O'REILLY®
Report

The Cyber Resilience Reckoning

A New Strategy for Surviving the Agentic Threat Landscape

Govind Rangasamy

The Cyber Resilience Reckoning

A New Strategy for Surviving the Agentic Threat Landscape

Govind Rangasamy

O'REILLY®

The Cyber Resilience Reckoning
by Govind Rangasamy

Published by O'Reilly Media, Inc., 141 Stony Circle, Suite 195, Santa Rosa, CA 95401.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*https://oreilly.com*). For more information, contact our corporate/institutional sales department: 800-998-9938 or *corporate@oreilly.com*.

Acquisitions Editor: Simina Calin
Development Editor: Michele Cronin
Production Editor: Jonathon Owen
Copyeditor: Paula L. Fleming
Cover Designer: Susan Brown
Cover Illustrator: Ellie Volckhausen
Interior Designer: David Futato
Interior Illustrator: Kate Dullea

September 2025: First Edition

Revision History for the First Edition
2025-09-19: First Release

979-8-341-65986-5

[LSI]

Table of Contents

Acknowledgments

I would like to thank Katherine Demacopoulos for her fantastic effort across all aspects of the book, Anna Griffin for recognizing the need for the book in the market and sponsoring it, and Chris DiRado for his security expertise and assistance.

To my wife, Bhuvana, and amazing sons, Pranav and Sanjit—you remind me every day what truly matters. Your love, curiosity, jokes, and hugs gave me the joy and energy to keep going.

CHAPTER 1

The Cyberattack Is Coming. Can You Recover Confidently?

An organization's lifeblood flows through its networks, applications, and data stores. These vital systems face mounting threats as cyberattacks become more frequent and sophisticated. Cybercriminals are already deploying sophisticated AI-enhanced tools, but a far more dangerous threat is emerging. Agentic artificial intelligence (AI), which can reason, plan, and act autonomously, will revolutionize cybercrime tactics, making attacks more scalable and efficient.

Unlike traditional ransomware, which follows preprogrammed scripts, agentic AI can adapt its strategy in real time, learning from defensive responses and evolving attacks faster than human defenders can counter.

High-profile incidents illustrate that no industry or geography is immune. We've seen examples ranging from the poisoning of a Florida water treatment plant and the 11-day shutdown of Colonial Pipeline to paralyzing school districts with ransomware attacks and the wholesale encryption of hotel and casino systems.

Even the most hardened perimeter defenses and advanced threat intelligence programs, like the Financial Services Information Sharing and Analysis Center (FS-ISAC) and the US Cybersecurity & Infrastructure Security Agency (CISA), are only part of the story. Today's adversaries are well funded, extraordinarily patient, and laser focused on disabling the target organization's ability to recover.

Understanding Today's Cyberthreat Landscape: The Industrialization of Cybercrime

Just a few years ago, ransomware attacks were largely confined to a handful of sophisticated hacking groups. Today, *ransomware as a service* (RaaS) platforms are available on the dark web for any moderately skilled criminal to rent. The RaaS business model operates like any corporate software as a service (SaaS) platform, offering subscription fees and revenue-sharing arrangements that make cybercrime accessible to anyone willing to pay.

These criminal enterprises provide turnkey access to polished extortion frameworks, complete with call centers for victim support, negotiation assistance for ransoms, and even "guarantees" of data deletion if victims refuse to pay. This industrial-scale approach has led to an explosion in attacks. Criminals no longer need to develop their own malware. They can simply choose from a menu: LockBit, REvil, DarkSide, Conti, BlackCat, and dozens more.

Each strain offers different specialized features designed to maximize damage and leverage. Modern ransomware variants routinely exfiltrate data before encryption to threaten public disclosure, systematically target and destroy backup systems to eliminate recovery options, and deploy payloads capable of wiping entire networks or cloud accounts within minutes.

Agentic AI Ransomware: When Attacks Think, Learn, and Adapt

The emergence of agentic AI marks a fundamental shift beyond traditional RaaS platforms. Unlike generative AI, which assists with specific tasks, *agentic AI* is proactive and can solve complex problems and make decisions autonomously. These AI agents don't simply execute preprogrammed attacks. They learn and adapt their strategies based on the specific environment they encounter.

For example, in controlled testing, Unit 42 researchers demonstrated an entire attack, from initial compromise to data exfiltration, in just 25 minutes. The speed differential is staggering: where human attackers had needed two days (median) to effect data

exfiltration, AI-assisted attacks accomplished the same objective approximately 100 times faster.

Dual Extortion and Beyond: Attacks That Go Deeper

Modern ransomware often follows a two-phase playbook:

1. Exfiltrate sensitive data.
2. Encrypt critical systems.

Even if a victim has off-site, second cloud region backups, the threat that sensitive data might be publicly exposed exerts immense pressure to pay. In late 2024, the LastPass breach exposed the encrypted vault backups of millions of users, and even though master passwords remained safe, the very fact that an attacker held a copy of each vault created a crisis of confidence.

Meanwhile, nation-state actors deploy ransomware not for profit but for strategic impact, for example, by shutting down the key operations of pipelines, utilities, healthcare systems, and government services at scale, directly affecting operational technologies.

The 2017 WannaCry ransomware attack offers a stark illustration. It severely disrupted hospital systems across the UK's National Health Service, leading to canceled medical procedures and diverted ambulances. This attack clearly demonstrates how cyberattacks can endanger lives when critical infrastructure fails.

The Ubiquity of Targets

The democratization of agentic AI through accessible platforms has shattered any illusion that certain industries might be immune from attack. Agentic AI can help plan attacks and then carry them out autonomously, making attacks more scalable and efficient while lowering the barrier to entry for cybercriminals:

- *Education:* In 2022, Vice Society held the Los Angeles Unified School District's data hostage, impacting over 1,000 schools and 600,000 students.
- *Energy:* That same year, Suncor Energy in Canada saw ransomware knock out its Petro-Canada card systems, stranding drivers at pumps nationwide.

- *Hospitality:* In 2024, BlackCat attacked MGM Resorts, silencing slot machines and reservations across 30 properties.
- *Retail and services:* By April 2025, Marks & Spencer in the UK had shuttered stores when DragonForce encrypted its business systems, and Mailchimp and SendGrid had suffered global phishing campaigns.

The moral is clear: our adversaries have both the tools and the incentives to strike anywhere. As digital transformation accelerates, connecting ever more devices, processes, and partners, the attack surface grows. The days when IT could wall off "critical systems" behind a fortress-like firewall are over. Every endpoint, cloud service, and third-party integration is a potential entry point.

Furthermore, the human interface itself, whether through a convincing AI-generated deepfake call or an employee using a compromised personal device for work, can become an effective gateway for an attacker.

Why Organizations Are Still Surprised by Attacks

Despite the mounting evidence of increased vulnerability to cyberthreats, many organizations continue to be caught off guard. This persistent vulnerability stems from deeply ingrained assumptions about cybersecurity that no longer match today's reality. Three critical blind spots—a prevention-only mindset, siloed teams and runbooks, and a belief in cloud resilience—leave even well-defended organizations exposed.

The Prevention-Only Mindset: Layers of False Confidence

Historically, cybersecurity has evolved through distinct waves; during each, cybersecurity experts believed they had finally mastered the protection of their organizations' information. Table 1-1 outlines this evolutionary trajectory.

Table 1-1. Evolution of cybersecurity through distinct waves

Era	Focus	False promise	Reality
1990s	Perimeter defenses	Firewalls and border routers would stop unauthorized access.	Attacks circumvented walls through phishing, social engineering, and insiders.
2000s	Email security	Scanning would eliminate malicious messages.	Malware hid inside legitimate traffic streams and attachments.
2005+	Network security	Monitoring would detect anomalies.	Sophisticated threats appeared harmless until they triggered breaches.
2010+	Endpoint protection	Antivirus would block execution.	Fileless malware and zero-day exploits bypassed signature detection.
2015+	Identity security	Zero trust would allow only authenticated access.	Stolen tokens, API keys, and misconfigured permissions created gaps.
2020+	Cloud security	Providers would handle security.	Attacks targeted misconfigured cloud permissions and container registries.

Despite the rising tide of cyberthreats, many organizations remain locked in a prevention-first mindset. We invest heavily in next-gen firewalls, endpoint detection and response (EDR), security information and event management (SIEM) platforms, threat feeds, and red-team exercises, only to discover that these controls are necessary but not sufficient.

As soon as an attacker gains a foothold, whether through stealing credentials, executing zero-day exploits, phishing, or compromising the supply chain, the perimeter defense crumbles. Figure 1-1 illustrates the layered defense in which the cybersecurity industry places its trust.

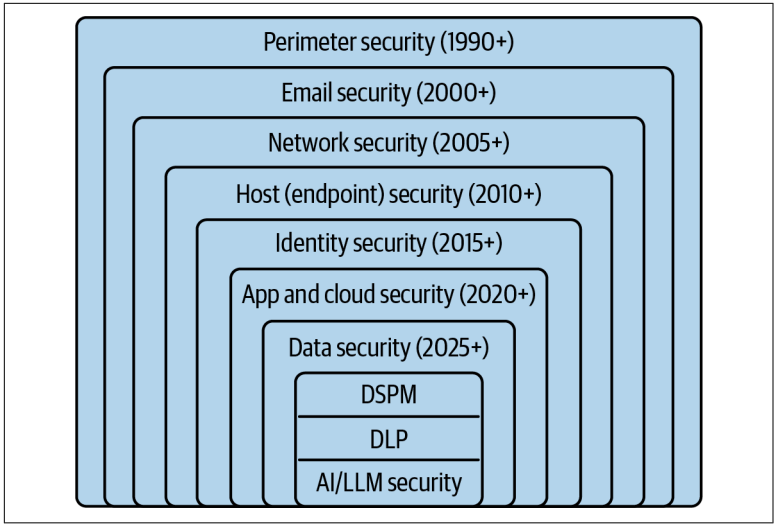

Figure 1-1. The layered defense of the cybersecurity industry

This raises a paradox: by focusing too heavily on prevention, we underinvest in recovery and particularly in recovery testing. We treat backups as a compliance checkbox rather than as a strategic asset, and we treat cyber recovery testing like an annual disaster recovery (DR) exercise—if we do it at all.

When the alarm sounds, we scramble to follow the guidance of ad hoc runbooks only to find them stale, incomplete, untested, and incompatible with today's dynamic environments.

Siloed Teams and Disconnected Runbooks

Cybersecurity, cloud operations, application development, enterprise architecture, and business continuity teams often work in silos, each with its own processes, tools, and priorities.

The result? Policies and runbooks live in PowerPoint decks, Word docs, and ticketing systems and are rarely, if ever, actually used during a real crisis. Connections and dependencies between applications, network configurations, identity systems, GitHub repositories, container registries, database servers, and data backup copies are seldom documented.

The first time you try to rebuild is the worst time to discover a missing piece.

Illusions of Cloud Resilience

Many CIOs, CTOs, and CISOs believed that moving to the cloud under the pretext of digital transformation would magically solve the organizational recovery problem. Hyperscale providers advertise multizones, regions, snapshots, replicated copies, and native backup tools that promise to dramatically cut recovery times. To achieve proper resilience, teams often need to piece together over a dozen tools and services.

Scale amplifies the problem. Organizations are not properly managing their data backup posture across all their cloud accounts. The "shift left" culture, which gives developers greater operational responsibility, has arguably created more risks than it has mitigated.

Hyperscale providers keep releasing more services and tools to ease the self-service model, but this very model has led to broken processes, ultimately leaving organizations at even greater risk.

These gaps become evident under pressure. In one recent case, a financial services firm tried to fail over to a second region after a simulated breach but discovered that data encryption keys and identity roles had not been replicated. The "region-failover" script failed, stranding the recovery site in an unusable state.

Cloud alone is not a cure-all. It demands fully tested, application environment–wide rebuilds to help make sure every configuration, every credential, and every object is in place.

Resilience is impossible without recovery

The harsh reality is that most organizations have built their entire cybersecurity strategy on a dangerous illusion: that they can prevent every attack. This prevention-first mindset creates a false sense of security that crumbles the moment an attacker breaches the perimeter.

Without a proven ability to recover quickly and completely, even the most sophisticated defenses become meaningless. This is because resilience isn't about avoiding failure—it's about bouncing back from it.

Redefining resilience as "rebuild confidence"

If resilience means anything, it means confidence. Confidence that you can switch the lights back on when they go out. Cyber resilience is not firewall uptime, nor is it patch cadence. It is our ability to recover business-critical application services (e.g., customer portals, payment systems, production lines, or electronic health records) in minutes or hours rather than days or weeks.

The motto "Rebuild is imperative" demands three core shifts:

- *From backup to full application environment rebuild:* The focus must expand beyond making file copies and block snapshots to having ability to rebuild every application component (i.e., network, compute, storage, identity, and particularly their dependencies) to get entire critical services up and running.
- *From infrequent DR drills to regular recovery testing:* Instead of once-a-year failover exercises, teams must run monthly automated "rebuild drills" in cloud accounts isolated from production.
- *From siloed playbooks to cross-functional recovery as code (RaC):* Security, cloud, architecture, development, and DR teams must agree on shared runbooks that are versioned as code and tested in unison.

What if RaC could be created automatically and updated regularly?

The Cost of Complacency

When minutes of downtime can cost an organization thousands, it's sobering to realize what a full day offline truly means. Even a single day can unleash devastating consequences: millions in lost revenue, crippling regulatory fines, and severe (often irreparable) reputational damage.

Retailers are shuttered, unable to sell; manufacturers are disabled, unable to ship; and hospitals are paralyzed, unable to access vital patient records. Every minute lost can mean a customer enraged, a partner betrayed, and a brand critically damaged.

By contrast, organizations that adopt consistent rebuild strategies report slashing their average recovery times from 48 hours (about

2 days) to less than 2 hours and a confidence that even the most sophisticated attack cannot keep them down.

That is the return on investment of rebuild: not just saved dollars but saved *trust*. The next chapter will dig deeper into developing a Rebuild function.

CHAPTER 2

The Illusion of Safety: Why Old Recovery Methods Fail and Why the NIST Cybersecurity Framework Must Embrace *Rebuild*

You can't put out a forest fire with a teacup of water.

—Adapted from the wisdom of wildland firefighters

In the late 19th century, engineers built towering levees between frontier towns and the Mississippi River, confident that sheer mass could hold back any flood. Families picnicked atop these embankments, believing that the battle against the water was won. Yet when spring thaws unleashed unprecedented torrents, the levees cracked like eggshells and water poured through, drowning the town.

The lesson to be learned from their sodden foundations: no defense, however grand, is impregnable when it rests on flawed assumptions.

Today's digital infrastructure faces a similar crisis of misplaced confidence. The cybersecurity and backup industries have built their own levees, each convinced of its solution's adequacy:

- The cybersecurity industry erects barriers: perimeters, firewalls, endpoint defenses, cloud security, and identity frameworks.
- The backup industry focuses on data vaults, tape archives, snapshots, and replication strategies.

Yet high-profile breaches and ransomware incidents have shattered this illusion. No matter how high the walls, attackers persistently find a way through. Likewise, traditional backup methods prove woefully inadequate when attempting to restore fully compromised ecosystems after modern ransomware strikes. This chapter will explore why conventional backup approaches are insufficient and why a Rebuild function can substantially improve recovery outcomes.

The Dangerous Complacency of Backup and Data Protection

For over 50 years, we've believed that just protecting data was sufficient. Data protection systems have emphasized cost-effective, long-term retention, not protection of the entire application ecosystem. Therefore, when organizations attempt restoration after a ransomware event, they frequently discover a cascade of critical failures that render their backups nearly useless.

For example, tape vaults preserve files but restoring servers takes days. Disk-based backups improve speed but live on vulnerable networks. Disaster recovery sites promise seamless failover but consistently fail due to stale procedures, configuration drift, and knowledge gaps.

The fundamental flaw isn't in the data protection—it is the illusion that backing up data equals the ability to restore functioning systems. True recovery requires not just offline data storage but also comprehensive system restoration capabilities that most backup strategies fail to deliver.

Modern Ransomware: The Death of "Backup and Restore"

Ransomware attacks don't merely discover backups—they hunt them down first**,** making certain that your Plan B is compromised before Phase Two begins.

Once the attackers gain domain admin rights, they systematically dismantle the organization's recovery capabilities by disabling or deleting snapshots, tampering with backup retention policies, and corrupting supposedly immutable vaults. They even compromise

the orchestration layers that manage these systems, making every potential recovery path a dead end.

The threat extends beyond simple destruction. Dual-extortion gangs have perfected a calculated approach: they first steal sensitive data to threaten public disclosure, then encrypt what remains to paralyze operations. This two-pronged assault maximizes leverage, as organizations face both operational shutdown and reputational catastrophe.

The Mirage of Immutable Storage

Immutable storage is designed so that once written, snapshots cannot be altered. Yet attackers have developed sophisticated counterstrategies that expose the technology's fundamental limitations. Cybercriminals hijack the management plane or control layer to alter immutability policies and settings before snapshots complete, effectively neutering the protection before it takes effect. They also exploit configuration vulnerabilities to grant themselves the power to purge or re-encrypt archives, turning the organization's own security controls against it.

Even when vaults remain technically secure, their scope of protection remains dangerously limited, safeguarding data but leaving network configurations, microservices meshes, and identity trees completely exposed to attack.

Data Security Tools: Helpful yet Incomplete

Recent advancements in security have introduced sophisticated tools like data security posture management (DSPM) for comprehensive data visibility, data loss prevention (DLP) for monitoring data movement, and AI-supported security tools that enable intelligent threat detection and response.

While these technologies represent significant progress in cybersecurity capabilities, they focus primarily on prevention and detection rather than on comprehensive recovery, leaving organizations vulnerable when attackers successfully breach their defenses.

Cloud Recovery's Siren Song and Hidden Shoals

Hyperscale cloud providers promised infinite capacity, instant snapshots replicated across regions, and self-service DR pipelines. Many organizations, told their data would be safer, migrated terabytes within weeks. However, these promises masked fundamental gaps that only became apparent when organizations needed their recovery capabilities most.

Beyond Traditional Backup

Legacy methods fail to capture the complex dependencies that span cloud services; the inevitable configuration drift that occurs in network and identity realms; and the malware or misconfigurations that hide within containers, serverless functions, or application libraries. Without *golden copies* (fully scanned, multi-component, point-in-time, clean application and data copies), restores are built on sand.

The old approach of "Let's hope the backup works" is as risky as betting your business on a single, untested parachute.

The Crucial Missing Piece: Regular, Comprehensive Rebuild Testing

Both defense-based cybersecurity and traditional data backup strategies remain fundamentally incomplete. The solution lies in regular, real-world testing of complete rebuild capabilities.

Rebuild testing represents a fundamental shift from hoping that backups work to testing their effectiveness with comprehensive validation. This approach reconstructs the entire digital environment exactly as it existed at a known-clean moment, providing holistic environment recovery that goes far beyond simple data restoration.

The process involves rewinding every layer of the organization's infrastructure—not just its data but also its network configurations, compute resources, identity frameworks, containers, serverless configurations, and API gateways—so that the restored environment mirrors the original.

Most critically, Rebuild testing incorporates comprehensive scanning for malware, vulnerabilities, configuration drift, and unauthorized alterations that may have infiltrated the environment before the

snapshot was taken. This validation step transforms backup restoration from a leap of faith into a verified, secure recovery process that organizations can trust when their survival depends on it.

The NIST Cybersecurity Framework's Blind Spot

The National Institute of Standards and Technology's (NIST's) Cybersecurity Framework provides an elegant structure across six core functions: Identify, Protect, Detect, Respond, Recover, and Govern. However, its Recover function is dangerously misunderstood, creating a critical blind spot that leaves organizations vulnerable even when they believe they're protected.

Recover versus Rebuild: Two distinct functions

NIST's Recover function focuses on restoration to get systems back to a functional state after an incident. This approach treats recovery as damage control, emphasizing speed over validation. Organizations restore from backups, execute DR scripts, and celebrate when applications appear to be running, often without verifying the integrity or completeness of what they've restored.

Rebuild, by contrast, represents a paradigm shift toward reconstruction with confidence. Rather than simply restoring what was lost, the Rebuild function is designed to create a verified, clean application environment from known-good components. It's the difference between patching a damaged wall and constructing a new one from trusted blueprints. Both walls may appear functional, but only one can maintain structural integrity.

Where Recover falls short

In practice, Recover is often no more than a series of cursory compliance activities that provide minimal assurance of actual recovery capability. Organizations conduct annual tabletop DR exercises that test procedures on paper but never validate actual system restoration. They perform sporadic database restores and periodically test virtual machine (VM) spins that touch only fragments of their infrastructure, while ignoring the complex interdependencies that modern applications require.

Most critically, organizations rarely perform full-scale, full-application restores with regular validation to verify that all components match snapshot data and function together as an integrated

system. Table 2-1 identifies where NIST's current framework falls short of modern requirements and shows how Rebuild addresses these critical gaps.

By elevating Rebuild to its own pillar—a living extension of Recover—we focus defenders on a more stringent test of resilience: the ability to go back to a specific point in time, select a safe golden copy, and rebuild the application from the network to the data layer on demand.

Table 2-1. Gaps in NIST's current framework and how the Rebuild function fills the gap

Function	Traditional strength	Current gap	How Rebuild fills the gap
Identify	Asset inventories, BIAs	No regular recovery confidence	Historical catalog of golden copies
Protect	IAM, encryption, firewalls	Cannot stop zero-day or insider attacks	Immutable snapshots for Rebuild artifacts
Detect	SIEM, XDR, UEBA	Alert ≠ assured restoration	Automated Rebuild tests triggered by event feeds
Respond	IR playbooks, quarantine	Playbooks rarely validate full restore	Integrated Rebuild runs as part of response
Recover	Backup and failover scripts	Partial, manual restores; untested runbooks	Code-defined orchestration of full-environment Rebuild
Rebuild			*Regular, automated point-in-time rebuild drills*

The Extended Framework

To maintain resilience, we must add the Rebuild function to the NIST framework. Recover remains the policy, the plan, and the postmortem—the strategic framework that defines what should happen when disaster strikes. Rebuild is the living engine that can make recovery a proven reality rather than a theoretical possibility, as shown in Figure 2-1.

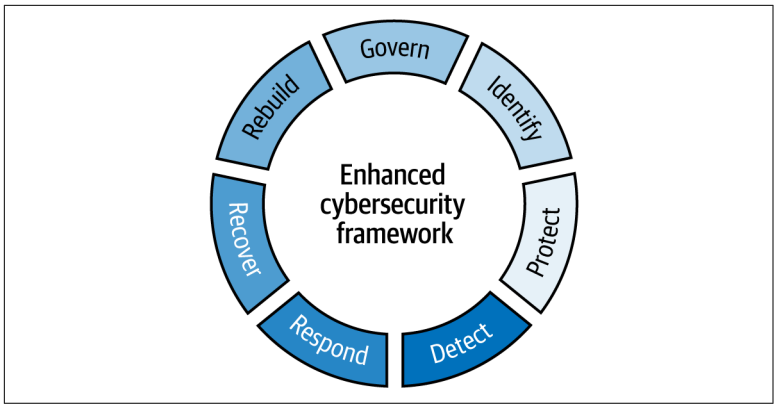

Figure 2-1. The NIST framework with the addition of the Rebuild function

The Rebuild function has a number of elements:

Point-in-time infrastructure recovery
: This captures comprehensive point-in-time snapshots of the infrastructure components, including network configurations, compute resources, application images, and their interdependencies. Unlike traditional backups that focus on data, these snapshots are designed to re-create the infrastructure context that applications need to function properly.

Golden copies
: These are point-in-time images that have undergone comprehensive scanning for malware, misconfigurations, and security vulnerabilities and thus provide validated, clean recovery points. These aren't just data copies—they're verified-clean snapshots of entire application stacks that you can trust implicitly, eliminating the fear that restoration might reintroduce the very problems you're trying to escape.

Recovery as code
: RaC transforms ad hoc restoration procedures into automated, versioned, and repeatable processes. Instead of relying on outdated runbooks, RaC treats rebuild procedures as living software that evolves with your infrastructure so that recovery capabilities improve over time rather than degrading through neglect.

Rebuild transforms "I hope the backup works" into "I have confidence this rebuild will work." That's because you've tested it, refined it, and proven it works dozens of times before you actually need it.

CHAPTER 3

Unleashing the Rebuild Advantage: Testing the Unexpected in the Cloud

In the early 20th century, automobile safety improved dramatically, not merely because cars became sturdier but because crash tests became central to the vehicle development process, transforming safety from a mere hope into an assured outcome. Similarly, in the early days of aviation, pilots believed that mastering flight meant building stronger planes.

Yet, when disaster struck midair, it was not the strength of the aircraft alone that saved lives but rather the pilot's ability to recover from unexpected events. No matter how well engineered the plane, survival often hinged on thorough, rigorous training through simulations of every conceivable emergency.

Today, digital resilience faces a similar transformative moment. Organizations must move beyond assuming that backups and cybersecurity defenses will work. A missing element, and perhaps the most critical element, is regular testing. Regular testing enables organizations to reliably rebuild their digital environments after catastrophic failures. This chapter will focus on how organizations can implement continuous rebuild testing.

Real Stories About Backups That Didn't Work and Recovery Plans That Failed

Recent incidents reveal the fragility of our digital ecosystems despite the considerable investments that organizations have made in cybersecurity and data backups. The scale and intensity of ransomware attacks have surged dramatically in just the past few years, proving that even comprehensive disaster recovery plans fail when they haven't been tested against the realities of modern cyberthreats.

Consider how major enterprises have suffered in the last few years when their recovery plans, despite meticulous documentation and significant investment, failed spectacularly.

MGM Resorts and Caesars Entertainment
: These companies faced debilitating attacks in late 2023, causing extensive disruptions. Despite their comprehensive DR plans, both companies struggled to restore critical business functions promptly. Robust documentation existed, but actual restoration failed due to untested dependencies, outdated configurations, and missing integrations between data backups and application recovery.

National Health Service (NHS), London, UK
: A 2024 attack with Qilin ransomware exposed a stark reality: nearly a million patient records and critical health data systems were compromised, despite robust backups. The NHS had to learn the hard way that restoring databases alone was insufficient; without verified restoration procedures for applications, identities, and network architectures, the backups proved futile.

AWS S3 Bucket Attacks
: *In this sobering 2025* scenario, Codefinger ransomware targeted the company's cloud storage buckets, rendering traditional cloud backup strategies ineffective. Many organizations, despite regular cloud backups, have woken up to discover that these backups are locked behind ransomware encryption. Such incidents highlight the critical need for an entirely new testing approach to restore procedures.

These stories illustrate a troubling reality: our traditional assumptions around data protection and cybersecurity are incomplete. The exponential rise in sophisticated ransomware, coupled with

increasingly prevalent cloud-region outages, demands a new discipline in digital resilience: testing full-environment rebuilds regularly, comprehensively, and rigorously.

Beyond Traditional Recovery: The Modern Challenge of Cloud Rebuilds

Rebuilding cloud environments involves much more than simply restoring data from backups. Modern applications operate in highly dynamic ecosystems, composed of numerous cloud-native services, that are constantly changing through multiple, simultaneous DevOps pipelines. Now these same pipelines are becoming increasingly AI driven.

Hidden Dependencies and Configuration Drift

Each DevOps pipeline can independently update configurations, deploy microservices, and adjust security policies, exacerbating the risk of configuration drift and obscuring critical dependencies. Teams that use continuous integration and continuous deployment (CI/CD) tools like AWS CodeDeploy CodePipeline, and CodeBuild frequently alter environments without clear visibility into the impact of their changes. Thus, they often create hidden vulnerabilities or overlooked dependencies, making a complete and accurate rebuild difficult in a crisis.

The Minimum Viability Strategic Framework: Making Rebuild Achievable

The practical reality facing most organizations is that rebuilding everything simultaneously is neither practical nor necessary. This is where the concept of *minimum viability* (or "minimum viable company") becomes the strategic bridge between the Rebuild theory and successful implementation.

The minimum viability framework recognizes that organizations need a keen understanding of their most critical assets and what it takes to restore them to operational status. Rather than attempting the overwhelming task of testing complete environment rebuilds, organizations can implement the Rebuild function incrementally by focusing on what truly matters for business survival.

Organizations can implement the Rebuild function effectively by tiering their applications and services using established frameworks like ISO 22301 (business continuity management systems) or NIST's business impact analysis guidelines. These frameworks help organizations categorize systems based on their operational criticality:

- *Mission critical:* Systems you can't do anything without (e.g., Active Directory, order management system, patient care systems). These applications form the foundation of minimum viability—without them, the organization cannot function at all.
- *Business critical:* Systems needed for the full recovery of operations (e.g., email, accounting, supply chain management). These enable expanded operational capacity beyond basic survival.
- *Non-critical:* All other systems that support full functionality but aren't essential for immediate business continuity.

This tiered approach transforms the Rebuild function from an overwhelming "everything must work" challenge into a strategic, phased recovery process. Organizations can achieve minimum viability by focusing Rebuild testing first on business-critical systems, then expanding systematically to mission-critical and non-critical applications.

This approach dramatically reduces initial recovery time objectives (RTOs) for essential business functions while maintaining the goal of complete environment restoration.

The Complete Rebuild: Metadata, Automation, and Orchestration

Effective rebuilding involves capturing all relevant metadata—not just the application data but also detailed configurations, resource dependencies, identity and access management (IAM) policies, networking topologies, and API endpoints.

It's crucial that organizations replicate this comprehensive metadata securely and immutably across multiple cloud regions or isolated accounts to minimize single points of failure and enhance security against ransomware and unauthorized changes.

Rebuilding must leverage automated *infrastructure as code (IaC) techniques*, integrating previously fragmented recovery processes

into executable automation pipelines. This approach enables resilience operations to remain adaptable, consistent, and verifiable.

Centralizing and continuously updating the recovery code allows teams to manage resilience proactively rather than respond to crises reactively. Therefore, a comprehensive rebuild process means not just restoring data but also orchestrating the entire cloud environment seamlessly and consistently.

Operationalizing the Rebuild Function: Point-in-Time Infrastructure Recovery and Recovery as Code

Commvault's Cloud Rewind (formerly Appranix) addresses critical weaknesses in traditional DR practices with two innovative concepts for the on-demand rebuilding of application environments: point-in-time infrastructure recovery (PITR) and RaC.

Historically, organizations struggled with fragmented recovery runbooks, with each runbook managed independently by security, application, enterprise architecture, and backup teams. In a crisis, business continuity teams incurred significant delays as they pieced together these scattered recovery documents, leading to prolonged downtime.

Point-in-time infrastructure recovery resolves this by providing an automated, comprehensive snapshot of an entire digital ecosystem—not merely the data but the complete application stack, microservices, serverless functions, identity and access configurations, network topologies, and their dependencies.

By regularly capturing these complete point-in-time states, based on the given policies, PITR (Figure 3-1) enables organizations to maintain validated, comprehensive golden copies that are readily available for both minimum viability restoration and full application stack rebuilds.

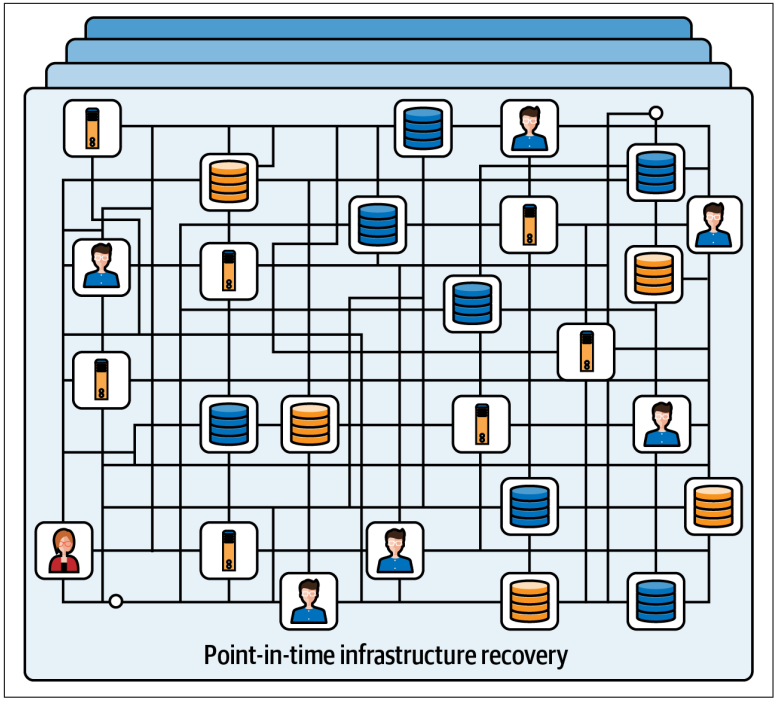

Figure 3-1. Illustration of how PITR provides a comprehensive snapshot of an entire cloud application ecosystem

The ability of PITR to capture tiered snapshots becomes particularly valuable for minimum viability strategies. Organizations can configure policies that prioritize business-critical applications for more frequent snapshots, provide mission-critical systems with verified recovery points, and maintain baseline protection for non-critical applications.

This tiered approach enables rapid restoration of minimum viable operations while maintaining comprehensive protection across the entire environment.

Leveraging hyperscale cloud platforms like AWS, Azure, or Google Cloud enhances this capability further. The vast, on-demand compute resources and built-in isolation features of hyperscale clouds allow organizations to execute frequent large-scale tests with ease and efficiency. These platforms simplify complex recovery validations, turning expensive and infrequent disaster recovery exercises into routine, cost-effective rebuild exercises.

Complementing PITR, *recovery as code* transforms recovery into a unified, automated process that operationalizes the Rebuild function. Rather than maintaining separate, cumbersome runbooks, RaC embeds all necessary recovery steps directly into executable automation pipelines (Figure 3-2).

Version-controlled code serves as a single source of truth for recovery, aligning security teams, architects, application developers, and backup specialists around a centralized, consistent process. This code-driven approach integrates regular rebuild testing seamlessly into daily DevOps workflows, significantly reducing operational complexity.

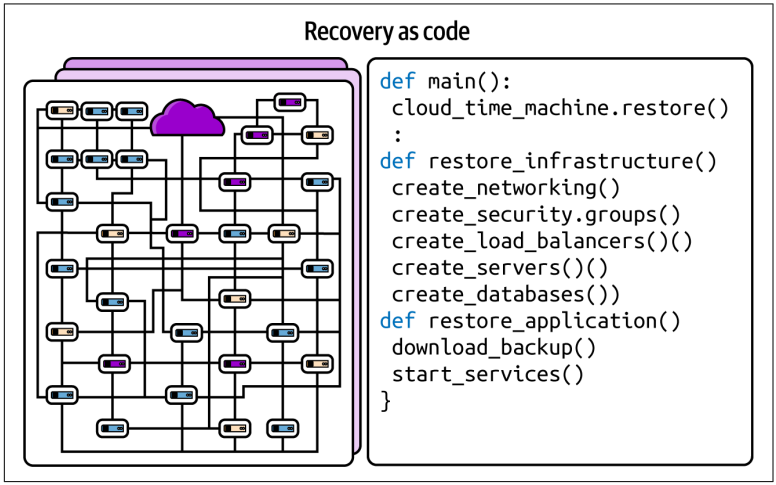

Figure 3-2. PITR hand in hand with RaC, which embeds all recovery steps in executable automation pipelines

Using Rebuild Strategically

RaC can be structured to support minimum viability workflows, with separate automation pipelines for business-critical, mission-critical, and non-critical application tiers. This enables organizations to execute rapid minimum viability restoration while simultaneously preparing for full environment recovery, making the Rebuild function both strategic and practical.

Together, PITR and RaC fundamentally change organizational resilience. These innovations empower organizations to move from uncertainty and reactive crisis management to a proactive, demonstrable capability, significantly reducing recovery times, simplifying

compliance requirements, and building unmatched trust with stakeholders. Regular rebuild testing thus becomes not only achievable but also a strategic imperative for modern digital resilience.

The Business Value of Regular Rebuild Testing with Cost Optimization

Adopting regular Rebuild testing with PITR and RaC fundamentally transforms organizational resilience by making the Rebuild function both economically viable and strategically valuable. Rather than relying on uncertain recovery plans and fear-driven crisis management, organizations can achieve clear, measurable, and demonstrable recovery capabilities.

Consider a healthcare organization facing a ransomware attack. Rather than waiting for complete infrastructure restoration, minimum viability planning helps enable rapid restoration of patient care systems, emergency department operations, and critical communication capabilities. The organization can restore its administrative systems, scheduling platforms, and reporting tools later, without impacting patient care.

This tiered approach delivers a number of key measurable business benefits.

Enhanced Operations and Customer Trust

By prioritizing revenue-generating systems in minimum viability planning, organizations can resume core business operations quickly while maintaining customer confidence. This trust becomes a competitive advantage, particularly in industries where digital reliability directly impacts customer relationships.

Regulatory and Compliance Benefits

Minimum viability planning helps organizations meet regulatory requirements for maintaining essential services during disruptions. Automated rebuild testing simplifies compliance processes by regularly generating comprehensive evidence, making audit preparations more efficient.

Organizations can demonstrate adherence to regulatory standards like the Health Insurance Portability and Accountability Act

(HIPAA) and industry frameworks like SOC 2, ISO 27001, and the Digital Operational Resilience Act (DORA) with minimal manual effort, providing auditors immediate insights into resilience capabilities.

Cost and Time Reduction

Automating rebuilds also substantially reduces complexity and cost. Traditional disaster recovery tests are expensive, disruptive, and prone to human error. By adopting RaC, organizations can effectively rebuild and test on demand using hyperscale cloud programmability and deployment models.

They can automate intricate Rebuild processes, transforming scattered manual runbooks into streamlined, repeatable code-based procedures. This automation not only reduces operational overhead and eliminates manual testing expenses but also promotes consistency and reliability across every test.

Organizational Confidence and Trust

Perhaps most significantly, regular Rebuild testing engenders unparalleled organizational confidence and trust. Regularly validated recovery capabilities offer leadership teams clear, demonstrable assurance of their readiness to deal with disruptions.

Regulatory bodies, customers, and partners gain reassurance that the organization is proactively mitigating risks and can swiftly recover from any cyberattack or cloud failure. This trust becomes a strategic advantage, setting resilient organizations apart in a world increasingly defined by digital threats and disruptions.

Making Rebuild Testing Practical: Infrastructure and Validation Methods

Having established the strategic framework for minimum viability, the question becomes: How do organizations actually execute regular rebuild testing at scale for critical applications? The answer lies in two enablers that make frequent testing both affordable and realistic:

- First, *hyperscale cloud platforms* provide the elastic infrastructure needed to spin up complete testing environments on demand.
- Second, *chaos engineering principles* help these tests simulate real-world failures rather than predictable scenarios.

Together, these approaches transform Rebuild testing from a costly annual exercise into a routine operational capability.

Using Cloud Platforms as a Powerhouse for Testing

Cloud platforms such as AWS, Azure, and Google Cloud offer an ideal environment for regular rebuild testing due to their unmatched flexibility, scalability, and affordability. Unlike traditional data centers, hyperscale clouds allow organizations to spin up fully isolated sandbox environments instantly, execute rigorous testing, and tear them down without affecting production.

This flexibility removes the cumbersome barriers traditionally associated with disaster recovery testing, enabling more frequent, thorough, and meaningful tests.

A significant advantage of these platforms is their vast on-demand compute and storage resources, which allow organizations to quickly scale resources up or down according to specific testing needs. The strategic use of spot instances provides space compute capacity at drastically reduced costs—often 70% to 80% less than typical on-demand prices—allowing for increased testing frequency without additional budget constraints.

The cloud environment also enables simulation of partial infrastructure failures, testing of cross-regional failover capabilities for business-critical systems, and validation that minimum viability restoration procedures work under various failure conditions—all within cost-effective, isolated testing environments.

Regular rebuild testing, once seen as complex and costly, is now practical and accessible, enabling what used to be an occasional compliance activity to become a core organizational capability.

Chaos Testing: Building Resilience Through Intentional Failure

Chaos testing is the practice of deliberately injecting controlled disruptions into a system to uncover hidden weaknesses and validate resilience under realistic conditions. This approach intentionally introduces failure scenarios, such as infrastructure outages, network latency, or unexpected resource spikes, to test whether systems continue to function reliably under stress.

Unlike standard DR drills, which often simulate predictable scenarios, chaos testing thrives on unpredictability, continuously challenging assumptions and revealing blind spots in application and infrastructure resilience.

In the context of rebuild tests, chaos testing becomes particularly critical. Because production environments evolve constantly, new services are deployed, configurations change, and workloads fluctuate. Traditional static rebuild tests quickly become outdated.

Cloud platforms enable organizations to run adaptive rebuild tests that reflect the dynamic nature of modern applications. By integrating chaos engineering principles with rebuild practices, tests evolve proactively, mirroring production complexity and continuously adjusting to changes.

Chaos testing becomes even more valuable when applied through the lens of minimum viability. Rather than testing random failures across entire environments, organizations can focus chaos experiments on business-critical systems to understand how failures could cascade and impact minimum viability restoration.

For example, during a minimum viability restoration exercise, chaos testing might deliberately disable Active Directory services to validate that backup authentication systems can maintain essential business operations. Or it might simulate network partition failures between critical microservices to confirm that minimum viability applications remain functional even when dependent services are unavailable.

Making Regular Rebuild Tests Happen: From Theory to Concrete Results

Establishing regular rebuild testing transforms organizational resilience from theoretical reassurance into concrete, measurable capability. Just as software development now incorporates regular quality assurance (QA) practices, digital resilience must similarly adopt regular rebuild testing. CIOs and CISOs can expect tangible outcomes: demonstrable recovery capability; measurable risk reductions; and clear alignment among security, cloud operations, and recovery teams.

Implementing Tiered Testing Approaches

A structured approach begins by scheduling monthly or quarterly rebuild test days that incorporate both minimum viability and full environment restoration scenarios. These events must be carefully planned and include regular recovery drills, minimum viability exercises, and controlled chaos experiments.

Minimum viability test days

Organizations should conduct separate exercises focused specifically on restoring business-critical systems within defined time windows. These tests validate that essential business functions can be restored quickly, typically within hours rather than days.

Success metrics for these tests include time to restore identity services, time to bring core business applications online, and verification that minimum viable operations can be sustained while full recovery continues.

Full environment rebuild testing

It's also important to conduct broader exercises that test complete infrastructure restoration, confirming that mission-critical and non-critical systems can be successfully restored after minimum viability is achieved. These tests validate the organization's ability to return to full operational capacity.

Defining Roles and Responsibilities

Clearly defining roles and responsibilities becomes crucial for effective testing, particularly when balancing minimum viability priorities and full environment restoration.

Security teams

Security teams validate that recovered environments are rigorously scanned and cleared of vulnerabilities, ransomware signatures, and misconfigurations. During minimum viability restoration, these teams prioritize the validation of business-critical systems while conducting comprehensive security assessments of the broader environment in parallel.

Cloud operations and applications teams

These teams focus on infrastructure provisioning, configuration alignment, and orchestrating comprehensive rebuilds from the environment PTIRs. They manage the technical execution of both minimum viability restoration and full environment rebuilds, confirming that infrastructure dependencies are properly sequenced and that restored services meet functional requirements.

Recovery teams

Overseeing the entire rebuild process, recovery teams provide coordination across all groups and accurate documentation of results. They manage the transition from minimum viability to full operational capacity and coordinate between different testing scenarios.

Measuring Success and Demonstrating Value

To assess the effectiveness and demonstrate the value of rebuild testing, clear metrics must be established and communicated. These metrics go beyond simple recovery time objectives (RTO) and recovery point objectives (RPO). They include:

Minimum viability metrics
: Time to restore entire business-critical systems and associated dependencies, success rate of minimum viability procedures under stress conditions, and ability to sustain essential operations during full recovery

Comprehensive recovery metrics
: Overall recovery time for complete environments, success rate of full rebuild procedures, and verification that all systems function correctly after restoration

Resilience testing metrics
: Frequency and thoroughness of the chaos tests conducted, the number of vulnerabilities or misconfigurations identified and resolved through testing, and improvement in recovery performance over time

Business impact metrics
: Reduction in potential revenue loss during cyber incidents, improvement in customer trust metrics, and demonstration of regulatory compliance through regular testing evidence

CIOs and CISOs should expect regular reports highlighting these metrics, and these metrics will provide transparency and concrete evidence of improvement over time.

Regular rebuild testing provides concrete evidence of recovery capabilities to all stakeholders—from security teams that wish to validate threat mitigation to executives who seek to demonstrate operational resilience.

Long a best practice, regular Rebuild testing becomes a foundational element of enterprise risk management. By integrating minimum viability principles with comprehensive testing, organizations can confidently demonstrate their ability to restore essential business functions rapidly while maintaining complete digital resilience. This capability has become essential for organizational survival as cyberthreats continue to evolve and intensify.

CHAPTER 4

Your Key to Victory: Implementing the Rebuild Function for True Cyber Resilience

True cyber resilience hinges on one critical capability that extends beyond traditional cybersecurity frameworks. As we established in Chapter 2, the NIST Cybersecurity Framework requires a seventh function: Rebuild. This function transforms uncertainty into confidence, making the "Sure hope the backup works" approach obsolete by proving recovery capabilities through regular testing.

The minimum viability approach outlined in Chapter 3 makes this transformation achievable. Rather than attempting to test everything simultaneously, organizations can implement Rebuild systematically, starting with business-critical systems and expanding to complete environments. This strategic framework turns an overwhelming challenge into a manageable, step-by-step process that delivers immediate value while building toward comprehensive resilience.

The Agentic Threat: Why Speed of Attacks Changes Recovery Requirements

Experts predict that we could be living in a world of agentic attackers as soon as this year, with AI agents representing an attractive prospect to cybercriminals because they're much cheaper than

hiring professional hackers and can orchestrate attacks more quickly and at far larger scale than humans.

Agentic ransomware represents a collection of AI bots that perform all the steps needed for successful ransomware attacks, but faster and better than human operators. These systems don't just accelerate existing attack methods—they fundamentally change the game by operating at machine speed with machine learning capabilities.

The implications for recovery are profound. In nearly one in five cases, data exfiltration now takes place within the first hour of compromise. Traditional backup and recovery approaches, designed for human-speed threats that provided days or weeks of warning, become obsolete when attacks move from reconnaissance to encryption in minutes.

This is why the fast, automated Rebuild function becomes even more critical in the agentic era. Only through regular, automated testing can organizations stay ahead of adversaries that learn and adapt at superhuman speed.

The Strategic Path Forward: From Minimum Viability to Complete Resilience

The journey toward Rebuild confidence follows the minimum viability framework that makes comprehensive testing achievable. Organizations begin by identifying their business-critical systems, those essential for minimum viability during a crisis. They then implement PITR snapshots and RaC automation for these priority systems first.

Success with minimum viability creates a foundation for expansion. Organizations can then extend Rebuild capabilities to mission-critical systems like accounting and supply chain management, followed by non-critical applications. Each expansion builds on proven processes and growing organizational expertise, and ultimately comprehensive Rebuild capability is achieved across the entire digital ecosystem.

The hyperscale cloud platforms discussed in Chapter 3 make this progression economically viable. Spot instances reduce testing costs by up to 90% compared to on-demand pricing, according to the

official AWS and Azure documentation, while elastic infrastructure enables frequent validation without impacting production systems.

Chaos testing confirms that testing scenarios reflect real-world failure scenarios, building genuine confidence in recovery capabilities.

Overcoming Implementation Challenges

The minimum viability approach systematically addresses common organizational roadblocks to Rebuild implementation. Indeed, when leaders express concern about costs, the business case becomes compelling. Rapid restoration of revenue-generating systems pays for itself with the first incident it helps prevent.

Consider the real-world examples from Chapter 3. MGM Resorts and Caesars Entertainment had comprehensive DR plans, yet both organizations struggled with restoration because they lacked tested Rebuild capabilities. The minimum viability approach could have facilitated the rapid restoration of core gaming and hospitality operations while complete recovery proceeded in parallel, minimizing business disruption and customer impact.

Similarly, NHS London's experience with Qilin ransomware demonstrated that robust backups mean nothing without verified restoration procedures for applications, identities, and network architectures. A minimum viability strategy would have prioritized patient care system restoration, allowing critical healthcare operations to continue while comprehensive rebuilding addressed administrative and support systems.

Budget and Resource Constraints

RaC automation collapses fragmented runbooks into unified, version-controlled pipelines. Security, cloud operations, application, and recovery teams collaborate on the same code base rather than maintaining separate, siloed documentation. This consolidation eliminates the need for separate drills across teams while providing consistency and reducing manual effort.

Return on investment becomes clear when organizations measure minimum viability restoration times. Reducing RTO from 48 hours to 2 hours for business-critical systems produces immediate value. Automating evidence collection for SOC 2, ISO 27001, and DORA

compliance reduces audit overhead while providing ongoing validation of recovery capabilities.

Leadership Buy-in and Organizational Alignment

Nothing builds leadership support like demonstrable capability. Organizations can show live rebuild metrics through dashboards, place audit reports at executives' fingertips, and demonstrate recovery completion in under an hour rather than in days or weeks.

The minimum viability approach makes the business impact immediately visible. When leadership sees that core revenue-generating systems can be restored rapidly and reliably, funding and organizational support follow naturally. Each successful minimum viability test builds confidence for broader Rebuild implementation.

Looking Ahead: When Rebuild Becomes Standard Practice

The future of cyber resilience is already emerging, driven by the reality of agentic threats. As agentic AI becomes more capable, security teams will delegate more tasks to autonomous agents with minimal instructions, allowing systems and networks to keep up with constantly evolving threat tactics. In five years, the Rebuild function will be as fundamental to cybersecurity as the current NIST framework functions. Minimum viability planning will be standard practice, with organizations maintaining tested recovery procedures for business-critical systems as rigorously as they maintain financial controls.

This transformation will reshape how organizations approach digital resilience:

- *AI-enhanced Rebuild capabilities:* Future Rebuild systems will leverage AI to automatically identify configuration drift, predict potential failure scenarios, and conduct systematic chaos testing that anticipates novel attack vectors before they're deployed. These agentic recovery systems will work alongside human operators to autonomously take on routine tasks, augment human decision making, and automate workflows.
- *Speed-matched response:* As attacks accelerate, Rebuild capabilities must accelerate correspondingly. Organizations will

implement AI-powered recovery systems designed to execute complete environment restoration faster than agentic attackers can adapt their strategies.

- *Competitive advantage:* Organizations that can demonstrate rapid, reliable recovery will gain significant competitive advantages. Customers and partners will prefer vendors that can demonstrate reliable service continuity. In regulated industries, proven Rebuild capabilities will become a requirement for maintaining licenses and certifications.

The organizations that embrace the Rebuild function today—starting with minimum viability and scaling to comprehensive coverage—will become those that not only survive tomorrow's attacks but emerge stronger from them. They will transform cyber incidents from business-threatening disasters into manageable operational challenges.

Your Path Forward: From Hope to Confidence

The choice facing every organization is clear: continue to hope that traditional backup and recovery methods will suffice against modern threats or begin to build a demonstrated Rebuild capability that provides genuine confidence.

The minimum viability approach makes this choice actionable. Begin with a single business-critical system. Implement PTIR snapshots and RaC automation. Conduct chaos testing to validate restoration under stress. Measure and demonstrate the results.

Success with one system builds the foundation for systematic expansion. Each additional system benefits from growing organizational expertise, proven processes, and established automation. The journey from backup hope to rebuilding confidence accelerates as capabilities mature.

Test your minimum viability first, then expand systematically. Recover with confidence through proven Rebuild capabilities. Thrive by making resilience a competitive advantage.

The cyberthreats of tomorrow will be more sophisticated, more persistent, and more devastating than today's attacks. Organizations that wait for perfect solutions or ideal conditions will find

themselves unprepared when survival depends on recovery speed and reliability.

Your journey starts with understanding your business-critical systems and implementing the minimum viability framework. The technology exists. The methodologies are proven. The only question is whether you'll begin now or wait until the next attack forces your hand.

Choose confidence. Choose Rebuild. Choose to thrive amid uncertainty.

About the Author

Govind Rangasamy is the founder and CEO of Appranix, now part of Commvault, and a Forbes Technology Council author. A serial entrepreneur with extensive experience in enterprise cloud management, Govind founded Appranix to revolutionize infrastructure-centric resilience models, which he believes are inadequate for today's distributed, dynamic cloud applications. He regularly contributes to Forbes and is a frequent podcast guest and conference speaker on cloud resilience.