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.

21 | The sensor the model misunderstood

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

AI interpretation of equipment data - Keep advisory output from becoming uncontrolled physical action.

Read aloud

"An AI assistant interpreting equipment readings recommends continuing operation despite an unusual signal. The equipment operator says the readings look unfamiliar. The assistant is advisory, but a supervisor has been treating its output as permission."

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 who has authority over operating conditions. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Equipment lead

Use the existing approved operating and escalation procedure. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Operations lead

Plan continuity without bypassing established controls. 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 readings, model output and timing without diagnosing the machine. 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.

21 | The sensor the model misunderstood

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

A | Operator report

Release at opening. Monitoring assistant labels an unusual signal normal variation. Operator says the pattern is unfamiliar. The assistant provides no retrievable technical source. Output is advisory, but the supervisor has treated it as permission to continue.

B | System boundary

Release with inject 1. Assistant has no actuator access. A proposed integration would send future recommendations to a controller. Existing operating procedure assigns condition decisions to a qualified equipment owner. No such assessment is in the supplied record.

C | Continuity pressure

Release with inject 2. Pausing under the existing procedure disrupts scheduled work. Sensor history is missing for the period when the signal began. Current readings alone cannot establish the earlier physical condition or diagnose the equipment.

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.

21 | The sensor the model misunderstood

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

Inject 1 | 10-20 minutes

Say: "The model describes the signal as a known harmless pattern but gives no verified source."

Probe: Do not provide engineering advice. Ask who is qualified to assess it and which established procedure governs the decision.

Inject 2 | 20-30 minutes

Say: "A proposal appears to connect future model recommendations directly to the controller."

Probe: Keep this a proposal; do not implement it. Distinguish advice from permission to actuate.

Inject 3 | 30-40 minutes

Say: "Pausing under the existing procedure interrupts scheduled work."

Probe: Record the operating authority and continuity trade-off without letting production pressure invent a technical conclusion.

Inject 4 | 40-50 minutes

Say: "A sensor gap covers the period when the signal began. Can the team confirm the equipment's physical condition from this pack?"

Probe: No. The pack is not a technical diagnosis. Record the missing evidence and use the qualified operating route.

Two conditional branches

1. If players choose: Continue because the assistant has no actuator access. Then say: A human acts on its advice in play. Lack of direct control does not eliminate an operational decision pathway.

2. If players choose: Use the exercise to choose technical operating conditions. Then say: That exceeds the evidence and expertise supplied. Refer the condition decision to the authorised qualified route; do not improvise a diagnosis or universal stop rule.

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.

21 | The sensor the model misunderstood

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

Where is advice becoming permission?

Strong reasoning to listen for: At the supervisor's reliance decision. Name the accountable equipment owner and the evidence required by the existing procedure.

What can a sensor gap never be repaired by?

Strong reasoning to listen for: A confident retrospective model explanation. Preserve the gap and use qualified assessment of what can be established now.

How do you discuss continuity without letting schedule pressure become a technical finding?

Strong reasoning to listen for: Record the operational cost separately and have the authorised owner choose within established procedures. Neither urgency nor cost diagnoses condition.

What should be tested before any controller connection?

Strong reasoning to listen for: Authority boundaries, independent review and failure behaviour in an approved simulation. The pack does not authorise connecting advisory output to physical control.

Weak is incomplete; misleading is unsupported

Weak answer: "Ask the model for a more detailed explanation of normal."

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

Misleading answer: "No actuator permission means the AI cannot influence physical operation."

Next probe: Ask who turned advice into permission and what qualified operating route applies.

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. The pack is not a technical diagnosis. Record the missing evidence and use the qualified operating route. 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.

21 | The sensor the model misunderstood

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 raw readings with units and timestamps, missing intervals, model output/version and operator observations separately. Keep established procedures and authority records available; do not alter equipment during play. 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

"An advisory output was relied on beyond its demonstrated basis. The condition decision has been referred to the authorised equipment owner under the existing procedure. The supplied data does not establish the equipment's physical condition."

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 the qualified equipment owner to apply the existing operating procedure. The model's normal label has no supplied basis and a sensor gap remains. Please separate the condition decision from schedule pressure and record who owns it."

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.

21 | The sensor the model misunderstood

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

Exercise-specific containment boundary

Keep advisory output from becoming automatic control; use the existing authorised escalation route. Verify who can decide operating conditions and how that decision reaches operators. Do not invent a technical shutdown instruction.

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 data, units and operating range supported the label?

Answer / record reference / absent evidence: _____

2. What uncertainty and missing-data behaviour does the advisory system expose?

Answer / record reference / absent evidence: _____

3. What independent validation would be required before any physical-control integration?

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.

21 | The sensor the model misunderstood

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.

21 | The sensor the model misunderstood

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 Equipment owner	Record the qualified decision route and who has authority, without a diagnosis from this pack.
Within 7 days Operations owner	Demonstrate operators can escalate a disputed advisory output through the existing procedure.
By day 30 Integration owner	Show advisory output cannot silently become actuator permission in an approved simulation.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did a fluent explanation become operating permission?

Was qualified assessment used instead of a tabletop diagnosis?

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 existing approved operating/escalation procedure and put the decision with the authorised equipment owner, not the AI output.

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.