FENIX24 · ANNUAL RESEARCH REPORT

State of Recoverability Report 2026

Security tells you what can be attacked. Recoverability tells you what survives.

FENIX24

FENIX24

Executive summary

Cybersecurity has spent two decades organized around a single idea: keep them out. That idea is being retired as the measure of success. Gartner now advises security leaders to define success around resilience rather than prevention, on the grounds that resilience, unlike absolute security, can be tested, practiced, measured, and improved. Regulators are now writing recovery obligations into law, insurers are pricing recovery posture into premiums, and boards have stopped asking if the organization is secure and started asking how long it would be down.

The annual probability that a given organization experiences a significant cyber event has nearly quadrupled since 2008, by the Cyentia Institute's measure, and roughly 3,000 significant incidents now become public every quarter. When an outcome becomes an operating condition, managing its consequences matters more than promising to prevent it.

What the shift has exposed is a measurement vacuum. Organizations can produce a number for nearly every dimension of resistance, such as endpoints covered, vulnerabilities patched, or alerts triaged. Almost none can produce a defensible number for the dimension that now defines the outcome, which is how many hours or days stand between a destroyed environment and an operating business. The plans exist. The evidence doesn't.

Two gaps, taken together, describe the state of recoverability in 2026. The first sits between the resilience organizations claim and the recovery they can demonstrate; it appears in every dataset in this analysis, from government surveys to insurance claims. The second is newer and opening fast. It sits between knowing what would prevent recovery and being able to do something about it. The industry is getting steadily better at finding these gaps, but finding a gap and closing it are different disciplines, and only the second one restores a business.

The industry has agreed that recovery is the measure of resilience. It has not yet built the means to measure it.

SECTION 01

The new measure of success

~4x

The annual probability of an organization experiencing a significant cyber event has nearly quadrupled since 2008.

Source: Cyentia Institute, IRIS 2025

Breach normalization is now official policy at every institution that prices risk. Markets, regulators, and customers increasingly treat incidents as an ordinary cost of operating rather than proof of outright negligence. Gartner has formalized this stance by urging security leaders to build strategy around limiting impact, maintaining operations, and recovering quickly. From a business outcome perspective, in Gartner's analysis, mitigation has become functionally equivalent to prevention. The executive suite has internalized the same view. In Accenture's mid-2026 survey of 1,000 CEOs and CISOs, 87% describe cyber disruption as a recurring operating reality, and 72% say preventing all of it is no longer realistic.

Fifteen years, 566% more incidents in the public record.

~450

significant security incidents entering the public record per quarter, 2008

~3,000

per quarter, 2024

Source: Cyentia Institute, IRIS 2025

Follow the spending against the outcomes and the reason for the reversal is plain. Security investment has climbed for a decade, incident frequency climbed with it, and the cost of incidents climbed fastest of all. Cyber events now consume eight times more of a victim's annual revenue than they did fifteen years ago, by the Cyentia Institute's analysis. A discipline that optimizes only the probability of impact, while leaving the severity of impact unmanaged, will show exactly this curve.

FENIX24

The growth is ransomware-shaped.

<1%
of all reported incidents were ransomware, 2015

52%
of monthly reported cyber events, on average, by 2023

The attack type driving the incident record's reacceleration is the one whose defining harm is downtime.

Source: Cyentia Institute, IRIS Ransomware Study

Prevention cannot be fully proven, because a quiet quarter is often indistinguishable from a lucky one. Recovery can. It has objectives that are met or missed, sequences that work or stall, and timelines that can be rehearsed against a clock. Resilience earned its place as the organizing strategy precisely because it can be measured, and that is the standard it sets for every organization now adopting the word. Resilience that has never been measured is a position, not a capability.

TAKEAWAY

The definition of winning has changed from "nothing happened" to "it happened and we were operating again in hours." Very few organizations have instrumented the second definition.

SECTION 02

Downtime is the P&L event

>50%

of the value of large cyber claims is business interruption. The largest cost of a cyber incident is the not-operating.

Source: Allianz Commercial, Cyber Security Resilience 2025

The most reliable witnesses to the true cost of cyber incidents are the institutions that pay for them, and their claims data has converged on a verdict. Across the world's largest independent claims datasets, compiled by Allianz Commercial, Munich Re, and NetDiligence, business interruption is the dominant loss component, exceeding ransom payments, forensics, legal fees, and notification costs combined.

Two of the world's largest claims datasets agree: business interruption is the biggest loss.

Source: Allianz Commercial, 2025 · Munich Re, 2025

Business interruption is also the one loss component that behaves differently from all the others. Forensics, legal, and notification costs are roughly fixed once an incident occurs. The BI line is a variable, and the variable it tracks is recovery speed. Allianz's claims analysts note that interruption losses correlate directly with the quality of response and continuity execution; indecision and poor coordination extend them. Every hour of recovery time purchased in advance is claim value that never accrues. No other security investment has that direct a relationship to the largest number on the loss sheet.

FENIX24

A claim with business interruption costs 650% more.

650%

Average cost premium of claims with a BI component vs. those without, five-year data. Ransomware drove 81% of SME claims that had one.

Source: NetDiligence Cyber Claims Study 2025, 2020–2024 data

Inside a claim, the interruption is the bill.

Averages for claims with a reported business interruption component. Note: the two panels use different scales.

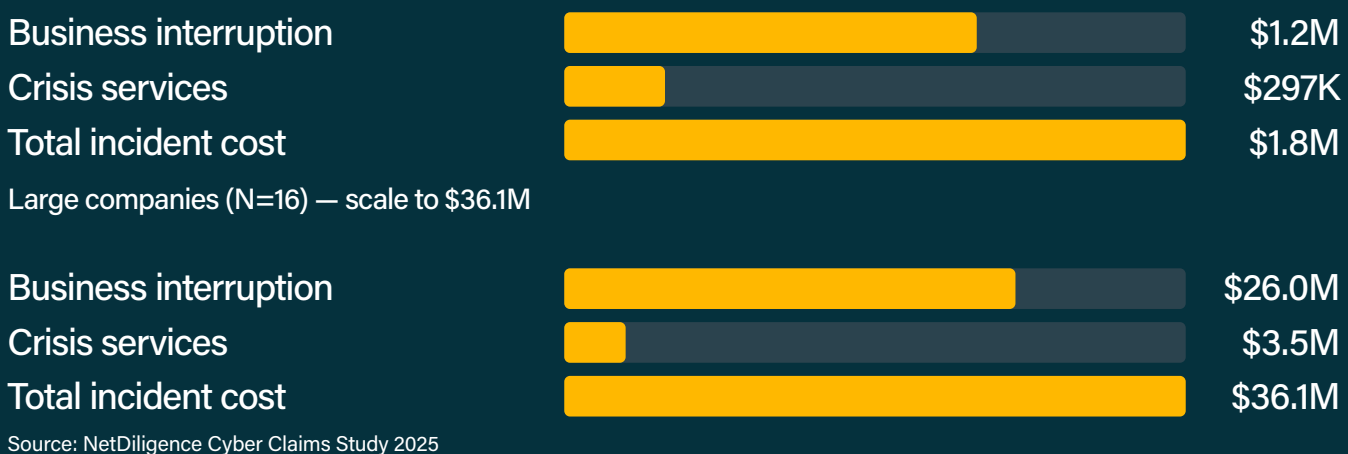

Source: NetDiligence Cyber Claims Study 2025

The official numbers understate the loss, and say so. The FBI's IC3 logged 3,600 ransomware complaints in 2025 with $32 million in reported losses, and its own methodology notes the total generally excludes lost business, time, wages, files, equipment, and third-party remediation. It also counts only the incidents victims report to the FBI. The Cyentia Institute, estimating across a fuller incident population with operational costs included, put annual ransomware losses near $95 billion for 2023. The two figures measure different things. But the categories the headline number excludes by design are the same downtime costs the claims data above shows dominate the bill, which means a budget case anchored to reported losses is missing its largest line item.

FENIX24

The most quoted ransomware number excludes most of the cost of ransomware.

$32M

ransomware losses reported to the FBI's IC3 in 2025, a total its methodology says generally excludes downtime, lost business, and remediation

~$95B

the Cyentia Institute's independent estimate of annual ransomware losses for 2023, operational costs included

Source: FBI IC3 2025 Annual Report · Cyentia Institute IRIS Ransomware Study

TAKEAWAY

Measure cyber exposure in hours of interrupted operations, not records exposed. The claims data already does.

SECTION 03

The preparedness illusion

85% / "largely untested"

of large UK businesses have a continuity plan covering cyber. The same government survey found plans "largely untested or incomplete."

Source: UK Cyber Security Breaches Survey 2025/26

On paper, the enterprise is ready. In the UK government's Cyber Security Breaches Survey, 85% of large businesses hold a continuity plan that covers cyber, three-quarters keep a formal incident response plan, and nearly nine in ten maintain formal cyber policies. The same survey then interviewed the organizations behind those numbers and found plans that lived in staff handbooks, had never been exercised within anyone's memory, and in some cases were tested only because a post-incident review forced the issue. The survey's own conclusion is that an untested plan creates a false sense of preparedness that is itself a risk.

In ISACA's global survey of practitioners, 43% expect an attack on their organization within the year while only 41% are confident in their team's ability to respond to one, a profession bracing for impact it doesn't believe it can absorb.

The gap persists because testing is structurally unrewarded. A full restore exercise is disruptive and cross-functional, with an owner in no single department. Restore order, backup strategy, and recovery target validation each sit in a seam between security, infrastructure, and the business. A capability that lives in a seam has no one whose job depends on exercising it. Accenture's 2026 C-suite survey puts numbers to the vacancy. Nearly three-quarters of organizations file cyber resilience under security or IT, fewer than half govern it as a shared business outcome, and when a thousand CEOs and CISOs were asked who primarily owns end-to-end business survivability, no single answer cleared 35%.

Documentation, by contrast, is cheap, and nearly every compliance framework has previously accepted it as evidence. ENISA's EU-wide data shows where that leads. Investment is flowing, driven heavily by NIS2, yet business continuity ranks as the second hardest requirement to actually implement, and nearly a third of European organizations completed no security assessment of any kind in the past year. Regulation is succeeding at producing artifacts, but artifacts are not capability. Regulators are reaching the same conclusion. The newest rules ask organizations to prove a recovery works, not to file a plan that says it would, and documentation alone is starting to fail the audit. What they are reaching for instead is **continuous recoverability validation**, proof produced on the cadence the environment changes rather than the cadence the audit calendar sets.

Restore paths, backup survivability, and recovery targets tested against the live environment on an ongoing cadence, not an annual one.

Compliance is trusted. Recovery is untested.

87%
of organizations believe regulatory compliance ensures cyber resilience

35%
regularly test resilience through real recovery and continuity exercises, not just audits

Source: Accenture, Redefining Cyber Resilience Pulse Survey, 2026 (n=1,000 C-suite)

The scale of the misestimate only shows up when expectation meets an incident. Eighty-one percent of executives in the same Accenture survey assume critical downtime in a serious incident would run ten days or less, but Accenture's analysis puts average recovery at three to six months. Fenix24's engagements read the same way from the field. When recovery time objectives are documented at all, they typically promise 24 to 48 hours. Across more than 800 clients, only four came close to that benchmark for partial business operations, and none were at full operational capacity until several weeks after the incident.

Executives budget days. Recoveries may take months.

≤10 days
assumed ceiling for critical downtime in a serious incident, per 81% of executives

3–6 mo
average actual recovery, in Accenture's analysis

Source: Accenture, Redefining Cyber Resilience Pulse Survey, 2026

The default cadence is annual, but environments don't drift annually.

How often a cyber risk assessment is performed. This measures risk assessment frequency, the broadest and most common check organizations run. Recovery validation is rarer still.

Source: ISACA State of Cybersecurity 2025 (global)

There is also a category error hiding inside most of those plans. Response planning and recovery planning are different disciplines. A response plan assigns roles, escalation paths, and communications. A recovery plan defines restore order, system dependencies, backup survivability, and who holds decision rights at 3 a.m. when the options are all bad. Plans that answer the first set of questions get counted, in every survey above, as if they answered the second.

FIELD DATA · WHAT ORGANIZATIONS BELIEVE GOING IN

> The beliefs organizations most often carry into a recovery, which are later proven to be false.
>
> *"My backups are good." "Everything is backed up." "We know our applications and the order they restore in." "We know what is on our network." "Everything must be patched and hardened before the business can operate again." "The ransomware will keep spreading even after containment."*

Source: Fenix24 field data & recoverability intelligence, 500+ ransomware recoveries

A recovery plan that has never been executed is a hypothesis. An organization that has never run one has never seen its own number.

TAKEAWAY

One audit question separates the two groups. When was the full restore path last run, end to end, against current recovery targets? "Never" and "we simulated it" are the same answer.

SECTION 04

Recovering against an adversary

65%

of initial access in real-world incidents is identity based. Attackers log in more often than they break in.

Source: Palo Alto Networks, Global Incident Response Report 2026

Disaster recovery, as a discipline, carries one assumption in its foundations: the system failed and the environment around it is honest. A flood does not steal credentials. A power outage does not hunt for backup repositories. The recovery infrastructure, in every traditional DR model, sits safely outside the event.

Ransomware inverts each of those assumptions, and the incident data shows how methodically. Two-thirds of intrusions begin with identity, per Palo Alto Networks' 2026 global incident data, and credential compromise has ranked as the most common intrusion technique for a decade across the Cyentia Institute's long-run analysis. That statistic has a recovery consequence most continuity plans never absorb. The accounts, directory services, and administrative paths an organization would recover with occupy the same identity plane the attacker just took. Recovery teams routinely face restoring an environment using credentials they can no longer trust, inside a directory that may itself be the compromise. Fenix24's field data makes the coupling concrete. Active Directory figures in effectively every recovery our team of breach experts runs, typically as the first major system compromised, and backup infrastructure is usually joined to the same directory. When infrastructure management answers to the enterprise identity provider, restoration begins with rebuilding identity, or redesigning authentication outright, before the first system can safely return.

FENIX24

FIELD DATA · THE IDENTITY PLANE

99.2% of clients arrive with no documented identity recovery plan. Of the plans that did exist, none survived contact with the threat actor.

95% have no meaningful multifactor controls on critical infrastructure consoles, while only 15% lack sufficient controls at network ingress. The front door is guarded. The management plane is not.

94% run backup systems joined to the production directory the attacker takes first.

A fifth of the first 48 hours of a typical engagement goes to the identity plane alone, standing up one authentication source healthy enough to prove positive control. Rebuilding infrastructure to minimum viability runs 72 hours or more. An assumption that can be retired is that a compromised directory rarely has to be rebuilt from nothing. Organizations brace to greenfield the entire domain when a targeted, evidence-led restoration of the existing one is usually faster and safer.

Source: Fenix24 field data & recoverability intelligence, 500+ ransomware recoveries

Identity is the first constraint. Time is the second. Average breakout time from initial access to lateral movement fell to 29 minutes in 2025, per CrowdStrike's global threat data, with the fastest observed at 27 seconds and data leaving one victim within four minutes of entry. Any recovery model that depends on convening a call, locating a runbook, and escalating through management operates on a timeline the adversary abandoned years ago.

The attacker's clock.

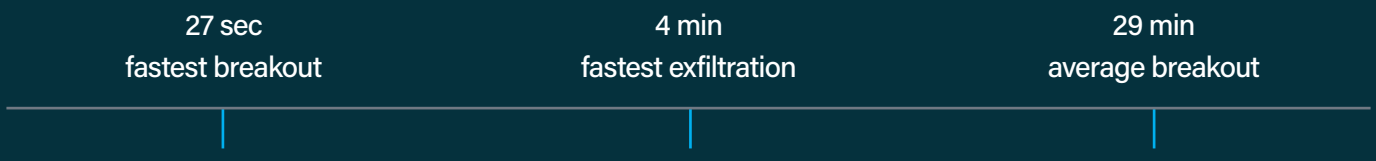

Source: CrowdStrike 2026 Global Threat Report

Backup infrastructure is targeted early and specifically, for a reason that is economic rather than technical. An organization that can restore has no reason to pay. Destroying the restore path is what converts an IT incident into a board-level extortion event, which is why backup existence and backup survivability have become different properties. A repository reachable with the same credentials, on the same network, under the same directory as production is, from the attacker's chair, just another target with a more valuable payload. And survival is not the last hurdle. In 38% of Fenix24 engagements where backups came through the attack intact or nearly so, they still could not carry the recovery. The data was too old to resume operations on, incomplete or corrupted long before the attack, the wrong type for the workload, or slower to restore than rebuilding outright. Some arrive labeled immutable while running on infrastructure incapable of immutability. Existence, survivability, and usability are three different properties, and only the third one recovers a business.

The attack itself keeps moving. Exfiltration appeared in 40% of large claim value in early 2025, in Allianz's claims analysis, nearly double the prior year, while encryption rates fell to a six-year low. A growing share of incidents never locks a screen. They steal data, redirect funds, or poison trust in identity systems. Recovering from "we cannot trust our directory" is a different project than recovering from "our files are encrypted," and an operational resilience program has to withstand both.

Extortion is shifting from encryption to exfiltration.

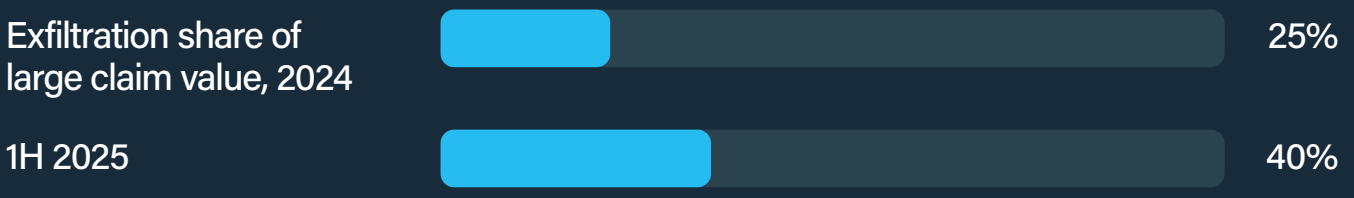

Meanwhile, encryption rates fell to their lowest level in six years.

Source: Allianz Commercial, Cyber Security Resilience 2025

TAKEAWAY

Stress test recovery assumptions against an adversary, not an outage. If the backups can't survive a privileged attacker and the restore path depends on the compromised identity plane, the stated RTO measures nothing.

SECTION 05

The unit of recovery is the business service

6% → 15%

Supply-chain business interruption more than doubled as a share of large cyber claim value in a single year.

Source: Allianz Commercial, 2024 vs. 1H 2025

Servers restore. Businesses recover. Between those two statements sits the map most organizations don't have.

"Recoverability" and "resilience" get used as if the words are interchangeable. They are not. A server can be recoverable. So can a database, a backup job, or an image. None of that means the business is, because a business does not run on components. It runs on services, and a service is recovered only when everything it depends on is back, in order, and trusted. That is **operational recoverability**, and it is the level at which resilience is decided. An organization can restore a hundred machines and still be unable to invoice a customer, pay its people, or ship an order.

The ability to restore the business function, not just the systems beneath it. Every dependency back, in order, trusted, proven under real conditions.

Recoverability without the operation attached is a server count. Resilience without recoverability is a policy. The capability is where the two meet, proven under real conditions, not assumed.

A business service is a chain of dependencies: applications, databases, identity, network paths, backup infrastructure, cloud platforms, third-party integrations, and the people who operate them. The service returns when the chain returns, in order, and not a moment sooner. An ERP restored to a clean hypervisor produces nothing while the logistics integration it feeds remains down; a payroll system is not recovered until the identity provider that authenticates it is trusted again. Recovery time is set by the slowest necessary link, including the links nobody documented. Sometimes the slowest link is physical. Storage runs short in 82% of Fenix24's engagements, leaving restored data nowhere to land without overwriting the forensic record, and in 38% the network cannot move data at recovery scale, a backbone adequate for daily operations facing petabytes of restoration. And sometimes the link cannot be rebuilt at all, because it is a legacy operating system or application version that can no longer be procured once the original is gone.

FENIX24

FIELD DATA · THE MISSING MAP

> The number of organizations that arrived at a recovery holding a complete picture of their applications and dependencies was zero. The closest approximations either sat in a configuration database that went down with the attack, or were assembled mid-recovery once the breach forced the business to decide what it needed back first, and were then revised as restoration progressed.

Source: Fenix24 field data, 800+ client engagements

The chains increasingly run outside the walls. Third parties now figure in roughly 30% of breaches, double their share from 2021, by the Identity Theft Resource Center's count, and Allianz's claims data shows contingent business interruption from supplier and vendor events jumping from 6% to 15% of large claim value in a year. The executives furthest along have already reoriented. Among CEOs of highly resilient organizations in the World Economic Forum's 2026 outlook, 78% name supply chain and third-party dependencies as their top remaining challenge. Hardening your own environment relocates the risk; it doesn't eliminate it. Accenture's 2026 survey measures the blind spot directly. Sixty-one percent of leaders say sustaining critical operations depends on external partners, and only 38% report a clear view of their value chain and its most critical dependencies.

The risk is moving into the supply chain.

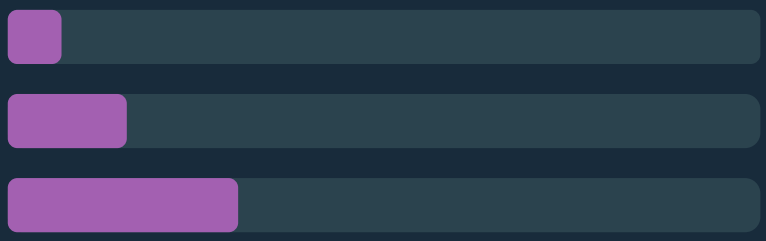

Source: Allianz Commercial 2025 · Identity Theft Resource Center 2025 Annual Data Breach Report

Cyber events sit at the top of DRI International's global ranking of operational threats, with general IT disruptions close behind. A defective software update or a cloud region failure also stops revenue, and also demands an answer to what comes back first, so the dependency knowledge described above pays out in those scenarios too. But the two conditions are related, not interchangeable. An outage takes systems down. An attack takes systems down and corrupts the credentials, backups, and directory the restoration would run on. A recovery capability built for the adversarial case absorbs the ordinary outage along the way. One built only for outages meets an adversary unprepared.

TAKEAWAY

Take the single most revenue-critical business service and demand its complete dependency map, third parties included. If the map doesn't exist, no recovery timeline attached to that service means anything.

SECTION 06

Insurance is already pricing recoverability

+40% / −5%

U.S. cyber claims rose nearly 40% in 2024 while premium rates fell 5%. Underwriters aren't getting generous. They're getting selective.

Source: NAIC 2025 Cybersecurity Insurance Market Report

Two curves crossed in the U.S. cyber insurance market in 2024, and their divergence says more about the future of resilience than either number alone. Claims volume climbed almost 40%, to nearly 50,000, by the NAIC's count. Rates fell an average of 5% in the fourth quarter, the first decline in seven years. Rising losses and falling prices coexist only when underwriters believe they can tell good risks from bad ones, and the NAIC credits the softening to sustained investment in controls that underwriters view favorably. Organizations that can evidence their security and recovery posture are buying better terms, higher limits, and lower retentions. Proof of recovery has become a financial instrument.

Claims up 40%. Rates down 5%. Underwriters are sorting, not softening.

+40%
cyber claims, 2024
(~50,000 reported, U.S.)

−5%
average premium rates,
Q4 2024, the first quarterly
decline in seven years

Source: NAIC 2025 Cybersecurity Insurance Market Report

What the policy cannot do is the part that gets discovered mid-incident. Extortion demands reached $150 million in NetDiligence's recent claims record, and individual payments reached $75 million, with fifty separate ransoms of eight figures or more. Coverage at that scale transfers capital, but not capability. Buyers lean on it anyway. Sixty percent of organizations treat cyber insurance as a substitute for resilience rather than a financial backstop, per Accenture's 2026 survey, and total disruption costs can run ten times the payout once downtime, lost revenue, and reputational damage are counted.

FENIX24

A POLICY RESTORES	A POLICY CANNOT RESTORE
Capital, after the claim is adjusted.	Domain controllers Trust in a compromised identity plane A validated clean backup The sequence in which business services return A single operating hour

The claim check funds the recovery. It cannot accelerate one. Recovery speed was fixed months before the incident, by the survivability of the backups, the currency of the dependency map, and the existence of a rehearsed restore path. The insurers' own data closes the argument. Business interruption is the majority of what they pay, per both Allianz and Munich Re. Allianz also reports that interruption losses scale with response and continuity execution. The industry that prices this risk for a living has concluded that recoverability is the variable. Buyers should read their renewal questionnaires as the market signal they are.

TAKEAWAY

Use the renewal as a catalyst. The evidence underwriters now ask for and the evidence an actual recovery requires are the same evidence.

SECTION 07

From board attention to board evidence

64% / 19% / 87%

64% of organizations report meeting minimum resilience requirements, and 19% report exceeding them. In a separate global survey, 87% of C-level executives call their protection inadequate.

Sources: World Economic Forum 2026 · Munich Re 2025

Board engagement, the perennial ask of the security profession, has arrived. In highly resilient organizations, 99% report board involvement in cybersecurity, per the World Economic Forum's 2026 global outlook, half with regular briefings and nearly half with formally defined oversight roles. Regulation arrived alongside it, and it increasingly names recovery specifically. In the EU, NIS2 has become a primary driver of security investment, per ENISA. In the U.S., SEC rules require public companies to disclose a material cyber incident within four business days of determining it is material, and to describe the board's oversight of cyber risk in the annual report, which makes recovery time information investors read. New York's amended DFS cybersecurity regulation goes further for the financial firms it covers, requiring business continuity and disaster recovery plans built for cyber disruption, tested annually, with annual tests of the ability to restore critical systems from backups.

What all that attention has produced is harder to read, because the only measures in circulation are self-reported. Nearly two-thirds of organizations say they meet their minimum resilience requirements, and fewer than one in five say they exceed them, in the World Economic Forum's 2026 survey. Yet in Munich Re's separate global survey of C-level executives, 87% judge their own organization's protection inadequate, and in ISACA's global survey of security professionals, just 56% believe their board gives the issue appropriate priority. The divergence runs vertically as well. In Accenture's matched 2026 survey of a thousand CEOs and CISOs, chief executives are more than twice as likely as their security leaders to believe regulatory compliance alone ensures resilience, and 96% of CISOs say their mandate now spans enterprise resilience and crisis leadership while 72% of CEOs agree. These are different surveys asking different populations different questions, and that is the finding. The industry's picture of its own resilience is assembled from opinion, and the picture changes with the question. Evidence would not. The 2025 ransomware attacks on major UK retailers, cited in the WEF's report for the operational disruption they caused at household names, are what the distance between reported confidence and tested capability looks like when it surfaces.

FENIX24

The confidence spread.

64%
report meeting minimum resilience requirements

WEF GLOBAL CYBERSECURITY OUTLOOK 2026

19%
report exceeding them

SAME SURVEY

87%
of C-level executives call their protection inadequate

MUNICH RE GLOBAL SURVEY, 2025

Three separate surveys of different populations, presented as a spread, not a computed gap.

The spread exists because most of what boards receive is claim rather than evidence. Framework alignment, maturity tiers, and policy attestations describe intentions. Gartner's prescription points at the alternative, urging impact thresholds tied to mission-critical value chains, defining how much disruption is tolerable and where recovery time decides the outcome. Those thresholds are only writable for business services that have been mapped, and only credible for recovery targets that have been tested. The endpoint of the convergence is a **recoverability score**[1], sitting in board materials where the attestation used to. Tracked across the factors that actually govern restoration time and watched over quarters, that score becomes a **recoverability index**[2], a trend line a board can read the way it reads any other measure of operating health. Governance is converging, from a different direction, on the same requirements the claims data and the incident data reached in earlier sections.

The AI wave sharpens the demand without changing it. In the same WEF survey, 87% of respondents rank AI-related vulnerabilities the fastest-growing cyber risk, and CEOs of the most resilient organizations rank it first among their concerns. New dependencies and new uncertainty, same governing questions. What has to keep working, what can fail safely, and how fast does the business return when a dependency gives way?

[1] A defensible, evidence-backed number for how fast the business comes back, tracked over time like any operating metric.

[2] A composite measure of recovery readiness across the factors that govern restoration time, tracked continuously so the direction of travel is visible before an incident, not after.

TAKEAWAY

The board request coming next year is not a security update. It is proof of recovery, with evidence and a trend line, and it cannot be assembled retroactively.

SECTION 08

The shift toward recoverability

> Security tells you what can be attacked. Recoverability tells you what survives.

Every gap this analysis has measured comes back to the same missing capability. Organizations cannot produce a continuously validated, evidence-backed answer to what happens after impact. Which business services return. In what order. In how many hours. With what confidence.

Twenty years of security investment built precise instrumentation for resistance. The industry can count endpoints, score vulnerabilities, and rank exposures in real time, because it decided those numbers mattered and built the systems to produce them. It never built the equivalent for recovery. Recovery was filed under IT operations, governed by documents, and measured by attestation, which is how an industry that can describe its attack surface in granular detail ended up unable to say, with evidence, how many hours stand between a destroyed environment and an operating business.

That asymmetry is now correcting, and every force in this report is pushing the same direction. Insurers price recovery posture because business interruption is the majority of what they pay. Regulators are writing tested recovery into law and beginning to reject the paper plan. Boards have moved from asking whether the organization is secure to asking how long it would be down. The analyst consensus has redefined the objective around surviving impact rather than avoiding it. Independent of one another, the institutions that price risk, enforce it, govern it, and study it have arrived at a conclusion recovery specialists reached years ago from inside the incidents: resistance is necessary, but recovery is what keeps the business standing afterward. The next decade of security investment follows that recognition, into instrumentation for recovery as rigorous as the instrumentation built for defense.

Building it demands two things the current model separates. The first is measurement, a continuous, evidence-backed reading of what the business could actually restore, and how fast, replacing the annual attestation that satisfies an auditor and predicts nothing. The second is execution, because measurement alone changes no outcome. An assessment does not harden a backup repository. A dependency map does not isolate the identity tier a restore depends on. A finding, however precise, has never brought a business service back online. The organizations that get this right will treat recoverability the way the industry learned to treat detection: instrumented, operated, and exercised continuously, owned as a business outcome rather than delegated as an IT task.

What that measurement has to account for is not a mystery. The failures that decide recovery timelines are consistent and they repeat across industries and incident types. They are conditions that were present before the attacker arrived and would have been visible to anyone measuring for them.

FENIX24

FIELD DATA · WHAT ACTUALLY SLOWS A RECOVERY

The recurring blockers, in the order they most often decide the timeline.

1. No documented recovery plan or system prioritization
2. Insufficient storage capacity to recover into
3. Undocumented application dependencies
4. Backups that report success but aren't operationally recoverable
5. Compute and storage too slow for at-scale recovery
6. Network bandwidth too thin for recovery-scale data transfer
7. Waiting too long to engage restoration experts
8. Organizational friction and decision paralysis
9. Poor backup implementation fundamentals
10. Flawed recovery philosophy: red/green, clean-room, full-remediation-first

Two of these are purely physical and routinely overlooked: storage runs short in 82% of engagements and bandwidth in 38%, because a network sized for daily operations was never sized to move petabytes back at once. The rest are decisions. The most expensive is treating full forensic remediation as a precondition for resuming any operation at all, when resuming limited operation under containment is often the faster, safer path. Accountability for the business risk stays with the organization; it cannot be handed to a vendor along with the backups.

Source: Fenix24 field data & recoverability intelligence, 500+ ransomware recoveries

Every item on that list is knowable in advance. Storage capacity, backup usability, dependency maps, identity recovery paths, and decision rights can be measured, scored, and corrected before an incident rather than discovered in the middle of one. That is the discipline the industry is now standing up, and it deserves a name it does not yet have in common use. Fenix24 calls it **recoverability intelligence**, the continuous, evidence-based measurement of what a business can restore and how fast, built to the same standard the industry already holds for measuring how it can be attacked. Detection told organizations what was happening to them. Recoverability intelligence tells them what they would get back, before the day they have to find out.

> The discipline of measuring, continuously and from live conditions, what a business could actually restore and how fast: its dependencies, backup survivability, recovery sequence, and achievable recovery times.

That day is no longer hypothetical for anyone. The question every organization now faces is no longer how to avoid the attack. It is what happens when the environment goes dark. Can operations come back before the business is gone? Prevention was always the easier problem to instrument, and recovery was always the one that decided whether the company lived. The reckoning underway across insurers, regulators, boards, and analysts is the market correcting a measurement error decades in the making.

The correction rewards the organizations that move first and punishes the ones that wait, because recoverability cannot be bought in the moment it is needed. It is built in the quiet beforehand or it is not there at all, and the environment keeps changing no matter who is watching. Every quarter a company treats recovery as a document instead of a tested capability is a quarter its real posture drifts further from what its plan claims, until an attacker collects the difference. The organizations that internalize this will absorb the worst day their business will ever face and keep operating. The ones that do not will meet it the way this report's incident data describes, learning what they could actually recover at the only moment it can no longer be changed. Resistance decides how often that day comes. Recovery decides whether there is still a business standing on the other side of it. The industry is finally measuring the right one.

You've spent twenty years measuring what can be attacked. Start measuring what will survive.

FENIX24

Find out what has to come back first for your business to survive. Then find out if it can.

Fenix24's Resiliency Intelligence Assessment defines the systems your business cannot operate without and tests your recovery posture against real RPO and RTO targets, using live recoverability intelligence surfaced by Argos99 from your environment. You leave with evidence of where you stand, the gaps that would decide your recovery timeline, and a prioritized roadmap for closing them. That is where provable resilience starts. Every gap closed from it is measurable progress your board, your regulator, and your insurer can see.

LEARN MORE ABOUT THE ASSESSMENT

Recoverability Intelligence POWERED BY ARGOS99

ABOUT THIS ANALYSIS. The State of Recoverability synthesizes research published 2024–2026 by Allianz Commercial, Munich Re, the NAIC, NetDiligence, the Cyentia Institute, Accenture, the FBI's Internet Crime Complaint Center, the UK Government's Cyber Security Breaches Survey, ENISA, the World Economic Forum, Gartner, ISACA, Palo Alto Networks, CrowdStrike, the Identity Theft Resource Center, DRI International, and the Business Continuity Institute. All statistics are attributed to their original sources, and regional scope is noted where research covers a specific market. Analysis is Fenix24's, informed by more than 500 ransomware recovery engagements. No Fenix24 client data appears in this report.