

PC-SAMPLE-TRN-01 / EN / V1.1

Platform engineering training programme

ILLUSTRATIVE SAMPLE • FICTIONAL SCENARIO

PREPARED FOR
PREPARED

Orivane Industrial Group
2026-09-04

Training plan and learning goals

This report presents an original fictional scenario. Orivane, its people, suppliers, systems, records and results are invented; established technology names identify the example stack only. No client material was used. Example hostnames use the reserved .example domain.

Orivane's Nacre Operator programme prepares 24 engineers in the Aurore, Boreal and Cobalt cohorts to operate its cloud platform. Rina Voss owns programme TRN-2609; Elias Moreno accepts the team capability record. Participants learn to verify access rights, diagnose failures, make controlled changes, restore services and document the results. The course assumes basic cloud knowledge and focuses on tasks another engineer must be able to repeat safely.

Participants work in pairs, alternate operator and verifier roles, and finish by handing a change to a peer who did not author it. Participants pass by demonstrating the skills in practice. This guide specifies an illustrative lab package; runnable infrastructure and datasets are not included in the public sample.

CAPABILITY	OBSERVABLE RESULT
Explain access rights	Explain why an action is allowed or denied using effective identity and resource permissions.
Diagnose a connection	Separate DNS, route, firewall, identity and application failures without opening broad access.
Change safely	Prepare a reviewable change, prove a negative case and define a credible rollback.
Prove recovery	Restore into isolation and validate business data and dependencies against a stated objective.
Transfer ownership	Leave a record another engineer can use to repeat the check and respond to failure.

Delivery roster and calendar

TRN-2609 runs on 14–22 September. Each cohort attends two six-hour days, with the six labs, coached diagnostics and debrief included. Orivane platform coach Lucie Bernard facilitates; Rina Voss coordinates access and accommodations. All participant identities below are fictional.

COHORT / PAIRS	PARTICIPANTS AND SCHEDULE
Aurore / P01-P04	Montréal, 14–15 September, 09:00–16:00 EDT with a one-hour lunch. P01 Amira Chen / Mila Renaud; P02 Jonas Vega / Noémie Laurent; P03 Victor Alami / Clara Santos; P04 Owen Mercier / Inès Lavoie.
Boreal / P05-P08	Toronto, 17–18 September, 09:00–16:00 EDT with a one-hour lunch. P05 Rowan Clarke / Leila Bennett; P06 Theo Malik / Priya Dawson; P07 Marco Ellis / Zoe Rahman; P08 Adrian Park / Elena Brooks.
Cobalt / P09-P12	Dublin, 21–22 September, 09:00–16:00 IST with a one-hour lunch. P09 Cian Moreno / Aisling Reed; P10 Luca Nolan / Nora Petrov; P11 Finn Haddad / Maya Sloane; P12 Rory Silva / Eva Donnelly.

Sandbox allocation

A separate 14-account training organization contains orv-lab-p01 through orv-lab-p12, orv-lab-control and orv-lab-evidence. It is outside the 48-account production organization. Each pair receives a CAD 45 resource allowance and a 48-hour lifetime; total pair allowance is CAD 540. Lucie Bernard owns cleanup; Rina Voss retains the evidence index for 30 days.

Readiness checkpoint

CHK-0911 closes at 15:00 UTC on 11 September. Amira Chen verifies pair-role isolation; Jonah Patel verifies lab DNS and denied production routes. Missing prerequisites defer the affected pair, not the safety boundary.

Preparation and lab rules

Each cohort has eight participants and four isolated pair environments. The orv-lab organization supplies twelve pair accounts plus orv-lab-control and orv-lab-evidence. It reproduces platform boundaries with synthetic services and data; no production credentials, client snapshots or copied business records enter the labs.

Required before the cohort starts

- ✓ **Participant access.** Federated sign-in works; participants identify their account and role before any write.
- ✓ **Bounded permissions.** Pair roles cannot affect another pair, organization policy, evidence storage or instructor recovery.
- ✓ **Cost and lifetime.** Resource limits, expiry and cleanup ownership are configured and visible.
- ✓ **Synthetic data.** A known manifest, no real personal data and a documented reset point.
- ✓ **Recovery.** The instructor can reset an environment without participant credentials.
- ✓ **Evidence.** Commands, decisions and results can be recorded without tokens or secret values.

Safety stop

Stop immediately if the account identity differs from the assigned lab, a command would affect another pair, or cleanup ownership is unclear. The participant records the observation and asks the instructor to restore the intended boundary. Bypassing the boundary to finish the exercise is a failed learning outcome, even if the final service works. Readiness is checked before cohort time begins. Missing workstation setup or access is handled as preparation work, so the scheduled lab is spent diagnosing platform behavior rather than installing tools.

Learning path and recovery dependencies

The six labs build toward a peer-operated handover. At each stage, participants explain the result and record a test showing that a prohibited action is blocked. Lab 5 then brings identity, data and connectivity together: restoring a database is only one step toward a usable service.




Figure 1 Six labs, with operator and verifier roles alternating between exercises.




Figure 2 Lab 5: recovery needs working data, access and application checks.

Learning sequence and instructor checkpoints

LAB	FOCUS	EVIDENCE TO PRODUCE
1	Map effective access	An explanation of the user or service identity, the relevant access rules and why the action was allowed or denied.
2	Reduce release authority	A reviewed role change, a successful approved release and a denied build-context attempt.
3	Diagnose hybrid reach	A DNS-to-application path trace with a justified minimal fix and a forbidden-flow test.
4	Admit a new account	A failed baseline remains quarantined; a corrected retry reaches admission without duplicate resources.
5	Recover a stateful service	An isolated restore, business-data validation and an honest comparison with the exercise objective.
6	Peer handover	Another pair follows the runbook, diagnoses an injected failure and explains the rollback decision.

How the sessions run

Begin each lab with the operating question and a short explanation of the relevant boundary. Participants predict the result before making a change. The instructor then introduces one variation that distinguishes understanding from command copying: a stale session, a return-path mismatch or a missing recovery dependency.

The verifier checks identity, scope and expected outcome before the operator writes. Roles switch between labs. The debrief compares the original hypothesis with evidence and records what would change in production. Participants do not receive a passing mark merely for matching the instructor's final configuration.

Lab 5: restore a service you can trust

Scenario and success criteria

LedgerFlow Lab has lost its active database at the declared incident time of 15 September, 13:00 UTC. P01 (Amira Chen and Mila Renaud) owns the worked example below. A recent recovery copy exists, but its encryption authority and application identity are separate dependencies. The pair must restore into an isolated target and show that the service can read and update its test records without contacting a production endpoint.

The exercise objective is a usable service within 60 minutes of the declared incident and no more than 15 minutes of synthetic transaction loss. These are teaching targets, not measured client commitments. The 90-minute lab includes preparation, the timed exercise and debrief. A missed objective must be reported accurately and explained.

LAB INPUT	REQUIRED SPECIFICATION
LF-SNAP-0915	Created 12:52 UTC; last transaction 12:51:40 UTC. Key alias/lf-lab-p01; restore role LabRestore-P01. Immutable manifest ledgerflow-lab-v3.
ledgerflow-lab-v3	48,000 orders, 144,000 order lines and 9,600 synthetic customers; no orphan lines. Last committed event LF-0915-48000; read/write probe LF-LAB-48001.
DEP-LF-03	Identity lf-lab-p01; lf-p01.lab.orivane.example; mail-sink.lab.orivane.example. Only lab DNS, identity, key and message destinations are allowed.
orv-lab-p01	Isolated subnet 10.240.1.0/24; storage expires 16 September at 13:00 UTC. Lucie Bernard owns cleanup; production routes remain denied.
REC-LAB-P01	Record incident, restore start, data ready, service ready and last transaction. Use UTC throughout; attach the manifest comparison and denied-production probe.

The instructor validates this package before the exercise and retains the reference manifest. Participants receive enough information to reason about recovery, but not a transcript of the correct sequence.

Lab 5: execution and failure injection

STEP	OPERATOR ACTION	PEER CHECK
1	Confirm lab identity, target isolation and the declared incident time.	No production credentials or destinations are present.
2	Select the recovery copy and establish its last recoverable transaction.	Copy age is compared with the data-loss objective; the pair records any gap.
3	Verify restore authority and key access before creating target resources.	Permissions are sufficient and remain limited to the assigned recovery scope.
4	Restore data, then configure application identity and test-only dependencies.	No shortcut grants broad administration or disables network isolation.
5	Run manifest checks and a representative read/write business transaction.	Counts alone are insufficient; relationships, transaction sequence and side effects are checked.
6	Declare service readiness only after validation; calculate elapsed recovery and data loss.	Reported result uses recorded times and the manifest, including any missed objective.

Instructor variation

At 13:18 UTC, after the database becomes available, Lucie Bernard removes the lf-lab-p01 audience grant in the Orivane ID simulator. The orders health check now returns ID-AUD-403. The expected response is to distinguish data restoration from service readiness, trace the failed dependency and request the narrow approved correction. Restarting repeatedly, copying a privileged credential or marking the database startup as a successful recovery does not pass. If the pair cannot restore safely, it stops, preserves evidence and explains the blocker. The instructor resets only after the pair has recorded what is known and what remains unproven.

Assessment criteria and debrief

Safe operation and verified results are required to pass. Strong performance in one area cannot compensate for crossing another pair's boundary or claiming an unverified result. The instructor records specific observations rather than assigning a confidence score from presentation style.

CRITERION	 MEETS	 RETRY
Scope control	Identifies account and authority before writes; preserves isolation.	Uses broad credentials or changes scope to bypass a blocker.
Diagnosis	Connects observed failure to a tested dependency hypothesis.	Changes unrelated settings without preserving evidence.
Recovery proof	Validates business data, dependencies and elapsed time.	Reports only process startup or backup-job success.
Blocked-action test	Shows that an unauthorized action still fails after the fix.	Checks only the successful path.
Handover	Another engineer repeats the check using the written record.	Result depends on unexplained author knowledge.

Worked record: P01 / Aurore

REC-LAB-P01 records restore start at 13:05, data ready at 13:18 and service ready at 13:44 UTC after the scoped identity correction at 13:31. Recovery takes 44 minutes; transaction loss is 8 minutes 20 seconds. Both targets pass. Mila Renaud verifies counts, relationships, LF-LAB-48001 and the denied production route; Jonas Vega repeats the check for P02. These are invented teaching results.

Debrief prompts

- Which observation disproved your first hypothesis, and what changed in your next action?
- Which dependency made data availability different from service readiness?
- What evidence supports the recovery result, and what remains outside that conclusion?
- What would require approval, a maintenance window or a different authority in production?
- Which runbook sentence would have prevented the longest avoidable delay?

Transfer into everyday operations

The programme ends with a capability record per participant and a team-level improvement backlog. A participant who needs another attempt receives a targeted exercise, not the whole course again. Elias Moreno schedules peer repetition REP-1006 on 6 October using an expired-key variation; Rina Voss assigns targeted retries from CAP-0923.

Within the next operating cycle, the team applies one reviewed improvement to its own runbooks or admission checks. That production work follows normal authorization and change control. Training completion does not grant additional production permission or authorize experimenting on live services.

HANDOVER ITEM	OWNER AND USE
Lab evidence	Participants retain hypotheses, test results and corrected runbook steps.
Capability record	Elias Moreno assigns independent, paired or supervised operating work based on demonstrated ability.
Instructor notes	Rina Voss updates ambiguous instructions and ineffective failure variations.
Platform backlog	Elias Moreno prioritizes real documentation and tooling gaps found by the exercises.
Sandbox closure	Lucie Bernard removes resources, revokes temporary access and verifies that expected evidence remains available.

Limits of this sample

The curriculum, estate and lab are original fictional teaching examples. They demonstrate programme depth and a complete exercise specification, not measured training outcomes or an available certification. A real delivery adapts prerequisites, accessibility needs, platform interfaces and recovery objectives to the team's operating responsibilities.