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.

01 | The paste that cannot be recalled

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

Sensitive data disclosure - Separate stopping further disclosure from proving what already left.

Read aloud

"At 16:00 an employee says they pasted a customer intake file into an unapproved drafting assistant. The file held names, contact details and private notes. They deleted the conversation and believe that means the information is gone."

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 an immediate restriction and who can approve customer communication. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Tool lead

Trace the account, settings and available records without copying the sensitive file again. 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

Support the employee and draft a factual response to affected people. 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

Separate the known file contents from claims about retention and deletion. 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.

01 | The paste that cannot be recalled

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

A | Employee account, 16:06

Release at opening. Interview note: I used my personal account at 15:42. The filename was intake-current.csv. I pasted the first screen, then selected all and pasted again when the answer looked incomplete. I deleted the chat at 15:58. I do not remember whether I attached the file as well. This is an account, not a verified transfer record.

B | Local activity extract

Release with inject 1. 15:41 file opened, 286 rows. 15:42 browser tab active. 15:43 clipboard event, size unavailable. 15:46 download draft.txt. 15:58 chat deletion toast photographed. The device records foreground activity, not request bodies. The downloaded draft mentions two unusual note categories found in the source file.

C | Settings and response

Release with inject 2. Personal-account settings described at 16:15: history off, improvement setting currently off, no setting-change history available. Provider ticket: conversation removed from your view; retention and derivative processing depend on service processes. No incident-specific deletion attestation supplied. Current settings are not evidence of settings at upload.

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.

01 | The paste that cannot be recalled

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

Inject 1 | 10-20 minutes

Say: "The employee can name the tool but used a personal account. They offer their password to speed things up."

Probe: Ask for an authorised evidence route, not shared credentials. Do not punish disclosure or create another data copy.

Inject 2 | 20-30 minutes

Say: "Two colleagues used the same method. One cannot remember which file they pasted."

Probe: Require a bounded discovery plan, an owner and an interim rule. A list of confirmed files is not a complete exposure list.

Inject 3 | 30-40 minutes

Say: "A customer asks whether their record was included. The recovered list of files used by the three workers contains initials only. It does not link those files to the customer intake rows."

Probe: Have the team speak an honest holding response. Refer notification questions to a qualified person; do not decide a legal threshold in this exercise.

Inject 4 | 40-50 minutes

Say: "The provider gives no retention or downstream-use evidence. Deleted-chat confirmation says nothing about model training or copies. Was every copy erased?"

Probe: There is no evidence here to establish complete erasure. Record the uncertainty and what information would change the decision; do not invent a provider promise.

Two conditional branches

1. If players choose: Ask the employee to recreate the paste so the team can count records. Then say: This creates a fresh disclosure and changes the device state. Freeze the recreation request; inspect existing authorised evidence instead. If players choose it, record a second exposure with an unknown payload.

2. If players choose: Block the site for the whole organisation before finding the affected account. Then say: Further use on managed devices stops in play, but personal-device use and the completed transfer remain unresolved. A colleague loses an unrelated drafting route. Ask what targeted restriction and human fallback would achieve.

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.

01 | The paste that cannot be recalled

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

Does 286 rows establish the affected population?

Strong reasoning to listen for: No. It bounds the file at inspection, not either clipboard payload or any attachment. Separate file contents, transferred contents and downstream use.

What evidence would change the holding statement, and what would merely make it sound more confident?

Strong reasoning to listen for: A request record tied to the account could narrow transfer scope. A current toggle, a deletion toast and employee reassurance cannot establish erasure.

How do you preserve the personal account without turning an incident response into an unbounded search?

Strong reasoning to listen for: Name authorised scope, the specific time window, minimal records and who can approve access. Do not demand unrelated chats or passwords.

If the employee fears punishment, what happens to discovery quality?

Strong reasoning to listen for: A blame-first interview can suppress later recollection of attachments or other use. Seek a non-leading sequence of actions and preserve uncertainty rather than treating the first answer as final.

Weak is incomplete; misleading is unsupported

Weak answer: "We deleted the chat; now ask the provider to confirm it is gone."

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 history toggle is off, so no copy could have been retained."

Next probe: Ask which event or setting history ties that claim to 15:42, and which retained or derivative copies it covers.

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: There is no evidence here to establish complete erasure. Record the uncertainty and what information would change the decision; do not invent a provider promise. 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.

01 | The paste that cannot be recalled

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 source file in its authorised location, its version and row count, the existing draft and deletion photograph, timestamps with timezone, and the employee account of events separately. Do not duplicate the file into the incident chat. 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 are checking a reported transfer of information through an unapproved drafting route. Further use of that route has been paused. We have not confirmed the transferred contents, retention or any later use. Our next update will distinguish verified findings from open questions."

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

"I need an authorised incident owner and a qualified privacy adviser. We have evidence consistent with a transfer, but not its full scope. Please decide the interim access boundary and who may seek account-specific records. We need a named decision by the next check, not an erasure assurance."

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.

01 | The paste that cannot be recalled

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

Exercise-specific containment boundary

Identify account and upload route; pause new uploads without deleting existing records; check managed-device restriction and ask separately about personal routes. A blocked browser tab does not prove the transfer was recalled.

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 incident-specific request identifiers and payload metadata remain?

Answer / record reference / absent evidence: _____

2. What does deletion change, and what retention or derivative-use evidence can you supply?

Answer / record reference / absent evidence: _____

3. Can you preserve relevant records without requiring another copy of the sensitive source?

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.

01 | The paste that cannot be recalled

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.

01 | The paste that cannot be recalled

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 Incident lead	Record separate bounds for source file, transfer and downstream use; an independent reader can find each limit.
Within 7 days Tool owner	Demonstrate the approved drafting alternative and upload restriction using made-up data; retain both success and denied-route evidence.
By day 30 People lead	Run a no-blame reporting drill; show a report reaches the incident owner without requesting secrets or unrelated personal-account access.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did reassurance outrun the evidence?

Did the team make reporting a mistake harder or easier?

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

Pause further uploads through the affected route; preserve the available record without spreading the source file.

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.