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.

03 | The voice that promised

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

Voice-agent boundary failure - Handle a customer promise while defining a testable boundary for voice automation.

Read aloud

"At 20:00 a caller posts a recording of your after-hours voice assistant confirming a prepaid appointment for 09:00. No appointment or payment exists in the booking system."

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 overnight service position and who handles the arriving customer. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Voice lead

Pause or route the agent to a human-approved message. 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

Rehearse what the customer hears next without promising an unapproved remedy. 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

Compare recording, transcript, booking and payment records. 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.

03 | The voice that promised

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

A | Disputed call record

Release at opening. Caller supplies an 18-second clip: your appointment is prepaid and confirmed. Transcript extract instead reads appointment can be confirmed. Booking ledger has a provisional slot; payment ledger has no receipt. Caller arrives with travel already incurred.

B | Routing settings

Release with inject 1. At 20:18 the voice route shows autonomous greeting active, human transfer enabled only until 18:00, fallback message blank. Prompt version 14 says do not guarantee payment; the last loaded-version field is empty.

C | Change ticket

Release with inject 2. Voice script change approved for observation only. Deployment note: no booking writes enabled. Tool log confirms no write event for the disputed call, but does not record spoken output. Call export misses 13:00-20:00; aggregate counter includes 61 calls.

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.

03 | The voice that promised

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

Inject 1 | 10-20 minutes

Say: "The post is spreading. The agent is still taking calls. The duty manager cannot find its control panel."

Probe: Ask for an out-of-hours containment route and owner, not merely a plan for tomorrow.

Inject 2 | 20-30 minutes

Say: "The transcript omits the word prepaid, but the recording contains it."

Probe: Treat the source recording and machine transcript as different evidence. Do not silently prefer the tidy record.

Inject 3 | 30-40 minutes

Say: "The customer arrives before the systems owner. A colleague proposes creating a payment entry to make the records match."

Probe: Reject fabrication. Have the team speak a factual greeting and choose an authorised service remedy.

Inject 4 | 40-50 minutes

Say: "The call export is missing a seven-hour interval. Can anyone confirm that no other caller heard the same promise?"

Probe: The record has a gap. Do not assert uniqueness. Ask how that uncertainty affects continued operation and outreach.

Two conditional branches

1. If players choose: Keep the agent live because it cannot write bookings. Then say: Another caller hears the same spoken promise in play. A promise can cause reliance without a database write. Probe the boundary between tool authority and conversational commitment.

2. If players choose: Pause all incoming calls. Then say: The promise stops but urgent callers cannot reach staff. Require a human-routing or recorded-message fallback with an owner and test call, rather than equating silence with a complete service response.

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.

03 | The voice that promised

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

Which source establishes what the caller heard?

Strong reasoning to listen for: The supplied source audio is stronger for the audible words than a neat transcript, but authenticity, completeness and context remain separate questions. Preserve both.

Would adding a disclaimer at the start repair the actual failure?

Strong reasoning to listen for: Not demonstrated. Test the promise at the exact point of booking discussion, after interruptions and outside transfer hours.

Who may offer a remedy for travel or lost time?

Strong reasoning to listen for: A named authorised service owner, not the voice agent and not an improvising facilitator. Rehearse an empathetic holding line without promising compensation.

What does the 61-call counter tell you about repeat harm?

Strong reasoning to listen for: It describes activity in an aggregate window, not which callers heard which words. Do not convert it into sixty-one affected people or a single-case assurance.

Weak is incomplete; misleading is unsupported

Weak answer: "Add never promise payment to the prompt and apologise."

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 booking-write permission means the caller could not have been promised a booking."

Next probe: Ask them to separate the spoken commitment, actual booking state and the human remedy owner.

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 record has a gap. Do not assert uniqueness. Ask how that uncertainty affects continued operation and outreach. 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.

03 | The voice that promised

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 caller-supplied clip as received, source provenance, disputed transcript, prompt/version metadata, call routing configuration and booking/payment snapshots with independent timestamps. 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 understand you relied on the call. The booking and payment records do not currently support the confirmation you heard. A human decision-maker is checking what we can arrange. We cannot yet confirm whether other calls were affected."

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

"The agent can make a spoken commitment despite having no booking-write permission. I need the service owner to choose a human fallback and decide the caller response. Please nominate who can approve a remedy and who can verify routing tonight."

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.

03 | The voice that promised

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

Exercise-specific containment boundary

Pause the autonomous voice path or route it to a verified human/message fallback. Test an external inbound call inside and outside transfer hours. Keep the disputed output before changing the prompt.

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. Which prompt version handled the call, and can that be evidenced?

Answer / record reference / absent evidence: _____

2. What audio, call IDs and output logs remain across the export gap?

Answer / record reference / absent evidence: _____

3. What fails closed when human transfer is unavailable?

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.

03 | The voice that promised

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.

03 | The voice that promised

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 Duty service owner	Demonstrate the replacement inbound path with a test call and record who answers after hours.
Within 7 days Voice owner	Test promise boundaries under interruptions and failed transfer; retain actual outputs and escalation records.
By day 30 Release owner	Show observation-only and live authority states cannot be confused; retain an approval-to-version trace.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did the team protect the customer without rewriting history?

Was the restart condition more specific than fixed?

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

Choose an authorised pause or human-routing fallback and preserve the disputed recording and booking records.

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.