

ResOps: The Future of Resilient Business in the Era of AI




From Disruption to Continuous Trust

AI is reshaping business at unprecedented speed, but it is only one of several forces increasing operational fragility. Cloud complexity, SaaS growth, identity sprawl, expanding third-party dependencies, and tightening regulations have created an environment where disruptions outpace traditional response models.

Organizations are investing heavily in GPUs, models, and automation, but the practices required to operationalize the resilience of these capabilities have not kept pace. As a result, errors spread faster, attack surfaces widen, and risks to data integrity continue to grow.

The result is a widening resilience gap. Innovation is accelerating faster than the protections needed to support it. Most enterprises are unprepared for this shift. Accenture reports that 77% lack essential data and AI security practices, and 90% lack the maturity required to defend against AI-enabled threats¹.

1. [State of Cybersecurity Resilience 2025](#), Accenture



The Resilience Gap

Traditional resilience models were built for an era of strong perimeters and isolated recovery processes. That approach fails today, when disruptions move at machine speed and systems span hybrid, multi-cloud, SaaS, edge, and third-party environments.

These interconnected ecosystems create new dependencies, new failure modes, and reduced visibility. AI accelerates the challenge even further – decisions and errors propagate instantly, and attackers are now targeting data integrity as much as availability.

Siloed functions: Security, identity, and recovery remain fragmented, slowing response times and creating blind spots.

Distributed data: Hybrid, multi-cloud, and edge ecosystems have expanded the attack surface.

AI acceleration: Machine-speed operations mean threats cascade faster than humans can respond.

Reactive posture: Recovery is still treated as a last resort – slow, partial, and costly.

According to the Cloud Security Alliance, only 2% of businesses have implemented cyber resilience across all critical areas. Meanwhile, 42% of executives cite cloud-related threats as their top concern – the same threats security leaders feel least prepared to defend against².

The result is a widening **resilience gap**: the distance between how fast technology evolves and how slowly most enterprises can respond when something goes wrong.

² Cloud Security Alliance

ResOps: A New Operating Model for Resilience

ResOps (Resilience Operations) brings together disciplines that have traditionally operated in silos – security, identity, data protection, recovery, and governance – into a unified operational model. While many vendors (including Commvault) contribute important components, the model itself is broader than any single technology or domain.

At its core, ResOps is an ongoing loop that strengthens resilience across the organization: safeguarding data, detecting risks, restoring trust quickly when incidents occur, and feeding clean, verified data back into the business.

ResOps Across the Enterprise



These functions exist across the industry today. But ResOps elevates them into an operational discipline that aligns teams, increases accountability, and enables resilience to be continuously proven rather than assumed.

In that sense, ResOps introduces a new operational persona: **one focused on addressing resilience risks holistically and efficiently across the enterprise.**

Discover & Protect: Securing Data at the Source

This is the foundational pillar of ResOps: knowing what data exists, where it resides, how it is classified, and whether appropriate protections and governance are applied.

What It Is

Discovery and protection form the foundation of ResOps. Organizations must continuously know what data they have, where it resides, and how it is being used. Protection means applying governance and control as data is created, accessed, moves, and evolves – so it remains trustworthy throughout its lifecycle.

Why It Matters

Whether protecting customer data, meeting regulatory compliance, or powering AI innovation, enterprises face a common challenge: trusted, governed, and secure data in real time.

This is especially critical for organizations adopting retrieval-augmented generation (RAG) to extend large language models with proprietary data. Without discovery and protection – knowing what data exists, where it resides, and how it's classified – RAG pipelines can inadvertently expose or contaminate sensitive information. Unclassified or unprotected data can flow directly into vector databases, becoming an attack surface for data poisoning or privacy leakage.

Continuous discovery, classification, and protection create the trusted, governed data which fuels AI systems. This is foundational for responsible AI and resilient operations.

How to Achieve It

Unified discovery and classification:

Continuous scanning across hybrid, multi-cloud, and edge environments, classifying and labeling data at creation.

Policy-driven protection:

Apply security, backup, and retention policies automatically at the data layer – enforcing encryption, masking, access controls, and recovery readiness as data is created, moved, and stored.

Proactive governance:

Continuously monitor and enforce data access, backup, and retention policies for consistent compliance across environments.



Detect: Observing Change in Real Time

Detection within ResOps extends beyond security alerts. It includes monitoring data integrity, identity behavior, system changes, and any deviation that could compromise resilience.

What It Is

Detection is about always-on observability, identifying anomalies, attacks, and unintended changes in identities and the data they access across the enterprise, including critical copies of data.

Why It Matters

Business operations – including both human and AI-driven decisions – rely on trusted data. Every second that data is inaccurate, unreliable, or compromised puts the business at risk. Detection shortens that window, allowing resilience to operate at speed.

How to Achieve It

AI-driven anomaly detection:

Monitor data and identity behaviors in near-real-time, surfacing deviations rapidly to minimize risk.

Auditing and visibility:


Gain insight not just into what changed, but who – or which digital agent – made the change.

Automated correlation:

Separate normal business activity from indicators of compromise, reducing false positives and alert fatigue.

Actionable intelligence:

Detection feeds directly into response playbooks, so insight translates instantly into action.



Recover: The Fastest, Cleanest, Most Complete Recovery

Recovery is a core pillar of ResOps, but in the broader industry model it is no longer limited to restoring backups – it includes restoring full business services with confidence, speed, and integrity.

What It Is

Recovery is no longer a passive backup plan. It is the active, operational core of resilience. In ResOps, recovery means restoring the full enterprise – fast, clean, and complete.

Why It Matters

Modern operations – from AI systems to customer applications – depend on data that is complete, accurate, and uncompromised. When recovery falls short, customer confidence falters and regulatory compliance is put at risk. True resilience means restoring the business with integrity and assurance, every time.

How to Achieve It

Continuous testing:

Test recovery paths daily, so readiness is proven and trusted.

Synthetic Recovery:

Surgically remove compromised data and automate precise recovery, resulting in rapid, clean, complete recovery.

Automated orchestration:

Automate runbooks and workflows, for testing, forensics, and production validation by enabling isolated, rapid recovery that minimizes the need for manual processes and intervention.

Enterprise-scale restoration:

Restore the fastest, cleanest, most complete version of the business across workloads, clouds, and data estates.

Importantly, traditional metrics that would have been used to measure restoration fall short in the context of modern recovery, driven by the ResOps model. Recovery time objectives and recovery point objectives, for example, give no attention to the fact that data needs to be clean of corruption before it is restored to production environments.

New metrics are now evolving – and being adopted by modern enterprises – that are appropriate to the task at hand. Mean Time to Clean Recovery is one metric that a board of directors can use to “ask the right question” of its ResOps personnel in relation to a business’ ability to recover from a cyberattack.

The Case for a Unified Approach to Resilient Operations

Whether protecting against ransomware, supporting compliance, or enabling AI innovation, fragmented ecosystems cannot deliver resilience at enterprise scale. Security tools, identity systems, and recovery processes were never designed to work together. The result is duplication, inconsistency, and gaps that attackers exploit.

Unified across data:

True resilience requires consistency across structured and unstructured data, SaaS applications, endpoints, and cloud-native services. Without it, organizations multiply risk and cost through fragmented tools and redundant processes – different teams solving the same problem in different ways, each adding complexity instead of control.

Unified governance:

By extending consistent policies across today's production workloads and tomorrow's AI stacks, governance dramatically reduces risk. It shrinks the surface area of exposure as data sprawls across clouds, regions, and edge locations.

Unified security, identity, and recovery:

Bringing the disciplines of data security, identity resilience, and cyber recovery together increases the integrity and trust of critical data. Centralizing what is otherwise a fractured set of tools for recoverability and resilience across the enterprise.

Unified intelligence:

Continuous intelligence increases the availability of critical data. By applying AI-driven insights, organizations can trust that their data is always available – enabling the cleanest, fastest, most complete recovery when it matters most.

ResOps as Competitive Advantage

ResOps is not simply an evolution of existing practices, it is a new category of enterprise operations purpose-built for the AI era. Just as DevOps redefined how software is built and delivered, ResOps establishes the operational discipline required to safeguard, govern, and restore trust in data-driven businesses.

It unifies what historically has been fragmented – security, identity, data protection, governance, and recovery – into a single, continuous resilience fabric. This integration creates a new class of operational capability: one that continuously validates data integrity, preserves the reliability of AI systems, and enables the business to withstand and recover from disruption at machine speed.

In a world where AI decisions propagate instantly and digital ecosystems are deeply interconnected, resilience can no longer be episodic, siloed, or assumed. It must be engineered, automated, and proven continuously.

That is the defining purpose of ResOps.