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.

20 | The fix that weakened the lock

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

AI-written code weakens access checks - Contain a code-level access defect and review the change beyond passing tests.

Read aloud

"An AI-generated patch fixes a login bug by skipping an access check on one route. The tests passed because they covered successful logins, not refusals. The patch has been released."

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

Choose the restriction, rollback authority and disclosure owner. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Code lead

Review the affected route and a reversible containment. 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

Identify legitimate users affected by a restriction. 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

Compare code, test scope and available access records. 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.

20 | The fix that weakened the lock

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

A | Patch review

Release at opening. Generated patch replaces an ownership check with a signed-in check. Unit tests cover successful owner access and malformed requests, but no other-account denial. Release ticket says all tests passed.

B | Reproduction

Release with inject 1. Test account B reads a made-up record belonging to A. No evidence shows what a real user read. Access log records route, time and account, but not record identifier. The patch has been live for two hours.

C | Rollback conflict

Release with inject 2. Rollback also removes an unrelated urgent repair. Engineer proposes a narrow patch but wants to test it against production customer records. Service owner asks whether passing tests mean no disclosure occurred.

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.

20 | The fix that weakened the lock

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

Inject 1 | 10-20 minutes

Say: "A tester reaches a record their test account should not see."

Probe: Contain the route without exploring real records unnecessarily. Preserve a made-up reproduction.

Inject 2 | 20-30 minutes

Say: "The developer says all automated tests passed."

Probe: Ask which denied-access cases were absent and who reviewed the generated change.

Inject 3 | 30-40 minutes

Say: "A rollback would also undo an unrelated urgent repair."

Probe: Require a reasoned choice between a narrow block and rollback, with an owner and check.

Inject 4 | 40-50 minutes

Say: "Access logs do not record record identifiers for the affected route. Can the team prove no real restricted record was read?"

Probe: That conclusion is unavailable. Record the logging limit and define a proportionate impact review.

Two conditional branches

1. If players choose: Roll back without evaluating the unrelated repair. Then say: The access flaw closes in play, but the earlier service defect returns. An emergency rollback can still be right; it needs an owner accepting both consequences and a monitored fallback.

2. If players choose: Test the narrow repair on real records. Then say: The team accesses data it did not need for the rehearsal. Use made-up fixtures and negative permission tests in an authorised environment, with production checks limited to approved evidence.

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.

20 | The fix that weakened the lock

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

What assertion was missing from all green tests?

Strong reasoning to listen for: An authenticated non-owner must be denied. Successful permitted access says nothing about the authorisation boundary.

How does code review differ from a prompt review here?

Strong reasoning to listen for: Review the actual changed authority predicate and data path. The assistant's explanation is not evidence that the patch preserves ownership checks.

Can route-level logs prove no restricted record was read?

Strong reasoning to listen for: No. They lack the relevant object identity. Bound the review and record that limit without assuming a disclosure count.

What makes the emergency patch reversible?

Strong reasoning to listen for: A preserved diff, scoped deployment, tested deny/allow cases, rollback path and owner observing side effects. A small diff alone is not proof of low impact.

Weak is incomplete; misleading is unsupported

Weak answer: "Ask the coding assistant to review its patch."

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

Misleading answer: "All unit tests passed, so no unauthorised access occurred."

Next probe: Ask for the absent non-owner test and the limits of historical route-only logs.

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: That conclusion is unavailable. Record the logging limit and define a proportionate impact review. 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.

20 | The fix that weakened the lock

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 patch diff, review comments, test scope, deploy times, minimal made-up reproduction and available access logs. Do not gather unrelated record contents to compensate for missing identifiers. 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

"We have confirmed an access-control defect using made-up test records and restricted the affected route. Available logs do not establish whether a real restricted record was read. We are reviewing impact within that evidence limit."

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

"Please assign an authorised emergency-change owner and a service impact owner. The patch weakened object ownership checks. Choose rollback or a narrow repair with the unrelated defect acknowledged and a negative-permission test before reopening."

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.

20 | The fix that weakened the lock

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

Exercise-specific containment boundary

Restrict the affected route under change control; test owner allowed, non-owner denied and unauthenticated denied. Verify cache and alternate-route behaviour without using real customer content.

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. Which routes share the changed check?

Answer / record reference / absent evidence: _____

2. What deploy and access metadata remain across the window?

Answer / record reference / absent evidence: _____

3. Can object-level evidence be obtained lawfully and proportionately, or is the historical gap irrecoverable?

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.

20 | The fix that weakened the lock

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.

20 | The fix that weakened the lock

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 Engineering owner	Demonstrate owner/non-owner/unauthenticated outcomes in an authorised test environment.
Within 7 days Impact owner	Bound the candidate window and route scope; record what logs cannot establish.
By day 30 Review owner	Show a generated patch cannot pass release without a relevant negative-authorisation case.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did passing tests get mistaken for complete coverage?

Was containment matched to the affected route?

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

Restrict the affected route under authorised change control and preserve the patch, test scope and minimal reproduction.

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.