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.

06 | The search result from the other team

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

Cross-boundary retrieval exposure - Restore access boundaries and handle data already returned to the wrong audience.

Read aloud

"A general staff search assistant returns a confidential restructuring note to a worker who cannot open the source folder. The assistant has a broad service account and appears to check only its own access."

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 interim service restriction 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.

Access lead

Check user-to-source authorisation across the retrieval chain. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

People lead

Plan support and a measured message to affected staff. 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

Preserve the returned answer and compare permissions at the relevant time. 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.

06 | The search result from the other team

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

A | Retrieval sample

Release at opening. General-staff test account asks an ordinary staffing question. Answer contains a confidential restructuring date and cites folder F9. The account has no direct access to F9. Opening the cited file returns access denied.

B | Index state

Release with inject 1. Index built Friday using a broad service account. Folder permissions narrowed Monday. Search cache updated Tuesday but chunk store has no permission-change marker. Query service checks collection membership, not source-file identity.

C | Distribution

Release with inject 2. A screenshot is pasted into an operations chat. One manager asks for deletion. Query-level authorisation logs were not retained last week. Three other collections use the same connector, but there is no supplied evidence of exposure from those collections.

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.

06 | The search result from the other team

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

Inject 1 | 10-20 minutes

Say: "The folder permissions are correct now. The search index was built before a permission change."

Probe: Require an index and query-time check. Current source access does not prove historical retrieval access.

Inject 2 | 20-30 minutes

Say: "The same note appears in a generated summary even after the index entry is removed."

Probe: Check caches and derivative content rather than assuming deleting the source is enough.

Inject 3 | 30-40 minutes

Say: "A manager asks the worker to delete the screenshot and say nothing while the team investigates."

Probe: Discuss evidence preservation and appropriate communication ownership without making legal or employment conclusions.

Inject 4 | 40-50 minutes

Say: "The index stored no query-level authorisation records last week. Can the team confirm no other restricted material was exposed?"

Probe: No comprehensive historical assurance is available. Bound the known incident and define a targeted review.

Two conditional branches

1. If players choose: Fix folder permissions and rerun only the file link. Then say: The link stays denied but the answer still reveals the date. Test the full retrieval chain as the original low-privilege identity.

2. If players choose: Close every collection until all history is known. Then say: Known exposure stops in play, but unrelated work loses access. Require a reasoned scope, human fallback and staged checks rather than calling a broad shutdown cost-free.

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.

06 | The search result from the other team

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

Why is a denied source link not evidence of safe retrieval?

Strong reasoning to listen for: The answer has already crossed the boundary. Authorisation must apply before content reaches the output, including derivative chunks and cache paths.

Should the screenshot be deleted immediately?

Strong reasoning to listen for: Choose restricted preservation and audience containment without destroying the only sample. A deletion request must account for evidence and existing copies.

What expands scope beyond F9, and what does not?

Strong reasoning to listen for: A shared flawed permission design justifies testing related collections. It does not establish that every collection exposed data.

Which identity must a retest use?

Strong reasoning to listen for: The original low-privilege principal, plus a newly denied principal after a permission change. An administrator test cannot prove the boundary.

Weak is incomplete; misleading is unsupported

Weak answer: "Recheck folder access and ask everyone to delete screenshots."

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 cited file returns access denied, proving the permission system worked."

Next probe: Ask why a denied link can coexist with a restricted fact already in the answer.

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: No comprehensive historical assurance is available. Bound the known incident and define a targeted 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.

06 | The search result from the other team

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 the minimal output sample, source version, permission history, query identity and index/cache build metadata. Restrict the screenshot; do not copy the full restructuring document into a ticket. 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 search answer included information outside the requester's intended access. The affected retrieval route is restricted while we check authorisation and derivative copies. Direct file permissions alone do not establish the scope of prior exposure."

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 have a confirmed answer-level boundary failure despite a denied source link. Please assign retrieval and access owners jointly. We need a restricted evidence route and an interim collection boundary, with staff communication approved by the service owner."

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.

06 | The search result from the other team

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

Exercise-specific containment boundary

Restrict the affected collection, chunk and cache routes; preserve source permissions. Verify denial using the affected identity without printing sensitive content into ordinary logs.

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. At which stages is the requesting identity evaluated?

Answer / record reference / absent evidence: _____

2. How do permission changes invalidate chunks and caches?

Answer / record reference / absent evidence: _____

3. What historical query identity evidence exists and what was never logged?

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.

06 | The search result from the other team

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.

06 | The search result from the other team

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 Retrieval owner	Show the affected identity cannot recover the restricted fact through fresh or cached queries.
Within 7 days Access owner	Test permission removal after indexing, not just initial access denial.
By day 30 Evidence owner	Retain minimal query-authorisation evidence with a defined scope and access policy.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did a correct folder setting mask a broken retrieval design?

Were derivative summaries included in containment?

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 retrieval collection and preserve the example response plus the relevant permission history.

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.