White Paper

halcyon

22 Days of Hell

Ransomware recovery is a three-week fight.
Most organizations are set up to lose the battle.

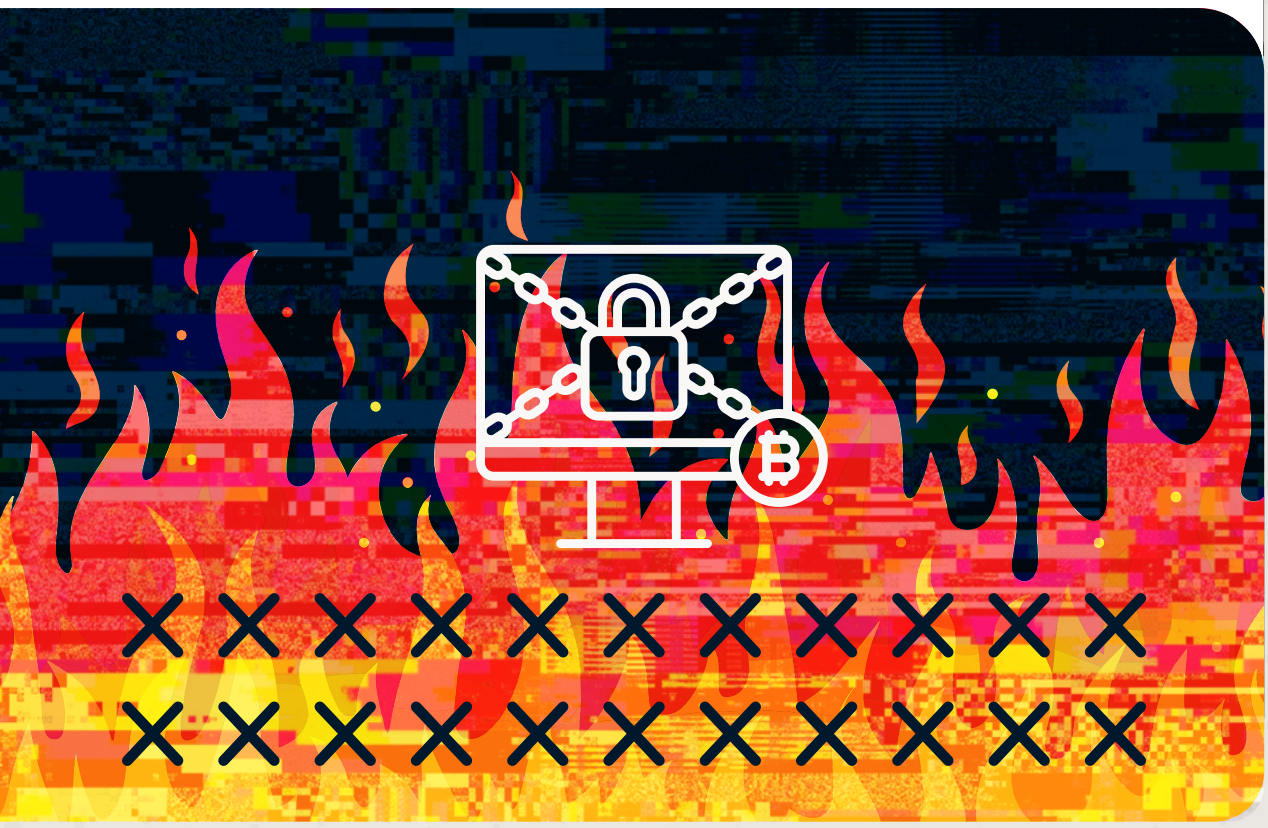

Executive Summary

Most organizations believe that recovering from a ransomware attack takes a matter of hours. The reality is that the average organization takes around 22 days to recover, and large enterprises often take longer.

The gap in perception matters because it provides organizations with a false sense of security, that with secured backups and a tested recovery plan that they'd be able to easily survive a ransomware attack. However, when the real event unfolds over weeks, the "trusted" recovery plan does not survive contact with it, and every day of unplanned downtime carries a cost that can threaten the survival of the business.

Key Takeaways

- **The perception gap is real, and preparation does not close it.** Organizations with tested backups, documented playbooks, and cyber insurance still land inside the 22-day average.
- **Recovery is not one step but a long, sequential process.** Restoring a business means fully evicting the attacker, verifying that backups are clean and free of the compromise, and rebuilding systems in an order dictated by technical dependency.
- **The delay is driven by a collision in priorities, not from technical failures or poor planning.** Throughout recovery, the methodical work that safety requires runs directly against the business's pressure to restore operations immediately. That pressure forces shortcuts which extend the recovery it was meant to shorten and resurfaces later as reinfection, rework, or legal and regulatory liability.
- **Better prevention alone is not the answer, because breach should be assumed.** Determined attackers will get in. Treating recovery as the safety net is a losing position, because the recovery itself is the costly, weeks-long event that damages the business more so than the attack itself.
- **The real answer is cyber resilience.** It means having a plan and ability to detect an attack while it is underway, contain and evict the threat actor before the damage spreads, and prevent or stop encryption as it's happening to ensure that full organizational recovery is not required. Resilience of this kind keeps the business running through an attack, rather than depending on the recovery cycle to bring it back afterward.

Ransomware grew again to 48% of all breaches, up from 44% the previous year.

Verizon 2026 DBIR

Introduction

Ask a CISO whether a hired red team could break into their network and nearly all of them will say yes. Industry data backs that up: 94% of organizations that run red team exercises experience some level of successful penetration, and SANS research found that nearly 60% of ethical hackers can breach a corporate environment in under five hours. A real attacker has the same tools and skills as the red team, but with more time, more motivation, and no rules of engagement. If nearly every security leader concedes that a determined attacker will eventually get in, then the question that matters is not whether the defenses will hold. It is what happens when they fail.

Most organizations believe they are ready for that moment. They have an incident response plan, documented recovery procedures, and the testing they believe proves both will work. Every year, teams run ransomware tabletop exercises where a facilitator walks through a scenario, the security team talks through containment steps, IT describes the restore process, and someone confirms the backups are tested and immutable.

Then the real attack happens, and the numbers tell a different story. According to Gartner, the average organization takes 22 days to recover from a ransomware attack while Total Assure estimates recovery time to be 38 days for enterprise environments. These are not outlier figures from poorly prepared organizations. They are averages across thousands of incidents, including organizations that had documented incident response plans, tested backups, and cyber insurance in place when the attack hit.

The gap between the tabletop and the real incident is not a knowledge gap or a process planning gap. It is structural. Simulations run while the business is still up, so the security team works through the plan with full focus and no competing demands. A real incident takes the business down, and the pressure to get it back up competes with the team's ability to actually address the attack.

This paper details what actually happens during the worst 22 days of an organization's life, and why the time it takes to recover is nearly impossible to improve. The timeline is not driven by weak tools or bad planning but by the structure of recovery itself: work that cannot be safely rushed, colliding with business pressure that demands it be rushed anyway. Backups, however well tested, are not the safety net they appear to be, and treating recovery as the fallback plan is itself the mistake. Attackers will get in, but an intrusion only becomes a 22-day catastrophe when it is allowed to spread unchecked. As this paper will examine, the goal should not be to find ways to improve the recovery process and time it takes. The stronger position is shifting focus from recovering after the damage to mitigating the attack itself: detecting it early, containing its spread, and evicting the attacker before business operations are ever disrupted.

22

Days of Downtime

The average organization takes 22 days to recovery from a ransomware attack.

Gartner

halcyon

Prepared Organizations Still Take 22 Days

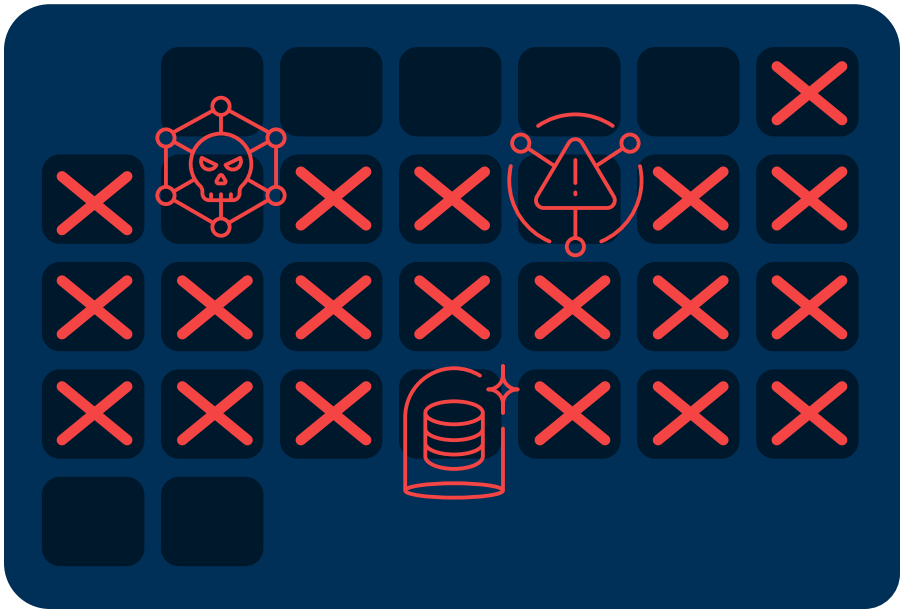

The natural response to the numbers in the introduction is to assume they describe someone else but the data proves otherwise.

The first problem is that documented preparation and functional preparation are not the same thing. Veeam's 2025 Ransomware Trends survey of 1,300 organizations found that 98% of ransomware victims had a playbook, but less than half had the key technical elements that playbook depends on, including backup verification and frequency standards (44%) or a pre-defined chain of command (30%). The consequences of those gaps show up directly in the recovery timeline. Organizations that have not built a process to validate backups before restoring them either do not know how to do it when the incident hits, take far longer to do it than the plan assumed, or skip it entirely under pressure and restore the attacker right back into the environment. Organizations without a defined chain of command lose control of the response to whoever has the most organizational authority in the room, which during an active incident is almost always a business leader pushing for restoration speed rather than a security leader holding the line on process.

The second problem runs deeper. Even organizations that are genuinely prepared still get caught inside the 22-day average. Preparation determines how well the response starts, but it does not shorten the work the response actually requires. The investigation takes as long as the attacker's spread demands, and business pressure builds regardless of how good the plan was. This is why confidence and outcomes come apart so sharply: in the Halcyon 2026 Security Leadership Survey, 90% of security leaders rated their security solutions as sufficient or mostly sufficient and 99% were confident in their ability to detect attacks, yet nearly half still experienced moderate to significant disruption from a successful attack.

There is one preparation decision worth singling out. Modern ransomware, now supercharged by AI, can find and exploit a way in faster than most teams can patch, which means keeping attackers out entirely is no longer a realistic plan. But relying on backup and recovery to rebuild the entire environment afterward is its own trap, one this paper spends the following pages exposing. Instead, organizations need a plan and the capability to detect an attack in progress, contain it, and evict the attacker before it spreads far enough to force a full recovery. The goal is not to prevent every intrusion or to recover faster after a catastrophic one, but to stop an incident from ever growing into the weeks-long ordeal this paper describes.

The 4 Phases of Recovery

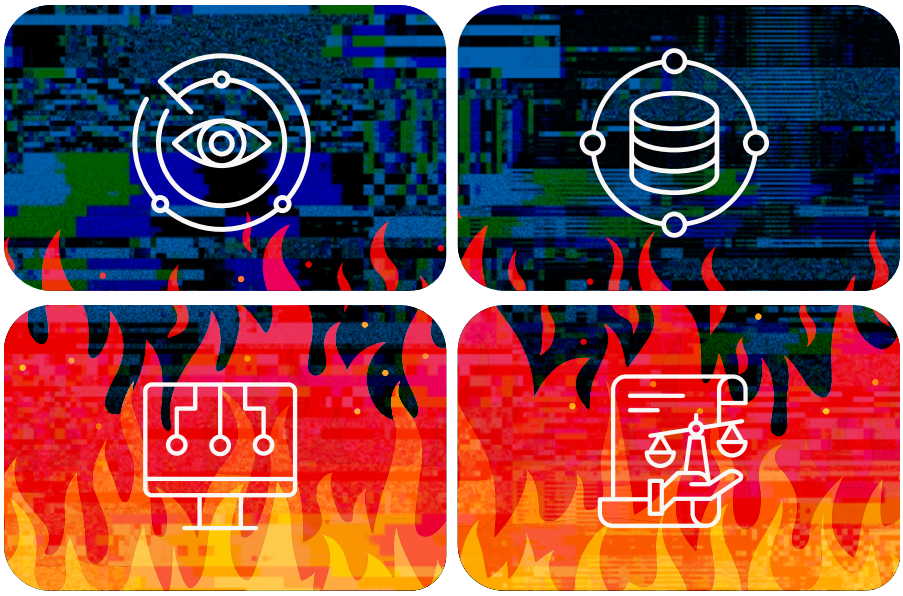

Real-world ransomware recovery moves through four distinct phases with a foundational reassessment of the environment running throughout which perpetually assesses the situation and helps inform future decisions and actions.

Situational Reassessment: The continuous evaluation of what is happening in the environment, where it is happening, and how severe it is at any given moment. It begins at first detection and continues until the incident is fully resolved.

- **Phase 1: Addressing the Attack and Legal Requirements.** Three workstreams running in parallel: investigation (1A), containment and eviction (1B), and forensic, legal, and notification obligations (1C). This is the most variable phase in the entire timeline, resolving in hours or running for weeks depending on how far the attacker got.
- **Phase 2: Backup Verification.** Confirming restore points predate the intrusion and are free of compromise before anything is restored. This is where the majority of recovery failures occur.
- **Phase 3: Tiered System Restoration.** Systems return in dependency order, not business priority order, with every credential rotated and every system validated before it returns to production.
- **Phase 4: Hardening, Regulatory Work, and Normalization.** Closing the attack vector and working through the regulatory and legal obligations that extend months or years beyond the technical recovery.

Phase 1

Assess the Environment, Investigate the Incident, Contain and Evict the Attacker, Disclose the Event.

The first phase of recovery begins the moment an incident is discovered and covers the full security workload that follows: investigating the incident, containing the attack, evicting the attacker, and managing the legal obligations that come with it, including forensic preservation, breach disclosure, and cyber insurance notification. Not every incident requires all of that. If the team catches the attacker early enough that no systems are taken down and no data is touched, the response may close at containment and eviction without triggering the full weight of legal, regulatory, and disclosure obligations. Once the incident reaches critical systems or shuts the business down, the length of this phase, and with it much of the total recovery time, comes down to the scope of the attack and how far the attacker managed to spread before being detected.

To start the phase, the security team first needs to understand how the attacker got in, where they have been, what they have touched, and when the intrusion began. An attacker present for two weeks before encryption triggered means the last two weeks of backups may all contain the attacker's persistence mechanisms.

During the above investigation, the security team is simultaneously working to contain and evict the attacker. Compromised systems need to be isolated to stop further lateral movement, attacker access has to be cut off, persistence mechanisms have to be identified and removed, and compromised credentials have to be rotated. Eviction is not complete until the team can confirm with reasonable certainty that the attacker has no remaining access to the environment, no remaining footholds, and no way back in.

In addition, the team must also consider the legal and notification workstream. Insurance carriers require notification within 24 to 72 hours once a severe incident has been identified. Outside legal counsel needs to be engaged immediately, because attorney-client privilege covers decisions made with counsel present and matters for every piece of documentation that follows. Regulatory frameworks across SEC, GDPR, HIPAA, NYDFS, CIRCIA, and applicable state laws all impose notification requirements that start a fixed clock on deadlines set by the regulator rather than the organization. Initial filings have to go out based on what is known now, with updates following as the investigation progresses.

These three sub-phases do not run in a linear sequence, they run in parallel and they feed each other constantly. For example, the investigation surfaces a compromised system, and containment isolates it. The containment work reveals a persistence mechanism the investigation had not yet found, which prompts the investigation to look elsewhere for the same pattern. The legal workstream needs information the investigation has not yet produced, and the answers it does get reshape what has to be disclosed and when. Every action in one sub-phase changes the conditions the other two are operating under.

Underneath the three sub-phases sits a continuous reassessment of the environment. Think of it as the background phase, or phase 0. From the moment something anomalous is detected, the security team is asking the same set of questions on repeat: what is happening, where, how severe is it, and how does the current picture differ from the one they had an hour ago. Without continuous reassessment, the response runs on stale information, and stale information during an active incident produces the kinds of shortcuts that lead to reinfection and restarts the entire recovery process.

What this phase actually looks like in practice is not a series of completed steps. It is a loop: investigate, reassess, contain, reassess, take the next action, reassess again and so forth. The loop runs until the security team can confirm that the attacker is fully out, the scope is fully understood, the persistence mechanisms are all closed, and the legal and notification work is in motion. That confirmation is the gate that opens the next phase of recovery.

2026

recovered less than 75% of their affected data

Veeam 2026 Data Trust and Resilience Report

halcyon

Where the Business Pressure Compounds the Problem in the First Phase

When the business goes down, the CISO's responsibility expands well beyond the security function. Leadership begins asking questions the security team has to answer, not because those questions are security questions, but because the security team is the only group with visibility into what is actually compromised and what is actually safe. Four of these decisions show up in nearly every major incident, and each one pulls the team away from the active security work in ways that extend Phase 1.

Whether to leave "unaffected" systems connected and running.

The attack is identified and the security team moves to contain it. The correct instinct is to start isolating systems and cutting the connections the attacker could use to spread. But across the rest of the business, most systems are still up and running, and from where those teams sit, nothing appears wrong with the systems they depend on. When security moves to take those systems offline, the business pushes back, because what looks like a precaution to the security team looks like self-inflicted downtime to everyone else.

The security team's position is the one the business does not want to hear. A partially compromised environment is not a safe environment, and early in an incident "unaffected" is an assumption, not a confirmed status. The only safe move is to isolate broadly and confirm what is clean afterward. The joint CISA, FBI, NSA, and MS-ISAC #StopRansomware Guide instructs that if several systems or subnets appear impacted, organizations should take the network offline at the switch level rather than isolate only the systems confirmed to be compromised. Mandiant's containment guidance reaches the same conclusion, recommending broad network and endpoint isolation to halt propagation and noting that until the investigation is complete it is difficult to determine which systems and credentials the attacker is using.

However, in practice and faced with a certain, quantifiable operational loss on one side and a risk it can describe but not yet prove on the other, leadership frequently decides to keep the business running, and the systems that should come offline stay online. Each one remains a path the attacker can use to move laterally and rebuild footholds, so containment cannot close around a perimeter the business has insisted on keeping open, and every system the team is not allowed to isolate is one it has to keep re-checking as the investigation widens. Keeping operations running buys real time in the moment, but it borrows against a longer and less certain recovery, because the incident cannot end until containment and eviction are complete.

2026

of organizations hit by ransomware in the past year said that the cybercriminals attempted to compromise their backups during the attack.

Sophos

halcyon

Which systems can be brought back online during the incident.

The security team found a compromised system, isolated it, and worked to contain and evict the attacker. As the CISO and security team move on to the next compromised system, business leaders see a clean system sitting idle and want it turned back on. Proper security processes say to leave it alone until the entire incident is resolved and the attacker is confirmed removed from the whole environment. Business operations say the system is critical to keep the company running and that every minute it stays down is lost revenue. Who wins?

The answer **should be** security, because eviction from one host is not the same as a system that is safe to operate while the incident is live. Reconnecting it creates exposure through three distinct channels:

- **Network:** Plugging the system back into a network where the attacker still holds footholds elsewhere re-exposes it to the same lateral movement that compromised it in the first place.
- **Identity:** Even if the clean system was reconnected to an isolated part of the network, it still authenticates against the same Active Directory, service accounts, and Kerberos infrastructure, which may still be compromised, so the attacker can reach it again through valid or forged credentials no matter how clean the machine is. Mandiant treats this as a core containment problem, calling for identity-store integrations to be severed during the response and attacker-used credentials treated as unknown until the investigation is complete.
- **Containment:** To box the attacker in, the team severs trust relationships, replication, and site-to-site connectivity; reconnecting one system can re-establish those links and turn it into a bridge back into the isolated environment, re-exposing the very systems the team worked to protect.

Reconnection can also destroy forensic state. Powering a system through reboots and resuming activity overwrites the volatile memory, live process data, and network artifacts that investigators depend on, **which is why CISA warns that powering systems down forfeits evidence held in volatile memory, and why responders treat rebooting before capturing memory as a critical mistake. That lost evidence underpins the insurance claim, the regulatory filing, and any legal process, and even a genuinely safe reconnection still costs a per-system forensic check that pulls analysts off active eviction.**

2026

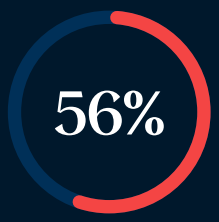

of attacks succeeded in encrypting data, up 6 points from 2025.

Sophos State of Ransomware 2026

halcyon

Building workarounds for inaccessible business systems.

The systems are down and the business has stopped running. Leadership and department heads are scrambling to get any business function back online, losing revenue and customer trust by the second, and all eyes turn to the security team. Turning systems back on or leaving untouched systems online are not safe options, as addressed above, so what is the next best one? Can the team build something to replace them – a temporary environment, a cloud-based stand-in, a manual process – so the business can keep functioning while the response continues?

While building and connecting new systems are not typically security functions, most organizations do not have enough internal development and infrastructure staff to design, stand up, and securely connect replacement systems on no notice in the middle of a crisis. The people who can are already consumed, and the security team is often the only group that both understands the environment well enough to build a safe replacement and can judge what is safe to connect while the attacker is still active, so the work lands on them. Even in the organizations that do have the staff to build, security cannot simply hand the task off, because any machine spun up during an active incident becomes part of containment. NIST's incident response guidance treats temporary workarounds that keep critical services available as a containment measure, and containment-phase systems must be hardened against giving the attacker a new path in before they go live. That makes oversight and validation of every workaround a security responsibility regardless of who does the building.

When Norsk Hydro was hit by LockerGoga in March 2019, the attack rendered most of the company's systems unusable and Norsk Hydro took the drastic step to isolate their IT infrastructure and revert to manual operations with pen and paper and manual production procedures. A MITRE case study documents employees doing work by hand for weeks before IT systems were able to be brought back online. They are held up as a model response because they demonstrated a successful response by doing the things the attack demanded: standing up and sustaining alternative operations under fire while the security work continued in parallel.

Most organizations do not have the preparation or the people that made manual operation work for Norsk Hydro, and the 2021 Conti ransomware attack on Ireland's Health Service Executive shows what that looks like. The independent post-incident review the HSE commissioned from PwC found the organization had no documented incident response plan and was heavily reliant on third parties, including the military, to give its response any structure and their recovery and rebuild still took four months.

2025

The use of backups to restore encrypted data is at the lowest rate in six years, used just in 54% of incidents.

Sophos State of Ransomware 2025

What makes this demand harder than the others is that it is not unreasonable, but the work still falls on the same small team already running containment, investigation, and eviction. Every hour spent designing a temporary environment and validating that it will not become a new way into the network is an hour not spent removing the attacker. That work carries its own risk, because infrastructure built quickly under pressure is the kind most likely to ship with a misconfiguration or an overlooked exposure. The team ends up carrying a second full job, one it cannot refuse and cannot fully delegate, at the exact moment its attention is most needed elsewhere.

Business operations decisions and communications.

The technical response is underway, the team is deep in containment and investigation, and the CISO should be leading it all. Instead, the CISO is being pulled away because the rest of the organization has questions: Should the company pay the ransom, what does the board need to hear and how often, what does the company say publicly and when? These are business, legal, and public-relations decisions that depend on facts that only the CISO and security team hold.

When a ransom demand is on the table, leadership wants to know whether paying is the fastest way back to operations, and that decision belongs to the CEO, CFO, board, and legal counsel. However the questions the decision turns on are all technical and fall to the CISO and their team: Whether the attacker's decryptor would actually work, whether data was stolen and how much, whether the attacker still has access, and how long recovery would take with the key versus without it.

The honest version of that analysis rarely supports paying. Decryptors are often slow and incomplete, and paying does not reliably end the incident. CrowdStrike's 2025 research found that 93% of ransomware victims had their data stolen anyway, 83% experienced another attack, and 45% failed to fully recover. Payment also carries legal exposure: The U.S. Treasury's Office of Foreign Assets Control has warned that facilitating a payment to a sanctioned entity can itself violate sanctions law, which means the security team's attribution work (identifying who the attacker actually is) can determine whether paying is even lawful.

2025

Insufficient incident response and recovery planning was cited as a top challenge by 22% of organizations overall

World Economic Forum Global Cybersecurity Outlook 2026

halcyon

Communication also falls on the CISO, because the security team is the only group that knows what is actually true at any given moment. Internally, the board and executives want continuous updates on scope, impact, and timelines. Externally, regulators, customers, and the market need statements, but nothing can go out until it is confirmed accurate, since anything that later proves false becomes its own crisis. This carries legal weight too, because public companies must disclose a material cybersecurity incident to the SEC within four business days of determining it is material, and that determination depends on what the investigation has found. Every statement has to be grounded in facts the team is still assembling, and getting one wrong creates a second problem on top of the one they are already fighting.

These decisions share a common structure: the CISO carries them without owning them. The authority to decide sits with executives, the board, and legal counsel, but the technical basis for every decision, and the blame if any underlying fact turns out wrong, sits with the security team. Underneath that is a quieter pressure, because the CISO's own job, and increasingly their personal legal exposure, often rides on the outcome of an incident they are simultaneously managing and reporting on. More than half of CISOs fear being fired over a breach on their watch, and prosecutions like the Uber case and the SEC's charges against the SolarWinds CISO have made that exposure real. Asking one person to serve as the neutral technical input for decisions that may end their career is a structural conflict, not a matter of character

Where Phase 1 Leaves Your Organization

Each of these decisions and activities is difficult on their own. What makes Phase 1 the determining factor in the entire recovery is that they all land at once, on the same small team. How long the team takes to work through all of it is the single largest variable in how long the organization stays down.

When Phase 1 finally ends, the recovery clock does not start over. However long it took to confirm scope and verify eviction becomes the foundation every later phase builds on, and the exhaustion and any shortcuts taken under pressure carry forward with it. The organization tends to believe the hardest part is over once the attacker is out, but what waits beyond it is the highest-pressure, highest-failure decision in the recovery.

2026

Only 44% of highly resilient organizations simulate cyber incidents or plan recovery exercises with ecosystem partners.

World Economic Forum Global Cybersecurity Outlook 2026

halcyon

Phase 2

Backup Verification and Clean Restore Point Identification

This phase exists because backups cannot be trusted by default. Attacking backups is part of the standard ransomware playbook, since the entire leverage of an attack depends on the victim being unable to simply restore and walk away. Veeam's own research found that backup repositories were targeted in 96 percent of attacks and successfully breached in 76 percent of them. The verification process is looking for three distinct things: that the backup still exists, that the data has not been tampered with, and that the backup does not carry the same vulnerabilities that let the attacker in or the backdoors they planted after gaining access.

What the Verification Process Involves

According to CISA's #StopRansomware Guide, the verification process must confirm the environment is free of the attacker's footholds before rebuilding, restore and validate from a recovery point confirmed to predate the intrusion, and take care not to re-infect clean systems on the way back to production. Backup vendors like Veeam formalize the middle step with restoring backups to a clean room, an isolated environment where a candidate backup can be examined before it is trusted, first before proceeding with the restore process.

Step one is identifying a candidate recovery point. Using the intrusion timeline from Phase 1A, the team finds the most recent backup that predates the attacker's entry, since any backup taken after that date may already carry the compromise. CISA's guidance is direct that this cannot be skipped, because a ransomware event is often evidence of an earlier, unresolved compromise, and restoring without first accounting for the attacker's footholds reintroduces them. The candidate cannot be chosen by backup date alone, because timestamps can be altered and the entry date itself is a forensic finding that may still be firming up.

Step two is restoring that candidate into isolation. Following Veeam's clean room approach, the backup is restored into a quarantined environment with no path to production, so it can be examined without risk.

Step three is validating the restored data inside that environment. It is scanned with malware and behavioral detection tools and compared against a known-clean baseline to surface planted tooling, modified files, or persistence mechanisms a surface integrity check would miss.

Step four depends on the result. A backup that passes becomes a validated recovery point, and restoration can begin. One that fails sends the team back to step one to select an earlier candidate, and the sequence repeats. The further back the team is forced to go, the more recent data is lost at each step, which is why a failed candidate is not just a delay but a direct cost to how much of the business can be recovered.

Unfortunately, there is no fixed duration for this process because the time is driven by how far the team has to search and how much they have to scan. Elastio reports that much of an organizations downtime is spent hunting for a clean recovery point rather than restoring data, and that about 58 percent of organizations recover into an infected state because they restored from a compromised backup. The time it takes scales with the environment: a small estate with immutable, pre-tested backups can clear verification in days, while a large enterprise with many backup sets and a long dwell time can spend far longer.

Where Business Pressure Compounds the Problem

The same organizational pressure that shaped Phase 1 returns here, and in some ways it is stronger, because the attacker is gone and the only thing now standing between the business and resumed operations is the security team asking for more time. From the outside, verification looks like caution rather than necessity.

How Recent and How Fast?

The pressure on verification arrives as two questions, asked in sequence. First, how recent of a restore point can the business use? Second, how fast can the team verify it, or rather, can verification be skipped altogether to start restoring sooner?

The first question is about data loss. Verification helps to confirm whether a backup is clean and in doing so reveals how far back the last clean one sits. When a long dwell time has poisoned the recent recovery points, the nearest verified-clean backup may be weeks or months old, and everything created since is lost if the team restores from it. Leadership, finance, and operations push back hard against that loss, because it is immediate and easy to measure while the risk of a hidden compromise in a newer backup is not. The pressure, then, is to pick a more recent backup and treat it as clean. A backup can pass a malware scan while still carrying the unpatched vulnerability or the hidden access the attacker used to get in, but that danger does not show up the way lost weeks of records do.

The second question follows from the first. Verification takes days the business does not want to spend once eviction is already confirmed and production is waiting. The pressure mounts to compress the work: scan less thoroughly, skip the clean room, or trust a recent backup outright. This is not a rare lapse. Veeam found that 63 percent of organizations did not quarantine or sandbox their backups before restoring, skipping isolation and scanning under pressure to restore quickly directly back into production. Real world examples show the consequences of skipping this step. In one case documented by Malwarebytes, an organization hit by Royal ransomware rebuilt its systems from backup and believed it had recovered, but endpoint monitoring installed afterward immediately flagged returning ransomware, outbound traffic to a Cobalt Strike command-and-control server, and inbound remote access attempts. The rebuild had restored operations and the attacker along with them, and a deeper hunt then turned up compromised domain admin accounts, a domain controller, and a server that the recovery had left in the attacker's hands.

External Forces That Stall Verification

The verification clock is not always in the security team's control. External forces can directly impact or control whether a restore is necessary or what the process of recovery and restoration looks like.

Global median dwell time across 2025 investigations rose to 14 days, up from 11 days in 2024.

Mandiant M-Trends 2026

The first is the insurance carrier and the incident response firm it brings in. Once a claim is opened, most policies require the organization to use the carrier's approved panel firm and to cooperate fully with its investigation; engaging your own vendors or failing to cooperate can itself void coverage. That hands the carrier significant control over the pace of the response. Its forensic investigation is not a neutral exercise either, it is scrutinizing how the attacker got in and whether the organization's attested security controls were actually in place, because those findings determine whether the claim is paid. Only about one in four closed cyber claims results in a payout, so the incentive to find grounds for denial is real. For the security team, this means parts of the response, from containment and eviction to starting recovery can stall while the carrier's process runs.

The second is the ransom decision, which competes directly with verification for the same clock. While leadership weighs whether to pay, backup verification may be paused or deprioritized, because if the attacker's decryptor turns out to be the faster path, the work of identifying and validating a clean restore point is effort the organization may not need. Negotiations alone commonly run eight to ten days, and if the decision ultimately lands on not paying, that time is gone and the team starts verification later than it otherwise would have. The decision to pay does not avoid the delay either. Decryptors are frequently slow, incomplete, or unreliable, with one analysis finding that 40 percent of organizations that paid still failed to recover all their data. An organization that pays, attempts decryption, and discovers it cannot fully recover is then back to restoring from backup, losing days or week of time in an already lengthy recovery process.

Where Phase 2 Leaves Your Organization

A backup restored without verification reintroduces the attacker, collapses the gains of containment and eviction, and sends the organization back to the start with less data and less time than it had the first time. The clock is not entirely the team's own, either: an IR firm or a pending ransom decision can delay when verification begins or runs. Done properly, and despite these pressures, the phase ends with the one thing the organization can finally trust: a confirmed clean restore point.

22

Second Handoff

Median handoff time between initial-access brokers and ransomware operators collapsed to 22 seconds in 2025.

Mandiant M-Trends 2026

Phase 3

Tiered System Restoration

With a confirmed clean restore point in hand, the work finally turns to rebuilding. Phase 3 is where systems come back online, and it is the first phase where the recovery progress can be seen and measured by the rest of the organization. What makes this phase deceptively difficult is that bringing systems back is not a single action, but an ordered sequence dictated by technical dependency rather than business urgency.

Phase 3 is also where the calendar cost becomes most visible. Restoration is not fast even when everything goes right. Industry recovery data puts core systems coming back within several days when clean backups exist, but full operational restoration, including secondary systems and edge cases, commonly runs two to four weeks for most organizations, and longer for large or complex environments.

How Restoration is Actually Done

Restoration follows dependency order, not importance. A system cannot come back before the systems it relies on, regardless of how critical it is to the business, because a restored application that cannot reach its database, authenticate its users, or resolve network addresses is not actually restored.

In practice the sequence is well defined. Following CISA's guidance, restoration proceeds from a predefined critical-asset list that accounts for the systems each asset depends on, and the recovery community has converged on a consistent order built around that principle. The identity plane comes first. Active Directory is rebuilt starting from the forest root domain, with credentials, Kerberos tickets, and service accounts reset from a clean baseline. Core network services come next, DNS, DHCP, and the authentication services that the rest of the environment needs to communicate, restored from clean images rather than from backups that might carry the attacker's persistence. Only once that foundation is confirmed working does restoration move to Tier 1 business systems in priority order, then Tier 2 and Tier 3 systems in turn. Done this way, core systems are typically back within several days.

The reason the restoration order is strictly defined is that improper restoration order can create compounding technical issues down the chain. An application restored before its database, identity, or DNS is ready may appear to come up and then behave unpredictably, and the symptom usually shows up in the system that was restored too early rather than in the foundation that is actually missing, causing the team to troubleshoot the wrong system. Dependencies can also run in circles, where, as DXC notes in its ransomware recovery guidance, Active Directory depends on a database and that database in turn depends on Active Directory, so the team has to identify and break the loop before either can come back. These types of issues are what makes the process slow, and the slowness is a function of the work itself rather than a lack of urgency.

Where Business Pressure Compounds the Problem

Phase 3 brings the same pressure as every earlier phase: business departments want their systems back. What's changed is that progress is now visible. As foundational systems come back online, every leader has a reason their function should be next.

That pressure pushes toward restoring out of order, and doing so carries a real cost even though the backups are clean. Every system pulled forward out of sequence diverts the team from the systems that should have come first, so giving in to the pressure does not just risk one broken system, it slows the entire restoration.

A second pressure arrives once systems are actually back. The restore point predates the attack, so every system returns missing whatever data was created between that point and the moment of encryption. The business unit that just regained its system is the first to feel the gap, and it turns to the security team to close it. But missing data is rarely something the team can recover. Unless a copy survived in a SaaS platform, a partner system, or some other untouched source, the only remaining options are to reconstruct it by hand or accept the loss, and both are business decisions, not technical ones.

$4.9M

Global average cost of a data breach.

IBM Cost of a Databreach 2026

Where Phase 3 Leaves Your Organization

Phase 3 is the phase that finally looks like recovery. Systems are coming back, the business can see them, and the pressure shifts from getting anything working to getting everything working faster. The danger in that visibility is that it makes the dependency-driven pace feel like obstruction, and the weeks restoration legitimately requires stack on top of however long Phase 1 and Phase 2 already ran. By the time the bulk of systems are back, the organization is often a month or more into the incident, and the people who lived through it are ready to call it finished.

Phase 4

Hardening, Regulatory Work, and Normalization.

Phase 4 begins while Phase 3 is still underway and continues long after the systems are back. What remains is hardening the environment against a repeat incident, closing out the regulatory obligations that began with the notifications filed during the response, and absorbing the legal and insurance consequences that follow a breach. This is also where every shortcut taken under pressure in the earlier phases comes due, as the gaps created during containment, verification, and restoration resurface as specific liabilities.

What Phase 4 Actually Involves

The first work of Phase 4 is hardening, and it starts before restoration is even finished. Every incident has a root cause – the vulnerability, misconfiguration, or exposed credential – that let the attacker in, and bringing systems back without closing it simply rebuilds the environment that was breached. This is the work that turns a recovered environment into a more defensible one, and it is the reason recovery is not the same as resolution. A business that restores operations without hardening has not ended the incident. It has only reset the conditions for the next one.

The second front is regulatory. The notifications filed during the response, to regulators, state attorneys general, and affected individuals, are not the end of the organization's obligations but the start of them. Once a breach is reported, the relevant regulators can open investigations into how it happened and whether the organization met its legal duties before and during the incident, and those inquiries run on the regulators' schedule, not the organization's. They demand evidence, documentation, and answers, often months or years after the team that lived through the incident has moved on to other work. For example, the HHS Office for Civil Rights can take years to resolve contested findings under HIPAA investigations. The organization is left defending the decisions it made under pressure during the incident, in a venue where the speed and improvisation that the crisis demanded are judged against a standard of what should have been done.

The third front is legal, and it arrives fastest of all. Plaintiffs' firms monitor breach notifications and now routinely file class-action lawsuits within days of a disclosure. These cases turn on whether the organization's security was reasonable and whether it acted appropriately once breached relitigating every decision made during the incident, often for years. Blackbaud, the cloud software provider hit by ransomware in 2020, is the clearest illustration. Despite real preparedness before the attack, including tabletop exercises and strong cyber insurance, they faced dozens of class actions consolidated into multi-district litigation alongside concurrent federal and state investigations, a tail that took nearly five years to resolve.

Where the Work Compounds on a Depleted Team

These activities of hardening and relitigating of the incident and response activity compound on top of the restoration process and the day to day work that is finally resuming. Every activity requires time and attention that the team often cannot balance, and the strain leads to burnout and departures–some voluntary, some forced.

The CISO, after an incident severe enough to cause major downtime, is frequently held accountable and let go as the fall guy. The people who ran the response are exhausted and some choose to leave. Each departure removes a piece of the institutional knowledge the hardening and investigations depend on, and the responsibilities fall to colleagues who are already overextended.

The effect compounds: a smaller, fatigued team is asked to carry more work across more fronts all at once which can lead to new security gaps.

Where Phase 4 Leaves Your Organization

By the time Phase 4 winds down, the business has long since returned to normal, but the organization is not the same as it was before the attack if it fully survived. The systems are back and being hardened, but the team that recovered them is smaller and depleted, the legal and regulatory consequences may still be years from resolution, and the true cost of the incident, counted in money, time, people, and disruption, is far higher than the downtime alone ever suggested. The attack that looked like a 22-day outage is, in reality, an event the organization absorbs for years.

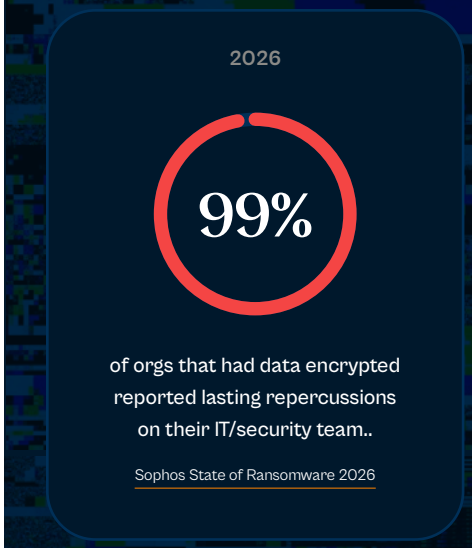

The Real Answer Is Not Faster Recovery

Everything in this paper points to a single, uncomfortable conclusion: recovery is not a reliable backstop. Even when backups are safeguarded and usable, recovering an entire organization takes weeks, and that downtime is enormously expensive. For some organizations it is fatal. Knights of Old, a 158-year-old UK transport company, had cyber insurance and expert responders on site the morning after its attack, and still collapsed entirely, putting 700 people out of work, because the ransom was unpayable and its backups and recovery systems had been destroyed. Lincoln College paid the ransom and got its data back, then closed its doors after 157 years anyway, undone not by the attack itself but by the cost and time required to rebuild. One company could not afford to recover. The other recovered its data and still could not survive the recovery.

Faster recovery is rarely achievable, and relying on backups as a surefire answer to a catastrophic attack is a gamble most organizations lose. The better investment is cyber resilience: the ability to contain an attack, limit its spread, and evict the attacker before the business is significantly disrupted. Catching an attack early changes the entire equation, because restoring a handful of critical files or a single endpoint is feasible in a way that restoring an entire organization is not.

Halcyon is built for exactly this. The platform is designed on the assumption that attackers will get in and focuses on taking away every mechanism ransomware needs to do damage once they do. Halcyon Data Exfiltration Protection (DXP) stops attackers from stealing data before it leaves the environment, File Resilience prevents encryption from taking hold, fleet inoculation limits the blast radius when a system is compromised, and Key Material Capture ensures encrypted data can be recovered without ever involving the attacker. All of it is backed by the Ransomware Operations Center (ROC), a 24/7 team of ransomware experts who monitor customer environments, identify attacks as they unfold, and contain and evict the threat before it can disrupt the business. Together, the platform and the ROC keep organizations operating through attacks that would otherwise force them into the unwanted recovery process detailed in this paper.

The 22 days is not a number to improve. It is a situation that can and should be avoided, by ensuring the business never has to depend on a lengthy recovery to survive.

Detect. Disrupt. Defeat Ransomware™

Ransomware will come for your business. Halcyon makes sure it doesn't win. Ransomware has evolved into a human-led operation, with real attackers using AI-powered tools to evade your defenses, destroy your backups, and shut your business down. Halcyon adds a ransomware-native defense layer that catches what generalist tools were never designed to stop, preventing data extortion and making encryption virtually impossible. And if something gets through, we recover in minutes, not weeks. Turning what could have been a business-ending event into a non-event.

For Security Leaders and Their Organizations, Halcyon Provides:

Prevention: AI and behavioral models trained exclusively on ransomware stop threats before execution while preventing attackers from disabling endpoint defenses, including the EDR tools already deployed in your environment.

Data Exfiltration Protection: Detection of bulk data movement stops data theft and double extortion before stolen data becomes negotiating leverage, while shielding security operations from tampering.

Recovery: Unique key capture and decryption technology reduces downtime from weeks to hours, without relying on backups, delivering the sub-24-hour recovery capability Halcyon experts identify as the threshold for fundamentally changing your ransomware posture.

24/7 Ransomware Operations Center (ROC): Our team of ransomware experts handles the fight against ransomware for you, at no additional cost, continuously hunting and continuously responding, at machine speed.

The Halcyon Guarantee: Reduce risk with our comprehensive ransomware warranty and recovery services.

With Halcyon, organizations can:

- Eliminate ransom payments
- Ensure operational continuity
- Prevent data extortion
- Close the speed gap between AI-powered attacks and AI-powered defenses
- Reduce blast radius through platform-level containment
- Maintain the recovery posture the board needs to see

Learn more at halcyon.ai or schedule a demo today.