HARVARD BUSINESS REVIEW

Digital Article / Cybersecurity and Digital Privacy

Your Company Needs a Plan for the First 72 Hours After a Cyberattack

Five steps to ensure the core of your business can get back up and running.

by Harpreet Sidhu, Jason Lewkowicz, Rouzbeh Hashemi, and Yusof Seedat

Published on HBR.org / October 12, 2026 / Reprint H09C0G

Audshule/Stocksy

Most companies operate with one of two assumptions about cyber disruption: that strong enough defenses will prevent it, or that a well-written business continuity plan will manage it when it happens. Our research, drawn from a global survey of 1,000 senior executives across 20 industries in 13 countries, suggests both assumptions are dangerously wrong.

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

AI-enabled adversaries have sharply accelerated the pace and scale of cyberattacks, with attacks increasing by 89% year-over-year in 2025, according to a recent CrowdStrike report. As autonomous, AI-driven techniques become standard practice, attackers can scan environments, identify vulnerabilities, and move laterally across systems in minutes. That speed has compressed breakout time—the interval it takes for an attack to spread beyond its initial entry point—from 62 minutes in 2023 to just 29, drastically reducing the window for response before disruption cascades across the enterprise. As AI models become more capable, the response window will keep narrowing, turning today's difficult or expensive techniques into routine attacks and leaving organizations even less time to detect and contain a breach. At the same time, the cost of discovering exploitable vulnerabilities has already collapsed to almost nothing.

Yet according to our research, most (73%) of executives still believe predict-and-control security approaches are sufficient and persist with a strategy that experience disproves. While 81% expect that if a cyberattack struck today, their critical systems would be restored within 10 days our analysis of real-world cyber disruption cases places the actual average at three to six months. And while almost every company now has a business continuity plan, the evidence suggests many of these plans are not built for the realities of cyber disruption: business interruption remains the single largest driver of cyber losses, accounting for more than half of all cyber claims.

The gap between recovery expectations and reality exposes a structural misunderstanding of resilience. Traditional business continuity plans often assume that companies can restore operations quickly from fundamentally trustworthy systems.

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

Those that recover fastest start from a different premise: in a cyber crisis, the business may need to operate before its technology environment can be fully trusted. Instead of attempting to restore everything at once, they identify and preserve the minimum core operations needed to sustain trust, maintain continuity, and support decision-making, even when essential platforms like phone lines, Wi-Fi, internet, communication networks, and information flows remain unavailable or compromised for extended periods.

These companies pinpoint what we call their "minimum viable enterprise": the smallest set of decisions, applications, and processes needed during extreme events. By preparing this essential core, they can recover much faster than companies that struggle to determine what to recover first, who has the authority to make decisions, or what manual fallbacks to activate when digital tools fail.

Protection and Resilience

After a global IT outage in 2024, for example, one U.S. airline we advised was largely operational by the next day because it had already defined the essential operations, decisions, and fallbacks needed to keep moving during disruption. AI can make this approach more dynamic and precise, helping organizations continuously map their environments and dependencies, identify the systems, data, and connections that support the essential operating core, and keep that understanding current as the business evolves. Combined with zero-trust controls that restrict lateral movement and limit the blast radius, this allows recovery to focus on what matters most rather than everything at once.

The problem is that many organizations still equate resilience with stronger protection. Fifty-four percent of executives in our survey believe strong protection controls—such as firewalls, multi-factor

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

authentication, and encryption, among others—are enough to ensure business continuity after an attack. But protection and resilience serve different purposes. Protection lowers the probability of disruption; resilience determines whether the business can continue operating when disruption occurs.

The activities matter less than the purpose they serve. In a conventional recovery plan, identifying critical functions, mapping dependencies, setting recovery objectives, and running exercises help determine what to restore first. In a cyber crisis, that is not enough. Most leaders (81%) already recognize managing cyber disruption involves operating under degraded system conditions. Companies may have to operate while systems are unavailable, data is unreliable, infrastructure is compromised, and adversaries remain active. The goal is to keep the business running during disruption and restore the highest-priority functions first.

Cyber resilience now requires the same discipline that companies apply to planning for other catastrophic disruptions: identifying, in advance, what could bring the enterprise to a halt and redesigning around those points of failure. Toyota, for example, has treated recovery as the start of an operational redesign ever since an earthquake in 2011. The company identified the 500 parts most likely to halt production, diversified suppliers, and built inventory buffers, then continued updating those measures as its supplier network evolved. As a result, when an earthquake struck Kumamoto in 2016, the disruption that might once have lasted months was contained within days. The lesson: Leaders must know before a crisis what can bring the business to a halt and what must keep running when broader operations are impaired.

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

That shift from recovery planning to operating-core design is what distinguishes the companies that recover fastest. In our research and advisory work, five steps stand out.

1. Define the essential operating core, not just critical processes. The first move is to determine what the enterprise must keep doing when its normal operating environment can no longer be trusted. Each business unit should ask one question: if all applications were down for 72 hours, which outcomes would create the greatest harm if they stopped?

This exercise goes beyond a faster version of a traditional business impact analysis (BIA). A BIA helps organizations catalogue essential activities, estimate the cost of an interruption over time, and set recovery priorities. This exercise identifies the essential operating core that must keep functioning when technology systems may be compromised and the normal operating environment can no longer be trusted.

It also forces leaders to make the hardest trade-offs before a crisis begins: which outcomes must be preserved, which processes can move to manual fallbacks, and which services can pause without causing irreversible harm. Put simply, a BIA shows what disruption would cost; this exercise defines what the enterprise must still be able to do once disruption is underway.

The 72-hour question moves the conversation from "What technology platforms do we run?" to "What can we not afford to stop?" The focus should be on enterprise outcomes, not technical components: customer experience, employee safety, regulatory obligations, and essential service delivery. It should also clarify what must be stabilized immediately and what can be sustained or brought back online over the following weeks.

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

For instance, many hospitals that recovered quickly from the 2024 global IT outage mentioned above had already determined that their non-negotiable outcome was patient safety. They knew which clinical applications were essential, which were non-essential, and what manual protocols to activate when digital tools failed. Hospitals without those fallbacks for imaging, patient records, and emergency communications faced far more severe clinical risk.

2. Predefine the decisions that determine recovery.
Once leaders define the essential operating core, the next question is who can make the decisions that keep it running under pressure. In a cyberattack, the limiting factor is often not technical capability but decision-making speed: who has authority, what conditions trigger action, which trade-offs are acceptable, and what the organization can stop.

Those choices should be made before disruption, not during it. Companies need to predefine recovery sequencing, decision authority, escalation triggers, and trade-offs across business units. Which capabilities recover first? Who can pause a service? Who decides whether to shift to manual operations? When should leaders communicate with customers, regulators, suppliers, or employees? These are not administrative details. They are the operating model for crisis conditions.

That sequencing should distinguish between two stages of recovery. The first, stabilization, happens on day one. Its purpose is to protect what cannot be lost: customer-facing capabilities, employee safety, and legal or regulatory obligations. The second, sustenance, unfolds over the following weeks. It focuses on restoring revenue, rebuilding operational capability, and reestablishing the processes that keep the enterprise running. Across both stages, every major decision should balance three priorities at once: trust, continuity, and compliance.

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

Only 43% of executives told us that cyber resilience is fully governed as a business outcome shared across enterprise leadership. That leaves many organizations treating recovery as a technical issue until the moment it becomes an enterprise issue. Companies that recover fastest do the opposite. They make decision rights a permanent part of the operating model.

The 2024 global IT outage mentioned above showed why that matters. One U.S. airline we advised assembled the right operating teams and IT experts within an hour because it had already defined who needed to act, what decisions had to be made, and in what sequence. It cancelled a manageable number of flights and was largely operational by the next day. Other airlines took days to recover, cancelling thousands of flights, stranding millions of passengers, and incurring hundreds of millions of dollars in losses. When crew-tracking systems failed, the problem was not merely technical; it exposed the absence of pre-defined operating decisions.

3. Expose the dependencies that could stop it.
After defining what must keep running and who can act, leaders need to identify what could still stop the operating core from functioning. Leaders should map dependencies between applications, data, infrastructure, suppliers, and third-party services that could take down the outcomes the business cannot afford to lose.

As digital ecosystems grow more complex, these hidden failure points become harder to see. Sixty-nine percent of organizations told us their operations are increasingly dependent on complex internal and third-party digital ecosystems. Yet only 38% have a clear view of their value chain and its most critical interconnections. That gap leaves companies exposed to failures they do not see until those failures have already spread.

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

That is what happened when a ransomware attack on a Chicago-based software provider left half the nation's car dealerships struggling to operate. Many dealerships had not understood how dependent they were on the same shared infrastructure provider. When that provider went down, dealerships reverted to pen and paper for weeks, exposed by a dependency they had not seen until it failed.

The same principle applies inside the enterprise. Leaders need to know not only which applications support critical functions, but also which identity systems, data feeds, communications channels, cloud services, suppliers, and operational workarounds those functions rely on. The goal is to find the fragile points before a crisis forces improvisation.

4. Build a clean environment to run the essential core.
Once leaders know what the operating core depends on, they need a validated place from which to run or restore it. This is where cyber recovery differs most sharply from conventional continuity planning. After a ransomware or wiper attack, the primary technology estate may no longer be trustworthy. Systems that appear intact may have been compromised weeks before the attack was detected. Rebuilding from them can reinfect the business.

A clean recovery environment provides a known-good foundation. It should be built on validated and immutable backups, tested under real conditions, and kept separate from the infrastructure it may need to replace. Its purpose is not to restore every application. Its purpose is to support the capabilities, decisions, data, and communications required to preserve trust, continuity, and compliance during the first days of disruption.

The 2019 ransomware attack on, a medical clinic in California, demonstrated how quickly the absence of a clean recovery environment becomes existential. When ransomware infected the clinic's servers in

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

August of that year, it encrypted not only patient health records but also the clinic's backup hard drives. With no trusted foundation from which to restore operations, the clinic determined that rebuilding its medical records was impossible and permanently closed its doors in December 2019. A dedicated recovery environment which is kept separate from the infrastructure it may need to replace turns that kind of irreversible loss into a recoverable one.

Leaders should set recovery time objectives through the same lens. Instead of asking how quickly they can restore every system, they should ask how quickly each essential capability must become available. For the most vital functions, the benchmark is typically restoration within 24 to 72 hours. Without that clock, stakeholders are left in ambiguity.

An Australian building-materials manufacturer that had stored its recovery data in a protected backup repository that ransomware could not reach illustrates how a clean environment can expedite recovery. After an attack deleted hundreds of thousands of files and modified thousands more, the company identified its mission-critical data first and began restoring the last working copies from the protected repository. The business was operating at 100% efficiency within an hour; it subsequently reported zero data loss, no ransom payment, and full recovery in less than 24 hours.

5. Test under real scenarios and reassess as the business evolves. The final move is to prove that the essential operating core can function under pressure. Only 35% of organizations regularly run real recovery and continuity exercises rather than relying on audits or documentation. That gap carries real consequences because resilience depends on whether the business can sustain critical operations when systems, data, people, suppliers, and communications are under stress.

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

The best exercises do not simply ask whether systems can be restored. They test whether the enterprise can preserve its core capabilities when systems fail, such as which services recover first, who has authority to act, what manual fallbacks should be activated, how to handle regulatory communications, and how leaders should balance customer trust, continuity, and compliance. While compliance audits show whether controls exist, recovery exercises show whether the company can keep operating when those controls fail.

These exercises also need to be repeated as the business changes. The essential operating core changes as products, technologies, partnerships, and operating models evolve, as do the capabilities that deserve priority and the interconnections that support them. Without regular reassessment, companies may find themselves prepared to protect an enterprise that no longer reflects how the business works.

From Recovery to Resilience

Continuity planning remains essential, but leaders must redesign it for continuously evolving cyber conditions. The belief that an organization can simply restore its way out of a major breach is a dangerous fiction. The question is no longer *if* a disruption will occur, but *when*—and leaders who wait for a crisis to define their response are already behind.

These five moves can help organizations shift from reactive recovery to proactive resilience, but only if leaders make hard choices before disruption begins: which outcomes deserve priority, who can act, and how the enterprise will operate when its normal environment cannot be trusted.

Building this resilient core requires a fundamental change in operating model and leadership mindset. By defining, protecting, and rehearsing the operating core the business cannot afford to lose, companies can

Copyright © 2026 Harvard Business School Publishing. All rights reserved.

navigate the inevitable with confidence. In an era of continuous cyber disruption, advantage will belong to organizations that know, before the crisis begins, what must keep running and how to sustain it when the rest of the enterprise is impaired.

This article was originally published online on October 12, 2026.

HS

Harpreet Sidhu is the Global Cybersecurity Lead at Accenture.

JL

Jason Lewkowicz is the Global Cyber Resilience and Defense lead at Accenture.

RH

Rouzbeh Hashemi is the Global Data and Infrastructure Resilience lead at Accenture.

YS

Yusof Seedat is the Global Research and Thought Leadership Lead at Accenture.

Copyright © 2026 Harvard Business School Publishing. All rights reserved.