O'REILLY®
Report

Compliments of

Commvault®

ResOps

An Executive Guide

Sustainable Resilience for Cybersecurity and Operations Teams

Jason Cray, Ben Herzberg & Michael Thelander

ResOps: An Executive Guide

Sustainable Resilience for Cybersecurity and Operations Teams

Jason Cray, Ben Herzberg, and Michael Thelander

O'REILLY®

ResOps: An Executive Guide
by Jason Cray, Ben Herzberg, and Michael Thelander

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: Jill Leonard
Development Editor: Andy Kwan
Production Editor: Destiny Baitinger
Copyeditor: Vanessa Moore
Proofreader: Krsta Technology Solutions
Cover Designer: Susan Brown
Interior Designer: David Futato
Interior Illustrator: Kate Dullea

August 2026: First Edition

Revision History for the First Edition
2026-08-26: First Release

This work is part of a collaboration between O'Reilly and Commvault. See our statement of editorial independence.

979-8-295-90057-0

[LSI]

Table of Contents

Introduction: When Assumptions Crumble

In September 2025, ransomware forced Jaguar Land Rover (JLR) to halt global manufacturing. Assembly lines stopped. Supply chains froze. The Cyber Monitoring Centre estimated the cost to the UK economy at approximately £1.9 billion (roughly US $2.4 billion). JLR recorded its lowest monthly production output in 73 years. Whatever recovery capabilities existed, they were not enough. The operational path back proved far harder than anyone had planned for.

JLR is not alone. A 2023 report showed that 90% of organizations experienced at least one major cyberattack, with 56% of all downtime incidents now directly attributed to ransomware and phishing. The average recovery time from a ransomware incident was 75 days in 2024. The global average cost of a data breach has reached $4.99 million as of 2026; an insider security incident, $19.5 million. It's safe to say every organization behind these numbers had funded security, but that very few had operationalized the processes to achieve proven recovery.

While recent events and actuarial analysis are important, this publication is not about the attacks that happened last year or last month. It's about being ready for tomorrow's more potent and likely more successful AI-enabled attacks. It's about an approach to cyber resilience that allows CISOs and CIOs to answer a new set of questions.

If we were hit by a ransomware attack tonight—or if our cloud provider went dark, or a critical system failed—could we recover? How long would it take? And how do we know? Do we have proof?

Not only are these the questions their boards increasingly ask, but also the ones regulators now demand. They're the ones that no one in their organization seems to be able to answer with confidence.

Most executives, when asked whether their organization could survive a major disruption, answer with plans and policies. They give assurances that backups are running and disaster recovery (DR) systems are in place.

But plans are not proof. Policies are not performance. And the gap between what organizations believe about their resilience and what is actually true has never been wider or more consequential.

No one argues that the threat landscape has fundamentally changed in terms of impact, targets, and frequency. Ransomware attacks now routinely target backup infrastructure itself, and AI-enabled systems create new data and new attack surfaces at a pace that utterly outstrips our ability to protect them. And regulators—from the EU's Network and Information Security Directive 2 (NIS2) and Cyber Resilience Act (CRA) to auditors of the Cybersecurity Framework (CSF 2.0) from the National Institute of Standards and Technology (NIST)—now demand operational resilience measured not by documentation, but by demonstrated outcomes.

At the same time, the approaches organizations traditionally rely upon have not kept pace. DR focuses on backups. Cybersecurity focuses on detection and prevention. Business continuity focuses on risk registers and manual runbooks. Each discipline does its job within its own silo. None of them answers the questions that matter: Can we recover? And can we prove it?

The AI Paradox

In April 2026, Anthropic announced Project Glasswing, a coalition with AWS, Google, Microsoft, and more than fifty other organizations. The project was built around Claude Mythos Preview, an unreleased frontier AI model that can autonomously discover—and develop exploits for—zero-day vulnerabilities across major operating systems, web browsers, infrastructure, and configurations. According to early reports, the model identified tens of thousands of critical vulnerabilities that traditional tools have missed for years.

Anthropic's blog suggests that only a fraction of vulnerabilities found by Mythos have been patched. Even though there are many open questions about the exploitability of these flaws, adding thousands of new findings to an already overloaded remediation process does not make organizations more resilient. It makes them more aware of how exposed they are, without necessarily giving them the capacity to do anything about it.

We call this the AI Paradox: the same AI capability that accelerates defense also accelerates offense, but it does not do so at the same rate. Attackers are fast, operating without the constraints of legacy systems, technical debt, or customer patch cycles. For them, AI compresses the window between vulnerability discovery and active exploitation from weeks to negative seven days, meaning a fix for the attack won't likely be developed until a week after the event. For organizations without the operational discipline to prioritize, contain, and recover with confidence, faster threat detection becomes faster threat saturation.

AI-enabled vulnerability discovery raises the urgency of resilience. It does not replace it. Finding the vulnerability is step one. Knowing which critical services it threatens, containing the blast radius, recovering cleanly, and proving that recovery works, that is a different discipline entirely. It is not only the discipline most organizations have not yet built, but an area ResOps (Resilience Operations) specifically aims to address.

The defenders that will benefit most from Mythos and similar frontier AI models are those that already have the operational foundation to act on what the tools surface: a defined minimum viable company (detailed in Chapter 3), mapped dependencies, tested recovery playbooks, and a resilience backlog that can absorb new findings and prioritize them against known impact tolerances. For every other organization that assumes they can keep the threats out, the deluge of exploits will find them unprepared and searching for a resilience strategy.

The Questions That Matter Now

For decades, the majority of security and IT operations playbooks were built on a comforting assumption: if we prevent enough attacks, detect them fast enough, and maintain sufficient backup copies, we will be resilient.

That assumption no longer holds.

Modern enterprises operate in tightly coupled, highly automated ecosystems where a single failure can cascade instantly across systems, clouds, SaaS applications, and third-party dependencies. The blast radius of any disruption—technical, cyber, or operational—is larger than most organizations realize and faster-moving than any manual process can contain.

At the same time, adversaries have adapted. They have learned that the fastest way to maximize leverage is not to defeat your security controls, but to corrupt your recovery capabilities. When backup infrastructure is compromised before ransomware detonates, the organization is not just attacked. It is left with no reliable path back.

Against this backdrop, *three questions define the modern resilience challenge*—and they are the organizing framework for this guide:

Can we recover?
: Not theoretically. Not just per the plan. But end-to-end, under real-world pressure, including all the dependencies, cloud services, and third-party systems that our critical operations *actually* require.

How long will it take?
: Not just the recovery time objective (RTO) documented in a plan. The actual, tested, proven time to restore a set of minimum viable operations our business needs to function cleanly.

Can we prove it?
: Not to ourselves. But to our board, our regulators, our customers, and our cyber insurance carrier. Can we prove it with documented evidence from realistic exercises, instead of just assurances from the teams responsible for the systems themselves?

These are the questions that ResOps is designed to answer. This publication explains what ResOps is, why it matters, how it works, and how to start your organization on a path that allows you to answer these questions with confidence.

CHAPTER 1
The Resilience Gap

> *To be able to ask a good question is already half of the solution to a problem.*
>
> —J. Robert Oppenheimer

For years we built resilience on the hope that prevention and clean backups would carry the business through its worst day. They don't, not on their own. Survival turns on a harder question: when the disruption comes, will we be ready?

Most organizations cannot answer that question with confidence. Not because they lack investment, but because they lack evidence.

This evidence gap runs across the entire industry and highlights several additional gaps: between what we *think* will happen and what we now *know* will happen; between our budgets and our expectations; between the vastness of our attack surface and our certainty of where the edges are. These are all gaps that affect our understanding of our level of resilience, of our resilience posture. NIST, in their advisory SP 800-160 Vol. 2 Rev. 1 on "Developing Cyber-Resilient Systems," defines cyber resiliency engineering as the process designed to:

> architect, design, develop, implement, maintain, and sustain the trustworthiness of systems with the capability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises that use or are enabled by cyber resources.

Failures of resilience are not hypothetical. They emerge predictably from a set of failure patterns that most enterprises share, regardless of industry or size:

Invisible dependencies
: Critical services rarely fail alone. They fail because of cascading dependencies that no one fully mapped. A SaaS application depends on a cloud identity provider. A payment system relies on a third-party API. A database shares infrastructure with a dozen other services. When one node fails, the ripple is immediate and often invisible until it is too late. (This is covered in detail in Chapter 3.)

The backup-as-insurance mentality
: Most organizations have backups. Fewer know whether those backups are clean, recent, or complete enough to actually restore operations. Fewer still have tested whether the restoration process works end-to-end, under pressure, at the speed the business requires.

Siloed accountability
: Security teams own detection. IT operations teams own infrastructure. Business continuity teams own plans. Nobody owns the end-to-end answer to whether the organization can actually recover from a major disruption. That gap does not belong to one team, which means in practice it belongs to no one.

AI and automation amplification
: As organizations deploy generative AI and agentic systems, they create new data sprawl, new dependencies, and new failure modes faster than governance and recovery capabilities can adapt. AI-driven processes that touch customer service, financial reporting, or core operations become critical services overnight, often without recovery planning in place.

Why Current Approaches Miss the Mark

The disciplines most organizations rely on today were each designed to solve a different problem. These are outlined in Table 1-1. Each is necessary. None, by itself, was designed to answer the executive resilience question end-to-end.

Table 1-1. Comparing ResOps to business risk and continuity initiatives

Discipline	Primarily answers	Common focus	Typical output
Backup and DR	Do we have recoverable copies and recovery procedures?	Systems, data, infrastructure recovery	Backup jobs, DR plans, restore procedures
Cybersecurity	Are we monitoring for, detecting, preventing, and containing various threats?	Controls, monitoring, incident response	Alerts, incidents, containment actions, reports
Business continuity	How will the business operate through disruption?	Business processes, communications, workarounds	Business impact analyses (BIAs), continuity plans, crisis playbooks
ResOps	Can critical services recover within business-defined impact tolerances?	End-to-end service recoverability	Evidence-backed service resilience indicators (SRIs) and resilience backlog

The gap between these disciplines is where resilience failures live. It is the space between "we have copies" and "we can recover." Between "we detected the attack" and "we restored clean operations." Between "we have a plan" and "we have proof."

That gap is organizational, not technical. It cannot be closed by any single team working within its existing mandate. It requires a shared operating discipline, a common language, and a common standard of evidence.

In ResOps, recovery means more than getting systems to power on. It means restoring critical services cleanly, completely, and within business-defined impact tolerances.

Service Resilience Indicators

Traditional resilience metrics—uptime percentages, RTOs, recovery point objectives, and mean time to recovery—are useful but incomplete. Some describe historical performance. Others describe planning assumptions. Put all four together and they still fall short of the one proof that matters: that a critical service can be recovered cleanly, end-to-end, within the tolerance the business has defined.

Service resilience indicators (SRIs) are a different category of measurement. Where traditional metrics often describe objectives or past performance, SRIs measure evidence-backed preparedness. For each critical service, SRIs answer four questions (covered in detail in Chapter 4). The answer to these questions establishes our evidence-based confidence in recovering the service within its business-defined impact tolerance.

SRIs transform resilience from a belief into a posture, one that can be tracked over time, reported to the board, and used to drive investment decisions. Consider the difference between these two statements from the same organization:

- We believe we can recover the payments service in four hours.
- Last quarter, we restored the payments service in 3.2 hours from a verified clean recovery point, against a four-hour, business-defined impact tolerance. The service is within tolerance.

The 3.2 hours to a verified clean recovery has a name: mean time to clean recovery, or MTCR (covered in Chapter 3). The difference is not optimism versus pessimism. It is belief versus evidence. Boards, regulators, and insurers can act on the second statement. They cannot govern with the first.

SRIs are the operational output of a functioning ResOps program. Chapter 4 examines them in full, including how they are calculated, reported, and used to make investment decisions.

CHAPTER 2
What ResOps Is

Any fool can know. The point is to understand.
—Albert Einstein

You cannot put ResOps to work until you can say precisely what it is, where its edges are, and where it sits relative to everything around it. That precision separates a discipline you operationalize from one you only discuss.

An Operating Discipline, Not a Project

ResOps is the operational discipline that unites security, infrastructure, IT operations, business continuity, and business owners around critical services, resilient design, and continuous validation, so organizations can withstand disruption, recover within business-defined impact tolerances, and prove that capability with evidence.

That definition carries four characteristics that distinguish ResOps from every other resilience-related practice:

Unites teams
: ResOps is inherently cross-functional. It spans security, infrastructure, IT operations, business continuity, business leadership, and the teams responsible for AI and data systems. No single team owns resilience end-to-end. ResOps connects distributed accountability into a shared operating model.

Centers on critical services
: Not all systems are equal. ResOps begins with a disciplined definition of which services are truly critical: what we call the minimum viable company (MVC), which is the smallest version of the organization that can deliver its core value, serve customers, generate revenue, and meet its most urgent legal, regulatory, safety, or mission obligations.

Relies on continuous validation
: Resilience is not a state you achieve. It is a posture you establish and continuously verify through repetitive stress testing. ResOps replaces reliance on the annual DR tests with an ongoing program of exercises, real-world simulations, clean recovery validation, and evidence collection.

Produces evidence, not assurances
: The output of a ResOps program is not a policy document. It is a set of SRIs: measurable, auditable indicators, backed by evidence, that show whether critical services can recover within business-defined impact tolerances.

This is a fundamentally different mandate than any of the existing disciplines. Backup and DR teams are accountable for copies and recovery procedures. Security teams are accountable for detection and response. Business continuity teams are accountable for plans and crisis coordination. ResOps creates accountability for outcomes: the actual ability of the organization to withstand and recover from a bad day.

What ResOps Is Not

Precision matters because ResOps intersects with existing practices:

Not a product
: ResOps cannot be purchased. Technology platforms, including data protection and cyber recovery tools, can enable and accelerate a ResOps program. The discipline itself is organizational.

Not DR modernization
: DR is a component of ResOps, but ResOps goes further. It validates end-to-end recoverability, including dependencies, rebuild pathways, and human readiness, not just whether backup jobs are being completed.

Not a compliance exercise
: Regulatory frameworks increasingly require a demonstrated operational resilience. A well-functioning ResOps program produces the evidence those frameworks demand, but compliance is a byproduct, not the objective.

Not a one-time project
: ResOps is a continuous operating rhythm. Organizations do not complete ResOps any more than they complete financial planning or security operations.

If any discipline is centered on a point in time or on a singular outcome, it's not ResOps.

Where ResOps Fits

ResOps does not displace existing disciplines. It unites them under a common operating and accountability framework:

For cybersecurity
: ResOps assumes that even the best defenses will sometimes fail, and validates the ability to recover when they do. It ensures the security team's mandate expands from "detect and contain" to "contain and recover."

For IT operations
: ResOps gives infrastructure, architecture, and operations teams a structured framework for designing, testing, and improving recovery capabilities, moving recovery measures from tribal knowledge to documented, proven playbooks that work under real pressure.

For business continuity and governance, risk, and compliance (GRC)
: ResOps turns continuity planning and regulatory expectations into proven SRIs. It does not replace business continuity management (BCM) or GRC; it gives teams stronger proof that plans can be executed and tolerances can be met.

For business leadership
: ResOps translates technical recovery capabilities into business language: SRIs, impact tolerances, and MVC definitions that executives and boards can understand, govern, and fund.

What matters is not where ResOps lives on the org chart, but that it has a named owner, executive sponsorship, the mandate to produce evidence, and funding to support it.

CHAPTER 3
How ResOps Works

> *What we can control is our performance and our execution, and that's what we're going to focus on.*
>
> —Bill Belichick

ResOps provides executives with three things they often lack: clarity that the business can survive, visibility into current resilience posture, and evidence to prove that posture to anyone who asks.

It operates through three phases, in order: Visibility, Control, and Confidence. The order is how you read the model, not a sequence you finish one stage at a time. Governance, exercises, and improvement all run continuously.

Before Anything Else: Governance

ResOps depends on governance. Because resilience crosses organizational boundaries, the authority to set priorities, assign accountability, and make investment decisions must be established before the operating model can work. Governance provides the structure for making those decisions across functions rather than leaving each team to optimize for its own priorities.

Table 3-1 shows how that foundation carries through the three phases of a ResOps program.

Table 3-1. Three phases of a ResOps program, the executive questions each phase seeks to answer, and eight steps that get you there

ResOps phases	Questions to answer	ResOps Steps
Visibility: governance and planning	• What is our goal? • What must we protect? • What can we tolerate?	Step 1: Set the charter, impact tolerances, and funding Step 2: Define MVC, map dependencies
Control: design, architecture, and function	• Can we contain the damage? • Can we bring the service back clean? • Who has authority when it breaks?	Step 3: Contain the blast radius and protect recovery paths Step 4: Engineer assured, clean recovery Step 5: Integrate crisis operations and decision authorities
Confidence: reporting, validation, and improvement	• Can we prove it works? • Are we within tolerance on current evidence? • Are we getting better?	Step 6: Validate through realistic exercises Step 7: Measure with SRIs and RPS Step 8: Feed lessons back into the cycle and drive continuous improvement

Phase 1: Visibility

The Visibility phase requires an organization to know three things: what services the business cannot function without, how much disruption each can absorb before the harm is unacceptable, and whether there is evidence to verify those services are within those limits. Most organizations think they know the answer to the first question, few have formally answered the second, and even fewer can answer the third with evidence.

Step 1: Set the charter, impact tolerances, and funding

The charter establishes the authority and accountability that ResOps requires. ResOps needs a named leader with the authority to set priorities across organizational boundaries and report outcomes to the board. In most organizations, the CISO owns resilience outcomes as resilience moves under the security mandate, while the CIO owns operational recovery. The charter should establish a governance rhythm:

- Policy that prioritizes funding based on gaps against tested impact tolerances
- Monthly reviews that surface gaps requiring investment decisions
- Quarterly board reporting on resilience posture, customer impact, regulatory exposure, and financial risk

Without a charter, exercises produce findings that go nowhere. With it, every gap has an owner, a timeline, and a business consequence if it is not closed. The charter also sets the funding policy that everything downstream depends on: not a one-time budget line, but a standing rule that ties investment to the size of the gap between current recoverability and what the business has said it can tolerate. That is what turns the tolerance setting work in Step 2 into decisions instead of documentation.

Step 2: Define MVC, map dependencies

If we lost everything tonight, what comes back first?

That question, formally answered by technology leaders, business owners, and legal and compliance teams together, produces the MVC: the smallest version of your organization that can serve customers, generate revenue, and meet its most legal, regulatory, safety, or mission obligations. Every service in that definition needs a named business owner, and every service carries a business-defined impact tolerance: how much downtime, data loss, degraded performance, incorrect output, or manual workaround the business can absorb before unacceptable harm occurs.

But disruption has to be defined correctly first. Most organizations define it as downtime. That misses most of what actually hurts. A few examples:

- A payment platform processing at 20% of normal volume is not *running*.
- A reporting system that's online but producing bad data is not *functioning*. The dashboard says green, but the organization is already bleeding.

Disruption is defined by business outcomes, not system status. And it rarely starts at the service level. A Tier 0 service can appear fully operational while a dependency underneath it is already failing. By the time the disruption reaches the surface, the business impact has already begun. Impact tolerances must be built around the full picture: service behavior, dependency health, and business output together. Get this definition wrong and every decision that follows is calibrated against the wrong target.

Tolerances are expressed in business terms: revenue impact per hour, transaction thresholds, customer service levels, regulatory obligations—not uptime percentages.

Many organizations express MVC priorities in tiers:

Tier 0
: Disruption of even a few hours creates unacceptable harm

Tier 1
: Recoverable within one to two days

Tier 2
: Extended downtime is tolerable

The thresholds vary. The discipline is the prioritization. Revisit it regularly. A generative AI system that was a pilot last year may be Tier 0 today.

Once critical services and tolerances are defined, the next task is to map what each one depends on. Every critical service depends on infrastructure, providers and relationships over which it does not have full control, as well as on other internal applications and services, some of which may be at a lower priority. Dependency mapping makes that exposure visible before an incident does.

The findings are usually uncomfortable. Services assumed to be self-contained share infrastructure with others. A single cloud identity provider can take down applications across the enterprise when it fails. These are the cascades that a service-level status indicator misses until the business impact has already begun. Mapping the full dependency chain tells you where recovery could fail and where resilience investment will matter most.

Phase 2: Control

The Visibility phase tells you where you stand. The Control phase is about how you design the organization to remain functional when things go wrong. The questions here are not technical; rather, they are organizational:

- Who decides?
- Who acts?
- What has been preapproved?
- How does the business absorb a blow without losing the ability to function?

Step 3: Contain the blast radius and protect recovery paths

When something breaks, the decision window to act is minutes. Every minute spent determining who has authority to act is a minute the disruption continues to spread.

Containment is the organizational capability to limit damage before it cascades: isolating systems, disabling integrations, shifting to reduced-capacity operations. These are business decisions with business consequences, and they cannot be delegated in the moment to just anyone who is available. Step 5 covers how that authority is predetermined.

Containment has to protect what recovery will need. Isolating a compromised system can stop the blast radius from spreading, but if that action destroys or cuts off recovery points, logs, or forensic evidence, it may also eliminate a clean path back. Blast-radius decisions and recovery-path decisions have to be made together, or they undermine each other.

Step 4: Engineer assured, clean recovery

Recovery plans are usually written under the assumption that people executing them will have current information, familiar systems, and time to think. Real disruptions provide none of those.

Assured recovery is the discipline of designing for real conditions, recovering clean when detection might have failed. This includes having named owners for every critical service; documented procedures tested under real-world chaos, by the teams who will actually execute them, not written for someone else to follow; and people who have practiced enough to perform under pressure, not just read a playbook.

For critical services, engineering recovery requires protected recovery points, isolated recovery environments, trusted rebuild paths, validated data integrity, and documented sequencing for identity, infrastructure, applications, and dependencies. These are not optional refinements. They are the difference between recovery that works and recovery that reintroduces the problem.

It also means recovering clean and within tolerance. Speed without integrity is not recovery. The definition of success is not "systems are back on." It is "systems are back on, data is verified, dependencies are confirmed healthy, and the service is within its defined impact tolerance."

Step 5: Integrate crisis operations and decision authorities

Containment and recovery only work if the organization knows in advance who makes decisions and who will act. Authority for containment actions must be preapproved: defined ahead of time, with clear conditions and clear acceptance of the trade-offs involved. In a real disruption, the response that says "we need to escalate" is not an answer. It is a failure mode.

Most organizations run normally through monitoring teams, engineering groups, and routine processes. When disruption crosses a threshold, a different mode is required: integrated command authority, cross-functional coordination, and a shared real-time picture of what is happening and what has been decided. If your team has never practiced that shift, they will improvise it under the worst possible conditions.

ResOps closes that gap through predefined triggers, clear command structures, cross-trained teams that do not depend on specific individuals, and joint procedures tested by the people who will actually use them. A recovery that works once simply proves a scenario. A recovery that works repeatedly, under different conditions and with different people, proves capability. That capability is not a property of a system; rather, it is a property of an organization that has practiced until the response is predictable regardless of who is in the room.

Phase 3: Confidence

The Control phase focuses on designing your systems for survival. The Confidence phase communicates success, building on proof that the design works and supports the organization's accountability to keep proving it.

Step 6: Validate through realistic exercises

An untested recovery plan is a hypothesis. Exercises are how you turn that into evidence. ResOps testing is safe, continuous, and scoped. It is not an annual high-risk event. It is a regular program of exercises calibrated to the maturity of the program and the criticality of the services being tested. The exercise portfolio can include executive tabletops, technical simulations, cleanroom restores, limited failover tests, and third-party dependency exercises.

Every exercise produces two things: evidence that updates SRIs and posture across the full service and dependency picture, and findings that are entered into a resilience backlog that the team uses for ongoing prioritization. The score tells you where you stand against your tolerances. The findings tell you what to fix.

Step 7: Measure with SRIs and RPS

There are two measures that make resilience verifiable, not just felt: SRIs and resilience posture score (RPS).

SRIs are the readings, whether each critical service is provable within tolerance, based on the latest evidence across the service and dependency picture. They cover four things: whether the service can recover, whether its data would restore clean, where recovery would break, and how current the evidence is.

The recoverability reading worth naming is mean time to clean recovery (MTCR). Traditional recovery metrics understate cyber problems because they focus on restoration targets rather than the full path back to a trusted state. In a cyber event, the organization must identify the compromise, determine the blast radius, validate a clean recovery point, and confirm that restored data and dependencies can be trusted. MTCR captures the time from disruption onset to a critical service running again at a verified clean state. If it exceeds the business-defined impact tolerance, the service is not resilient, no matter how fast the final restore step appears.

SRIs roll up into the RPS, a single per-service score on three bands. Read it as a status:

Green
: Within tolerance and within cadence

Yellow
: A dependency is degrading, a cadence is slipping, or a tolerance window is narrowing (Yellow is when you act)

Red
: You cannot currently prove this service can recover within its defined tolerance (Red is a material risk disclosure, not a technical finding)

Together, SRIs and RPS turn "we think we're ready" into a score the business can act on. Chapter 4 goes into additional details on SRIs and the SRI dashboard.

Step 8: Feed lessons back into the cycle and drive continuous improvement

Though this step is listed last, it is not the end of the process. It is the central pivot point that transforms ResOps from a program into an operating discipline.

SRI posture data makes that loop legible at the executive level. Boards and regulators do not need to understand the technical details of a recovery exercise. They need to see which services are within tolerance or which are not, what the financial and regulatory exposure is, and whether the organization is getting better or worse over time. Monthly council reviews surface gaps that require cross-functional decisions about investments. Quarterly board reporting translates SRI trend data into the business language of tolerance attainment and risk trajectory.

Every gap identified through testing or monitoring enters the resilience backlog: a prioritized list tied directly to the enterprise risk register, remediated in order of how far current recoverability falls short of defined impact tolerances—not by technical severity, but by business impact.

As Figure 3-1 shows, there are five stages of the ResOps discipline:

Discover
: Find and classify sensitive data across your estate.

Protect
: Establish the right protection policies and enforce them, instantly flagging drift.

Detect
: Continuously hunt for anomalies and threats, in real time, as they appear in protected data.

Recover
: Rapidly recover only what's clean and assured, at any scale, and at machine speed.

Restore
: Verify data is threat-free before it goes back to production, so resilience runs on a trusted foundation.

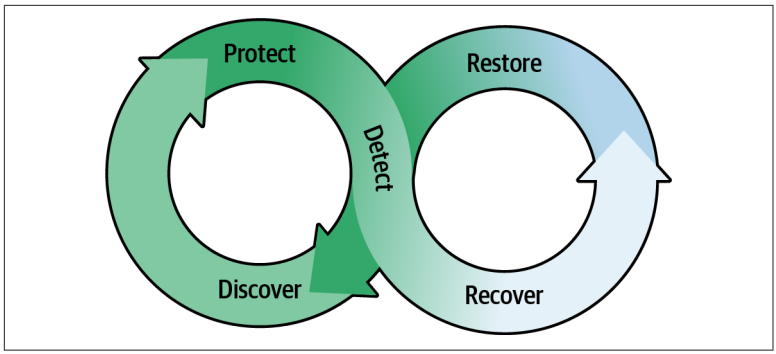

Figure 3-1. ResOps disciplines can maintain continuous resilience as critical data cycles through five stages from discovery to restoration

These steps create a process that, when integrated with both operations and security, enables resiliency to transform from an aspiration to a continually reinforced reality. In this process, tolerances are defined. Posture is monitored across the full service and dependency chain. Gaps are found, logged, fixed, and closed. Evidence accumulates. The board sees progress and confidence over time, not just the current state.

Resilience is not something you achieve. It is something you continuously prove.

CHAPTER 4
Proving It Works

What can be asserted without evidence can be dismissed without evidence.
—Christopher Hitchens

Chapter 3 described service resilience indicators (SRIs) as the live instrument of a mature ResOps program. Resilience that cannot be shown in numbers an executive trusts is not being managed; it is being hoped over. This chapter turns the instrument into control.

Boards are no longer satisfied with assurances. Research from Gartner in early 2026 titled *Cybersecurity Trend: AI and Cyber Resilience Redefine the CISO's Remit* made the shift explicit: cyber protection alone no longer meets executive expectations, and board directors view cyber risk as a direct threat to them personally, as well as to shareholder value. The conclusion for CISOs and CIOs was clear. The mandate has moved from prevention to resilience: minimizing business harm when attacks succeed, not just trying to stop them.

That shift requires a different kind of measurement, one that SRIs are well built for.

Why Traditional Metrics Fail Executives

Traditional resilience metrics are useful but incomplete. Some, like uptime percentages and mean time to recovery, describe past performance. Others, like recovery time objectives (RTOs) and recovery point objectives (RPOs), describe planning targets.

None of them, in isolation, proves that a critical service can be restored from a verified clean recovery point, with dependencies intact, within the tolerance the business has defined.

The sharpest version of this problem is RTO. An RTO tells you what you are aiming for. It does not tell you whether you will achieve it. A board member cannot evaluate "Our RTO for the payments service is four hours." They can evaluate "Last quarter we exercised recovery of our payments service, restored from a verified clean recovery point in 3.2 hours, against a four-hour, business-defined impact tolerance. The service is within tolerance." The difference between those two statements is not technical precision. It is evidence. The first is an assertion. The second is governable. SRIs are how you produce it.

What Are SRIs?

SRIs are evidence-backed indicators that show whether critical services can recover cleanly, end-to-end, within business-defined impact tolerances. They translate technical recovery evidence into a form executives, boards, regulators, and insurers can understand and govern.

An SRI is not a raw metric. Each is an evidence-backed indicator tied to a service's impact tolerance. Together, drawn from recovery exercises, clean recovery validation, dependency health, cadence compliance, and unresolved gaps, the SRI set for a critical service addresses four key points:

Recoverability
: Has this service been successfully recovered in a realistic exercise, within its defined impact tolerance, within the required cadence?

Data integrity
: Was the service restored from a verified clean recovery point, with data integrity confirmed? This is what MTCR measures: the full span back to trustworthy operation, not just the speed of the final restore step. A service that comes back fast from a compromised state has not recovered. The SRI counts the recovery only when the recovered state can be trusted.

Tolerance confidence
: Did recovery meet the business-defined impact tolerance? How current and credible is the evidence that this service holds inside its limits?

Fragility
: Even when recovery works, is the service quietly weak? Are its critical dependencies mapped, monitored, and in the exercise cadence? A service can recover fast and still be fragile if the dependencies it relies upon have never been tested.

SRIs are maintained per service, aggregated by tier, and tracked over time. They provide a posture view that executives can govern and use to make investment decisions.

The Executive Dashboard

A ResOps SRI dashboard presents resilience posture in terms boards can act on (see Table 4-1). The model is simple by design:

Green
: The service has been exercised within cadence. Recovery has been completed within the business-defined impact tolerance. Clean recovery from a verified recovery point has been confirmed. Critical dependencies are mapped, tested, and not carrying unresolved recovery gaps.

Yellow
: The service is approaching the edge of its exercise cadence, a dependency is degrading, or a recent exercise revealed a gap currently under remediation. Yellow is an early warning, not a failure. It is the signal that demands action before Red arrives.

Red
: The service has not been successfully exercised within tolerance, or a known gap places recovery capability at risk. The organization cannot currently prove it can recover this service within its defined tolerance.

Table 4-1. ResOps uses an SRI dashboard to provide executives and team leaders with at-a-glance insight into the resilience of their critical services

Service	Tier	Impact tolerance	Last exercise	Recovery result	Dependency status	Resilience posture
Payments	0	4 hours	Q2	3.2 hours	No material dependency gap	Green
Identity and access	0	2 hours	Q2	Not achieved	MFA dependency gap	Red
Customer CRM	1	24 hours	Q1	18 hours	SaaS export risk	Yellow

The dashboard is a risk disclosure tool. When a Tier 0 service shows Red, that is a board-level conversation that calls for determination of the specific gap, the remediation plan, the financial and regulatory exposure while it stays open, and who owns closing it. The board cannot meet its governance obligation on resilience without that conversation. The SRI dashboard is what makes it possible.

How SRIs Drive Investment, Risk Visibility, and Board Accountability

The practical power of SRIs is that they connect recovery capability directly to business decisions:

Investment decisions
: A resilience backlog prioritized by SRI impact gives CIOs and CISOs a defensible, evidence-based case for where to spend. Not "we need more budget for DR," but "three of our Tier 0 services are Yellow or Red, here are the specific gaps, and this is the remediation cost versus the business exposure if we do not close them." That is a conversation a CFO can engage with.

Risk visibility
: SRI trend data shows whether the organization's resilience posture is improving, degrading, or holding steady. That trajectory is a direct input to enterprise risk management, regulatory reporting, and cyber insurance negotiations. The cyber insurance market is increasingly moving toward evidence of recoverability, tested backups, and operational resilience rather than questionnaire-based assertions alone. Organizations that can show consistent SRI improvement over time enter renewal discussions with stronger evidence than organizations relying on static policy documents.

Board accountability
: Quarterly SRI reporting gives boards the ability to fulfill their governance obligation on resilience, not by reviewing technical documentation, but by assessing a business-level view of recovery readiness: which services are within tolerance, which are not, what the financial and regulatory exposure is, and what is being done about it. Boards increasingly expect not just a snapshot of current posture but evidence of systematic improvement over time. SRI trend data provides exactly that.

SRIs as an Organizational Language

Beyond their function as a measurement tool, SRIs do something subtler and equally important. They give disparate teams like security, infrastructure, IT operations, business continuity, and business leadership a shared language for resilience. They provide a metric that everyone can see, report on, and be accountable to.

That shared language is what turns resilience from a technical function into an organizational capability. When the CISO, the CIO, and the CFO can look at the same dashboard and have the same conversation about the same services, ResOps has achieved something most organizations have never had: alignment on what resilience actually means and what it will take to prove it.

The evidence does not just satisfy a regulator or an insurer. It changes the organization's relationship with its own resilience: we stop *believing* we might be ready, and we start *knowing* whether we are or not.

CHAPTER 5

Start and Scale Your ResOps Initiative

The hardest thing about getting started, is getting started.
—Guy Kawasaki

ResOps begins with an executive decision. Recoverability of critical services will be governed as a business outcome, owned and measured like any other. Everything in this chapter follows from that decision. Without it, the steps that follow are useful but fragile. With it, they compound.

Two barriers stop most ResOps programs before they gain traction. The first is ownership. When no single person is accountable for driving the discipline across organizational boundaries, it fragments back into siloed activity within weeks. The second is funding. When a program spans multiple departments, each team's budgetary instinct is the same: yes, this is important, but it belongs in someone else's line item.

The order matters. An organization that solves ownership but struggles with funding will find a way forward. An organization that funds a program nobody owns will watch the investment dissolve. Start with governance.

Establish Minimum Governance

If no one owns the outcome, no one owns the gap. ResOps governance requires four things to be in place:

- An executive sponsor with board-level visibility
- A ResOps owner with cross-functional authority
- An accountable business owner for each critical service in scope
- A recurring forum where SRI posture and the resilience backlog get reviewed, prioritized, and either funded or explicitly accepted as risk by whoever holds that authority

The CISO should be the accountable sponsor, supported by the CIO, which is the direction the Gartner research cited at the beginning of Chapter 4 points to. You don't have to follow this exactly. What matters more is an owner with cross-functional authority, service owner participation, evidence production, and board reporting.

In Chapter 3, we discussed the governance rhythm this structure follows—monthly reviews and quarterly board reporting on resilience posture, tolerance attainment, and business exposure. The goal here is to form the structure that makes that rhythm possible. Start it early, even before the program has full coverage. The accountability it creates keeps the program from fragmenting.

Fund the Outcome, Not the Silo

Once ownership is established, the funding conversation changes. The question shifts from "Who pays for this?" to "How do we redirect what we're already spending toward a shared outcome?" ResOps does not begin by asking for more money. It begins by asking for existing investments to point at the same target. When gaps require new funding, the resilience backlog makes the business case with evidence.

Most large organizations are already running overlapping investments across security operations; backup and disaster recovery; infrastructure; cloud resilience; business continuity; governance, risk, and compliance (GRC); and cyber insurance readiness. Each is often owned by a different team, and each often produces outputs the others cannot easily use. Figure 5-1 illustrates how ResOps does not add another silo to that landscape; rather, it coordinates existing efforts around a shared evidence standard.

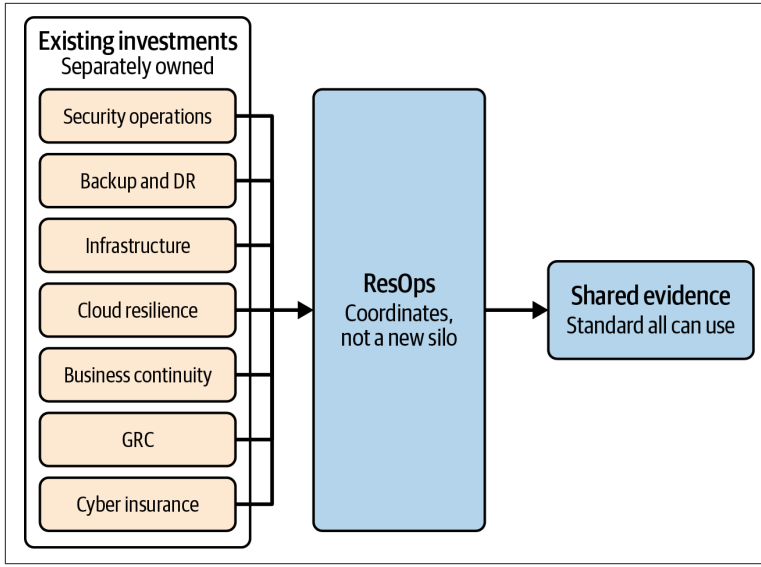

Figure 5-1. A high-functioning ResOps program may not be free but it does add efficiency to existing programs while centralizing the evidence each program relies on

The resilience backlog turns that coordination into investment logic. It shows which critical services are outside tolerance, why, what it costs to close the gap, and what exposure remains if the gap is accepted. That is a conversation a CFO can engage with. It is also the mechanism that keeps funding decisions tied to evidence rather than assumption.

Don't Forget GRC

Regulation is moving beyond documented plans toward demonstrated resilience. The Digital Operational Resilience Act (DORA) in EU financial services, and Network and Information Security Directive 2 (NIS2) in other essential and important sectors, both raise expectations for tested operational resilience, accountable governance, and evidence that critical services can withstand disruption. Both carry meaningful consequences for organizations that cannot demonstrate compliance with resilience, governance, and incident reporting obligations. In the United States, the US Securities and Exchange Commission (SEC) cybersecurity disclosure rules, evolving sector-specific requirements, and frameworks such as NIST CSF 2.0 are moving the conversation in the same direction.

Cyber insurance underwriting is reinforcing the same standard: organizations that can show evidence of tested recovery capability enter policy renewal discussions in a materially stronger position than those relying on questionnaire-based assertions alone. ResOps gives organizations a repeatable way to produce that evidence for regulators, insurers, and the board.

Start with One or Two Tier 0 Services

Choose one or two services from your Tier 0 classification. Confirm the business owner. Define and validate the impact tolerance. Map the critical dependencies. Assess what recovery evidence exists today. Run one realistic exercise. Produce the first SRI. Add the findings to the resilience backlog.

The goal is not perfection. The goal is evidence.

First 90 Days

One service, one proven recovery, ninety days. Let's discuss this in order:

Days 1 to 30
: Establish governance. Name the executive sponsor and ResOps owner. Form the monthly ResOps council. Select one or two Tier 0 services. Identify and confirm business owners. Define initial impact tolerances.

Days 31 to 60
: Build visibility. Map critical dependencies for the selected services. Identify what recovery evidence currently exists. Surface obvious recovery blockers. Define initial SRI criteria.

Days 61 to 90
: Produce evidence. Run one realistic recovery exercise. Measure the mean time to clean recovery (MTCR). Produce the first SRI dashboard. Build the initial resilience backlog. Report findings and investment needs to leadership.

The first exercise almost always produces a result worse than expected. That is not failure. That is the moment the organization stops guessing and starts governing.

Early Wins That Build Trust

A well-scoped initial engagement produces early wins that build organizational momentum before the program reaches full scale:

The MVC conversation
: Bringing security, operations, and business leaders together to define the MVC is often the first time those groups have formally agreed on what matters most. That alignment has value independent of the technical work that follows.

The dependency revelation
: Dependency mapping consistently surfaces previously unknown risks: services that depend on inadequately protected infrastructure, or third-party relationships that represent single points of failure. Finding those risks before disruption is an immediate, concrete win.

The first SRI dashboard
: Even an initial dashboard that shows a Red or Yellow status sends a clear message to leadership: resilience is now being measured, not just managed. That visibility changes the organizational conversation.

Scale Through Maturity

Scaling ResOps does not mean testing everything at once. It means expanding the same evidence-producing rhythm, service by service, starting with the MVC and moving outward by business criticality. The cadence matters as much as the coverage: each new service should enter the same cycle of ownership, tolerance definition, dependency mapping, exercise evidence, SRI reporting, and backlog remediation.

ResOps programs are not built overnight. They are developed through a progression of maturity stages, each building on the capabilities of the previous one. Most organizations that create and sustain a ResOps program go through four phases of maturity. Aligning cross-functional teams on where you are today, and why, is a good start. Table 5-1 shows a maturity model that provides a roadmap for that progression.

Table 5-1. The four phases of ResOps maturity

Maturity level	Description	Visibility and control	Confidence and testing	Executive signal
0 Fragmented recovery	Capabilities exist but are fragmented. Critical services are undefined. Recovery is manual and relies on individuals.	No isolation by design. Dependencies are incomplete or stale. Tolerances are undefined or RTO and RPO only.	Infrequent, component level. Ad hoc restores with unpredictable outcomes.	Leaders cannot prove recoverability. Recovery may succeed, but only through improvisation and luck.
1 Basic discipline	Critical services are identified and basic recovery objectives set. Ownership exists but evidence is incomplete.	Dependency mapping is partial. RTO, RPO, and MTCR may be defined but not validated. Basic backup is in isolation only.	Periodic and largely technical, with limited involvement from other teams. Runbooks are rarely exercised realistically.	Leaders believe they can recover, but confidence rests on plans more than proof.
2 Integrated ResOps	ResOps is established as a cross-functional operating discipline. Tolerances and SRIs are defined for critical services.	Dependencies are mapped and maintained. Isolation is applied purposefully. Gaps flow to a managed backlog.	Scenario-based exercises and drills produce actionable findings covering critical services.	Leaders get a consolidated posture view and can prove recovery for known scenarios.
3 Evidence-driven resilience	Resilience is continuously validated and managed as a business capability. Evidence guides investment and change.	Mapping is continuous and rebaselined on change. Isolation is designed in, so a compromise reaches far fewer systems.	Frequent, automated, tied to tolerances. Includes detection-failure drills. Rebuild-from-scratch is feasible.	Leaders answer, with proof, whether services recover within tolerance. Resilience is measured, governed, improving.

Most organizations beginning a ResOps program find themselves at Level 0 or Level 1. It is a starting point. The goal is not to reach Level 3 immediately, but to make deliberate, evidence-based progress and report that progress to the board at each stage.

Common Pitfalls

Organizations that struggle to build momentum in ResOps programs typically encounter one or more of the following problems:

Tool-first thinking
: ResOps is an operating discipline, not a technology deployment. Organizations that begin by evaluating tools before defining their MVC, mapping dependencies, or establishing accountabilities find that their investments do not produce the expected outcomes. The tools serve the discipline. Not the other way around.

Compliance mindset
: ResOps is the process that produces the evidence regulatory frameworks increasingly demand. But when it is treated as a compliance exercise, it produces documentation rather than capability. The output of a compliance-oriented program is a binder full of policies. The output of a genuine ResOps program is a set of SRIs backed by exercise evidence.

Lack of ownership
: ResOps requires a named owner with the authority to convene cross-functional teams, set priorities, and report outcomes to leadership. Without that ownership, the program fragments back into siloed practices within weeks.

Scope creep
: The impulse to solve everything at once is understandable but counterproductive. A ResOps program that tries to address fifty services simultaneously typically makes meaningful progress on none of them.

Declaring victory after one exercise
: A single successful recovery proves a scenario. It does not prove capability. Repeatability, under different conditions and with different people, is what distinguishes a capable organization from one that got lucky once.

The Executive Checklist

Every executive sponsor should be able to answer these questions before declaring that a ResOps program has genuinely launched:

- Is there a named executive sponsor with board-level visibility?
- Is there a named ResOps owner with cross-functional authority?
- Has the MVC been defined, with named business owners for each critical service?
- Have one or two Tier 0 services been selected as the starting scope?
- Have impact tolerances been defined by business owners and accepted through the governance process?
- Have dependencies been mapped for each selected service?
- Has at least one realistic recovery exercise been completed?
- Has the first SRI dashboard been produced and reported to leadership?
- Does a resilience backlog exist, tied to enterprise risk and funding decisions?

A program that can answer yes to all nine has moved from assumption to evidence. What that evidence enables, and what it demands of the organization over time, is what Chapter 6 addresses.

CHAPTER 6

Conclusion: Answers You Can Prove

> *You cannot change your destination overnight, but you can change your direction overnight.*
>
> —Jim Rohn

This handbook began with three questions: If a disruption hit tonight, could you recover? How long would it take? Can you prove it? Those questions are not going away. They will be asked by an increasingly broad range of voices: your board, your regulator, your cyber insurer, your customers, and even your competitors.

One more question runs beneath all three. In our Introduction we noted how Mythos and AI-powered tools compress the window between vulnerability discovery and active exploitation, radically shifting our calculus on the importance of both remediation and rebuilding.

Attackers are not constrained by legacy systems, technical debt, or customer patch cycles. The question is no longer whether disruption will come. It is whether the organization is designed to absorb it.

The organizations that answer these questions well are not the ones with the most sophisticated technology or the most comprehensive plans. They have done the work: defined what must be protected, mapped what it depends on, tested whether it can actually be restored, measured the gap between current capability and required tolerance, and committed themselves to closing that gap systematically over time.

That is the ResOps journey: the movement from assumption to evidence.

What Changes Will You See When ResOps is Established?

ResOps is equal parts people, process, and technology, often driven by an overarching change in philosophy. With this many moving parts it can be hard to discern real traction. But when ResOps takes hold in an organization, five shifts begin to become self-evident sooner rather than later:

Plans become evidence.
: The organization stops asserting what it believes about its resilience and starts proving what it knows.

Systems become critical services.
: The focus moves from managing infrastructure to protecting the business outcomes that infrastructure enables.

Recovery targets become tested evidence.
: Business owners define impact tolerances, teams validate recovery against realistic exercises, and SRIs show whether services are within tolerance.

Siloed activity becomes governed accountability.
: Security, infrastructure, IT operations, business continuity, and business leadership align around a shared outcome and a shared dashboard.

One-time testing becomes continuous improvement.
: Every exercise produces findings. Every finding enters the resilience backlog. Every gap has an owner, a timeline, and a consequence if it is not closed.

Know Before You Go

Before an organization embarks on its ResOps journey, or scales an existing program, four questions must be answered honestly:

Which services, if disrupted, would cause the most harm?
: If you cannot answer this with confidence, the first task is not to build recovery capabilities. It is to perform the business impact analysis and MVC work that produces the answer.

What are the impact tolerances for those services?
: If the tolerances have never been explicitly defined by business owners, they are not real tolerances. They are assumptions. Assumptions often struggle when the incident arrives.

Who is the named owner for each service's resilience outcome?
: If the answer involves multiple teams with shared accountability and no single owner, the first governance task is to establish clear ownership. Shared responsibility without named ownership becomes no accountability.

Have you actually tested recovery of your most critical service, under realistic conditions, against current data, with a measured result?
: If the answer is "not recently" or "not quite that way," the organization does not know whether its resilience posture is real or theoretical.

If you can answer all four confidently, with evidence, you have a foundation. Build from there. If you cannot, start there. Everything else is secondary.

The Executive Move

Chapter 5 ended with a launch checklist that could tell your team whether it was ready to start a ResOps program. An organization that can answer yes to those nine questions has not "succeeded at ResOps." It has simply crossed the threshold from assumption to evidence. The work that follows is to keep that evidence current, govern the gaps it reveals, and improve recoverability over time.

The executive move after reading this handbook is simple: ask for proof.

Not a policy. Not a DR plan. Not a presentation slide saying backups are healthy. Ask for evidence that one or two Tier 0 services can recover cleanly, end-to-end, within business-defined impact tolerances. The ongoing dialogue to have with your ResOps and technology teams (as opposed to the point-in-time checklist from the previous chapter) will sound like the following:

- What are our Tier 0 and Tier 1 services, and who is the named business owner for each?
- What is the business-defined impact tolerance for each critical service?

- When was recovery last tested under realistic conditions, and what was the MTCR?
- Which dependencies remain untested or unmapped?
- Which services are Green, Yellow, or Red on the resilience posture dashboard?
- What is in the resilience backlog, and what is the remediation status?
- What decision do you need from me?

That last question matters most. It transforms ResOps from an organizational program into an act of governance led by an executive.

If your teams cannot answer those questions today, that is not failure. (Remember, it's a discussion, not a checklist.) Treat it as a starting point. Give them 90 days, one or two Tier 0 services, and the mandate described in Chapter 5. Ask them to define tolerances, map dependencies, run a realistic recovery exercise, produce the first SRI dashboard, and return with the resilience backlog. That is how ResOps becomes real.

The Future of ResOps

ResOps is a young discipline, but its direction is clear. It will become critically important through at least three different lenses—AI, regulatory, and market differentiation:

AI will accelerate both sides.
: It will help defenders discover weaknesses, automate testing, and model recovery scenarios faster than ever. It will help attackers find paths to disruption with the same speed. The result is a compressed window between exposure and impact, and a rising premium on organizations that have already practiced recovery.

Regulatory expectations will harden.
: DORA, NIS2, SEC cybersecurity disclosure rules, and emerging AI governance frameworks all point toward the same expectation: demonstrate operational resilience, do not merely document it. Organizations building ResOps capability now will be ahead of that curve.

Resilience will become a market differentiator.

> Customers, partners, insurers, and investors are becoming more sophisticated about resilience posture. The ability to demonstrate proven recovery capability through SRIs, independent testing, or regulatory attestation will increasingly separate organizations that inspire confidence from those that invite doubt.

As organizations become more deeply dependent on AI-driven systems, more exposed to sophisticated and automated attacks, and more accountable to regulators who demand demonstrated rather than theoretical resilience, the discipline will mature rapidly and become a standard feature of enterprise operating models.

Figure 6-1 illustrates the value of ResOps across the organization.

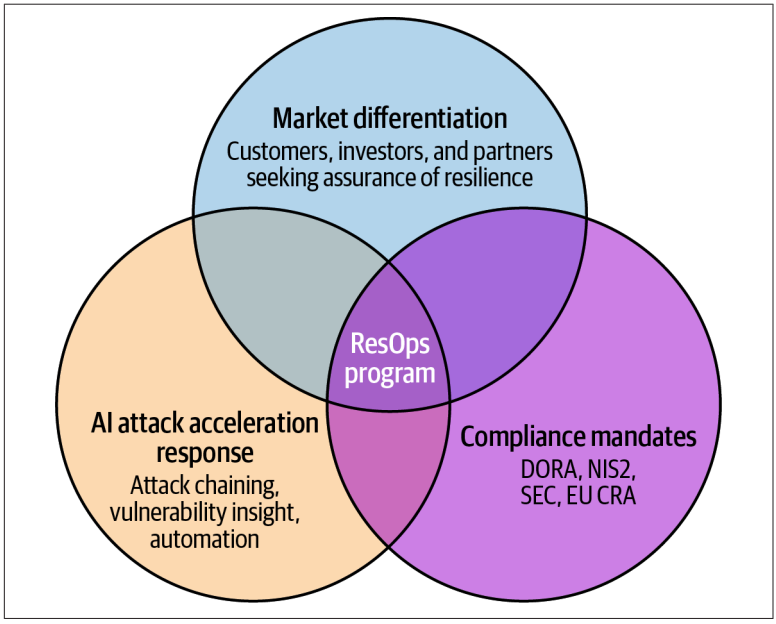

Figure 6-1. Your ResOps program can serve as a cornerstone for CISOs to build AI responses, for GRC teams to guide cross-compliance initiatives, and marketing teams to define your differentiation

A Final Word

Disruption is inevitable. The severity of its impact on your organization doesn't have to be. Mythos and other frontier models make disruption not only inevitable but perhaps immediate. ResOps is the discipline of deciding, in advance, what must survive; proving that it can; and improving until the evidence matches the tolerance the business requires. It does not replace security, disaster recovery, or business continuity. It makes them work together around the outcome that matters most during a bad day: the ability to keep the business functioning and recover cleanly.

If you take one thing from this handbook, take this: resilience is not what your plans *say*. It is what your organization *can prove*. Under pressure. Start with one or two Tier 0 services. Define tolerance. Map dependencies. Test clean recovery. Produce the first SRI dashboard. Build the backlog. Then ask again next quarter whether the evidence has improved.

In a world where shocks are inevitable and recovery windows are shrinking, advantage belongs to the organizations that have already practiced recovering.

That is the dual promise of ResOps: resilience you can prove and answers you can rely on.

About the Authors

Jason Cray is a leading data protection specialist with more than 26 years of experience in the IT industry, including two decades dedicated to data protection. He specializes in data backup, disaster recovery, and cyber resilience, advising organizations on how to safeguard critical information against evolving threats. With deep expertise in cyber recovery planning and execution, Jason has guided companies in building and operationalizing robust strategies to reduce the impact of cyberattacks.

Ben Herzberg is an experienced technical leader with an extensive history of securing data and enhancing cyber resilience. His professional experience in development, research, and security includes roles such as the CTO of Cynet and head of research at Imperva. Ben is also the chief scientist for Satori, a Commvault company.

Michael Thelander leads product marketing at Commvault, where his work focuses on cyber resilience strategy, data protection and recovery, and the emerging discipline of Resilience Operations (ResOps). He has spent almost 20 years in the cybersecurity industry in product management and marketing leadership roles. Michael's writing and analyses have appeared in *SC Magazine*, *Cyber Defense Magazine*, *Cyber Security: A Peer-Reviewed Journal*, and other publications.