Evidence Over Hope:

The Executive Case for Resilience Operations

Critical services go down
When attacks, outages, or failures hit – the question is how fast and cleanly you can recover.

Recovery plans fail under pressure
Plans tested annually with low-stress scenarios routinely fail when real disruptions hit.

READINESS REPORT ISSUE 4 APRIL 2026

READIVERSE

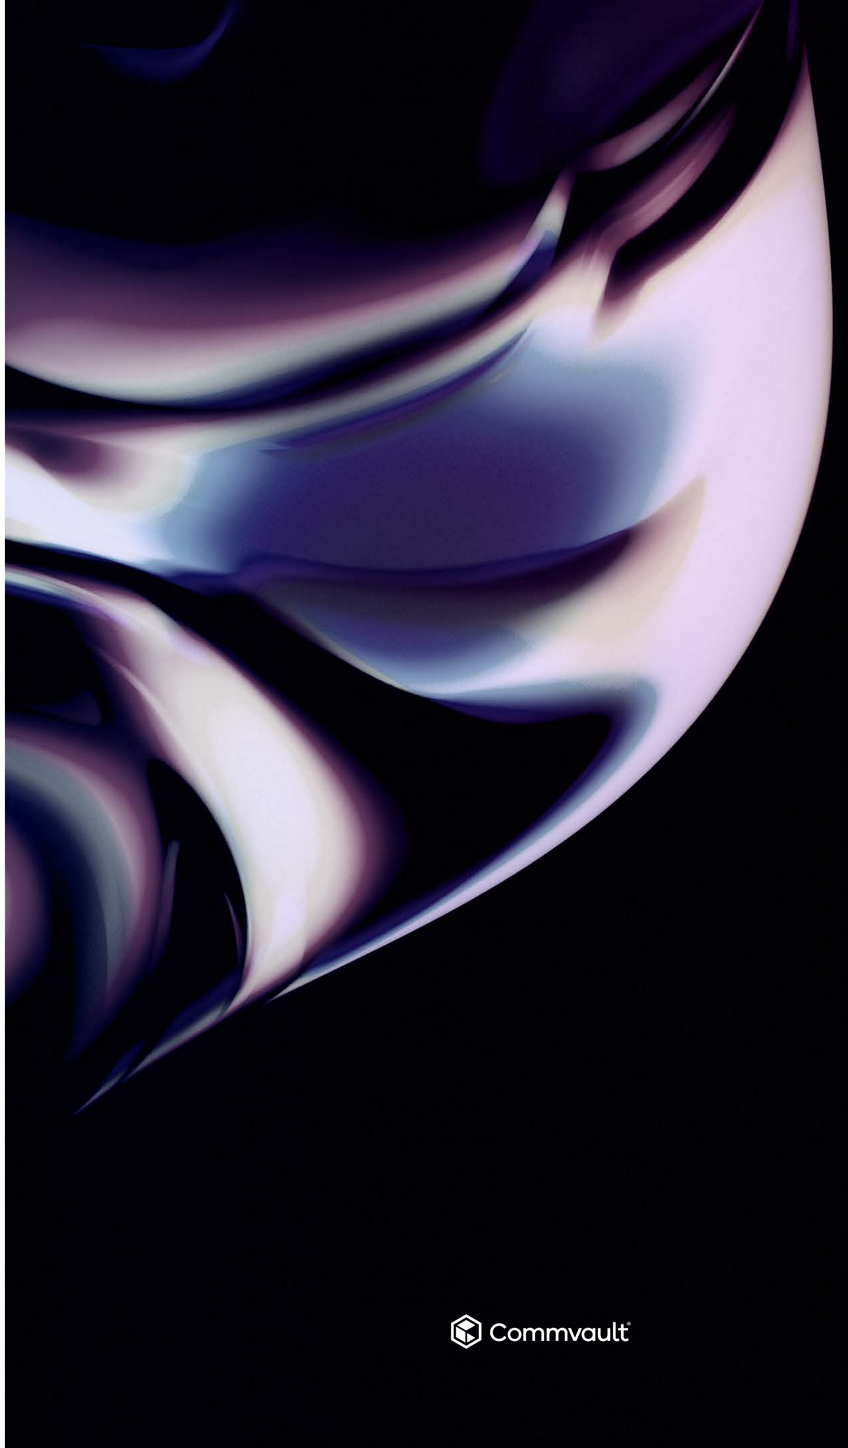

Commvault

Executive Summary

Three questions sit at the center of every significant cyber event, outage, or disruption an enterprise faces:

- Can we recover our critical services?
- How long will it take?
- Can we prove it before the incident happens?

Cyber security answers: Are we under attack? Backup and disaster recovery answer: Do we have copies? Neither answers the question that actually matters: Can we actually come back for each critical service, right now?

That gap is the problem Resilience Operations (ResOps) solves.

ResOps is the operational discipline that unifies security, infrastructure, and business operations around a single, measurable outcome: the ability to withstand disruption and restore critical services within defined impact tolerances.

It is designed to help replace assumptions with evidence, support the shift from annual disaster recovery tests to continuous validation, and enable the transition from fragmented ownership to accountable leaders who can answer the recoverability question at any moment.

This report aims to present the executive case for ResOps, seeks to explain how it differs from every other operational discipline, and offers CIOs and CISOs a practical framework for standing up the function starting with the right structure, the right team, and the right tolerances.

The central insight

Every other operational discipline optimizes a positive workflow. ResOps is the only discipline built for the moments when the system is not operating as intended. Its success criterion is surviving failure.

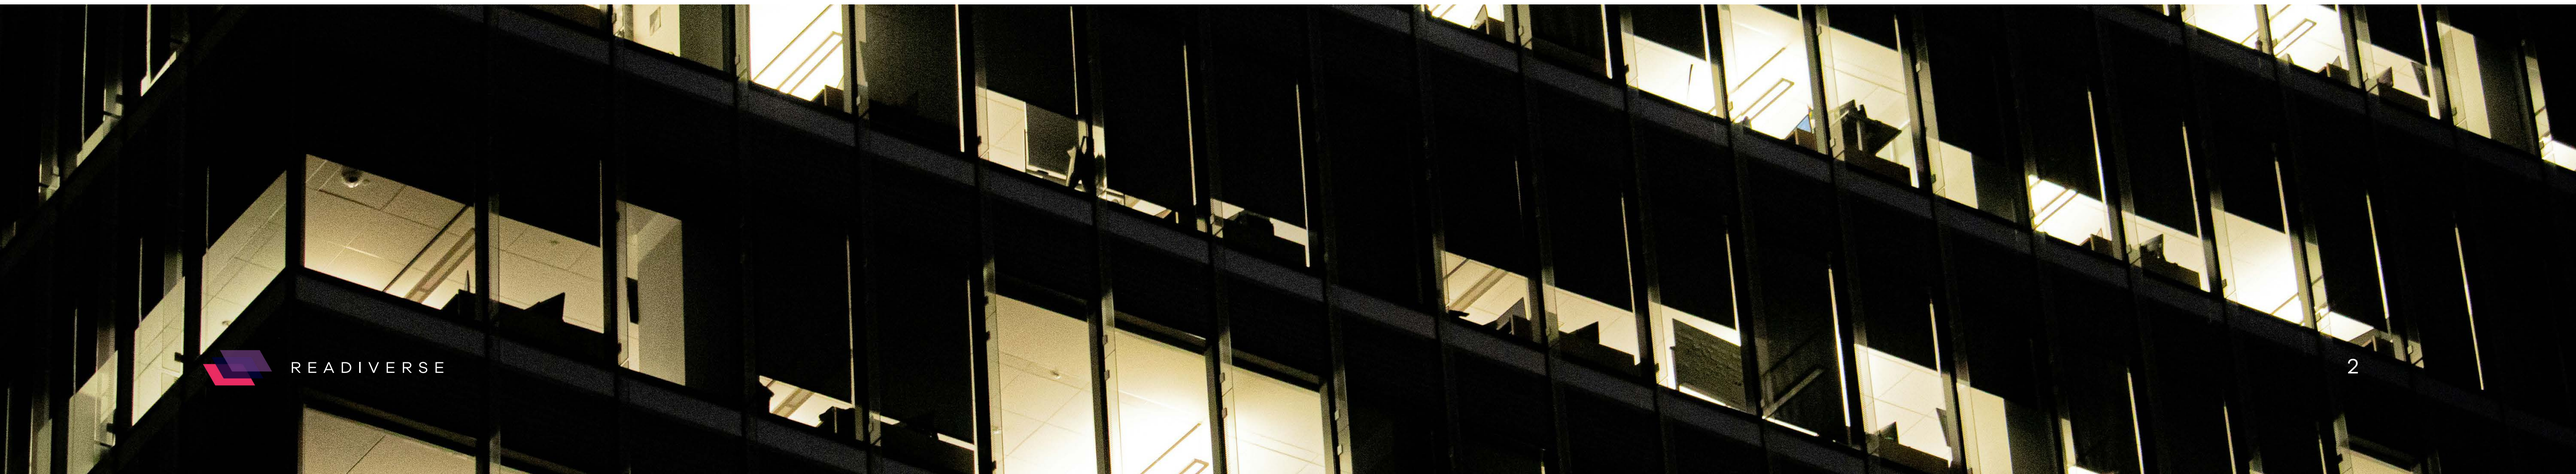

The Disruption Reality

Disruptions now move faster than organizations can react. Outages cascade across on-premises, cloud, and SaaS ecosystems in seconds. Cyberattacks that once took weeks to propagate now unfold in minutes. Insider threats, talent gaps, and brittle supply chains expose hidden single points of failure that don't appear on any risk register until they activate.

And then there is AI – which accelerates both mistakes and attacks while introducing entirely new failure modes: model drift, data poisoning, prompt exploitation. The organizations building the most sophisticated AI capabilities are often the same ones creating the most novel recovery challenges.

What emerges from this landscapeis a critical gap between two realities most organizations live in simultaneously:

Reality		What it Means
We have backups.	→	Backups confirm data exists. They say nothing about whether that data is clean, complete, or recoverable within the window that matters.
We have disaster recovery plans.	→	Plans describe intent. Most were written during normal conditions, are tested annually in low-stress simulations, and have never been validated against a real disruption at scale.
We have security tools.	→	Security tools detect and contain. They cannot answer whether the business can recover the specific services customers and regulators depend on.
We have compliance documentation.	→	Regulators now expect demonstrated operational resilience, not documented processes. The gap between documentation and proof can be a significant source of regulatory risk.
We measure recovery time objectives (RTO) and recovery point objectives (RPO).	→	Speed and recency of recovery are necessary but insufficient measures. They assume recovery data is clean and trustworthy. In an era where attackers specifically target backup infrastructure, that assumption is no longer safe.

The consequence of living in this gap is not theoretical. When real disruptions occur, organizations with fragmented resilience discover in real time what their recovery architecture can and cannot do. The organizations that recover fastest are not necessarily those with the most tools, but those that treated recoverability as an operational discipline before the incident arrived.

Why ResOps Is Different

ResOps is not a rebranding of backup. It is not an extension of security operations or IT operations. It is a fundamentally different discipline, because its success criterion is fundamentally different.

Every other operational discipline optimizes a positive workflow:

Discipline		Optimizes For
DevOps	→	Speed and reliability of software delivery
SecOps	→	Detection and containment of threats
IT Ops	→	Stability and availability of systems
DataOps	→	Quality and flow of data products
AIOps	→	Automated insight and event reduction
CloudOps	→	Cost and performance management of cloud infrastructure

All of these assume the business is functioning normally. Their goal is to improve efficiency, reliability, automation, or visibility while the system is operating as intended.

ResOps is the only discipline built for the moments when the system is not operating as intended. It is designed to optimize recovery under stress. It is intended to assume total disruption and seeks to design rebuild pathways. It is built to rely on evidence from realistic tests, not confidence from annual simulations. It aims to demonstrate an ability to meet impact tolerances in real scenarios. It seeks to manage the boundary between normal operations and crisis-state operations. And it is designed to span the entire enterprise – not a single team or domain.

The defining question

ResOps exists to answer one question with evidence rather than assumption: “Can we actually come back – for each critical service – right now?”

What ResOps Requires

Building a ResOps capability requires three foundational commitments that other operational disciplines don't demand:

Accountable ownership at the executive level.
Not a shared responsibility across teams, but named owners who define impact tolerances, establish budgets, and can answer the recoverability question for specific critical services.

Continuous, realistic validation.
Not annual disaster recovery tests in controlled environments, but regular exercises that simulate real scenarios – ransomware, cloud region failure, key SaaS provider loss, AI system compromise – in isolated environments help that protect business continuity.

Evidence over assertion.
Not confidence based on documentation, but measurable indicators that show whether each critical service can be restored within defined tolerances, updated continuously as environments change.

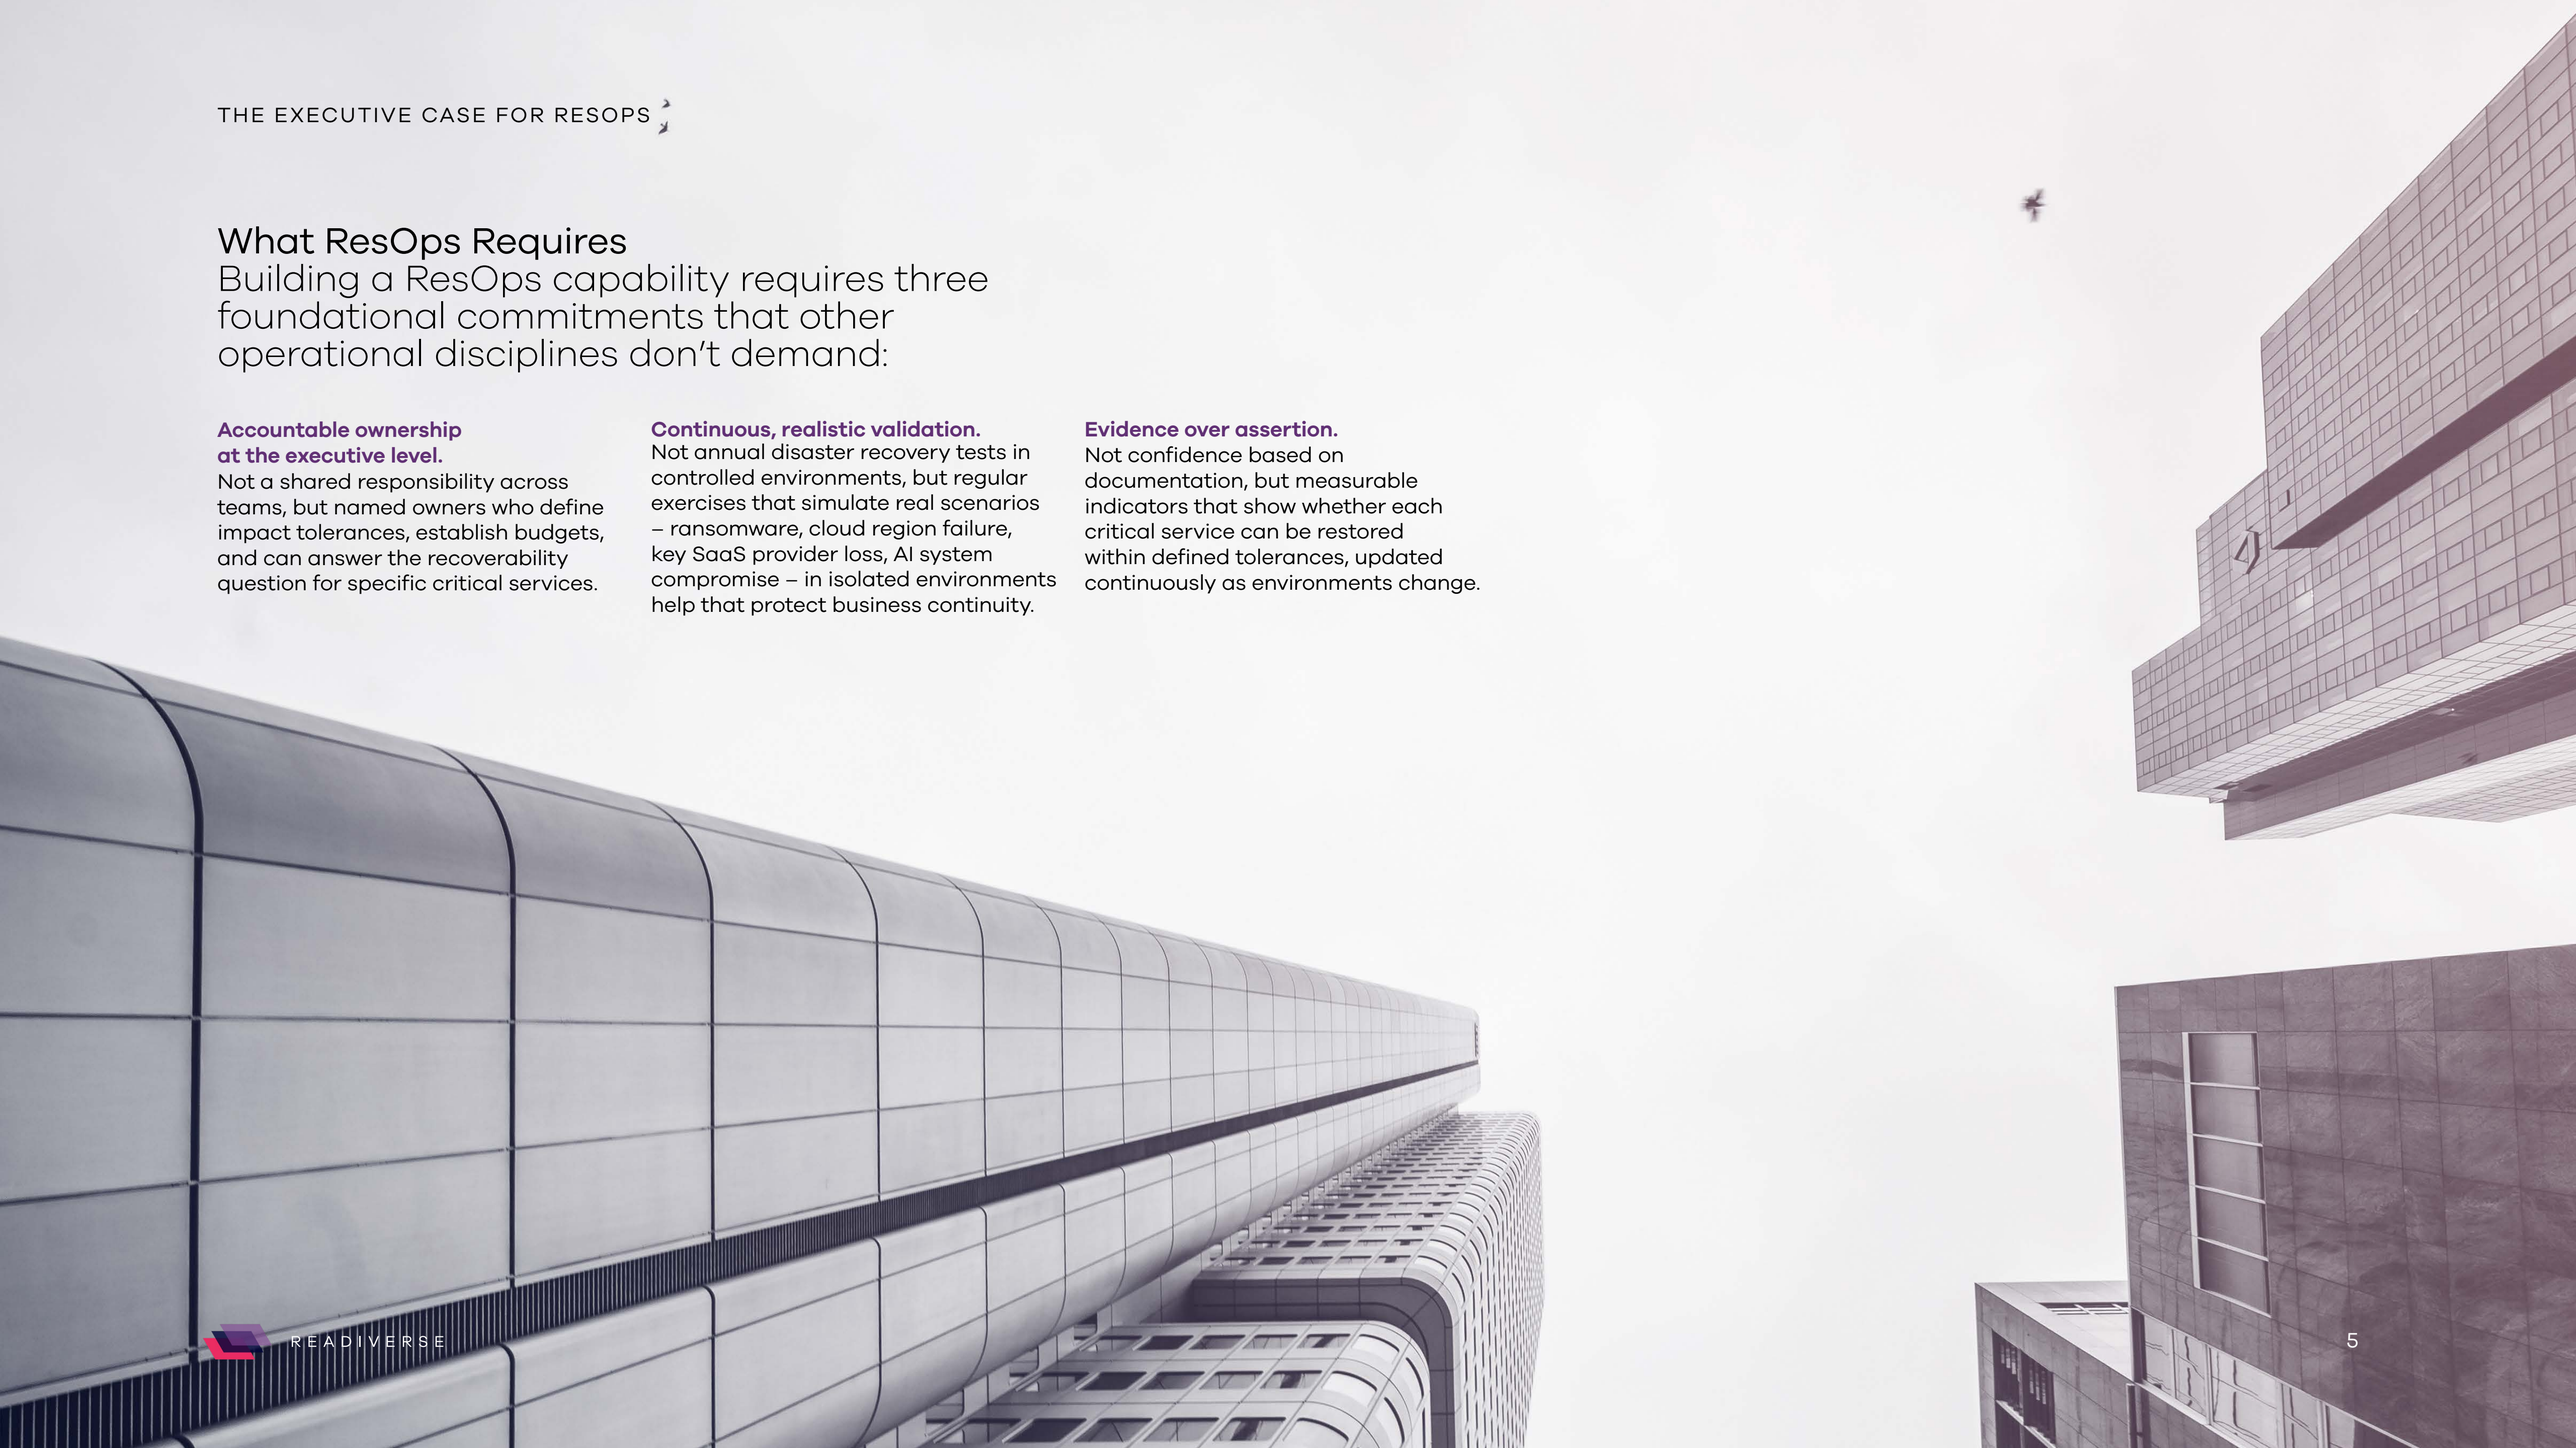

ResOps in Practice: The Three Imperatives

A complete ResOps model spans five domains: resilience governance, recovery planning, recovery architecture, recovery assurance, and resilience measurement.

For CIOs and CISOs standing up this capability, three imperatives cut across all five domains and define the foundational work that must come first. They are not a roadmap for the full ResOps journey – but without them, the rest of the model has nothing to stand on.

1. Define What Must Survive

The first act of ResOps leadership is not technical. It is definitional. Before investing in recovery capabilities, organizations must answer a question that is surprisingly difficult to answer well: Which services, if disrupted, would threaten the business at an existential or material level?

This is the critical services inventory: a deliberate mapping of the services, processes, and capabilities that define the enterprise's ability to serve customers, meet regulatory obligations, and generate revenue. Most organizations have versions of this list in various documents. ResOps makes it operational.

For each critical service, impact tolerances define the boundary between acceptable and unacceptable disruption:

Maximum tolerable downtime:
How long can the service be unavailable before the impact becomes unacceptable?

Maximum tolerable data loss:
What is the furthest-back recovery point the business can accept?

Minimum viable service:
What degraded operating mode is acceptable during recovery?

These tolerances are business decisions, not technical assumptions.

2. Design for Controlled Degradation

Once critical services and their tolerances are defined, the engineering objective is clear: Design each service so that when disruption occurs, the failure is contained, not cascading.

This is resilience in depth – multiple layers of protection applied across technology, process, and people, so that when one layer fails, the next absorbs the impact without total system failure.

Technology resilience:
Identity protection, data portability, service provider failover, infrastructure-as-code to enable clean rebuilds.

Process resilience:
Documented recovery procedures, pre-defined decision rights during incidents, communication protocols that work when tools and people are degraded.

People resilience:
Cross-training for critical functions, backup authority structures, capability to execute under pressure without key individuals.

The objective is predictable behavior under stress. When disruption occurs, the organization responds according to design – not improvisation.
The test of whether this pillar is in place is simple: If the three people who know how to recover a critical system are unavailable, does recovery still happen?

3. Prove It Continuously

Resilience that has not been tested is not resilience. It is a gamble.

ResOps treats validation as a continuous operating rhythm, not an annual event. The goal is to regularly prove that the organization can meet its impact tolerances through controlled exercises – designed to expose weaknesses without disrupting production operations.

A mature exercise portfolio includes scenarios at every level: executive tabletops that walk leadership through major disruption decisions, technical exercises where engineering teams step through recovery processes end-to-end, game days that validate failover and runbooks, red team exercises that simulate real attacks, and bi-annual live recovery drills that restore from clean, isolated recovery points.

The output of each exercise is not a report to file. It is a prioritized resilience backlog – specific improvements, owned by specific service teams, tracked through normal governance processes until resolved.

CXO Tool: Structuring Your ResOps Council

ResOps is not a function that can be delegated to a single team. It requires coordinated ownership across the enterprise, anchored by a governing body with executive mandate and cross-functional representation.

The ResOps Council is the mechanism that makes this coordination operational. It provides the governance structure, the accountability model, and the regular rhythm that transforms ResOps from a program into an operating model.

Council Structure: Three Tiers
Consider structuring your ResOps Council across three tiers, each with distinct responsibilities:

1. Strategic Tier Sets Direction and Accountability

Head of ResOps
Owns the single question: "Are we recoverable right now, for our most critical services?" Typically reports to CIO or CISO. Acts as the bridge between Cybersecurity, IT Ops, and Backup and Disaster Recovery.

CIO or Delegate
Confirms ResOps aligns with technology strategy, investment priorities, and enterprise risk posture.

CISO or Delegate
Validates alignment between resilience posture and security architecture, particularly around containment and isolation controls.

2. Design and Validation Tier Builds and Tests

ResOps Architect
Designs resilience architecture and reference patterns. Owns the integration model between security tools, backup platforms, disaster recovery tooling, and cloud providers.

ResOps Engineer
Automates the ResOps lifecycle: discovery, data collection, verification, testing, and evidence generation.

ResOps Analyst
Translates raw signals into a clear recoverability posture view by business service, by service-level agreement, by risk. During incidents, acts as the recovery navigator.

3. Business and Governance Tier Owns and Assures

Business Resilience Owners (3–5)
Represent specific critical services or business lines. Define what "recovered" means in business terms. Participate in exercises and sponsor remediation work.

Governance and Risk Lead
Translates regulatory expectations into concrete ResOps controls. Prevents new projects from going live without meeting minimum resilience criteria.

Council Operating Rhythm

The council is only as effective as the cadence it keeps. Establish a governance rhythm that tracks progress against defined impact tolerances and escalates systemic issues before they become incidents:

Cadence		Purpose
Monthly Council Review	→	Track resilience backlog progress, surface emerging gaps, unblock systemic issues, and review exercise results.
Quarterly Board Reporting	→	Cover impact tolerance attainment, key resilience risks, and the financial and regulatory implications of current gaps.
Quarterly Executive Tabletop	→	Walk leadership through a major disruption scenario – without touching production – to validate decision rights, escalation paths, and cross-functional coordination.
Bi-Annual Live Drill	→	Restore from clean, isolated recovery points to validate that impact tolerances are actually achievable, not just theoretically designed for.

Key principle

The ResOps Council's governance message should be consistent and cultural: Resilience is not a checklist. It is a property of how the business operates, designs, and learns.

CXO Tool: The Impact Tolerance Framework

Impact tolerances are the decisions that drive everything else in the ResOps model. They translate business risk appetite into the technical and operational requirements that engineering teams, backup architects, and recovery teams build against.

Setting them is not a technical exercise. It is a business conversation: one that requires CIO, CISO, and business leadership alignment on what the organization can and cannot absorb when disruption occurs.

1. Identify Your Critical Services

Begin with the services, processes, and capabilities that are non-negotiable. These are not determined by technical complexity or infrastructure investment. They are determined by business impact.

For each business unit or domain, answer: If this service were unavailable for 24 hours, what would happen to customers, revenue, regulatory standing, and the organization's ability to operate?

Tier 0 – Existential:
Service loss threatens the organization's ability to operate as a business (e.g., core transaction processing, identity and access systems, safety-critical operations).

Tier 1 – Critical:
Service loss causes material customer impact, significant revenue loss, or regulatory exposure within hours.

Tier 2 – Important:
Service loss is operationally disruptive but can be managed manually for 24–72 hours.

Tier 3 – Standard:
Standard IT services and internal productivity tools with lower urgency.

2. Define Tolerances for Each Critical Service

For Tier 0 and Tier 1 services, define three dimensions of tolerance. These become the measurable targets the ResOps model builds toward and validates against:

Maximum Tolerable Downtime
How long can this service be unavailable before the impact becomes unacceptable to customers, regulators, or the business?

Maximum Tolerable Data Loss
What is the furthest-back recovery point we can accept? What data loss volume or time range crosses a line we cannot accept?

Minimum Viable Service
What degraded operating mode is acceptable during recovery? Can we serve customers at reduced capacity while full recovery proceeds?

3. Close the Gap Between Tolerance and Reality

Once tolerances are defined, test them. The most common discovery organizations make when they first define formal impact tolerances is that their current recovery capabilities do not meet them – not because of poor engineering, but because tolerances were never formally defined and used to drive investment decisions.

The gap between defined tolerances and demonstrated recovery capability becomes the ResOps backlog: the prioritized set of improvements that the council tracks, funds, and governs until the organization can answer the recoverability question with evidence.

The governance principle

Policy ties funding to tested impact tolerances. Resilience investments should be prioritized by how far current recoverability falls short of defined business requirements – not by technical complexity or tool vendor roadmaps.

Where to Start: A 90-Day Readiness Path

ResOps is not built in a quarter, but the foundational steps that determine long-term success can begin immediately. The organizations that build effective ResOps capabilities typically share a common pattern in their early months.

Phase		Priority Actions
Days 1–30: Define	→	Convene an initial ResOps working group with CIO, CISO, and two or three business unit leaders. Identify your Tier 0 and Tier 1 critical services. Draft initial impact tolerances for your top five critical services.
Days 31–60: Map	→	Map dependencies for your top critical services across applications, data stores, infrastructure, identities, and third parties. Identify single points of failure. Assign Business Resilience Owners for each Tier 0 service.
Days 61–90: Test	→	Run your first executive tabletop exercise against a realistic scenario. Validate whether your current recovery capabilities match your defined impact tolerances. Generate your first ResOps backlog from the gaps you find.

The first tabletop exercise is often the most clarifying moment in a ResOps journey. It surfaces the gap between what leadership believes about recovery capabilities and what the technical reality actually is. That gap, made visible and measurable, becomes the foundation of the ResOps program.

The mindset shift

Hope is not a strategy when outages cascade in seconds and recovery windows shrink to minutes. The organizations that recover fastest are not those with the most tools – they are the ones that treated recoverability as an operational discipline before the incident arrived.

Measuring Resilience: Service Resilience Indicators

Executives need a consistent, measurable view of resilience health – one that surfaces risk before disruption occurs and tracks improvement over time. Service Resilience Indicators (SRIs) are the mechanism ResOps uses to make that possible.

SRIs translate the technical reality of recovery capability into business-relevant signals. They are evidence-backed, drawing on validation results and dependency mapping to show readiness, fragility, and where investment is required.

Recoverability score:
The confidence, based on tested evidence, that a critical service can be restored within its defined impact tolerances.

Data integrity score:
Validation that recovery points are clean, complete, and free from corruption or malware – the foundational input to Mean Time to Clean Recovery.

Dependency health:
The resilience posture of the services and systems that a critical service depends on.

Tolerance confidence:
The gap, if any, between defined impact tolerances and demonstrated recovery performance.

Anchoring the measurement framework is a metric the board can hold: **Mean Time to Clean Recovery (MTCR)**. Where RTO measures how fast a system is restored and RPO measures how recent the recovered data is, MTCR measures how long it takes to restore from data that is verifiably clean – making integrity a measurable recovery objective, not an assumption.

This matters because modern attackers do not simply disrupt operations; they target the backup and recovery infrastructure organizations depend on to come back. An organization that restores quickly from compromised data has not recovered; it has reinfected itself. MTCR makes that risk visible and measurable. It is the missing link between cybersecurity and business continuity, and it belongs alongside RTO and RPO as a core recovery metric reported at the board level.

These indicators provide a single dashboard that gives the ResOps Council, CIO, CISO, and board a consistent view of resilience health – updated continuously, not quarterly. When the dashboard shows a gap, the resilience backlog shows the path to closing it.

The goal is simple visibility: For any critical service, at any moment, the organization can answer whether it is recoverable and produce the evidence to prove it.

Conclusion

The disruption landscape facing enterprises today is not going to simplify. AI accelerates the speed of both innovation and compromise. Cloud complexity continues to grow. Regulators are raising the bar on operational resilience. Boards are asking harder questions about recoverability than they were two years ago.

ResOps is the response to that reality – not as a new tool or a new team, but as a new operating model. One that treats resilience as a measurable, governed, continuously validated property of the business. One that replaces hope with evidence.

The organizations that build this capability now will be better positioned to recover faster, protect customer trust more effectively, and face their next disruption with evidence rather than assumption. That advantage is real, it is durable, and it begins with a decision to treat recoverability as an operational discipline, not an afterthought.

Related Resources on Readiverse

ResOps Introduction Paper
"ResOps: The Future of Resilient Business in the Era of AI"

Foundational overview of the ResOps discipline for organizations beginning their resilience operations journey.

ResOps Implementation Framework
"ResOps in Practice"

Detailed practitioner guidance across the five ResOps domains, with implementation actions for ResOps Architects and Engineers. Note, this is intended to be a starting point. We know that there are many other disciplines that would fit within the core ResOps domains, and we look forward to getting feedback and input from the broader development community to further flesh this out.

From Minimum Viability to Operational Resilience

Co-authored with Deloitte. Goes beyond recovery readiness to the harder question: how do you prove it? Covers the ResOps maturity model, impact tolerances, service resilience indicators, and Mean-Time-to-Clean-Recovery — the metrics that make resilience a continuous discipline, not a periodic exercise.

Available at readiverse.com/resops