White Paper

halcyon

A Guide to Ransomware Resilience

Five Questions Every Business Leader Should Ask

A non-technical guide for CFOs, CIOs, GRC, and Business Continuity leaders.

The Real Risk of Ransomware

Let's start with the assumption most executives make. It might be the most expensive mistake in enterprise risk management.

The assumption goes something like this: we're probably not a primary target. We're not a hospital, we're not critical infrastructure, and we aren't big enough for ransomware actors to go after. Our security team is competent, and we'd know if something was wrong.

That's a reasonable assumption. It's also precisely wrong, and the gap between those two things is where organizations lose.

Modern ransomware operations don't select victims the way many imagine they do, with a sophisticated actor studying your industry, evaluating your defenses, deciding you're worth the effort. That model describes only a small fraction of attacks. The majority of ransomware intrusions today begin with automated tools scanning the internet for exposed vulnerabilities, weak credentials, and misconfigured systems, at industrial scale, continuously, with no regard for industry, size, or brand recognition. Your organization isn't singled out; it'll typically be discovered. And if it presents an opening, it gets exploited.

The organizations hit hardest in the last three years weren't the ones that ignored security. Many had capable teams, modern tools, and recent audits. What they didn't have was resilience: **the specific capability to survive and recover when an attack hits.** Commonly, prevention alone fails more often than security briefings suggest.

The organizations hit hardest in the last three years weren't the ones that ignored security.

This isn't a conversation for your CISO alone. They're managing the prevention side of this problem and, for the most part, managing it well. This is a conversation for the people who will own the consequences when prevention isn't enough:

- The CIO who has to explain why recovery took three weeks
- The GRC lead who has to face the regulator
- The CFO who has to answer for the financial impact
- The board member who has to face the shareholders

None of these five questions requires a technical background to ask, but all of them require someone with authority and accountability to answer. The organizations that ask them before an incident are the ones that are more likely to survive.

By the time ransomware detonates, the operators behind it have typically spent days, sometimes weeks, inside your environment.

The Framing That Changes Everything

Security professionals are trained to think about ransomware as malware, malicious code that encrypts files and demands payment. That framing is partially accurate, but strategically insufficient.

Malware is a containment problem: you detect it, isolate it, remove it, and move on. The damage is bounded and isolated to a single event.

Ransomware isn't a containment problem. By the time ransomware detonates, the operators behind it have typically spent days, sometimes weeks, inside your environment. They've mapped your network, located your backups, identified your most operationally critical systems, read your cyber insurance policy, and chosen their moment. When the encryption begins, it doesn't hit one system; it hits everything simultaneously, timed for maximum impact.

What follows is a full business disruption event:

- Operations halted
- Finance systems locked
- Customer / patient data inaccessible
- Leadership making decisions under extreme time pressure with incomplete information
- Board members and regulators starting to ask questions
- Legal counsel suddenly very involved

Security is about keeping attackers out. Ransomware resilience needs to assume a campaign will get through despite your best efforts. Further, whether your organization can survive it in the time it takes to recover.

The CISO is responsible for the security measures that let the attacker in. But who owns what happens afterward?

Most executives carry a quiet assumption that even if ransomware gets through, recovering from backups is an exit ramp; the data says otherwise. Others might think that paying the ransom, while not an ideal option, is still a safety net. However, most organizations that pay a ransom are forced to pay a second ransom, and many still don't get their data back even if a key is provided.

Organizations that built in ransomware resilience didn't face that choice. These five questions are about what it takes to be a ransomware-resilient organization.

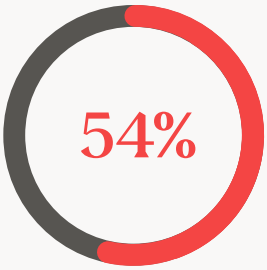

The overwhelming majority (over 90%) of ransomware victims had backups in place. However, **only 54%** were able to successfully restore their data from backups and that took over three weeks on average.

Source: Gartner Japan, Yano 2025

Question 01

What would ransomware actually cost your business? Most organizations are surprised when they calculate it.

Many executive teams have a rough intuition about ransomware costs: there's a ransom demand, you pay it or you don't, and either way the security team cleans things up. The real number, when organizations finally sit down and calculate it, is almost always a shock.

That's because the ransom is rarely the largest cost; in many cases, it isn't even close.

When a ransomware event causes disruptions, operations stop, not some operations, but all of them dependent on the affected systems. For most organizations today, that means ERP systems, financial platforms, customer databases, communications infrastructure, and supply chain tools going dark simultaneously. Every day those systems are unavailable is a day of lost productivity, deferred revenue, and mounting operational costs.

Now think about how long recovery actually takes. Industry data consistently shows that the average ransomware recovery takes three weeks or more. That's for organizations with reasonable backup infrastructure that doesn't itself get compromised. Three weeks of full or partial operational disruption at your organization's revenue run rate is a figure worth calculating when creating a ransomware resilience strategy.

Then layer in what the ransom figure doesn't capture: forensic investigation to determine how the attacker got in and what they accessed, legal fees as counsel navigates breach notification requirements and potential litigation, regulatory costs if customer or employee data was exposed, and the longer-tail expense of rebuilding customer trust that, for some organizations, never fully returns.

The picture that emerges when you add these costs together rarely resembles what leadership assumed ransomware would cost. For mid-sized organizations, total ransomware impact commonly runs into the tens of millions. For larger organizations, the number can exceed nine figures.

Take your organization's daily revenue, multiply it by the 22 days it takes to recover on average. Add your estimated legal and regulatory exposure. As rough as that number is, it is closer to your actual ransomware cost than anything currently in your risk register.

Question 02

If ransomware hits, who in the organization has the authority to coordinate response across IT, legal, finance, and communications simultaneously? Have they ever run a ransomware-specific response exercise?

This is the question that exposes how organizations are actually structured for a ransomware event versus how they think they are.

Here is what the first six hours of a ransomware incident actually requires.

It is 11 PM on a Sunday. Systems are going dark across your environment:

1. Your IT team is trying to understand the scope
2. Legal counsel needs to know whether breach notification timelines have started
3. Your CFO needs to understand the financial exposure
4. Your communications team needs to prepare for the possibility of public disclosure
5. Your board needs a briefing
6. Regulators may need to be notified within 72 hours of determining materiality.

Every one of those tracks is running simultaneously. Every one of them requires decisions. And every one of them needs a single point of authority that can move fast, cut across functional boundaries, and make calls without a three-day approval chain.

That is typically not a CISO role; a CISO typically manages the technical containment. What is needed is a cross-functional crisis leadership role, one that touches finance, legal, operations, communications, and executive governance all at once.

The organizations that navigate ransomware well don't just have the right technology. They have a named response owner with pre-authorized cross-functional authority, a tested playbook that covers each track, and the muscle memory from having rehearsed it before the incident rather than during it.

The organizations that navigate ransomware well don't just have the right technology, they have a named response owner with pre-authorized cross-functional authority.

Few organizations have all of those things. Not because their leadership is negligent, but because ransomware has been treated as a technical incident response problem. The moment it detonates across a business, it becomes a non-technical response situation.

The goal isn't to manage a ransom negotiation. The goal is to recover so fast that the negotiation never becomes the conversation. That requires a response infrastructure, owned at the executive level, built specifically for ransomware, not adapted from a general incident response template.

The goal isn't to manage a ransom negotiation. The goal is to recover so fast that the negotiation never becomes the conversation.

Question 03

In virtually every major ransomware attack, data is stolen before a single file is encrypted. What is your plan for when the attacker threatens to publish it?

There's a version of ransomware that most business leaders have in their heads: the attacker locks your files, you restore from backup, you're back online, you move on.

That version of ransomware effectively stopped being the standard playbook around 2019. What replaced it is considerably harder to recover from, and it almost never appears in business continuity plans.

Modern ransomware operations run in two stages.

In the first stage, which happens silently over days or weeks before the encryption event, the attacker identifies and exfiltrates:

Your most sensitive data	Intellectual property
Customer or patient records	Employee personal info (PII)
Financial statements	Contracts

Anything that has value then becomes leverage

In the second stage, they encrypt your systems and present two separate demands: pay to get your systems back, and pay again, or pay more, to prevent the stolen data from being published.

Here's what makes it categorically harder to manage than a systems outage: you can't restore your way out of it. A backup solves an encryption problem; it doesn't un-steal data. Once your customer records are in the hands of an attacker threatening public disclosure, your exposure exists independent of your recovery timeline. Your systems can be back online and you're still negotiating over whether your most sensitive information gets published, sold, or weaponized.

- For the CFO, this means a different category of financial exposure: potential regulatory penalties for data exposure, notification costs for every affected customer or employee, and civil liability that can run for years after the incident.
- For the GRC and legal team, it means breach notification obligations that trigger within days of discovering a potential exposure, regardless of whether you've contained the incident or determined exactly what was taken.
- For the board, it means a reputational crisis that unfolds publicly, at a pace the organization doesn't control, with the attacker holding the editorial calendar.

Pull up your business continuity plan. Find the section that addresses what happens when an attacker has your most sensitive data and is threatening to publish it. If that section doesn't exist, and in most organizations it doesn't, you don't have a ransomware recovery plan. You have a system restoration plan. Those aren't the same thing, and the difference will become apparent at the worst possible moment.

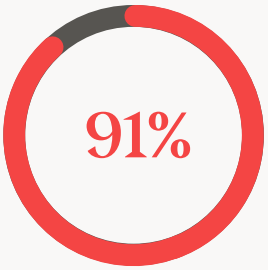

of ransomware attacks include data exfiltration.

Source: Arctic Wolf 2025 Threat Report

Question 04

Does your cyber insurance policy actually cover what you think it does?

For many executive teams, cyber insurance is how they reduce the risk from ransomware. The logic is reasonable: we have coverage, so if the worst happens, we're protected. It's in the risk register, the premium is paid, so it's safe to move on.

The problem is that the gap between what most organizations think their policies cover and what they actually cover has grown substantially as insurers have responded to an industry-wide ransomware crisis.

The policies organizations are renewing today aren't the policies they bought three years ago. Many organizations don't realize how significantly the terms have shifted until they're mid-incident, trying to trigger coverage.

There are four coverage assumptions worth testing directly against your current policy language.

1. **The first is the ransomware sub-limit.** Many cyber policies carry a headline coverage number, say $10 million, with a substantially lower sub-limit for ransomware-specific events. If your policy has a $2 million ransomware sub-limit inside a $10 million policy, you have $8 million of assumed coverage that doesn't exist for your most likely catastrophic event. That's not a technicality; it's a material financial exposure hiding in plain sight.
2. **The second is the notification window.** Virtually every policy requires you to notify your insurer within a defined window after discovering a covered event, commonly 24 to 72 hours. Ransomware incidents are chaotic in the early stages. If your organization spends three days trying to understand the scope before anyone calls the insurer, you may be out of compliance before coverage is triggered. Insurers use these windows, and they use them aggressively.
3. **The third is the security controls requirement.** Modern cyber policies include attestations, representations that you maintain specific security controls as a condition of coverage. These aren't always validated at policy issuance. They're validated at claim time. If an insurer determines your controls fell below the attested standard at the time of the incident, they have grounds to deny or significantly reduce the claim.
4. **The fourth is business interruption methodology.** The largest cost in most ransomware events isn't the ransom or the forensic fees. It's the revenue disruption during recovery. How that disruption is calculated varies enormously between policies. Two organizations with the same headline coverage can receive very different payouts for the same incident based on policy language alone.

Before you move on:

- What is your ransomware sub-limit?
- What is your notification window?
- What security controls does your policy require you to maintain as a condition of coverage?

You need to know what your policy actually covers before an incident occurs.

Question 05

If ransomware caused material harm to your business, what would your answer be when the board and regulators ask: What precautions were in place and how long it will take to recover?

In 2023, the SEC adopted rules requiring public companies to disclose material cybersecurity incidents within four business days of determining materiality, and to disclose their processes for assessing and managing cybersecurity risk annually. The FTC has pursued enforcement actions against companies whose security practices fell below what regulators deemed a reasonable standard of care. State attorneys general have brought actions under breach notification laws that show limited patience for organizations that treated security as a checkbox rather than a managed risk.

The regulatory landscape isn't moving toward less scrutiny of how organizations handle ransomware risk. It's moving decisively toward more.

The question of what precautions your organization had in place is no longer hypothetical. It's a question you will answer, on record, if a material incident occurs. And the answer that used to be sufficient, we had antivirus, we had backups, we had a security team, is being tested in courtrooms, regulatory proceedings, and shareholder meetings with increasing regularity.

The regulatory landscape isn't moving toward less scrutiny of how organizations handle ransomware risk. It's moving decisively toward more.

What regulators, boards, and plaintiffs' counsel are beginning to ask isn't only whether an organization had security tools. It's whether leadership understood the specific risk ransomware posed to their business, whether they assessed whether their existing protections were adequate against that specific risk, and whether they took action proportionate to what that assessment revealed.

Those are governance questions, not technical ones. They belong to the CFO, CIO, the board, and the GRC function, not the CISO. Answering them credibly requires documented evidence that the organization understood ransomware as a business continuity risk, evaluated the gap between their existing capabilities and genuine resilience against that risk, and made a reasoned, informed decision about how to address it.

Organizations that have done that work have a defensible answer. Organizations that treated ransomware as IT's problem, hoping that reasonable security hygiene would be enough, are learning that hope isn't a governance strategy.

What Do All Five Questions Have in Common, and What Does Ransomware Resilience Actually Require?

Read across these five questions and a pattern emerges that most organizations find uncomfortable: they already have tools, plans, and coverage that they were told would protect them. And the data says those tools, plans, and coverage aren't delivering what organizations believe they are.

What those outcomes reveal is a ransomware resilience problem. Not a security failure, not a budget failure, but a fundamental mismatch between what organizations have invested in and what ransomware resilience actually requires. Understanding that mismatch is the only way to understand why building genuine resilience requires something different from what you already have.

Why your security stack isn't a ransomware resilience strategy

Your endpoint detection and response (EDR) tools were built for a broad threat landscape. Ransomware operators know this, and they engineer around it deliberately. They use valid credentials and legitimate admin tools to move through your environment, disable or bypass EDR before detonating, and then execute. By the time your EDR fires an alert, the leverage is already established.

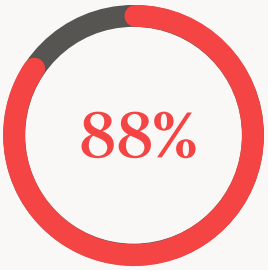

of organizations were confident their security could prevent ransomware.

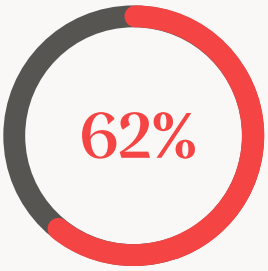

reported major operational disruption when it arrived.

The confidence was near-universal. The outcomes were not.

Source: Halcyon's Business Risk Report

> "EDR solutions have not been designed to tune into the specific behaviors and approaches of ransomware attacks... the amount of undetected ransomware calls into question the effectiveness of those solutions."
>
> – Omdia Technical Validation

Adding more to an EDR-centered stack doesn't build ransomware resilience. The architecture is not suited to the threat. More investment in the wrong architecture produces more of the same outcome.

Why your backups aren't a ransomware resilience strategy

Backups are the most trusted exit ramp in ransomware response, and ransomware operators target them first specifically because of that. Industry data shows attackers compromise backups in 94% of ransomware attacks. They corrupt them, encrypt them, or delete them before detonation. When your team reaches for the restore, it isn't there. Or worse, restoring from it reintroduces the threat.

Even when backups survive intact, 41% were given a recovery key but still had to rebuild their systems entirely. They paid. They still didn't recover cleanly. And they were still in a three-week rebuild.

Backups solve for accidental data loss. Ransomware resilience requires something they were never designed to provide: recovery capability that functions even when the adversary has specifically targeted your ability to recover.

Why your incident response retainer isn't ransomware resilience

An IR retainer is a post-incident service. By the time your IR firm is engaged, scoped, and on-site, the attacker has been in your environment for an average of eleven days. The exfiltration, the backup compromise, the lateral movement: all of it happened before anyone made a phone call.

IR firms are essential for forensics, containment, and notification compliance. They're not ransomware resilience. Resilience means the organization survives the attack at an acceptable cost. IR means the organization understands what happened after it didn't.

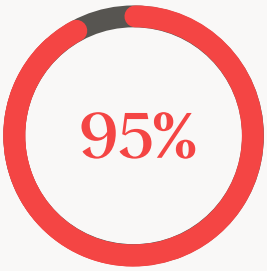

of organizations running modern enterprise security stacks still failed to thwart ransomware attacks over the past year.

Source: Omdia survey of 400 orgs

Why every quarter without ransomware resilience has a calculable cost

Small and mid-sized organizations aren't safer from this exposure; in fact, they're **more exposed**. The 2025 Ransomware Evolution Report found that smaller organizations are targeted at nearly four times the rate of large enterprises, precisely because their defenses are weaker and their recovery capacity is lower. The organizations that believe their size makes them less relevant to ransomware actors are the ones being targeted most aggressively.

The cost of ransomware resilience is fixed and knowable in advance. The cost of a ransomware event without it is uncapped, and the data consistently shows that organizations underestimate it by an order of magnitude until they experience it directly.

What ransomware resilience actually looks like

Halcyon is purpose-built for the specific failure modes that the five questions in this guide have surfaced, the ones your existing security investments were never designed to address.

Where EDR misses ransomware because it treats it as generic malware, Halcyon is trained exclusively on ransomware behaviors. Omdia's independent validation found that when tested against live attacks using the same technique used to kill EDR tools before detonation, Halcyon detected and isolated the threat in seconds.

Where backup-dependent recovery takes weeks and fails 78% of the time when a ransom is paid, Halcyon captures encryption key material in real time during an attack, enabling full data recovery in minutes without relying on backups and without paying a ransom. Omdia specifically validated this capability against Akira ransomware and confirmed complete recovery without a ransom payment.

$4.88M

Average ransomware breach cost across all industries, before the long-tail costs of regulatory action, litigation, and reputational damage that 58% of victims reported as a consequence of data exfiltration.

Source: Halcyon Ransomware Research Center

Before Halcyon, only 7% of users described themselves as confident in their ransomware resilience. After deployment, that number transformed to 99%. Not because Halcyon replaces their security stack, but because it closes the specific gap their stack was never built to cover.

Here's what that looks like in practice. A mid-market logistics company was hit on a Thursday evening. The attack followed the standard playbook: valid credentials, lateral movement, backup targeting, detonation timed for the weekend. Halcyon detected the encryption attempt in its earliest seconds, captured the key material in real time, and isolated the threat before it could propagate. By Saturday morning, operations were fully restored. No ransom was paid. No customers were notified. No board crisis. The COO's assessment: most employees came in Monday without knowing anything had happened.

The same attack pattern, at a peer organization running a conventional security stack and backup-dependent recovery, resulted in 22 days of operational disruption and a total recovery cost that exceeded $4 million. That's before regulatory exposure was calculated. The difference wasn't the attack. It was the ransomware resilience infrastructure in place before the attack arrived.

Ransomware resilience isn't a security posture metric. It's a specific, measurable capability: how fast can this organization recover from ransomware, with what data integrity, at what cost, without paying a ransom? Every organization should be able to answer that question, most can't. Not because they lack security investment, but because none of their security investments were built to answer it.

Halcyon is that answer. The only platform built from day one around the question that ransomware asks: When ransomware gets through everything else, can your business survive it?

The same attack pattern, at a peer organization running a conventional security stack and backup-dependent recovery, resulted in 22 days of operational disruption and a total recovery cost that exceeded $4 million.

halcyon

How Resilient Is Your Organization, Right Now?

The five questions in this guide define what ransomware resilience requires. For most organizations, working through them surfaces gaps that aren't visible in a security dashboard or a compliance report. Knowing the gaps exist is the first step. Understanding what closing them looks like in practice is what turns awareness into action.

Halcyon is purpose-built for exactly the failures this guide has surfaced. It actively protects your Volume Shadow Service from the destruction that precedes every major ransomware detonation. It captures encryption key material in real time so your organization can recover clean data without relying on backups that may have been compromised and without negotiating with the attacker holding your systems hostage. It wraps around your existing security stack, your EDR, your endpoint protection, your backup infrastructure, and closes the specific gap none of them were built to close: what happens when ransomware gets through everything else.

The best way to understand what that looks like for your organization is to see it. The Halcyon Ransomware Resilience demo walks you through the platform in the context of a real attack scenario, from the moment ransomware begins executing, through detection and key capture, to full recovery. No slides, no abstract capability claims, just the platform doing what it was built to do against the threat it was purpose-built for.

See the Ransomware Resilience Demo