

IT@Intel: Green Computing at Scale

Intel's innovative disaggregated server architecture and data center facilities design and operation reduce e-waste and increase energy efficiency—helping to protect the environment and produce significant cost savings

Intel IT Authors

Shesha Krishnapura
Intel Fellow & Chief Technology Officer

Ty Tang
Senior Principal Engineer

Vipul Lal
Senior Principal Engineer

Shaji Achuthan
Senior Staff Engineer

Table of Contents

Executive Summary	1
Background	2
Reduce e-Waste and Refresh Costs with Disaggregated Servers	3
Have Faith, Move Fast	5
Disaggregated Servers Do Not Increase Component Failure Rates	6
Improve Sustainability through Innovative Data Center Design and Processes	6
Conclusion	10
Related Content	10

Executive Summary

Do what's right for the business *and* what's right for the environment. That is the dual goal behind many of Intel IT's data center transformation initiatives. We have found that these goals are often closely linked: reducing total cost of ownership (TCO) can help us also reduce total cost to the environment (TCE).

Adopting disaggregated servers enables us to refresh only certain server components like the CPU and the memory without refreshing other components like networking, drives, fans, power supplies, cables or chassis. This approach to server refresh has not only provided a 44 percent cost savings, but has significantly reduced e-waste as well.

Our innovative approach to data center cooling reduces our electrical costs and contributes to one of the world's highest data center efficiency ratings—a power usage effectiveness (PUE) ratio of 1.06. These same approaches also have reduced fresh water usage by up to 61 million gallons per year. In addition, using fuel cells as a dual source of energy has allowed us to eliminate equipment such as power backup systems, which simplifies infrastructure design and lowers data center construction costs, and also reduces our carbon footprint.

Through disruptive technologies, innovative data center design, optimization of data center operational processes and higher resource utilization, we continue to look for ways to transform our data centers to better meet Intel's—and the world's—needs.

“Do what's right for the business *and* what's right for the environment.”

—Shesha Krishnapura
Intel Fellow & CTO

Background

Intel IT plays a crucial role in helping all of Intel meet its goals to create a *responsible, inclusive and sustainable world, enabled* through technology (RISE) goals. We continually seek ways to support Intel's efforts to be a global leader in sustainability and enable Intel's customers and others to reduce their environmental impact through our actions and technology.¹

One way in which we can contribute to Intel's sustainability is through data center transformation. The world's data centers currently consume 1–3 percent of the world's electrical supply,² and the world produces upwards of 50 million metric tons of e-waste annually.³ When Intel IT talks about data center strategy and data center transformation, we consider both total cost of ownership (TCO) and total cost to the environment (TCE).

The TCE terminology was first introduced to the industry in 2018 through a collaborative outreach effort between Supermicro, Intel and NASA. In the public video release, titled "[Mission: Green Computing](#)," Intel fellow and Intel IT chief technology officer Shesha Krishnapura said, "When you talk to the CIOs of the world, what you hear is digital transformation and most of the focus is on the application side. But what these people miss is the transformation at the infrastructure level."

"When you talk to the CIOs of the world, what you hear is digital transformation and most of the focus is on the application side. But what these people miss is the transformation at the infrastructure level."

— Shesha Krishnapura, Intel Fellow & CTO

Profile and Growth of Intel's Computing Environment

As of Q3 2021, Intel's data centers consist of over 360,000 servers (which includes over 3.3 million high-frequency Intel® Xeon® processor cores), over 577 petabytes (PB) of storage and more than 662,200 network ports, with an overall 103-megawatt (MW) data center capacity.

Four major business functions drive Intel's data center requirements, as listed in Figure 1. To support these four major functions, we deploy diverse data center infrastructure capabilities. Intel's silicon designs have become more complex and innovative. To support this complexity, 95 percent of our total data center capacity is devoted to supporting the Design Computing function, which has steadily grown at least 30–40 percent annually (across computing, storage and networking demand). Building, maintaining and "greening" Intel's computing infrastructure—which has grown in excess of 677x since 2003—is one of Intel IT's primary focus areas.

¹ For more information on Intel's RISE goals, visit <https://www.intel.com/content/www/us/en/corporate-responsibility/2030-goals.html>.

² Study with 3% as source: Bawden, T. "Global warming: Data centres to consume three times as much energy in next decade, experts warn." Independent. January 23, 2016. <https://www.independent.co.uk/environment/global-warming-data-centres-consume-three-times-much-energy-next-decade-experts-warn-a6830086.html>

Study with 1% as source: "Study: Data Centers Responsible for 1 Percent of All Electricity Consumed Worldwide." <https://www.datacenterknowledge.com/energy/study-data-centers-responsible-1-percent-all-electricity-consumed-worldwide>

³ UN environment programme, January 24, 2019, "UN report: Time to seize opportunity, tackle challenge of e-waste." <https://www.unep.org/news-and-stories/press-release/un-report-time-seize-opportunity-tackle-challenge-e-waste>

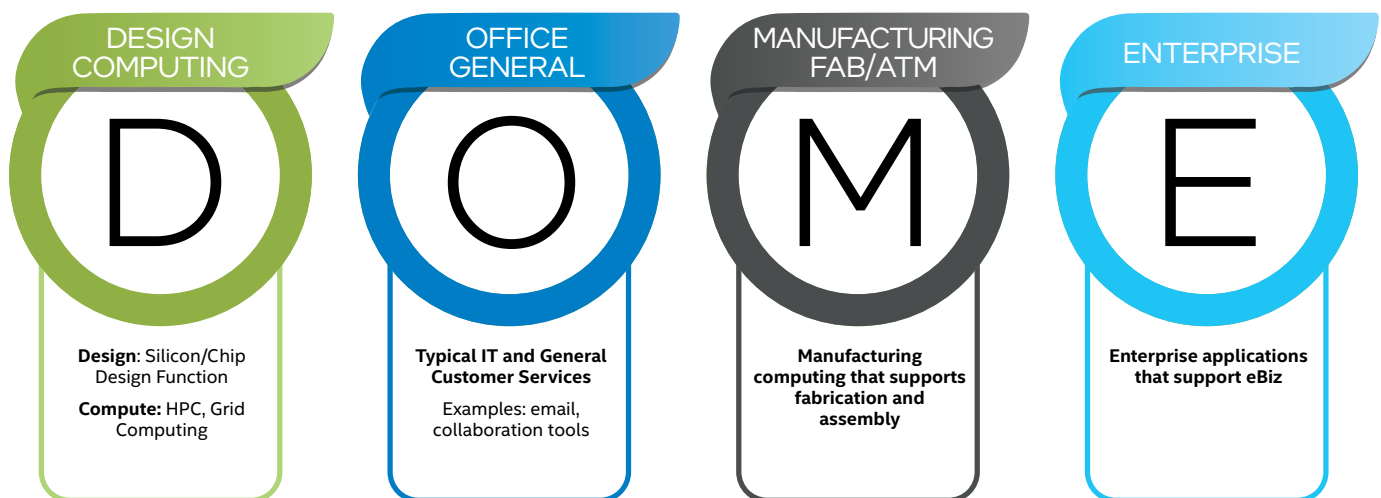


Figure 1. Major business functions supported by Intel's data centers

Intel IT Data Center Transformation Strategy

We operate our data center service like a factory by applying breakthrough technologies, solutions and processes to achieve industry leadership.

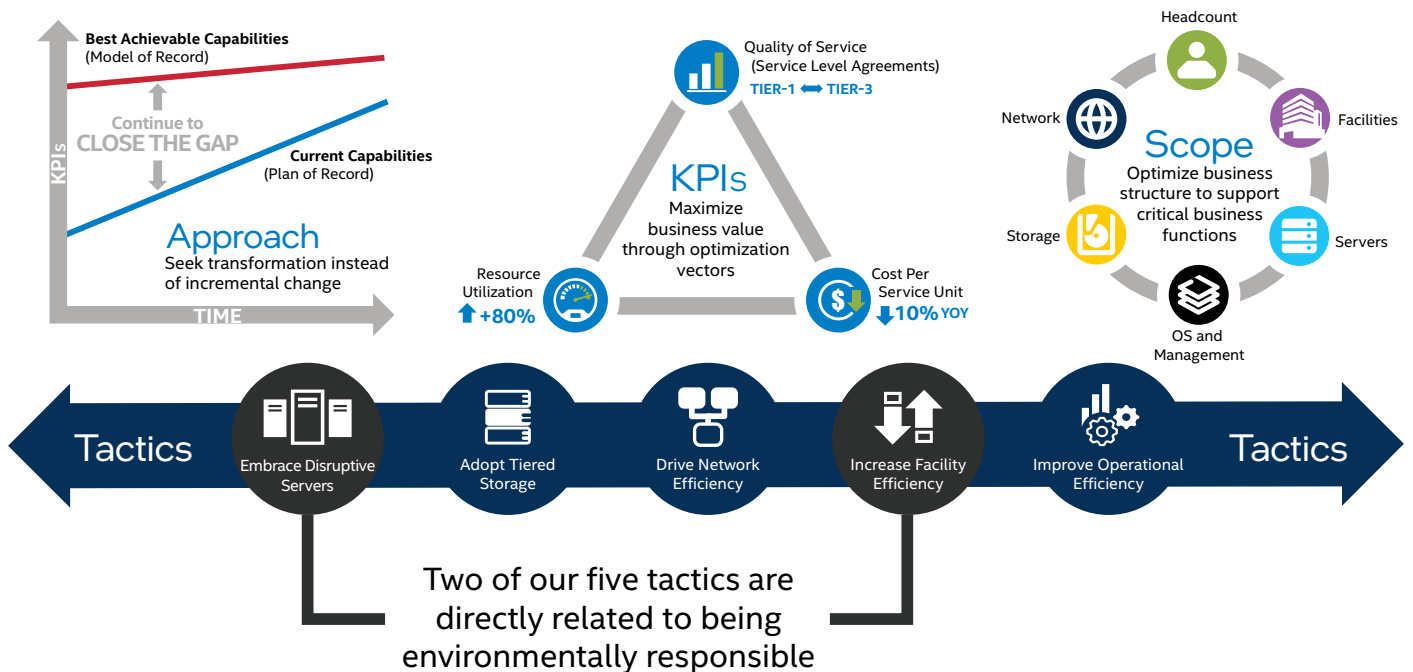


Figure 2. Disruptive server technology and facility efficiency are two of the underlying pillars of Intel IT's data center strategy.

Snapshot of Intel IT's Data Center Strategy

Intel IT runs Intel data center services like a factory, affecting change in a disciplined manner and applying breakthrough technologies, solutions and processes. This strategy (illustrated in Figure 2) enables us to optimally meet Intel's business requirements while providing our internal customers with effective data center infrastructure capabilities and innovative business services.⁴

One goal of our evolving data center strategy is to increase operational efficiency. Building on previous investments and techniques, our data center strategy has generated savings exceeding USD 3.8 billion from 2010 to 2019. But an equally important goal that drives our data center strategy is being environmentally responsible. Two of the five pillars of our data center strategy—adopting disruptive server technology and increasing facility efficiency—are directly related to both goals.

The rest of this white paper describes how implementing these two pillars substantially reduces e-waste and enables us to design data centers with some of the lowest PUE ratings in the world.⁵

Reduce e-Waste and Refresh Costs with Disaggregated Servers

Basic blade server design uses a shared power supply, fans, chassis, chassis controller and network; however, each blade has its own CPU with associated DRAM, along with direct attached storage (DAS) supported by either SAS or SATA drives and controllers. Intel IT refreshes data center servers every four years to take advantage of improvements in the Intel Xeon processor: more cores, better performance per core or more DRAM per core. But historically, we have needed to replace the entire server—even though many components such as the chassis itself plus cables, power supplies, network switches, fans and I/O components such as SSDs and SAS drives still have many years of useful life remaining.

It seemed that this approach to server refresh represented a needlessly high TCO and TCE. Why replace so many server components that do not change from one processor generation to the next? Why fill the recycling center or landfill with perfectly good drives and components? In 2016, these questions prompted Intel IT to reimagine server design, leading to the first server innovation in more than a decade: the disaggregated server.


⁴ For additional details about Intel IT's data center strategy, refer to the IT@Intel white paper, "Data Center Strategy Leading Intel's Business Transformation."

⁵ For more information, refer to the IT@Intel white paper, "Intel IT: Extremely Energy-Efficient, High-Density Data Centers."

The idea behind the disaggregated server is disarmingly simple: **Decouple the CPU/DRAM module from the NIC/drives module on the motherboard.** Redesigning the server to be modular enables us to upgrade the CPU/DRAM module to increase compute and memory performance and/or capacity while retaining the other components that are not ready for end-of-life. This results in faster technology adoption, which in turn puts new technology at our Design engineers' fingertips.

The new disaggregated design (see Figure 3) makes server refresh a whole new experience. Instead of spending many hours on a refresh, we can now simply remove a few screws, slide the CPU/DRAM module out, and install the new CPU/DRAM module. This module connects to the PCIe slot, which supports the multi-generational drives (SAS drive, SATA drive, or Intel® SSDs, including NVMe drives).

Example of a 1-Socket Disaggregated Server



Example of a 2-Socket Disaggregated Server




Figure 3. The disaggregated server architecture is characterized by a CPU/DRAM module and a NIC/drives module that can be refreshed independently of each other and of the rest of the server components.

Since the first disaggregated server design in 2016, we have continued to evolve the concept. We currently have deployed more than 220,000 disaggregated servers, using 13 different blade designs including both single-socket and two-socket servers. We use 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 performance cores, high bandwidth or high IOPS storage needs and the ability to add in accelerator cards on demand.

“The disaggregated server architecture is a perfect fit for our data centers. Just like when a homeowner upgrades lighting, replacing only the bulbs with the most energy-efficient ones without replacing the entire lighting fixture, Intel IT prefers to upgrade just the compute modules with the latest technologies without replacing the entire server infrastructure.”

— Shesha Krishnapura, Intel Fellow & CTO

Examples of Disaggregated Server Products

Below are examples of the disaggregated server products that are deployed in Intel® data centers. The servers and chassis are from Supermicro Computer, Inc. with Intel IP on disaggregated server architecture.



1-Socket Intel® Xeon® Processor-based Disaggregated Server (Entry-Level)



2-Socket Intel® Xeon® Scalable Processor-based Disaggregated Server



1-Socket Intel® Xeon® Scalable Processor-based Disaggregated Server



1-Socket Intel® Xeon® W-2200 Processor-based Disaggregated Server



1-Socket Intel® Xeon® W-3200 Processor-based Disaggregated Server

Environmental Benefits

Our internal testing indicates that ordering refresh components for a disaggregated server can save up to 86 percent in volume (meaning fewer boxes to ship and to store and stage) and up to 82 percent in shipping weight (meaning lower shipping costs). Avoiding disposal of still-useful components such as power supplies, SAS drives and fans, combined with fewer shipping materials and less fuel and time spent transporting the new parts can contribute to less e-waste. We estimate that we

reuse more than 50 percent of system components. Reuse also contributes to a smaller carbon footprint. Furthermore, when we selectively refresh a server's components, we can often support a more circular economy, reusing the disposed components as replacement spares for the remaining server fleet.⁶

Operational Benefits

Refreshing a 3U/14-blade or 6U/28-blade chassis by replacing all the blades but keeping the chassis itself along with the networking switch, power supply and fan modules saves 17 percent compared to a full-acquisition (rip-and-replace) refresh. But with disaggregated servers installed in the data centers, it is possible to refresh only the CPU/DRAM module, saving 44 percent in cost compared to a full-acquisition refresh (see Figure 4). (These results are based on internal testing at Intel and serve as an example only.)

In addition, there is no need to reinstall the OS or spend time replacing parts unnecessarily. In our internal tests, we determined that disaggregated servers represent a 77 percent reduction in technician time due to far fewer handoffs and required skill sets.

When viewed over a three- to five-year refresh cycle, the disaggregated server design can deliver, on-average, higher performance and more efficient servers at lower costs than a traditional rip-and-replace model by allowing data centers to independently adopt new and improved technologies. Also, the disaggregated servers we have installed are designed for advanced airflow and cooling. The ambient temperature for these servers can be as high as 104°F (40°C). Green computing features such as this give Intel IT the opportunity to operate our data centers more efficiently.

Have Faith, Move Fast

Concept Engagement to Full Large-Scale Production Delivery in Five Weeks

When Shesha Krishnapura, Intel Fellow and Intel IT Chief Technology Officer, first presented his idea for the disaggregated server in 2016, his idea was met with skepticism.

"An organized skepticism is part of the process of innovation," said Krishnapura. "Others must try to prove that your idea is not good enough."

But when colleagues looked closely at his design, they became convinced of its worth. With this backing, in June 2016 Krishnapura approached one of Intel's suppliers and told them he had a simple—but groundbreaking—idea.

"It is rerouting the motherboard, moving some of the components from the left side to the right side and adding this connector," Krishnapura told the supplier. "The cost should be minimal, and we should be able to do this very fast."

"Very fast" does not even begin to describe the pace at which the new design came to life. Krishnapura says the supplier used its vertically integrated full-service capabilities and collaborated closely with Intel IT to deliver a solution—an optimally tuned, high-quality product with full supply chain and large-scale delivery support—in five weeks. Within a few more weeks, several thousand of the new servers were installed and running Intel® silicon design jobs.

"With 280 Intel® Xeon® E or 100 Intel Xeon W or Intel Xeon Scalable processor-based server blades packed into a nine-foot rack, the high-density, high-efficiency and disaggregated architecture is a game changer," said Krishnapura. "For the first time it allows for the independent refresh cycles of the server compute modules."

⁶ Circular systems employ reuse, sharing, repair, refurbishment, remanufacturing and recycling to create a closed-loop system, minimizing the use of resource inputs and the creation of waste, pollution and carbon emissions (en.wikipedia.org/wiki/Circular_economy).




Figure 4. Refreshing the CPU/DRAM module in a disaggregated server saves at least 44 percent in cost compared to a full-acquisition (rip-and-replace) server refresh. Based on Intel internal testing, March 2017.

Disaggregated Servers Do Not Increase Component Failure Rates

Higher ambient temperature and selective component replacement does not increase AFR.

It may seem to someone who is not familiar with disaggregated servers that selectively replacing only the CPU and/or DRAM, while keeping most of the other server components, may lead to a higher annual failure rate (AFR). In addition, Intel IT's decision to run our data centers at a higher ambient temperature may also contribute to a higher AFR (see "Innovative Approaches to Data Center Cooling" for more information). But according to our data, our AFR has not risen significantly since putting these approaches to data center transformation into action.

We collected data based on analysis of 207,956 HPC servers from May 2019 through June 2020 and grouped them into two categories: a newer server group and an older server group. 114,252 (55 percent) of these servers were one to four years old, while 93,704 (45 percent) were four to eight years old.

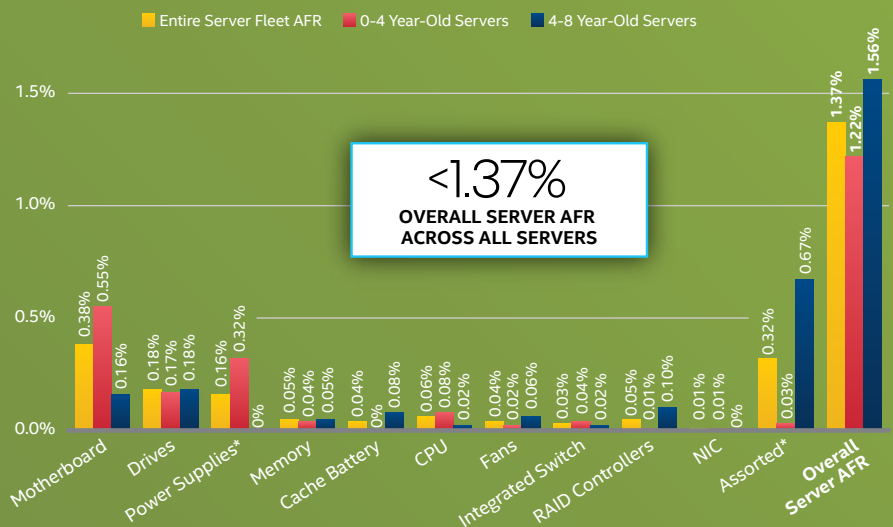
We calculated the following AFRs for each server group, revealing that the component failure for servers more than four years old is only marginally higher than for newer servers. The overall server AFR across all servers was 1.37 percent.

- 1.22% AFR for servers 0-4 years old
- 1.56% AFR for older servers

The following figure shows the AFR by component across the entire dataset:

This data shows that selectively replacing components is an environmentally conscious decision that can lower total cost to the environment without harming Intel's total cost of ownership. What's more, the AFR data was collected from two data centers—one with a power usage effectiveness (PUE) ratio of 1.06, running at 91°F (32.7°C), and one with a PUE of 1.49 running at a much lower temperature. This data also proves that a higher ambient temperature, if kept within approved server specifications, does not hurt the life of the components based on our experience.

Annualized Failure Rate by Server Component



Number of Components
Overall components analyzed across 207,956 servers

Component Type	Across Entire Server Fleet	0-4 Year-Old Servers	4-8 Year-Old Servers
Motherboard	207,956	114,252	93,704
Drives	229,930	123,216	106,714
Power Supplies	48,413	23,518	24,895*
Memory	1,090,710	599,405	491,305
Cache Battery	207,956	114,252	93,704
CPU	223,050	119,033	104,017
Fans	97,941	49,335	48,606
Integrated Switch	17,787	7,495	10,292
RAID Controllers	207,956	114,252	93,704
Network Interface Card	207,956	114,252	93,704

* For 4-8 year-old servers, power supply failures are reclassified as "assorted (loose cables, component reseat)".

Improve Sustainability through Innovative Data Center Design and Processes

The energy efficiency of a data center is traditionally measured by its PUE, which is a ratio of the IT load (the power consumed by computing equipment as opposed to cooling and other overhead) to the total power consumption of the data center. Ideally, the PUE would be 1; that is, all the power consumed is used by the computing equipment.

The energy efficiency of a data center directly impacts both TCO and TCE. For example, for every 1 MW of power consumed by a data center, reducing the PUE from 1.5 to 1.06

saves over USD 385,000 in annual utility costs.⁷ Possibly more importantly, this same PUE reduction is responsible for 2,726 avoided metric tons of CO₂ per year, which is equivalent to carbon sequestered from 3,559 acres of forest in a year.⁸

As we pursue our overall strategy to reduce carbon footprint and e-waste while still meeting data centers' dramatically increasing energy demands, disaggregated servers are only one portion of a broader picture that includes innovative approaches to data center cooling, alternative sources of energy and data center processes and resource optimization.

⁷ PUE and cost savings can be calculated at [42u.com/measurement/pue-dcie.htm](https://www.42u.com/measurement/pue-dcie.htm).

⁸ US EPA AVERT Calculator epa.gov/energy/greenhouse-gas-equivalencies-calculator

Innovative Approaches to Data Center Cooling

Intel IT is continually refining data center design to increase density and efficiency.⁹ Since the 1990s, our data centers have evolved through three generations.

- **Gen 1A (1990s)**—Characterized by forced chilled air from the ceiling, with 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 more than 2.0.
- **Gen 1B (1990s)**—Data centers that used chilled air from the row end had a PUE of ~1.4.
- **Gen 2A (early 2000s)**—These data centers had a raised floor and forced chilled air or hot/cold air segregation kept density at 42U, but power consumption delivered to the racks increased to 15 kW, resulting in a lower PUE of ~1.4.
- **Gen 2B (mid-2000s)**—Further improvements such as the addition of chimneys to the chilled air from the row end and hot/cold air segregation enabled us to deliver 30 kW to the racks and lowered the PUE to ~1.18.
- **Gen 3A (2013)**—One of our modern data centers uses free air cooling along with hot/cold air segregation and extremely high rack density to achieve an industry-leading PUE of 1.06, with an extreme rack density of 60U and up to 43 kW/rack.
- **Gen 3B (2015)**—Another modern data center uses close-coupled evaporative cooling and hot/cold air segregation to achieve a similar PUE, rack density and kW/rack.

Historically, data centers run at cool temperatures, from 64°F (17.8°C) to 68°F (20°C). But many newer server platforms allow for a wide range of thermal operating conditions—from 32°F (0°C) to 104°F (40°C). The specific equipment we are using is designed to operate between 41°F (5°C) and 95°F (35°C). We recognized an opportunity to use innovative cooling designs along with operating our data centers at a higher temperature—91°F (32.7°C)—to

⁹ For information about Intel's efforts to also improve manufacturing facilities' efficiency, read the "Reducing Carbon Through Energy Conservation in Manufacturing" white paper.

reduce or eliminate expensive CRAC units and chillers. **In the eight years we've been running data centers at the higher temperatures, we have observed no increase in component failure.**

Our innovative cooling approaches and extremely high rack densities have enabled us to **meet increasing demands for compute and storage resources over the years while reducing our data center space footprint by 22 percent** (see Figure 5). The following two sections describe free-air cooling and close-coupled evaporative cooling at a high level. For more information, refer to the IT@Intel white paper, "Intel IT: Extremely Energy-Efficient, High-Density Data Centers."

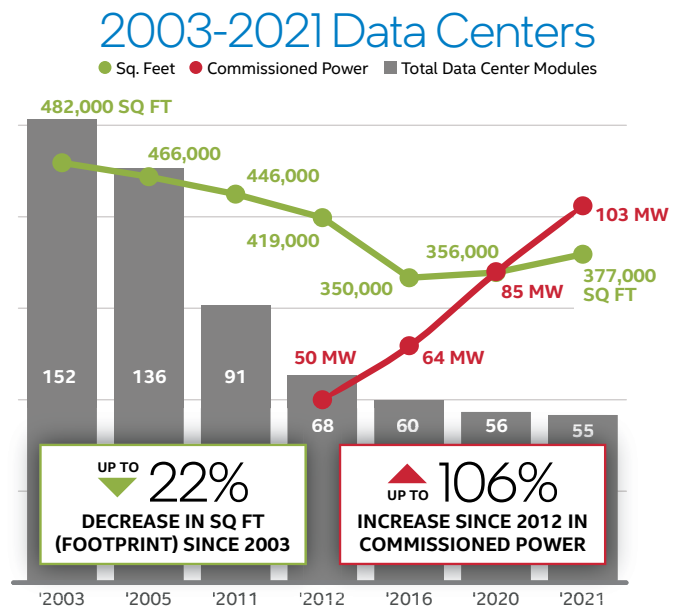


Figure 5. Innovative data center designs have enabled us to decrease data center square footage while increasing power density and capacity.



A close-coupled, evaporative cooled, Gen 3B data center hot aisle with 60U racks with up to 280 disaggregated servers/rack and up to 43 kW/rack.

Free-Air Cooling

One of our modern data centers is located in a building that was originally used as a manufacturing facility. For this building, we chose a flooded-air design that provides supply air to the cold aisles by flooding the room with air at a slightly positive pressure. This design allows any amount of air required by any server to be available. The building was already equipped with rooftop air handlers, which we repurposed so that they now provide the airflow filtering and supplemental cooling. The supply-air demand for the room is 542,324 cubic feet per minute, and the facility more than meets this by providing 572,000 cubic feet per minute. We also reconfigured the existing rooftop supply-air plenums to connect to the room's overhead supply-air plenum, which is large enough to accommodate two school buses parked side by side (see Figure 6). If the outside temperature exceeds 90°F (32°C), we augment free-air cooling by using the available chilled water supply. Doing so ramps up the supplemental cooling before it is actually needed.

This facility also uses alternating hot aisles (exhaust air) and cold aisles (supply air). Airtight doors at the end of each hot aisle and the translucent air segregation enclosure above the server racks are made of a material similar to what is used for heavy-duty greenhouse roofs and walls. Hot aisles average 110°F (43°C) in the winter and 125°F (52°C) in the summer. The range of temperatures for the supply and exhaust air in the room makes the environmental conditions acceptable for the few staff that work occasionally in the room.

With free-air cooling, other innovations such as hot/cold air segregation and state-of-the-art Intel Xeon processor-based servers enable us to achieve a rack power density of up to 43 kW and a PUE of 1.06.




Figure 6. Alternating hot aisles (exhaust air) and cold aisles (supply air) provide efficient air segregation, and the exhaust-air plenum's louvers are positioned parallel to prevailing winds so that hot air can easily exit the building.

Close-Coupled Evaporative Cooling

In our other modern data centers, we use a recycled-water-based cooling solution that maximizes the wet-side economizer concept and pushes the limits of evaporative cooling for the climate zone, combined with a cooling coil placement near the data center heat load (see Figure 7). A two-inch pipe carries slow-moving water at a low pressure between the two racks. Thin copper coils combined with fans at the top of the racks transfer the hot air from the rack through the coils to the water in the pipe. The cooler air comes right back in front of the rack. The warm water and cool water from outside exchange heat through the plate heat exchanger. In short, this system is similar to blowing on a hot cup of tea to cool it down. It is a simple technique that is easy to construct and manage and is significantly less expensive to operate than traditionally cooled data centers. We estimate that a 20 MW data center that uses close-coupled evaporative cooling can save 61 million gallons of fresh water annually.

Like the free-air-cooled facility, the close-coupled evaporative cooling configuration allows us to operate 144 racks (72 on each side) at a rack density of up to 43 kW and achieve a PUE of 1.06. When we first constructed the close-coupled evaporative cooling data center in 2015, we had a single 31 MW module. Since that time, we have added two additional modules (10 MW and 8.5 MW), for a total capacity of 49.5 MW.

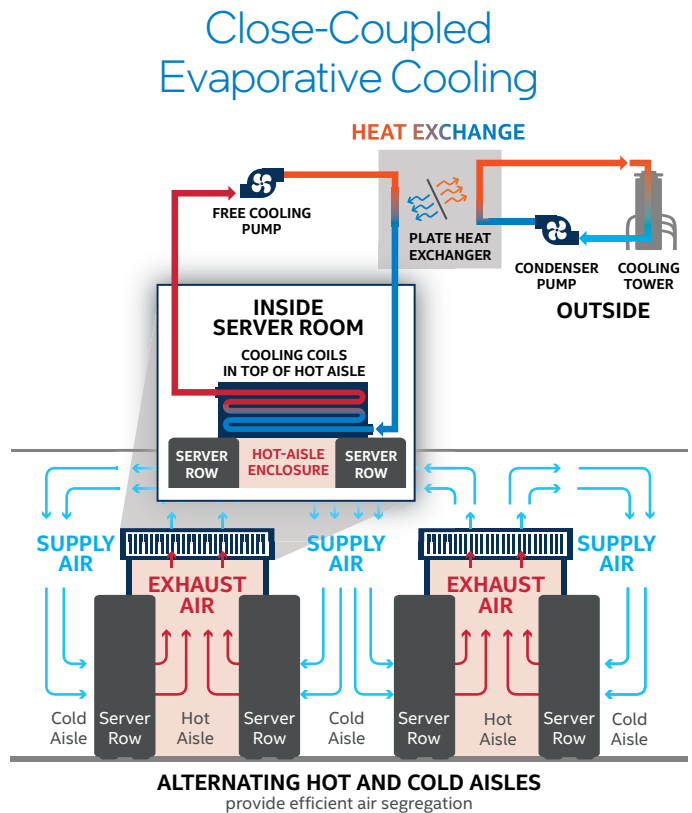


Figure 7. The close-coupled evaporative cooling design uses water-cooled coils to transfer heat away from the servers.

Advantages of a Custom Rack Design

Intel has 71 percent more capacity per rack floor space than standard industry designs.

We are maximizing our investment by increasing the density rather than scaling out, and we have more than doubled our previous server capacity per row. We now provide 71 percent more capacity per rack floor space than standard industry designs.

Besides employing a highly dense data center rack space, we use advanced small form-factor design servers for extremely high rack capacity, coupled with the highest performance per core. The combination of higher performance per core and high server rack density enables us to offer much more compute capacity throughput than traditional data centers of similar footprint.

The disaggregated servers allow us to choose from a broad range of Intel® Xeon® Scalable, Intel® Xeon® W and Intel® Xeon® E high-performing SKUs for our silicon design workloads. We selectively upgrade some of the servers with the latest high-performance CPU SKUs for critical-path workloads to maintain the best possible performance per core for our electronic design automation (EDA) workloads with optimal TCO.

We increased the rack height to 60U and decreased the rack width to less than 20 inches. This design resulted in an overall cooling density of 1,100 W/square foot and a rack power density of up to 43 kW/rack—1.5x more than what we have delivered for high-density computing in the past, and 5x greater than the industry average³—a density that the industry has referred to as “impossible” and “remarkable” for air-cooled facilities.

³ Uptime Institute, December 7, 2020, “Density is rising,” <https://journal.uptimeinstitute.com/rack-density-is-rising>

Alternative Sources of Power

In several proofs of concept, fuel cell technology has proved to be a reliable, stable, efficient and sustainable source of electricity for Intel's data centers. Intel IT joined forces with Intel's Corporate Services and Global Supply Chain groups to enable fuel cells to meet Intel's redundancy and capacity expectations.

For some of Intel's data centers, fuel cells are the primary source of power (as opposed to the power grid). To date, we have fuel cell installations at multiple global Intel data center locations. Fuel cells are now part of Intel IT's plan of record for new data center design and expansion scope.

Our use of fuel cells has provided many benefits, including the following:

- Efficient use of real estate
- Reliable and high-quality power
- Scalable modular installations
- Affordability
- Faster deployment time
- Lower CO2 emissions, which helps reduce Intel's TCE

For more details about our experience with and best practices for fuel cell installations, read the IT@Intel white paper, “[Fuel Cells – An Alternative Energy Source for Intel's Data Centers.](#)”

Data Center Processes and Resource Optimization

Beyond the use of disaggregated servers, innovative approaches to cooling and alternative power sources, we take advantage of every opportunity to optimize data center practices to benefit both Intel and the environment.

Intel IT's global commitment to wise water use—through a combination of reuse and recycling, conservation and community-based restoration—is a primary focus area. In 2020, Intel's largest water recycling plant reached a milestone of 1 billion gallons (that's 26 minutes of water flowing over Niagara Falls). The recycled water goes into scrubbers, cooling towers and other equipment, creating a huge loop of reclaim and reuse. At some sites, we also recycle water by sending it from the data center to a nearby manufacturing facility. Over the last decade, our global water conservation and related efforts saved an estimated 44 billion gallons of water (19 hours of water flowing over Niagara Falls).¹⁰

Over the past several years, our sustainable water management efforts, such as recycling and reuse, have enabled us to procure over 7.1 billion kWh of green power, which is equivalent to 82 percent of Intel's global electricity use. Our new goal is to achieve 100 percent of our global energy use through green power by 2030.¹¹ Intel also participates in extensive initiatives to reduce the company's carbon footprint.¹²

¹⁰For the full press release, visit <https://www.intel.com/content/www/us/en/newsroom/news/water-facility-milestone-billion-gallons-recycled.html>.

¹¹Intel and the Environment

¹² See “Reducing Carbon Through Energy Conservation in Manufacturing”

Intel's commitment to being environmentally responsible is underscored by the company recently signing the "EU Climate Neutral Data Centre Pact," which is part of the European Green Deal, committed to making Europe climate neutral by 2050. Intel has identified specific data center modules in Europe that will pursue initiatives relating to energy efficiency, clean energy, water conservation, circular economy and circular energy.

Conclusion

Intel's business needs will continue to increase, and we will likely continue to see double-digit growth in the demand for compute and storage. And yet, the global challenges facing the IT industry, such as climate change, e-waste and responsible water usage continue to mount as well. "Green computing" is more than just pursuing the lowest possible PUE; it encompasses a broad range of initiatives and concepts. By adopting disruptive technologies like disaggregated servers and constantly thinking outside the box about how we operate our data centers, Intel IT is proving that the concepts of operationally efficient (TCO) and environmentally friendly (TCE) are often more closely linked than many people may suspect.

Related Content

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

- [Data Center Strategy Leading Intel's Business Transformation](#) white paper
- [Fuel Cells – An Alternative Energy Source for Intel's Data Centers](#) white paper
- [Hyperscale High-Performance Computing for Silicon Design](#) white paper

For more information 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, and 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.

Follow us and join the conversation:

- [Twitter](#)
- [LinkedIn](#)
- [#IntelIT](#)
- [IT Peer Network](#)

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

Intel IT Contributors

Shruthi Balasubramanyam, Site Reliability Engineer
Paul Budhai, Data Center Manager
Toan Cao, Data Center Mechanical Engineer
Nissim Hamu, Data Center Electrical Engineer
Prasad Pusuluri, Site Reliability Engineer
Harish Srinivasappa, Site Reliability Engineer
Michelle Vincenti Urpí, Principal Electrical Engineer
Jenny Wu, Site Reliability Engineer

Acronyms

AFR	annual failure rate
CRAC	computer room air conditioning
KPI	key performance indicator
PUE	power usage effectiveness
RISE	responsible, inclusive and sustainable world, enabled (through technology)
TCE	total cost to the environment
TCO	total cost of ownership



Intel technologies may require enabled hardware, software or service activation.
 No product or component can be absolutely secure.
 Your costs and results may vary.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.
 Other names and brands may be claimed as the property of others. 0821/WWES/KC/PDF