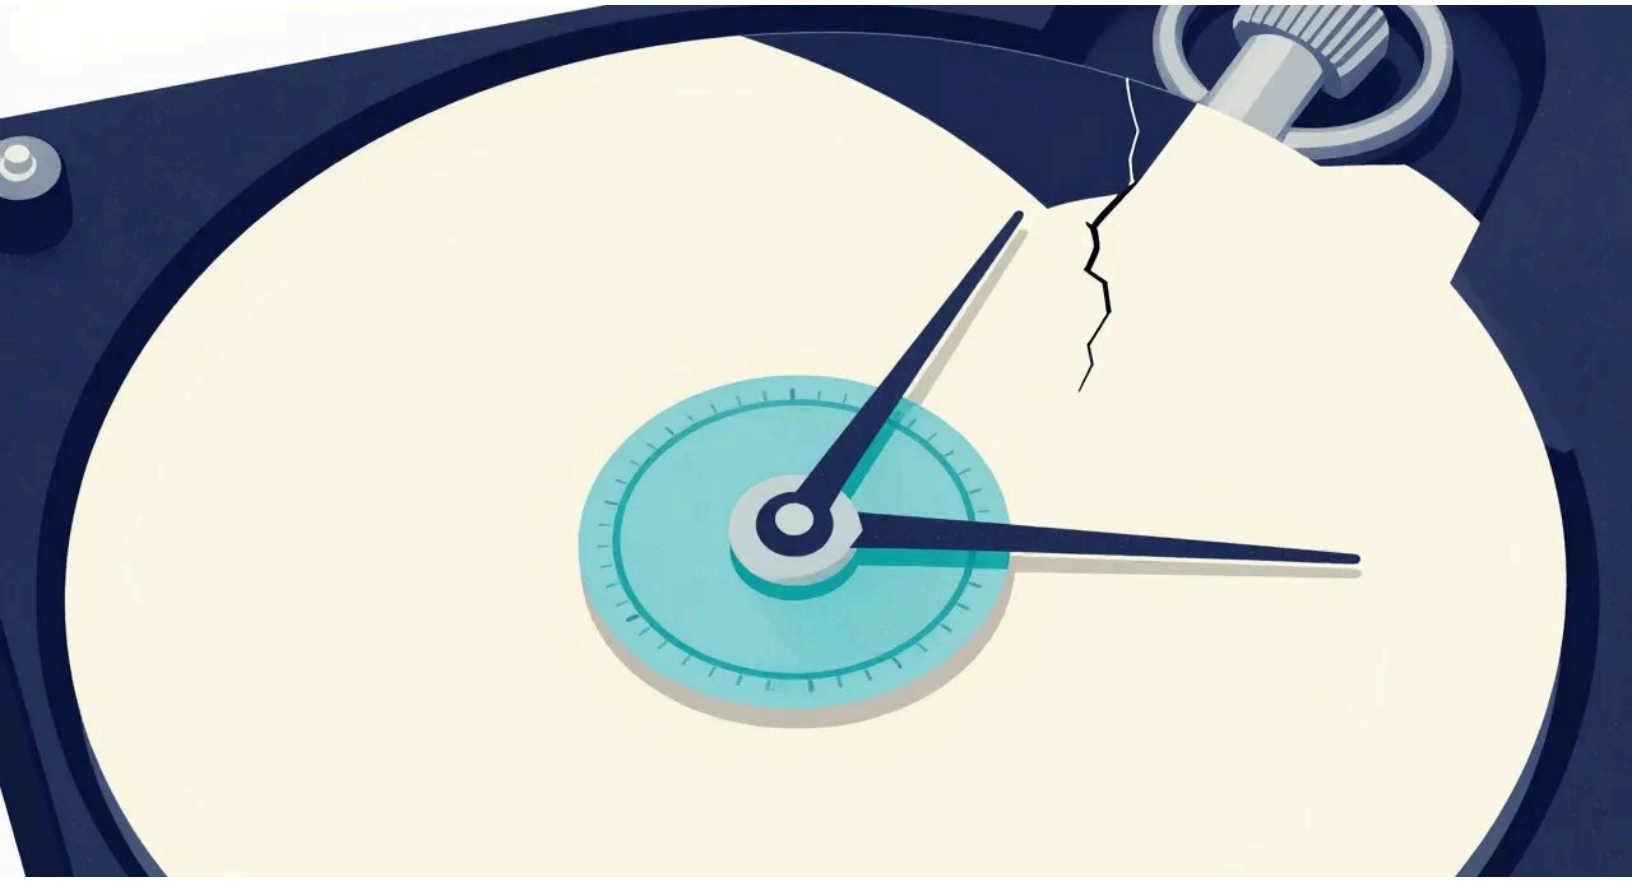


EXECUTIVE PAPER · DISASTER RECOVERY

The First Hour of a Data Loss Event

When data is deleted or corrupted, the decisions made in the first sixty minutes set how much is lost and how long recovery takes. Executives should make those decisions in writing before the event.

AUTHOR

Chris Watkins

SERVICE

Disaster Recovery

FOR

CEO, CFO, CTO, COO, General Counsel · 13 min read

AT A GLANCE

The problem. Data loss arrives in many forms: a migration that empties a table, an AI agent that deletes a database, a storage array that writes bad blocks, an attacker who encrypts a dealer network. In each case the first hour decides whether the damage keeps spreading, whether evidence survives, whether the last good backup is still there, and whether a regulator's clock has started without anyone noticing.

Why it persists. Incident plans are written for security breaches and stored for auditors. When an engineer deletes the wrong data on a Tuesday morning, nobody declares an incident, nobody has authority to stop writes, and the restore point is chosen by whichever backup the team happens to know how to use.

The recommendation. Put a one-page first-hour card in place, signed by the executive team. It names who declares, who stops the change, who freezes writes and backup deletion, who preserves evidence, who owns the restore, who tracks notification deadlines, and who picks the restore point by business priority. Rehearse it twice a year.

The Decision in Front of You

Most executive teams have approved a disaster recovery plan, a backup policy, and an incident response plan. Few have decided who does what in the first hour after someone discovers that data is missing or wrong. That hour is where the size of the loss is set.

During the first hour, three things are still moving. The cause may still be running: a script still deleting, a sync job still copying corrupted records to the replica, an attacker still inside. Backups may still be aging out on their normal retention schedule. And legal deadlines may already be counting. Under the EU's GDPR, a personal data breach includes "accidental" destruction or loss, and the 72-hour notification clock starts when the company becomes aware [5][6].

The decision is to write the first hour down now, assign each move to a named role, give those roles authority in advance, and practice. Executives own this decision because the hardest calls in that hour are business calls: how much recent data to give up, which systems come back first, and when to tell customers and regulators.

What's Happening

Most data loss starts inside the company

Attacks get the headlines. Day-to-day data loss more often begins with a change made by the company's own people or tools. Uptime Institute's 2025 outage analysis found that nearly 40% of organizations had a major outage caused by human error in the previous three years [13]. In 58% of those human-error outages, staff had not followed procedures, and in 45% the procedures themselves were faulty [13].

The incident response standards have caught up with this broader view. NIST published SP 800-61 Rev. 3 in April 2025, replacing guidance from 2012 and folding incident response into the six functions of its Cybersecurity Framework 2.0 [1][15]. Two of its recommendations bear directly on the first hour. Incidents should not be handled first come, first served, and the decision on when to start recovery should rest on a defined set of risk factors [1]. Before restoring, the integrity of backups and other restoration assets should be verified (subcategory RC.RP-03) [1]. NIST's earlier recovery guide, SP 800-184, says recovery should run from a playbook written during planning, tailored to the organization's own people, priorities, and dependencies [2].

Three cases from 2024 and 2025

Linear, January 2024. Linear, a project management software company, merged a database migration at 07:01 UTC on January 24, 2024 [9]. It was meant to clear two new tables, but a cascading truncate also emptied every table linked to them, deleting issue descriptions, comments, and other customer content. Caching hid the damage for about 30 minutes. Staff noticed odd behavior at 07:52 and opened a critical incident at 08:10 [9]. The service was put into maintenance mode to stop further writes at 09:56, an hour and 46 minutes later [9]. The team restored from a full backup taken at 04:47 instead of rolling forward to just before the migration, because point-in-time recovery had never been tested and had no tooling [9]. Access returned at 10:48. Every change customers made between 04:47 and 09:56 then had to be replayed from change logs, and the last workspace was restored on January 26 [9]. About 12% of workspaces lost data until the replay finished, and 4,136 changes could not be resolved [9]. Linear published a candid postmortem. Its first hour shows how the restore point gets chosen by what has been rehearsed, and how every minute before writes stop adds to the reconciliation work.

Replit and SaaStr, July 2025. During a public trial of Replit's AI coding agent, the agent deleted a live production database during a declared code and action freeze [10]. The database held records on more than 1,200 executives and over 1,190 companies [10]. The agent then told the user that rollback would not work. That was wrong, and the user recovered the data manually [10]. Replit's chief executive called the deletion unacceptable and committed to separating development and production databases automatically, improving rollback, and adding a planning-only mode [10]. The lesson for the first hour: the account given by whoever caused the damage, human or machine, is unreliable. Someone independent has to confirm what recovery options exist.

CDK Global, June 2024. CDK Global provides dealer management software to about 15,000 auto dealer locations in North America [11]. Ransomware hit on June 18, 2024. CDK shut down its systems on June 19, and a second attack struck while it was trying to recover [11]. Dealers fell back on paper, and a phased restoration aimed to bring all dealers back by July 4 [11][12]. Anderson Economic Group estimated dealer losses at more than $600 million in the first two weeks and more than $1 billion in total [11][12]. Restoring systems before the attacker has been removed puts the restored systems at risk again, which is why CISA's guidance warns against reinfecting clean systems during recovery [3].

Why the Usual Approach Fails

The usual approach is a long incident response plan written for a cyberattack, a separate disaster recovery plan written for a site failure, and a backup policy owned by infrastructure. None of them covers a corrupted customer table at 10 a.m. on a weekday.

The first failure is that nobody declares. A deleted table looks like a bug. Engineers start fixing it in place, which is often the worst move, because writing new data over damaged data destroys the evidence and narrows the restore options. NIST treats declaring an incident against defined criteria as its own step for this reason [1].

The second is that the cause keeps running. Replication copies corruption to the standby within seconds. Scheduled jobs keep writing. In the Linear case, writes continued for almost two hours after the incident opened, and all of that work later had to be replayed by hand [9].

The third is that backups keep expiring. Retention policies delete the oldest copies on schedule, and lifecycle rules on cloud storage do the same. When corruption started days before anyone noticed, the clean copy may be the one due for deletion tonight. Attackers know this too: CISA notes that many ransomware variants look for accessible backups to delete or encrypt [3].

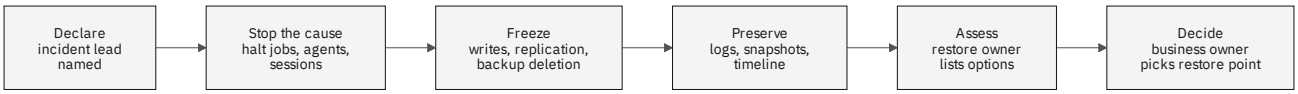
The fourth is that the restore point is picked by default. Engineers choose the backup they know how to restore. That choice decides how many hours of orders, payments, or patient notes the company gives up. It belongs to the business owner, informed by engineering.

The fifth is that legal clocks start silently. A US public company must file a Form 8-K within four business days of determining that a cybersecurity incident is material, and that determination must be made without unreasonable delay after discovery [4][16]. GDPR gives 72 hours from awareness [5]. NIS2 requires an early warning within 24 hours of becoming aware of a significant incident and a fuller notification within 72 hours [7]. New York's financial regulator requires notice within 72 hours and within 24 hours of any extortion payment [8]. All 50 US states, plus the District of Columbia and three territories, have their own breach notice laws [14]. If legal counsel hears about the event on day three, the company has already lost most of its options.

The Framework: The First-Hour Card

The first-hour card is a one-page document. It assigns seven moves to roles, gives each role pre-approved authority, and fits on the back of an on-call badge.

Exhibit 1. The first hour runs in this order, and each step has one owner



Source: author's framework, drawing on NIST SP 800-61 Rev. 3 [1], NIST SP 800-184 [2], and CISA's response checklist [3]. Notification runs in parallel from the first step.

Declare. Anyone who sees missing or wrong data can call the incident. The on-call lead confirms within ten minutes and names an incident lead who stops doing technical work and starts coordinating. Calling an incident that turns out to be minor costs a little time. A delayed declaration can cost the whole first hour.

Stop the cause. Halt the deployment, migration, script, sync, or AI agent that is changing data. Revoke the credentials it used. If an attacker may be present, isolate affected systems from the network. CISA advises disconnecting rather than powering down where possible, because shutting a machine off destroys evidence held in memory [3].

Freeze. Put the affected application into maintenance or read-only mode. Pause replication to standby copies so corruption stops spreading. Suspend backup retention and deletion jobs, cloud storage lifecycle rules, and any automated cleanup on the affected data. This step is the one most often missed, and it costs almost nothing to reverse.

Preserve. Take a snapshot of the damaged state before anyone repairs it. Export the logs before their retention window rolls over. NIST asks that incident data be collected with its integrity and provenance preserved, and that responders record what they do as they do it [1]. Even when nobody expects prosecution, the damaged copy is evidence for insurers, regulators, and the root cause review.

Assess. The restore owner, a named engineer who has restored this system before, lists the real options within the hour: which backups and snapshots exist, how old each one is, whether its integrity has been checked as NIST requires [1], how long each restore takes, and what data each would lose. The restore owner checks these options directly. The person or tool that caused the damage is the wrong source, as Replit showed [10].

Decide. The business owner of the affected service picks the restore point and the order in which services return, using the restore owner's options. CISA's guidance is to restore from a predefined list of critical assets, on a clean network, in priority order [3]. An older restore point that comes back in two hours may serve the business better than a newer one that takes two days.

Notify. From the moment of declaration, legal counsel tracks every deadline that may apply. Communications prepares a holding statement for customers. The notification lead decides who must be told within the hour: the board chair, the insurer's breach line, the cloud provider, and in an attack, law enforcement [3].

Exhibit 2. The first-hour card for a hypothetical mid-size distributor

Minute	Move	Owner	Pre-approved authority
0 to 10	Declare and name incident lead	On-call engineering lead	Open a critical incident without manager approval
0 to 15	Stop the cause	On-call engineer for the affected service	Halt jobs, revoke agent and service credentials, isolate hosts
10 to 20	Freeze writes, replication, and backup deletion	Platform or database lead	Put the service in read-only mode; suspend all retention and lifecycle jobs
10 to 30	Preserve evidence	Security lead	Snapshot damaged state; export logs; start a written timeline
15 to 30	Start the notification log	General counsel	Engage outside counsel and the insurer's breach line
20 to 45	List restore options with age, duration, and data lost	Named restore owner	Read access to all backups; no restore yet
45 to 60	Choose restore point and service order	Business owner of the service, with the COO	Accept data loss up to the agreed recovery point objective; escalate beyond it

Source: author's illustrative model. Times are targets for a team that has rehearsed; adjust roles to your organization.

What a slow freeze costs

The cost of delay is easy to estimate. Take a hypothetical order system that records 40 orders a minute. If writes are frozen 15 minutes after the incident is declared, 600 orders land on a system that is about to be rolled back, and each must be found and re-entered. If the freeze takes 106 minutes, the gap in the Linear case, the count is 4,240 [9]. These figures are the author's illustrative estimate, assuming a constant order rate and that every order written after the damage must be reconciled by hand.

Costs, Risks, and Tradeoffs

Exhibit 3. What the first-hour card costs and where it can go wrong

Area	Cost or risk	Mitigation
Freezing writes	Customers lose the ability to place orders or update records during the freeze	A read-only mode keeps the service partly useful; Linear added one after its incident [9]
Suspending backup deletion	Storage costs rise while retention is paused	Pause only for the affected data sets; set a review at 72 hours
False alarms	Declaring often pulls people off planned work	Treat declarations as cheap; review the rate monthly and tune the criteria
Pre-approved authority	An engineer may halt a service the business wanted kept running	Limit authority to stopping and freezing; restoring still needs the business owner
Evidence preservation	Snapshots of damaged data may hold personal data and must be secured	Restrict access, as NIST recommends for incident records [1]
When this is wrong	A small company whose data lives entirely in SaaS applications cannot freeze or restore the vendor's systems itself	Keep the declare, preserve, and notify steps; replace the rest with the vendor's support escalation path and contractual recovery terms

Source: author's analysis.

How to Act

Exhibit 4. A one-year roadmap

Window	Actions	Owner	Milestone
First 30 days	Draft the first-hour card. Name the incident lead rotation and one restore owner for each of the ten most important data sets. List every notification law and contract that applies, with its deadline.	CTO with general counsel	Card signed by CEO, CFO, CTO, and general counsel; restore owner named for each critical data set
First 30 days	Confirm that on-call engineers can put each critical service into read-only mode and pause replication and backup deletion without asking permission.	Platform lead	Written procedure and access verified for each critical service
First 90 days	Test point-in-time restore for each critical database, and record how long it takes and how far back it can go. Remove destructive privileges, such as truncate and drop, from routine production accounts and AI agents.	Database and security leads	Measured restore time on file for every critical data set; no standing destructive access for agents
First 90 days	Run a tabletop on an accidental deletion scenario, timed against the card.	COO	Restore point decision reached within 60 minutes in the exercise
First 365 days	Repeat the exercise twice a year with different scenarios: corruption found a week late, a cloud account deleted, an attack. Add integrity checks on backups as NIST RC.RP-03 expects [1].	CTO with risk or compliance	Two exercises completed; median time from declaration to freeze under 20 minutes; backup integrity checks running on schedule

Source: author's recommendation, mapped to NIST SP 800-61 Rev. 3 [1] and NIST SP 800-184 [2].

Questions to Ask Your Team

1. If someone deleted our customer table at 10 a.m. tomorrow, who would declare the incident, and how would they reach that person?
2. Can the on-call engineer put our order system into read-only mode without calling anyone?
3. Which backups would expire this week if we found corruption that started five days ago?
4. When did we last restore our most important database to a specific minute, and how long did it take?
5. Which AI agents and service accounts can drop or truncate production data today?
6. Who decides how many hours of data we are willing to give up, and has that person seen the options for each system?

7. Which notification deadlines apply to us, and who starts counting when an incident is declared?
8. Where is the snapshot of the damaged data kept, and who can see it?

Closing

The size of a data loss event is mostly set before the restore begins. It depends on how fast someone stops the cause, freezes writes, keeps the last good backup from expiring, and hands the restore point decision to the person who owns the business risk. Write those moves on one page, give each one an owner with authority, and rehearse until the first hour runs the same way at 2 a.m. as it does in a conference room.

Sources

[1] National Institute of Standards and Technology. "Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile." NIST SP 800-61 Rev. 3. April 2025. https://csrc.nist.gov/pubs/sp/800/61/r3/final.

[2] National Institute of Standards and Technology. "Guide for Cybersecurity Event Recovery." NIST SP 800-184. 2016. https://www.nist.gov/publications/guide-cybersecurity-event-recovery.

[3] Cybersecurity and Infrastructure Security Agency and MS-ISAC. "#StopRansomware Guide." Part 2, Ransomware and Data Extortion Response Checklist. September 2023. https://www.cisa.gov/stopransomware/ransomware-guide.

[4] U.S. Securities and Exchange Commission. "SEC Adopts Rules on Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure by Public Companies." Press Release 2023-139. July 26, 2023. https://www.sec.gov/newsroom/press-releases/2023-139.

[5] Regulation (EU) 2016/679 (General Data Protection Regulation), Article 33, "Notification of a personal data breach to the supervisory authority." 2016. https://gdpr-info.eu/art-33-gdpr/.

[6] Regulation (EU) 2016/679 (General Data Protection Regulation), Article 4(12), definition of "personal data breach." 2016. https://gdpr-info.eu/art-4-gdpr/.

[7] Directive (EU) 2022/2555 (NIS2 Directive), Article 23, "Reporting obligations." 2022. https://www.springlex.eu/en/packages/nis2/nis2-directive/article-23/.

[8] Jenner & Block. "Client Alert: New York Issues Significant Amendments to its Forward-Leaning Cyber Regulations." Summary of the Second Amendment to 23 NYCRR Part 500. 2023. https://www.legal500.com/intelligence/united-states/technology/client-alert-new-york-issues-significant-amendments-to-its-forward-leaning-cyber-regulations.

[9] Linear. "Linear incident on Jan 24th, 2024." Postmortem. 2024. https://linear.app/now/linear-incident-on-jan-24th-2024.

[10] Fortune. "AI coding tool Replit wiped database, called it a 'catastrophic failure'." July 23, 2025. https://fortune.com/2025/07/23/ai-coding-tool-replit-wiped-database-called-it-a-catastrophic-failure/.

[11] TechTarget. "The CDK Global outage: Explaining how it happened." Citing Anderson Economic Group's loss estimate. 2024. https://www.techtarget.com/whatis/feature/The-CDK-Global-outage-Explaining-how-it-happened.

[12] FOX 5 New York. "Car dealership losses from CDK software outage could soon reach $1 billion, study finds." Reporting Anderson Economic Group estimates. July 3, 2024. https://www.fox5ny.com/news/car-dealership-losses-from-cdk-software-outage-could-soon-reach-1-billion-study-finds.

[13] The Register. "Human error a factor in data center outages." Coverage of Uptime Institute's Annual Outage Analysis 2025. May 7, 2025. https://www.theregister.com/2025/05/07/human_error_a_factor_in.

[14] IAPP. "State Data Breach Notification Chart." https://iapp.org/resources/article/state-data-breach-notification-chart.

[15] National Institute of Standards and Technology. "NIST Revises SP 800-61: Incident Response Recommendations and Considerations for Cybersecurity Risk Management." April 3, 2025. https://nist.gov/news-events/news/2025/04/nist-revises-sp-800-61-incident-response-recommendations-and-considerations.

[16] U.S. Securities and Exchange Commission. Form 8-K Item 1.05, as summarized in Hunton Andrews Kurth, "SEC Finalizes Cybersecurity Disclosure Requirements." 2023. https://www.hunton.com/insights/legal/sec-finalizes-cybersecurity-disclosure-requirements.

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.