



intel®

2023-24 Intel IT Annual Performance Report

The Next Evolution of IT



Insights From Our CIO

Intel is evolving quickly, and Intel IT is a key driver in that evolution. As the architects of Intel's IT landscape, we are the strategic, trusted technology partner for Intel's business. Our focus is on innovation, AI, and optimization as well as the IT experience of our workforce, our end users, and Intel's customers.

We run IT as a business. That means we focus on value-based outcomes and invest in and prioritize technical projects *only* when they provide business value by increasing revenue, reducing costs, or providing savings.

We can recognize business value only when we are in tune with Intel's business units. Over the years, we have forged critical alliances with enterprise stakeholders while adhering to strategies and operations that prioritize the organization's overarching goals. But there is more to be done, and Intel IT is evolving to meet the challenges on the horizon. As Intel's new CIO, I work with the IT team to ensure that IT is a beacon of innovation, providing the essential technology and resources that will spearhead Intel's transformation and continued growth. As we continue into 2024, it is paramount that we cultivate deeper partnerships across business units. Embracing a "team-of-teams" ethos, our organizational blueprint includes efficient integration, high standards, clarity, and accountability.

A mindset of "One Intel" and "One IT" is crucial to our acceleration of Intel's IDM 2.0¹ evolution. We formed the IT IDM 2.0 Acceleration Office (IAO), working closely with business units and functional teams as they bring a new internal foundry model to life and start reporting separate profit and loss statements.

The IAO is now embarking on a massive upgrade to Intel's enterprise resource planning (ERP) system to robustly support Intel's products and the Intel foundry business units. This brings an incredible opportunity to reduce costs, with a goal to remove about 40% of our legacy applications over the next three years. Other IAO focus areas will include supply chain planning and data analytics.

The IAO has collectively saved Intel USD 3 billion, and more is to be gained. The IAO will help the internal foundry model restore Intel's healthy margins, establish a competitive cost structure, and enhance the culture and incentives required to deliver process and product leadership.

I am proud of Intel IT's agility, innovation, resilience, and focus on operational excellence. Intel IT has responded to the largest priority of a huge company with alacrity by taking a leadership role. Intel's former president and CEO, Andy Grove, captured the excitement and energy that is coursing through Intel IT's veins when he said, *"A corporation is a living organism; it has to continue to shed its skin. Methods have to change. Focus has to change. Values have to change. The sum total of those changes is transformation."*

In 2024, Intel IT will build on that evolutionary concept by delivering the modern, trusted, and resilient technology solutions required to enable Intel's IDM 2.0 transformation.

Motti Finkelstein

Corporate Vice President and Chief Information Officer

¹ "IDM 2.0" is a major evolution of Intel's integrated device manufacturing (IDM) model. For more information, read Intel's press release from March 2021, "Intel CEO Pat Gelsinger Announces 'IDM 2.0' Strategy for Manufacturing, Innovation and Product Leadership."

In this issue

- 4 Constructing the Intel IT Council and the Joint Technology Council
- 5 Using AI and Analytics to Improve IT's Operational Efficiency
- 9 Building a Cloud-to-Edge Infrastructure
- 11 The Continuing Evolution of Information Security
- 13 Maximizing Client Productivity and Optimizing User Experience
- 15 Driving Sustainability through Technology
- 17 Reinvigorating a Customer Zero Mindset
- 19 Delivering the Next Evolution of Intel



~2,350
IT Employees



SUPPORT

~124K
Employees



115
Intel Sites




IN

53
Countries


Infrastructure

■ Storage ■ Servers




IT Spending

■ Per Employee ■ Against Revenue




IT Managed Devices

■ PCs ■ Smart Devices



Our IT Environment



USD 1.4 B

Business value delivered through AI solutions




USD 7.5 B

Data center strategy cumulative cost savings



USD 72 M

Supply chain business value delivered in U.S. tariff savings



USD 239 M

Infrastructure capital cost reduction



~1.4 M Devices

Monitored for security compliance



13-Week Reduction

Pre-silicon design cycles

Data as of January 1, 2024. Employee count represents an average of beginning and end of year.



Leadership

Constructing the Intel IT Council and the Joint Technology Council

After the IT organization was realigned into the Office of the CTO, we established the IT Council to accelerate business transformation, increase security, and fortify technology alignment and enablement across all Intel business units, while providing ample focus in our critical IT areas.

The Intel IT Council leadership team supported three organizational pillars: Information Security and Privacy; Corporate IT and Digital Transformation; and Engineering IT and Developer Cloud Services. The IT Council also coordinated closely with Manufacturing, Supply Chain and Operations, and Technology Development leadership to enable their acceleration of IDM 2.0.

The IT Council provided experienced leaders with deep and diverse knowledge to scale IT businesses and services to help them meet growing business information and computing needs. The IT Council's quick and efficient decisions to align objectives, priorities, innovation, and investments will lead IT into the future.

As Intel's business matures, the IT Council has evolved, beginning with Intel naming a CIO as the Council leader at the end of 2023. We are also morphing the IT Council into a Joint Technology Council, which includes leaders from all of Intel's business units to enable better relationship-building and prioritization of strategic investments.

"I treat IT as if it's an engineering function – we're an engineering company and so IT should engineer our way into a new future of technology."

– Greg Lavender, Executive Vice President, Chief Technology Officer

The Joint Technology Council will underscore Intel IT's evolving culture and mindset. In 2024, we will continue to develop our leadership role through both councils and focus on maximizing our partnerships and interactions with the business units. We will continue to drive standards and consistency across platforms, architecture, data, and analytics to continue Intel's IDM 2.0 journey.

Explore

- [Creating a System of Innovation](#)
- [A Systematic Approach to Encouraging Innovation](#)

Driving Innovation at Intel

Intel IT's evolution requires new ideas, new energy, and new mindsets. That's the idea behind our "system of innovation." We believe innovation is a skill, not an accident. We take a systematic approach to bring innovative ideas to the surface and help leaders to follow through on every idea.

Seven foundational principles allow workers to explore ideas and reduce technical debt. In addition, we reserve 20% of team capacity for continuous improvement efforts through an initiative called Keep Improving IT.

To build Intel IT employees' innovation skillsets, we hold quarterly events called Innovation Days (InnoDays), where the primary rule is "There are no bad ideas." During InnoDays, teams have 24 hours to ideate, experiment, and achieve success—and learn from failure. So far, InnoDays have resulted in innovations that save at least 172 hours per week (9,000 hours per year) by implementing 69 new ideas.



Operations

Using AI and Analytics to Improve IT's Operational Efficiency

AI is the key to transforming IT and the foundation to Intel's evolution. As Intel's CTO Greg Lavender says, Intel IT must apply AI solutions "to everything, from how hybrid processors distribute work, to spotting early yield issues on a factory line, to mining the web for new sales opportunities." Our CIO adds that AI will "transform us as IT and transform Intel to be more efficient. Many things we do today should be automated, should be resilient, and should just work."



Our AI journey began almost 20 years ago, and since then we have conducted more than 100 use cases and proofs of concept across all functions. In 2023, our AI efforts contributed more than USD 1.4 billion in value to Intel.²

To expand our existing AI implementations with Intel's business units and complement our "AI Everywhere" initiative and the Intel IT AI Center of Excellence, we have collaboratively launched "AI Inside." We expect this programmatic approach to developing innovative AI use cases across the company to achieve a 30-40% improvement in productivity, quality, and asset utilization over the next 18-24 months.

² This value is measured for all AI solutions in production for 2023 outcomes. The different solutions either contribute to Intel revenue (via product improvements or sales impact) or to cost saving (via cost of development or manufacturing, which impacts factory yield).

Collectively, these initiatives will help scale our current AI solutions, drive efficiency, and expand emerging capabilities. Starting with four areas (silicon design, test and validation, software development, and manufacturing) we will accelerate AI benefits by concentrating on common technology stacks.

With our successes in embedding AI within our most critical business processes and building it into Intel's products, we plan to rapidly expand our AI efforts with visibility and sponsorship from our top leaders.



"To take the whole company into the future, we need a more futuristic IT. One, for instance, that applies AI solutions to everything from how hybrid processors distribute work to spotting early yield issues on a factory line to mining the web for new sales opportunities."

– Greg Lavender, Executive Vice President,
Chief Technology Officer



Increasing Scalability and Repeatability

AI solutions are integral to Intel's business-critical activities, so there is a tight focus on model quality and preventing performance degradation. Given the number of AI models we have in production, the cost of manually monitoring them for degradation and retraining them would prevent our teams from undertaking any new projects.

To enhance repeatability and scalability, we have integrated industry-standard, automated quality control principles and tools throughout the model lifecycle. Our vision is for AI solutions to understand the quality state of their functions and respond independently to that state without any human intervention. Achieving total automation is a work in progress; however, our MLOps solution has enabled us to develop, deploy, and support hundreds of AI model quality gates with low cost and effort—allowing us to scale AI as fast as the business needs it.

“We are building an enterprise AI environment to address our current and future needs for productivity, quality, and optimization.”

– Motti Finklestein, Corporate VP,
Chief Information Officer



Transforming Intel's Manufacturing Processes

We continue to evolve Intel's manufacturing processes by adopting various technologies like Industrial Internet of Things (IIoT) edge compute, machine vision coupled with AI/ML, and data analytics. These innovations improve equipment predictive maintenance and inline quality inspections, which in turn help manufacturing engineers improve cost, quality, and throughput time.

End-of-Line Inspection

Manual end-of-line inspection is time consuming and highly dependent on yield analysis engineers' experience and skill. It is also limited in the sense that only a portion of the product can be inspected per lot. We've changed all that by improving wafer yields with an AI-driven, automated pattern classification solution that alerts yield analysts to quickly address the problem—preventing more wafers from being affected.

Our AI algorithms enable examination of 100% of the wafers in every lot and detect baseline patterns with >90% accuracy. The combination of professional experience with AI allows yield analysis engineers to support more products, use knowledge captured collectively across fabs, and shorten the time to resolution through root-cause analysis. The solution runs in a private cloud powered by Intel® Xeon® Scalable processors and open-source software.




Inline Inspection

We use computer vision and machine learning to perform inline inspection during assembly and test processes. Inline inspection can catch defects and tool issues much earlier than end-of-line inspection. Our solution can also detect micrometer-sized defects that are hard to see even with a microscope, as well as defects that are impossible to detect after the process is complete.

High-resolution cameras take multiple images per second; the images are then analyzed by a machine-learning model at the edge. If defects are detected, the solution can sound an alarm or even stop the tool. The solution, which includes Intel® Core™ i9 processors, Intel Xeon Scalable processors, and Intel® ARC A770 discrete GPUs, is being deployed to all Intel assembly and test factories.

This inline inspection solution is already saving Intel up to USD 2 million annually. It also reduces business risk, improves product quality, and frees engineers from tedious manual offline inspection. We anticipate that the solution will eventually extend to additional processes, multiplying these business benefits.



Powering Digital Twins and Intel® Automated Factory Solutions

The fabrication of semiconductors includes stringent requirements and consumes substantial time and budget. Optimizing the fab processes can help improve quality and save costs.

Intel® Automated Factory Solutions (Intel® AFS) digitize the concept of modeling, raising it to new, dynamic levels. Now, instead of simple physical replicas, manufacturers can build Digital Twins (DTs), which are datasets that simulate not only the physical attributes of entities (such as shape, color, and size), but also less tangible characteristics (such as strength, elasticity, conductivity and more). Plus, once created, DTs of different objectives can be combined into DT systems that mimic their real-world counterparts' behavior. Intel has used DTs to significantly elevate per-person productivity and reduce unit throughput time while maintaining outstanding product quality.

The Intel AFS team will leverage Intel's expertise in manufacturing automation software like DT technology and provide it as commercial products for other manufacturers around the world.

For example, Sony Semiconductor Manufacturing Corporation (SCK) collaborated with Intel on their software projects, Intel® Factory Pathfinder and Intel® Factory Recon.

Within six months of live operation, SCK produced overwhelmingly positive results. They achieved an impressive 25% reduction in the transport traffic on the interconnects between fabs while maintaining previous levels of manufacturing output.

Our work on Digital Twins, including the development of the Intel AFS software suite and our collaboration with SCK was awarded a [2024 CIO 100 award](#).



Explore

- [Optimizing Factory Performance with DT Technology](#)
- [Intel and Sony Advance DT Technology for Manufacturing](#)
- [Intel Operations Recon](#)

Factory Equipment Availability and Productivity

We have developed a standardized IIoT infrastructure that is highly scalable and can be used for a variety of use cases. The infrastructure takes advantage of Intel® IoT Gateways equipped with Intel® CPUs. The gateways enable us to use automation to connect devices and data, resulting in an efficient, data-driven factory environment that embodies smart manufacturing.

We use this infrastructure in Intel's assembly and test factories to predict pump motor and blower failures and detect material flow leaks. These three use cases alone are saving Intel hundreds of thousands to millions of dollars. In the near future, we plan to deploy at least nine more IIoT use cases in assembly and test facilities, including using laser shaft analysis to detect misalignment and to prevent toxic odors from affecting factory workers and operations.

In the subfab environment beneath the main fab floor, the use of the IIoT infrastructure goes beyond mere equipment monitoring.

The Intel IoT Gateway infrastructure enables closed-loop control to collect data from the equipment and send instructions back to the equipment as needed, such as to shut down a tool or sound an alarm. Our Intel IoT Gateway-based solution is 10x less expensive than the proprietary vendor solution we had been using. Reduced scrap and less factory downtime can provide additional savings—up to USD 16 million per factory.



Augmenting our IIoT-based equipment monitoring capabilities, we also worked with Intel's Virtual Factory experts to create a powerful analysis method that lets managers quickly traverse an entire factory or several factories to identify underutilized equipment. Intel's high-volume factories involve thousands of tools, each running from one to 50 operations, which equates to tens of thousands of tools and operations needed to monitor, analyze, and adjust. Identifying tools that are operating below their optimal capacity and finding the root cause is vital, but knowing where to focus engineering efforts is a monumental challenge.

Explore

- [Transforming Intel with AI](#)
- [Autonomous Quality in AI Model Productization](#)
- [Computer Vision and AI for Inline Inspection](#)
- [Expanding Low-Cost IIoT Manufacturing Use Cases](#)
- [Fault Detection for Factory Equipment](#)
- [Improve Manufacturing Yield Analysis with AI](#)
- [Analytics and Factory Equipment Runtime](#)
- [Using ML to Characterize Database Workloads](#)
- [Minimizing Manufacturing Data Management Costs](#)

The team built upon an existing platform to create and implement this new data analysis *within 48 hours of the initial request*. Twenty-four hours later, the second version was released to production, and the final version was released with all enhancements ready for "[Copy Exactly](#)" distribution within another 72 hours. The analytics solution rapidly increased production at one of Intel's factories, generating significant value in the first quarter of usage. Manufacturing managers and engineers can use this analytics solution to instantly identify areas of need in a production line and develop solutions to maximize tool availability, which can increase business value for every facility around the globe.



Exploring the Future of GenAI at Intel

Intel IT is building a Generative AI (GenAI) framework that has already delivered value to Intel's business units. The ultimate goal is to provide a unified One Intel GenAI platform for enterprise-wide application, and we anticipate that 75% of Intel employees will use this GenAI platform in their work. Recently, the number of platform users grew from 500 to 1,600 in a single month. As the number of users grows, we are also establishing governance processes to help ensure that the GenAI platform will be used responsibly and safely. Below are three GenAI projects from Intel's AI Solutions Group:

- **Virtual Technical Assistant (VTA)** uses large language models to provide question-answering and reasoning capabilities. The VTA derives information from multiple data sources and scales the dissemination of technical knowledge. This enhances both seller and customer experiences by quickly providing accurate answers to technical inquiries, which accelerates the design pipeline and increases efficiency and productivity.
- **Co-Design** brings the power of GenAI technologies to Design Engineering processes and tasks. Co-Design will help with tasks such as answering questions about Intel's technical design documents; generating, explaining, and debugging code; and creating validation and test artifacts like assertion rules and test programs.
- **Account Pulse** is a new AI capability in Sales CRM that can help sellers to save time when managing their accounts. Powered by large language models, Account Pulse summarizes the seller's accounts and opportunities based on Sales CRM data. Combined with knowledge sharing and decision support, Account Pulse's improved visibility of customers allows sellers to reduce the time they spend writing and editing summaries for stakeholders within and outside the Sales organization.



Infrastructure

Building a Cloud-to-Edge Infrastructure

Intel IT works closely with Intel's business units to understand their objectives and processes so we can identify the solutions they need to achieve their goals. A key aspect of Intel IT's cloud strategy is to place workloads where they make the most sense. This may include using several public cloud providers, while keeping some workloads in our private cloud so we can process data at the edge to decrease latency and bandwidth usage for latency-sensitive workloads.

Our Design Engineering data center environment provides the vast majority of our installed data center capacity. This internal hosting of data center workloads has lowered our unit costs by 3.36x. By applying innovative data center techniques to our internal hosting, we have achieved unit-cost levels that are significantly lower than if we had hosted all of our workloads using a public cloud infrastructure. Our workloads and ability to achieve high server utilization are particularly well suited for private cloud investment. From 2010-2022, we have garnered combined capital and operational savings in excess of USD 7.5 billion, which helps fuel our continuous innovation cycle. Another benefit of our Design Engineering data center environment is its

extreme efficiency—running at a Power Usage Effectiveness (PUE) of 1.06. This not only lowers Intel's energy bills, but also helps meet Intel's sustainability goals.



Edge computing, such as using computer vision analytics to detect product defects, can help lower latency and bandwidth usage. This helps to prevent overloading Intel's networks.

Even with the advantages of internal hosting and the power of edge computing, we also foster innovation and agility by continuing to use the public cloud (see the "[The Expansion of the Intel® Developer Cloud](#)").

Scaling SDN and Automation

Demand for data center and manufacturing network capacity will increase as Intel's business evolves, and we'll need even more capacity as Intel's foundry business ramps up. Additionally, business pressures often require new capacity to be available in production environments within 24 hours. We have long recognized the potential of software-defined networking (SDN)

to help meet these challenges and have deployed SDN in our data centers and in Intel's factories. SDN's flexible, open interface makes it easy to integrate additional business-driven automation to meet our growth and timeline needs. Our data-driven network architecture and strategy are intentionally created to help provide the performance levels and network availability that Intel, and ultimately Intel's customers, require to be successful.



"Our concept was to use data center technologies and bring them to the manufacturing floor. We've had to swap the [networking infrastructure] that exists, which is classic Ethernet, and put in SDN. I've upgraded a whole factory from one code version to another code version without downtime for factory tools."

– Robert Colby, Principal Engineer,
Intel IT Infrastructure

CIO100
A W A R D S

Data Center SDN

We are deploying a new data center SDN architecture that uses a leaf-spine underlay with overlay networks. Our adoption of SDN and automation architecture has provided numerous benefits to Intel:

- Reduced network provisioning times for entire server racks from 8 hours to 2 hours.
- Improved reliability as performance-related incidents were reduced by 70%.
- Improved efficiency by increasing data center workload capacity by 25% without adding staff.
- Increased network flexibility by adopting open, standards-based SDN and orchestration layers.

Manufacturing SDN

Our groundbreaking deployment of SDN for the factory network garnered Intel the [CIO 100 award](#) in 2023. SDN enables us to complete new factory network builds faster, with 85% less headcount by using automated scripts. Cisco Application Centric Infrastructure switches powered by the Intel Xeon processor D-1500 and D-1600 families contribute to our fast factory networks, with 100x more efficient on-switch memory utilization for security implementation compared to the classic factory network.

Explore

- [Our Data Center Strategy](#)
- [Data Center Networks with SDN](#)
- [Transforming Industrial Manufacturing with SDN](#)
- [Data Center Facilities Risk Management](#)
- [Cloud CaaS – Wins, Trends and Strategies](#)

The Expansion of the Intel® Developer Cloud

Intel's Software and Advanced Technology Group (SATG) embodies Intel's vision for software and advanced technology by delivering open and secure software platforms, driving industry standards, and ensuring that software is a powerful and competitive differentiator for the company. An important aspect of SATG is the availability of developer cloud services, which can change how we engage with internal and external software development ecosystems.

In 2022, SATG worked with Intel IT to transform Intel® Developer Cloud (DevCloud) to be more scalable and broader in scope. The infrastructure supporting today's Intel DevCloud was developed by several Intel IT teams who determine which data centers could optimally host the services, set up the network, and manage load balancing.

The expanded Intel DevCloud, first available in 2023, provides access to various Intel® Toolkits and libraries; AI frameworks and tooling; and AI foundation models. Internal and external

developers can use these resources to accelerate AI development using Intel® architecture-optimized software on the latest Intel® Xeon® processors and GPU compute. By the end of 2023, the Intel DevCloud was being used by 10,000 unique active customers every day. DevCloud is akin to a digital Galapagos Archipelago; it's a natural, unique environment for accelerated IT evolution.

With the Intel DevCloud, software engineers can get hands-on experience with the latest Intel® products and build AI skills, access and evaluate pre-release Intel platforms and associated Intel architecture-optimized software stacks, and accelerate AI deployments with the latest machine-learning toolkits.

Intel IT development teams use the DevCloud to test and develop solutions using the latest Intel® hardware and optimized software without downloading or configurations. In other words, Intel IT built the Intel DevCloud, and now we're using it to "engineer our way into the future."

Security

The Continuing Evolution of Information Security

As cyber threats evolve, Intel's Information Security (InfoSec) group adapts to the new threat landscape. InfoSec's "defense in depth" strategy uses automation and security orchestration to detect and mitigate 99% of known cyber threats. The InfoSec group can then focus on hunting down the remaining 1% before they can harm Intel or its customers.

InfoSec's continuous improvement of Intel's security position helps keep Intel secure and make it safe for Intel to go fast.



Modernizing IAM

Intel's business units and employees interact with many internal and external B2E, B2B, and B2C entities. Hundreds of applications and portals require users to sign in, which involves extensive identity access management (IAM).

Intel IT's IAM team designed and deployed a scalable, resilient microservices-based identity-as-a-service platform to seamlessly migrate more than 200 applications from a legacy single-sign-on platform to a public cloud-based identity provider. This transformative initiative harnessed the agility and scalability of microservices to accelerate application migrations, decrease maintenance efforts, and optimize resource utilization.

The success of the identity-as-a-service security transformation spawned from a "One Intel" partnership between InfoSec, Supply Chain IT, IT Corporate Functions, and business units like the Sales and Marketing group.

This project delivered several benefits:

- Reduced technical debt across all Intel business units by unifying dozens of login tools that had proliferated over more than a decade.
- Enhanced Intel's overall security posture while allowing new applications to launch quickly.
- Made it easy for end users to log in.
- Reduced system maintenance, freeing IT staff to work on future innovations.

Because of its positive impact, the project received an internal Intel award for outstanding achievement.



One Intel



Advancing Intel's Security Posture

While much of our InfoSec technology is on-premises, InfoSec has also adopted cloud-based solutions. The evaluation of CrowdStrike began about six months prior to the onset of COVID-19, primarily for endpoint detection and response and next-generation antivirus protection. When the pandemic sent 130,000 Intel employees to work from home using 171,000 endpoints, InfoSec accelerated and expanded its adoption of CrowdStrike technology, condensing a deployment that might have taken months to three weeks.

Intel co-engineered a solution that combines CrowdStrike's hardware-enhanced exploit detection technology and Intel® Threat Detection Technology (Intel® TDT) to proactively detect fileless malware and other sophisticated attacks that may be hiding in memory. This new capability offloads the scanning of memory to the GPU or the graphics card that's integrated with Intel's SoC, enabling memory to be scanned without impacting the performance of the actual compute platform. This solution helps prevent even more malicious activity on endpoints, improving the focus on the most important threats.

Securing Our Global Manufacturing Supply Chain

Today's global semiconductor supply chain is digital and complex; it exists in an environment where manufacturing resilience drives increased investment of time and technology. Intel is accelerating the adoption of accountability standards across industries through leadership and collaboration with its suppliers, stakeholders, consortia, and fellow travelers.


From inbound materials to product design and development, manufacturing processes, and outbound finished goods, the entire supply chain lifecycle presents risks that must be mitigated. Intel must be able to attest to the security of its products and software. The InfoSec group recommends a layered, in-depth defense strategy for securing the global manufacturing supply chain. One important aspect is for Intel and other semiconductor manufacturers to adopt two standards that can improve a company's manufacturing risk posture: SEMI E187 (Specification for Cybersecurity of Fab Equipment) and SEMI E188 (Specification for Malware-Free Equipment Integration).

In addition, Intel is a member of the SEMI Cybersecurity Consortium and encourages other manufacturers to join.

The semiconductor future is bright, and Intel is poised for success. More companies may benefit by following Intel's supply chain security efforts and working to develop a new level of structured cooperation to address common challenges.

"We now see CrowdStrike as more than a strategic supplier. CrowdStrike is a key innovation and collaboration partner in our fight against cyber adversaries."

— Brent Conran,
Chief Information Security Officer



Explore

- Advancing Intel's Security with CrowdStrike
- Protectors Spotlight



Intel: America's Most Cybersecure Company

Consider these disturbing statistics: As of 2022, the manufacturing industry saw the highest share of cyberattacks³ among leading industries worldwide; in 2023, the global average cost of a data breach was USD 4.45 million.⁴

Intel's Information Security (InfoSec) team keeps Intel's cybersecurity as comprehensive and effective as possible. From investing in a data lake that can help detect anomalies in network behavior, to automating and accelerating threat responses, to training employees who are the company's first—and best—defense against many cyberattacks.

Not only is Intel's approach to cybersecurity effective, but it's also getting noticed. Forbes named Intel the "Most Cybersecure Company in America" out of 200 ranked companies. The ranking considered a wide range of factors, such as network and application security, malware vulnerability, regularity of patches, robustness of cybersecurity personnel—and even hacker chatter about possible exploits. Also considered was whether their performance in various areas increased or decreased over time.

³ Statista, February 2023, "Distribution of cyber attacks across worldwide industries in 2022."

⁴ IBM, "Cost of a Data Breach 2023."



Productivity

Maximizing Client Productivity and Optimizing User Experience

We're committed to providing Intel's employees with the best possible IT experience. We support an inclusive, diverse, and flexible environment so employees can maximize their productivity wherever they work.

We aim to deliver a consistent and highly mobile-capable experience that improves productivity, creates a seamless experience across platforms, and delivers a stable workspace. Recent initiatives include using client health telemetry to proactively identify issues and inform a "smart refresh," embracing Windows 11, and exploring advanced support technologies such as chatbots.

Gaining Insight to Client Health Using Telemetry

If an employee's PC crashes or exhibits substantial performance degradation, that directly affects the employee's ability to do their job of helping Intel thrive. Intel IT is committed to predicting and preventing as many client issues as possible and providing quick and easy access to IT support. Over the last couple of years, we have significantly increased our use of client health telemetry to gain insight into device health and provide fixes and recommendations to PC users. We can also use the telemetry data to optimize PC refresh based on actual client health instead of a rigid, calendar-based approach to PC refresh.


Digital experience management (DEM) tools with predictive telemetry enable Intel IT to go beyond simply fixing employee PC issues, and instead prevent them from occurring in the first place.

IT administrators can use a variety of tools to analyze telemetry data, interpret it, and create action plans.

Our use of telemetry data is delivering impressive results:

- In 2023, the automation and self-healing aspects of the DEM platform and IT Chatbot capabilities helped avoid more than 18,000 end-user IT incidents and provided 135,000 hours of productivity savings through self-healing capabilities and user-assisted fixes.
- In 2022, customers' satisfaction with their PCs jumped by 27.4% compared to 2020 levels.
- Employees can self-select the best-fit device based on personas.

Embrace redefined, interactive worker experiences and advanced client health telemetry




Taking Advantage of Windows 11 for Optimized PC Performance and User Experience

Historically, we have used manual, time-intensive IT processes to help ensure PC users could easily and efficiently upgrade their OS. This manual approach included 9-12 months of designing and testing the custom interface. Using Windows Update for Business (WUfB) to upgrade to Windows 11 simplified and automated the upgrade process without creating a new interface. Once application readiness testing and the pilot were complete, we rolled out most Windows 11 in-place upgrades in 13 weeks.

By upgrading to Windows 11 and taking advantage of cloud-managed, native upgrade features such as WUfB, we can:

- Maximize user experience with the hybrid architecture available on 12th Generation and newer Intel Core processors.
- Deliver a stable and reliable platform.
- Enable high-velocity OS upgrades (13 weeks instead of 9-12 months).
- Reduce technical debt by retiring custom upgrade tools and processes.
- Decrease help desk calls by 80%, compared to the Windows 10 21H2 in-place upgrade.

13-Week Deployment Strategy for Upgrading to Windows 11



Optimizing the Employee Experience with AI

Intel IT constantly pursues new ways to enhance our user support capabilities. We have begun evaluating how AI can enhance productivity by using the Microsoft Copilot AI-powered chatbot, which is a complimentary feature in Intel's Microsoft contract. The AI-powered chatbot uses commercial data protection to summarize, analyze, extract, and visualize information from web pages, PDFs, and other information sources. Our goal is to make the chatbot accessible on the web for all employees in Windows 11 through the Edge browser. This chatbot could potentially help employees perform their daily tasks—such as creating, editing, and summarizing content—more efficiently. For example, a simple prompt could collate relevant information from across the organization or quickly summarize meeting outcomes.

AI-based chatbots that use NLP and Generative AI offer an opportunity to mimic conversations with real people in both voice and text formats. This ability could provide employees with improved self-help IT options. We plan to evolve our former rule-based chatbot technology to use self-learning/self-understanding dynamic language models to reach IT support resolutions more quickly.

Chatbots could also help optimize processes and workflows for support ticket and service desk data. For example, based on previous successful ticket resolution data, the chatbot could identify the best IT staff member to provide in-depth technical support depending on the specific topic. Or, the chatbot could perform a detailed review of support tickets to identify application or environment problems previously hidden by incorrect classification. In the future, we plan to explore deploying voice-enabled chatbots (voicebots). Our adoption of these new technologies will be balanced with careful consideration of privacy, security, and ethical compliance.

Explore

- [PC Client Telemetry Insights](#)
- [Improving Thunderbolt Intelligence](#)
- [Enterprise-Wide Windows 11 Upgrade](#)
- [Benefits of Windows 11 OS Upgrade in the Enterprise](#)

Sustainability

Driving Sustainability through Technology

Sustainability continues to be a major concern for companies, with increasingly stringent regional mandates going into effect. Intel has steadfastly committed to a sustainable future and has taken action through systems change, technological innovation, and global collaboration. As part of that commitment, Intel is driving sustainability requirements into every product across the company.

Intel conducted a global study to see how technology leaders can impact sustainability as part of their digital transformation. The study revealed that CTOs are uniquely positioned to drive positive change for both their companies and the environment. In fact, 84% of CEOs and chief sustainability officers agree that CTOs have the potential to become the greatest drivers of sustainability in an organization.⁵

⁵ Intel, June 2023, "Introducing the Sustainable CTO."

Intel launched the Sustainable CTO initiative to help CTOs and other technology leaders navigate technology for the future of the business. This initiative will apply the full potential of technology to tackle increasingly intricate and interconnected global dilemmas that will gain prevalence in the coming decade and beyond.



"The new role of the sustainable CTO will play a crucial part in an organization's transition to net zero, but CTOs can't navigate the challenges they will face alone. We aim to help CTOs on their sustainability journey and outline the roadmap for organizations to achieve their sustainability goals."

— Motti Finkelstein, Corporate VP, Chief Information Officer

Going Green

Intel IT is responsible for Intel's global technology infrastructure and services, and it's our obligation to reduce our environmental impact and help Intel attain its sustainability goals. Here are some Intel IT data center achievements:

- Highly efficient world-class PUE.⁶
- Recycling water to cool data centers.
- Deploying fuel cells as an alternative energy source.
- Reusing data center heat for office spaces and manufacturing processes.

The Leadership in Energy and Environmental Design (LEED) organization recently awarded

⁶ Intel, August 2021, "Green Computing at Scale."

Intel the **LEED Gold certification** for two HPC data centers, recognizing Intel IT's efforts to address carbon, energy, water, waste, transportation, materials, and health quality concerns.



Intel IT is proud to pursue sustainability through our operations and technology. We aim to reach net-zero greenhouse gas emissions, net-positive water use, and zero waste to landfills by **Intel's stated deadlines**.

We are also dedicated to helping Intel's customers and partners use our technology to improve sustainability. Developing energy-efficient products with sustainable manufacturing processes can help Intel empower others to reduce their energy use and carbon emissions.



Our experience enables us to provide feedback to regulatory bodies such as the EU Energy Efficiency Directive (EED) and help guide industry-best practices. Some of Intel IT's leaders are also members of the Green Software Foundation and are active participants in the Climate Neutral Data Centre Pact.

Sustainability is integral to Intel IT's mission of building a better future through computing. We believe in the potential for technology to help create a more sustainable world. And we are determined to lead by example to drive sustainability not just through our words but also with our concrete actions. We encourage others in the industry to join us in this critical effort.



Lowering Carbon Footprint through Application Optimization

Intel IT is using Intel® Granulate™ software to optimize the performance of our Cloudera and Databricks deployments. Intel Granulate software autonomously optimizes application runtimes, thereby improving application performance at the core, which leads to more efficient use of existing infrastructure, reduced CPU utilization, and lower compute costs. When running at peak performance, applications reduce CPU utilization and memory, leading to a minimized carbon footprint and up to 45% cost reduction.⁷

⁷ <https://granulate.io>

Exploring Fuel Cell Technology to Augment the Power Grid at Data Centers Globally

In several proofs of concept, fuel cell technology has proven to be a reliable, stable, efficient, and sustainable source of electricity for Intel's data centers. Fuel cells help eliminate the need for backup generators or UPS and augment Intel's sustainability initiatives. Additionally, the total cost of ownership for fuel cells is lower than traditional backup systems. Intel IT joined forces with Intel's Corporate Services and Global Supply Chain groups to deploy fuel cells to meet Intel's redundancy and capacity expectations.

For some of Intel's data centers, fuel cells are the primary power source (as opposed to the power grid). To date, we have installed fuel cell

technology at Intel data centers in Bangalore, India, as well as Folsom and Santa Clara in the U.S. Fuel cells are now part of Intel IT's plan of record for new data center design and expansion scope.

Our experience with fuel cells has yielded many benefits, such as reliability, scalability, affordability, lower carbon footprint than traditional UPS batteries and diesel generators, and better data center PUE.

We have continued to extend our use of fuel cell technology as an alternative energy source for data centers. Today, we use fuel cells in both tier-1 and tier-3 data centers, with plans for an additional fuel cell capacity of 32 MW in our Santa Clara facility.

To further improve our carbon footprint, we are exploring the use of using hydrogen instead of natural gas to power our fuel cells.



Fuel cells at Intel, Bangalore

Explore

- [Sustainable CTO Initiative](#)
- [Fuel Cells as Energy Source for Data Centers](#)
- [Intel IT's Data Center Strategy](#)



Innovation

Reinvigorating a Customer Zero Mindset

In 2022, Intel CTO Greg Lavender pulled Intel IT under his leadership umbrella. He saw IT as a partner and early adopter—a “customer zero” that innovates, delivers real value, and will help drive the company into the future. With this in mind, the IT Practitioners Lab (ITPL) was formed.

Our Vision

The ITPL’s vision is to provide user feedback throughout the product lifecycle to help improve quality and expedite adoption into the marketplace. This initiative involves internal users reviewing early product concept proposals, evaluating pre-launch technologies, and ultimately adopting Intel’s products, services, and technologies into our production environments. Customers will have real, tangible proof points from IT’s experience with new Intel offerings that can aid them in their own evaluation and adoption. Unlike external customers, the ITPL is a safe harbor—Intel IT can provide practitioner-level feedback, and the business units can share ideas much earlier without worrying about IP leaks.

For new products, business units can engage early with the ITPL to get help formulating design and modifying feature requirements to make the product desirable, easier to adopt, and more valuable. The ITPL also assists with proofs of concept to demonstrate how Intel IT delivers value after implementation. As products mature, the ITPL can check the feature roadmap and help verify that new features are useful and will add value.



Business Outcomes

ITPL’s positive business outcomes have enormous potential, including increased product value through improved features, better usability, and ease of implementation; fewer post-launch customer problems by identifying issues earlier in the product lifecycle; and sales and marketing enablement through Intel IT proof points.

Intel IT has always worked with Intel’s product teams to test new Intel® hardware and provide

feedback. But now, with Intel’s renewed focus on software, we are taking a more programmatic approach to our technology evaluation. In 2023, the ITPL evaluated both hardware and software, focusing on trust and confidential computing (security); AI software and services on Intel® Developer Cloud; and data center performance and optimization.

Plans for the future include continued expansion of the ITPL’s scope and services as well as driving additional revenue and product quality for Intel.




Optimizing Application Efficiency

Intel Granulate software is an autonomous workload optimization solution that continuously profiles and learns specific characteristics of each job, and then optimizes resource scheduling. The result is increased efficiency and improved throughput of big data workloads—both on-premises and in the cloud.

The ITPL conducted two proofs of concept with Intel Granulate software—one with our on-premises Cloudera platform and one with our Databricks platform in the cloud.

Intel Granulate Software and Cloudera

For Cloudera, during a week of testing in our production environment, Intel Granulate software optimized more than 1,000 unique jobs—equivalent to approximately 42,000 executions of Apache Spark jobs and applications. We saw an average of 25% reduction in memory utilization and an average of 44% reduction in CPU utilization.⁸



44% Better
CPU Utilization

⁸ Testing configuration information can be found in the white paper, "Granulate™ Optimizes Memory and CPU Utilization in Intel IT's Cloudera Platform."




Intel Granulate Software and Databricks

For Databricks, during a four-week proof of concept, we achieved a 23% improvement in vCore utilization, a 17% increase in data throughput, and reduced CO2 emissions.⁹

Based on these results, we will gradually deploy Intel Granulate software across our remaining Databricks environment to gain its benefits, such as running more jobs concurrently without adding memory or compute infrastructure. Additionally, as our platform grows, we can avoid capital and operational expenses in the future.

⁹ Testing configuration information can be found in the white paper, "Granulate Optimizes vCore Utilization & Data Throughput in Intel IT's Databricks Platform".

Databricks Proof of Concept Using Intel Granulate Software



Looking Ahead

Delivering the Next Evolution of Intel

As Intel charts a new course into the future, Intel IT must continue to evolve to support the changing needs of the organization. To achieve that, we will capitalize on new technology innovation and continue to build a solid foundation for Intel's delivery of the next evolution of its IDM 2.0 model. The main key to ensuring such success is having skillful people so that we can execute at a torrid pace, since nothing can happen without the best talent.

We made transformative progress in 2023, creating efficient operations and making successful inroads to transition to the internal foundry model, and we will continue to build on that momentum. In 2024, we will expand our AI capabilities into all core business processes with a focus on creating impeccable quality and reliability of AI systems and other opportunities as they present themselves. We will use our strategic data assets to unleash more accurate analytics and intelligent decision-making to enable our strategy while cost-effectively scaling our global business operations.

We will focus on maturing, scaling, and integrating our solutions and services. All of us across IT must strive to approach our daily work and ambitious goals with an innovative mindset, which will enable us to provide leadership that can rapidly actualize Intel's digitization and transformation.

Our success in enabling Intel's business depends on a tight connection and trust with the business units. We will *know* what the business units want and will use our expertise and influence to help them prioritize projects. Once priorities are set, we will provide solutions and technology that support the business units' outcomes and needs.


Finally, we will focus on operational excellence. We must deliver on our commitments and do it well, and we will continue to make our systems and processes more resilient.

I remain confident in our ability to deliver on IT's critical evolutionary role and to exceed the expectations of the enterprise in this year and beyond. The same "One Intel" mindset that propelled us through 2023 will continue to drive Intel forward as an iconic company with a long and prosperous future.

What are your own goals and what type of future are you evolving toward? [Let's start a discussion.](#)

Intel CIO Motti Finkelstein quotes
Winston Churchill

"To improve is to change; to be perfect is to change often."



Stay Updated

Keep an eye on the future by joining us on our journey at intel.com/IT.

Read More

Learn more from our IT leaders and subject matter experts in our [Thought Leadership community](#).

THE INFORMATION PROVIDED IN THIS PAPER IS INTENDED TO BE GENERAL IN NATURE AND IS NOT SPECIFIC GUIDANCE. RECOMMENDATIONS (INCLUDING POTENTIAL COST SAVINGS) ARE BASED UPON INTEL'S EXPERIENCE AND ARE ESTIMATES ONLY. INTEL DOES NOT GUARANTEE OR WARRANT OTHERS WILL OBTAIN SIMILAR RESULTS. INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS AND SERVICES. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS AND SERVICES INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

All rights reserved. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation in the U.S. and/or other countries. Other names and brands may be claimed as the property of others. © Intel Corporation 0424/WWES/KC/PDF