RESOLVE

EXECUTIVE PLAYBOOK

INFRASTRUCTURE RESILIENCE · DISASTER RECOVERY

The Disaster Recovery *Execution* Playbook

How infrastructure, network and security leaders turn recovery plans into orchestrated, continuously tested and provable outcomes.

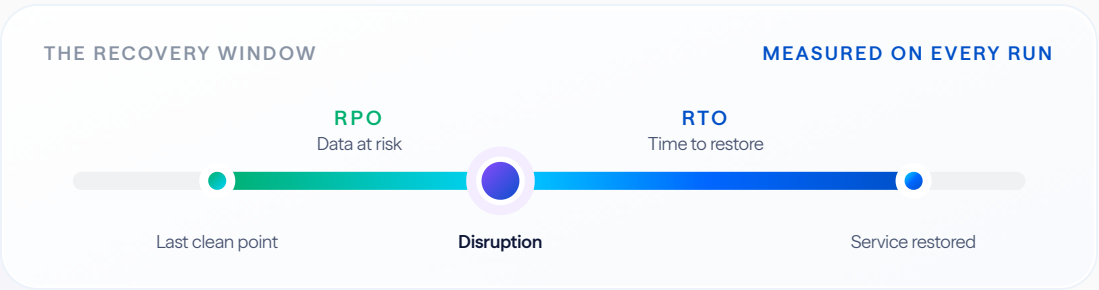

PUBLISHED
October 2026

READING TIME
30 minutes

OVERVIEW

Recovery Plans Succeed Or Fail In Execution

Most enterprises have a recovery plan, replication in place and a backup platform they trust. The weak point sits between the alert and the restored service, where dozens of steps across storage, compute, network, identity and security teams must happen in the right order, under pressure, often at 2 a.m. This playbook gives leaders a model for running those steps as orchestrated, tested and self-documenting workflows.

EXECUTIVE SUMMARY

An hour of downtime now costs most large enterprises more than $300,000, and procedural failure drives the majority of human-error outages. Leaders who convert Tier 1 runbooks into executable, continuously tested workflows shorten recovery, prove RTO and RPO on demand, and hand auditors evidence the moment a test completes.

Key Findings

01 **Downtime is a seven-figure event.** More than 90% of mid-size and large enterprises put one hour of downtime above $300,000, and 41% place it between $1 million and over $5 million.

02 **Procedures are the failure point.** Nearly 40% of organizations suffered a major outage caused by human error in three years, and 85% of those trace to procedures that were skipped or flawed.

03 **Backups are a primary target.** Attackers tried to compromise backups in 94% of ransomware attacks. When they succeeded, median recovery cost rose to $3 million from $375,000.

04 **Handoffs set the pace.** Recovery spans many tools and teams, so a plan moves only as fast as its slowest manual handoff, approval or validation step.

Recommendations

01 **Make Tier 1 runbooks executable.** Convert them into orchestrated workflows with approval gates where business risk is decided.

02 **Test recovery every month.** Run end-to-end tests in isolated environments and record achieved RTO and RPO on every cycle.

03 **Validate before you promote.** Make clean-point and integrity checks a mandatory step in every restore.

04 **Generate evidence as you execute.** Capture every step, timestamp and approval for DORA, ISO 22301 and SOX.

05 **Automate the way home.** Orchestrate failback and DR cost cleanup with the same rigor as failover.

PLANNING ASSUMPTION

By 2028, enterprises that run Tier 1 recovery as orchestrated, continuously tested workflows will restore critical services in a fraction of the time required by peers who execute runbooks by hand.

PLANNING ASSUMPTION

By 2028, auditors and regulators will expect recovery evidence captured by the systems that execute recovery, at the moment each step runs.

ANALYSIS

01 The Recovery Gap

Recovery investment has concentrated on data protection: replication, snapshots and immutable backups. Those controls determine whether data survives. The speed and safety of the return to service depends on how well hundreds of dependent steps are sequenced, approved and verified.

The research points the same way from three directions: downtime is expensive, outages are increasingly caused by procedure, and adversaries now aim directly at the recovery path.

Exhibit 1. Three Signals That Recovery Execution Is A Board-Level Risk THIRD-PARTY RESEARCH

90%+

of mid-size and large enterprises say one hour of downtime costs more than $300,000.

Source: ITIC 2024 Hourly Cost of Downtime Survey

85%

of human-error outages stem from staff not following procedures or from flawed procedures.

Source: Uptime Institute, Annual Outage Analysis 2025

94%

of ransomware victims say attackers tried to compromise their backups; 57% of those attempts succeeded.

Source: Sophos, State of Ransomware 2025

Exhibit 2. When The Recovery Path Is Compromised, Every Outcome Gets Worse RANSOMWARE VICTIMS, 2025

Source: Sophos, The Impact of Compromised Backups on Ransomware Outcomes. Survey of 2,974 organizations hit by ransomware. Bars drawn to a common scale within each panel.

IMPLICATION Recovery speed is now an operational capability to be engineered, measured and audited with the same discipline applied to uptime. The fastest recovering organizations treat the recovery process itself as software: versioned, executable and tested.

DIAGNOSIS

02 Six Places Recovery Breaks Down

Post-incident reviews of failed or slow recoveries tend to name the same handful of causes. Each one is a gap between what the plan says and what a team can execute under pressure, and each one closes when the step becomes an orchestrated action with a defined owner, input and check.

FAILURE POINT	WHAT IT LOOKS LIKE	BUSINESS CONSEQUENCE	RESOLVE COUNTERMEASURE
Runbook Drift	Documents describe last year's environment. Hostnames, IP ranges and credentials have changed since the last edit.	Steps fail on first contact and engineers improvise during the most critical hour.	Runbooks become executable workflows validated on every scheduled test run.
Manual Sequencing	Dependency order lives in senior engineers' heads. Application tiers start before their databases are ready.	Cascading restarts, extended outages and a long tail of broken integrations.	Dependency-aware orchestration across database, application, network and DNS layers.
Decision Latency	Approvals are chased across phone bridges while leaders wait for data on impact and data loss.	Minutes become hours before anyone authorizes failover.	Approval gates delivered in Teams or Slack with the data-loss window and impact attached.
Unverified Restores	Restored systems go live without integrity, malware or application checks.	Reinfection, corrupted data and a second outage on top of the first.	Automated clean-point selection, scanning and integrity validation before promotion.
Evidence Scramble	Audit evidence is assembled after the fact from chat logs, tickets and memory.	Findings, regulatory exposure and weeks of rework after every test.	Every step, approval and result is time-stamped and logged as it executes.
Forgotten Failback	The return to primary is improvised days later, and standby capacity keeps running.	Extended risk exposure and DR infrastructure spend that never switches off.	Orchestrated failback with staged traffic return and automated resource cleanup.

LEADER ACTION Ask your team to walk through the last recovery test step by step and mark each step as automated, scripted or manual. The manual steps on the critical path are your first automation backlog.

FRAMEWORK

03 The Resolve Recovery Execution Model

Every recovery workflow runs through five phases. Resolve carries the sequencing, integrations and checks at machine speed, while people make the calls that carry business risk with the facts in front of them.

Exhibit 3. Five Phases Of An Orchestrated Recovery RESOLVE FRAMEWORK

Detect & Correlate

Monitoring, SIEM and health checks feed one correlated trigger that starts the workflow.

Decide

The incident commander gets impact, data-loss window and recommended path in Teams or Slack.

Orchestrate

Resolve executes recovery across storage, compute, network, DNS and identity in dependency order.

Validate

Synthetic transactions, integrity checks and malware scans confirm each tier before the next starts.

Report & Improve

Evidence, timings and approvals post to ServiceNow and GRC, and gaps feed the next version.

HOW THE RESOLVE PLATFORM MAPS TO THE MODEL

BUILD

Jarvis

Describe a recovery procedure in plain language and Jarvis assembles the workflow for review and testing.

START FAST

The Automation Vault

Prebuilt workflows and building blocks for failover, restore and validation patterns.

CONNECT

Integrations

API and CLI reach across cloud, colocation and on-premises, from ServiceNow and Splunk to AWS and Azure.

KNOW

Knowledge Agent

Surfaces runbook context, service dependencies and prior incident history in the moment.

EXECUTE

Automation Agent

Runs recovery steps across domains with logging, retries and rollback paths built in.

GUIDE

Agent Assist

Guides engineers through decision points and gives leaders a conversational view of status.

ASSESSMENT

04 The Recovery Execution Maturity Model

Use the five levels to place your program and set a realistic next step. Level 3 for Tier 1 services is achievable within a quarter once the first workflows reach production.

L1 Documented

HOW RECOVERY RUNS
Runbooks live in wikis and PDFs. Engineers execute every step by hand from a bridge call.

TESTING & EVIDENCE
Annual tabletop exercises. RTO and RPO are stated in the plan and rarely measured.

NEXT MOVE
Inventory Tier 1 services and their dependencies. Identify the manual steps on the critical path.

L2 Scripted

HOW RECOVERY RUNS
Teams own scripts for their own domain. Handoffs and sequencing between teams stay manual.

TESTING & EVIDENCE
Partial failover tests by domain. Evidence is collected by hand after each test.

NEXT MOVE
Connect the scripts into one cross-domain workflow with a single trigger and approval gate.

L3 Orchestrated

HOW RECOVERY RUNS
Tier 1 failover runs end to end from one trigger, in dependency order, with approval at defined gates.

TESTING & EVIDENCE
Quarterly full tests. Achieved RTO and RPO are captured automatically on each run.

NEXT MOVE
Schedule non-disruptive tests monthly and add clean-point validation to every restore.

L4 Continuously Validated

HOW RECOVERY RUNS
Recovery workflows are versioned, tested monthly in isolated environments and updated with every change.

TESTING & EVIDENCE
Audit-ready evidence packs generated on every run and mapped to control frameworks.

NEXT MOVE
Link detection directly to recovery workflows and extend coverage to Tier 2 services.

L5 Adaptive Resilience

HOW RECOVERY RUNS
Correlated events start recovery automatically. Agents select the recovery path from live telemetry and present it for approval.

TESTING & EVIDENCE
Continuous validation across tiers, with live RTO and RPO posture reported to leadership.

NEXT MOVE
Extend the model to cyber recovery, third-party outages and business process continuity.

SELF-CHECK If your team cannot state the achieved RTO from your last Tier 1 test within five minutes, your program is operating at Level 2 or below, regardless of what the plan document says.

USE CASES

05 Eight Recovery Workflows To Orchestrate With Resolve

Automated step Human approval gate

01 Ransomware Recovery To A Verified Clean Point

CYBER RECOVERY

TRIGGER

SIEM or EDR correlation detects mass encryption behavior on production servers.

ORCHESTRATED RESPONSE

1. Quarantine affected segments and disable compromised accounts in Active Directory and Okta.
2. Lock backup retention and deletion policies to protect immutable copies.
3. Find the last snapshot that predates the first indicator of compromise.
4. Restore into an isolated environment and run malware and integrity checks.
5. **Gate:** the security lead reviews scan results and approves promotion.
6. Promote, rotate credentials and generate the incident evidence pack.

SYSTEMS IN SCOPE

SIEM, EDR, firewalls, Active Directory, Okta, backup platform, ServiceNow, Teams

WHY IT MATTERS

Attackers target backups in 94% of ransomware attacks, so clean-point validation decides whether recovery holds.

KPI TO TRACK

Time to verified clean restore

02 Cloud Region Failover

CLOUD RESILIENCE

TRIGGER

Datadog or LogicMonitor health checks breach thresholds across an AWS or Azure region for a sustained window.

ORCHESTRATED RESPONSE

1. Confirm the event across provider status, synthetic checks and application telemetry.
2. Scale warm standby capacity and verify replica lag against the RPO target.
3. **Gate:** the incident commander approves failover with the data-loss window displayed.
4. Promote the standby database and shift DNS and traffic manager weights to the standby region.
5. Run synthetic transactions and update the status page and major incident record.

SYSTEMS IN SCOPE

AWS, Azure, Datadog, LogicMonitor, DNS, ServiceNow, Slack or Teams

KPI TO TRACK

RTO attainment by service tier

03 Virtual Infrastructure And Storage Site Recovery

DATA CENTER

TRIGGER

Loss of a primary vSphere cluster, storage controller, datastore or site.

ORCHESTRATED RESPONSE

1. Verify replication state and select the recovery point for each protected group.
2. Fail over storage volumes and validate paths before compute starts.
3. Recover VMs in bulk and in dependency order: database, then application, then web.
4. Re-IP workloads, update DNS and reapply firewall and load balancer configuration.
5. **Gate:** application owners confirm each tier before the next tier starts.

SYSTEMS IN SCOPE

VMware vCenter, ESXi, Veeam, IBM SVC, F5, DNS, firewalls, CMDB

IN PRODUCTION

A Fortune 500 global payments provider runs bulk Windows VM recovery on Resolve, and a global travel technology provider automates storage controller failover.

KPI TO TRACK

Recovery time per application tier

04 Network Core And WAN Failure Rerouting

NETWORK

TRIGGER

SolarWinds, LogicMonitor or Netcool detects a core router, circuit or SD-WAN path failure.

ORCHESTRATED RESPONSE

1. Run CLI diagnostics on interface state, BGP neighbors and error counters across affected devices.
2. Push the pre-approved configuration that moves traffic to the backup path.
3. Validate latency, packet loss and application reachability from multiple sites.
4. Open the carrier ticket with diagnostics attached and track it to closure.
5. **Gate:** the network lead approves return to the primary path after a stability window.

SYSTEMS IN SCOPE

Routers and switches via CLI, SD-WAN controller, SolarWinds, Netcool, ServiceNow

IN PRODUCTION

A global telecom carrier runs 200,000+ closed-loop network automations a day on Resolve.

KPI TO TRACK

Network MTTR and share of incidents closed automatically

05 Continuous DR Testing And Audit Evidence

GOVERNANCE

TRIGGER

A monthly schedule, or a change event such as a new application tier entering production.

ORCHESTRATED RESPONSE

1. Build an isolated test environment that mirrors production networking.
2. Execute the full recovery workflow and measure achieved RTO and RPO against targets.
3. Tear down test resources to contain cost.
4. Compile the evidence pack with timestamps, approvals and logs mapped to DORA, ISO 22301 and SOX controls.
5. **Gate:** the service owner reviews any missed target and accepts the remediation plan.

SYSTEMS IN SCOPE

Hypervisor or cloud APIs, backup platform, ServiceNow, GRC platform

WHY IT MATTERS

Monthly tests catch runbook drift before a real event does.

KPI TO TRACK

Tier 1 test coverage in the last 90 days

06 Database Corruption And Point-In-Time Rollback

DATA

TRIGGER

A backup job fails, an integrity check fails, or a deployment writes bad data to a production database.

ORCHESTRATED RESPONSE

1. Place the application in maintenance mode and halt writes.
2. Identify the last consistent point from backup history and transaction logs.
3. Restore it to a side-by-side instance and validate row counts, checksums and key records.
4. **Gate:** the data owner approves cutover and the scope of transactions to replay.

SYSTEMS IN SCOPE

SQL Server and other database platforms, PowerShell, CI/CD pipeline, ServiceNow

IN PRODUCTION

A global IT services provider automates SQL Server backup monitoring, restore and refresh, part of a 20 to 40% workload reduction across accounts.

KPI TO TRACK

Achieved RPO versus target

07 Identity Provider Or Critical SaaS Outage

THIRD PARTY

TRIGGER

Okta, Entra ID or a business-critical SaaS platform becomes unavailable.

ORCHESTRATED RESPONSE

1. Confirm the outage against provider status and internal authentication telemetry.
2. **Gate:** the security lead approves activation of break-glass access.
3. Issue time-bound break-glass credentials with full session logging.
4. Reroute help desk intake and post user communications in Teams or Slack.
5. On restoration, revoke break-glass access and reconcile queued changes.

SYSTEMS IN SCOPE

Okta, Entra ID, privileged access tooling, Teams, Slack, ServiceNow

WHY IT MATTERS

Uptime Institute attributes roughly two-thirds of publicly reported outages to third-party providers.

KPI TO TRACK

Break-glass access revoked on schedule

08 Orchestrated Failback And DR Cost Control

RETURN TO PRIMARY

TRIGGER

The primary site or region is declared healthy after a failover.

ORCHESTRATED RESPONSE

1. Reverse replication and confirm the primary is fully synchronized.
2. **Gate:** change advisory approves the failback window.
3. Return traffic in stages with validation at each increment.
4. Decommission temporary capacity so standby spend stops.
5. Close the incident with a post-incident report and updated workflow version.

SYSTEMS IN SCOPE

Replication tooling, AWS, Azure, DNS, load balancers, ServiceNow

KPI TO TRACK

Days in DR state and residual DR spend

BUILD IT IN JARVIS

"When Datadog reports failed health checks across us-east-1 for more than five minutes, scale the standby environment in us-west-2, confirm replica lag is under five minutes, and ask the incident commander in the Major Incident Teams channel to approve failover. On approval, promote the standby database, shift DNS weights, run synthetic checks and log every step to the ServiceNow major incident."

Example prompt. Jarvis turns a plain-language description into a workflow your team reviews, tests and versions before it runs in production.

FIELD EVIDENCE

06 Recovery Workflows Resolve Customers Run Today

ACROSS THE RESOLVE CUSTOMER BASE

Recovery Automation Proven At Carrier Scale

200K+ closed-loop network automations every day at a global telecom carrier

13M+ automation executions a year at a UK telecom carrier

99% faster MTTR delivered for Resolve customers

100% of overtime eliminated at Syniverse

GLOBAL TELECOM MANAGED SERVICES

Fully Automated 4G Network Recovery

Multi-vendor radio network faults are detected, diagnosed and recovered end to end across element managers.

- Automated 4G site and service recovery
- Alarm correlation from Netcool into Remedy

Netcool, Remedy, Huawei Element Manager, Nokia NetAct

$97M saved annually, with ROI inside six months on the first contract

LEADING EUROPEAN TELECOM OPERATOR

Analytics-Triggered Network Service Restoration

Big data analysis triggers healing actions that connect to the OMC and execute commands to restore network service.

- Closed-loop restoration across multi-vendor RAN
- Automatic log collection for root cause

BMC Remedy, IBM MQ, Micro Focus, CMDB via REST

90% of L1 NOC operations fully automated

GLOBAL TRAVEL TECHNOLOGY PROVIDER

Storage And Load Balancer Failover

Core infrastructure failover runs as validated workflows in an operation handling 45,000 tickets a month.

- IBM SAN Volume Controller failover
- Load balancer pair failover, sync and validation
- Failed backend node restart and validation

IBM SVC, F5, Palo Alto, Cisco, IBM MQ

60% of incident types addressed by the first 16 runbooks

GLOBAL IT SERVICES PROVIDER

Database Backup, Restore And Refresh

Backup health and database restores run as repeatable actions across hundreds of client accounts.

- SQL Server backup monitoring
- Database backup, restore and refresh
- Critical server health checks

SQL Server, PowerShell, ServiceNow, Exchange

20–40% workload reduction across accounts

UK CABLE AND BROADBAND OPERATOR

Outage Impact And Service Restoration

Planned and unplanned network outages flow into CRM with customer impact attached, so field and network teams act on the right outages first.

- HFC outage automation
- Outage toolbox with custom impact analysis
- Power outage checks on Netcool events

Netcool, Remedy, Active Directory, U2000 Element Manager, CMTS

34% fewer tickets a month, about 6,500, and €3.8M saved annually

FORTUNE 500 GLOBAL PAYMENTS PROVIDER

Bulk VM Recovery And Replication Monitoring

Server recovery and replication health run as automated operations workflows across open systems.

- Bulk Windows VM recovery
- Replication error detection on data servers
- Remote validation of server reboot logs

SCOM, ServiceNow, Splunk, Netcool, SSH, MIR3 notification

11 automated operations workflows in the library

GLOBAL RETAILER

VMware Snapshot And Backup Operations

Snapshot and backup operations become self-service workflows with consistent execution across the VMware estate.

- Create, list, revert and delete VM snapshots
- Add or remove VMs from backup groups
- Provision Veeam proxies and repositories

VMware vCenter, ESXi, Veeam, Active Directory, easyVISTA

66 infrastructure and IT workflows in the library

US FIBER AND TELECOM PROVIDER

Emergency Service And Line Restoration

Critical service alarms trigger validated restoral runbooks directly on switching equipment.

- e911 alarm validation and restoral
- Line equipment restore with diagnostic validation
- Restoral exclusion checks and timed notifications

Netcool, DMS switches, GHEM

32 network workflows in the library

THE PATTERN Customers start with the high-frequency recovery actions: restart, revert, single-component failover and restore. Once those run reliably as workflows, the same building blocks compose into full site, service and region recovery.

Customer use cases and results from the Resolve customer use case library. Results are customer-reported. Customers are described by industry.

SCENARIO

07 Region Down At 02:14: Two Ways The Night Unfolds

A primary cloud region degrades in the middle of the night. The infrastructure, replicas and runbook are identical in both timelines. The difference is how the recovery steps execute.

Manual Runbook — 3h 51m

- **02:14** Alerts from three monitoring tools land in separate queues.
- **02:41** On-call engineer acknowledges and starts triage.
- **03:05** Bridge call assembled with database, network and app teams.
- **03:30** Runbook located. Last updated 14 months ago, two hostnames are wrong.
- **03:52** Leadership approves failover without a confirmed data-loss window.
- **04:10** Database promotion runs by hand, step by step.
- **05:20** DNS changed. Cached records keep some users on the failed region.
- **06:05** Service restored. Evidence reconstructed from chat logs over the next week.

Orchestrated With Resolve — 24m

- **02:14** Correlated health-check failure starts the region failover workflow.
- **02:16** Standby capacity scales. Replica lag verified at 40 seconds, inside RPO.
- **02:19** Approval request reaches the incident commander in Teams with impact and data-loss window attached.
- **02:24** Commander approves from a phone.
- **02:27** Database promoted. DNS weights shifted with low TTLs set in advance.
- **02:31** Synthetic transactions pass across checkout, login and API tiers.
- **02:38** Service restored. Status page, ServiceNow record and evidence pack complete.

9.6x faster return to service in this scenario

5 min from alert to a fully informed approval request

0 hrs spent assembling audit evidence after the event

Illustrative scenario. Timings are modeled for planning purposes and will vary by environment, architecture and approval policy.

MEASUREMENT

08 The Recovery Scorecard

Report these seven measures to leadership every quarter. Each comes straight from workflow execution data, so the scorecard updates itself as tests and real events run.

METRIC	DEFINITION	WHY LEADERS CARE	TARGET
RTO Attainment	Share of tests and real events that restored service within the target RTO	The single clearest measure of recovery readiness	95%+ Tier 1
RPO Attainment	Measured data loss window compared with the approved target	Quantifies business exposure in every event	100%
Tier 1 Test Coverage	Share of Tier 1 services tested end to end in the last 90 days	Untested recovery is assumed recovery	100%
Runbook Currency	Share of recovery workflows executed successfully in the last quarter	Detects drift before a real event does	90%+
Time To Decision	Minutes from correlated alert to approved recovery path	Decision latency is often the longest step	Under 15 min
Time To First Update	Minutes from alert to the first stakeholder and customer communication	Protects trust while recovery runs	Under 15 min
Evidence Freshness	Days since the last complete, audit-ready evidence pack per Tier 1 service	Keeps auditors and regulators satisfied on demand	Under 35 days

WHERE EACH NUMBER COMES FROM

EXECUTION LOGS

Every Workflow Run

Timestamps on each step feed RTO and RPO attainment, time to decision and time to first update.

TEST SCHEDULER

Every Recovery Test

Scheduled and change-driven tests feed test coverage and runbook currency.

EVIDENCE PACKS

Every Audit Record

Generated packs mapped to DORA, ISO 22301 and SOX feed evidence freshness.

ACTION PLAN

09 The 90-Day Recovery Execution Roadmap

Start with two workflows that prove value quickly: continuous testing for one Tier 1 service and one failover pattern. Expand from there on a predictable cadence.

Days 0–30

Baseline

- Inventory Tier 1 services, owners and dependencies.
- Load existing runbooks so the Knowledge Agent can surface them.
- Record achieved RTO and RPO from the most recent tests.
- Select two pilot workflows from The Automation Vault.

EXIT CRITERIA

Baseline scorecard published to leadership.

Days 31–60

Orchestrate

- Build pilot workflows with Jarvis and Vault building blocks.
- Define approval gates, owners and escalation paths.
- Run the first automated test in an isolated environment.
- Connect evidence output to ServiceNow and the GRC record.

EXIT CRITERIA

First measured, evidence-backed test complete.

Days 61–90

Prove And Expand

- Schedule monthly tests across all Tier 1 services.
- Add the ransomware clean-point recovery workflow.
- Report the first quarter-over-quarter scorecard.
- Set the Tier 2 backlog and quarterly expansion cadence.

EXIT CRITERIA

Tier 1 operating at Level 3 or above.

FOR THE BOARDROOM

Six Questions Your Board Will Ask

Q1

When did we last recover our most critical service end to end, and how long did it take?

Q2

What is our measured data loss window today, and who approved it?

Q3

If attackers reach our backups, how do we confirm a restore point is clean?

Q4

Which recovery steps depend on knowledge held by one or two people?

Q5

Can we produce audit evidence for our last recovery test within the hour?

Q6

How long does failback take, and what does each day in DR cost us?

EVIDENCE

10 Evidence And Methodology

1 **ITIC 2024 Hourly Cost of Downtime Survey.** Independent web survey of more than 1,000 firms worldwide, fielded November 2023 to March 2024. Source of the $300,000 per hour and 41% findings. calyptix.com/ITIC-2024-Hourly-Cost-of-Downtime

2 **Uptime Institute, Annual Outage Analysis 2025.** Source of the human-error, procedure, IT and networking, and third-party provider findings. businesswire.com/Uptime-Annual-Outage-Analysis-2025

3 **Sophos, The Impact of Compromised Backups on Ransomware Outcomes.** Survey of 2,974 IT and cybersecurity professionals at organizations hit by ransomware. Source of backup targeting, recovery cost and recovery time findings. sophos.com/impact-of-compromised-backups

4 **Resolve customer use case library.** Source of the field evidence in section 06 and the In Production notes in section 05. Workflow counts reflect use cases recorded per customer. Results are customer-reported, and customers are described by industry.

5 **Regulatory and standards context.** EU Digital Operational Resilience Act (DORA), ISO 22301 business continuity management, and NIST SP 800-34 contingency planning guidance inform the governance use case and scorecard.

HOW TO READ THIS PLAYBOOK

Key Findings draw on third-party research cited above. Recommendations, the execution model, the maturity model and scorecard targets reflect Resolve's point of view and field experience. Planning assumptions are forward-looking judgments intended to support strategy and budgeting.

ABOUT THE SCENARIO

The Region Down scenario in section 07 is illustrative. It holds infrastructure, replicas and runbook content constant and varies only the execution method, to isolate the effect of orchestration on time to recovery.

NEXT STEP

Turn Your Recovery Plan Into Running Workflows

Start with one Tier 1 service. Resolve helps your team build, test and prove its first orchestrated recovery workflow, then scale it across infrastructure, network and security.

Get Started

Explore The Automation Vault

RESOLVE

© 2026 Resolve. This playbook is provided for planning purposes. Third-party research is cited in section 10. Customer results are customer-reported and vary by environment. resolve.io