rubrik

THE FOUR LEVELS OF CLOUD CYBER RESILIENCE: AN IT LEADER'S GUIDE

AN INTRODUCTION FOR THE MODERN IT & SECURITY LEADER

If you're an IT or security leader, you've likely recognized that the traditional playbook for data protection wasn't designed for the realities of the cloud. The old strategies were built for a world with physical boundaries—recovering from predictable events like fires and hardware failures. Today, however, the threat isn't a faulty server. It's a sophisticated adversary who understands your cloud environment and is actively working to dismantle your ability to recover. And in the era of AI, this threat is only growing larger.

Recent threat intelligence paints a stark picture. The H2 2025 Google Cloud Threat Horizons Report notes that threat actors are now refining their tactics to specifically target "cloud-native backup and recovery for modern cyber threats." Preventing a victim's ability to recover isn't an accidental byproduct of an attack; it's a primary objective.

Microsoft's research on threat actors like Storm-0501 shows exactly how this plays out. These groups have evolved, moving beyond simple on-premises attacks to pivot directly into cloud environments. They compromise the identity fabric, escalate privileges, and then methodically execute their attack, which includes the calculated destruction of your recovery capabilities.

In an environment where your infrastructure is in the cloud, a single compromised credential can unravel everything. The challenge has shifted from an operational task—restoring a lost file—to a question of business survival. Your board isn't asking if your backups completed successfully. They're asking, "If our cloud accounts are compromised, can the business survive? And how fast can we be back online?"

Answering that question requires a solution built to withstand an intelligent adversary, not just a simple mistake. It demands a new level of honesty about your true recovery posture.

This guide is a tool for that self-assessment. Find out where you stand, and more importantly, how to get to the next level.

CLOUD RESILIENCE MATURITY SCALE

Increasing cloud resiliency level

LEVEL 1

THE UPTIME OPTIMIST

LEVEL 2

THE CONVENTIONAL DEFENDER

LEVEL 3

THE DURABLE VAULT

LEVEL 4

THE RESILIENT ENTERPRISE

THE UPTIME OPTIMIST

This is the default state for many organizations moving to the cloud, defined by legacy assumptions and a dangerous misunderstanding of modern threats.

The Mindset: Complacency Through Durability

"The cloud is a fortress. With 'eleven nines' of durability from providers like AWS and Azure, our infrastructure is safer than it ever was on-premises. Our job is to protect against simple user error—an accidental deletion or a bad deployment. We have basic snapshots for that."

This mindset fundamentally confuses infrastructure durability (preventing data loss from hardware failure) with cyber resilience (surviving a malicious attack), leading to a dangerous sense of complacency.

In Practice: The Technical Reality

The team relies on a patchwork of manual processes and basic, homegrown automation. This patchwork extends beyond just scripts to include separate documentation, processes, and monitoring for each component, creating critical security blind spots.

A system administrator might take manual snapshots before a major change, or a simple PowerShell/Bash script runs nightly to snapshot key VMs. These scripts are often brittle, lack error handling, and run with overly permissive credentials. For object storage, they've checked the "enable versioning" box as a quick way to ensure they have multiple versions of a file—just in case.

These safeguards may come in handy to address operational issues, but they do little to protect data against malicious attack. Critically, all these copies and scripts operate within the same production cloud account, using the same set of admin credentials and making them vulnerable to anyone who has access to those credentials.

The Critical Flaw: A Single Point of Failure

The entire model is built to withstand an accident, not an adversary. The fatal flaw is the **single security domain**. An attacker who compromises admin credentials doesn't care about infrastructure durability; they care about the data. And when it's all in one place, they see one unified target. They can disable versioning, delete all object versions, wipe out every snapshot, and destroy the production environment using the same set of stolen keys. Furthermore, the system lacks any intelligence about the health or recoverability of the data itself. The organization can't confidently answer: "Is this backup even usable?"

THE BLUEPRINT FOR MATURITY: HOW TO LEVEL UP

Own the Risk

Make the mental shift from "the cloud provider protects it" to "we protect our data in the cloud." Acknowledge that the threat is now a credentialed attacker, not a failing disk.

Centralize Control

Ditch the ad-hoc scripts and manual clicks. Consolidate all data protection under a single, policy-driven platform to gain visibility and consistency.

Architect for Attack

Ask the critical question that defines the next level: "If our primary cloud admin credentials are stolen, can our backups survive?"

THE CONVENTIONAL DEFENDER

An organization at Level 2 has professionalized its backup operations. They have a centralized tool and can pass an IT audit, but they've solved the operational challenge, not the security one.

The Mindset: Confidence Through Compliance

"We have a formal, centralized backup system with defined RPOs and RTOs that meets our compliance requirements. All our jobs are monitored, and we can prove to the auditors that we have 30 days of retention. We are following industry best practices."

This mindset conflates operational tidiness with security. The focus is on successfully creating backups, not on ensuring their survivability during a security breach.

In Practice: The Technical Reality

The team uses a centralized tool, either a native service like AWS Backup or a third-party platform. Policies are applied consistently, and basic security hygiene is typically in place, with MFA and SSO enforced on cloud accounts managing the backup services.

However, the real vulnerability lies beneath the surface: While there may be a corporate-level team managing what gets backed up, they often lack full visibility into everything that exists across their environment.

Each business unit or application team makes their own decisions about backup priorities. Because each business unit is measured on how well it performs its main function, not how well it backs up its data, backup decisions become driven by cost concerns rather than business risk.

The Critical Flaw: Credential Compromise and Lack of Immutability

This model fails for several reasons. First, without complete discovery and data classification across business units, the corporate-level backup team cannot confidently ensure comprehensive backup coverage across the organization.

And because each business unit is backing up its data without clear, over-arching guidance, they're left to make their own decisions on what constitutes critical data and systems. This lack of guidance combined with a focus on other business priorities means backups may not get the attention they deserve. Since backups are not free, they are often one of the first line items teams choose to minimize or skip entirely.

This decentralized, cost-first approach creates dangerous protection gaps that remain invisible until it is too late.

The model also doesn't do anything to address the fact that an attacker with compromised admin credentials can pivot from the production environment directly to the backup system. Since cloud native backups typically reside in the same cloud account as the protected data, compromising the admin account gives attackers access to both.

Crucially, the backup data itself still isn't truly indelible. While snapshots are immutable by design (their contents cannot be modified once taken), they can still be deleted by a privileged user. The same holds true for versioned objects in cloud storage: The versions themselves are immutable, but an attacker with sufficient privileges can simply delete all versions.

An attacker doesn't even need to delete backups manually. They can use stolen credentials to log into the backup console, change retention policies from 30 days to just one day, and then let the system gracefully delete itself over time. To security monitoring tools, these actions look like normal policy management, not an attack, so they don't trigger an alert. By the time the organization discovers the compromise and attempts recovery, their backup history has been systematically erased through what appears to be routine operational changes.

This distinction matters because immutability alone doesn't protect against a determined adversary who compromises admin credentials. They don't need to modify your backups or even directly delete them. They just need to change the rules that govern their lifecycle.

THE BLUEPRINT FOR MATURITY: HOW TO LEVEL UP

Establish Data Classification and Centralized Governance

Implement a corporate-wide data classification framework that identifies critical systems and workloads across all business units. Create a centralized backup governance team responsible for defining and enforcing protection standards based on business criticality, not individual team budgets.

Gain Enterprise-Wide Visibility

Deploy tooling and processes that provide a single pane of glass across all business units and cloud environments to identify protection gaps before an incident. Know definitively what is and isn't being backed up.

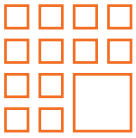

Isolate the Recovery Domain

Your backup data must live in a completely separate security context—a different cloud account with its own, separate identity management. This is the first step to creating a true logical air gap.

Demand Architectural Immutability and Indelibility

Your backup data must be truly immutable and indelible, meaning it cannot be prematurely deleted, modified, or encrypted, even by an administrator with root-level privileges, for a set period.

THE DURABLE VAULT

At this level, an organization can confidently say their backup data will survive a full compromise of their production environment. They have achieved data durability, but not true cyber resilience.

The Mindset: Data Durability Equals Resilience

"We assume our production environment will be fully compromised, so we've built a vault. Our backups are logically air-gapped, immutable and indelible, and use a separate set of credentials. An attacker can't touch our data. When an attack happens, our job is to begin the forensic process to find a clean recovery point in our secure vault."

This mindset correctly prioritizes data survivability but makes a dangerous assumption: that the ability to *preserve* data is the same as the ability to *restore the business*.

In Practice: The Technical Reality

The team has successfully implemented the blueprint from Level 2. Their backup system and data reside in a dedicated "bunker" cloud account with its own, separate identity management, enforced multi-factor authentication (MFA), and strict role-based access control (RBAC). Access is a "break-glass," highly-audited procedure.

The data is architecturally indelible. Not only are the snapshots immutable by design, the organization has also implemented retention locks and privileged access controls that prevent deletion or policy manipulation, even by an administrator with root-level privileges, for a defined period. Additionally, they've established centralized data classification and governance, ensuring critical systems across all business units are identified and protected according to corporate standards rather than individual team decisions.

All the right security boxes have been checked to ensure the data will survive.

The Critical Flaw: The Reactive Recovery Bottleneck

The model breaks down the moment an attack is discovered. This is because a durable vault, while safe, lacks sophisticated threat or data context. It knows that a backup was taken, but it doesn't know what's inside it. It has operational metadata but no security intelligence.

A critical misconception emerges here: Organizations assume their enterprise security tools can help with recovery, not just prevention. Endpoint detection and response platforms, cloud security posture management tools, and threat detection services are excellent at identifying malicious attackers in real-time across production environments. However, these proactive security tools cannot scan backup data directly. They monitor live systems, running processes, and active network traffic—not the static, point-in-time snapshots stored in your backup vault.

This creates a dangerous blind spot. Just as many organizations mistakenly believe cloud providers can recover deleted data, they also assume their security stack will accelerate recovery during an attack. The reality is that when ransomware strikes, those security tools have already done their job in the production environment—but they offer no insight into which backup snapshots are clean, when the attack began, or what the optimal recovery point is.

The core problem is one of context and time: When a complex incident occurs—whether it's a subtle, slow-and-low attack carried out via a compromised identity or a massive operational mistake—traditional forensic methods cannot distinguish between a malicious act and an operational error, leading to weeks of recovery time.

This absence of context forces the team into a frantic, manual, and error-prone investigation:

The Hunt for Clean Data

The recovery process begins with manually mounting and scanning terabytes of backup data, often involving the slow rehydration of full backups just to inspect file integrity.

Fragmented Forensics

The team must perform forensic analysis across dozens of disparate systems and attempt to reconcile fragmented metadata from siloed backup repositories to piece together what happened.

Security Tool Blindness

The organization's existing security infrastructure—designed to protect production workloads—provides zero visibility into the backup repository. So, teams are forced to start threat analysis from scratch in the backup data—without the proper tools to investigate that data—during the most critical moment.

This reactive process means that when the pressure is highest, you are unable to answer the most important questions:

- Which backup copies are clean and which contain dormant malware?
- When did the attack actually begin across your cloud environment?
- What's the optimal recovery sequence to minimize business disruption?
- Has sensitive data been compromised, and what are your regulatory obligations?

You have a vault full of safe data, but no way to quickly find the last-known good recovery point. And even if you could find the last-known good copy, the team still has to figure out what, in all that data, needs to be restored and the correct way to restore it. This process can take days or even weeks, all while the business is down.

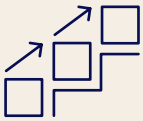

THE BLUEPRINT FOR MATURITY: HOW TO LEVEL UP

Automate Intelligence, Not Just Tasks

The goal is to have threat intelligence ready at a moment's notice. The platform should be able to instantly pinpoint the likely start of an attack and identify the blast radius—not just automate backups.

Shift from Reaction to Proaction

Don't wait for an attack to begin your analysis. Your data security platform must be able to continuously analyze your backups for indicators of compromise *before* a crisis.

Bridge the Security Visibility Gap

Recognize that your production security tools cannot see into backup data. Implement a data security platform with native threat detection capabilities that continuously scan backup snapshots for malware, encryption patterns, and indicators of compromise—providing the recovery intelligence your existing security stack cannot deliver.

Orchestrate Application Recovery

Evolve from restoring single servers to building automated, push-button runbooks that can recover entire, multi-tier application stacks in the correct sequence.

THE RESILIENT ENTERPRISE

This is the state of modern cyber resilience, where recovery is a core, tested, and automated business function. The organization doesn't just survive an attack. It recovers predictably and rapidly.

The Mindset: From Reactive to Ready

"Recovery is not a reactive event. It's a proactive capability. The moment of crisis is the wrong time to start doing forensics. We continuously analyze our data so that when an attack happens, we can instantly define the blast radius and identify clean recovery points. Our job is to execute on a plan that has already been created, not to create one in real time."

In Practice: The Technical Reality

At this level, organizations use an integrated data security platform that directly solves the reactive bottleneck of Level 3. The entire security strategy shifts from reactive recovery to proactive preparation, performing the heavy lifting of threat analysis and recovery planning as part of normal operations. This is achieved through a set of natively integrated capabilities:

- **Continuous, Inline Threat Detection**
  This isn't a scheduled scan. It's a constant analysis that happens as backups are created. The platform looks for signs of encryption and is alerted to suspicious modifications, additions, or deletions. This allows it to quickly determine the blast radius of an attack, identifying the exact workloads—down to the file level—that have been compromised.

- **Proactive Threat Hunting**
  The platform continuously scans backups for Indicators of Compromise (IOCs), using file hashes, patterns, or YARA rules. This allows the team to identify which point-in-time backups contain malware, quarantine those snapshots to prevent reinfection, and ensure they are never used in a recovery.

- **Integrated Sensitive Data Discovery and Classification**
  Building on the centralized data classification framework established at Level 2, the platform natively discovers and classifies sensitive information across all workloads in real-time. This provides continuous visibility into data proliferation across business units, enables risk-based protection policies, and ensures the highest level of security is applied to their most critical assets *before* an attack.

- **Orchestrated Recovery with Built-In Testing**
  The platform transforms recovery from a manual, untested hope into an automated, validated capability. Pre-built recovery plans define the exact sequence for restoring multi-tier applications—specifying boot order, network dependencies, and which verified clean snapshots to use. Most critically, these plans can be tested regularly in isolated environments without disrupting production, allowing teams to discover and fix recovery gaps before an incident occurs. This is a fundamental shift: At lower maturity levels, organizations discover their backup systems don't work when they desperately need them to. At Level 4, recovery plans are continuously updated and validated, and RTOs are proven facts, not aspirational goals.

These capabilities are not bolted-on third-party tools. They are woven into the fabric of the platform, bridging the security visibility gap that exists with traditional production-focused security tools. While endpoint detection, cloud security posture management, and threat detection services excel at protecting live environments, they cannot scan backup data. This integrated data security platform extends threat intelligence directly into the backup repository, which can also integrate with trusted sources like Mandiant to incorporate real-time threat intelligence from the broader security ecosystem.

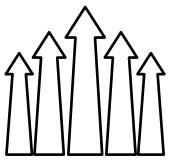

The Strategic Advantage: Predictable Recovery

This proactive model succeeds because it turns the chaotic, unknown timeline of a Level 3 recovery into a predictable, measurable business process. Meeting this level of sophistication solves the four critical questions that paralyze Level 3 organizations during an incident:

- **Which backup copies are clean?** Continuous threat detection has already identified and quarantined snapshots containing malware.

- **When did the attack begin?** Inline analysis pinpoints the likely start of the attack across your entire cloud environment.

- **What's the optimal recovery sequence?** Pre-validated orchestration plans specify the exact boot order and dependencies for multi-tier applications.

- **Has sensitive data been compromised?** Integrated data classification reveals exactly which workloads contained regulated or sensitive information.

Intelligence is then made actionable through recovery orchestration. Instead of manual scripting, the team creates pre-validated cyber recovery plans for critical applications that specify boot order, networking, and which verified clean recovery points to use.

These plans aren't theoretical. They're tested regularly in non-disruptive isolated environments, proving that the entire backup-to-recovery pipeline actually works. This orchestration allows the team to track progress during an actual incident, generate detailed reports for stakeholders, and execute recovery with confidence rather than hope. It transforms cyber resilience from a defensive insurance policy into a strategic business asset.

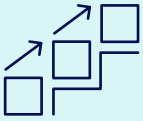

THE BLUEPRINT FOR MATURITY: HOW TO LEVEL UP

Test Relentlessly

Use your orchestration plans to conduct regular, non-disruptive recovery tests to prove your RTO and make sure your recovery plans stay up to date.

Integrate with the Business

Use APIs to connect your data security platform to your CI/CD pipeline, ensuring new applications are born resilient and that the pipelines themselves are protected.

Stay Informed

Ensure you have a way to stay up to date on the latest threat intelligence, so you can protect your organization against the next wave of attacks.

WHERE DO YOU STAND?

THE PATH FORWARD

Being honest about your organization's current maturity level isn't an academic exercise. It's the first step toward building a durable cyber resilience strategy. The difference between levels isn't just about technology. It's about mindset, architecture, and preparation.

Use this guide to start a strategic conversation with your team. Identify your critical flaws, understand the blueprint for maturity, and begin the journey toward predictable, confident recovery, so you'll be ready when (not if) the time comes.

rubrik

Global HQ
3495 Deer Creek Road
Palo Alto, CA 94304
United States

1-844-4RUBRIK
inquiries@rubrik.com
www.rubrik.com