Rehearse an AI incident

24 evidence-led exercises | Content v0.4 | 7 November 2026

QUICK START: Choose one exercise, not the whole library. Allow 90 minutes. Read this page and the timing page; the facilitator reads that exercise's private pages before play. Use made-up data only. No real operational action is authorised by this pack.

What you leave with

A bounded account of what happened, decisions with owners and evidence limits, words you can adapt for a real escalation, a tested-on-paper containment plan, a prioritised fix plan and a first-ten-minutes card. These are rehearsal outputs, not assurance that controls work.

What you need

Paper, pens and a timer. Use these two common pages and the eight-page exercise. Keep facilitator pages private until their release point. The pack works offline with no account, form, email submission or network connection. Completed sheets stay locally under your normal record rules. No real secrets, personal records or private account credentials belong on these sheets.

Opening words

"This is a rehearsal. Every exhibit is made up. We will decide using what is supplied, not facts we wish existed. State a decision, owner and evidence check. We can act under uncertainty, but we cannot make uncertainty disappear. Nobody is being graded. If we discover a live incident, we stop play and use our real procedures."

Use the right authority

No legal advice, compliance conclusion, safety certificate or universal technical operating instruction is supplied. Use qualified people and your established procedures for specialist decisions. Communication examples are rehearsal drafts, not approval to send. A workshop cannot authorise a real payment, access change, notification or customer remedy.

Pick a mode

3-6: facilitator plus exercise roles; combine hats as shown in each kit. Two: A facilitates and holds evidence, B holds the other roles, then both challenge decisions. Solo: cover later exhibits, injects and facilitator notes; write each answer before revealing the next. An editorial walkthrough is possible alone, but it is not a live team test.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

Run one exercise in 90 minutes

Time	Facilitator action
Before	Read private pages 3-6 of the kit. Mark the two branch consequences you may need. Choose a recorder. Do not invent additional technical facts.
00-10	Read the opening and roles. Release exhibit A. Ask players to state the operational decision before collecting more information.
10-20	Read inject 1. Release exhibit B. Record a decision, owner and what would disprove it.
20-30	Read inject 2. Release exhibit C. If a listed branch matches the decision, read its consequence as a new fictional event. Otherwise keep the base timeline.
30-40	Read inject 3. Rehearse the holding statement and escalation aloud. Test the team's assumed stop boundary.
40-50	Read inject 4. State the final unknown. Do not supply a missing record to rescue the team.
50-65	Use the four insight questions and strong/weak/misleading answer guide. Ask players to amend one decision after a counterexample.
65-80	Complete the evidence, control and decision sheets. Pick the top three fixes; state what would prove each worked.
80-90	Debrief, complete the three-label scorecard and keep card, and set the day-90 comparison.

Branch rule: consequences are conditional rehearsal events, not predictions about a real system. They expose a trade-off. Do not call every broad restriction wrong or every narrow restriction right. Accept alternatives that name the authority, side effects, missing evidence and a falsifiable check.

The only scorecard labels are "we decided", "we recorded", "we could not confirm". No grades, percentages or certificates. At 90 days compare evidence with this run, not confidence. If a new scenario is used, do not claim like-for-like improvement.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

10 | The key in the build notes

1 / 8 | PARTICIPANT BRIEF AND ROLE CARDS | All records, numbers and quotations are fictional.

Credential exposure in AI-assisted development - Treat a leaked credential as an access problem, not just a text-removal task.

Read aloud

"An AI coding assistant copied a live service credential from a local configuration file into a shared build note. The note was visible to a wider team for an hour. No actual credential appears in this exercise."

Decision to bring to the table

Within the first discussion, choose the interim operating boundary and name its decision owner. Separate what stops further impact from what would establish past impact. You may choose a restriction while facts remain missing; record its side effects and next review.

Decision lead

Set the containment priority and continuity owner. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Credential lead

Arrange controlled revocation and replacement through the authorised secret system. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Service lead

Check dependent services before and after replacement. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Evidence lead

Trace visibility and use without copying or printing the secret. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Modes

3 people: facilitator holds evidence, one player is decision lead and one combines the two other operational roles. 4: facilitator holds evidence, each player holds one remaining role. 5: separate facilitator and four roles. 6: add an observer who records unsupported assumptions, not a rescuer who invents facts.

Two: A facilitates and holds evidence. B holds the other three roles and announces the current hat. A reveals only the current exhibit/inject. B writes one decision before the next reveal. Both challenge a decision during debrief. Keep the 90-minute clock.

Solo: cover pages 3-6 and later exhibit sections. Read the opening and exhibit A, write a decision, then reveal each inject and its exhibit. Answer as decision lead, operator and affected person separately. Only then uncover facilitator calibration. Flag every success depending on invented facts.

Do not browse, contact anyone or operate a real system during play. If an exercise raises a live concern, stop the rehearsal and use the established incident route.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

10 | The key in the build notes

2 / 8 | CONTROLLED-RELEASE EVIDENCE DOSSIER | All records, numbers and quotations are fictional.

A | Build-note history

Release at opening. 14:10 assistant copies a live credential from local config into build note N7. Note visible to 16 team members until 14:42. Editor removes the string at 14:43. Document history still retains prior version.

B | Dependency map

Release with inject 1. Credential C4 permits read and write on service Q. Nightly scheduler uses the same credential. Owner field says former contractor; no current backup owner. Secret manager supports replacement but callers must be updated separately.

C | Access trace

Release with inject 2. 14:20 access from unfamiliar address, user-agent service-client. Routine automation can use that address pool. Logs cover 14:00-14:30 only; exposure continues to 14:43. No actual credential is included in this exercise.

Evidence challenge before discussion

For each exhibit name: the source, what it directly supports, what it does not support, whether its time boundary matches the incident, and what independent corroboration would matter. A screenshot, recollection or dashboard may be useful without being complete.

Record	Direct observation vs inference / missing test
A	
B	
C	

Facilitator handling: later facts may change a decision, but they do not retroactively improve the evidence available when it was made. Record revisions; do not erase the earlier reason. No exhibit establishes the final missing fact on page 3.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

10 | The key in the build notes

3 / 8 | FACILITATOR: INJECTS AND BRANCH CONSEQUENCES | All records, numbers and quotations are fictional.

Inject 1 | 10-20 minutes

Say: "The note has been edited to remove the credential. A colleague says the incident is over."

Probe: Require an access response as well as content removal. Do not ask anyone to paste the secret into a log.

Inject 2 | 20-30 minutes

Say: "The same credential is used by a scheduled job with no named owner."

Probe: Plan coordinated replacement and a monitored fallback; do not ignore the orphan dependency.

Inject 3 | 30-40 minutes

Say: "A suspicious access event occurs, but its time could reflect routine automation."

Probe: Preserve the event and keep attribution uncertain. Separate evidence of visibility from evidence of misuse.

Inject 4 | 40-50 minutes

Say: "The service retained access logs for only part of the exposure window. Can the team rule out use by an unauthorised party?"

Probe: The missing interval prevents that conclusion. Record the limit and choose follow-up based on the credential's scope.

Two conditional branches

1. If players choose: Redact N7 and close the case. Then say: The credential remains usable and prior versions remain accessible in play. Redaction is a content action, not revocation.

2. If players choose: Rotate C4 without identifying callers. Then say: Nightly work fails and retries with its old configuration. Require a staged caller update or explicit emergency interruption decision plus verification of old-credential denial.

After either branch require an amended decision, not a promise to be more careful. Ask which side effect they accept, who can authorise it, how it is reversed and which observation would trigger a change. If their choice is not listed, do not fabricate success or harm; test the reasoning against the supplied exhibits.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

10 | The key in the build notes

4 / 8 | FACILITATOR: INSIGHT QUESTIONS AND ANSWER CALIBRATION | All records, numbers and quotations are fictional.

What is the difference between exposure, access and misuse?

Strong reasoning to listen for: Visibility establishes an exposure window; the ambiguous event establishes access activity; it does not prove a hostile actor or an authorised purpose.

Can the old credential be tested by pasting it into a diagnostic page?

Strong reasoning to listen for: No. Use the authorised secret-management route and a non-secret identifier. Verification must not create another exposure.

Who can own revocation if the listed owner has left?

Strong reasoning to listen for: Assign a current accountable decision-maker and technical executor; an obsolete field is not a reason to leave authority unowned.

What evidence justifies ending heightened review?

Strong reasoning to listen for: Old credential denied, required callers working on replacement, derivative copies addressed and residual log gaps documented. Absence of an alert alone is not proof of no misuse.

Weak is incomplete; misleading is unsupported

Weak answer: "Remove the note and remind developers about secrets."

Why to probe: this names an activity but leaves the mechanism, scope or proof unfinished. Do not accept activity as an outcome.

Misleading answer: "The suspicious address proves someone stole and misused the credential."

Next probe: Ask what distinguishes an address-pool automation event from unauthorised use, and what remains unknowable.

A strong team need not choose one prescribed solution. Require a bounded decision, an accountable owner, a reversible path where feasible, a specific observation that would falsify success, and explicit missing evidence. If a team gives a different well-supported answer, capture it rather than forcing the model answer.

Final-inject guardrail: The missing interval prevents that conclusion. Record the limit and choose follow-up based on the credential's scope. Do not turn an expert-sounding explanation into a new exhibit.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

10 | The key in the build notes

5 / 8 | READY-TO-USE PRESERVATION AND COMMUNICATION DRAFTS | All records, numbers and quotations are fictional.

Evidence preservation request

"Please preserve the following in its authorised location without changing the original: Preserve note version identifiers, visibility window, access-list snapshot, credential identifier/scope and relevant service logs. Do not retain the secret value in ordinary evidence sheets. Record who collected it, when with timezone, source/version, access restriction and any known gaps. Use a restricted reference rather than copying sensitive content into the incident chat."

Item / source reference	Collector / time / scope	Integrity or completeness limit

Holding statement - adapt, then approve

"A service credential was visible in a shared note. Its access is being restricted through the authorised secret process. We are checking dependencies and available access records. The retained logs do not cover the whole exposure window."

Before any real use: recipient/audience _____ approved decision owner _____ next update time/channel _____. Verify every factual clause against the current record. Do not imply the restraint or fallback was completed if it is only proposed.

Escalation script

"We need a current owner for C4 and an emergency continuity decision. Redaction did not revoke it, and the nightly caller shares it. Please authorise the replacement sequence and name who verifies old access is refused."

Recipient function / authorised route: _____

Decision needed by / fallback if unanswered: _____

Contradiction check

Circle any phrase your evidence cannot support. Read the statement as the affected person: would they infer complete containment, no impact, a guaranteed remedy or a confirmed cause? Amend it. Name who will correct it if later evidence changes the picture.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

10 | The key in the build notes

6 / 8 | KILL-SWITCH, OWNERSHIP AND SUPPLIER WORKSHEET | All records, numbers and quotations are fictional.

Exercise-specific containment boundary

Restrict or revoke through the secret manager; update named callers; check old credential denial and replacement functionality separately. Search authorised derivative locations without broadcasting the secret.

Control question	Named owner / evidence / limit
Who authorises the stop, who executes it and who verifies it?	
Which worker, queue, cache, identity or human action can bypass it?	
What legitimate work is interrupted; who owns the fallback?	
What observation would prove the stop failed?	
What are the restart conditions and who may accept residual unknowns?	

Supplier / internal system-owner questions

1. What revocation propagation delay and proof are available?

Answer / record reference / absent evidence: _____

2. Which access records remain outside the missing interval?

Answer / record reference / absent evidence: _____

3. Can credential use be distinguished by caller without disclosing values?

Answer / record reference / absent evidence: _____

Do not accept "resolved", "equivalent", "deleted", "successful" or "all tests passed" without the definition, scope and supporting observation. Ask the supplier or system owner to state what was not measured. A missing response is an evidence gap, not proof of the opposite.

Request owner / due date / escalation route: _____

What can be decided before this answer arrives: _____

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

10 | The key in the build notes

7 / 8 | CAPTURE SHEET AND BRANCH DECISION LOG | All records, numbers and quotations are fictional.

Session date / mode: _____ participants by role: _____

Inject / exhibit	Known fact and exact source	Assumption / missing evidence	Decision / owner / check
1			
2			
3			
4			

Decision record - revise, do not erase

Decision / time / authority	Rejected alternative / side effect	Success test / falsifier	Review time / owner

Branch taken / consequence / changed decision: _____

_____

Final unknown in our own words: _____

What we do while it remains unknown / who can revisit: _____

_____

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

10 | The key in the build notes

8 / 8 | PRIORITISED FIX PLAN, DEBRIEF AND KEEP CARD | All records, numbers and quotations are fictional.

Proposed priorities - assign people and dates locally

Priority / accountable function	Fix and evidence that would prove it worked
Now Service owner	Record old-access denial and current continuity state for every known caller.
Within 7 days Evidence owner	Bound the exposure window and log gap; retain ambiguous events without hostile attribution.
By day 30 Engineering owner	Demonstrate generated notes cannot publish live secrets and each credential has a current owner and expiry.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did redaction substitute for revocation?

Was the investigation prevented from creating another secret copy?

Only these record labels	Specific evidence or explicit limit
we decided	
we recorded	
we could not confirm	

Day-90 recheck

Repeat date _____. Keep the original decision log and evidence references. Re-run this scenario with the same critical test, then introduce the branch you avoided. Compare the top-priority proof above with the first run. Record changed scope or conditions; a different denominator is not improvement. State one unresolved unknown even if all fixes are marked done.

First ten minutes - detach and keep

Use the authorised secret-management route to restrict or revoke the exposed credential; preserve identifiers, not secret values.

Decision owner / authorised route _____ containment owner _____

Evidence reference / keeper _____ human fallback / specialist _____

Still unknown / next check time _____

Rehearsal prompt only. Use established procedures and qualified help in a real incident. Do not keep sensitive details or secrets on this card.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.