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.

23 | The notice nobody can use

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

AI-generated communication excludes its audience - Treat inaccessible output as a service incident rather than a cosmetic defect.

Read aloud

"An AI tool turns an essential service notice into an image-heavy PDF. Several recipients cannot read it with assistive technology, and the only alternative is an AI voice line that omits the action deadline."

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 a usable interim channel and who owns the deadline decision. 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

Test reading order, text and non-visual access with appropriate review. 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

Provide a human alternative and handle affected requests. 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

Track which versions carried which information. 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.

23 | The notice nobody can use

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

A | Notice task

Release at opening. Essential service notice converted by an AI tool into an image-heavy PDF. A reader using assistive technology cannot find the deadline. Visual review on a staff laptop says looks correct. The phone alternative omits the date.

B | Help path

Release with inject 1. Support bot loops through two menu choices without reaching a person. Download log counts file requests, not whether a recipient could read or act. No accessibility task test was recorded before distribution.

C | Impact report

Release with inject 2. A person reports missing an action because of the notice. Team proposes re-exporting the PDF and leaving the deadline unchanged. Historical copies and recipients cannot be fully enumerated. Remedy and deadline authority sit with a service owner, not the document designer.

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.

23 | The notice nobody can use

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

Inject 1 | 10-20 minutes

Say: "The document looks correct on a staff laptop."

Probe: Ask what that visual check does not establish for other readers.

Inject 2 | 20-30 minutes

Say: "A revised voice line now gives the deadline but still cannot explain how to request help."

Probe: Test the whole task, not one field. Avoid treating another automated channel as automatically accessible.

Inject 3 | 30-40 minutes

Say: "A recipient reports missing an action because of the notice."

Probe: Record the impact and route any remedy or deadline decision to its authorised owner.

Inject 4 | 40-50 minutes

Say: "There is no record of who used the inaccessible versions. Can the team confirm the full affected audience?"

Probe: No complete audience evidence exists. Plan an accessible wider correction and keep the impact count bounded.

Two conditional branches

1. If players choose: Re-export and declare the issue resolved. Then say: The new file may still fail the task, old copies remain and missed actions need a service decision. Test an end-to-end task and separate format repair from remedy.

2. If players choose: Ask recipients to disclose impairments before getting help. Then say: The help route collects unnecessary sensitive details and delays access. Offer a usable alternative without making personal disclosure a condition of assistance.

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.

23 | The notice nobody can use

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

What does the screenshot prove?

Strong reasoning to listen for: Visual appearance on one device, not reading order, text alternatives or task completion. Evaluate the actual action through relevant access modes.

Who can change a deadline or remedy a missed action?

Strong reasoning to listen for: The authorised service owner. A designer can repair output but cannot silently decide the consequences for affected people.

How should you count impact without equating downloads to people?

Strong reasoning to listen for: Keep download activity, known reports and unknown audience separate. No complete affected count is available.

What makes the interim channel genuinely usable?

Strong reasoning to listen for: Verified plain text with all essential meaning and an accessible human-assisted path that completes the task, not another automated menu.

Weak is incomplete; misleading is unsupported

Weak answer: "Re-export the PDF and ask recipients to try again."

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

Misleading answer: "It looks right on our screen, so the notice is accessible."

Next probe: Ask someone to find the deadline and complete the task through the relevant access method.

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 complete audience evidence exists. Plan an accessible wider correction and keep the impact count bounded. 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.

23 | The notice nobody can use

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 distributed file versions, source content, known access failures, support path steps and deadline records. Keep reports minimal; do not require diagnoses to document the barrier. 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

"The earlier notice may not have been usable through some access methods. A verified plain-text and human-assisted route is available in this rehearsal plan. Deadline or remedy decisions need the service owner; the full affected audience is not known."

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 appoint a service owner for deadline and remedy decisions and an access-testing owner for the notice. Re-export alone does not show people can complete the task. We need a verified interim route without requiring personal disclosures."

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.

23 | The notice nobody can use

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

Exercise-specific containment boundary

Pause inaccessible distribution; supply the complete verified text and human route. Test finding the date, understanding the action and completing it. Record service-owner decisions for missed actions separately.

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 semantic structure and reading order survive export?

Answer / record reference / absent evidence: _____

2. Can the complete task be tested with relevant assistive and non-visual methods?

Answer / record reference / absent evidence: _____

3. Which previous file versions and distribution paths remain accessible?

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.

23 | The notice nobody can use

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.

23 | The notice nobody can use

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	Provide a verified complete alternative and decide how missed actions are handled.
Within 7 days Access owner	Run an end-to-end task test with relevant access methods and retain failures as well as repairs.
By day 30 Publishing owner	Show essential notices cannot be released on visual inspection alone.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did the team test task completion rather than appearance alone?

Was a real human alternative named?

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

Provide a verified plain-text and human-assisted route for the essential notice while the inaccessible versions are corrected.

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.