intel

2024-25 Intel IT Annual Performance Report

IT – The Resilient Change Agent

Insights From Our CIO

Winston Churchill once said, "To improve is to change, so to be perfect is to have changed often." This recognition of continuous improvement, and that change is an integral part of progress, resonates with me now more than ever as Intel charts a new course into the future. It is no secret that we have faced some challenging times recently. Yet, through this, Intel IT has maintained our path and adapted with creativity and collaboration.

Good leaders steer through uncertainty and navigate difficult moments with clarity and purpose, while appreciating the teams and strategies that drive success. We must stay open to new opportunities, learn from every experience, and believe in our ability to be resilient—to adjust and succeed. "The secret of change is to focus all of your energy, not on fighting the old, but on building the new."[1]

Resiliency is essential to enterprise success, and IT plays a critical role in building it across technology, operations, company culture, and the workforce. It is more important than ever that we listen to our partners, Intel's business units (BUs), to enable them to develop the best products for Intel's customers. My goal for IT is to deliver shared value through a One Intel mindset characterized by active listening, financial prudence, and risk management.

Intel IT has worked hard to forge critical alliances with enterprise stakeholders while adhering to strategies and operations that prioritize the organization's overarching goals. IT is now at a place where we are trusted advisors. We must hold ourselves accountable to the commitments we make so that we continue to build on that trust.

This strategy isn't new; we continue to maintain our focus on innovation, AI, and optimization as well as the IT experience of Intel's workforce and customers. We're always exploring new ways to do things more efficiently, both in the business and in IT.

We accomplished a great deal in the past year, such as enabling more than USD 4 billion in IT business value. In 2025, we have an opportunity to build on what we have achieved; we will maintain our One Intel mindset, values, and principles as we move forward. They will guide us as we continue building our strong business partnerships and deliver on commitments. Focus areas include critical priorities like AI as well as operational excellence through self-service, automation, and validation while pacing the change to our investments. Ultimately, our efforts are building a resilient workforce with the right leadership and skill sets to help build the new Intel.

Motti Finkelstein
Corporate Vice President and Chief Information Officer

[1] Dan Millman, "Way of the Peaceful Warrior"

In this issue

~2,607 IT Employees **SUPPORT** ~100K Employees **AT** 55 Intel Sites **IN** 53 Countries

Infrastructure

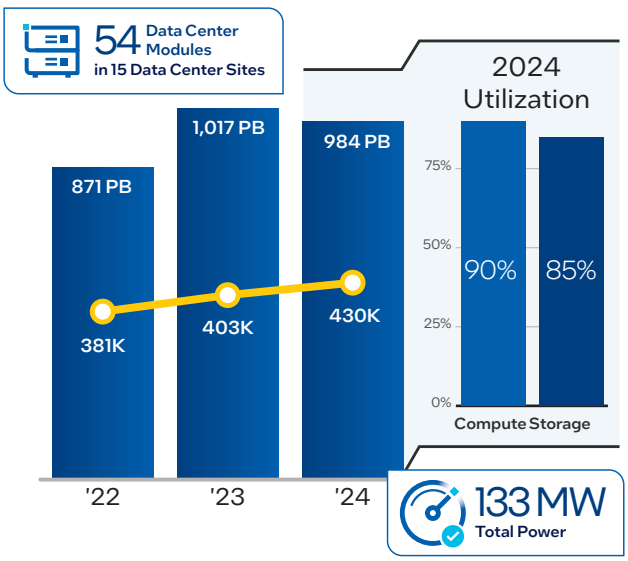

IT Spending

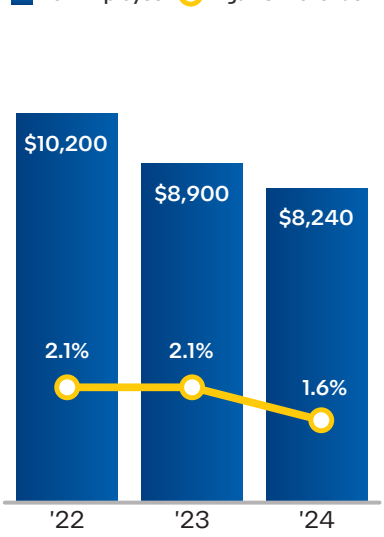

IT Managed Devices

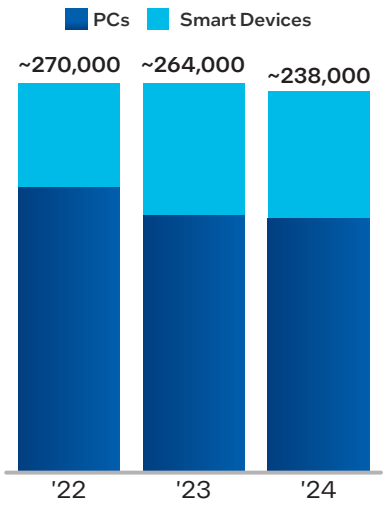

Our IT Environment

>30K
Monthly users and >7K personal assistants created on iGPT marketplace

>USD 133 M
In tech debt cost avoidance

~1.4 M Devices
Monitored for security compliance

USD 9.6 B
Data center strategy cumulative cost savings

USD 4 B
Business value enabled

19-Week Reduction
Pre-silicon design cycles

Data as of January 1, 2025. Employee count represents an average of beginning of the year.

Enabling the Business to Thrive

Our highest priority is to empower Intel's product leadership and its Foundry business. Through an increasingly tight connection with Intel's business units (BUs), we've built strong business relationships and have engaged in several "customer zero" projects to accelerate market adoption of Intel's products. In support of Intel's digital transformation, we are focused on harnessing AI for bold business opportunities and delivering high-quality, accessible data.

30,000
Monthly Users

7,000
Personal Assistants

Delivering the Power of AI to Employees

We are in the early stages of realizing the full potential of GenAI to streamline processes and enhance efficiency at Intel. The expansion of our "AI Inside" initiative across the company has already produced significant improvements in productivity, quality, and asset utilization, and we're just getting started.

Our AI strategy follows a three-fold approach. We start with foundational tools, drive the utilization of that foundation across the business to deliver value, and then optimize the value of AI by teaching employees how to use it effectively.

Providing Foundational Tools for Responsible GenAI Innovation

We have developed an enterprise-grade, scalable GenAI platform. The platform consists of iForge, the backend engine that runs and connects models and data sources, and iGPT, the frontend that provides users with intuitive no-code AI tools so that all employees can build and use AI assistants for specific use cases. IT and BUs can use the platform to build what they need to solve problems and quickly adapt to a rapidly changing business landscape.

Optimization
Unleash the workforce to revolutionize Intel

Utilization
Secure the environment supporting the business

Foundation
Provide the right tools

Already, 30,000 monthly users have created more than 7,000 personal assistants and hundreds of GenAI use cases, including productivity, cost reduction, cost mitigation, and incremental revenue. By funneling GenAI use cases into our enterprise GenAI platform, we help to ensure safe, responsible use of AI using approved, internally built capabilities.

Now we are on a mission to infuse agentic AI across the company. Use cases include improving customer experience, automating legal contract reviews, and creating a OneAI bot that can easily adapt to a wide variety of purposes. We are incorporating additional knowledge assistants that can quickly answer "How Do I..." questions. Additionally, our system architects are upskilling to support Intel's transformation with AI-powered solutions and architectures—all based on one standard AI architecture.

Intel IT is driving a strategic shift toward leveraging Intel® silicon products for internal efficiency and external market presence, for both GenAI and classic AI applications like machine learning. We are expanding the use of Intel® Xeon® processors in our retrieval-augmented generation (RAG) systems while investigating Intel® Gaudi® accelerators for GenAI. Intel's factories are migrating GPU-based classic deep-learning solutions to affordable, readily available Intel Gaudi accelerators for model training.

Identifying Use Cases with High Business Value

We use an AI dashboard to clearly track all AI use cases across Intel. The dashboard has two purposes:

- Prioritize the highest-value projects.
- Help guide decisions about new capabilities that can be used across multiple use cases.

Intel IT AI Dashboard

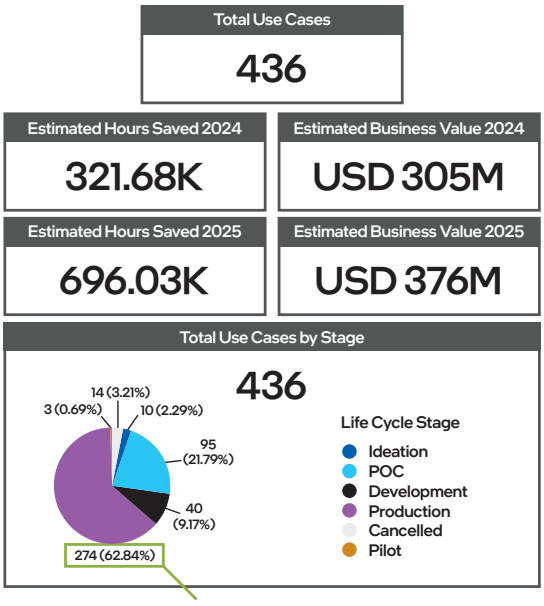

274 AI use cases are currently in production

In other words, the dashboard helps us achieve managed, intelligent growth of AI. Our 2025 focus areas are developing a roadmap for future capabilities; defining additional use cases; providing data that is reusable, vectorized, and optimized; and driving employee awareness.

Factors that affect AI use case approval include cost of implementation, redundancy with other solutions, usage, and value. The Information Security and Privacy teams are also involved in the approval process, using a secured governance model to help ensure AI solutions built and used inside Intel follow our security, privacy, and other regulations for responsible AI requirements.

Infusing AI into Enterprise Architecture

For enterprise architecture (EA), the integration of AI and GenAI presents a transformative opportunity to automate processes and enhance productivity. AI and GenAI can automate tasks across EA phases, from initial architecture vision to implementation governance.

We are currently piloting an AI EA agent that can discover and summarize existing architecture artifacts—such as policies, principles, standards, business process flows, and solution architecture diagrams—making complex architecture documents more accessible to both technical and business audiences.

Our goal is to standardize the diverse formats of architecture artifacts to enable effective image and visual inferencing. We will be able to upload a design document and receive automated feedback to improve the quality and consistency of architectural deliverables, as well as connect architecture artifacts to application portfolio and issue management systems. Eventually, automation can extend to architecture governance processes through self-service capabilities.

By embracing AI for EA, we can enhance the productivity of our architects and help ensure that deliverables are more accessible and interpretable to business stakeholders, ultimately driving greater alignment between IT and business needs.

Instilling a Culture of AI Utilization and Innovation

We've built the foundational tools and identified priority use cases; now it's time to expand employees' knowledge so they can use the tools in their daily work. Our AI Everywhere program is driving culture, awareness, and training across the company.

We use a survey to determine the baseline of employees' AI awareness and savviness then offer AI training, including meetups and training sessions. Our trainings present a message of reassurance to employees who may have reservations about this emerging technology. To quote our CIO, "GenAI is designed to empower employees in their jobs by enhancing decision-making, productivity, and innovation. It's a collaboration tool and no substitute for human experience, judgment, or creativity."

Explore

- IT@Intel: Democratizing the Use and Development of Generative AI Across Intel
- Scaling GenAI: How Intel IT Is Building an Enterprise-Ready Generative AI Framework
- IT@Intel: Training Machine-Learning Models on Intel® Gaudi® Accelerators

Rolling Out AI PCs to Support AI Inside

We foresee the proliferation and adoption of many AI apps and features in the next 18 to 24 months. This includes our GenAI platform, our AI Assistant Builder, and many AI capabilities that are being integrated into the commercial software applications that employees use every day.

However, if the PCs in our environment struggle to run AI workloads, then having AI tools at the fingertips of our employees isn't going to enable AI to reach its full business value potential.

As part of our regular PC refresh cycle, we have deployed 25,000 AI PCs. Unlike traditional laptops, AI PCs are equipped with Intel® Core™ Ultra processors—an eXtensible processing unit (XPU). The XPU consists of a CPU, a GPU, and a neural processing unit (NPU).

Our AI PCs are delivering excellent results in the areas of security, productivity, creativity, and collaboration. Here are some of the advantages AI PCs bring to Intel's workforce:

- Employees can use GenAI application features to improve the user experience.
- Employee productivity is enhanced through AI automation of mundane tasks, such as meeting minutes and emails.
- The XPU can improve the PC's battery life, reducing the need for frequent charging while reducing energy consumption.
- AI workloads run locally instead of in the cloud, reducing latency and network load and mitigating the risk of data-in-motion security breaches.

We anticipate that by making AI PCs part of our regular PC refresh cycle, we can reduce costs. We can do this by eliminating the need for a mid-cycle refresh and reducing IT support requirements thanks to AI-driven self-diagnosis and self-repair capabilities.

IT as Customer Zero: Partnering to Deliver Customer Value

Intel's business units (BUs) see IT as a partner and early adopter—a "customer zero" that innovates and delivers real value. With deep product expertise and thought leadership in technology infrastructure, Intel IT can provide practitioner-level feedback throughout the product lifecycle to deployment, helping improve quality and expedite adoption into the marketplace. We review early product concept proposals, evaluate pre-launch technologies, and ultimately adopt Intel's products, services, and technologies into our production environments. Our experience with Intel's offerings provides customers with real, tangible proof points that can aid them in their own evaluation and adoption.

Unlike external customers, Intel IT is a safe harbor—we can provide direct feedback, and the BUs can share ideas much earlier without worrying about intellectual property leaks. For new products, BUs can engage early with us to get help formulating design and modifying feature requirements to make the product desirable, easier to adopt, and more valuable. We also assist with proofs of concept to demonstrate how Intel's products deliver value after implementation. Here are two examples of our collaboration with Intel's BUs and sharing our findings with the industry:

- Improving electronic design automation (EDA) with Intel® Xeon® processors. Intel's Design engineers run more than 273 million compute-intensive batch jobs every week, each taking from a few seconds to several days to complete. Based on our 2024 benchmarking, 4th and 5th Generation Intel Xeon Scalable processors deliver significant improvements in throughput and per-core performance for Intel® silicon Design workloads.[2] Upgrading our servers to these newer processors enables significant server consolidation, helping avoid data center construction costs and reducing power consumption. Read our white paper, "IT@Intel: Increasing EDA Performance and Throughput with the Intel® Xeon® Processor Scalable Family," for more details.
- Collaborating on Intel® Connectivity Analytics software. We worked with Intel's Client Computing product group and Cisco to develop a commercial product that addresses real IT network management issues. We use Intel Connectivity Analytics to improve the Wi-Fi network user experience and network management efficiency to expedite troubleshooting, reduce mean time to repair, and reduce overall network management costs. See the "Increasing Employee Productivity with Device Experience Optimization" section and read our white paper, "Optimizing and Troubleshooting Wi-Fi Networks Using Client Analytics," for more details.

[2] Testing by Intel IT as of January through March 2023 through February 2024. See Tables 2-4 in the white paper, "Increasing EDA Performance and Throughput with 4th and 5th Generation Intel® Xeon® Scalable Processors" for configuration details and workload run times.

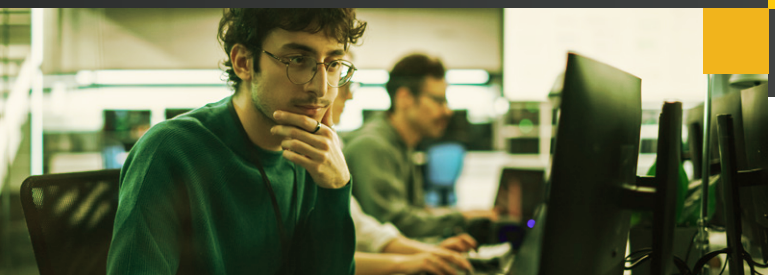

Building a Strong Data Foundation

High-quality, well-governed data supports Intel's transformation, is critical to our AI and GenAI initiatives, and is pivotal to improving operational efficiencies. Our data strategy team collaborates as a single entity across five main pillars: master data, data quality, governance, analytics, and readiness to enable enterprise-wide innovation.

Master Data for a Single Source of Truth

As Intel moves forward with its transformation, with Foundry and Products as separate reporting entities, we needed to implement a single source of the truth. We are replacing legacy systems with the SAP Master Data and Governance tool to create a centralized hub that governs material, customer, vendor, reference, and finance master data for both Intel Foundry and Intel Products. The combination of master data governance and role-based access controls helps us control data access between and within the reporting entities and provide "arms-length" data security between Intel Foundry and Intel Products.

Benefits of master data include:

- The ability to better understand the business and integrate multiple systems using consistent data.
- Harmonized and secure data for improved operational efficiency.
- Stabilized operations using real-time analytics to improve data quality and reduce manual interventions.
- Increased transparency of data maintenance bottlenecks and audit logging.

Data Quality to Engender Trust

Our data team is committed to creating high-quality data to avoid inefficient processes and missed opportunities. Our approach includes the following:

- Identify data quality issues prior to testing.
- Perform multiple rounds of data profiling to assess data, using leading data-quality tools and AI-based methods.
- Validate data integrity and fitness to business processes with system and user testing.
- Perform ongoing data quality monitoring and data observability.

Here are just a few of our success stories:

- We tied our primary data quality tool into our transactional system and started measuring data quality in only a month.
- In the second phase of our enterprise resource planning (ERP) transformation, we achieved a 98% data quality pass rate, compared to the goal of 80%.
- Less than 10% of total defects found during system testing were related to data quality issues.

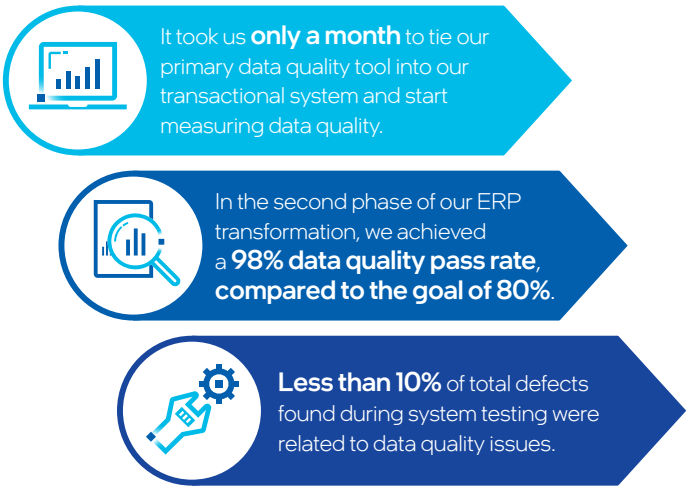

Data Governance to Turn Data into a Strategic Asset

Our approach to data governance uses a federated data ownership model that is applied across Intel Foundry and Intel Products. Our standards, methods, and tools assist in maintaining ongoing data hygiene. For example, we've created and documented 10,000 data definitions, aligned across business and IT, enabling data catalog tools.

We focus on data literacy and promoting standardization and reuse. Our Enterprise Data Governance Office engages with key leaders who are driving our transformation strategy to help ensure data is fit-for-purpose to business needs. Twenty-four data domains are aligned to support the strategy, with corresponding stewards helping to maintain the integrity of various datasets by embedding data governance into operations teams.

We are now evangelizing federated data governance to the BUs, helping them see that our tools to automate data governance enhance the resiliency of the entire enterprise. Our data governance tools also help protect the data by establishing system rights and privacy protections, so users only see information they need to know.

Data Analytics Supported by a Modern Data Lakehouse

Late in 2023, Intel IT completed the first phase of the modernization of our data analytics platform by migrating a large BU's data to a cloud-based data lakehouse solution. This transformation established a single source of truth and enabled the creation of reusable data assets.

We also developed a metadata-driven Cloud Ingestion Framework (CIF) that includes reusable data ingestion and transformation templates. This helps deliver actionable insights in a structured and governed manner and enables self-service analytics, reporting, and decision-making.

We completed the data modernization initiative within just four quarters. This complex project involved moving over 4,000 objects, supporting more than 3,000 users, and integrating data from over 100 sources. The migration delivered significant improvements, such as the following:

- 65% decrease in view processing times.
- 18% reduction of object footprint.
- 50% drop in incidents.
- 58% decrease in job failures.

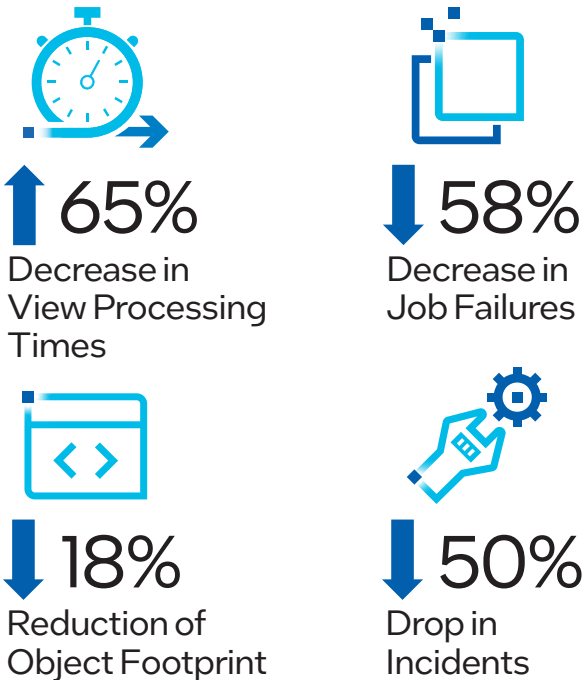

The modern data lakehouse is now enabling us to support the transformation of the overall Intel business model, establishing Intel's Product and Foundry businesses as separate entities, with a strict separation of foundry and product data. This data segregation leverages robust security measures to help protect customer intellectual property (IP) and foster trust with customers.

In support of Intel's transformation strategy, we are organizing data into three super zones that have specific security access roles:

- Intel Foundry super zone contains only foundry data and is accessible only to the Foundry business.
- Intel Product super zone contains only product data and is accessible only to the Product business.
- Corporate super zone contains shared services data like Finance and HR, and will have restricted access based on a need-to-know basis.

This granular access control helps ensure that the business can comply with regulatory requirements, enhance data governance, and commit to data security-critical factors to maintain customer confidence and support secure, collaborative analytics environments.

The data lakehouse provides the following business benefits:

- Accelerates onboarding of new use cases.
- Enables cross-domain analysis while allowing for domain-specific customizations.
- Equips business users with the ability to define, maintain, and apply business logic.
- Minimizes time to insight for better operational decision-making.
- Provides a catalog of crucial details about our data.

Data Readiness to Enable Transformation

It takes considerable data engineering to transform data from legacy customized structures into modern, industry-standard data structures. We are:

- Establishing robust tracking and validation processes for both automated and manual data loads.
- Conducting practice data load rounds (such as five each for each release) to help ensure seamless operation during the final cutover.

So far, we have over 56 million master and transactional data records loaded, with over 98% yield on average across 30 datasets. During our first and second releases, we achieved 100% and 90% data load success, compared to targets of 90% and 80%, respectively. And, out of the 22 Supply Chain Planning objects, 77% were 100% loaded into our digital supply chain and omni-channel commerce fulfillment platform during the first cycle of integrated testing, beating our goal by 17%.

56 Million
Master and Transactional Data Records Loaded

98% Yield
On Average Across 30 Datasets

Partnering with the Business—and the Ecosystem—to Drive Enterprise-wide Data Excellence

We've established scalable tools and processes to measure, manage, and govern data, and have developed multi-year rolling roadmaps for each data excellence pillar. We are now ready to start scaling our data strategy across the enterprise.

Aligning with our CIO's commitment to active listening, we are developing plans to articulate to the BUs the art of the possible from a data perspective. We are also taking advantage of close relationships with ecosystem partners who provide different perspectives and collaborative solutions.

Explore

- IT@Intel: Transforming Siloed Manufacturing Data into Unified Insights
- IT@Intel: Modernizing Enterprise Data Analytics Using Databricks in the Cloud
- IT@Intel: Master Data – Managed!

Driving Operational Excellence

Operational excellence means delivering positive business outcomes while ensuring security, regulatory compliance, and operational continuity. In 2024, we saw impressive reductions in incident volumes, successfully implemented advanced monitoring solutions, and maintained strong availability measures among several other key performance measures. In 2025, we continue to raise the bar on operational excellence through developing our OneIT Resiliency Program, conducting ongoing data center and factory modernization, and increasing financial discipline to bend the cost curve.

Developing a Framework for End-to-End Resiliency

Intel IT plays a crucial role in establishing Intel as the benchmark of excellence in the tech industry. In 2024, we launched the OneIT Resiliency Program to cultivate operational excellence, creating end-to-end resiliency across three pillars: technology, operations, and culture and organizational entities.

We conferred with several advisory and research firms and learned that a comprehensive IT resiliency framework and maturity model had not yet been developed. Therefore, we created the industry's first end-to-end, easy-to-use IT resiliency framework, which formalizes the application and infrastructure development cycle and includes foundational processes that cut across the entire cycle.

Next, we built an organizational resiliency maturity model and roadmap. The model clearly defines five levels of resiliency, while the roadmap enables developers to know what is necessary to move from one level to the next. To further simplify the design of resilient systems, we established detailed design standards for key IT pillars. The combination of the framework, model, and standards helps IT teams to build resilient systems and establish a roadmap of deliverables to continuously improve their organizational maturity.

While our work on resiliency will continue, our framework, design principles, and resiliency standards are already making a difference. For example, a previous IT incident in 2023 required eight days to fully recover our systems. In contrast, when a major security software vendor inadvertently caused a major disruption in 2024, we fully recovered all Tier 1/Tier 2 applications within a few hours and had 100% of applications operational within 1.5 days. We also made excellent progress on increasing application resiliency. Among our Tier 1 applications, 85% had zero risk, and the remaining applications closed all risks within six months.

In addition to improving infrastructure resiliency throughout 2025, we will continue to partner with BUs to strategically assess and streamline Tier 1 and Tier 2 applications, optimizing operational efficiency and enhancing overall business performance. Our resiliency program underscores the importance of proactively identifying risk and creating a culture that identifies and escalates risk.

Explore

- IT@Intel: IT Resiliency Drives a Resilient Enterprise

Continuing Our Data Center Modernization

Intel IT operates 54 data center modules at 15 data center sites. These sites have a total capacity of 133 megawatts, housing more than 454,000 servers that underpin Intel's vast computing needs. We are constantly enhancing our data center strategy to continuously improve IT operational efficiency and environmental responsibility. Building on previous investments and techniques, our data center strategy has generated savings exceeding USD 9.6 billion from 2010 to 2023.

Intel IT's data centers are delivering impressive results. Through the use of advanced infrastructure software and high-performance computing (HPC), we have revolutionized our data center operations to boost efficiency and sustainability. We've also reduced costs by implementing strategies such as disaggregated server design and enhancing international WAN links. Following are details about some of our most recent data center modernization efforts.

More Resilient, Sustainable Data Centers

Intel recognizes the growing importance of sustainability. From 2012-2024, we have saved over 1.91 billion KW hours compared to industry-standard data centers. We are also expanding the use of technologies that can reduce carbon emissions in our data centers.

Fuel Cell Technology

To meet additional HPC scale needed for Intel Products and Intel Foundry, we are increasing our use of fuel cell technology, with plans to add 31 MW of fuel cell capacity to an existing 6 MW fuel cell installation deployed over ten years ago.

Advanced Cooling Technologies

Data centers have historically run at low ambient temperatures, but new technology has relaxed thermal operating constraints, which lowers air conditioning costs and reduces or eliminates expensive computer room air conditioning (CRAC) units and chillers. In ten years of running data centers at higher temperatures, we have observed no increase in component failure.

Our latest data centers utilize close-coupled evaporative cooling using recycled water, allowing us to operate at an annual average of 1.06 Power Usage Effectiveness (PUE), with 54 KW power per 60U rack. Combined with increased rack densities, we are meeting the increasing demands for compute and storage while reducing our data center footprint.

Resilient Design Zones

Our HPC environment underpins the crucial silicon design work of Intel's engineers. To increase resiliency, we introduced Design zones into the Design hub computing environment, successfully scaling multiple zones to provide increased scale and reliability in a cost-effective manner.

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.

Using Software-Defined Networking to Keep Up with the Pace of Business

As Intel's business expands, the demand for data center network capacity has surged by over 25% annually, with new capacity needing to be operational within 24 hours. As early as 2014, we acknowledged the potential of software-defined networking (SDN) to help address these challenges.

We believe that our selection of open, standards-based technologies for constructing underlay and overlay networks, combined with an open-source orchestration layer, has provided us with the flexibility to adapt to evolving business needs and realize the value of a broader ecosystem.

25%
Increase in Network Demand

We have established a data center network architecture that extensively uses automation. The open-source platform provides the flexibility to integrate additional business-driven automation, enabling us to meet our growth objective and timeline requirements.

Our network architecture strategy relies on five pillars:

- **Scalability through standardization.** We maintain compliancy with our standards to enable automation and rapid scalability at large data centers.
- **Programmability.** This allows our workforce to quickly adapt to significant growth in network scale and facilitates full lifecycle provisioning of network infrastructure.
- **Security.** The network can be segmented using a common infrastructure to support various use cases and enhance data center security.
- **Resiliency.** Built-in resiliency helps ensure continuous operation of network functionality, facilitates rapid recovery, and maintains performance even in compromised conditions.
- **Supportability.** We strive to maintain the network's designed level of performance and availability. When issues arise, adhering to standards accelerates troubleshooting.

Over the past five years, we have migrated the majority (93%) of our data centers to a new SDN architecture based on a leaf-spine underlay combined with overlay networks, using industry-standard components and protocols. This has significantly improved network delivery times, reduced labor costs, and improved stability and security. We are delivering new network capacity 20% faster, with a 102% increase in network resource efficiency, a 70% reduction in performance-related incidents, and no major outages for design data centers for four years in a row.

Evolving Our Enterprise Architecture Processes and Practices

Intel's continued success hinges on quickly adapting to market disruptions and opportunities through innovation and business transformation. Intel IT is committed to contributing to that success by enhancing our EA processes. Current areas of effort include the following:

- Validating our EA practices and capabilities through additional industry-standard benchmarking.
- Improving business value and operational efficiency by distributing governance responsibilities.
- Defining a strategy for software and hardware asset management (more details are provided later in this section).
- Tracking emerging technologies as a critical part of our EA.
- Increasing the resiliency of our EA from end to end. See "Developing a Framework for End-to-End Resiliency" for more details.
- Infusing AI into our EA workflows. See "Delivering the Power of AI to Employees" for more details.

Software and Hardware Management

Intel IT and the BUs are working together to modernize Intel's business processes and supporting tools, including the ERP systems. We are reducing customization and transitioning to industry-standard business processes, including rationalization for more than 700 core connected ERP boundary applications. We have determined that we can migrate and eliminate more than 40% of these applications to reduce costs, standardize functions, and improve core ERP applications.

1,700+ Applications and **180+** Horizontal Platforms

spanning private cloud, public cloud, SAAS solutions, and 12K+ software products

The first phase of this transformation is in production. We have centralized Intel's finance capabilities in a new ERP system and new cloud-based AI and analytics data lakehouse (see "Data Analytics Supported by a Modern Data Lakehouse" for more information). Subsequent phases will bring business transformation to Intel's Product and Foundry businesses, which are also supported with new ERP systems.

All the EA and design work used to enable these efforts is coordinated through a technical work group (TWG) that includes experts from different segments, helping ensure that all new EA blueprints are vetted, gaps are identified and closed, and risks are managed.

The EA team defines the application rationalization strategy and establishes the roadmaps across all IT segments. Our application rationalization and software asset management efforts have seen the following successes:

- Eliminated 321 in-house applications.
- Removed 1,400 commercial software products from the environment.
- Generated USD 13 million in cost avoidance.

321
In-House Applications Eliminated

1,400
Commercial Software Products Removed from the Environment

$13 M
USD Generated in Cost Avoidance

Our hardware and software management process includes improving data quality (see "Building a Strong Data Foundation"). We have launched a cross-IT program led by the EA team to proactively manage Intel's hardware and software assets and optimize spending through intelligent asset management augmented by AI.

Tightly managing applications and infrastructure assets while optimizing EA processes with AI promotes innovation, reduces technical debt, improves business agility, and supports Intel's strategic growth imperative.

Explore

- IT@Intel: Data Center Strategy Leads Our Business Transformation
- IT@Intel: Validating and Evolving Intel IT's Multicloud Strategy
- IT@Intel: Scaling Intel's Data Centers with Software-Defined Networking and Automation
- IT@Intel: Fuel Cells – An Alternative Energy Source for Intel's Data Centers

Partnering with the Business: Joint Technology Council

The Joint Technology Council (JTC) was launched in 2024 to help propel Intel's strategy forward by enhancing our execution capabilities and fortifying the collaboration between IT and Intel's Business Units (BUs). The JTC members include the CIO, CISO, Global IT direct staff, leaders, enterprise stakeholders, business relationship managers (BRMs), and Technical Review Committee (TRC) leads. Our goals include refining investment roadmaps, fostering innovation, supporting Intel's top challenges, and executing decisions on investments.

In each of our monthly meetings, we dive into a specific topic and focus on initiatives in three main areas: Corporate, Product, and Foundry. For example, enterprise AI is a Corporate function, Software and Services (software sales) is a Product function, and Foundry Transformation and Operations Execution is a Foundry function.

The success of the JTC hinges on active participation, open communication, and a shared commitment to our goals. Every member brings a unique perspective that can help us challenge assumptions and collaborate to drive results. Together, we make sure that IT investments are strategic, our innovation is impactful, and our execution is flawless.

Supporting Intel Foundry through Factory Modernization

Intel needs smart factories both to manufacture its own products and to attract and retain Intel Foundry customers. Intel IT is creating intelligent, data-driven environments where AI, edge compute, and industrial IoT (IIoT) converge to optimize every step of production, including the following examples:

- Rolling out a private 5G network across guardhouses, power and water systems, cameras, robotics, and sensors.
- Extending our IIoT infrastructure based on Intel® IoT Gateways to closed-loop control of fault detection and classification (FDC) systems.
- Maximizing factory productivity by using IIoT infrastructure to reduce the risk of unexpected failures.

Our smart-factory building blocks—like AI, pervasive connectivity, and IIoT—are critical for scaling Intel's Foundry ambitions while continuously improving cost, quality, and business agility.

Increasing Operational Efficiency Using a 5G Private Network

Crucial factory operation services, such as water, power facilities, and security infrastructure use equipment that is often outdoors and difficult to connect to the corporate network. This lack of connectivity makes it challenging for equipment technicians to monitor, troubleshoot, and perform preventative maintenance.

We worked with Intel's Corporate Services (CS) team to connect critical factory-support infrastructure, including Internet of Things (IoT), robotics, operations, and security, to the corporate network using a 5G private network.

With this connectivity in place, CS technicians can tap into sensor data to enable predictive and preventative maintenance and to connect field devices directly to the corporate network to increase efficiency. Currently, our 5G private network program supports only 13 use cases—a fraction of the potential applications. Our 5G deployment spans five factories and is projected to generate a net present value (NPV) of USD 35 million over the next five years.

Extending Our IIoT Infrastructure to Include Closed-Loop Control of Subfab Equipment

Connecting the unconnected is a primary theme for Intel IT. The IIoT infrastructure that we have deployed in Intel's factories is based on open-source tools and Intel architecture. Factory equipment data is collected from sensors, analyzed, and formatted before being published through an Intel IoT Gateway to our central IIoT dashboard and sent to consuming applications. AI and machine learning at the edge analyze large volumes of data ingested from sensors to intercept events that might indicate a failure in real time, and to assist in factory equipment planning and utilization. This affordable and flexible IIoT infrastructure yields fast ROI. Our current NPV for IIoT in the factory is USD 232 million.

Our Current NPV for IIoT in the Factory is

USD 232 Million

We have also connected subfab equipment—such as pumps, gas abatement systems, and chillers—to the same Intel IoT Gateways. With this system, we gather data from a machine, analyze it, and then send commands back to the machine to modify its behavior, increasing factory efficiency by reducing downtime and scrap.

Our innovative open-source solution—the first of its kind—uses a fully standardized software and hardware stack and is reusable for many use cases. Our centralized repository of drivers makes it easy to build on existing knowledge and add new drivers. Using our IIoT-based FDC solution enabled us to use lower-cost, standardized infrastructure (compared to vendor-provided solutions), and modify vendor contracts, lowering overall Foundry costs by USD 70 million over five years.

Creating Significant Savings and Reducing Factory Equipment Downtime with IIoT

Building on previous successful implementation of IIoT use cases in Intel's fabrication facilities (fabs), we are expanding our IIoT use case portfolio to include Intel's Assembly and Test Manufacturing (ATM) factories. Here are some example use cases:

Predicting equipment failure

Pump seals, bearings, and shafts can wear out over time, leading excessive electricity usage, abrupt shut-downs, and potential damage to product. We deployed sensors to pumps and blowers to enable continuous vibration monitoring, then programmed the IIoT dashboard with specific thresholds so it will show an alert if a pump or blower is experiencing excessive vibration. Over four years, the estimated ROI is hundreds of thousands of dollars across several pumps. These savings result from both scrap avoidance in the thousands of units and noticeable electrical consumption savings. For blowers, the sensors prevented five cases of blower failure with an estimated avoidance value to Intel in the millions.

Optimizing Expensive Material Usage

Hydrofluoroethers (HFE) are a class of non-ozone-depleting organic solvents used by some equipment in ATM facilities as a coolant. However, it is expensive and difficult to procure. We deployed sensors to constantly monitor HFE flow rate, providing early detection of excessive HFE consumption caused by leaks in the HFE tubing and piping. This has resulted in a significant reduction in chemical waste and considerable savings on material purchases.

Improving Facility Maintenance with Robotic Automation

As of March 2025, we have deployed 14 robots equipped with multiple integrated IoT sensors. The robots enhance safety and operational efficiency by autonomously roaming Intel facilities to perform facility inspections and read values from instrumentation. This type of robotic automation delivers significant annual cost savings by enabling preemptive dispatch of repair and maintenance technicians based on real-time data. Since deployment, the robots have:

- Trained for 2,070 unique inspections across multiple facilities.
- Conducted 606,423 inspections.
- Covered a total of 2,870 miles.

We are conducting additional proofs of concept (PoCs), such as detecting acoustic leaks, reading gauges, and checking thermal status. So far, all these PoCs have resulted in nearly 100% accuracy.

Trained for
2,000+
Unique Inspections Across Multiple Facilities

Conducted
>606,000
Inspections Since January 2024

Covered
2,800+
Total Miles

Explore

- IT@Intel: Delivering Operational Efficiencies Using a 5G Private Network
- IT@Intel: Expanding Low-Cost Industrial Internet of Things (IIoT) Manufacturing Use Cases
- IT@Intel: Intel Cuts Downtime and Costs with Fault Detection Systems for Factory Equipment

Bending the Cost Curve

Typically, as revenue increases, operating expenses also go up. But through operational excellence, we can reverse this trend so that while revenue rises, operating expenses fall—this is called "bending the cost curve." By focusing on the right IT levers, we can drive down our fixed and variable unit costs to achieve the bend, capitalizing on both cost avoidance and cost-savings opportunities.

Our approach is two-fold:

- Eliminate all valueless technology activities.
- Increase investments in automation efficiencies.

The following figure illustrates our cost-curve model:[3]

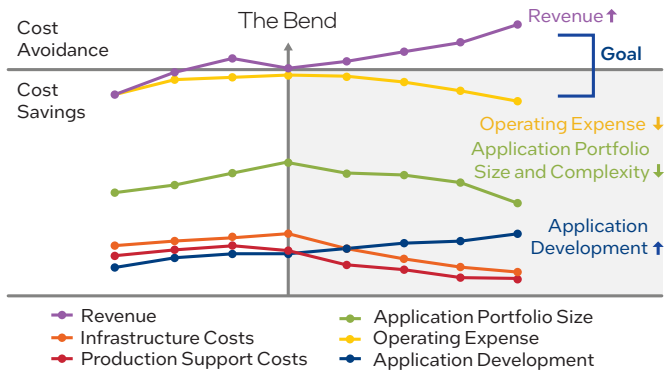

[3] Source for this cost curve graphic is Rubin Worldwide.

Here are some specific cost-curve drivers:

- Value-based project portfolio focused on initiatives with high ROI.
- AI Inside—in IT and across the enterprise.
- Infrastructure resiliency delivered through automation and proactive monitoring.
- Modern and frictionless IT services to enhance worker productivity.
- Application rationalization to decrease the complexity of our application portfolio.
- Workforce optimization through right-shoring, right-sizing, and upskilling.
- Transformative ERP programs to lay a scalable, robust foundation for the enterprise.
- Cost and IT services consumption transparency through IT showback, shadow spend visibility, and restricting IT commodities transactions through governance.

These activities are driving real business impact. By aligning our efforts with organizational priorities, enhancing cost transparency, and pursuing our business value goals, we are not only on track to meet our targets but also setting a standard for strategic IT investment. Together, we will continue to leverage technology as a catalyst, ensuring that every IT initiative translates into meaningful outcomes.

Embedded Consulting: Intel Flex Services

Intel Flex is an embedded consulting unit (ECU) that serves internal customers to help them solve their most challenging issues and accelerate results. Our professionals offer a wide variety of skills to help business units (BUs) achieve their business goals across four areas of practice: software engineering, DevOps as a service, business consulting, and outsourcing management.

Intel Flex has been around for more than 25 years, with a consistent record of customer-reported satisfaction beyond industry standards. In 2024, we achieved a partner satisfaction rating of 96, won 35 Division recognition awards, and enabled about USD 40 million in savings through our Lean Six Sigma Certification.

Two recent initiatives have been particularly impactful:

- **Intel® AI Assistant Builder.** Formerly known as Project SuperBuilder, this AI assistant simplifies the process of building and training a targeted on-device chatbot, enabling the creation of "lite" large language models (LLMs) for specific tasks. Whatever the use case, AI Assistant Builder can deliver fast and relevant responses. Various Intel partners have already rolled out solutions using versions of Intel AI Assistant Builder, and the solution won the "Best AI Debut" award at Mobile World Congress 2025.
- **Internal GitHub workflow runner system.** We designed, deployed, and now sustain an industry-leading solution that helps automate software workflows through the use of an internal GitHub. The solution has generated USD 4 million in annual cost savings, based on developers using a centralized service instead of development teams paying for individual virtual machines. The solution also enables productivity savings due to automation, reduced overhead, and increased re-use. Learn more about the Intel AI Assistant Builder.

Delivering an Amazing Employee Experience

We are committed to delivering modern and frictionless technology to enhance worker productivity. Our efforts span employee experience with devices, information, workplace and support, and collaboration. For example, our work with GenAI to provide chatbots and assistants enables employees to quickly find the information they need. (See "Delivering the Power of AI to Employees" for details on our GenAI work.) Another focus area is device experience, looking for ways to make device interaction more smooth, reliable, and enjoyable for employees, including more seamless client device management, better connectivity, and longer battery life.

Increasing Employee Productivity with Device Experience Optimization

Intel's employees rely on their client devices—primarily laptops—to complete their jobs, to the point that work comes to a halt when client device performance declines. That's why IT is on a continual mission to improve endpoint optimization and client lifecycle management.

Revisiting Client Management to Take Advantage of Cloud Capabilities

Thanks to maturing cloud-based client management applications, today's client build process is out-of-the-box and the new client management process is hosted in the cloud. We are gradually removing legacy tools from the entire fleet (as opposed to only new devices).

Historically, our OS upgrades were highly time-consuming. In contrast, our new, cloud-based, modern client management environment combines a number of components to support seamless client provisioning that improves the user experience (UX) and substantially reduces IT effort. Cloud-based Windows Update for Business (WUfB) has enabled us to retire manual task sequences and reduce engineering time from four weeks to just a few days. We completed most of the Windows 11 in-place upgrades in only 13 weeks once application readiness testing and the pilot project were complete. From a security standpoint, simplified compliance policies combined with conditional access can help ensure that devices are secure and up-to-date before they are allowed to access corporate resources.

Cloud-first tools have helped us to reduce our attack surface and cut PC setup issues by 38%. We calculate that the new processes will require 10,000 fewer person-days per year in end-user downtime and enable a 100% reduction in security patch preparation, which can save 10,208 user-facing patching days per year. These metrics underscore how moving to cloud-native client management helps us meet our two most important client management goals: increasing efficiency and improving employee productivity and UX.

Support Costs

- **38% Reduction** in Incidents Related to PC Setup
- **15% Reduction** in Engineering Effort vs. Legacy Process

User Downtime

- **10,000** Person-Days per Year Reduction in End-User Downtime

Security Patches

- **100% Reduction** in Security Patch Prep and Deployment Activities
- **10,208** User-Facing Patching Days per Year Saved

Managing Wi-Fi Networks Using Client Analytics

When connectivity issues affect the corporate Wi-Fi network, both productivity and UX suffer. Intel's Connectivity Solution group developed Intel® Connectivity Analytics technology and collaborated with Cisco and Intel IT to develop compelling use cases and solutions that utilize the tool. It is built into Intel® Wi-Fi adapters and requires no software installation or maintenance. As "customer zero," Intel IT

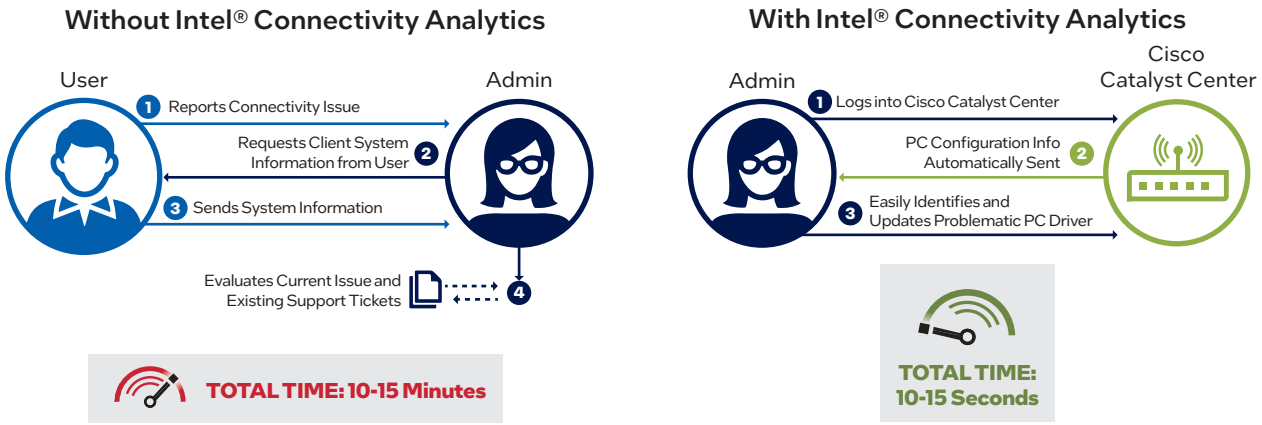

provided product feedback and practical applications of the technology in our enterprise environment and helped optimize the use cases to better define the analytics needed to achieve desired business outcomes. Intel Connectivity Analytics now improves the management of Intel's Wi-Fi network, reducing some client-side troubleshooting from upwards of 15 minutes down to just 10–15 seconds, and reducing the time to find the root cause for network issues from potentially days to often just seconds.

For example, locations with poor Wi-Fi coverage can be identified based on connectivity reports from clients. Prior to Intel Connectivity Analytics technology, discovering the cause of the issue using legacy network troubleshooting methods would have been time-consuming and required manual data analysis. Intel Connectivity Analytics also helps to quickly troubleshoot network issues by identifying co-channel interference so we can make changes to the channel or the AP power level.

We are closely collaborating with Intel's Connectivity Solution group to help expand the use cases for Intel Connectivity Analytics and further enhance AI for IT Operations (AIOps).

Using Benchmarking to Better Understand the End-User Laptop Experience

To support Intel's highly mobile employees, Intel IT monitors the UX on over 120,000 laptops.[4] PC experience benchmarking helps us evaluate how the Intel IT build affects power consumption and battery life in various scenarios; it also enables us to investigate and mitigate potential problems revealed by the benchmarking.

PC experience benchmarking isn't new at Intel IT, but our traditional tool did not provide insight into how the actual software applications used by Intel employees affected CPU and GPU usage. We now use Procyon O365, a utility that measures CPU and GPU power consumption for Microsoft Office 365. We also use the Intel® Battery Life Diagnostic Tool (Intel® BLDT) to help look even deeper into how applications affect PC battery life. Intel's Client Computing Group (CCG) uses similar benchmarking to identify the expected battery life for a laptop equipped with a new Intel® PC processor, determine if there are performance gaps in new products, and qualify systems so they meet customer expectations and provide a positive UX.

The new PC experience benchmarking toolset has revealed several software-specific issues that we mitigated to help improve battery life and employee experience.

A Better IT Support Experience with Conversational AI

This year we are transitioning our existing IT support bot, called AskIT, from a commercial solution to an internal iGPT-powered chatbot. By taking advantage of internal AI capabilities to deliver a more compelling solution, we are driving down costs and using conversational AI to deliver an improved UX.

This next-gen AskIT implementation has given employees better responses to help them solve their IT issues. As of mid-May 2025, over 11,000 employees have had more than 26,000 conversations with the new AskIT, and usage grows daily—resulting in a 30% reduction in the number of "How Do I..." support tickets so far. We expect to meet our goal of a 50% reduction by the end of the year.

11,000 Employees have had...

26,000 Conversations with AskIT, resulting in...

30% "How Do I..." ticket reduction

Explore

- IT@Intel: Modernizing Windows Client Management
- IT@Intel: Optimizing and Troubleshooting Wi-Fi Networks Using Client Analytics
- IT@Intel: More Accurately Benchmarking the End-User PC Experience

[4] The total number of laptops managed by Intel IT is much larger than this.

Shaping a Resilient Workforce

Intel IT's goal is to create an inspired workforce where everyone understands their purpose and has confidence in their skill. We can achieve this goal through continuous learning, adaptability, and integration of emerging technologies.

Creating Resilient, Inspired Teams

"In my view, the key to a successful IT organization is about automation, understanding, self-service, standardization, resiliency, and consistency."

This quote from our CIO underscores the importance of resiliency—and one aspect of driving IT and enterprise resiliency is creating a strong workforce that can keep innovating even in the face of disruption and rapid change. CIO.com stated that 69% of IT leaders had plans to upskill or reskill their current employees in the next 12 months to meet evolving business needs.[5]

[5] CIO.com, July 2024, "Foster a culture of learning with upskilling and reskilling"

Establishing a culture of resiliency starts from the top down, giving employees role models for contributing to the overall resiliency of the enterprise. We are also instilling employees with a OneIT mindset that views challenges as opportunities for growth.

In 2024, we formed a People Council, where leaders focus on the development and retention of top talent, talent trends, and strategic workforce planning. This structure holds management accountable for continuously focusing on our employees' development. To shape a resilient workforce, the People Council's strategy encompasses planning, attracting top talent, further developing the workforce's skills, retaining employees, and rewarding the workforce for high performance. Baked into all these phases is a commitment to building an inclusive culture to drive innovation and deliver business value, along with strong communication to keep managers and employees informed of available training opportunities.

Our recruitment strategies prioritize candidates with a proven track record of adaptability and problem-solving. For existing employees, we are establishing ongoing training and development programs that equip employees with the latest knowledge and skills. By building a resilient IT workforce, we can rest assured that they have the talent needed to navigate through challenges and capitalize on opportunities.

Currently, our retention strategy focuses on five top skills for employee development: AI, SAP, cloud, business relationship management, and data. In 2024, 75% of all investments in training and development were focused on these top five skills.

75%
of Training and Development is Focused on AI, SAP, Cloud, Business Relationship Management, and Data

We are also establishing a standardized, best-in-class set of data management and governance tools that the entire team can be trained to use. In that way, we have cross-team backup. It also helps with employee retention, because top employees want to work with top-of-the-line tools.

In summary, we are shaping a resilient workforce through accountability, trust, and the empowerment of skill development.

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 continue to capitalize on AI and various other new technology innovations to rapidly enable the BUs to deliver high-quality Intel products to market.

We made significant transformative progress over the last year, creating efficient operational excellence and making a successful transition to the internal Foundry model, and we will continue to build on that momentum. Moving forward, we will expand our AI capabilities into all core business processes with a focus on creating impeccable quality and reliability of AI systems. With a solid foundation in place, we will use our strategic data assets to cost-effectively scale our global business operations.

We will focus on maturing, scaling, and integrating our solutions and services. All of us across IT must approach our daily work and ambitious goals with an innovative mindset and an openness to new ideas, 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 BUs. We will know what the BUs need and want and will use our expertise to provide optimal solutions and influence informed prioritizations of projects. Once priorities are set, we will provide solutions and technology that support the BUs' outcomes and needs, on time and on budget.

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 self-service-enabled, agile, and resilient.

At the moment, times may seem difficult, but we need to be willing to accept that it's temporary, believe that there's a purpose, and that there's an end goal—to make Intel amazing again. I remain confident in our ability to deliver on IT's critical evolutionary role and to exceed the expectations of the enterprise, this year and beyond. The same "One Intel" mindset that propelled us in the past 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.

"The secret to change is to focus all of your energy, not on fighting the old, but on building the new."

- Dan Millman, "Way of the Peaceful Warrior"

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.