The Hidden Costs of Downtime

A $600 Billion Wake-Up Call

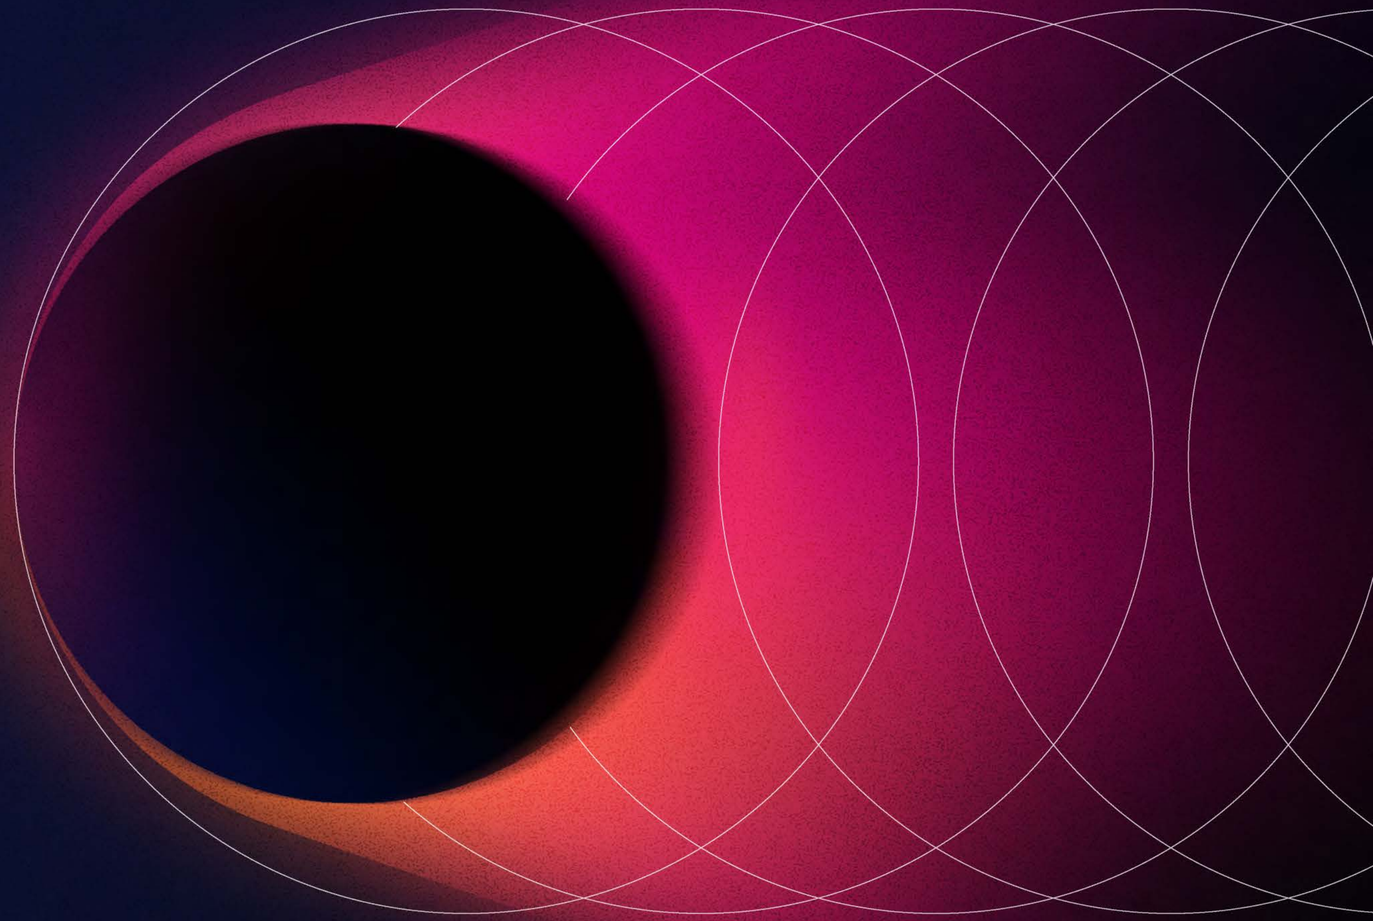

splunk>
a CISCO company

Contents

Executive foreword

Downtime is no longer a technical inconvenience; it's a business crisis. As digital ecosystems become more interconnected, every link in the chain — from third-party vendors to autonomous AI agents — introduces new risks, multiplying the consequences when things go wrong.

Our research reveals that the Global 2000's aggregate downtime costs have soared to $600 billion annually — up 50% from 2024. Alarm bells should be ringing in every boardroom.

The data also shows that outages come from many places, from infrastructure, application, and network failures to increasingly sophisticated cyberattacks. And although organizations are leveraging AI to reduce downtime, *every* technology leader we surveyed also admits that AI has caused at least one outage. Yet there is a path forward.

Downtime is inevitable; prolonged disruption is not. The most resilient organizations aren't the ones with the most tools or the boldest AI ambitions. They're the ones that align technology with business outcomes, empower their people with context, and design systems that bend — but don't break — under pressure.

Kamal Hathi
SVP and GM, Splunk

The bottom-line breaking point

The story is becoming all too familiar: A digital outage makes headlines, disrupting businesses, sometimes grounding flights, and even halting government services. Each incident is a stark reminder of the fragility of our digital infrastructure. For organizations, it should serve as a wake-up call — with a formidable price tag.

In partnership with Oxford Economics, Splunk calculated that the average cost of downtime[1] has reached a staggering $600 billion across the Global 2000[2], surging 50% in just two years[3]. On average, a company in this group now loses $300 million a year to unplanned outages and service degradation. This is a dangerous new financial threshold.

But the consequences of an outage extend far beyond immediate financial loss. Dips in stock price. Delayed product launches. Slow, corrosive brand damage. These are just some of the hidden costs that haunt an organization long after systems are back online, compromising long-term growth and competitive standing.

Organizations are taking notice. A resounding 79% of technology leaders consider downtime a higher organizational priority than just 12 months ago, fueled by a perfect storm of intense customer expectations for “always-on” services (74%), the higher cost of incidents (65%), rising cybersecurity threats (59%), and greater regulatory pressures (53%).

We also discovered that downtime is getting harder to evade. Today, outages stem from internal vulnerabilities, external cyber threats, AI-driven complexities, and growing interdependencies of third-party services. A failure in any link can set off a chain reaction across security, applications, infrastructure, and networks. No wonder organizations endure an average of 60 downtime incidents each year[4].

With the financial toll of downtime reaching a breaking point, the focus must shift from merely recovering after an incident to proactively building a foundation of digital resilience. This sort of predictive immunity will help minimize loss, create operational stability, and separate market leaders from those left struggling to keep up.

The rising cost of downtime

A single point of failure rarely stays contained. The ripple effect of an outage sends a $600 billion shockwave from the server room to the stock market, racking up fines, eroding brand trust, and stalling innovation.

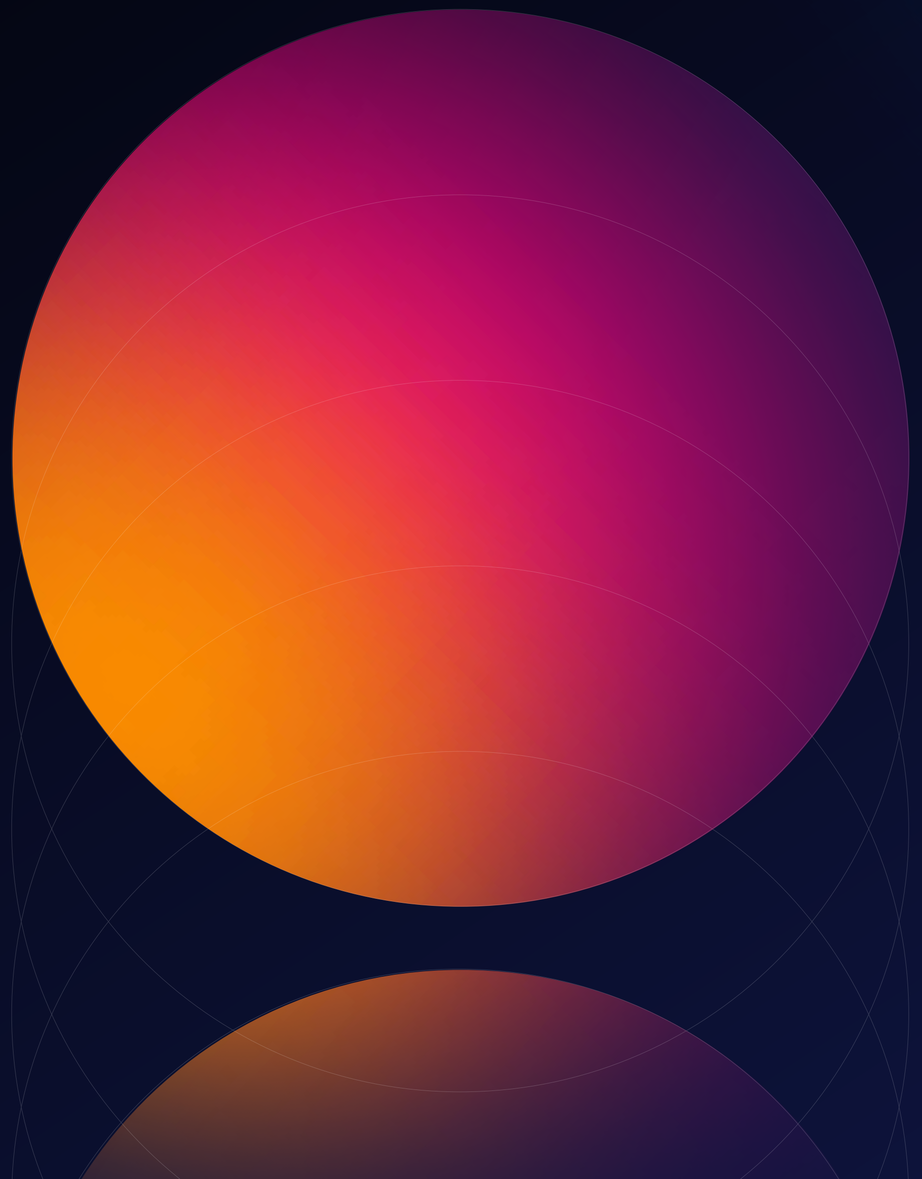

Downtime is an organization-wide problem. That's why we captured insights from security, ITOps, and engineering leaders on the front lines, as well as finance and marketing executives — including CFOs and CMOs — who see the brand and economic damage up close.

The aggregate cost of downtime for the Global 2000 is now $600 billion — up 50% in just two years. That works out to an average of $300 million per organization each year. The expenses add up quickly: Every minute of downtime costs $15,000, or over $900,000[5] an hour.

The majority of direct cost categories are higher in 2026 compared to 2024, signaling that downtime is becoming more expensive across the board. Lost revenue ($95M) is now nearly twice as high compared to 2024.

Regulatory fines and ransomware payouts nearly tripled according to the survey findings, reaching $51M and $40M, respectively. For technology executives, regulatory fines sting the most, with 57% calling them *very* or *prohibitively disruptive*.

Downtime direct costs have reached a fever pitch

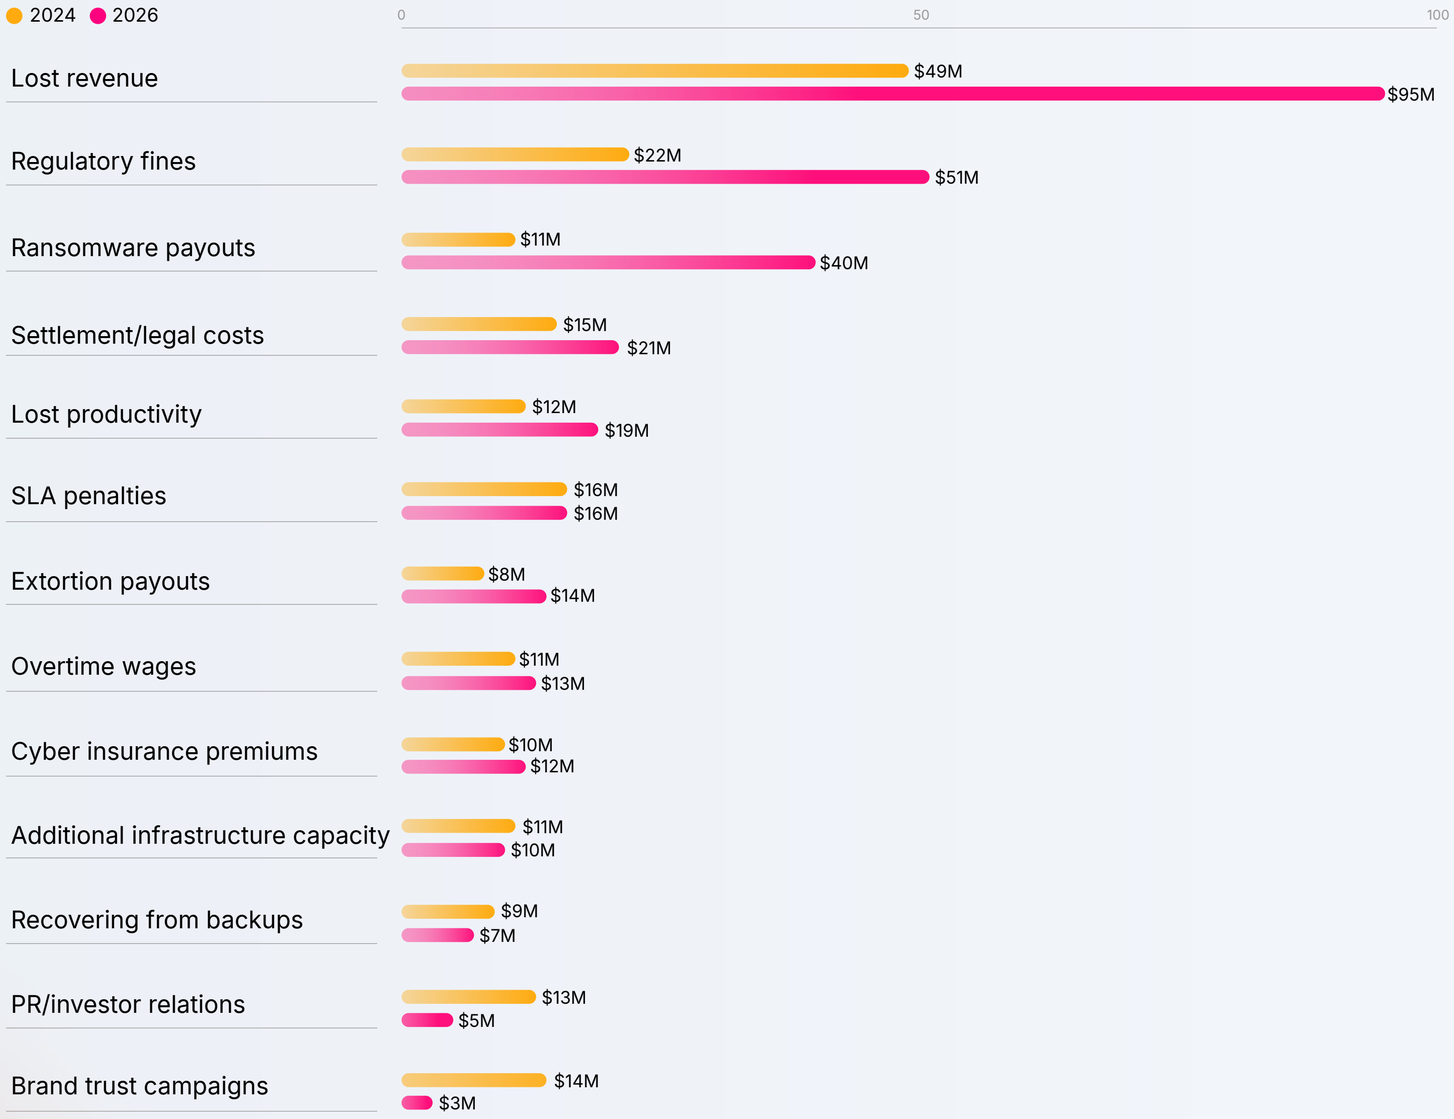

Cost per company in USD. Dollar amounts are rounded to the nearest whole number. Totals may not add up due to rounding of individual costs.

Digital resilience is no longer optional

Since 2023, fines levied by regulations like the EU's GDPR have escalated significantly, driven by more frequent enforcement. Meanwhile, the EU's Digital Operational Resilience Act (DORA) sets tough new standards for financial institutions and includes severe penalties for non-compliance. System downtime can trigger hefty penalties if an outage prevents access to essential services or personal data. Financial firms must meet new requirements for resilience testing, incident reporting, and managing third-party risk.

Alongside regulatory exposure, security incidents amplify the cost of downtime. While 62% of finance executives say they usually advise their CEO not to pay ransoms, ransomware payouts moved up from number eight to number three on the list of highest direct costs. Forty percent of finance executives admit they pay ransoms directly to attackers.

Ransomware has evolved from indiscriminate, opportunistic attacks into sophisticated, sustained operations. Experienced threat actors now conduct reconnaissance to identify revenue-critical systems and time their strikes for maximum impact and leverage. Attackers also preemptively encrypt or steal backup data, making restoration difficult or impossible. This strategic shift is rooted in financial savvy, as threat actors now calculate ransoms based on a victim's specific downtime costs. A multi-million-dollar demand seems reasonable compared to the cost of prolonged disruption.

But downtime costs aren't just higher on paper; they also feel more severe, with a growing number of technology executives viewing them as *very* or *prohibitively disruptive*.

Despite the rising costs of downtime and service degradation, Global 2000 companies are tightening their purse strings when it comes to repairing their public image in the wake of an incident. Average spending on PR/investor relations ($5M) and brand trust campaigns ($3M) has fallen significantly. According to Splunk SVP and GM Kamal Hathi, this is likely to prioritize immediate operational and compliance fixes. "It's a risky strategy," says Hathi, "as it ignores the growing impact of hidden costs like brand erosion. That kind of damage doesn't just delay recovery; it destroys customer trust that must be rebuilt from the ground up."

Direct costs: Disruption is growing more severe

Percentage of technology executives who consider a direct cost *very* or *prohibitively disruptive* (2024 versus 2026)

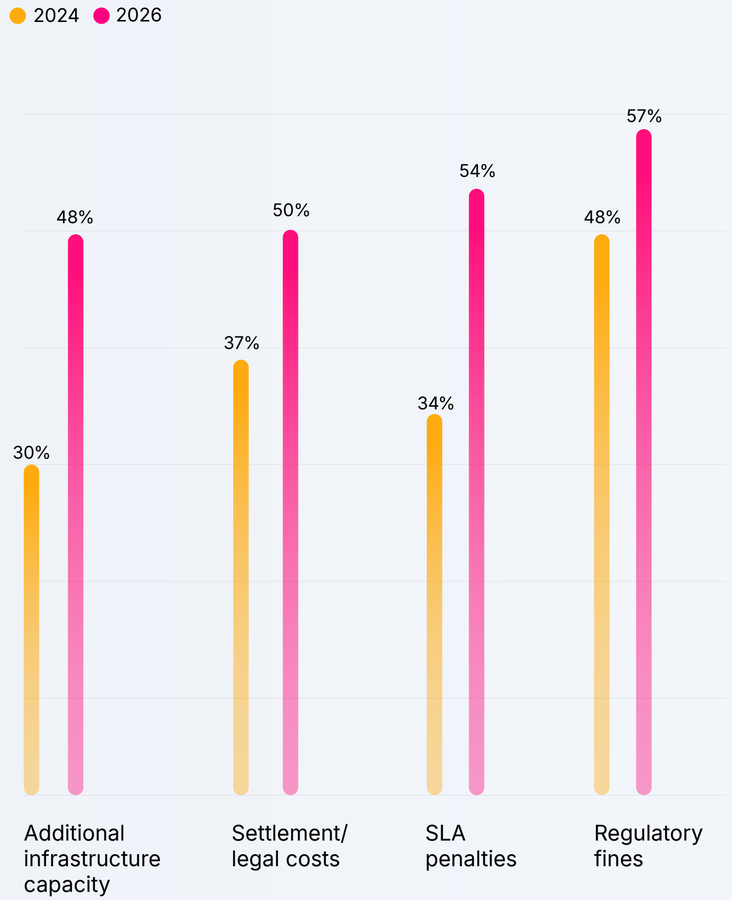

Downtime's blast radius is widening

Downtime triggers a chain reaction of consequences throughout an organization.

When services go dark, support queues surge. Frustration spills onto social media. Innovation stalls as security, ITOps, and engineering teams pivot to war rooms. The CMO prepares executive talking points and public statements, while the CEO and CFO monitor the falling stock price.

Even after an incident is resolved, this chain reaction continues: Engineering falls behind on the product roadmap. Finance must calm investors' nerves. Marketing measures a dip in brand perception and a spike in customer churn. It can take months of careful messaging and flawless service to rebuild the trust that was lost in mere moments.

Most often, the hidden costs of downtime take a human toll. From the SOC to the C-suite, survey respondents feel customers' frustration. A whopping 90% of technology leaders report increased demand for customer support; 76% of finance and 74% of marketing executives feel it as well. While it's possible to put a price tag on an outage, it's harder to calculate the lasting damage to an organization's most valuable asset: its people.

Hidden costs experienced most frequently

Percentage of technology executives who experienced the hidden cost at least once in the past year

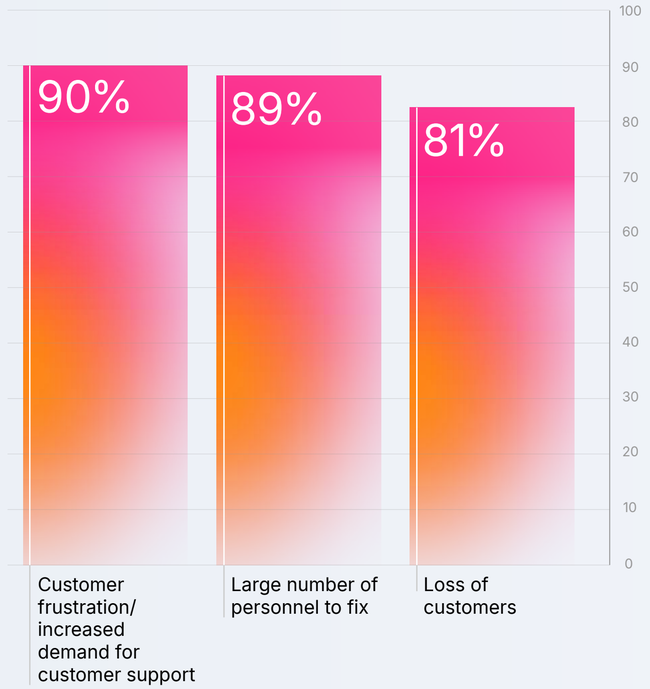

> It's the direct impact on your customers and employees that can turn a technical problem into a total business crisis.
>
> Kamal Hathi, SVP and GM, Splunk

The hidden costs of downtime hit harder

Compared to just two years ago, technology executives say the hidden costs of an outage are more severe. Public disclosure of a data breach is more than three times as disruptive and losing customers is more than twice as disruptive as in 2024. What's changed?

For starters, things got personal for executives. New SEC rules make leaders like CISOs and CFOs responsible for reporting security breaches. Failure to do so correctly could cost them financially and legally. Recent high-profile cyberattacks also demonstrated the serious consequences of a breach. These were not just security failures; they were business-crippling events.

In addition, users now face lower technical hurdles to change services, so patience for poor performance has worn thin. In competitive B2C markets, downtime and service degradation send customers fleeing to competitors. In B2B and regulated industries, they threaten future contracts. What once might have been a brief and recoverable dent in their reputation can now be an immediate loss of revenue. Organizations must tread lightly: 47% of technology leaders admit that customers are *often* or *very often* the first to detect downtime.

Hidden costs: Disruption is growing more severe

Percentage of technology executives who consider a hidden cost *very* or *prohibitively disruptive* (2024 versus 2026)

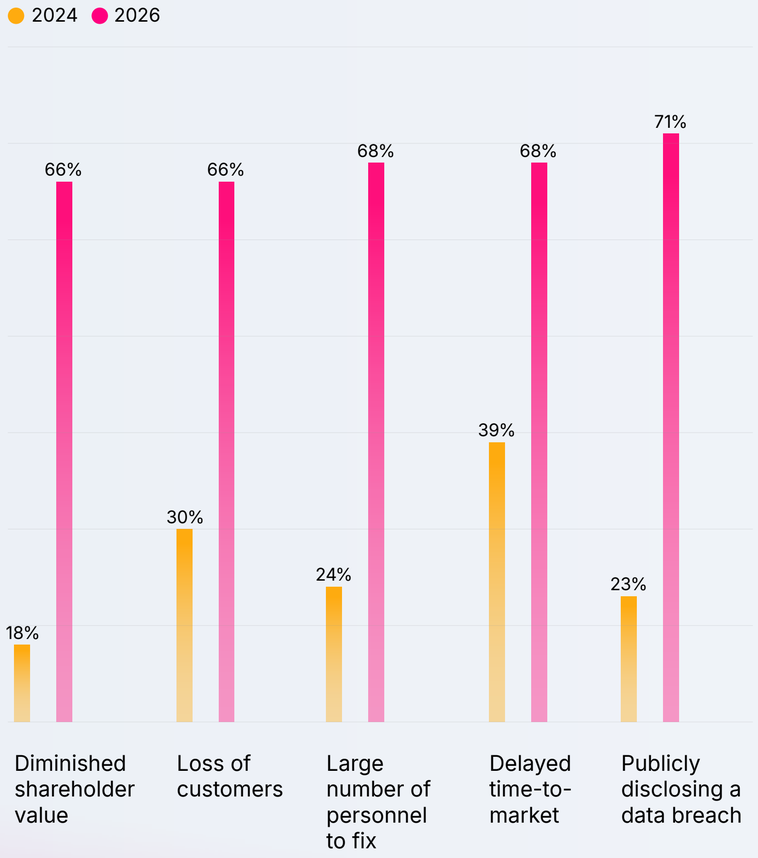

> **Watching companies suffer nine-figure losses makes the abstract risk of a cyberattack feel real and immediate.**
>
> Peter Sprenger, Field CTO, Splunk

Today, the perceived damage of many hidden costs has surged to a similar high-severity threshold, suggesting that organizations now view downtime as a systemic risk that afflicts every function equally.

Nearly 20% of marketing professionals report taking one full quarter for brand health to recover after an incident is remediated. Stock prices don't escape unscathed either. The survey reveals an average 3.4% drop after a downtime event[6]. In 2024, this figure was 2.5% on average, indicating that investors view operational resilience as a material risk factor, not just an IT issue. In fact, 49% of technology executives say shareholders have voiced concerns about downtime in the last 12 months, suggesting that downtime is now seen as an indicator of deeper organizational problems, such as poor risk governance, an underinvestment in modernizing digital infrastructure, and competitive vulnerability.

> **The gap between organizations who survive and those who thrive will be defined by their ability to remain resilient.**
>
> Jeetu Patel, President and Chief Product Officer, Cisco

The anatomy of an outage

The digital landscape is a perfect storm of internal vulnerabilities, external threats, and fragile third-party dependencies. With downtime origins spanning the entire technology stack, no domain is immune.

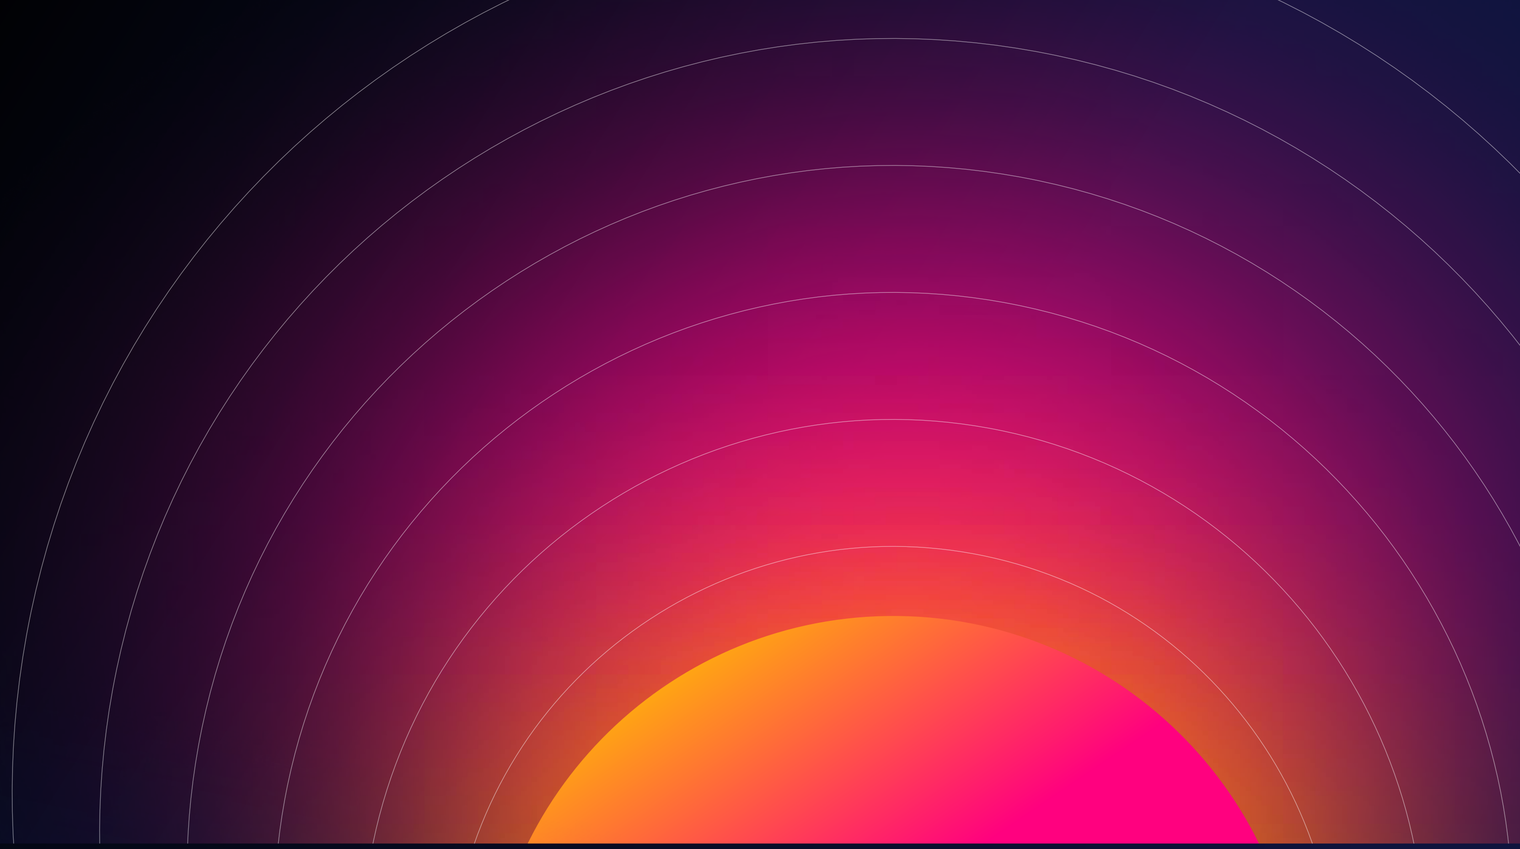

Organizations face an average of 60 downtime and service degradation incidents each year, originating throughout the entire technology stack — from network to applications and security. “The official count is 60,” says Greg Leffler, director of developer evangelism at Splunk. “But that only includes incidents organizations can detect. The true number is likely much larger, since countless smaller issues fly under the radar.”

While cyberattacks grab headlines, the data reveals human error — like software misconfigurations or application code errors — is still the main culprit of downtime. “This may be a direct result of the growing complexity of IT estates,” says Cisco ThousandEyes Principal Solution Analyst Mike Hicks. “As environments become more intricate and distributed, they create more blind spots, introduce more potential points of failure, and increase the likelihood of downtime-triggering mistakes.”

Disruption can strike from any direction

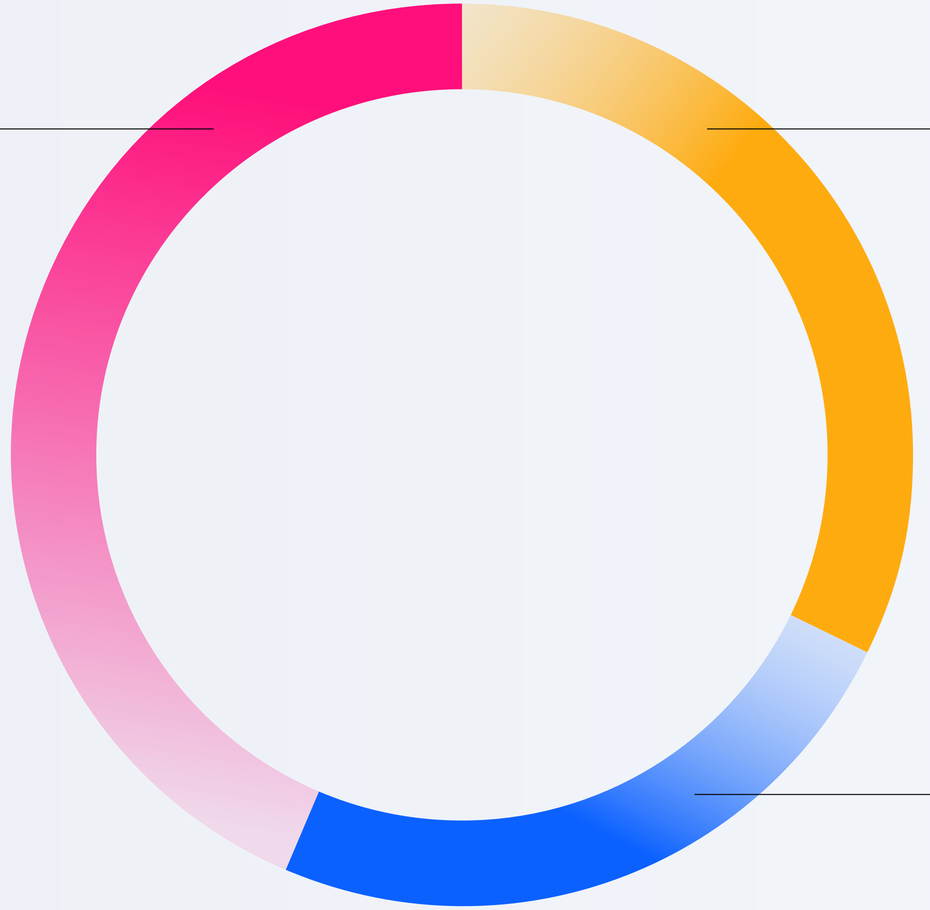

Percentages of the average total number of incidents across security-, application- or infrastructure-, and network- or IT environment-related causes, rounded to the nearest whole number.

The most common causes of downtime

1. ITOps-related human error
2. Security-related human error
3. Software failure (e.g. application bug)
4. Networking-related SaaS dependencies
5. Public cloud failures
6. Network congestion
7. ITOps-related third-party outage
8. Endpoint and IoT devices
9. Routing failures/ DNS errors
10. Phishing attack

Derived by number of causes queried multiplied by the average number of times experienced in the past 12 months.

Security incidents

While human error is most common, phishing and malware attacks often serve as the front door for the most severe types of unplanned downtime. They can introduce ransomware that locks up systems or allow attackers to sabotage infrastructure. Beyond the initial breach, organizations often endure prolonged outages when attempting to re-secure the network.

In particular, phishing has grown substantially more sophisticated and common; nearly half (49%) of security leaders now face these attacks *often* or *very often* compared to just 30% in 2024. In *The CISO Report*, virtually all CISOs (95%) said the growing sophistication of threat actor capabilities poses the greatest threat to their cybersecurity strategies.

The era of easily spotted phishing emails is over. With the help of AI, attackers craft highly convincing, personalized messages that can fool even the most vigilant employee and bring digital services to a grinding halt. Yet according to the *Cisco AI Readiness Index*, only 36% of organizations are highly aware of how malicious actors use AI to make these attacks more sophisticated.

Meanwhile, the perceived frequency of downtime caused by SaaS and other third-party application issues has nearly tripled since 2024, with 56% of security leaders now experiencing them *often* or *very often*. "When an issue originates from a SaaS or third-party application, security teams must navigate a complex web of cloud, on-prem environments, and other infrastructure they often have no visibility into," says Splunk Field CTO Peter Sprenger. "The problem is intensified by siloed ownership and a lack of visibility into critical dependencies."

But the greatest obstacle in resolving downtime isn't technical; it's organizational. Downtime doesn't recognize departmental boundaries, and yet investigations remain siloed. Security, ITOps, and engineering teams tackle the same crisis from different perspectives, with different data, and without shared context. In the event of a security breach, these siloes are true vulnerabilities.

This lack of shared visibility is perhaps why roughly one-third of security leaders admit downtime is *often* or *very often* initially misclassified as an IT issue. When a security breach is miscategorized as an IT incident, that delay gives attackers a critical head start. Every minute lost allows them to burrow deeper into the network, escalate privileges, and expand their foothold; exponentially increasing the cost, scope, and complexity of the event.

36%

of security leaders admit downtime is *often* or *very often* initially misclassified as an IT issue

ITOps and engineering incidents

After human error, software failure and third-party outages (API, CDN, SaaS) are the most common causes of application- or infrastructure-related downtime or service degradation.

Engineers are shipping code faster than ever. “Demand for accelerated innovation is placing immense pressure on development teams, forcing them to compromise on testing,” says Greg Leffler, director of developer evangelism at Splunk. “Code is being thrown into production to meet deadlines, often based on a ‘vibe’ rather than solid validation. This rush to deployment inevitably leads to more software failures, more system instability, and more disruptive downtime.”

How can ITOps and engineering teams manage these risks while simultaneously meeting business needs? According to Cisco ThousandEyes Principal Solution Analyst Mike Hicks, “The solution is the ability to see the whole picture when something goes wrong. The faster you can spot cause and effect, the faster you can fix the problem.”

Like their security colleagues, reliance on external providers has become a primary source of instability for ITOps and engineering teams, underscoring how the fate of an organization is tied to its vendors. Even applications that appear internally owned depend on third-party services like payment processors, authentication, and AI APIs. Visibility into unowned networks and external dependencies isn’t just a nice-to-have; it’s a prerequisite for digital resilience.

Third parties are increasingly to blame for downtime

Percentage of ITOps and engineering leaders who *often* or *very often* experience downtime caused by a third party (2024 versus 2026)

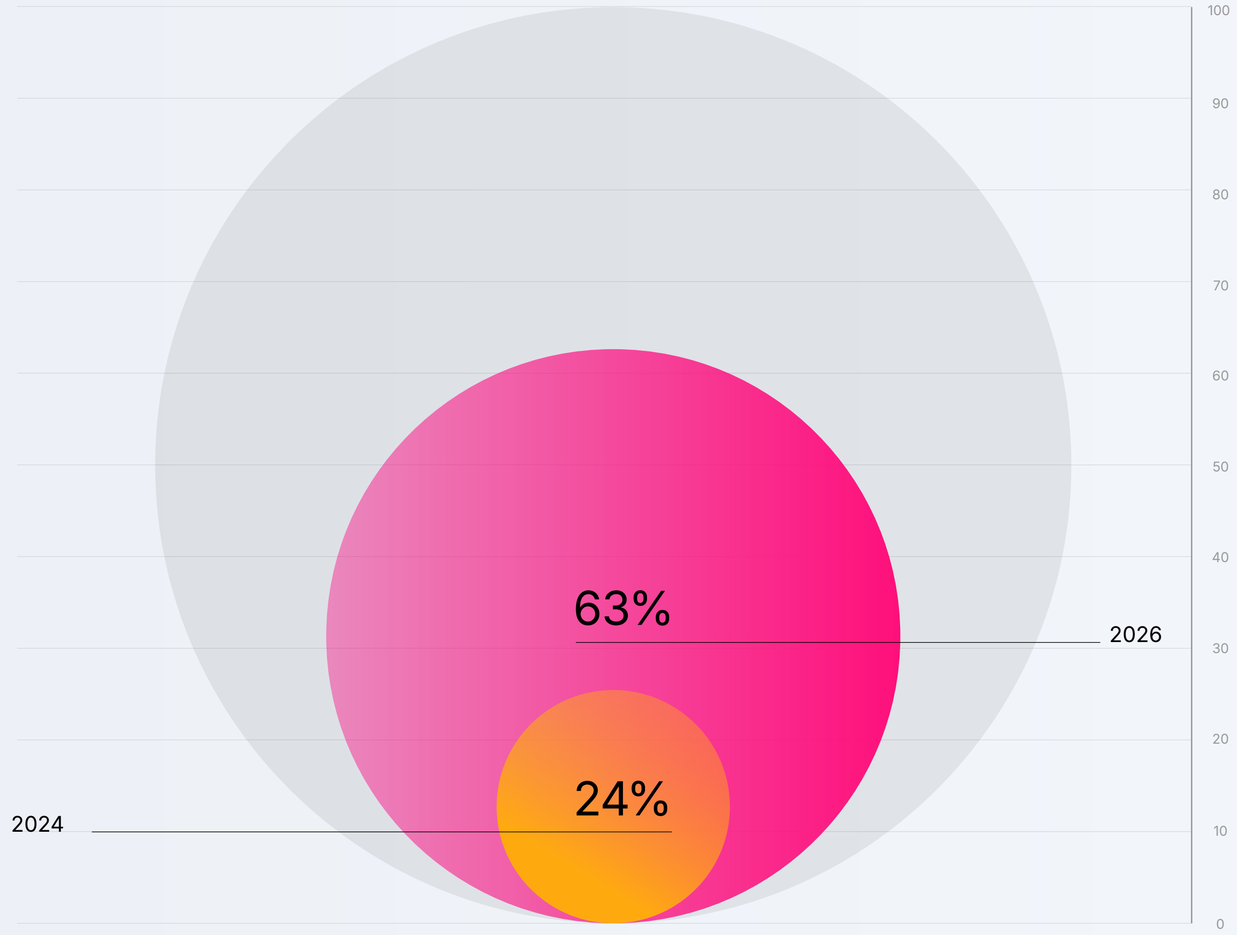

Networking incidents

The central paradox of the public cloud is that while its benefits are immense, its failures are the most painful. Among all network-related disruptions, public cloud failures cost businesses the most.

Beyond the cloud, downtime also stems from SaaS dependencies, network congestion, and challenges managing the proliferation of endpoint and IoT devices.

The growing reliance on SaaS means that companies are staking their customer experience on external service chains they neither manage nor fully understand. Without visibility into or control over these dependencies, they are oblivious to potential outages. The challenge is turning this "black box" of external services into something that can be proactively monitored and managed.

"Traditional monitoring stops at your firewall," says Greg Leffler, director of developer evangelism at Splunk. "End-to-end network visibility provides a map of the entire digital ecosystem — from your infrastructure to a vendor's data center, across multiple ISPs, to the end user — allowing teams to pinpoint whether a failure is internal, with a vendor, or somewhere in between."

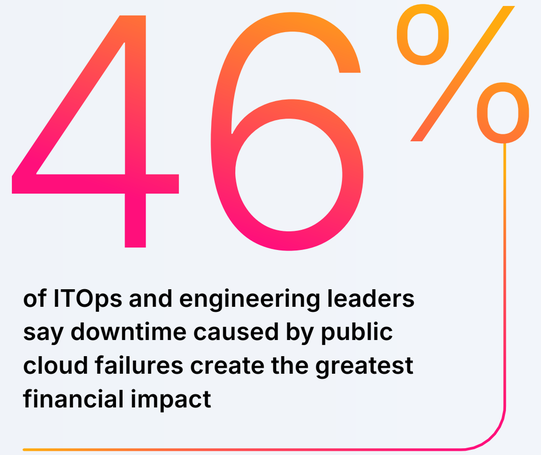

of ITOps and engineering leaders say downtime caused by public cloud failures create the greatest financial impact

Building resilience in the AI era

As downtime becomes more frequent, technology leaders are now prioritizing end-to-end visibility across all IT domains. And while AI speeds up diagnostics, human oversight remains essential for true digital resilience.

Thirty-one percent of technology leaders admit that despite significant tool spending, downtime is more frequent. To Cisco ThousandEyes Principal Solution Analyst Mike Hicks, this means either problems are evolving faster than solutions can address them, or organizations are investing in the wrong approach to resilience.

“Simply adding more monitoring tools cannot solve architectural complexity, which requires unified visibility and cross-domain context,” says Hicks. There’s also an expanding reliance on third-party services — which creates blind spots, drives up costs, and increases the likelihood of downtime.

This could explain why executives believe their budget is no longer sufficient. The percentage of technology leaders who call their cybersecurity tool budget adequate fell from 96% in 2024 to 61% in 2026, while the same sentiments for observability tool budgets dropped from 88% to 63%. Perhaps the status quo is no longer sustainable as rising costs and complexity are forcing companies to seek smarter solutions.

Downtime is persistent

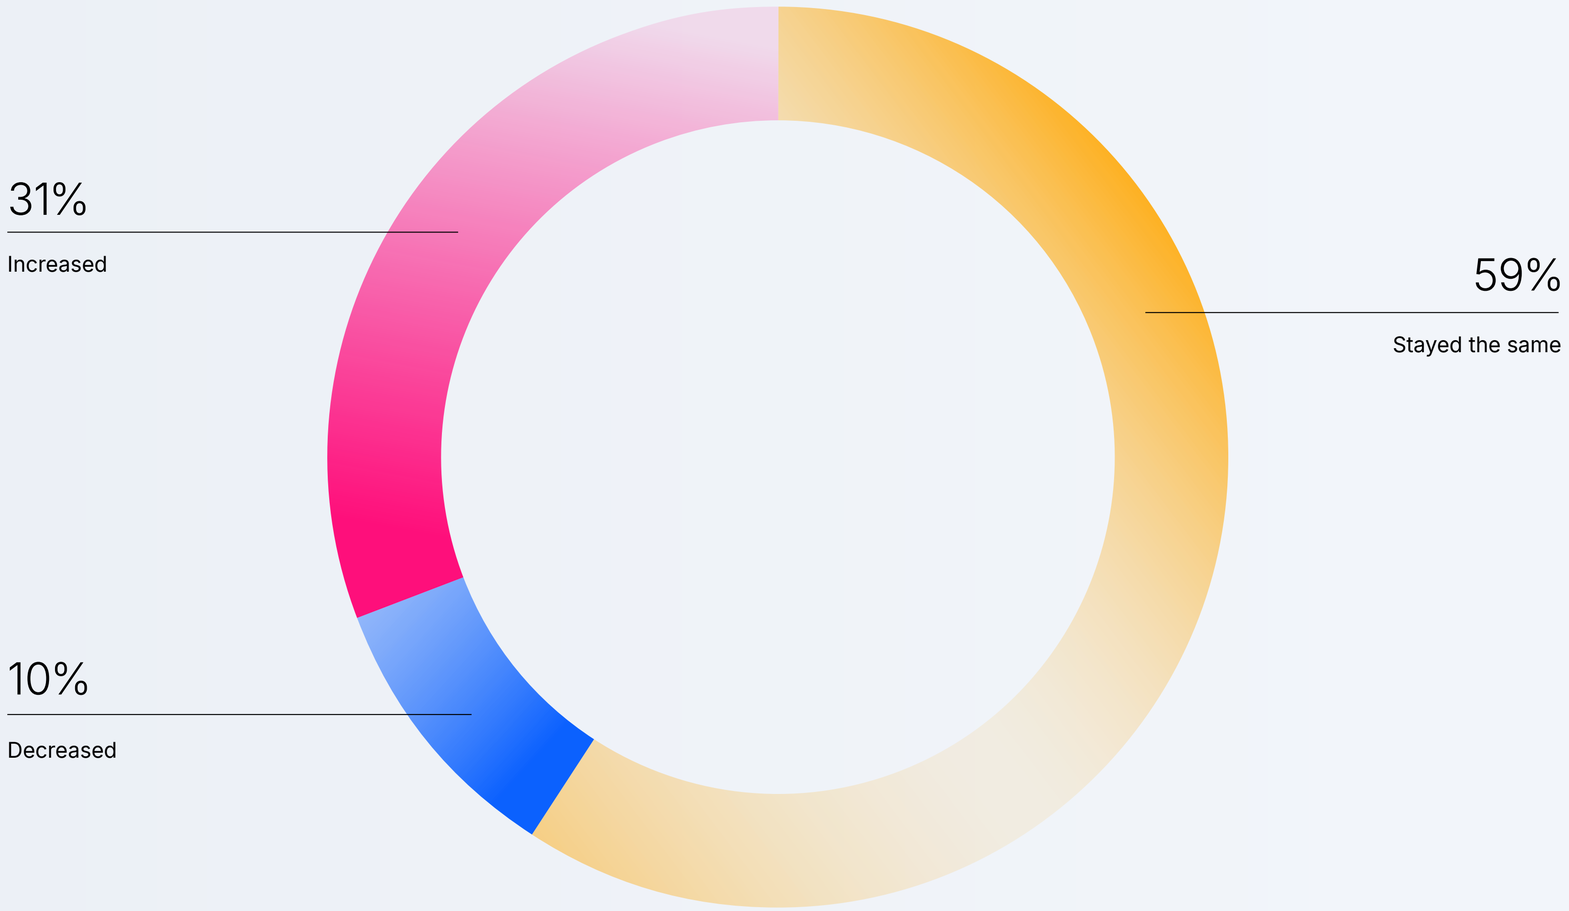

Organizations seek visibility to reduce downtime

ITOps and engineering leaders rank end-to-end observability as their top investment priority to improve infrastructure resilience.

Splunk Director of Developer Evangelism Greg Leffler sees this as a major shift. “In the past, ITOps and engineering primarily directed their budgets toward cloud services, data center upgrades, and disaster recovery.” Observability was not a major investment area. Today, this trend has reversed. Leffler continues: “Observability is now the top priority, often taking precedence over spending on the infrastructure itself.”

Automation ranks as the second-highest priority, suggesting a new approach. Marrying observability and automation may allow organizations to build more resilient operations and reduce manual intervention.

In all, technology executives recognize the need to visualize the entire dependency chain. In fact, among respondents with the lowest downtime costs[9], a whopping 98% say end-to-end visibility is *very* or *extremely important* for reducing downtime. Nevertheless, technology respondents admit complete visibility is rare across all IT domains.

Top investment priorities to build resilience

Percentage of ITOps and engineering leaders who identify the investment as a priority

AI tools are transforming downtime

Organizations are turning to AI as a solution to downtime and service degradation. Its ability to automate complex tasks, predict failures, and dramatically shorten response times makes it a leading contender to bolster digital resilience.

According to the survey, companies spend a median of $24.5M annually on AI tools that prevent and respond to downtime. The data also shows that the race to adopt agentic AI[10] is well underway. A striking 44% of technology respondents say they're already using it — signaling belief in its potential to prevent unplanned outages.

Across security, ITOps, and engineering, the most common use case for AI is incident triage and root cause analysis. For these tasks, 54% of respondents use generative AI[11], while 27% use AI agents. This addresses two critical pain points head on: slow investigations that prolong business impact and failure to locate the true source of an outage. Only 38% of technology executives report *always* finding the root cause of a downtime incident. This could be because some outages are too brief to warrant investigation. In other cases, teams could lack visibility into external dependencies or be so overwhelmed with alerts that they can only address the most critical events.

AI is transforming incident response in two distinct ways. Generative AI assistants help human investigators rapidly correlate data, summarize complex incidents, and suggest next steps to accelerate analysis. AI agents take it one step further: They can independently diagnose issues and execute common fixes, such as performing code rollbacks, while escalating more critical actions for human approval.

When it comes to minimizing the hidden costs of downtime, the data reveals a clear advantage for organizations using AI to manage incidents and orchestrate workflows. For instance, 74% of these "AI Workflow and Triage Experts[12]" avoided the need to publicly disclose a data breach last year, compared to just 54% of non-experts. They are also nearly three times more likely to say they've never lost customers due to downtime (42% versus 15% non-experts).

> **Technology leaders are prioritizing AI solutions that can analyze vast amounts of data, surface critical insights, and automate repetitive tasks — but leave the final decision-making and strategic oversight in human hands.**
>
> DJ Sampath, SVP, AI Software and Platform, Cisco

The AI paradox

Respondents' AI investment priorities suggest they are applying the technology to the right tasks. "AI is best leveraged for faster analysis, decision support, and the automation of lower-level tasks, but organizations should leave critical decision-making to humans," says Splunk Global Field CTO Cory Minton. "The goal is not to eliminate human involvement, but to make it faster, smarter, and more impactful. The future of digital resilience lies in this human-to-agent collaboration, where AI serves the expert, not the other way around."

While 56% of users say AI has reduced the overall risk of downtime, the technology, once again, proves to be a double-edged sword: Every technology leader surveyed admits that their organization has experienced some form of AI-related downtime. To make matters worse, 66% reveal that employees use unapproved AI tools (i.e. shadow AI) to assist with their jobs, introducing an added layer of risk into their environment. According to CISOs surveyed in *The CISO Report*, shadow AI is a top three AI-related concern, likely because it presents a direct challenge to the governance, control, and integrity of an organization's security operations.

Top AI investments aimed at reducing downtime

85%	AI-driven security automation
65%	AI-powered observability
47%	Predictive analytics for capacity/load management
37%	Network automation

> **The foundation for AI oversight is machine data — the logs, metrics, and traces that let humans see what an AI did, detect issues early, and correct course before small errors become outages.**
>
> Hanlin Fang, VP of Product Management, Splunk

“In my view, these figures reflect where we are in the AI adoption curve,” says Hanlin Fang, VP of product management at Splunk. Organizations are rushing to deploy AI systems without proper guardrails and oversight like clearly defined owners and escalation paths. To Hanlin, the organizations that successfully minimize AI-related downtime aren’t the ones with the most sophisticated technology. They’re the ones with humans in control with continuous monitoring and fast intervention when outcomes drift.

Roughly one-fourth of organizations have suffered attacks like prompt injections and data poisoning, proving that bad actors are actively probing AI systems for weaknesses, exposing critical gaps in security. Technology leaders are concerned. A significant majority (77%) believe that cybercriminals armed with generative AI will increase downtime at their organization, while 64% believe it already has.

AI-related downtime is becoming increasingly common

Percentage of technology executives who experienced downtime caused by the following AI-related issues

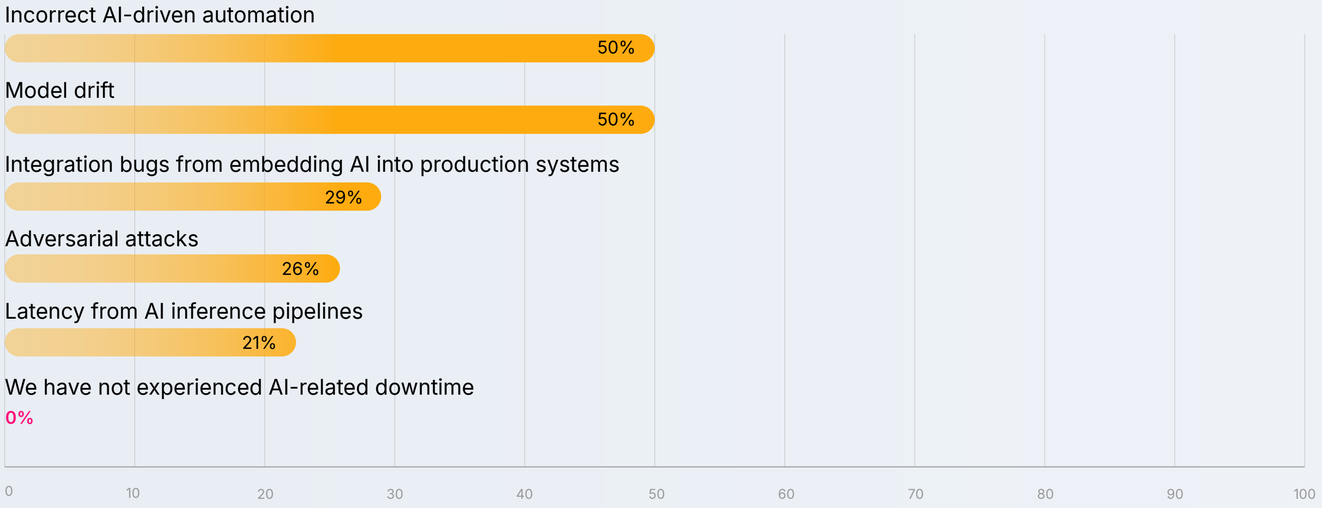

68%

of technology leaders worry their AI agents will behave unpredictably and cause downtime

Advancing beyond the outage

Resilience isn’t simply about keeping systems online. It’s about protecting trust, enabling growth, and ensuring the business can move forward — even when things go wrong. The findings in this report are evident: The most resilient organizations are bouncing back faster, learning from their mistakes, and recognizing the clear competitive advantage of predictive immunity. We’ve boiled the findings down to four key strategies to help you do just that.

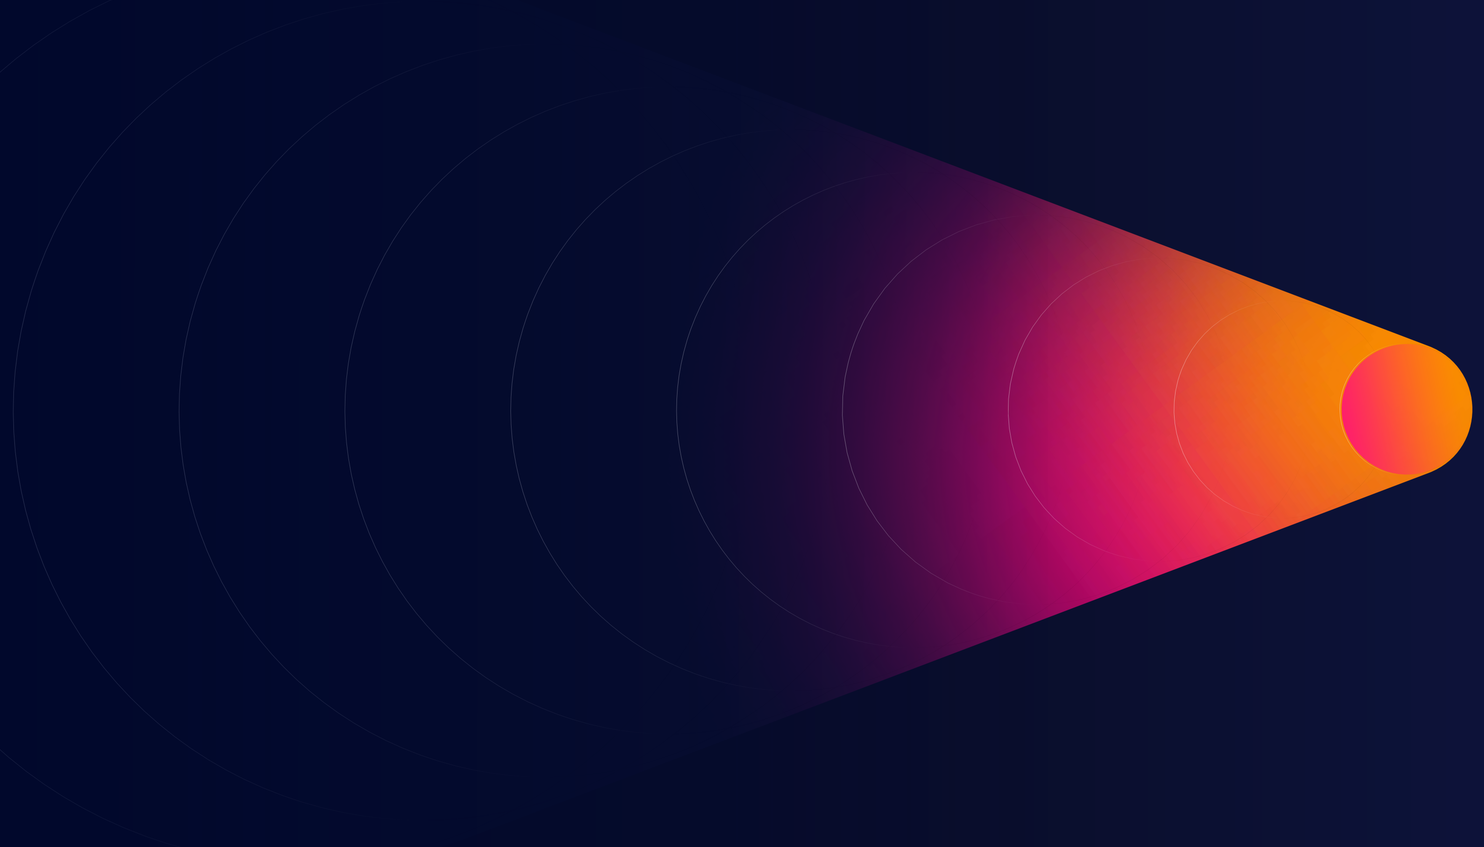

1. **Treat downtime as a business risk.** When outages are seen as only technical problems, their true impact is overlooked. Revenue dips, stock volatility, regulatory exposure, and brand damage demand executive attention. Translate downtime into business language — mapping incidents to profit impact, recovery timelines, and customer trust. Include downtime risk in board-level and executive risk discussions and align security, ITOps, engineering, finance, and marketing around shared resilience metrics.

2. **Design systems for humans.** Human error remains the leading cause of downtime in 2026. Complexity amplifies this reality — more tools, more services, more configurations, and more opportunities for mistakes. Reduce errors with safe-by-default changes, pre-flight checks, automated rollbacks, and guardrails for risky operations. Standardizing change management and deployment practices will ensure consistency, accountability, and controlled execution across teams. Resilient organizations don't expect perfection from people. When mistakes happen, fast detection and clear context beats blame — and shortens MTTR.

3. **Make detection and root cause analysis a team sport.** Downtime doesn't abide by org charts. Security, ITOps, and engineering teams frequently approach the same incident with different data, different tools, and different priorities. Digital resilience starts by connecting platforms, unifying data with the network path, and establishing a shared data source with shared context. A data fabric architecture provides an additional connective layer, strengthening collaboration and accelerating resolution by giving all teams access to information without costly movement or duplication.

4. **Use AI to accelerate insight — not sideline judgment.** AI has become a powerful ally in the fight against downtime. But left unchecked, it introduces new risks. When deploying AI to speed up incident detection, root cause analysis, or prioritization, always pair AI's analytical speed with expert human judgment and oversight. Require validation before remediation or automated actions are taken. At the same time, implement robust governance frameworks to manage shadow AI and secure AI-driven workflows. This model ensures speed never comes at the expense of trust, accountability, or safety.

> The key to transforming resilience from aspirational to operational is a unified platform that combines security, networking, and observability so organizations can detect problems earlier, respond faster, and recover before customers feel the impact.

Tom Gillis, SVP and GM, Infrastructure and Security Group, Cisco

Examining industry downtime expenses

The 2026 industry cost rankings[13] tell a story of shifting risk. The information services and technology sector ($402M) logged the highest downtime costs. When a major cloud provider, SaaS platform, or API service goes down, it's not just their revenue on the line; it's the revenue of every business dependent on their services. Their failures create a multiplier effect, causing cascading outages and massive SLA penalties across their entire customer base. Security failures for this sector are especially punishing: Financial losses from these lapses now account for more than a quarter of all downtime costs.

Energy and utilities ($364M) ranks second, reflecting the critical nature of its services and its reliance on digital systems to control physical infrastructure. The primary risk now is the convergence of IT and operational technology (OT), which exposes essential systems to new cyber threats, like a catastrophic attack on the power grid. According to Cisco's *State of Industrial AI Report*, poor IT/OT collaboration and increased AI adoption intensifies security challenges in this sector.

In contrast, while retail ($357M) once topped the list, recent existential threats like Black Friday and Cyber Monday outages may have forced the industry to invest heavily in multi-cloud strategies, CDN redundancy, and failover capabilities. Retail learned to architect for digital resilience out of competitive necessity.

For industries like financial services, manufacturing, and transportation and logistics, the cost of downtime is increasingly defined by legal and contractual penalties, not just lost revenue. This shift reflects intensifying regulatory enforcement, where broken SLAs and safety failures trigger severe financial and legal consequences.

Across all industries, system upgrades and damage control are consistently secondary concerns, accounting for only a small fraction of total downtime costs.

Downtime costs ripple through every industry

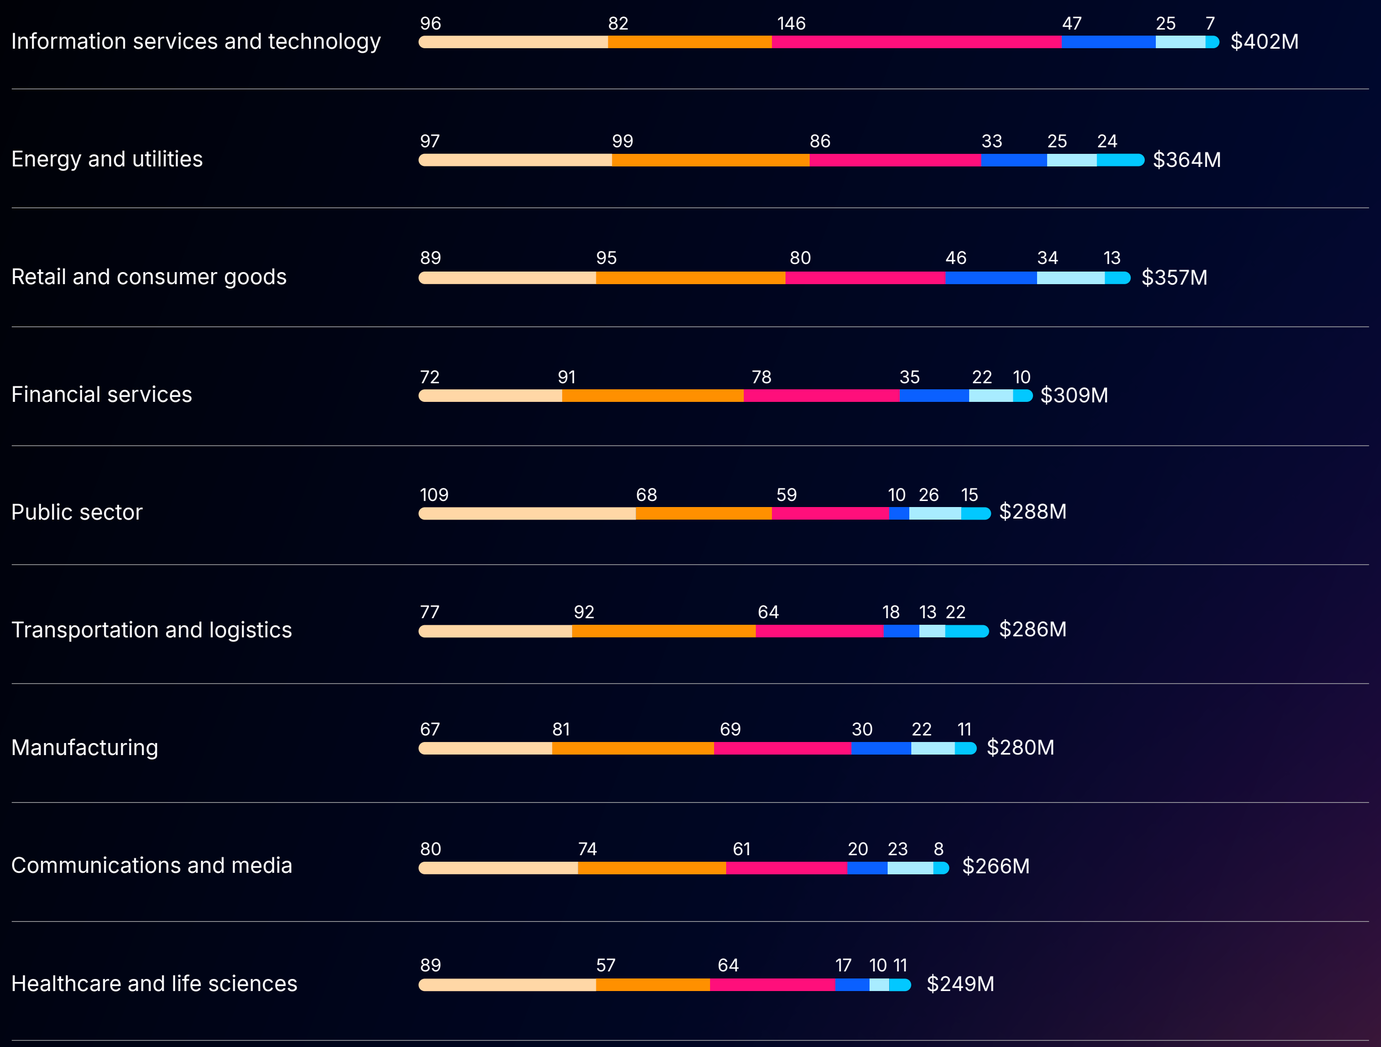

Cost category detail in $M:

Lost revenue	
Contractual/legal	Regulatory fines, SLA penalties, settlement/legal costs
Security lapse costs	Ransomware payouts, cyber insurance premiums, extortion payouts
Staffing/productivity	Overtime wages, lost productivity
Upgrade needs	Additional infrastructure capacity, recovering from backups
Damage control	Brand trust campaigns, PR/investor relations

Dollar amounts are rounded to the nearest whole number.
Totals may not add up due to rounding of individual costs.

Regional variations in downtime

Downtime is most expensive for organizations in EMEA ($354M)[14]. As the most heavily regulated region, these elevated costs are largely driven by the growing financial impact of compliance failures[15], which account for approximately $110M. While lost revenue ($87M) creates a substantial dent, it is only the third-largest cost driver in the region. Instead, security lapse costs[16] like ransomware payouts and cyber insurance premiums ($89M) surpass lost revenue, underscoring how regulatory exposure and cybersecurity risk compound the financial impact of outages in EMEA.

Lost revenue is the dominant cost driver for APAC ($104M), suggesting operational disruptions, supply chain dependencies, and market volatility directly affect revenue generation during outages. The region also suffers significant contractual/legal expenses ($75M), the second highest globally. While regulatory pressure may be lower than in EMEA, contractual commitments and service obligations carry a sizeable penalty when downtime occurs.

Companies in North America see most of their downtime costs go to security lapse expenses ($82M); cyber insurance premiums in this region tend to be particularly steep. Notably, staffing and productivity losses[17] ($39M) peak in North America, which generally sees the highest cost of labor globally.

Overall downtime expenses are lower and more evenly distributed in LATAM ($197M). However, lost revenue ($55M) and security lapse costs ($54M) still emerge as primary drivers, indicating that despite lower totals, organizations in the region face many of the same fundamental risks.

Across all regions, security lapse costs consistently represent a significant share of downtime expenses, signaling a universal rise in both the frequency and financial impact of cyberattacks. In contrast, upgrade needs[18] represent a relatively small share of downtime costs, indicating that, globally, organizations focus more on reacting to incidents than investing in proactive resilience.

Downtime costs transcend borders

Strengthen your digital resilience with Splunk

The CISO Report 2026:
From Risk to Resilience in the AI Era

From adopting agentic AI to measuring security ROI and navigating an expanding role, discover how 650 global CISOs are maintaining resilience and empowering their organizations to thrive in the AI era.

Download the report

Perspectives by Splunk
— by leaders, for leaders

Looking for more insights on executive strategy, cybersecurity, and observability trends? Learn how leaders tackle today's most pressing challenges including AI and the changing compliance landscape.

Get executive insights

Methodology

Oxford Economics fielded a hybrid survey using CATI (Computer Assisted Telephonic Interviewing) and online methods. The fieldwork captured responses from 2,000 executives from Global 2000 companies. Businesses from 20 countries are represented from APAC, EMEA, North America, and LATAM. Respondents hail from nine industry groups: financial services, retail and consumer goods, public sector, manufacturing, energy and utilities, healthcare and life sciences, information services and technology, transportation and logistics, and communications and media. Respondents come from technology (including security, IT, and engineering titles), finance (including Chief Financial Officers), and marketing functions (including Chief Marketing Officers).

How Oxford Economics calculated the costs of downtime

Oxford assessed the costs of downtime for the Global 2000 by adjusting survey responses to match the characteristics of the Global 2000 in 2025. They used survey responses to estimate downtime costs relative to revenues (i.e., as a percentage of revenues to adjust for differences in company size among respondents), calculated the typical cost per unit of revenue by taking the median of the previous metric for each country in the study sample, and then combined revenue data from fiscal year 2025 for each company in the Global 2000 with the median value for the corresponding country. This scaling process helped align the survey responses with the size and regional distribution of Global 2000 companies.

References

1 For the context of our survey, downtime was defined as: Any type of service degradation (such as latency/slowness), as well as service unavailability to end users of critical business systems.

2 “The Global 2000 ranks the largest companies in the world using four metrics: sales, profits, assets, and market value.” (Forbes)

3 Cost increase is a comparison of the 2024 and 2026 survey results derived from self-reported survey data.

4 Figure derived from self-reported survey data.

5 Figure based on average total number of incidents, average total cost of all incidents, and average MTTR of incident type.

6 Average stock price drop is derived from self-reported survey data from n300 respondents in finance roles. “On average, by what percentage does your organization’s stock price decrease after a single downtime incident?”

7 Issues relating to networking equipment, routing, DNS, internet paths between services, ISPs, endpoint and IoT devices.

8 Issues relating to servers, storage, databases, cloud instances, containers, application code, microservices, and APIs.

9 Defined as the 10% of survey respondents with the lowest overall downtime costs. Calculated by the sum of all direct cost categories divided by their company’s total revenue.

10 For the context of our survey, agentic artificial intelligence (i.e., agentic AI) was defined as: Artificial intelligence systems designed to act autonomously, make decisions, and perform tasks with limited or no direct human supervision.

11 For the context of our survey, generative artificial intelligence (i.e., generative AI or GenAI) was defined as: A kind of artificial intelligence that creates content, such as text, images, audio, or video, without human intervention or via human prompt. By utilizing algorithms and machine learning techniques, such as deep learning, generative AI can create outputs that resemble the ones produced by humans. ChatGPT is an example of generative AI.

12 “AI Workflow and Triage Experts” are defined by a scoring matrix applied to two categories (“Workflow orchestration across teams” and “Incident triage and root cause analysis”). Those with the highest scores represent the leader group.

13 All downtime costs listed represent the average per company in each industry.

14 All downtime costs listed represent the average per company in each region.

15 Regulatory fines, SLA penalties, settlement/legal costs

16 Ransomware payouts, cyber insurance premiums, extortion payouts

17 Overtime wages, lost productivity

18 Additional infrastructure capacity, recovering from backups

About Splunk

Splunk, a Cisco company, helps make organizations more digitally resilient. Leading organizations use our unified security and observability platform to keep their digital systems secure and reliable. Organizations trust Splunk to prevent infrastructure, application, and security incidents from becoming major issues, recover faster from shocks to digital systems, and adapt quickly to new opportunities.

Keep the conversation going with Splunk.

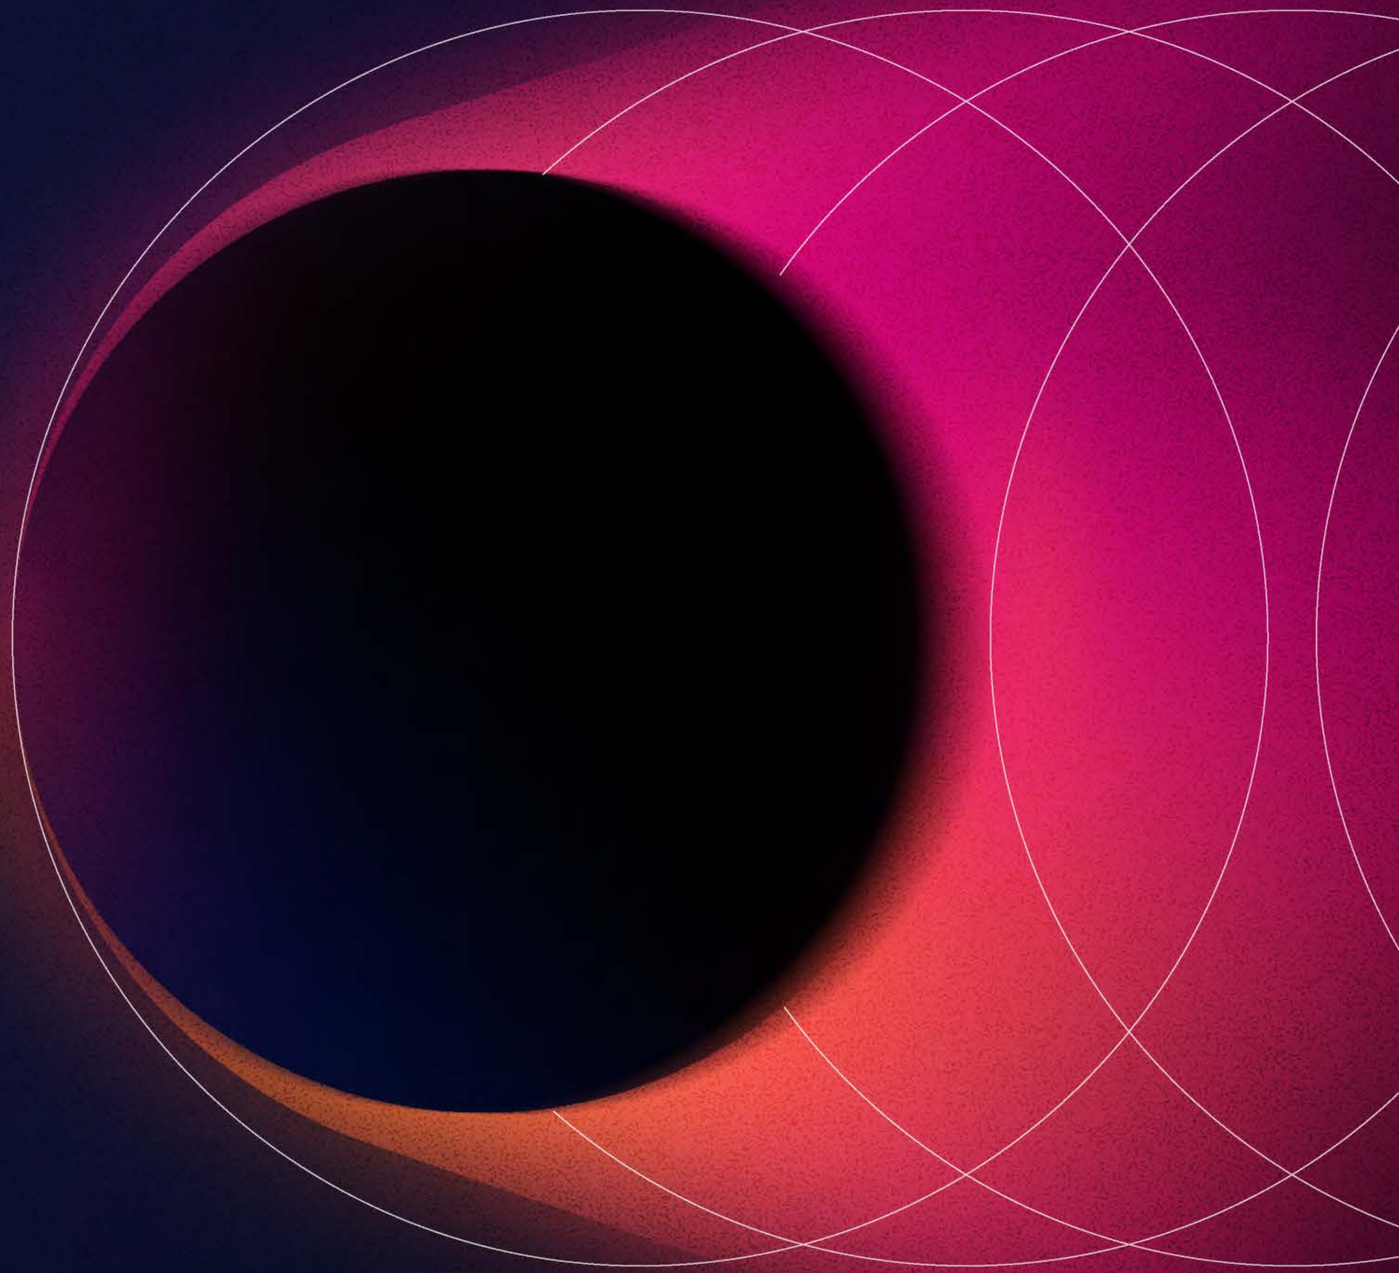

splunk> a CISCO company

26_CMP_report_The-hidden-cost-of-downtime-report-2026_v13