Rubrik Zero Labs

THE IDENTITY CRISIS

UNDERSTANDING & BUILDING RESILIENCE AGAINST IDENTITY-DRIVEN THREATS

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Rubrik Zero Labs and Wakefield Research recently surveyed 1,625 IT and security leaders globally to gain an understanding of their readiness to defend against and recover from identity-based attacks. This is in addition to the over **2.2 million snapshots** Rubrik scans daily in search of threats embedded in backup data.

As traditional network boundaries have dissolved amid cloud migrations, remote work adoption, and now agentic AI, identity is no longer merely a control layer. It has become the primary attack surface, which threat actors weaponize to gain access to IT environments and "live off of the land" over the course of an attack. The overwhelming majority of today's breaches are predicated on exploiting trust and valid credentials rather than circumventing network defenses.

Since almost all attacks include a human or non-human identity component—either for initial access, during privilege escalation, or in conducting lateral movement—it's no surprise that the vast majority of respondents (**90%**) to our survey agree that identity-based attacks represent the single largest threat to their organizations.

This report aims to quantify organizations' ability to withstand identity attacks, spotlighting critical areas of focus and projected response timelines.

Finally, we explore essential elements of identity resilience, including:

- The integration of real-time **visibility, response, and recovery** capabilities across all on-premise and cloud identity providers, accounting for both human and non-human identities
- Building high confidence in the organization's ability to **rapidly and reliably recover core identity infrastructure** itself to a pre-infected state
- The adoption of identity hardening and resilience-building through **zero trust principles** for limiting the attack surface and blast radius of successful identity compromise
- How identity resilience can be integrated into an organization's overall **cyber resilience lifecycle planning**

Only through this holistic approach can organizations harden their identity infrastructure and avert potential downtime, revenue loss, and reputational damage.

But first, a look at year-over-year trends from our previous survey.

2024 VS. 2025 SURVEY TRENDS

(Wakefield)

90%

had experienced a cyber attack in the past year

The same percentage that reported experiencing at least one in 2024

89%

of those who experienced a ransomware attack paid to recover their data or stop the attack

20%

reported experiencing more than 25 attacks compared to 18% in 2024

11% experienced 100 or more, a slight increase from 2024 (8%)

CONFIDENCE IN RECOVERY TIMES IS DECREASING

In 2025, only 28% believed they could fully recover from a cyber incident in 12 hours or less, compared to 43% in 2024

77%

reported managing larger cloud footprints than in the previous year

58%

believe Agentic AI will drive half or more of the cyberattacks they face in the coming year

58%

believe it would take 2 days or more to achieve full-service operations post-compromise

IDENTITY COMPROMISE:

THREAT ACTORS' TICKET TO LIVING OFF OF THE LAND

Most cyber incidents today involve some aspect of identity compromise. But identities should be thought of as tools, rather than targets themselves. Today, they are most often the means to a threat actor's intended end—whether that be surveillance, data theft, or extortion—than ends themselves.

86%

of basic web application attacks today **leverage stolen credentials**

(Verizon)[1]

79%

of CrowdStrike detections were "malware-free," meaning attackers were logging in, not deploying traditional malware

(CrowdStrike 2025)[2]

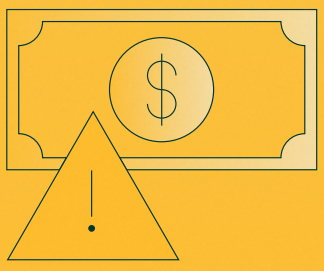

$4.67

MILLION (USD)

per breach

The average cost when attackers leverage compromised credentials

(IBM)[3]

As a threat vector, identity compromise allows threat actors to:

"Live off of the land," evading detection by abusing legitimate processes like admin tools, SaaS services, or even identity workflows themselves

Conduct subsequent attacks using previously compromised credentials

Maintain persistence within target environments, potentially by creating "shadow identity platforms"—rogue tenants or identity infrastructure outside the organization's legitimate governance and visibility.

[1] https://www.verizon.com/business/resources/reports/dbir/
[2] https://www.crowdstrike.com/en-us/resources/reports/global-threat-report-executive-summary-2025/
[3] https://www.ibm.com/reports/data-breach

Human identities are not the sole targets of compromise. Non-human identities (NHIs) are also under attack. These are commonly API tokens used to authenticate automated IT processes, certificates, containers, automation tools, service accounts, and AI agents.

By some measures, NHIs now outnumber human users by 82 to 1.[4] This drastically increases the attack surface for threat actors and will only increase with more widespread adoption of agentic AI.

Threat actors target human identities and NHIs for different reasons, depending on their aims. Understanding which identities are targeted for which purposes is critical to defending each category:

	Human Identities (Employee Accounts, Admin Logins, etc.)	**Non-Human Identities** (Service Accounts, API Keys, Tokens, etc.)
Primary Goal	**Initial access**, reconnaissance, and leveraging existing user-level privileges	Stealth, **persistence**, and high-level system access
Primary Risk	Susceptibility to **social engineering**	Susceptibility to **misconfiguration**, tendency to proliferate exponentially
Defense Evasion	**Harder to maintain:** Human accounts are usually protected by MFA, access policies, and behavior analysis like impossible travel	**Easier to maintain:** NHIs often lack MFA, have fewer monitoring controls, and their automated activity can easily blend into legitimate background noise
Privileges	Privileges are usually tied to a **specific role**, like an engineer or sales rep	Privileges are often **systemic and excessive**, like a service account that can read every database in the system
Persistence	**Easily revoked:** A human account can be locked out, the password rotated, and a session terminated in minutes during an incident	**Difficult to revoke:** NHIs may be long-lived, forgotten, or essential to critical operations, making rotation difficult and disruptive

[4] https://www.cyberark.com/resources/white-papers/2025-identity-security-landscape-executive-summary

THREAT ACTORS' IDENTITY OBSESSION

Identities–both human and non-human–have been central to a spate of recent high-profile cybersecurity incidents.

EntraID & n0Auth used to move laterally

In June 2025, EntraID, Microsoft's cloud-based identity and access management (IAM) service, was found to still contain a flaw that would allow threat actors to pivot from some compromised SaaS applications into an organization's core Microsoft 365 resources, permitting access to sensitive enterprise data.[5]

By manipulating a single mail attribute to match the address of the intended victim, a user could then gain access using a "Log in with Microsoft" feature belonging to a susceptible SaaS application. This attack against an application's identity assertion logic enables rapid lateral movement and unfettered access to core enterprise productivity tools like SharePoint and Teams.

Although the flaw was originally flagged in 2023, it's still believed to impact tens of thousands of SaaS apps. Since n0Auth requires tenant-side remediation, it's impossible to push a fix centrally. In this case, Microsoft's popularity greatly extends the blast radius for this application-level identity flaw. Attackers have continued to target Entra ID, often for the purpose of privilege escalation, throughout 2025.

ToolShed Attacks Target On-Premise SharePoint Servers

Critical vulnerabilities like those exploited in the July 2025 ToolShed attacks show how aspects of identity like key authentication can be manipulated by threat actors. In this case, suspected state-backed actors prioritized the theft of machine keys used to authenticate on-premise SharePoint servers belonging to high-value targets, likely to establish prolonged espionage campaigns from within compromised environments.[6]

This demonstrates the importance of NHIs for motivated threat actors looking to conduct ongoing attacks for purposes including espionage. Since keys are less ephemeral than tokens, it's important for security teams to prioritize their swift rotation following a security incident or the discovery of a CVE like the one enabling ToolShed attacks.

Scattered Spider Hacks Human Nature

The Scattered Spider group frequently impersonates IT or help desk staff to convince individuals to divulge credentials or bypass multi-factor authentication (MFA). Its success is fundamentally rooted in exploiting human psychology, particularly the responsiveness and inherent empathy of support teams that are often under pressure to resolve issues quickly. In August 2023, the group was allegedly able to cause Clorox $380 million in damages by calling the cleaning giant's third-party help desk and requesting a password set.[7]

Scattered Spider's success is not dependent on discovering zero-day software vulnerabilities, but rather on gaining initial access through social engineering and then abusing native tools like PowerShell and SaaS services. This means technical solutions alone are insufficient protection. Instead, organizations must invest in human-centric security, including continuous, adaptive security awareness training and identity infrastructure resilience-building.

[5] https://www.infosecurity-magazine.com/news/microsoft-noauth-flaw-2025/
[6] https://www.recordedfuture.com/blog/toolshell-exploit-chain-thousands-sharepoint-servers-risk
[7] https://www.cybersecuritydive.com/news/clorox-380-million-suit-cognizant-cyberattack/753837/

CONSEQUENCES OF IDENTITY-DRIVEN DISRUPTION

Compromising identity grants attackers the “keys to the kingdom,” allowing threat actors to live off of the land, abusing legitimate tools to conduct surveillance and exfiltrate data.

It's easy to see why a focus on identity has surged in popularity when evaluating the damage that can accomplished with a single successful compromise:

Impact Category	Outcome	Affected Organizational Assets
Data Breach	Exfiltration of PII, intellectual property, financial data, or other valuable data	Customer data, employee data, trade secrets, competitive advantage, financial records
Financial Loss	Direct fraud, incident response costs, legal fees, regulatory fines	Budget, revenue, legal resources
Reputational Damage	Loss of customer trust, brand erosion, negative publicity	Brand value, customer & partner relationship, market value
Operational Disruption	System downtime, service interruption, account lockouts	Revenue, IT infrastructure investment, security controls
Persistence & Privilege Escalation	Installed backdoors (service principals, federation mods), creation of shadow identity platforms	IT infrastructure, identity and access management, security controls
Compliance & Legal Risk	Non-compliance with data privacy regulations (GDPR, CCPA), lawsuits	Legal resources, compliance office
Supply Chain Compromise	Compromise of business partners and other third parties, leading to cascading attacks	Partner ecosystem, supply chain integrity, reputation

HOW ARE ORGANIZATIONS RESPONDING?

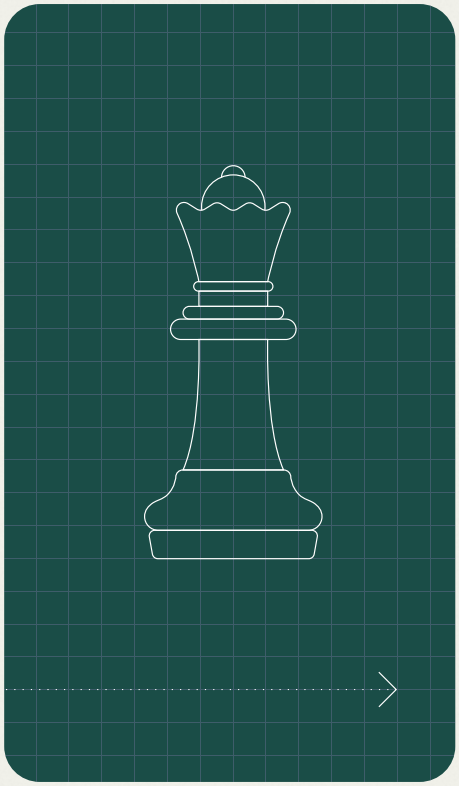

90%

of IT and security leaders surveyed agree that identity-driven cyberattacks are the top threat to their organizations.

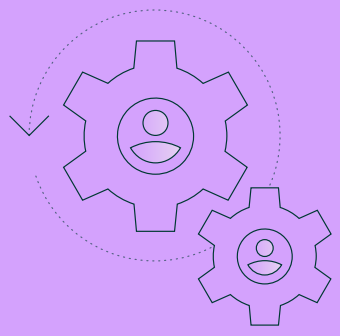

89%

plan to hire professionals specifically to manage or improve digital-identity management, identity infrastructure, and/or identity security in the next 12 months.

Which types of identities are you most concerned about being compromised:

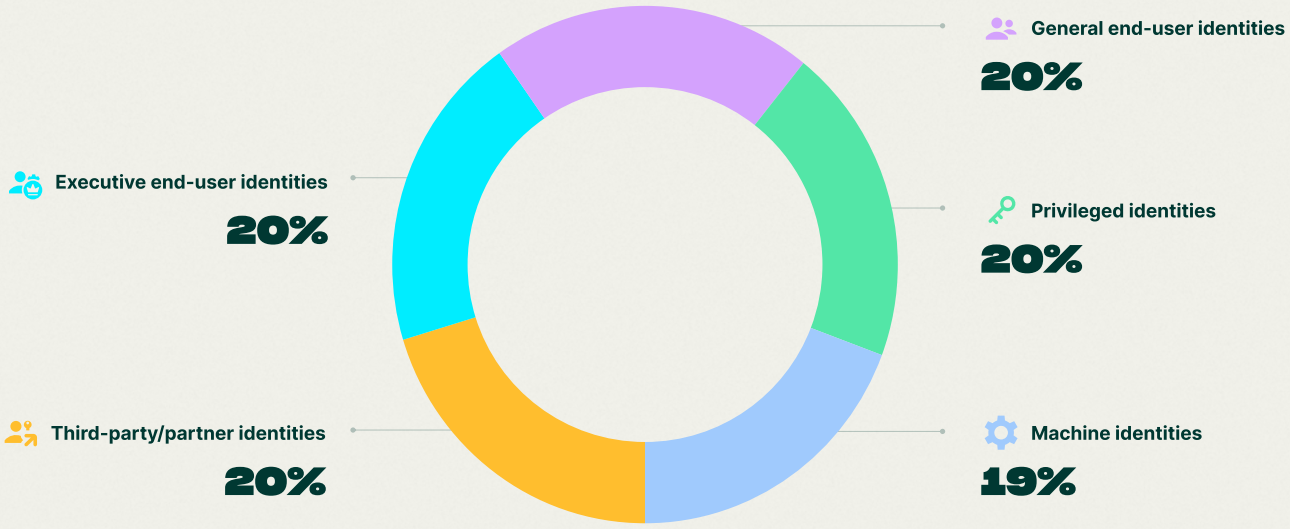

Where leaders once prioritized only privileged accounts, there is now a clear consensus that every identity is critical. This is acknowledgement that any compromised account—regardless of its initial permissions—is a dangerous foothold for an attacker to begin lateral movement and privilege escalation.

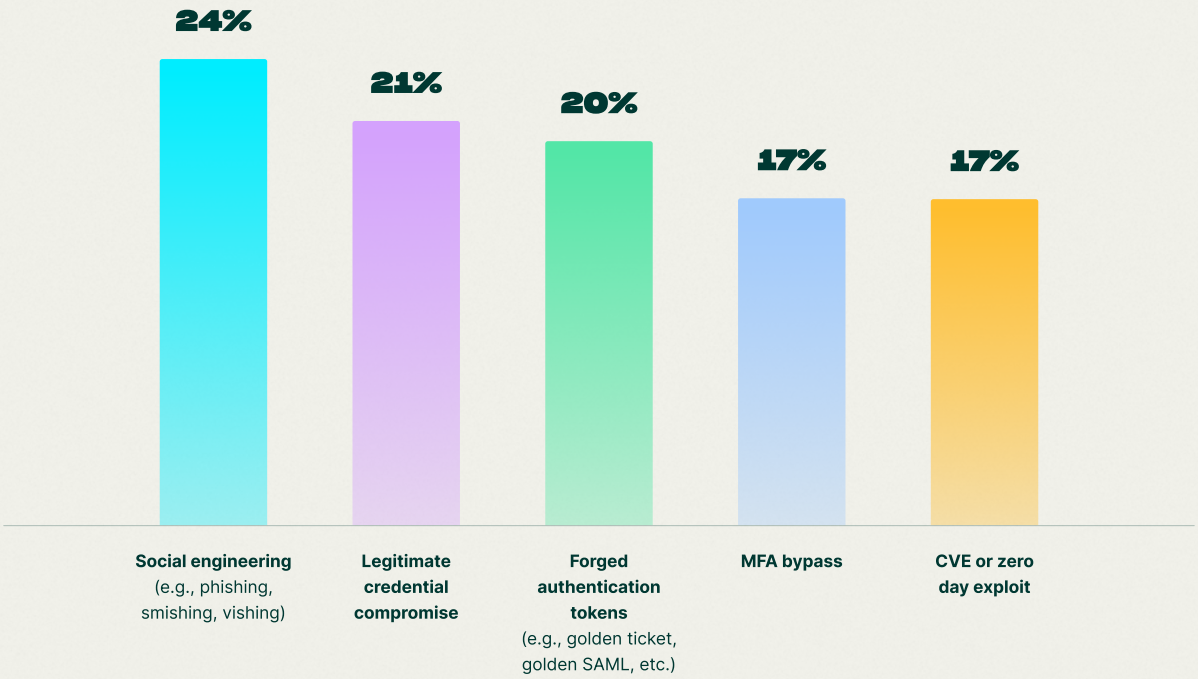

Data also points to shifting security priorities. While social engineering (24%) leads, MFA bypass (17%) and forged authentication tokens (20%) are now considered just as dangerous as zero-day exploits (17%).

This is a profound change. It shows leaders have recognized a new reality: the 'log in' attack powered by sophisticated identity forgery is now a more prevalent threat than the traditional 'break in' exploit.

Mandiant reported responding to more cloud breaches than ever before in 2024,[8] and cited identity solutions lacking sufficient security controls as the chief reason. As identity evolves as an attack vector, identity access management (IAM) solutions struggle to keep up with the latest techniques.

This is likely because:

Identity is multifaceted

Privileged access management (PAM), role-based access controls (RBAC), and API security are all sub-disciplines of IAM, each of which may be provided by a single vendor, expanding the attack surface.

New identity types are continuously emerging

AI bots and agents are variations of NHIs that did not exist five years ago, and APIs are being put to more uses than ever. Agents will create an explosion in the volume of machine-to-machine identities, which often lack the mature lifecycle management (i.e. provisioning, rotation, de-provisioning) applied to human accounts.

The way threat actors leverage identity is evolving

As endpoint and perimeter defenses increase in effectiveness, threat actors increasingly prioritize credential theft to “log in, not break in.”

The cloud introduces new risk

The complexity and granularity of cloud-native IAMs make them highly susceptible to human error and misconfiguration. Attackers actively scan for these misconfigurations, such as over-privileged roles or accidental public permissions.

[8] https://services.google.com/fh/files/misc/m-trends-2025-en.pdf

So it's no surprise that:

87%

of IT & security leaders are currently planning to change IAM providers or have already begun the process.

58%

of them are looking to do so for security reasons, suggesting the feature set for tackling identity-based threats is lacking in many single solutions.

60%

This is despite the fact that **60% had switched IAM providers within the last three years.**

Main reasons for switching identity providers:

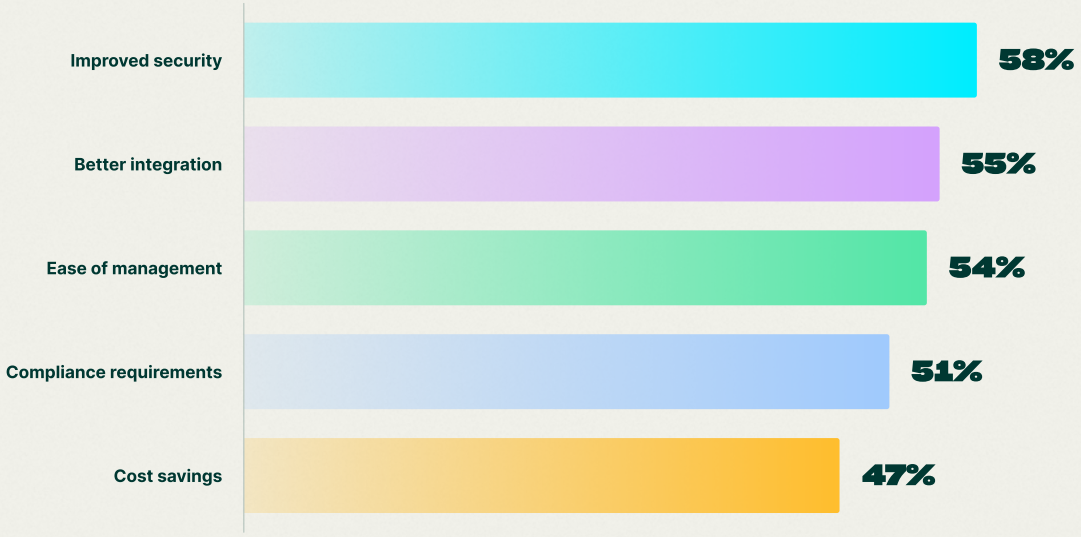

More concerning is the time it would take most to restore their identity infrastructure post-compromise. When the cost of downtime can reach **$6,000 per minute**,[9] organizations quickly face mounting expenses by relying on manual recovery processes (**54%**).

[9] https://www.isaca.org/resources/news-and-trends/industry-news/2024/centralized-services-and-their-impact-on-business-continuity

Reported time to recover identity infrastructure post compromise:

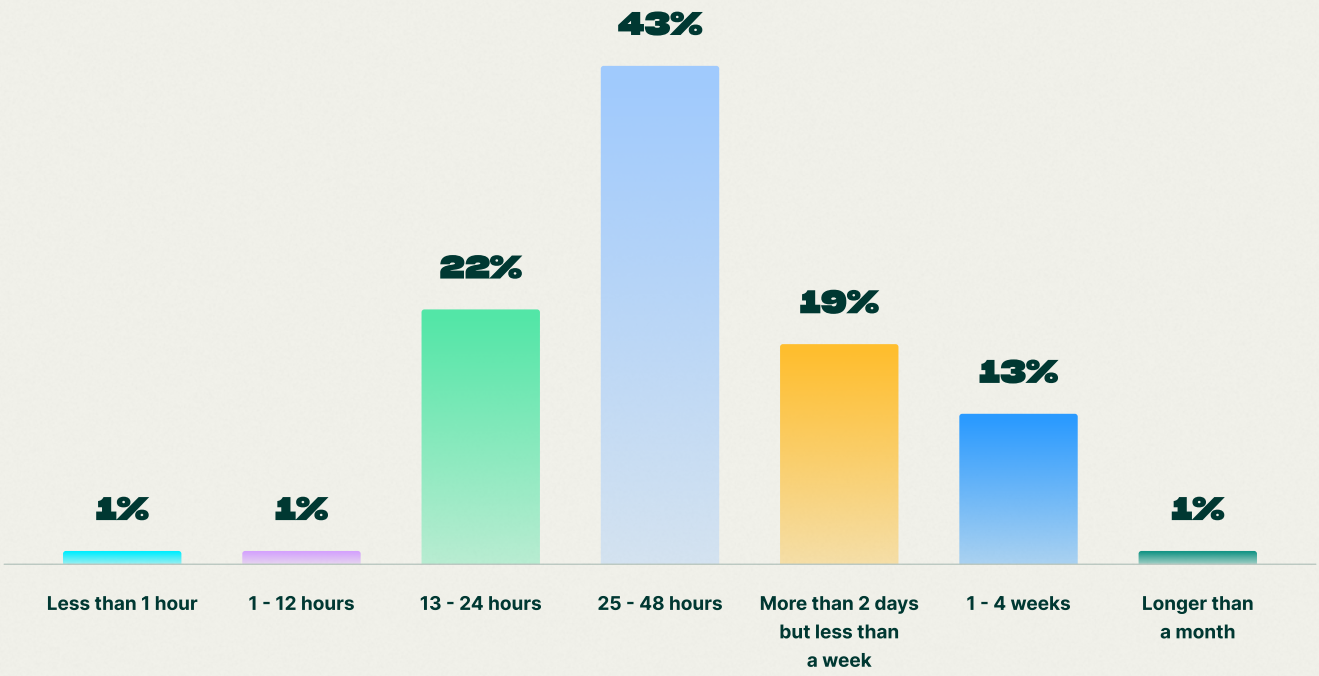

Processes requiring manual recovery procedures:

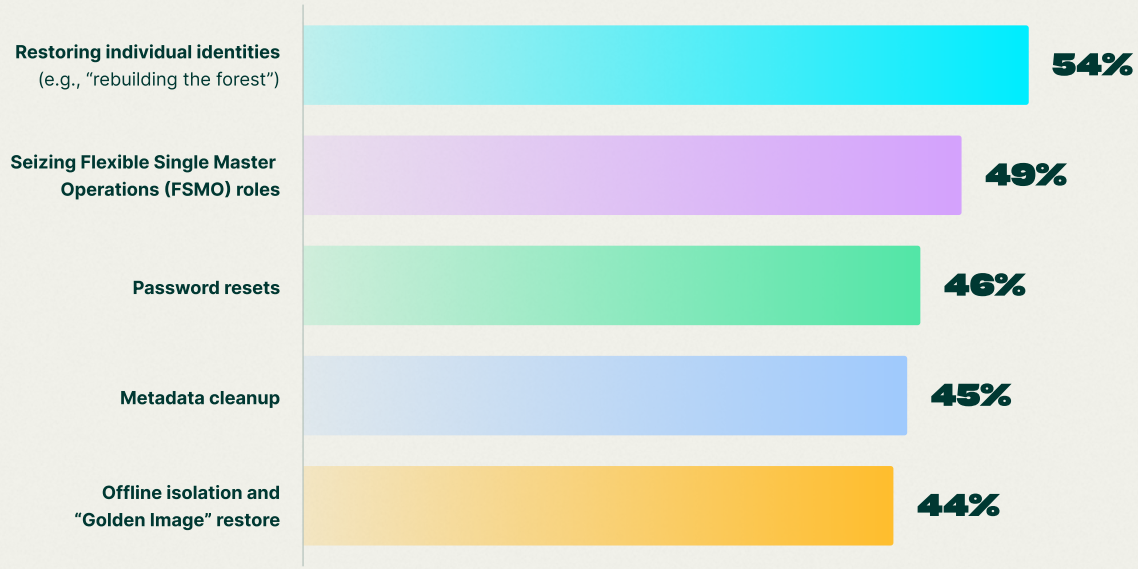

Interestingly, **89% of respondents have fully or partially incorporated AI agents into their identity infrastructure**. These new forms of NHIs will also need to be protected at massive scale and over the course of their entire lifecycle to prevent them from being weaponized. Already, more than half of respondents (**58%**) estimate that in the next year, **half or more of the cyberattacks they face will be driven by agentic AI**.

A failure to understand which threat actors target which identities—and why—is another factor hindering organizations' ability to measure and improve response and recovery times.

WHY DOES RECOVERY TAKE SO LONG?

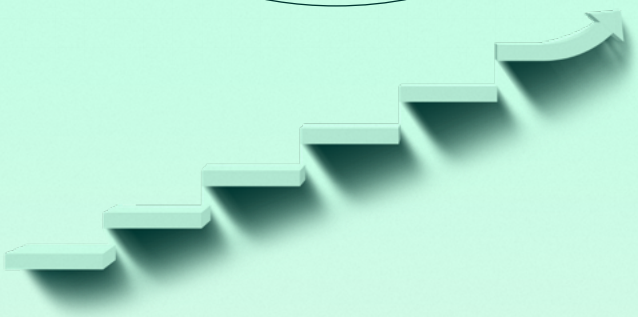

MTTR AS A DATA-DRIVEN ANSWER

Organizations often strive for cyber resilience, but struggle to quantify it. How do we measure the ability to anticipate, withstand, recover from, and adapt to cyber threats?

Recovery time objective (RTO) and recovery point objective (RPO) are traditional measures, but they fail to tell the full story. Maybe the RTO from an incident was four hours, but recovery actually took 32 days. An RPO may be measured in minutes but, in industries like banking, this could result in the loss of data pertaining to hundreds of thousands of transactions. In these examples, we know recovery operations teams overshot their objective significantly, but not why.

Moreover, much of the industry relies on a number of different metrics that can prevent a standardized understanding of recovery time across verticals. For some the clock stops when a threat is detected. For others, only when it's been eradicated completely. This makes it difficult to gain a true understanding of an organization's cyber resilience.

That's why Rubrik Zero Labs recommends **Mean Time to Recover (MTTR)** as an industry-standard, phase-aware, and data-driven metric for measuring an organization's ability to recover from a cyber incident.

THE NEW FRAMEWORK FOR DECONSTRUCTING MTTR

MTTR is not a single measurement. It is a multi-step process broken down into discrete phases to provide meaningful insights and identify areas for improvement. This model for recovery includes:

Detection	This is the starting point for any incident, whether it be a cyberattack or an accidental deletion.
Scoping	This phase involves the comprehensive effort to understand all the impacted systems and assets, as well as their dependencies. Understanding which applications and databases rely on one another is crucial.
Identify Clean Recovery Point	This is often the longest and most critical phase. It is the time it takes to find a trustworthy, malware-free backup from which to restore operations. Historically, this meant weeks or months setting up a dedicated "clean room" lab environment. Modern tools like cloud-based infrastructure and automation are significantly reducing this time.
Restore	This is the actual function of restoring data. If this takes multiple weeks, it could signal a throughput issue—a data-driven argument for investing in infrastructure to reduce high-risk recovery times. Given the amount of data organizations possess today, this phase could overtake identifying a clean recovery point as the longest recovery step.
Validation	How long does it take to ensure application health and data access? How long to ensure the necessary human and non-human identities are functioning and communicating correctly?

By measuring each phase individually, an organization moves away from a single, context-free number toward the granular insight needed to identify bottlenecks and make data-driven decisions.

Phased Metric	Clock Start	Clock Stop	Description
Mean-Time-to-Detect (MTTD)	Adversary impact timestamp or time of alert	Upon confirmation of known recovery scope	Driven by detection, anomaly analytics, and SIEM correlation
Mean-Time-to-Scope (MTTS)	Detection validated	A list of recovery objects is locked	Global Search, Data Classification, and SLA metadata drive this phase.
Mean-Time-Select-Clean-Snapshot (MTTCS)	Scope locked	A clean, uncompromised snapshot is identified	Immutability, threat scanning, and snapshot health scoring reduce this time
Mean-Time-to-Restore (MTTr)	Restore trigger	Data is made available to workload	Instant recovery, live mount, and cloud disaster recovery orchestration capabilities facilitate
Mean-Time-to-Validate (MTTV)	Data access ready	Application health is verified	Achieved through app consistency, automation playbooks, and post-restore malware scans

Total MTTR (Operational) = MTTD + MTTScope + MTTCS + MTTRestore + MTTValidate

WHY MTTR MATTERS

Organizations today often overlook the insights available within their backup data. MTTR turns this "crown jewel" into a tool not only for ensuring recovery but also for measuring resilience.

Ultimately, MTTR is one phase in a continuous maturity assessment cycle of:

1. **Determining resilience**
   Understanding current capabilities for withstanding an attack based on the most critical applications and likely attack scenarios via a well-designed scoring system akin to a risk assessment.

2. **Establishing a minimum viable business**
   By going beyond surveys and tiered application lists to truly understanding the essential applications and databases—as well as their dependencies—needed to maintain continuity by prioritizing the mission-critical.

3. **Validating recoverability**
   Through *true* crisis simulations—software-driven scenarios depicting realistic interactive drills—to produce measurable recovery time metrics in peacetime engagements that act as baseline for future improvement.

4. **Analyzing and improving**
   At each phase of the recovery effort to determine where bottlenecks are occurring and how they can be addressed. For example, decreasing mission-critical recovery data from 20TB to 2TB will significantly reduce throughput during recovery.

Crucially, by anonymizing and aggregating MTTR data, organizations can create powerful benchmarks against industry and regional peers. This can empower leaders to secure funding for cyber resilience and help break down silos by giving security, IT, and recovery operations teams a unified understanding of key performance indicators like recovery velocity and reliability.

INSTILLING IDENTITY RESILIENCE:

RECOMMENDATIONS AND RESPONSE CAPABILITIES

Identity should not be considered merely an asset to be protected. It should be treated as a primary control plane for all security decisions within a modern enterprise.

Investment must shift from solely endpoint and network-centric defenses to comprehensive identity governance, privileged access management, and advanced authentication solutions. A compromised identity now represents a direct—often undetected—path to an organization's most critical assets. This makes identity the most critical control point.

Given the volume of attacks involving identities, organizations must prioritize:

Identity visibility and recovery

If a cloud API key is compromised, which systems must be disconnected? If an Okta or Active Directory admin account is hijacked, can you quickly isolate the issue and force re-authentication for affected users? Answering and addressing these questions requires real-time visibility and recovery capabilities across hybrid identity environments.

Identity resilience-building

Resilience is the ability to bounce back quickly and confidently. Secure, offline backups of Active Directory or cloud directory data are critical for quickly rebuilding identity services if they are encrypted or wiped. By planning recovery steps in tandem with security, organizations drastically reduce the downtime and financial impact resulting from identity compromise.

Limiting attack surface and blast radius through zero trust principles

Identity should be considered the perimeter. Every access request, from inside or outside the network, must be continuously authenticated, authorized, and encrypted. Practically, this means enforcing strong MFA authentication, conditional access policies, and least privilege access by design.

ZERO TRUST APPLICATIONS IN IDENTITY SECURITY

Least privilege and role-based access control (RBAC)

All users and devices, regardless of location, are granted only the minimum access rights necessary to perform their specific job functions.

Just-in-Time (JIT) access

Granting elevated permissions or access only for the specific, limited duration that a user needs to complete a sensitive task, after which the permissions are automatically revoked.

Continuous verification

Constantly monitoring and re-evaluating a user's identity, device posture, and context (such as location and time) to maintain or revoke access throughout the entire session, in addition to strong MFA.

Microsegmentation

Networks should be subdivided to contain potential breaches and ensure that unauthorized users cannot easily traverse segments and move laterally in search of valuable resources to compromise.

THE CYBER RESILIENCE LIFECYCLE

True cyber resilience goes beyond cybersecurity by integrating **risk management, business continuity**, and **incident response** into a unified strategy. The goal is not just to prevent attacks, but to withstand them and bounce back quickly with minimal impact.

The Cyber Resilience Lifecycle

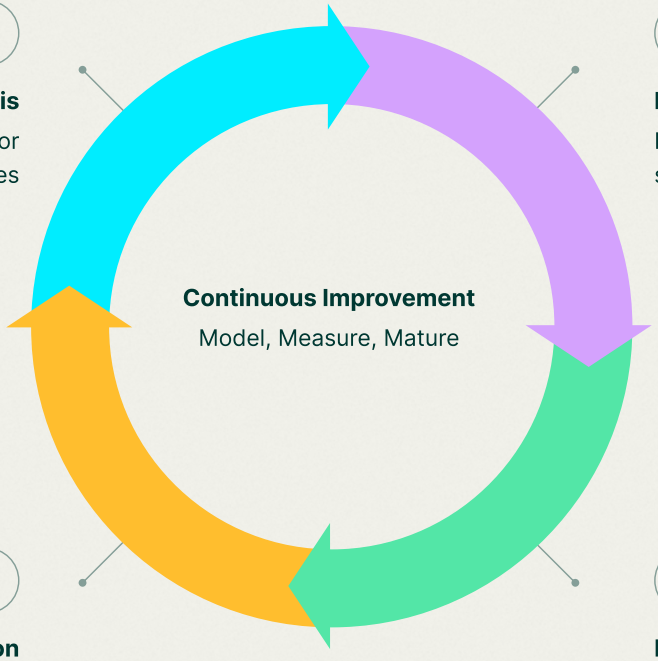

1

Resilience Assessment

Define and understand current security posture and risk profile

2

Establish MVB

Identify key assets and infrastructure to establish minimum viable business

3

Crisis Situation

Stress test defenses and security posture to measure outcomes

4

MTTR Analysis

Identify and fix bottlenecks for improved future outcomes

Continuous Improvement

Model, Measure, Mature

Aligned with NIST guidance[10] for anticipating, withstanding, recovering, and adapting to cyber attacks, the Rubrik Zero Labs framework seeks to operationalize resilience for practitioners and provide benchmarking and improvement metrics for senior leaders.

By integrating identity resilience planning into broader efforts to bolster cyber resilience, organizations take meaningful steps toward minimizing disruption, safeguarding business-critical assets, and earning the trust of valued stakeholders.

[10] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160v2r1.pdf

DATA AND METHODOLOGY

Rubrik Zero Labs is committed to providing practical, unbiased intelligence aimed at helping organizations reduce their data security risk.

To achieve this, we have included information from three main sources:

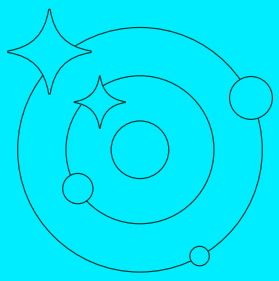

RUBRIK TELEMETRY

We employed Rubrik telemetry to gain insights into the typical organization's data environment and associated risks

INDEPENDENT RESEARCH

Perspectives from 1,600+ IT and security leaders through Wakefield Research

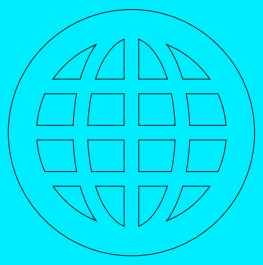

CONTRIBUTING ORGANIZATIONS

Research from respected cybersecurity organizations and institutions