



Countermove

Meridian Commerce

Disaster Recovery Exercise Report

EXERCISE TYPE

Disaster Recovery

SCENARIO

Ransomware or Wiper

DATE OF EXERCISE

5 October 2026

CONFIDENTIAL

This report describes the response capability of Meridian Commerce, including weaknesses identified during the exercise. It is intended solely for the recipient organization and its advisers, and must not be distributed further without written consent. Handle and store it accordingly.

Contents

Executive Summary	3
Key Takeaways	4
Prioritised Actions	4
Readiness Scorecard	6
Plan Assessment	8
How This Is Scored	10
Turn-by-Turn History	12



Meridian Commerce

Ransomware or Wiper

DATE

5 October 2026

SEVERITY

High

WHEN IT STARTS

Weekday, middle of the night

WHERE IT STARTS

Confirmed, awaiting declaration

PARTICIPANTS

CTO (DR Lead) · Head of Infrastructure · Lead DBA · CISO · Engineering Manager, Commerce · COO

OVERALL READINESS

87 %

TURNS COMPLETED

9

READINESS BAND

Operational

PLAN & DOCUMENTATION

83 %

TEAM EXECUTION

89 %

LIKELY EFFECTIVENESS

89 %

Overall readiness is the average of these three. Plan & documentation grades the organization's written basis, which exists before the exercise begins; team execution is what the participants did with it; likely effectiveness is whether that would have worked at the required scale and speed.

Executive Summary

Meridian Commerce conducted a Disaster Recovery tabletop exercise on 5 October 2026, with six participants covering the CTO, Head of Infrastructure, Lead DBA, CISO, Engineering Manager for Commerce Platform, and COO. The exercise tested the organisation's response to a ransomware and wiper-style attack across the primary AWS environment, working through detection, containment, DR activation, backup selection, clean-room validation, and formal closure. The organisation performed at an operational level of readiness, with real strengths and one significant structural gap that should be addressed before the next incident. The exercise revealed clear strengths in how the team activated: escalation followed the documented chain, the CTO invoked the plan with a written basis, and the recovery sequence was applied correctly and consistently throughout - including the rigorous application of the restore-source integrity rule that correctly ruled out two compromised snapshots. The most consequential gap the exercise exposed is the pre-existing and still-open breach of the Checkout and Payments RTO: the last measured restore time of three hours fifteen minutes already exceeds the approved two-hour target, and this incident confirmed that the gap is real under operational conditions. The absence of named deputies for the Lead DBA and CISO also created live pressure during recovery. Neither gap is an immediate crisis, but both need concrete remediation plans assigned and tracked to closure.

Key Takeaways

- The activation and escalation chain held under genuine out-of-hours pressure: the SOC reached the CTO by personal phone, the incident was formally declared before the DRP was activated, and the documented basis for activation was recorded in the ticket as the plan requires.
- The restore-source integrity rule was applied correctly and consistently throughout recovery, with the team ruling out two snapshots that postdated the earliest compromise indicator and switching to credible alternatives rather than accepting convenience over security.
- The Checkout and Payments RTO is structurally unachievable under current restore timings: open remediation item DR-114, which predates this incident, was confirmed live under exercise conditions and represents the organisation's most urgent DR risk.
- The absence of named deputies for the Lead DBA and the CISO created real bottlenecks - the CISO became a single point of failure for clean-room sign-off, and the Lead DBA's unavailability on day two blocked the Vault rebuild until a substitute key holder was approved out of process.
- The organisation has no on-call rota and no out-of-hours authority: every material decision on the first night required waking senior leaders by personal phone, and the gap between detection and the first authorised containment action was widening until the CTO was reached.

Prioritised Actions

Ordered by urgency. High marks something that would cause a recovery to fail or an obligation to be missed; medium a real weakness with a workaround; low is worth doing once the rest is in hand.

HIGH

Remediate DR-114 by redesigning the Checkout and Payments restore procedure to achieve the approved two-hour RTO, conduct a full-environment timed restore test before Q3 2026, and formally close the item with documented results signed off by the Head of Infrastructure and CTO.

The exercise confirmed that the measured restore time of three hours fifteen minutes is a real operational ceiling, not a test anomaly; under incident conditions with mandatory CISO validation added, the actual elapsed time was approximately eleven hours, making the approved RTO effectively unachievable without structural change.

HIGH

Name and document a standing deputy for both the Lead DBA and the CISO in the DR plan, confirm their acceptance of the role in writing, and include them in the next tabletop exercise so their authority is exercised before it is needed.

The CISO's lack of a named deputy created a single point of failure for all clean-room sign-offs during the exercise, and the Lead DBA's unavailability on day two blocked the Vault rebuild until an out-of-process substitution was approved; both gaps are already noted in the plan as unresolved.

HIGH

Establish a formal out-of-hours authority and on-call rota covering at minimum the CTO (or deputy DR Lead) and CISO, so that containment and DR activation decisions can be authorised within fifteen minutes of a critical alert without relying on personal phone calls to sleeping senior staff.

The exercise showed that the twenty-four to thirty-one minute gap between SOC detection and the first senior decision-maker being reached was a structural vulnerability; in a faster-moving incident, or one where the CTO's phone went unanswered, that gap could prevent timely containment.

MEDIUM

Add the Vault unseal key holder list and the two-of-three threshold to the plan's main body (not only an appendix), include a procedure for substituting a key holder when one is unavailable, and verify the list is current as part of the quarterly plan review.

The exercise showed that the Vault rebuild was blocked because the plan's key-holder arrangement had no contingency for one holder being unavailable; the team resolved it correctly but only through an unplanned substitution that required CISO approval without a documented process.

MEDIUM

Define and document a formal DR closure checklist - including the CTO's declaration trigger, CISO and COO countersignature requirements, open-item recording format, and the criteria for confirming the recovery VPC is decommissioned after failback - and attach it as a standing appendix to the DR plan.

The exercise produced a well-executed closure but the process was constructed in real time; the failback to the original production environment was recorded as 'next week' with no named owner, no stated completion criteria, and no checkpoint date, leaving the organisation in a non-redundant topology without a formal forcing function.

LOW

Add an explicit step to the recovery runbook requiring that all application servers provisioned during a DR recovery event are verified against the CrowdStrike host inventory before being placed behind a load balancer, with the CISO's delegate sign-off recorded in the ticket.

Two product-catalog servers were serving live traffic without EDR coverage for several hours on day two because they were provisioned late in the recovery sequence and the host-list check was not performed until the following morning; the gap was caught and resolved correctly but only by a manual audit.

Readiness Scorecard

	DOC	EXEC	EFFECT	
Activation & Escalation 10%	4	4	4	4.0
The SOC escalated to the CTO via personal phone at 02:24; the CISO formally declared a SEV1 incident at 02:33; the CTO activated the DRP at 02:41 with a documented basis - the known RTO breach risk from DR-114 and the untrusted production environment - exactly as the plan requires, naming the VP Engineering as deputy DR Lead.				
Roles & Deputies 10%	2	3	3	2.7
The team correctly named the VP Engineering as deputy DR Lead and substituted the VP Engineering as third Vault unseal key holder when the Lead DBA was unavailable; however, the plan names no deputy for the Lead DBA and no deputy for the CISO, gaps the team identified and logged as improvement items rather than resolving them within the exercise.				
Critical Services, RTO & RPO 15%	4	3	3	3.3
The team correctly referenced each service's approved RTO and RPO - including the known pre-existing breach on Checkout and Payments from DR-114 - briefed the COO with quantified data loss (roughly eight hours against a 15-minute RPO) and RTO miss, and recorded actuals in the closure report; Checkout and Payments ultimately recovered in approximately eleven hours against a two-hour RTO, reflecting the unresolved DR-114 gap rather than a failure of understanding.				
Dependencies & Recovery Sequence 15%	4	4	4	4.0
The team followed the plan's recovery sequence throughout: networking and identity first, then Aurora, then Checkout and Payments, then Product Catalog; they correctly caught that the Elasticsearch S3 snapshot postdated the compromise indicator and switched to an index rebuild from Aurora, and confirmed Order Fulfillment was reconnected with a clean EDR sweep before the COO's go/no-go.				
Backup, Restore & Data Integrity 15%	3	4	4	3.7
The team applied the plan's restore-source condition rigorously: the 00:00 Aurora snapshot was correctly ruled out once the CloudTrail review moved the earliest compromise indicator to 23:41, the 18:00 cross-region eu-west-1 snapshot was used instead, the tainted Elasticsearch snapshot was discarded in favour of an Aurora-sourced index rebuild, and the Aurora cross-region snapshot and test restore were completed against the documented runbook and verified before DR closure; documentation scores 3 rather than 4 because the pre-incident restore time for Checkout and Payments was already known to exceed its RTO (open item DR-114).				
Identity, Privileged Access & Clean Room 10%	3	4	4	3.7
Vault was sealed and all tokens and leases revoked at 03:08; recovery proceeded exclusively via break-glass accounts and freshly generated credentials; the CISO invoked the plan's 'CISO or delegate' language to name the senior security engineer as delegate for EDR sweeps, retaining sole sign-off authority; all three reconnection conditions (no persistence, credentials rotated, restore source predating 23:41) were documented and evidenced before each system-owner approved reconnection; Vault was subsequently rebuilt with a two-of-three unseal key threshold and interim credentials retired.				

DOC EXEC EFFECT

Validation & Reopening 10%

The CISO signed off identity, Aurora, and Checkout sequentially with evidence attached to the ticket and system-owner approval from the Head of Infrastructure, Lead DBA, and Engineering Manager for Commerce Platform; the flagged Checkout AMI was preserved for forensics and a replacement AMI verified via GitHub audit log provenance; the COO gave the formal business go/no-go at 13:02 and the storefront reopened at 13:34 with all three plan conditions met for each tier.

3 4 4 **3.7****Stakeholder Communication** 5%

The team shifted to the 'Meridian DR Bridge' Signal group immediately on identifying that Okta-backed tools could not be trusted; every material decision was logged in the incident ticket with time and approver; the COO was briefed with bracketed time-to-open estimates and updated when the 11:00 target slipped, and customer service was opened early to manage inbound volume.

3 3 3 **3.0****Supplier Preparedness** 5%

The plan documents AWS Enterprise Support, Stripe Slack Connect and status monitoring, and the colocation NOC SLA; the team referenced Stripe's external status page implicitly through the recovery sequence and noted the Elasticsearch cluster was self-hosted, but no supplier escalation actions appear in the team's submissions - AWS Severity 1 support, Stripe contact, and the colocation NOC were not exercised during this scenario.

3 2 2 **2.3****Documentation & Regulatory Compliance** 5%

The team identified the 72-hour GDPR notification clock from the moment of awareness, set a Thursday 12:00 checkpoint with Legal, confirmed the obligation at that meeting, assigned the General Counsel as owner with a Friday 18:00 deadline inside the window, and recorded the open item in the closure report; the DR closure report was countersigned by the CISO and COO and directed to the CEO and audit committee.

3 3 3 **3.0**

Doc = Documentation Exec = Execution Effect = Effectiveness, each 0-4. See "How This Is Scored" later in this report.

Plan Assessment

The exercise has been assessed against the organization's own disaster recovery plan, provided in advance. The two sections below are separate judgements: the first covers the plan as a document, the second how closely the team worked from it during the exercise.

The Plan as a Document

COVERED WELL

- The activation criteria and authority chain are precisely defined: the CTO holds sole activation authority with a clear two-step succession to the VP Engineering and then the Head of Infrastructure, and the required documented basis - which RTO is at risk and why normal recovery is insufficient - is specific enough to be actionable.
- The critical services table gives named RTO and RPO targets for every service tier, including the dependency on external Stripe for payments and the instruction to monitor status.stripe.com rather than attempt internal recovery.
- The recovery sequence is explicitly ordered with dependency logic named: identity before database before checkout, with Checkout and Payments correctly identified as depending on the Customer Account Database.
- The secure recovery requirements are unambiguous and directly actionable: three conditions must each be evidenced before reconnection, and the plan names the CISO or delegate as the signing authority, which proved workable under exercise pressure.
- Fallback communication arrangements are documented in sufficient detail: the Signal group name, the break-glass conference bridge, and the Facilities safe location for credentials are all specified.

GAPS IN THE PLAN

- The Lead DBA and the CISO have no named deputies in the plan; both roles carry sole-owner authority over critical steps (database recovery sign-off and clean-room validation respectively), and the absence of deputies was a live operational pressure point during the exercise.
- The Vault unseal key holder list and the procedure for substituting a key holder when one is unavailable are described only at a high level; the exercise showed this is insufficient when one of three holders is unavailable and a rebuild is time-critical.
- The plan records that the last measured end-to-end restore time for Checkout and Payments was three hours fifteen minutes against a two-hour RTO, and that DR-114 is open with a Q3 2026 target - meaning the plan itself documents a known structural breach of the most revenue-critical RTO with no interim mitigation stated.
- The plan does not specify a DR formal closure procedure: there is no checklist, no countersignature requirement, and no criteria for confirming the recovery environment is decommissioned after fallback to the original production infrastructure.
- The plan does not address the scenario in which a nightly or scheduled backup snapshot itself postdates the earliest compromise indicator - the exercise required the team to work this out from the restore-source condition language, and a brief worked example or explicit cross-reference to the compromise timestamp check would make the runbook more robust under time pressure.

How Well the Team Followed It

STRONG

The team worked from the plan throughout the exercise: activation authority, recovery sequence, restore-source conditions, secure reconnection criteria, fallback communication, and closure countersignature structure all reflected the plan's documented requirements. The single most consequential point at which the plan directly shaped the outcome was the restore-source integrity rule - the team cited it explicitly to rule out the 00:00 Aurora snapshot once the CloudTrail review pushed the earliest indicator to 23:41, and again to discard the Elasticsearch S3 snapshot that postdated the compromise window. Where the plan was silent - notably on the CISO deputy and the Vault key-holder substitution - the team improvised correctly and flagged the gaps as improvement items, which is appropriate rather than a failure of adherence.

FOLLOWED

- Activation authority and documented basis: the CTO activated the DRP with a written basis in the ticket referencing the specific RTO at risk and the untrusted production environment, exactly as the plan requires.
- Recovery sequence: the team followed the plan's ordered sequence - networking and identity, then Aurora, then Checkout, then Product Catalog, then Order Fulfillment - throughout the recovery, including holding downstream services until upstream dependencies were signed off.
- Secure reconnection: all three plan conditions were evidenced and recorded before any system was reconnected to the recovery network, with system-owner approval logged alongside CISO sign-off.
- Fallback communication: the team switched to the Meridian DR Bridge Signal group in turn one when Okta-backed tools were compromised, as the plan provides.
- Break-glass credential procedure: Vault was sealed and all issued credentials treated as compromised; recovery used only break-glass accounts and freshly generated credentials, consistent with the plan's secure recovery requirements.

DEPARTED FROM

- The plan names no deputy for the CISO; the team improvised by naming the senior security engineer as delegate for EDR sweeps under the 'CISO or delegate' language, which was a sound reading of the plan but went beyond what the plan explicitly provides.
- The Vault key-holder substitution - replacing the Lead DBA with the VP Engineering when the Lead DBA was unavailable - had no documented procedure; the team resolved it correctly but the decision required real-time CISO approval without a plan basis.
- The closure report structure, including CISO and COO countersignatures and distribution to the CEO and audit committee, was constructed during the exercise rather than drawn from a documented closure procedure in the plan.

How This Is Scored

Ten areas are assessed. Each is scored on three separate axes, and the area score is their average. Areas carry different weights, and the weighted total gives the overall readiness percentage on the first page.

The Three Axes

Documentation	Whether a written, approved and findable basis exists - a plan, runbook or procedure the team could actually point to during the exercise.
Execution	Whether the team carried the area out: decisions made, owners named, the defined sequence actually followed.
Effectiveness	Whether what they did would have worked in practice, at the scale and speed this incident demanded.

The 0-4 Scale

0 - Not established	No known solution, documentation or owner.
1 - Ad hoc	An approach exists, but it is person-dependent, undocumented or unapproved.
2 - Partially established	A process exists but is incomplete, outdated, untested or unclearly owned.
3 - Established	Documented, known and usable, with minor gaps.
4 - Validated	Established, and demonstrated convincingly during the exercise - the documented basis was known, the owner named, and it was applied correctly with no material gap left open.

Weighting

Weights reflect how much each area determines whether recovery succeeds at all. Restoring data and sequencing dependencies carry more weight than stakeholder communication - not because communication is unimportant, but because a failure there is recoverable and a failure in the former is not. Overall readiness is the weighted sum expressed as a percentage of the maximum possible score.

Readiness Bands

Critically Inadequate (0-39%)	Recovery would not succeed. Fundamental capabilities are missing rather than weak.
Major Deficiencies (40-59%)	Recovery would likely fail or overrun badly. Several core areas need building, not tuning.
Partially Operational (60-74%)	Recovery would probably succeed for some services, but not within the times the business expects.
Operational (75-89%)	Recovery is credible and likely to succeed, with identified gaps worth closing.
Robust & Validated (90-100%)	Recovery is documented, exercised and evidenced. Remaining work is refinement.

Critical Stop Criteria

Seven conditions are treated as disqualifying regardless of the weighted score. If any was observed during the exercise, the readiness band is capped at Major Deficiencies - a high average cannot compensate for one of these being absent.

What Each Area Covers

Activation & Escalation (10%)	Who can declare a disaster and on what basis, how that decision is escalated, and how quickly the plan is actually invoked once the trigger is met.
Roles & Deputies (10%)	Whether the roles the plan depends on are staffed and understood, and whether named deputies can act when the primary person is unreachable.
Critical Services, RTO & RPO (15%)	Whether business-critical services are identified and prioritised, with recovery time and recovery point objectives the business side has actually approved.
Dependencies & Recovery Sequence (15%)	Whether the order of recovery reflects real technical dependencies - identity and network before applications - rather than restoring whatever is loudest first.
Backup, Restore & Data Integrity (15%)	Whether backups exist, remain reachable when the primary environment is not, and can be restored to a known-good point within the required time.
Identity, Privileged Access & Clean Room (10%)	Whether identity services and privileged access can be re-established independently of the compromised environment, including a clean room where one is needed.
Validation & Reopening (10%)	Whether restored systems are verified before being returned to users, and whether system and business owners formally approve reopening.
Stakeholder Communication (5%)	Whether staff, customers, suppliers and management are kept informed on a defined cadence, through channels that survive the outage itself.
Supplier Preparedness (5%)	Whether critical suppliers are reachable under contract, understand their obligations during a disaster, and can be escalated to outside business hours.
Documentation & Regulatory Compliance (5%)	Whether decisions, timestamps and evidence are recorded as the incident unfolds, and whether regulatory notification duties are identified and met within their deadlines.

Turn-by-Turn History

Phase 1 of 5 · Declaration & activation Wednesday, Day 1, 02:17

Ransomware Encryption Detected Across Production Infrastructure

SITUATION

At 02:17 on Wednesday morning, the 24/7 SOC raised a critical alert after CrowdStrike EDR fired simultaneously on multiple production hosts: ransomware with wiper-style behaviour is actively encrypting files across the primary AWS environment. The SOC has confirmed propagation across the application server fleet and has observed encryption activity reaching the self-hosted Elasticsearch cluster nodes. The Aurora PostgreSQL primary instance is showing anomalous API calls consistent with credential abuse, and the SOC cannot yet confirm whether the database tier has been directly touched or whether the activity stopped at the application layer. No system has been isolated or shut down. No credentials have been rotated. Nothing has been contained.

The blast radius as understood right now: the application server fleet is compromised and partially encrypted, the three-node Elasticsearch cluster is affected, and the Aurora primary is at elevated risk with suspicious activity in its access logs. Checkout and Payments, Product Catalog and Search, and Order Fulfillment integration are all effectively unavailable or untrusted. The internal ERP and Google Workspace appear unaffected at this time, but lateral movement has not been ruled out. HashiCorp Vault's integrity is unknown - if it has been reached, every secret it holds must be treated as compromised. The earliest confirmed indicator of malicious activity in the SOC logs is 01:54, giving an approximate 23-minute window of active encryption before detection.

There is no on-call rota and no one holds out-of-hours authority tonight. The SOC is operational and has a clear picture of what is burning, but every senior decision-maker - the CTO, the VP Engineering, the Head of Infrastructure - is at home asleep and must be reached by ringing personal phones. The SOC has the technical picture; it does not have the authority to act on it. The clock is running: the Checkout and Payments RTO is 2 hours, and the last measured end-to-end restore time for that service was 3 hours 15 minutes - already beyond the approved target even before accounting for the mandatory CISO security validation that must precede any reconnection. Open remediation item DR-114, which was meant to close that gap, has a target of Q3 2026.

The SOC has escalated and is awaiting instruction. No incident has been formally declared. The DRP has not been activated. No external party has been engaged. The organisation is at the moment between detection and decision - and the gap between what is known and what has been authorised to happen is widening with every minute.

ACTIONS TAKEN

- The SOC analyst calls the CTO's personal phone and reaches them at 02:24; the VP Engineering and CISO are on the call by 02:31. The CISO declares a SEV1 incident at 02:33 and acts as Incident Commander for now.
- Containment first: the CISO authorises the SOC to network-contain every host with active encryption alerts through CrowdStrike, and to isolate the three Elasticsearch

- nodes. The Aurora primary is not shut down, but its security group is restricted to the DBA's bastion only. Every action is logged in the incident ticket with time and approver.
- At 02:41 the CTO activates the DRP. The documented basis in the ticket: Checkout and Payments cannot meet its 2-hour RTO by normal recovery, the last measured restore time was 3h15m (DR-114), and the production environment can no longer be trusted. The VP Engineering is named deputy DR Lead for the night.
 - Vault is treated as compromised until proven otherwise. We move coordination to the "Meridian DR Bridge" Signal group, since we can't trust Okta-backed Slack right now.

EVALUATION

WHAT WORKED

The team moved quickly and correctly in the right sequence: the SOC escalated to the CTO via personal phone, the CISO formally declared a SEV1 incident before the CTO activated the DRP, and the DRP activation at 02:41 is grounded in a documented basis - the known RTO breach risk from DR-114 and the untrusted production environment - exactly as the plan requires. Shifting coordination to the Meridian DR Bridge Signal group is the right call given that Okta-backed tools cannot be trusted.

MAIN GAP

The most significant gap is HashiCorp Vault: the team correctly flags it as compromised but takes no immediate action on it. Vault likely holds secrets underpinning every service in the recovery sequence - database credentials, API keys for Stripe, signing keys - and until its state is confirmed or its contents are rotated, you cannot safely use any credential it has ever issued. Who owns that assessment right now, and what is the action: take Vault offline, freeze secret issuance, or begin an emergency rotation? The decision needs an owner and a timestamp in the incident ticket, not a holding note.

GOING FORWARD

Going into containment and stabilization the team has a declared incident, an activated DR plan, and a named command structure - that is a solid foundation. The open question about Vault, and whether the Aurora primary can be trusted at all given the anomalous API calls, will dominate the next phase: you cannot begin any recovery sequence on a database tier whose access logs are still unexplained.

Phase 2 of 5 · Containment & stabilization Wednesday, Day 1, 03:05

Vault Still Live, Aurora Access Logs Escalate

SITUATION

The SOC has now pulled the full Aurora access log from the last two hours and sent it to the DR Bridge: between 01:57 and 02:14, an API key that traces back to HashiCorp Vault issued five IAM-authenticated calls to the Aurora cluster - including a `ModifyDBInstance` call that temporarily disabled deletion protection and one `CreateDBSnapshot` call with an unfamiliar snapshot identifier. The snapshot has since been deleted, but its existence in the log means the database tier must now be treated as directly touched, not merely at risk. Meanwhile, Vault is still running and still issuing credentials: no one has taken it offline, frozen secret issuance, or begun an emergency rotation in the 24 minutes since it was flagged as compromised. The Lead DBA is on the DR Bridge but has not yet received a confirmed instruction on whether to restore from the most recent Aurora snapshot or

to treat every snapshot taken after 01:54 as potentially tainted.

ACTIONS TAKEN

- Vault: the CISO orders it sealed at 03:08 and all tokens and leases revoked. We accept that this blocks every service that depends on Vault until it is rebuilt; recovery will use two break-glass accounts and freshly generated credentials, never anything Vault issued before tonight. Owner: Head of Infrastructure, logged in the ticket at 03:08.
- Aurora: we treat the database tier as touched. The Lead DBA's instruction is to restore from the cross-region snapshot in eu-west-1 taken at 00:00 - the last six-hourly snapshot before the first indicator at 01:54. Before committing to it, the DBA and the SOC check CloudTrail for the 72 hours before 01:54 for any earlier use of that Vault-issued key; if anything turns up, we fall back to the 18:00 snapshot. The DBA restores into an isolated VPC, nothing reconnects to production.
- Roles: the CISO hands the Incident Commander role to this week's security IC from PagerDuty, so the CISO can concentrate on clean-room validation, which only the CISO can sign off.

EVALUATION

WHAT WORKED

The team closed the Vault gap decisively: sealing it at 03:08 with all tokens and leases revoked, naming the Head of Infrastructure as owner, and logging the decision with a timestamp is exactly what the plan and the situation required. The decision to restore Aurora from the 00:00 eu-west-1 cross-region snapshot - with a CloudTrail check for earlier Vault key usage before committing - shows sound RPO reasoning and correctly keeps the restore target inside an isolated VPC before any production reconnection.

MAIN GAP

The most significant gap is the CISO role handover: the plan names no deputy for the CISO, so transferring Incident Commander duties to "this week's security IC from PagerDuty" needs to be explicit about what authority that person actually holds. The CISO retains the sole sign-off on clean-room validation before any system reconnects; the incoming IC cannot authorise that step, and if there is any ambiguity about that boundary tonight it will stall reconnection at the worst possible moment.

GOING FORWARD

Going into recovery and restoration the team has contained the blast radius, sealed the compromised credential store, and identified a credible Aurora restore point - a genuinely stable platform to recover from. The Checkout and Payments RTO clock is already past the approved two-hour window, so the recovery sequence must begin immediately and the known gap from DR-114 will determine how that conversation with the COO is framed.

Phase 3 of 5 · Recovery & restoration Wednesday, Day 1, 04:15

CloudTrail Check Delays Aurora Restore Decision

SITUATION

The CloudTrail review is taking longer than expected: the SOC has found two earlier uses of the same Vault-issued IAM key - a `DescribeDBSnapshots` call at 23:41 on Tuesday and a `ListBackups` call at 23:49 - both more than two hours before the confirmed

compromise window of 01:54. Neither call modified anything, but they confirm the attacker had read access to the Aurora backup inventory before the active encryption phase began. The Lead DBA must now decide whether the 00:00 snapshot can still be trusted as a restore source, or whether the 18:00 snapshot (Tuesday, Day 1) becomes the safe fallback - adding roughly 6 hours of data loss against the 15-minute RPO target. Meanwhile the Aurora restore into the isolated VPC has not yet started: the DBA has been waiting on the CloudTrail result and has not spun up the target environment, consuming over an hour of the recovery window with no restore work begun.

ACTIONS TAKEN

- The plan is explicit: the restore source must predate the earliest known compromise indicator. That indicator is now 23:41, so the 00:00 snapshot no longer qualifies, read-only or not. The CTO authorises the Lead DBA at 04:20 to restore from the 18:00 cross-region snapshot into the isolated VPC, starting now. We stop waiting on further analysis; the SOC keeps reviewing CloudTrail in parallel, and if it finds anything before 18:00 we reassess.
- The CTO briefs the COO at 04:30: roughly eight hours of orders and account changes (18:00 to 02:17) will be lost against a 15-minute RPO, and Checkout and Payments will miss its 2-hour RTO by several hours. The COO decides to keep a maintenance page up and opens customer service early at 07:00.
- The authority boundary is written into the ticket at 04:25: the PagerDuty IC coordinates the response and tracks actions; only the CISO can sign off clean-room validation and reconnection to production. Both acknowledged on the DR Bridge.

EVALUATION

WHAT WORKED

The team made the right call on the restore source: the plan's own language - restore source must predate the earliest known compromise indicator - is unambiguous, and the CTO authorising the 18:00 snapshot at 04:20 rather than waiting for further analysis is sound incident command. Briefing the COO at 04:30 with a clear, quantified picture of data loss and RTO breach, and getting a concrete business decision (maintenance page, early customer service opening) out of that conversation, is exactly what the COO's role exists for.

MAIN GAP

The most significant gap is that the team still has not started any of the other recovery sequence steps while the Aurora restore begins. Networking and DNS, and the identity layer, are prerequisites for everything downstream - Checkout and Payments cannot reconnect until identity is clean and DNS is pointed at recovered infrastructure. Over an hour of the recovery window was already consumed waiting on CloudTrail; who owns those parallel workstreams right now, and has the Head of Infrastructure actually begun spinning up the isolated network environment and preparing the identity recovery?

GOING FORWARD

Going into validation and business resumption the team has a running restore from a trusted source and a clear authority boundary documented - that is a credible foundation. The pressure point will be whether the CISO can complete clean-room validation fast enough to avoid extending the outage further, and whether the end-to-end restore time for Checkout and Payments - already measured at 3h15m under DR-114 - holds or

worsens given the credential rebuild from scratch.

Phase 4 of 5 · Validation & resumption Wednesday, Day 1, 07:40

Aurora Restored, But Validation Queue Backs Up

SITUATION

The Lead DBA reports at 07:40 that the Aurora restore from the 18:00 eu-west-1 snapshot completed at 07:22 into the isolated VPC, and initial row counts and schema checks look consistent. However, the CISO's clean-room validation queue has now stacked up: the identity layer (Okta and AD Connector) was rebuilt by the Head of Infrastructure and is awaiting CISO sign-off, the Aurora instance is awaiting CISO sign-off, and the first application server AMI for the Checkout and Payments tier has just been provisioned from a pre-incident immutable snapshot and is also queued for EDR sweep and sign-off - all three are sitting unvalidated at the same time. The CISO is one person with no named deputy, and the plan requires their explicit sign-off before any of these systems reconnect to production. Customer service opened at 07:00 as planned and is already fielding a high volume of calls; the COO is asking for a realistic time-to-open for the storefront.

ACTIONS TAKEN

- The plan's secure recovery requirement says "the CISO or delegate". The CISO uses that: at 07:45 she names the senior security engineer on shift as delegate in writing in the ticket, for running the EDR sweeps and collecting the evidence on all three items. The CISO keeps the sign-off itself.
- Sequence follows the plan's recovery order: identity first (Okta and AD Connector), then Aurora, then the Checkout application tier. The delegate starts the EDR sweep on the Checkout AMI in parallel so its evidence is ready when its turn comes. Each sign-off checks the three plan conditions: no persistence found by EDR, credentials rotated, restore source predating 23:41.
- The CTO gives the COO a time-to-open of 11:00, with 12:00 as the worst case, and says it will be confirmed at 10:00. Customer service uses that window, not a promise.

EVALUATION

WHAT WORKED

The team handled the CISO bottleneck well: invoking the plan's "CISO or delegate" language to name the senior security engineer as delegate in writing at 07:45 is exactly right, and sequencing validation in recovery order - identity, then Aurora, then the Checkout application tier - with the EDR sweep on the application tier running in parallel shows mature thinking under pressure. Giving the COO a bracketed time-to-open (11:00 confirmed by 10:00, 12:00 worst case) rather than a hard promise is the correct way to manage business expectations without overcommitting.

MAIN GAP

The most significant gap is that sign-off on each item still depends on a single person - the CISO - and the plan's three conditions for reconnection have not yet been documented as met for any of the three queued systems. The EDR sweep results, the credential rotation confirmation, and the restore-source date check each need an explicit record before the CISO signs: who holds those evidence items, where are they logged, and what happens at 10:00 if one of the three conditions has not cleared for the identity

layer?

GOING FORWARD

Going into the final stretch of this phase, the team needs to work through those sign-offs in sequence and get explicit system-owner approval - from the Head of Infrastructure for identity, the Lead DBA for Aurora, and the Engineering Manager for Commerce Platform for the Checkout tier - before the COO's business go/no-go for storefront reopening. The 10:00 checkpoint is the forcing function: be ready with documented evidence for each system or escalate the 12:00 worst-case to the COO before that window closes.

Phase 4 of 5 · Validation & resumption Wednesday, Day 1, 09:48

EDR Sweep Finds Residue on Checkout AMI

SITUATION

At 09:48 the senior security engineer delegate reports that the EDR sweep of the Checkout and Payments application server AMI has returned a low-confidence persistence indicator: a scheduled-task artefact in a temp directory that does not match any known application baseline. The identity layer and Aurora instance have both cleared their three-condition checks and are ready for CISO sign-off - but the CISO must now decide whether to sign off identity and Aurora and let the Engineering Manager for Commerce Platform stand up the Checkout tier on a replacement AMI, or hold everything while the artefact is investigated. The 10:00 COO checkpoint is twelve minutes away, and the 11:00 storefront target depends on the Checkout application tier being signed off and reconnected in the next hour.

ACTIONS TAKEN

- The CISO signs off identity and Aurora at 09:55: both cleared all three conditions independently of the Checkout tier, and the evidence (EDR reports, rotation log, snapshot timestamps) is attached to the ticket. The Head of Infrastructure and the Lead DBA give system-owner approval for their tiers. Both reconnect in the isolated recovery network; nothing is customer-facing yet.
- The flagged Checkout AMI is set aside and preserved for forensics, not deleted. The Engineering Manager, with the CISO's agreement, switches to the nightly AMI from the night before, taken at 02:00 on Tuesday, well before 23:41, and deploys the application code on top from a known commit in Git. The delegate sweeps it before anything else happens.
- At 10:00 the CTO tells the COO: 11:00 is no longer realistic; the new estimate is 12:30, confirmed by 12:00. The COO keeps the maintenance page and updates customer service.

EVALUATION

WHAT WORKED

The CISO's decision to sign off identity and Aurora independently at 09:55 - with evidence attached to the ticket and system-owner approval from the Head of Infrastructure and the Lead DBA - is the right call: the plan's three conditions were met for those tiers on their own merits, and holding them hostage to an unrelated AMI artefact would have delayed the entire recovery chain unnecessarily. Preserving the flagged AMI for forensics rather than deleting it shows disciplined evidence handling.

MAIN GAP

The most significant gap is the replacement AMI selection: switching to the Tuesday 02:00 AMI and deploying application code from a known Git commit is a reasonable approach, but the team has not stated what "known commit" means in practice - who identified it, how it was verified against the production release at the time of the incident, and whether the application code layer itself was in scope during the attacker's read access window from 23:41 onward. If the commit was pulled from a repository that could have been touched, the clean-room condition applies to it too.

GOING FORWARD

Going into the 12:00 confirmation window, the EDR sweep of the replacement Checkout AMI is the critical path item: if it clears, the CISO can sign off, the Engineering Manager can reconnect to the recovery network, and the COO has a defensible 12:30 storefront opening. The team should also confirm that Product Catalog (Elasticsearch) and Order Fulfillment integration are staged and ready so that business resumption does not stall immediately after Checkout goes live.

Phase 4 of 5 · Validation & resumption Wednesday, Day 1, 11:52

Checkout AMI Clears, Elasticsearch Blocks Reopening

SITUATION

The EDR sweep of the replacement Checkout AMI completes clean at 11:47, and the delegate confirms all three plan conditions are met: no persistence indicators, credentials rotated, restore source predates 23:41. The CISO is ready to sign off the Checkout tier - but the Engineering Manager for Commerce Platform has just flagged that the Product Catalog is not ready: the Elasticsearch cluster restore from its nightly S3 snapshot has been running since 10:30 and is still in progress, with an estimated thirty minutes remaining. Product Catalog depends on both Elasticsearch and Aurora, and Checkout depends on Product Catalog for search-driven landing pages; without it, customers can browse and pay only if the storefront is opened in a degraded mode with direct-link product access only. The 12:00 COO confirmation window is eight minutes away, and the 12:30 reopening target is now contingent on a business decision the COO has not yet been asked to make.

ACTIONS TAKEN

- Git commit: the Engineering Manager confirms the deployed commit is the signed release tag v4.18.2 from Monday. The repository is hosted on GitHub, outside the AWS account, with branch protection, and the audit log shows no pushes or tag changes since Monday afternoon. The CISO accepts that as provenance and signs off the Checkout tier at 12:06.
- Elasticsearch: before anything else, the Head of Infrastructure checks when the nightly S3 snapshot was taken. It ran at 01:00 on Wednesday - after 23:41 - so it does not meet the plan's condition. We stop that restore and rebuild the index from Aurora instead, which is the catalogue's source of truth. Estimated 45 minutes.
- The CTO puts both options to the COO at 12:00. The COO chooses not to open in degraded mode: customers landing on pages without search would generate more calls than a short further wait. Full reopening target 13:30, confirmed at 13:00, go/no-go by the COO.

EVALUATION

WHAT WORKED

The team resolved two significant open items cleanly this turn: the Git commit provenance check (signed release tag v4.18.2, verified via GitHub audit log outside the compromised AWS account) gave the CISO a defensible basis for signing off the Checkout tier at 12:06, and catching that the Wednesday 01:00 Elasticsearch S3 snapshot postdates the 23:41 compromise indicator before committing to it was sharp - stopping that restore and switching to an index rebuild from Aurora is exactly what the plan's restore-source condition requires.

MAIN GAP

The most significant gap is Order Fulfillment and warehouse integration: it has appeared on the "not yet confirmed" list since turn 5 and has still received no status update, owner action, or sequencing decision this turn. Its RTO is 8 hours from activation at 02:41, placing it in breach at roughly 10:41 - over two hours ago - and the COO's go/no-go at 13:00 will need to cover whether the storefront can open while Order Fulfillment remains unvalidated or degraded.

GOING FORWARD

The team has now met this phase's core objectives: identity, Aurora, and Checkout are signed off with documented evidence and system-owner approval, the COO holds the go/no-go decision, and the 13:30 full reopening target is a business-owned commitment rather than a technical guess. The exercise advances to return to normal operations, carrying the Elasticsearch rebuild timeline, the Order Fulfillment gap, and the forensic investigation of the preserved Checkout AMI as live pressure items.

Phase 5 of 5 · Normal operations Wednesday, Day 1, 14:05

Storefront Reopens, But Recovery Arrangements Persist

SITUATION

The COO gave go/no-go approval at 13:02 and the storefront reopened at 13:34 - Checkout, Product Catalog, and Aurora all serving live traffic from the isolated recovery VPC rather than the original production environment. Order Fulfillment was reconnected at 13:18 after a late but clean EDR sweep, with the warehouse integration team informed of the approximately eight-hour data gap; manual order reconciliation is underway. However, the team is now operating in a temporary state: nightly AMI backups have not been re-enabled, the six-hourly Aurora snapshot schedule has not been confirmed as running against the recovered instance, CrowdStrike EDR coverage across the rebuilt application fleet has not been verified as complete, and the Elasticsearch cluster is running from a freshly rebuilt index with no snapshot schedule in place. The DR situation has not been formally closed, HashiCorp Vault has not been rebuilt or re-enabled, and the team has not confirmed whether monitoring and alerting are watching the recovery environment at the same level of coverage as production.

ACTIONS TAKEN

- Workstreams, each with an owner and a deadline in the ticket:
- Backups (Head of Infrastructure and Lead DBA): nightly AMI jobs, the six-hourly Aurora snapshots with cross-region replication and an Elasticsearch snapshot schedule re-enabled today. Done means one successful run of each, plus a test restore of the

Aurora snapshot, by Thursday 18:00.

- Coverage (CISO's delegate): CrowdStrike host list compared against the rebuilt fleet, and monitoring and alerting checked against the production baseline, by Thursday noon.
- Vault (Head of Infrastructure, CISO approves): rebuilt from scratch on Thursday, new unseal keys split between three key holders, interim secrets moved in and rotated, break-glass accounts retired by Friday.
- Failback to the original production environment: planned for next week, not under pressure.
- Regulatory: the attacker held working credentials with access to customer data. The PagerDuty IC and Legal own the personal-data assessment under the incident response plan; the 72-hour clock runs from awareness at 02:17 Wednesday, so their decision is due by Saturday 02:17, with a checkpoint Thursday 12:00.
- The CTO closes the DR situation when backups are verified, Vault is rebuilt and coverage is confirmed. The closure report records RTO and RPO actuals (Checkout back after about 11 hours against 2; about 8 hours of data lost against 15 minutes), the evidence log and open items, DR-114 first.

EVALUATION

WHAT WORKED

The team has produced a well-structured return-to-normal plan with named owners, explicit deadlines, and clear done-criteria for every workstream - backups, EDR coverage, Vault rebuild, and regulatory notification each have an accountable person and a forcing function in the ticket. The decision to defer failback to the original production environment until next week rather than rushing it is sound: the recovery VPC is serving live traffic cleanly, and a pressured failback introduces unnecessary risk at this stage.

MAIN GAP

The most significant gap is the closure trigger itself: the CTO intends to close the DR situation when backups are verified, Vault is rebuilt, and coverage is confirmed - but those three conditions complete at different times across Thursday and Friday, with no stated process for the CTO to formally declare closure. Who initiates that call, what document records it, and does the COO need to countersign given that it was a business go/no-go that opened the storefront?

GOING FORWARD

The regulatory checkpoint at Thursday 12:00 is the other live dependency: if Legal's personal-data assessment concludes that a GDPR or PCI-DSS notification is required, the DR closure report will need to record that decision and its owner before it can be filed as complete. Make sure the CTO is briefed at that Thursday checkpoint - not just the PagerDuty IC and Legal - so the closure report reflects the full picture.

Phase 5 of 5 · Normal operations Thursday, Day 2, 10:35

Vault Rebuild Reveals Unseal Key Holder Gap

SITUATION

The Head of Infrastructure begins the Vault rebuild on Thursday morning and hits an immediate blocker: the plan calls for unseal keys to be split between three key holders,

but one of the three designated holders - the Lead DBA - is currently unavailable until Friday morning due to the overnight shift worked during the incident, and the Head of Infrastructure is not willing to proceed with only two key holders. Separately, the CISO's delegate has completed the CrowdStrike host-list comparison and found that two application servers provisioned late in the recovery sequence on Wednesday are not yet enrolled in the EDR platform and are currently serving live product-catalog traffic with no endpoint coverage. The Thursday 12:00 regulatory checkpoint with Legal is ninety minutes away, and the CTO has not yet been told whether a GDPR or PCI-DSS notification obligation has been identified.

ACTIONS TAKEN

- EDR gap first, because it is live: the two unenrolled servers are pulled from the load balancer now. The remaining fleet carries the catalogue traffic comfortably at this hour. They are enrolled, swept by the CISO's delegate, and only then put back. Owner: Head of Infrastructure, logged at 10:40.
- Vault: we don't run with a weaker split, and we don't wait for one person either. The CISO approves replacing the Lead DBA as third key holder with the VP Engineering, so the rebuild proceeds today with three holders and a two-of-three threshold, as designed. The change of holder is recorded in the ticket and the plan's appendix will be updated.
- The CTO joins the 12:00 checkpoint with Legal and the PagerDuty IC. Whatever Legal decides about notification goes into the closure report with its owner and deadline.
- Closure: the CTO declares the DR situation closed in the closure report once the three conditions are met, countersigned by the CISO for the security conditions and the COO for business operations. The report goes to the CEO and the board's audit committee.

EVALUATION

WHAT WORKED

The team resolved both live blockers cleanly and quickly: pulling the two unenrolled product-catalog servers from the load balancer before enrolling them in the EDR platform is the right sequence - no unmonitored host serves live traffic even for the duration of enrolment - and substituting the VP Engineering as third Vault unseal key holder, recorded in the ticket and flagged for plan update, unblocks the rebuild without weakening the two-of-three threshold. The closure countersignature structure - CISO for security conditions, COO for business operations, report to CEO and audit committee - is appropriate and fills the gap from the previous turn.

MAIN GAP

The most significant remaining gap is the Vault rebuild completion itself: the team has approved the key-holder change and unblocked the work, but has not stated a revised completion deadline or confirmed whether the interim break-glass accounts are still the live secret source. Until Vault is rebuilt and interim credentials are retired, every service running in the recovery VPC is dependent on a credential arrangement that was designed as temporary, and that risk needs a named deadline in the ticket before the DR situation can formally close.

GOING FORWARD

The regulatory checkpoint at 12:00 is the other forcing function: the CTO's presence

ensures the closure report will reflect Legal's notification decision - but the team should confirm at that meeting whether the 72-hour GDPR clock (expiring Saturday 02:17) has generated a concrete notification obligation, so the closure report is filed complete rather than with an open item against it.

Phase 5 of 5 · Normal operations Thursday, Day 2, 15:20

Vault Rebuilt, Closure Report Hits Final Gaps

SITUATION

The Head of Infrastructure reports at 15:20 that the Vault rebuild completed at 14:55: new unseal keys split across the three holders, interim secrets migrated and rotated, and the two break-glass accounts retired and logged. However, the CTO's closure report is now stalled on two items from the 12:00 Legal checkpoint: Legal has confirmed a GDPR personal-data breach notification obligation exists - the attacker's read access to Aurora between 23:41 and 02:17 is sufficient - but Legal has not yet named a controller representative to sign the regulator submission, leaving the notification owner blank in the closure report. Separately, the nightly AMI backup job ran successfully at 02:00 this morning, but the Lead DBA has not yet confirmed a successful Aurora cross-region snapshot run or the test restore that was set as the done-criterion for the backup workstream by Thursday 18:00, and that deadline is now less than three hours away.

ACTIONS TAKEN

- Notification: this belongs to the incident response process, not the DR closure. The CEO, as the controller's representative, signs the submission to Datatilsynet; Legal drafts it. Owner: General Counsel, deadline Friday 18:00, well inside the 72 hours that end Saturday 02:17. The DR closure report records it as an open item with that owner and deadline; we don't hold the DR situation open for it.
- Test restore: the Head of Infrastructure runs the Aurora test restore today, following the documented runbook, after confirming the first cross-region snapshot completed at 12:00. The Lead DBA reviews the result on Friday. We record that this only worked because someone else could follow the runbook: the plan names no deputy for the Lead DBA, and that goes on the improvement list.
- At 18:20, with backups verified, Vault rebuilt and EDR coverage confirmed, the CTO declares the DR situation closed in the closure report. The CISO countersigns the security conditions and the COO business operations, and the report goes to the CEO and the audit committee. Open items: the notification (General Counsel, Friday), fallback to production (Head of Infrastructure, next week), forensic analysis of the preserved AMI (CISO), DR-114 and naming deputies for the Lead DBA and CISO (CTO, in the plan revision). Lessons-learned review within two weeks.

EVALUATION

WHAT WORKED

The team closed the DR situation cleanly and in the right sequence: the GDPR notification is correctly handed to the General Counsel with a Friday 18:00 deadline inside the 72-hour window, leaving the closure report to record it as an open item rather than holding the DR situation hostage to it. The Aurora test restore - completed against the documented runbook by someone other than the Lead DBA - is a genuine validation of runbook quality, and the team correctly captured the deputy gap as an improvement item

rather than letting it pass unnoticed.

MAIN GAP

The most significant gap is the closure report itself: declaring the DR situation closed at 18:20 is correct given that all three conditions (backups verified, Vault rebuilt, EDR coverage confirmed) are met, but the report lists failback to the original production environment as an open item without a stated owner, timeline, or criteria for confirming the recovery VPC is decommissioned once failback completes. Until failback is done, the organisation is running in a non-redundant, non-standard topology - that risk needs a named owner and a checkpoint date, not just "next week."

GOING FORWARD

Going forward, the formal DR situation is closed and the organisation returns to monitored normal operations from the recovery VPC. The lessons-learned review within two weeks should prioritise DR-114 (the Checkout RTO gap that existed before this incident), the absent deputies for the Lead DBA and CISO, the absence of an on-call rota and out-of-hours authority, and the failback plan - all of which were live pressure points during this exercise.
