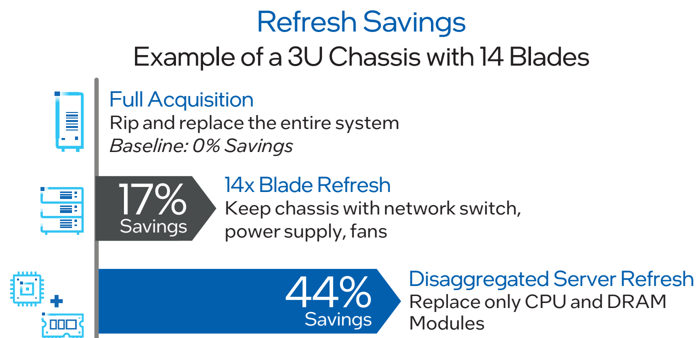

Figure 3. Refreshing the CPU/DRAM module in a disaggregated server saves at least 44% compared to a full-acquisition server refresh. Based on Intel internal testing, March 2017.

Our Data Center Evolution

Driving Up Density While Driving Down PUE

Gen 1 (1990s)
Characterized by forced chilled air from the ceiling and no hot/cold air segregation, these early data centers could accommodate 42U racks with a power consumption of 5 KW, resulting in a PUE of >2.0. Data centers that used chilled air from the row end had a PUE of ~1.40.

Gen 2 (early to mid-2000s)
With improvements such as raised-floor forced chilled air or hot/cold air segregation including chimney racks, density stayed at 42U, but power consumption delivered to the racks increased to as much as 30 KW, resulting in a lower PUE of ~1.18.

Gen 3 (2013 and beyond)
Our modern data centers use free air cooling or close-coupled evaporative cooling to achieve an industry-leading PUE of 1.06, with an extreme rack density of 60U and up to 43 KW/rack. Our latest data centers now support on average 54KW/rack with the same cooling technique.

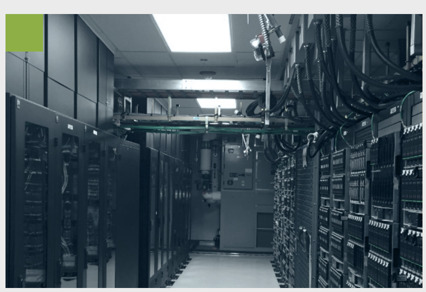

Read More: Extremely Energy-Efficient High-Density Data Centers Paper

Adopting Tiered Storage and Other Storage Techniques

A significant focus on effective utilization in our Design environment has enabled us to improve resource utilization from 46% to more than 80%. Our goal is to reach 85%.

Tiered storage is foundational to meeting our MOR goals. A four-tier approach to storage helped increase the effective utilization of storage resources, improve our performance to SLAs, and reduce TCO for Design storage. The tiers of Design storage servers are based on performance, capacity, and cost.

- Tier-1 servers have the highest performance and the least storage capacity to support tens of thousands of extremely high Network File System operations per second (NFSops) HPC jobs.
- Tier-2 servers offer medium performance but greater storage capacity; these are targeted to support thousands of intermittently high NFSops HPC jobs.
- Tier-3 servers provide lower performance but emphasize capacity.
- Tier-4 servers have the highest capacity but are used for low-frequency access and read-only archived data.
- We have initiated work to automatically tier unused blocks from these higher tiers to an on-premises object storage solution.

Our strategy has been updated to account for the computational scale of the site. This helps us to determine the appropriate performance level required for each tier and improves our ability to meet quality, SLA, and cost targets. Our automated systems monitor file server responsiveness and use that information to regulate the jobs through suspension and ramp controls. At the same time, the automated systems generate and analyze file access patterns to determine which jobs, users and files are experiencing the highest access rates. We selectively use storage QoS to isolate and mitigate the impact of very-high-IOPS workloads.

We have applied several other storage techniques to further enhance storage efficiency and reduce costs including scale-out storage, refresh cycles for storage and data reduction.

Scale-out Storage

We have executed 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. It also increases the effective utilization of space freed by using storage-efficiency technologies such as deduplication, compression, and compaction. We are performing storage scaling on-demand for read-only storage areas, which require extremely high access rates. We use mount options to increase attribute caching and avoid wasteful locking options on read-only areas. This reduces the storage load by more than 50% and improves job throughput. We have also enabled high-performance shared scratch spaces to meet the demand from our hyperscale EDA compute environment. As we march towards significantly higher compute scale, where the impact of storage overload is becoming more costly, we are shifting our bias towards achieving higher resiliency. This is achieved through increased redundancy and moderation of our storage capacity utilization targets.

Storage Refresh Cycle

To improve performance and reduce costs, we implemented an efficiency-based refresh cycle. This enables us to take advantage of storage servers with better performance and more efficient energy use. This approach has reduced 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 also reduce our data center footprint. It also helps us deliver better performance for our customers at a similar or lower cost per TB. Over the last few years, our refresh cycle has enabled us to shift from tape-based backup to disk-based backup with newer technology and architecture.

14 disaggregated servers in a 6U blade chassis with integrated network switch

This shift has made business continuity and rapid recovery from disaster a reality while reducing the backup cost and enhancing the SLA. We are also using this transition to further reduce our backup footprint. Our approach is to avoid backing up data for which it is more cost effective to regenerate it than to recover it from backup.

Data Reduction

The introduction of new storage to support company growth and our commitment to timely 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, compaction, and compression on our primary and backup storage servers. They have freed more than 220 PB of capacity, which we are making available for our users.

We continue to work closely with our internal design teams to achieve the following goals:

- Optimize their design flows to reduce the growth rate of their data and IOPS requirements.
- Dynamically adjust the allocations based on usage.
- Over-allocate capacity.

We have historically used efficient scanning algorithms to determine the age of files and then used that data to right-tier entire areas or subdirectories. We are now using block-level transparent data tiering to tier aged data to object storage. We combine the aging information with I/O activity to make more intelligent decisions to remove unused data within three to six months.

Increasing Facilities Efficiency

We used our new investment model to evaluate the number of data centers we currently have and the number we should have. The new investment model identified opportunities to reduce the number of data centers using techniques such as the following:

- Closing, retrofitting, or reclassifying data centers and improving efficiency.
- 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.

We have targeted 32 inefficient data centers since 2011. Our efforts have eliminated 66,375 square feet and converted 23,609 square feet of data center space to low-cost infrastructure rooms. This has saved Intel USD 25.45 million annually.

Figure 4 shows how we have consolidated our data center facilities from 2003-2025. We have reduced the total square footage by up to 18% and reduced the number of data centers from 152 to 53. Simultaneously, we increased our data center compute capacity and commissioned power by up to 166% from 50 MW to 133 MW over the last 13 years. From 2012-2024, we have saved over 1.91 billion KW hours compared to industry-standard data centers.

To reliably meet additional HPC scale needed for Intel Products and Intel Foundry, we are leveraging fuel cell technology to power our next Data Center expansion. The additional 31 MW capacity expands an existing 6 MW fuel cell installation already deployed at our Data Center hub since 2014.

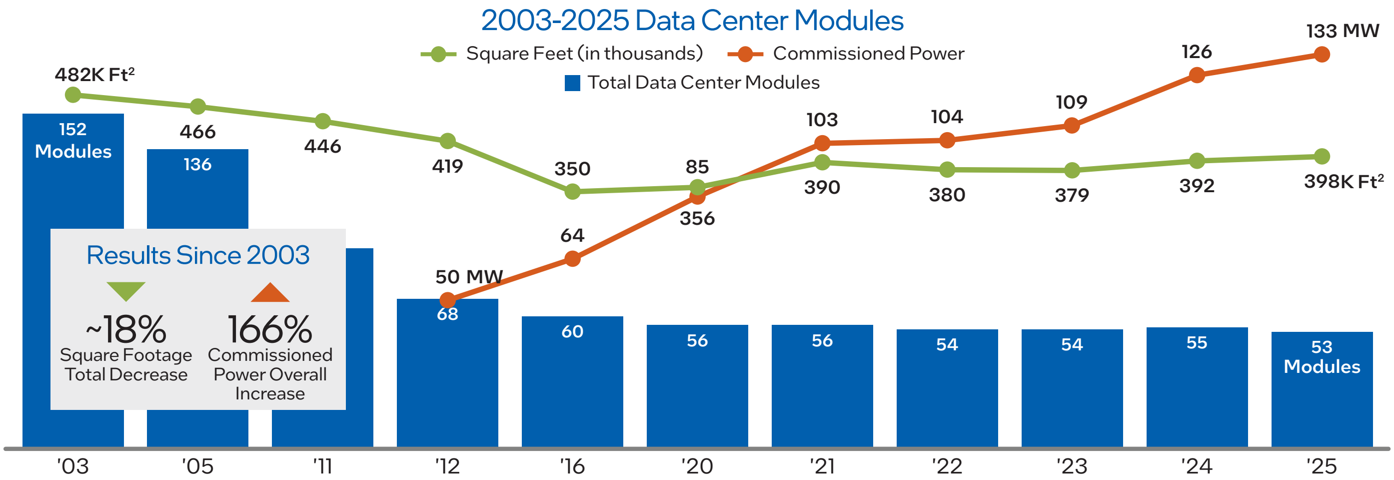

Figure 4. Innovative data center designs have enabled us to decrease data center square footage while increasing power density and capacity.

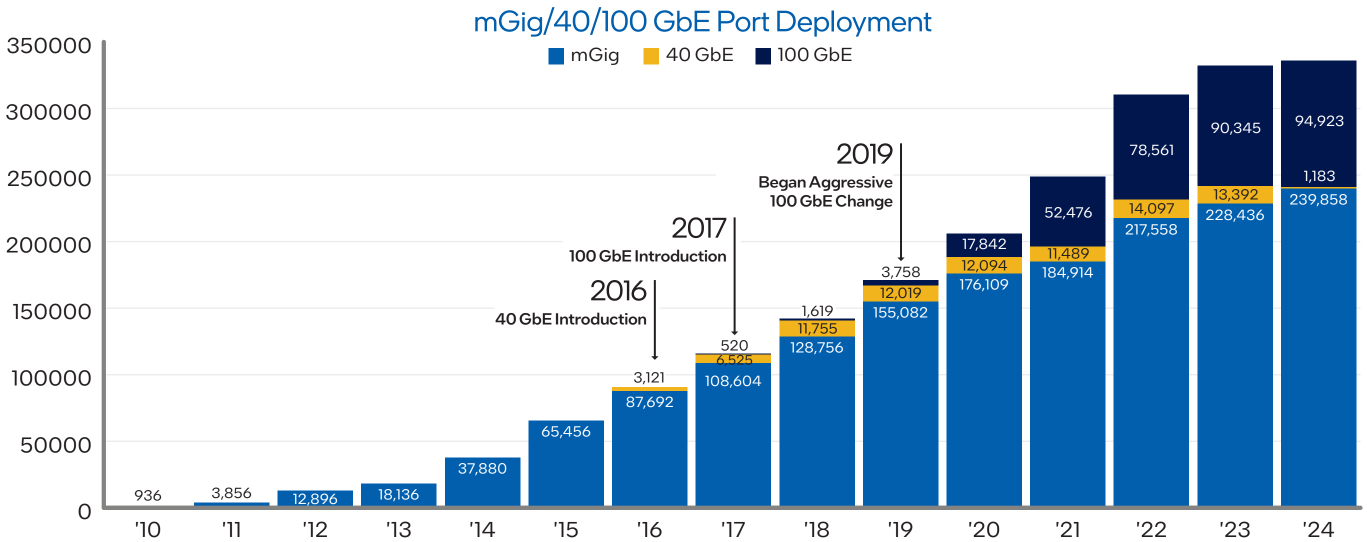

Figure 5. Implementing mGig/40/100 GbE data center fabric design accommodates current capacity growth.

Driving Network Efficiency

Data center growth is continually placing greater demands on Intel's network. In response, in 2010 Intel IT began to convert our data center network architecture to 10 GbE connections. Around 2015, we introduced 40 GbE to meet the inter-switch link capacity demand. In 2019, we started making 100 GbE pervasive within our data centers to keep up with the demand (see Figure 5).

To meet today's scale and capacity demand, we are now migrating the data center architecture to a leaf-spine architecture. We are also transitioning our switch interconnects to 100 GbE and multi-100 GbE. Our new 100 GbE data center fabric design accommodates our current annual network capacity growth of more than 30%.

In 2024, we have increased our 100 GbE capacity from 90,345 ports to 94,923 ports (see Figure 6). All the switch interconnects are being migrated to 100 GbE going forward. However, 1/10 GbE will continue to be a key part of the data center technology. We currently have deployed about 239,858 mGig ports. Our 40 GbE capacity decreased significantly in 2024 as we continue to migrate to 100 GbE network. Overall, we migrated our data center in 2024 from 71% to 86% on the Software-Defined Network (SDN) with 100 GbE fabric.

In addition to increasing the network capacity, we have also increased the effective utilization of mGig network ports (see Figure 7) over the last 15 years from 40% to 66% (1.65x increase). Higher utilization means we do not have to purchase additional ports to meet network capacity demand growth.

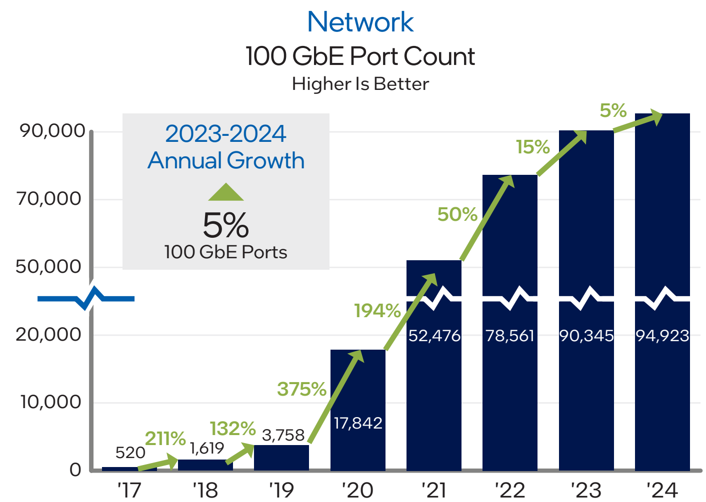

Figure 6. 100 GbE port count growth.

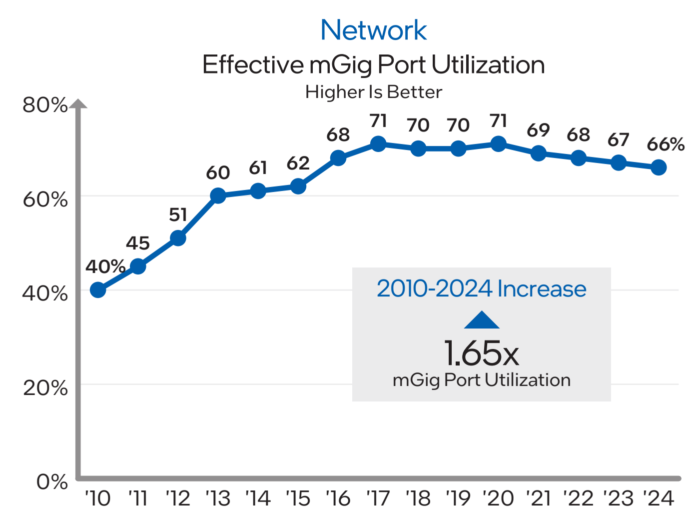

Figure 7. Effective utilization of network ports has increased by up to 1.65x between 2010-2024.

We are also focusing on improving data center stability. In the past, we used a large installation of layer 2-based technology. We have migrated to a layer 3-based network. This new architecture enables us to use all available bandwidth on primary and secondary paths at the same time. Therefore, we can 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. Using layer 3-based, scalable architecture within Intel's data center lets us plan for scale and resiliency. Also, we are using other technologies such as overlay, multi-chassis link aggregation and tunneling to extend layer 2 across data centers, over the layer 3 topology.

Due to the scale of the data center and new landings, we made zero-touch provisioning and automation a key part of the new architecture. With the new simplified modular design, each key building block has been converted into a module of the automation system. This approach allows us to provision the network within minutes with minimum effort. In addition, we can maintain consistency across the network and investigate anomalies.

We tend to adopt higher-speed network technology almost as soon as it is available in the market. We started adoption of 40 GbE in data centers in 2015 and adoption of 100 GbE technology in 2017, to keep pace with network demand. We will introduce 400 GbE capability in our data centers in 2025 to meet the future demands.

In 2015, we also made two key architecture changes within Design data centers. We reduced the oversubscription through the infrastructure and shifted from chassis-based switches to fixed form factor switches for better cost and upgrade efficiency.

With this move, we reduced the oversubscription from 8:1 to 6:1 on the compute side and 8:1 to 3:1 on the file server side. Over the same period, we transitioned 70% of our Design data centers to use fixed form-factor switches using a modular design. Now with the new leaf-spine architecture, we have maintained the same level of over-subscription ratios even though the file servers are transitioning to 100 GbE. This is possible by using 8x100 GbE interconnect links and 16x 100 GbE spine-to-universal spine links.

Achieving 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. Average annual growth rate of compute capacity demand over the last 15 years is up to 30%, while storage has grown annually at up to 34% (see Figure 8).

We expect the number of cores to continue to increase. We plan 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.

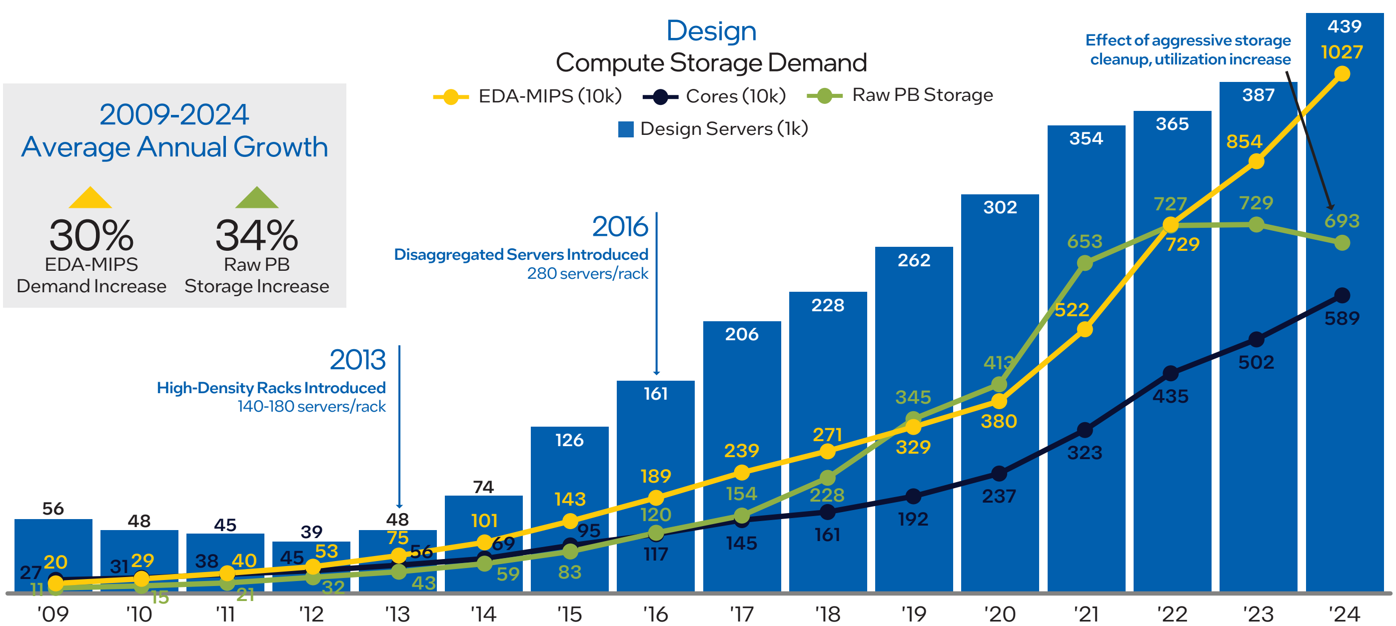

Figure 8. Despite continuing growth in compute and storage demand, our Design data centers are using powerful Intel technology to meet demand.

7th Generation of HPC

Designing Intel® microprocessors is compute intensive. Tapeout is the last stage 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 9, our HPC solution has enabled an up to 631x growth in tapeout compute capacity from 2005 to 2024. We are now using the 7th generation of our HPC solution and will continue to develop new HPC generations as Intel process technology advances. The figure also shows our commitment to quality. Through a disciplined approach to change management (running our data centers as if they are factories), we have reduced the number of compute issues that impact tapeout by 322x.

Intel Tapeout Computing Metrics

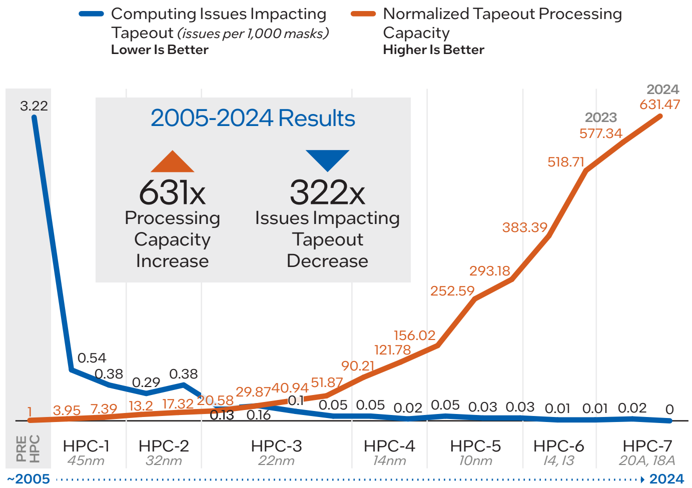

Figure 9. Our HPC solution, combined with disciplined change management, has steadily increased compute capacity and improved QoS.

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. This feature automatically and transparently intercepts all Design workloads running on two-socket batch servers and performs workload scheduling better than the default OS scheduling capability. Our tests have shown an average 17% improvement in Design performance on these two-socket servers. We are also deploying large-scale single-socket servers when possible. These servers do not need the NUMA-Booster feature and can further increase overall Design performance.

Increased Design Throughput Using SSDs as Fast Local Data Cache Drives

Intel silicon chip Design engineers 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.

Intel IT has deployed over 40 PB of SSD storage in over 20,000 servers as fast local data cache drives. This approach improves workload performance due to reduced network traffic and storage demand.

Optimizing Servers to Meet Compute Demand

Intel silicon design is continually increasing in complexity. To achieve concomitant faster time to market improvements, Intel IT provides a global framework for parallel hardware and software design of numerous System on a Chip platforms and IP blocks.

Matching single-socket servers and highly scalable server configurations in our data centers yields 25 to 30% faster product design and architecture validation processes. We use a global scheduling mechanism that pools compute capacity of over 427,000 servers at multiple sites around the world. In this way, our design hub provides scalable capacity and delivers optimal memory and compute capability in a shorter amount of time.

Since the first disaggregated server design in 2016, we have continued to evolve the concept. We currently have deployed more than 395,000 disaggregated servers, using 16 different blade designs including both single-socket and two-socket servers. We use the Intel Xeon processor E family, Intel Xeon processor W family and Intel Xeon Scalable processors. The various models are targeted to meet specific workload requirements, such as different memory capacity, throughput or number of performance cores, high bandwidth, high IOPS storage needs or the ability to add in accelerator cards on demand.

Intelligent Data Center Operations

We are evolving our data center operation with predictive analytics and intelligent algorithms to enhance overall resource efficiency, dynamic workload management capacities, services delivery, and operation costs. Improvements being implemented include:

- Real-time control to optimally utilize data center power, cooling, and density while achieving maximum performance from Intel® Xeon® processor-based servers.
- Data center workload management including resource prediction for incoming jobs, smart job placement during scheduling, dynamic runtime resource adjustment, idle job detection, and removal.
- Job profile-based auto-tuning for shared storage IOPS reduction and alerting of potential workload performance bottlenecks.

Enhanced User Experience across the Global HPC Design Community

We were able to successfully consolidate batch activities into our global compute hubs. Further consolidation was limited by the following:

- We did not want to negatively impact user experience for interactive users across the globe.
- We needed to provide local copies of critical data rapidly over the high-latency international WAN links.

We are now able to provide a game-changing remote interactive computing user experience by using User Datagram Protocol (UDP) instead of Transmission Control Protocol (TCP) for interactive jobs over the WAN. Using UDP has provided up to 4.5x faster response for computer-aided design (CAD) modeling. We have reached the stage where our international design team members have better user experience and higher throughput when working from home with systems in the US hubs than their local data centers. We also delivered up to 9x improvement in data transfer rates across the WAN through in-depth collaboration with internal and external technology experts.[3] This collaboration optimized the TCP stack, which can take full advantage of high-speed WAN links. The interactive computing and data replication improvements were achieved within existing WAN bandwidth. Combined, these achievements enable us to provide rapid turnaround through the hubs for the model build, design synthesis, layout and tapein cycle.

Design Zones Enable Highly Resilient Scaling at the Hubs

The dramatic increase in computing scale in a shared NAS environment with tens of thousands of compute servers can overwhelm the storage server. It can also introduce significant efficiency and reliability concerns when 10,000 or more such systems share the same NFS area and expect extremely high IOPS or throughput rates. We addressed this in our mission-critical tapeout environment. This environment runs parallel workflows that span the entire compute environment. We introduced the concept of partitioning the compute in the two major hubs into smaller, self-contained sites. Each site has its own NFS storage and management infrastructure. We worked with our tapeout team to update the tools, flows and work methods, along with IT software. As a result, we were able to scale while maintaining the efficiency and improving resiliency and scalability.

We later experienced the same scaling challenges for the rest of the HPC design environment in the hub. These issues were caused by the increased sharing on a higher scale and could not be addressed cost effectively or efficiently by the storage changes alone. We built on the tapeout "sites" concept to introduce design zones into the design hub computing environment. We successfully scaled multiple zones and achieved adequate separation to provide the necessary increased scale and reliability in a cost-effective manner. This is a challenging and ongoing effort. We must contend with decades' worth of legacy interdependencies across project and business units. These interdependencies use symbolic links and shared source files, tools, and flows. We expect that the profiling work that we are doing, combined with our efforts with containers, will enable us to achieve truly independent, scalable, and resilient zones without sacrificing efficiency or the agility to respond to peak computing demands.

[3] According to internal Intel IT measurements, February 2020

Design Computing Data Center Environment Improvement

Efficiency improvement examples and cost savings from 2010 through 2024

Computing

Intel IT innovations in the Design computing data center include disaggregated server innovation (44% savings during refresh); the NUMA-Booster solution (17% higher performance); SSDs (27% higher capacity at lower cost); faster servers (35% higher performance); single-day dock-to-production deployment and procurement efficiency; and dynamic adaptive algorithms for data center efficiency (150% higher rack density and auto alerting on workload bottlenecks).

Storage

We have improved Design computing data center storage efficiency by adopting innovative technology capabilities and increasing utilization.

Network

We adopted a multi-vendor strategy for our Design computing data center network, combined with a focus on the reduction of expensive maintenance costs associated with older equipment. As we adopt 100 GbE we are focusing on Intel® Silicon Photonics-based optics because that technology has a significant cost advantage over laser-based optics.

More Efficient Office and Enterprise Compute and Storage

Like our Design environment, the compute and storage demand in our Office and Enterprise environment are also growing quickly. Nevertheless, as shown in Figure 11 on the following page, we continue to meet that demand while maintaining the number of physical servers over the last three years.

From 2009 to 2017, we achieved an approximate 19x increase in the number of virtual OS instances. We also greatly increased average VM density per physical server—from 11 VMs in 2009 to 30 VMs in 2017 due to improved server platforms. In 2018, we implemented an aggressive VM reclamation strategy that led to a reduction of about 5,400 VMs. New workloads that were more cost effective to deploy on cheaper physical platforms than on a virtualized platform led to an increase in physical server counts.

In 2019, we brought additional existing virtualized workloads, VMs and hosts into our private cloud environment for centralized management, increasing the footprint by up to 1.77x. Process improvements and enhanced automation led to additional savings, and we are now deploying performance-based VMs.

From 2020-2024, we implemented an aggressive VM leasing and termination policy, reducing VM counts by 56% in that time frame.

2010 to 2024 Summary of Results

Our strategic approach has enabled us to deliver a data center infrastructure best suited to meet our complex and ever-increasing compute needs while transforming our cost structure. By applying the innovative data center techniques listed in this paper, we have achieved unit-cost levels that are significantly lower than if we were to host our workloads using public cloud infrastructure (Figure 10). Our workloads and our ability to achieve high server utilization are particularly well suited towards private cloud investment.

Over the 15-year period from 2010-2024, we have garnered combined capital and operational savings in excess of USD 11.41 billion, which help fuel our continuous innovation cycle.

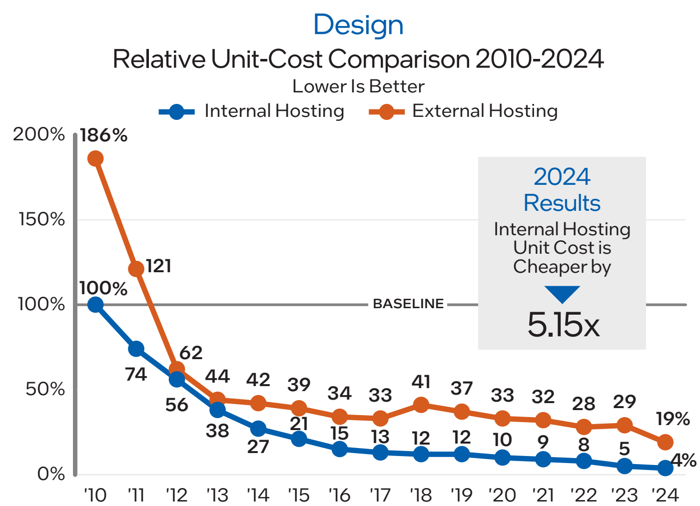

Figure 10. Unit cost including servers, storage, network, and operational costs shows private cloud hosting of our data center workloads is significantly less expensive than if we use public cloud services.

Reducing Unit Costs

The graphs on page 19 show that unit growth has continued to rise in both the Design and Office and Enterprise environments. Our investment model has enabled us to reduce unit costs in the Design environment by >96% and in the Office and Enterprise environment by 87%.

Before implementing our data center strategy, we spent only a quarter or less of our Design environment budget on servers—but the servers are what power Intel's business success. Our new investment model has enabled us to increase the spending on servers to >35%, while garnering more headcount efficiency through automation. A similar transformation is occurring in the Office and Enterprise environment, with much of the growth driven by newer analytics and security workloads.

Summary of Best Practices

Over the last decade, we have made many strategic investments and developed solutions to enable our data centers to be more efficient and to 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. Page 20 provides a summary of the best practices we have developed and the business value they have generated.

Intel IT Supercomputer: #81 in Top 500 (2015)

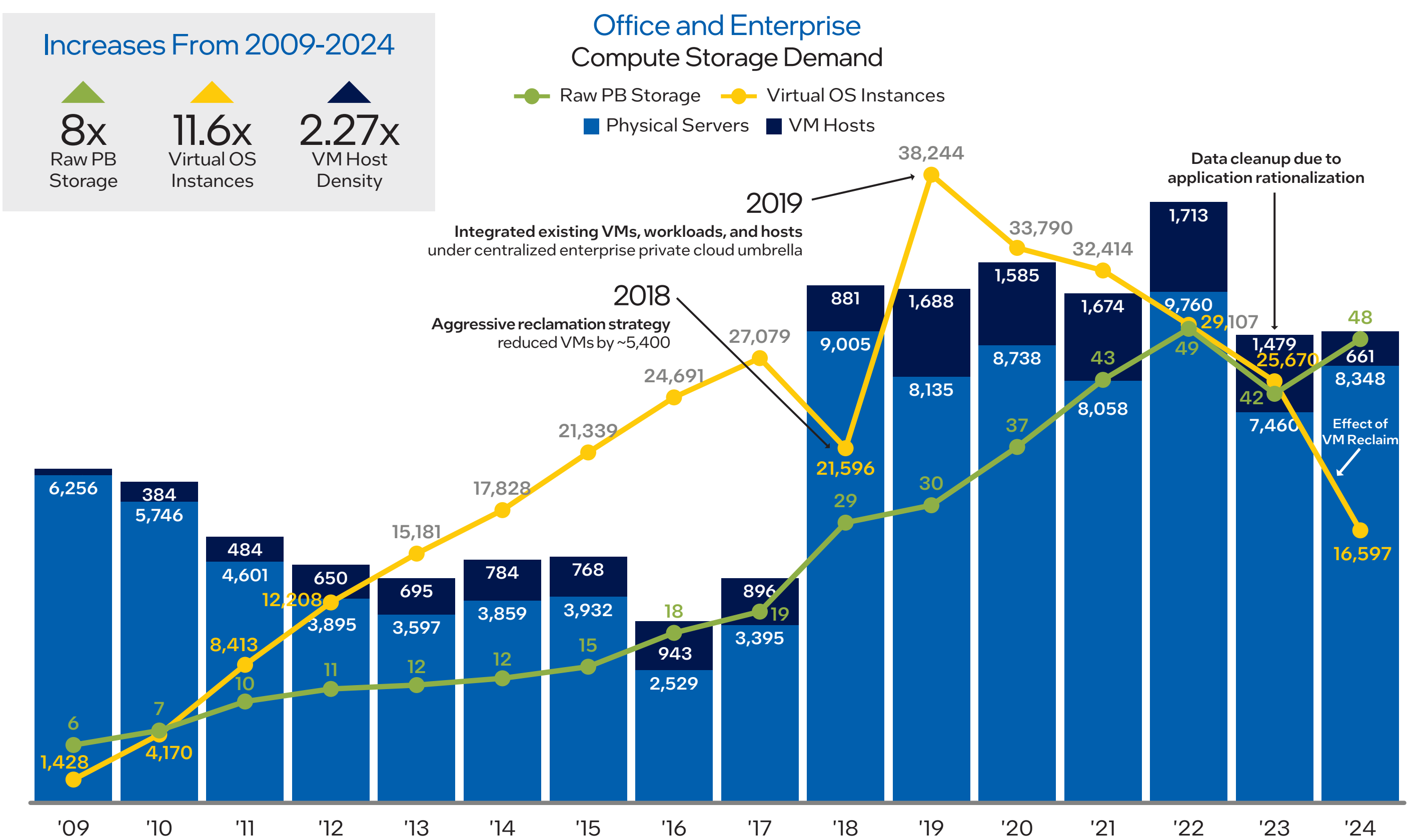

Figure 11. A high rate of virtualization combined with Intel architecture has enabled us to meet growing Office and Enterprise compute and storage demand while significantly decreasing the number of virtualization host servers.

2010-2024 Intel IT Data Center Strategy Results

Unit Growth Continues to Rise

Design

Environment Growth

Relative Performance 2010-2024

Higher Is Better

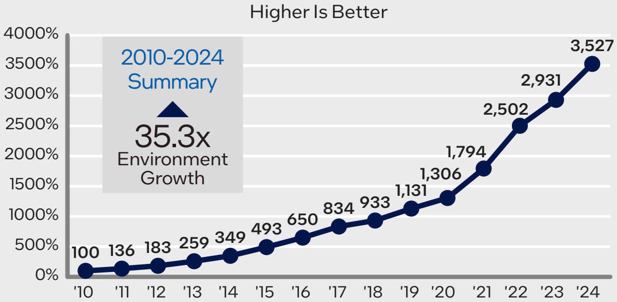

Total Spend vs. Per-Unit Cost

Relative Spending and Cost 2010-2024

Lower Is Better

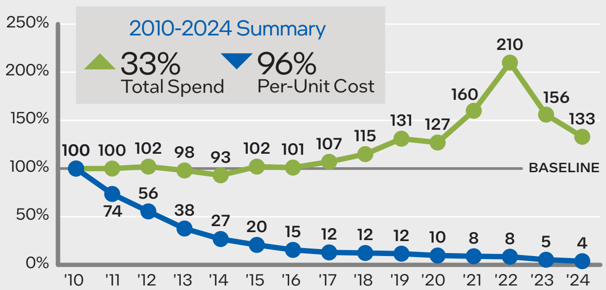

Cost Breakdown[a]

As a Percent of Total Cost

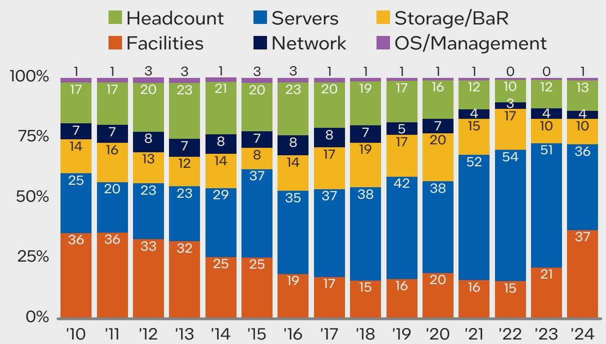

Change In Spending

Expanded View of Cost Breakdown

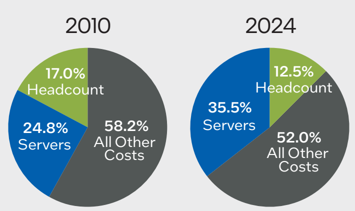

Office and Enterprise

Environment Growth

Relative Performance 2010-2024

Higher Is Better

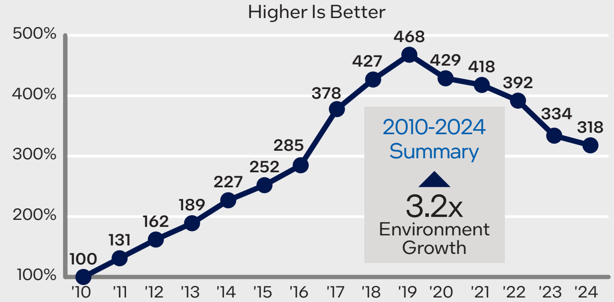

Total Spend vs. Per-Unit Cost

Relative Spending and Cost 2010-2024

Lower Is Better

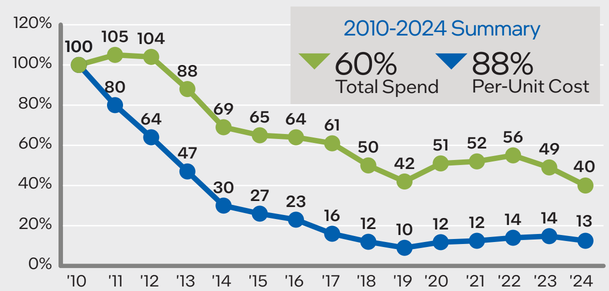

Cost Breakdown

As a Percent of Total Cost

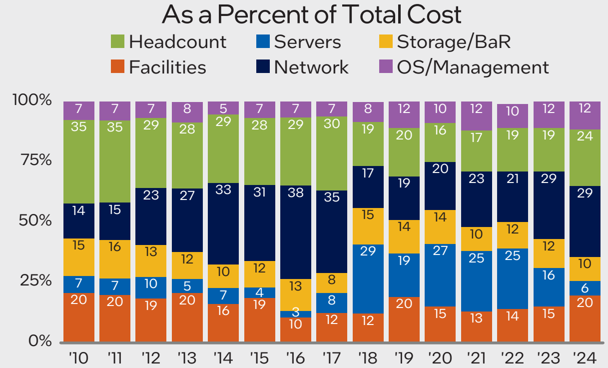

Change In Spending

Expanded View of Cost Breakdown

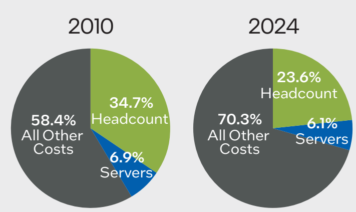

[a] Numbers in graphs may not add to 100 due to rounding.

2010-2024 Intel IT Data Center Best Practices

Applying Our MOR Approach Across Our Entire Infrastructure Stack

Servers

Adopted Disaggregated Servers
- Saved at least 44% over a full acquisition (rip-and-replace) refresh
- Reduced provisioning time (IT technician labor) by as much as 77%
- Decreased shipping weight of refreshed server material by 82%

Adopted Elastic Computing Services and Technologies
- Virtualized most of the Office and Enterprise servers
- Reduced the time it takes to provision a server from 90 days to on-demand provisioning using virtualization
- Enabled containers as a service

Enabled One-Day Dock-To-Production for Physical Servers
- Enhanced upfront planning and processes enhancement to order long-lead time items and rack readiness, which reduced the dock-to-production release from 10+ days to one day

Refreshed Servers Using The Latest Intel® Xeon® Processors
- Virtualization ratios of up to 60:1
- Reduced Design environment energy consumption by 10% annually between 2008 and 2013

Deployed SSDs as The Standard Local Disk in All New Servers
- Improved performance for I/O-intensive workloads and expected reduction of disk failure rates

Migrated Applications From RISC to Intel® Architecture
- Enabled significant savings and IT efficiencies
- Allowed us to realize the benefits of industry-standard operating systems and hardware

Deployed HPC
- Increased capacity (631x) and stability (322x) during HPC-7
- Saved USD 44.72 million net present value during HPC-1 itself

Enhanced Server Performance Through Software Optimization
- Increased Design job throughput up to 49%
- Delivered optimizations including disaggregated servers, NUMA-Booster, fast local data cache based on SSDs and high-frequency servers and optimal workload to platform pairing
- Boosted Design performance with smart jobs placement and scheduling, shared storage IOPS reduction, and idle jobs detection
- Significantly improved performance of data replication (up to 9x) and interactive jobs (up to 4.5x) over the WAN[a]

Storage

Refreshed and Modernized Storage Using The Latest Generations of Intel Xeon Processors
- Took advantage of innovative 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% in 2011-2012
- Reduced backup infrastructure cost due to greater sharing of resources
- Tiered backup solutions optimize backup costs, improve reliability

Right-Sized Storage Solutions Using a Tiered Model
- Provided storage resources based on business needs: performance, reliability, capacity, and cost
- Improved management of storage costs while still enabling easy access to necessary data
- Transitioned to scale-out storage to reduce operational complexity in tiering data
- Automated policy-based data migration between tiers

Implemented Thin Provisioning and Deduplication For Storage Resources
- Helped to control costs and increased resource utilization without adversely affecting performance
- Increased effective storage utilization in Design from 46% in 2011 to more than 80% now

Automatically Down-Tiered Inactive Blocks While Monitoring and Reclaiming Unused Data
- Policy-based down tiering of blocks that have not been recently accessed to reduce capacity demand rapidly and automatically for high-performance storage
- Continuously monitor and delete transient (non-IP) data that has not been accessed for six months or more based on customer expectations

Scaled Storage On Demand and Provided High-Performance Shared Scratch Spaces
- Enabled higher workload throughput for read-only storage areas that required high access

Network

Upgraded Data Center LAN Architecture to Support mGig/40/100 GbE
- Increased data center network bandwidth by 400% over three years, enabling us to respond faster to business needs and accommodate growth
- Increased network utilization from 40% in 2010 to 66% in 2024
- Eliminated spanning tree with multi-chassis link aggregation and Layer 3 protocol
- Reduced network complexity due to fewer NIC and LAN ports
- Decreased virtualized environment network cost by 18 to 25%

Opened The Data Center Network to Multiple Suppliers
- Generated more than USD 60 million in cost avoidance over five years with new network technology

Deployed Intel® Silicon Photonics Optical Transceivers
- For large-scale 100 GbE deployment, leveraged Intel® Silicon Photonics to significantly reduce the per-port cost

Facilities

Increased Data Center Efficiency
- Saved over 1.91 billion KW hours (from 2012-2024) compared to industry-standard data centers
- Added real-time control to optimally utilize data center power, cooling, and density while achieving maximum performance from Intel® Xeon® processor-based servers

Used a Tiered Approach to Redundancy, Availability, and Physical Hardening
- Improved matching of data center redundancy and availability features to business requirements
- Reduced wasted power by more than 7% by eliminating redundant power distribution systems within a data center

Retrofitted and Consolidated Data Centers Using a Modular Design
- Retrofitted a wafer fabrication plant to a high-density, high-efficiency data center modules with industry-leading PUE of 1.06
- Utilized free-air cooling and environmentally efficient evaporative cooling for maximum energy efficiency
- Avoided capital expenditures by not equipping the entire facility with generators

[a] Internal Intel IT, February 2020

Tactics

We are now applying our MOR approach across our entire infrastructure stack

Plans for 2025 and Beyond

Our data center strategy is continuously improving. We are always striving to close the gap between current achievements and the best possible scenario. To that end, we plan to continue to apply the MOR approach to our primary enabling tactics:

- **Embrace disruptive servers.** Deploy ultra-dense, power-optimized disaggregated 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. This will enable higher throughput more rapidly. Use strategic planning and storage protection technologies to deliver both backup and disaster-recovery coverage while reducing backup cost. Enhance automation to achieve fully autonomous performance and capacity management while providing greater visibility and control to our customers.
- **Increase facilities efficiency.** Use techniques such as higher ambient temperature for specific data center locations to take advantage of newer equipment specifications, which will help reduce cooling needs.
- **Drive network efficiency.** Continue to drive LAN utilization toward 75% and pursue an SDN to support agile, ultra-high-density data center designs. Continue to migrate to 100 GbE and 400 GbE with Intel® Silicon Photonics optics where appropriate and cost-effective, to meet network capacity demands. Drive the automation deeper into our day-to-day work.
- **Improve operational efficiency.** Increase the telemetry within the data center to improve the operational efficiency.

Conclusion

We are committed to providing a foundation for continuous innovation that will improve the quality, velocity, and efficiency of Intel IT's business services. To that end, 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 exceeding USD 11.41 billion from 2010 to 2024. Our data center transformation strategy is critical for Intel IT to stay competitive.

Key achievements include the following:

- Our breakthrough disaggregated server design allows independent refresh of CPU and memory without replacing other server components. This new design results in faster data center innovation and a minimum of 44% cost savings compared to a full-acquisition refresh. Along with this TCO reduction, the disaggregated server innovation enables significant TCE reduction (82% of material weight in a new server is removed with just a CPU-complex upgrade).
- One-day dock-to-production for new physical server deployment in our data center hub.
- We developed a system software capability called NUMA-Booster, which has saved millions while delivering additional usable server capacity.
- We deployed more than 40 PB of SSDs as fast local data cache drives. This increased workload performance due to lower network traffic and storage demand.
- Seven generations of HPC in our design computing environment created a 631x capacity increase and 322x quality improvement.
- We adopted new storage capabilities like deduplication and compression, accelerated storage refresh, focused on increasing utilization, removed unneeded data, and implemented policy-based tiering. All of these have resulted in getting additional usable capacity out of storage while reducing cost and providing higher performance.
- We deployed more than 94,000 100 GbE network ports, 1,000 40 GbE network ports, and 239,000 mGig network ports.

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:

- **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-cost 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. This enables us to maximize ROI and business value.

Related Content

If you liked this paper, you may also be interested in these:

- Fuel Cells: An Alternative Energy Source for Intel's Data Centers
- Green Computing at Scale
- Intel Takes On E-Waste With Disaggregated Servers
- Scaling Intel's Data Centers with Software-Defined Networking and Automation

For more on Intel IT best practices, visit **intel.com/IT**.

IT@Intel

We connect IT professionals with their IT peers inside Intel. Our IT department solves some of today's most demanding and complex technology issues. We want to share these lessons directly with our fellow IT professionals in an open peer-to-peer forum.

Our goal is simple: improve efficiency throughout the organization and enhance the business value of IT investments.

Visit us today at intel.com/IT or contact your local Intel representative if you would like to learn more.

Acronyms

DOME	Design, Office, Manufacturing, Enterprise
EDA-MIPS	electronic design automation MIPS
HPC	high-performance computing
KPI	key performance indicator
MIPS	meaningful indicator of performance per system
LAN	local area network
mGig	Multigigabit Ethernet
MOR	model of record
NFS	Network File System
NFSops	Network File System operations per second
NIC	network interface card
NUMA	non-uniform memory access
PUE	power usage effectiveness
QoS	quality of service
ROI	return on investment
SLA	service-level agreement
TCE	total cost to environment
TCO	total cost of ownership
VM	virtual machine
WAN	wide area network