

EXECUTIVE PAPER · DISASTER RECOVERY

RTO and RPO Are Business Decisions

Recovery targets decide how much an outage can cost and how much the company spends to prevent it. Executives should set them, sign them, and fund them.

AUTHOR

Chris Watkins

SERVICE

Disaster Recovery

FOR

CEO, CFO, COO · 12 min read

AT A GLANCE

The problem. Every recovery plan rests on two numbers per system: how long it can be down, and how much recent data the company can afford to lose. In most organizations IT picks those numbers. The result is either a target the business cannot afford or a target the systems cannot meet.

Why it persists. Nobody outside IT has priced an hour of downtime or a lost day of transactions for each system. Without those figures, engineers guess, and executives never see the guess or the bill attached to it.

The recommendation. Have each business owner estimate the cost of an outage over time through a business impact analysis. Choose a recovery tier for each system from a priced menu, and have the business owner and CFO sign it. Then require IT to build to the signed target, prove it in testing, and report any system that cannot meet it.

The Decision in Front of You

Two numbers govern disaster recovery. The recovery time objective (RTO) is the longest a system can stay unavailable before the harm becomes unacceptable. The recovery point objective (RPO) is how far back in time the restored data may be, which sets how much recent work the company is prepared to lose. Both sit inside a wider limit set by the business owner, the maximum tolerable downtime of the process, and the RTO must be shorter than it [1].

These numbers size the disaster recovery budget. A one-hour RTO with near-zero data loss needs a second running environment and continuous replication. A three-day RTO with a one-day RPO needs good backups and a tested rebuild procedure. In the illustrative model in Exhibit 2, the first design costs fifteen times as much as the second each year, and the choice repeats for every system.

The decision for leadership is who sets the targets. NIST SP 800-34 Rev. 1, the federal contingency planning guide, describes a curve where the cost of disruption rises the longer an outage runs while the cost of recovery rises as the RTO shrinks. The right target sits where the two curves meet [1]. The business knows the first curve and IT knows the second, yet most companies let one side draw both.

What's Happening

Confident targets, loose alignment

Veeam's Data Trust and Resilience Report 2026, based on more than 900 senior IT, security, and risk leaders, found 90% were confident they could recover within their RTOs. Only 69% said those RTOs were fully aligned with business continuity goals [2]. The vendor sells recovery software, so treat the

figures with caution.

Regulators now ask who set the number

In the United Kingdom, the Financial Conduct Authority's operational resilience rules (PS21/3) required firms, by March 31, 2025, to set an impact tolerance for each important business service and show they can stay within it in severe but plausible scenarios [3]. An impact tolerance is the most disruption a service can suffer before harm becomes intolerable. The FCA reported in 2024 that it had seen a wide range of tolerances with limited rationale for where intolerable harm begins, and it expects the board to approve the firm's self-assessment [3].

In the European Union, the Digital Operational Resilience Act (DORA) has applied in full since January 17, 2025 [4]. Article 12 requires financial entities to set recovery time and recovery point objectives for each function, taking into account its criticality and its potential impact on market efficiency [5]. Coverage of the related standards indicates Article 11 expects those objectives to rest on a business impact analysis and requires the continuity plans to be tested at least once a year [4].

In the United States, the FFIEC's Business Continuity Management booklet, issued in November 2019, places continuity under board and senior management oversight and builds the program on a business impact analysis that estimates the cost of disruption over time [6][7]. The booklet defines RTO, RPO, and maximum tolerable downtime more fully than its predecessor [7]. ISO 22301:2019, the international business continuity standard, also treats a business impact analysis that sets tolerable disruption periods and recovery objectives as a core requirement [8].

These rules bind financial firms, and the same reasoning fits any company whose revenue stops when its systems do.

Two recent cases

Delta Air Lines, July 2024. After the CrowdStrike software outage of July 2024, Delta recovered more slowly than other major U.S. airlines. It canceled about 7,000 flights and delayed more than 10,000 over five days, affecting about 1.3 million customers [9]. Its chief executive put the cost at about $500 million [10]. The outage disrupted Delta's crew-tracking system, which assigns pilots and flight attendants to flights, and the airline had to reset about 40,000 servers by hand [10]. The U.S. Department of Transportation investigated the slow recovery and closed the probe in 2026 without seeking penalties [9]. Public reporting does not say what recovery target Delta had set for crew tracking. The lesson is that the system that limits recovery is often a supporting one, and only the people who run the operation know which one it is.

Jaguar Land Rover, September 2025. JLR disclosed a cyberattack on September 2, 2025, stopped production at its major plants, and restarted in phases by October 8, roughly five weeks later [11]. It booked £196 million in incident costs in its July to September quarter, when it reported a loss before tax and exceptional items of £485 million against a £398 million profit a year earlier; tariffs and lower volumes also contributed [11]. The UK government approved a £1.5 billion loan guarantee to support the

supply chain [11]. The Cyber Monitoring Centre, an independent body, estimated the total cost to the UK economy at £1.9 billion across more than 5,000 organizations [12]. For impact analysis, the lesson is that harm grows with time and spreads to suppliers, so a target set without that view understates what a week of downtime costs.

Why the Usual Approach Fails

When IT sets recovery targets alone, the first common failure is the uniform target. Someone writes "four-hour RTO, fifteen-minute RPO" for every system because it sounds responsible. Meeting it across the estate would need standby environments for payroll, the intranet, and the marketing site alongside the order system. When the quote arrives, the budget is cut, and the target stays in the plan anyway, now untrue.

The second is the reverse-engineered target. The team writes down what the current tooling happens to deliver, often "next business day" for everything. That number describes the technology. Nobody has checked whether the warehouse can ship without its system for a day, or whether a lost afternoon of orders can be re-keyed from email.

Both failures share a root cause. Downtime has never been priced per system by the people who own the revenue. NIST's guide assigns that task to process owners and business managers, working with the contingency planner to decide how much downtime and data loss each process can accept [1]. In the author's experience, the planner usually fills in the form alone.

Dependencies make it worse. Order taking depends on the order application, a database, an identity service, a network, and often a third-party platform. An RTO set for the application alone means little if the identity service recovers last, as Delta's crew-tracking delay showed [10].

The Framework: The Recovery Target Contract

Treat each recovery target as a contract. The business owner states what an outage would cost and signs for the risk, IT commits to build and prove the capability, and the CFO countersigns because the target sets the spend.

Exhibit 1. The Recovery Target Contract: business sets the target, IT builds and proves it



Source: author's framework, based on the business impact analysis process in NIST SP 800-34 Rev. 1 [1].

Impact. For each business service, the owner estimates the cost of an outage at one hour, one day, and one week, in revenue, penalties, customer loss, and regulatory exposure. The owner also estimates what an hour or a day of lost data would cost to re-create, and whether it can be re-created at all. Include suppliers and customers, as the JLR case shows [12].

Choose. IT publishes a short menu of recovery tiers, each with an RTO, an RPO, and an annual cost. The owner picks a tier. Cost and impact sit on the same page, which is the comparison NIST's cost balance curve calls for [1].

Sign. The business owner and the CFO sign the chosen RTO and RPO. The signature records who accepted the residual risk.

Engineer. IT builds every dependency of the service to the signed tier, sets backup frequency to the RPO, as NIST SP 800-53 expects [13], and records where the current design falls short.

Prove. Recovery is tested against the signed target, and any system that misses it goes to leadership with a date and a cost to fix. A target counts once a test has met it.

What a priced menu looks like

This illustrative model covers a hypothetical distributor's order system.

Exhibit 2. Illustrative recovery tiers for one order system: the middle tier costs least overall

Item	Tier 1: hot standby	Tier 2: warm standby	Tier 3: rebuild from backup
RTO / RPO	1 hour / 5 minutes	8 hours / 1 hour	72 hours / 24 hours
Annual cost of the capability	$600,000	$200,000	$40,000
Downtime cost in one major outage	$50,000	$400,000	$3,600,000
Expected annual downtime cost at 20% likelihood	$10,000	$80,000	$720,000
Total expected annual cost	$610,000	$280,000	$760,000

Source: author's illustrative estimate. Assumes $50,000 lost per hour of outage, one major outage with a 20% chance in any year, and downtime equal to the RTO. Excludes the cost of lost data, which would favor Tiers 1 and 2 further. Capability costs are placeholders; real quotes vary widely.

The model shows why IT cannot pick the tier alone. At $5,000 per hour of outage, Tier 3 becomes the cheapest choice for the same system. The per-hour figure comes from the business.

Costs, Risks, and Tradeoffs

Exhibit 3. Tradeoffs of business-owned recovery targets

Area	Cost or risk	Mitigation
Executive time	Owners must estimate impact for each service	Limit the first round to the top ten services and estimate in ranges
Uncomfortable answers	Signed targets expose systems that cannot meet them	Report each gap with a cost to close it and let the business fund or accept it
Higher spend	Honest targets for critical services may cost more than today's budget	Fund it partly by moving low-impact systems to cheaper tiers
Estimate quality	Owners may overstate impact to get better protection	CFO reviews the impact figures before signing
Drift	Targets go stale as the business changes	Re-sign annually and after major system or business changes
When this is wrong	A small company with a few SaaS tools has little to engineer; its recovery is set by vendors	Read the vendors' recovery commitments, decide whether they are acceptable, and plan manual workarounds

Source: author's analysis.

How to Act

Exhibit 4. A one-year roadmap to signed recovery targets

Window	Actions	Owner	Milestone
First 30 days	List the top ten business services and name an owner for each. Map each to the systems and vendors it depends on.	COO with CIO or CTO	Ten services, ten owners, dependency map drafted
First 90 days	Run a business impact analysis for each service: cost at one hour, one day, one week; tolerable data loss. Publish a priced tier menu.	Business owners with a continuity lead; IT for pricing	Impact figures and tier menu reviewed by the CFO
First 90 days	Owners choose tiers. Owners and CFO sign RTO and RPO for each service.	Business owners and CFO	Ten signed targets on file
First 365 days	Build each service and its dependencies to the signed tier. Align backup frequency to RPO. Test each service against its target and report gaps. Re-sign annually.	CIO or CTO with risk or internal audit	Every signed target met in a test, or a funded remediation date in front of leadership

Source: author's recommendation.

Questions to Ask Your Team

1. Who chose the recovery time and recovery point for our order system, and when?
2. What does one hour, one day, and one week of downtime cost on each of our top five services?
3. How much recent data could we lose on each of them, and could we re-create it?
4. Which of our written targets have we met in a test, and which have we never tried?
5. Which supporting systems, such as identity or scheduling, would recover last, and what waits on them?
6. Which vendors' recovery commitments are we relying on, and have we read them?

Closing

Recovery targets decide what an outage will cost and what the company spends to avoid it. Those are business choices with a price on each side. Have the people who own the revenue set the RTO and RPO for each system through an impact analysis, have the CFO sign them, and hold IT to building and proving exactly what was signed.

Sources

[1] National Institute of Standards and Technology. "Contingency Planning Guide for Federal Information Systems." NIST SP 800-34 Rev. 1, Section 3.2. 2010. https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-34r1.pdf.

[2] Veeam. "Veeam Report Reveals a Market-Wide Shift from Recovery Confidence to Proven Data Resilience." Press release, Data Trust and Resilience Report 2026. April 2026. https://www.albawaba.net/business/pr/veeam-report-reveals-market-wide-shift-1625554.

[3] Financial Conduct Authority. "Operational resilience: insights and observations for firms." May 28, 2024. https://www.fca.org.uk/firms/operational-resilience/insights-observations.

[4] sota.io. "DORA Article 11: ICT Business Continuity, Backup and Disaster Recovery for Financial Services." Secondary summary of Regulation (EU) 2022/2554. 2026. https://sota.io/blog/dora-article-11-ict-business-continuity-backup-disaster-recovery-financial-services-developer-guide-2026.

[5] European Parliament and Council. Regulation (EU) 2022/2554 (Digital Operational Resilience Act), Article 12. OJ L 333, 27.12.2022. Via SpringLex. https://www.springlex.eu/en/packages/dora/dora-regulation/article-12/.

[6] Compass IT Compliance. "The FFIEC Gives the BCP Booklet a Facelift." Summary of the FFIEC Business Continuity Management booklet. 2019. https://www.compassitc.com/blog/the-ffiec-gives-the-bcp-booklet-a-facelift.

[7] Protiviti. "FFIEC's Updated BCM Booklet Highlights Operational Resilience Concepts." December 3, 2019. https://blog.protiviti.com/2019/12/03/ffiecs-updated-bcm-booklet-highlights-operational-resilience-concepts/.

[8] RiskTemplate. "ISO 22301 Business Continuity: Requirements, Implementation, and How It Maps to Your BCP." Secondary summary of ISO 22301:2019. 2026. https://risktemplate-site.pages.dev/blog/2026-03-26-iso-22301-business-continuity-requirements-guide.

[9] Eckert Seamans. "DOT Closes Probe Into Delta's 2024 CrowdStrike-Related Meltdown." July 24, 2026. https://www.eckertseamans.com/stay-informed/blogs/aviation/dot-closes-probe-into-deltas-2024-crowdstrike-related-meltdown.

[10] NPR. "Delta's CEO says the CrowdStrike outage cost the airline $500 million in 5 days." July 31, 2024. https://www.npr.org/2024/07/31/nx-s1-5058652/delta-crowdstrike-outage-500-million-dollars.

[11] BleepingComputer. "Jaguar Land Rover cyberattack cost the company over $220 million." 2025. https://www.bleepingcomputer.com/news/security/jaguar-land-rover-cyberattack-cost-the-company-over-220-million/.

[12] Business Today. "Jaguar Land Rover hack triggers £1.9 billion economic loss, UK's worst cyber event yet: report." October 22, 2025. https://www.businesstoday.in/tech-today/news/story/jaguar-land-rover-hack-triggers-ps19-billion-economic-loss-uks-worst-cyber-event-yet-report-499156-2025-10-22.

[13] National Institute of Standards and Technology. "Security and Privacy Controls for Federal Information Systems and Organizations." NIST SP 800-53, control CP-9, via STIG Viewer. https://stigviewer.com/controls/800-53/CP-9.

The views in this paper are the author's own and do not represent any employer. Scenarios are hypothetical. This paper is not legal, financial, or compliance advice.