White Paper
May 2025

IT@Intel:
IT Resiliency Drives a Resilient Enterprise

Intel IT is building a robust resiliency framework across people, processes, culture, and technologies to deliver a best-in-class service experience for Intel business units and customers

Intel IT Authors

Madhurasmitha Chakravarthy
IT Sr. Director, Resiliency Co-Lead

Table of Contents

Executive Summary

Modern electrical grids incorporate redundancy and automated rerouting to maintain a reliable electrical supply during equipment failures or natural disasters. Plants have coping mechanisms to deal with adverse conditions like drought. Similarly, it is important that IT systems, including applications, infrastructure, platforms, and services, can withstand or recover quickly from difficulties. Resilient technology helps companies thrive, and when disruptions occur, enables them to emerge stronger than before.

In 2024, Intel IT launched our OneIT Resiliency Program, geared to instill resiliency across technologies, operations, and culture and organizational entities. In addition to defining these three pillars, we codified resiliency through a unique combination of framework, principles, standards, maturity model, and roadmaps. We focused first on applications, closing 100% of all resiliency gaps in Tier 1 applications by the end of the year. We also successfully weathered a major incident caused by one of our security software vendors: Despite widespread PC failure across the company, we restored functionality in just a few hours for Tier 1 and Tier 2 applications and under 1.5 days for all applications at Intel.

During 2025, we will extend our efforts to infrastructure and data centers. We hope that our resiliency framework, standards, blueprints, recipes, and maturity model will inspire other IT departments to pursue a similar path to IT—and enterprise—resilience.

Acronyms

IAO IT IDM 2.0 Acceleration Office

Business Challenge

Intel is on a mission to become the benchmark of excellence in the tech industry as a product company and as a foundry. IDM 2.0[1] represents one of Intel's most sweeping transformations and Intel IT plays a huge role in helping that transformation succeed. In particular, we have the responsibility of ensuring the resiliency of foundational IT systems and solutions because they help:

- Maintain business continuity.
- Uphold Intel's reputation for reliability.
- Support rapid innovation.

We define resiliency as the ability to adapt to unforeseen disruptions, whether they stem from process and technological failures or natural disasters. In 2023, we experienced a major IT incident that revealed resiliency gaps that were exacerbated by the following:

- Complexity in the integration and technology landscape, which is compounded by the fact that anyone at Intel, whether in IT or a business unit, can build an IT system.[2]
- Attrition of talent and knowledge due to insufficient alignment with career development or critical skills that enable resilient personnel and organizations.
- Declining process discipline due to lack of formalized emphasis on resiliency.

The importance of IT and enterprise resiliency is undisputed, and this incident made it clear we needed to improve our resiliency. However, systematically and holistically identifying existing risks so we can mitigate them wasn't straightforward. After conferring with several well-known advisory and research firms, it became clear that no one had yet developed a simple-yet-comprehensive IT resiliency framework and maturity model that was easy for IT teams to follow.

Solution Overview

Resiliency—the capacity to withstand or recover quickly from difficulties—is an essential attribute that enables individuals, organizations, and systems to navigate through challenges and emerge stronger. We developed the OneIT Resiliency Program, a strategic initiative designed to instill a culture of resiliency and help ensure that any IT capability is built to endure failure points or recover from them rapidly. For us, our OneIT Resiliency Program encompasses all IT and business unit applications, infrastructure, platforms, and services.

Our OneIT Resiliency Program focuses on three pillars:

- **Technology resiliency.** Create robust and secure systems that are capable of withstanding operational disruptions.
- **Operational resiliency.** Design processes with flexibility and redundancy to handle unexpected changes without compromising efficiency.
- **Cultural and organizational resiliency.** Encourage adaptability and problem-solving, and promote a mindset that views challenges as opportunities for growth and improvement. Foster a resilient workforce by investing in training and support.

The overall objective of the OneIT Resiliency Program is to cultivate operational excellence through discipline, unwavering standards, iterative improvements, and measured progress. We have specific objectives for each of the pillars, as shown in Table 1. The section, "A Deeper Dive into Our Resiliency Pillars" provides more details.

Table 1. Objectives for Each of the OneIT Resiliency Program Pillars

Resiliency Program Pillar	Objectives
Technology	▪ Develop resiliency gap plans for 100% of all Tier 1 and Tier 2 applications. ▪ Establish a technology resiliency standard for applications and infrastructure. ▪ Ensure we build out a resilient IT IDM 2.0 Acceleration Office (IAO) architecture.
Operational	▪ Establish a maturity model and resiliency roadmap. ▪ Drive focused improvement across key operational processes. ▪ Provide progress visibility through standardized dashboards.
Cultural and Organizational	▪ Enforce job role expectations and baseline training for resiliency. ▪ Establish governing structures to reinforce the culture of resiliency. ▪ Recognize and celebrate wins and learn from early adopters.

To achieve these objectives, we developed a comprehensive yet easy-to-use IT resiliency framework—the first of its kind in the industry. It provides a blueprint that connects the three pillars to help ensure that Intel IT's business units and Intel's customers receive best-in-class services. The framework uses language that both our IT engineers and business units can understand, so that risk can be clarified and enumerated. The framework formalizes the application and infrastructure development cycle (define, design, develop, deploy, and manage) and includes foundational processes that cut across the entire cycle, like impact management, change management, risk management, business continuity, incident management, and crisis preparation. See the section, "Capability Framework for Resiliency" for more details.

[1] "IDM 2.0" is a major evolution of Intel's integrated device manufacturing (IDM) model. For more information, read the press release.

[2] We define an IT system as a computer system, including all hardware, software, and peripheral equipment.

With the framework in place, we then built an organizational resiliency maturity model and resiliency roadmap. The model clearly defines five levels of resiliency (see Table 2), while the roadmap enables developers to know what is necessary to move from one level to the next. The combination of the framework and model helps coach IT teams to build resilient systems and establish a roadmap of deliverables to continuously improve their organizational maturity.

A Deeper Dive into Our Resiliency Pillars

Figure 1 illustrates our resiliency pillars. Intel IT is more than technology—it's our processes, our workforce culture, and how our teams work together. Collectively, these pillars provide a comprehensive and explicit strategy for achieving the IT resiliency Intel needs to thrive. The following subsections provide more details about each pillar.

Figure 1. Our OneIT Resiliency Program rests on three pillars.

Technology Resiliency

Technology is a broad term. It can refer to layers of technology that make up a particular solution. These include applications, platforms, middleware, servers, network infrastructure, and data centers. Technology can also refer to a string of solutions that comprise an overall workflow. Both elements need to be addressed as part of resiliency.

Resiliency at every layer. Defining resiliency standards for each layer enables us to assess the resiliency of an application, a server, or an entire data center. It is important to develop standards that help ensure end-to-end resiliency, instead of mere single-point resiliency. Consider the analogy of a symphony. For a successful performance, each individual within the orchestra has to play their part correctly, but the musicians must also perform in concert. It's the same in the IT world.

Institutionalizing resiliency standards for infrastructure, data centers, applications, and other technologies is the cornerstone of technology resiliency. This involves setting up robust protocols and guidelines that help ensure all technological assets are designed, implemented, and maintained to withstand disruptions. For instance, data centers should have redundant power supplies and failover mechanisms to ensure continuous operation. Applications should be built with scalability and security in mind, and be capable of handling increased loads. By institutionalizing these standards, we can minimize downtime, protect critical data, and deliver services seamlessly, thereby enhancing brand reputation and customer trust.

Resiliency throughout workflows. Let's consider the software development workflow. Developers write code, check the code into a source code management system, use a build system to build the image, use a tool to scan for code vulnerabilities, and then publish the software into a location for consumption. In the single-step view of resiliency, the application owners measure resiliency only at their step. So, what if the source code management system is down on Monday; the build system fails on Tuesday; on Wednesday, the scan tool isn't working; and on Thursday, the storage system fails.

Each of the individual processes has experienced only one failure during the week. However, the user of the entire process has faced failure four times in the same week—that is not a robust process. Instead of enabling resiliency and disaster recovery for single-purpose steps, we need to look at the end-to-end business process, so that the compounded risk is visible. For some processes, that may mean cooperation between IT and the business unit, because some steps may be managed by IT, and some may be managed by the business.

Operational Resiliency

For operational resiliency to become a reality, our existing processes must adopt a resiliency mindset. We should think from the perspective of resiliency every time we make a change, implement something new, or react to an incident.

Enhancing processes to ingrain resiliency into day-to-day work is the essence of operational resiliency. This involves re-evaluating and refining existing workflows to incorporate flexibility and redundancy. Whether it is incident management, change management, or risk management, these processes need to incorporate resiliency considerations.

For example, every time an IT team makes a change to an application, they need to consider the upstream/downstream systems and applications and coordinate with the relevant teams. Without this coordination, while the first team may consider the change successful, the downstream users may face failures. To drive operational resiliency, we are instituting new incident, impact, and change management processes in addition to establishing a governance structure and an Operations Council staffed by technology representatives. To bolster these efforts, we also conduct resiliency training and lunch-and-learn events.

Cultural and Organizational Resiliency

Unifying technology and operational resiliency with a sense of pride, ownership, and accountability for operational excellence is the goal of cultural resiliency. Closely related is organizational resiliency, which is about cultivating the necessary skills and talent for the adoption of our resiliency framework. Empowered, accountable teams build a culture of resiliency, and we invest in the development of a workforce that is skilled in resiliency practices and capable of implementing our framework effectively.

It's important that the resilient mindset starts from the top down—managers, team leaders, and other executive stakeholders set the tone for the rest of the workforce. We are using training programs, internal communications, and leadership initiatives to emphasize the value of resiliency and encourage a mindset that views challenges as opportunities for growth. When employees take pride in their work and feel accountable for the organization's accomplishments—and have the skills to succeed—they are more likely to engage in continuous improvement efforts and drive operational excellence.

While our governing structures can help foster the desired cultural and organizational resiliency, it's also important to incentivize and reward desired behaviors. So, we created a specific recognition package for resiliency. We have three cash award categories: quality, efficiency, and customer experience. Within those categories, we created specific awards like "Always-On Marvel" and "Productivity Prodigy," which are awarded according to evaluation criteria such as zero major incidents for the past six months or creating a permanent fix for a recurring problem with a high support ticket volume.

Additionally, 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 in resiliency. By building a resilient IT workforce, we can rest assured that they have the talent needed to navigate through challenges and capitalize on opportunities.

Roles and Expectations

When we examined organizational roles and expectations, we discovered that resiliency didn't factor into some of the role definitions, such as general manager and director. For other roles, like architect, service owner, and product owner, resiliency was already part of the role definition, but expectations needed to be enhanced. We also discovered certain resiliency-specific roles were missing, such as site reliability engineer. We also identified required training and other activities, such as forum participation, that would enable each role to best support end-to-end IT resiliency.

Let's consider the general manager's role as one example. We added expectations for this role, such as:

- Role-model and foster the resiliency culture.
- Appropriately fund resiliency initiatives.
- Support the teams to fulfill their resiliency mandates.
- Review and approve high-risk and high-impact changes.
- Recognize resiliency wins and talent.

We also added required training courses, such as major incident training, design thinking, digital leadership, business acumen training, change governance, and classes in stakeholder management and communications.

Geo Personnel Resiliency

Geo Personnel Resiliency is a facet of organizational resiliency that focuses on the resilience of our workforce. It helps ensure that employees can endure or recover from negative events without disrupting agreed-upon support. The goal is to understand risk exposure and prioritize actions to maintain this support.

For example, suppose there is a discrepancy between the support level for an application and its infrastructure. In this example, a software application has an expected support level of 24x7, yet the servers it runs on require only 16x5 support. What is the implication if a negative event occurs? Can our personnel provide the necessary coverage?

The coupling of strategic workforce decisions and a resiliency framework can address scenarios like this example misalignment. The personnel resiliency aspect entails understanding and evaluating the following:

- A catalog of available services and support levels, with varying permutations
- The current time-to-resolution rates
- Critical dependent skills
- Support risk exposure

Capability Framework for Resiliency

Resiliency is an essential aspect of the software development lifecycle that must be integrated into every phase to help ensure systems can withstand, adapt to, and recover from disruptions, whether they arise from technical failures, cyberattacks, or unforeseen events. As Intel increasingly relies on digital solutions, the ability to maintain service continuity and data integrity during adverse conditions becomes paramount. Resiliency must be embedded across all stages of the software development lifecycle—from initial definition and design to deployment and management—to safeguard business operations and achieve strategic objectives.

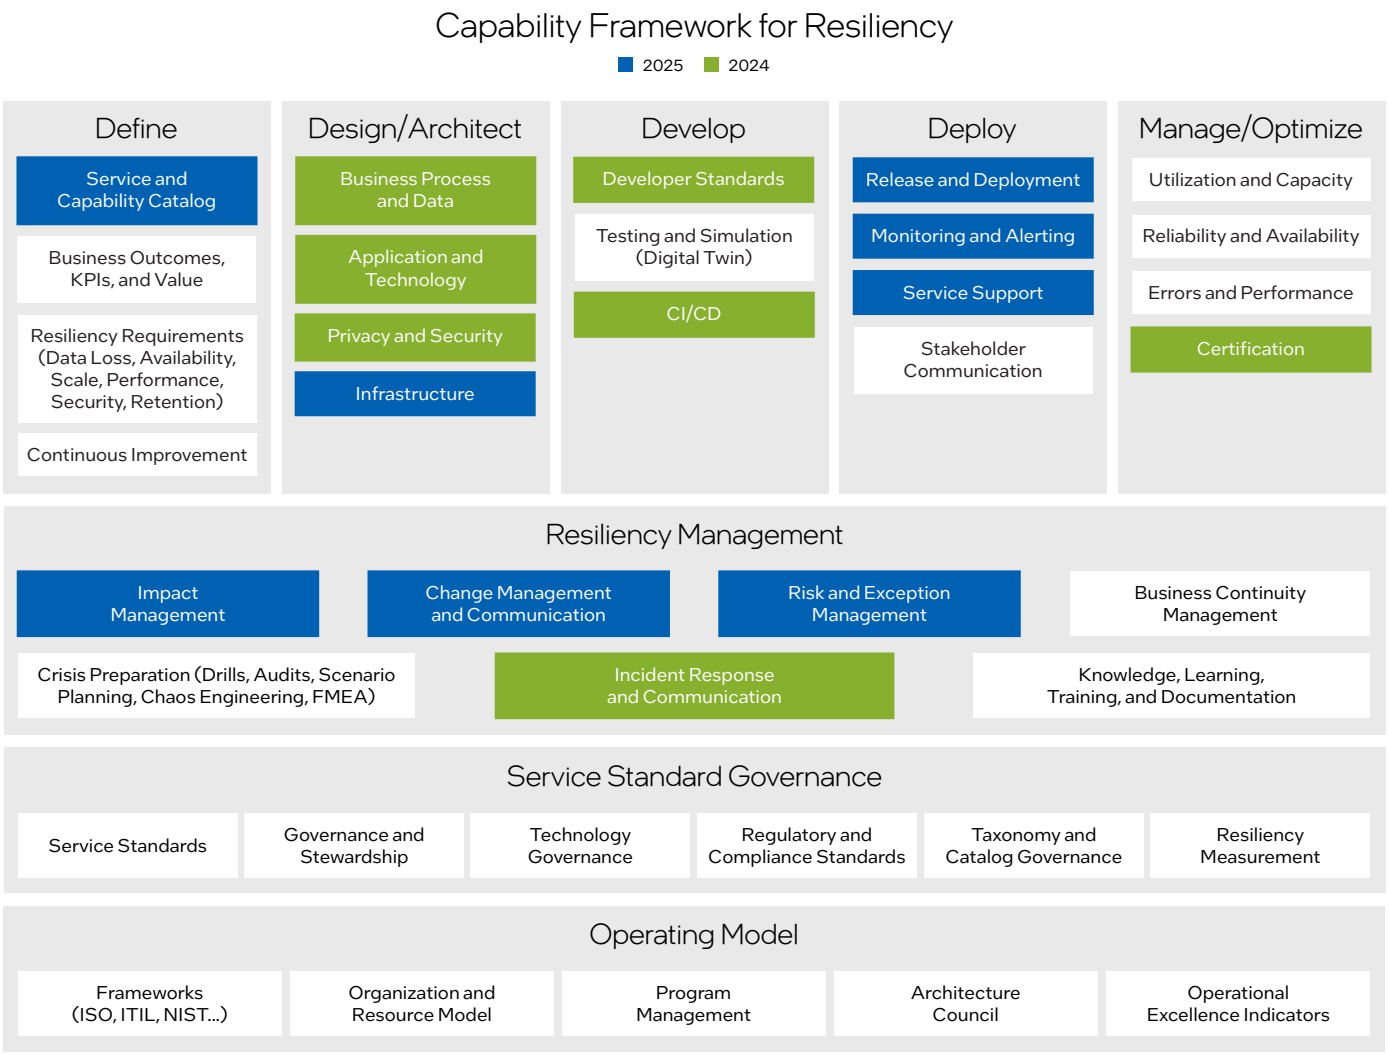

Figure 2. The resiliency framework provides a blueprint that helps ensure that Intel delivers best-in-class applications and services.

We have established a comprehensive capability framework for resiliency (see Figure 2) that defines minimum standards and includes a maturity model for each capability.

As mentioned earlier, the framework has several layers, and implementing it is a multi-year process. As the figure shows, in 2024 we focused mainly on applications—standards, tier and resiliency gap assessments, and certification, among other things. In 2025, our focus is on infrastructure resiliency, including heightened resiliency management.

At the top level, the framework follows the standard development pipeline—define, design, develop, deploy and measure. The Definition phase identifies into which business processes the application fits, like "order to cash" or "record to report" or "accounts payable." This phase also defines the business value and outcomes that matter to that particular process, as well as the resiliency requirements, which include the following:

- How much data loss can be tolerated for that business process?
- How much availability is required?
- How much does the application need to scale over the next few years?
- What type of data retention is needed?

Once the definitions are complete, the application can be designed and deployed, with the ability to measure how it delivers on the definitions over time.

The resiliency management, service standard governance, and operating model areas of the framework span the entire development cycle.

Resiliency Maturity Model and Roadmap

When we performed an industry scan to look for existing IT resiliency resources, we discovered bits and pieces that people had worked on, such as standards for risk management or a software resiliency methodology. But even the well-known research companies didn't have anything that resembled a comprehensive resiliency maturity model that could relate all the enterprise resiliency elements—we had to design our own. We hope that our work can inspire other enterprises can use a similar approach to enhance their own IT resilience.

Maturity Model

Our resiliency maturity model defines five levels of organizational maturity, ranging from Level 1, in which an IT organization takes a hit-or-miss approach to resiliency, to Level 5, where resilience is ingrained into the organization's DNA. Table 2 provides the specific characteristics for each maturity level.

As with most maturity models, there is a trade-off between the cost and benefit of moving up the maturity ladder. Ideally, we aim for a target maturity level score between "Level 3 - Managed" and "Level 4 - Optimized" because striving for "Level 5 - Integrated" is usually impractical and would diminish the return on investment as a whole. However, certain organizations at Intel with mission-critical importance may strive for Level 5, and some organizations may remain at Level 2 for some time.

Risk Assessment

Our goal is to increase IT resilience—but to do so, we need to know where we are currently. In 2024, we developed more than 20 risk categories for the technology resiliency pillar. Then, based on our tier definitions (see Figure 3), we evaluated several hundred IT apps against those risk categories. All evaluations and results are automated, so we can easily pull up dashboards that show risk assessment results (see Figure 4 on the next page for a sample resiliency scorecard). Out of all the application risks found, only 20% were considered high risk; the remaining had low impact.

Then we formulated a plan to resolve the priority risks using our resiliency, recovery, and operations design standards. We also developed resilient, highly available architectural blueprints and recipes to help developers meet resiliency goals. A key deliverable of the assessment was also to ensure that our IDM 2.0 transformation architecture was built out to necessary resiliency standards. The Technical Resiliency Checklist is now available as a part of the Intel Application Portfolio Management Solution, with integrated scoring and risk predictions.

We will complete a similar infrastructure assessment throughout 2025.

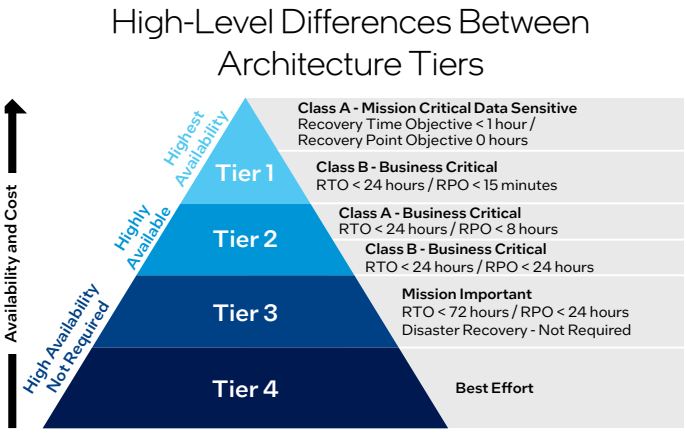

Figure 3. We formalized the differences between various application criticality tiers.

Table 2. Five Levels of Our Resiliency Maturity Model

Maturity Level	Objectives
1–Ad-Hoc	▪ Organizations lack a formalized approach to IT resilience. ▪ Responses to disruptions are reactive, and there's no consistent strategy or framework. ▪ There is minimal awareness of resilience importance.
2–Defined	▪ Organizations begin to establish standards, processes, and guidelines for IT resilience. ▪ They identify critical systems, create response plans, and allocate resources accordingly. ▪ They develop formalized incident response processes with defined roles and responsibilities. ▪ There is some awareness of resilience importance and leadership support.
3–Managed	▪ Organizations actively manage IT resilience. ▪ They have documented standards, guidelines, and plans; conduct regular risk assessments and audits; and invest in redundancy and failover mechanisms. ▪ They focus on improving incident response management, change and problem management, monitoring, and detection capabilities. ▪ There is a growing culture of resilience with leadership support.
4–Optimized	▪ Organizations optimize their resilience efforts. ▪ They continuously improve in hardening processes, proactive monitoring, and operations; they drive standards to execution, measure effectiveness, and foster operational excellence and discipline. ▪ Upstream/downstream application and capability dependencies are well-defined with integrated failure/recovery handling. ▪ They evaluate failure modes and failure response by using failure mode and effects analysis, chaos engineering, and site reliability engineering to continuously improve failure recovery operations with preventative measures. ▪ There is a strong resilience mindset growing across the organization.
5–Integrated	▪ At the highest level, IT resilience is fully integrated into the organization's culture and operations. ▪ They integrate resilience into their overall business strategy combined with decision-making, innovation, and risk management. ▪ They measure alignment with business strategy and adaptability as well as monitor innovation. ▪ Continuous learning and improvement, operational excellence, and discipline are institutionalized within the organization. ▪ Self-healing capabilities provide rapid recovery, preventing user impact. ▪ Observability, telemetry, and topology mapping are well integrated, establishing AIOps transformational capabilities. ▪ Resilience is ingrained into the organization's DNA.

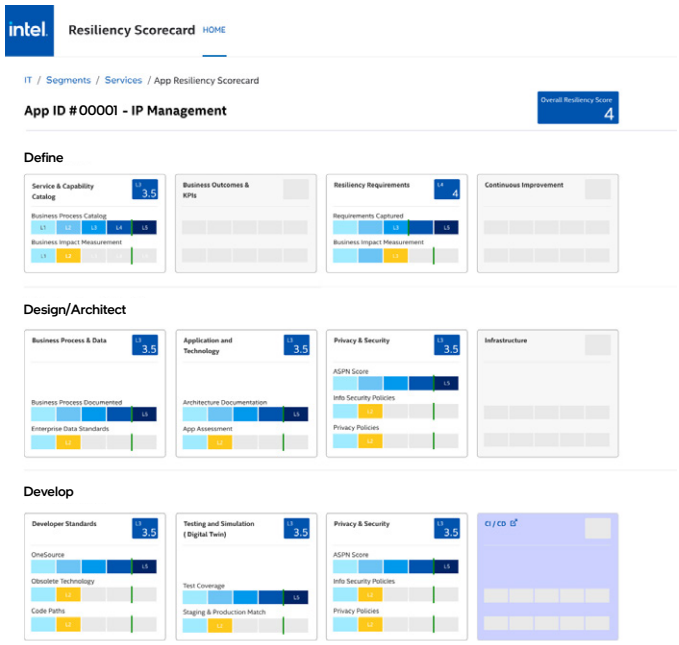

Figure 4. Our automated resiliency scorecard dashboard provides visibility into which resiliency gaps are associated with an application.

Resiliency Standards

To further simplify the design of resilient systems, we established detailed design standards for applications, databases, infrastructure, network, disaster recovery, and data protection. Here are a few examples of specific requirements:

- Application instances must be sized to handle full instance failure without degrading or impacting application performance within the DR/DC availability zone.
- Database instances must have affinity to server/host level with separation of instances on separate hosts.
- Servers must have redundant storage interconnects to separate fabrics.
- Redundant switch hardware must be installed in separate racks within the data center, preferably with separation of overhead fire/sprinkler proximity (that is, in separate rows).

These resiliency standards are integrated into our DevSecOps flow.

Summary of 2024 Achievements

In 2024, our primary focus was on strengthening the fundamentals of resiliency across all aspects of our organization while still driving changes across our strategic pillars. Our goal was to lay the groundwork and encourage teams to think about resiliency as they design and develop their solutions. For 2024, we made substantial progress in both technology and cultural resiliency.

- **Technology resiliency.** We established a technology resiliency standard for applications and infrastructure, and 100% of all Tier 1 and Tier 2 applications have resiliency gap plans in place.
- **Cultural resiliency.** We created governance structures, such as committees or working groups, to oversee and guide resiliency efforts across the organization. These structures include representatives from key functions and levels, to provide broad engagement and accountability. In addition, we developed policies and procedures that reinforced the importance of resiliency.

Results

Our resiliency journey has only just started, with plenty of work to accomplish in the coming 18 months. However, we can already tell that our framework, design principles, and resiliency standards are making a difference. For example, the IT incident in 2023 required eight days to fully recover our systems. In contrast, when a major security software vendor caused an inadvertent but major disruption in 2024, we fully recovered in under 1.5 days.

In 2025, we have already implemented several important resiliency-related deliverables:

- A business process to improve asset traceability.
- Unified compliance standards across resiliency, privacy and security, and commodities controls.
- A robust risk management solution that can support a dynamic risk model across enterprise solutions.

Next Steps

In addition to improving infrastructure resiliency throughout 2025, we will continue to partner with business units to strategically assess and streamline Tier 1 and Tier 2 applications, optimizing operational efficiency and enhancing overall business performance. We will also finish our data center resiliency audits, identifying risks and establishing plans to close the gaps.

Conclusion

Institutionalizing resiliency standards, enhancing operational processes, fostering a culture of pride and accountability, and cultivating the necessary skills and talent—these are ways we are building a robust framework that not only withstands disruptions but also drives continuous improvement and operational excellence. Our work to improve IT resiliency is essential for Intel's long-term success. Feel free to engage with the IT@Intel program to get more information about our resiliency design principles, standards, framework, and maturity model.

Related Content

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

- IT@Intel: Data Center Strategy Leading Intel's Business Transformation
- IT@Intel: Scaling Intel's Data Centers with Software-Defined Networking and Automation
- IT@Intel: Reliability Engineering Helps Intel Cut IT Manufacturing Systems Downtime in Half
- IT@Intel: Creating a System of Innovation
- IT@Intel: Data Center Facilities Risk Management

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

IT@Intel

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

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

Follow us and join the conversation on **X** or **LinkedIn.**

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