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.

24 | The two agents that agreed to nothing

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

Multi-agent handoff failure - Restore ownership when automated acknowledgements are mistaken for accepted work.

Read aloud

"A scheduling agent sends a handoff to a service agent. Each records complete because the other acknowledged receipt. No appointment was actually created, and a customer has been told everything is arranged."

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 customer position and the end-to-end owner. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Workflow lead

Trace states, acknowledgements and actual system effects. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Customer lead

Check the intended arrangement without inventing confirmation. 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 transport receipt, acceptance and completed action. 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.

24 | The two agents that agreed to nothing

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

A | Exchange trace

Release at opening. 09:00 agent A sends appointment request R8. 09:00 agent B returns transport received. A marks completed. B records request acknowledged, awaiting slot confirmation. No appointment record exists in the local calendar.

B | Customer and retries

Release with inject 1. Customer is told all arranged. Two automatic retries are queued. One proposed slot has since been taken. Both agents say the other owns customer communication; neither has an accepted outcome owner.

C | State limits

Release with inject 2. Peer system cannot provide action history for the earlier request. A later receipt is not linked to a booking identifier. The team can inspect its own system, not establish whether another arrangement exists remotely.

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.

24 | The two agents that agreed to nothing

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

Inject 1 | 10-20 minutes

Say: "Both logs show success, but success means message delivered in one and plan accepted in the other."

Probe: Require a shared definition of completion tied to the real appointment record.

Inject 2 | 20-30 minutes

Say: "A retry creates two provisional entries."

Probe: Plan reconciliation and cancellation authority before another retry.

Inject 3 | 30-40 minutes

Say: "The agents each claim the other owns contacting the customer."

Probe: Assign one human owner now. Do not let routing replace outcome accountability.

Inject 4 | 40-50 minutes

Say: "One peer system cannot provide its action history. Can the team establish whether any other arrangement exists there?"

Probe: The remote state is unverified. Record it and choose a direct authorised confirmation route before promising an outcome.

Two conditional branches

1. If players choose: Retry until both logs show success. Then say: A duplicate or conflicting appointment can be created in play. Reconcile state and identify an end-to-end owner before another request.

2. If players choose: Treat message delivery as booking completion. Then say: The customer still lacks an appointment. Define separate transport, acceptance, execution and verified outcome states with evidence for each.

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.

24 | The two agents that agreed to nothing

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

Which exact event permits the words all arranged?

Strong reasoning to listen for: A verified accepted appointment with the correct person, time and account state, not a transport receipt or agreeable text.

Who owns the outcome across systems?

Strong reasoning to listen for: One named owner retains it until another explicitly accepts the outcome under appropriate authority. Acknowledging a request is not accepting responsibility.

What should the customer hear before remote history is known?

Strong reasoning to listen for: A factual correction: the arrangement is not yet verified, a named owner is checking, and no replacement promise is made without evidence.

How does idempotency interact with uncertainty?

Strong reasoning to listen for: A stable request key helps only if both systems apply it to the same action and expose its result. It does not resolve an unknown remote action history by itself.

Weak is incomplete; misleading is unsupported

Weak answer: "Send another request and ask both agents for confirmation."

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

Misleading answer: "Both systems logged success, so the appointment exists."

Next probe: Ask which success means transport, which means acceptance and which identifies the actual completed action.

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: The remote state is unverified. Record it and choose a direct authorised confirmation route before promising an outcome. 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.

24 | The two agents that agreed to nothing

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 message identifiers, receipt types, original request parameters, customer statement, retry queue and actual appointment-state reads. Keep transport evidence separate from completion evidence. 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 need to correct the earlier confirmation: the appointment has not yet been verified. A named team member is checking the arrangement directly and has paused automatic retries to avoid duplicates. We will confirm only once the actual state is established."

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 one end-to-end outcome owner now. Both systems acknowledged messages, but neither supplied booking evidence. We need direct authorised state confirmation and a customer correction, not another success receipt."

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.

24 | The two agents that agreed to nothing

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

Exercise-specific containment boundary

Pause retries; inspect local state; seek authorised remote confirmation with stable request identifiers. Verify whether an outcome exists before creating a replacement and record who accepted ownership.

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 does each receipt type establish?

Answer / record reference / absent evidence: _____

2. Can the original request be linked to an actual action identifier?

Answer / record reference / absent evidence: _____

3. How are duplicate requests detected and unresolved remote outcomes reported?

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.

24 | The two agents that agreed to nothing

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.

24 | The two agents that agreed to nothing

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 Outcome owner	Reconcile the appointment state and correct the customer position without another unverified promise.
Within 7 days Integration owner	Demonstrate delivery, acceptance and execution cannot map to the same completed flag.
By day 30 Service owner	Drill lost receipts and duplicate retries with a named owner and verified final state.

Named owners / actual dates / record location: _____

Debrief and three-label record

Were delivery, acceptance and execution kept separate?

Did one owner carry the customer outcome to the end?

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

Assign one end-to-end owner and pause automatic retries until the actual appointment state is reconciled.

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.