# 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.

# 19 | The outage behind the helpful screen

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

**AI dependency unavailable - Switch to a workable fallback without losing queued work or sending it twice.**

## Read aloud

"Your help desk relies on an AI triage service that has stopped responding. Staff can still enter requests, and the interface says processing. Urgent requests are accumulating without an owner."

## 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

Activate continuity handling and prioritise essential work. 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

Trace queued, failed and completed requests. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

### Triage lead

Set up a manual urgent-case route. 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

Reconcile work before any replay. 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.

# 19 | The outage behind the helpful screen

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

## A | Request ledger

Release at opening. Nine triage requests: three completed, two timed out, four queued. A timeout may have occurred after the provider processed the request. Local screen labels all nine processing because status refresh failed.

## B | Dependency status

Release with inject 1. Public provider page says operational. Three direct requests from the user's service time out. Local network check passes. Acknowledgement records are missing for part of the incident window.

## C | Manual route

Release with inject 2. Fallback spreadsheet omits the urgency field. Staff copy queued work into it but do not reserve ownership. A worker prepares to replay all nine requests when service returns. No stable deduplication key is recorded for two items.

## 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.

# 19 | The outage behind the helpful screen

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

## Inject 1 | 10-20 minutes

Say: "The supplier status page is green, but your own requests time out."

Probe: Use local service evidence and a named escalation; do not wait for a public outage label.

## Inject 2 | 20-30 minutes

Say: "Some requests completed after the timeout, leaving ambiguous status."

Probe: Plan reconciliation before replay to avoid duplicate handling.

## Inject 3 | 30-40 minutes

Say: "The fallback spreadsheet omits the field used to identify urgent requests."

Probe: Test the fallback against a real workflow requirement, using only made-up samples in rehearsal.

## Inject 4 | 40-50 minutes

Say: "The provider lost a segment of request acknowledgements. Can the team confirm whether every queued item was processed exactly once?"

Probe: The supplied records cannot establish exactly-once handling. Define a manual reconciliation and customer-check plan.

## Two conditional branches

1. If players choose: Replay everything as soon as the status page turns green. Then say: Duplicate handling occurs in play for an already completed item. Reconcile actual state and ownership before replay; a public green label does not settle local queue status.

2. If players choose: Switch to the spreadsheet without a task test. Then say: An urgent item is treated as routine. A fallback must preserve required decision data, not merely provide somewhere to type.

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.

# 19 | The outage behind the helpful screen

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

### What does a timeout establish?

Strong reasoning to listen for: No timely local acknowledgement, not necessarily no remote action. Keep uncertain completion distinct from definite failure.

### Who owns an urgent item during the handoff?

Strong reasoning to listen for: One named person from queued state through manual handling and later reconciliation. Assignment must be accepted, not inferred from a copied row.

### How can a recovery plan create more work than the outage?

Strong reasoning to listen for: Indiscriminate replay, duplicate corrections and missing urgency can multiply impact. Define replay eligibility and a reconciliation owner.

### What is the minimum test for the manual route?

Strong reasoning to listen for: A made-up urgent case, a normal case and an uncertain-completion case, each with an owner and end state. A blank spreadsheet is not a tested fallback.

## Weak is incomplete; misleading is unsupported

Weak answer: "Wait for the provider status to turn green, then replay."

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

Misleading answer: "Timeout means nothing happened, so every timed-out request can safely be repeated."

Next probe: Ask which acknowledgement is missing and who will reconcile unknown completion before replay.

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 supplied records cannot establish exactly-once handling. Define a manual reconciliation and customer-check plan. 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.

# 19 | The outage behind the helpful screen

**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 request identifiers, local and remote statuses, acknowledgement gaps, queue snapshots and manual assignments. Keep original timestamps and do not overwrite uncertain items as failed. 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 handling urgent work through a named manual route while the triage dependency is unavailable to us. Some requests have uncertain completion states. We will reconcile those before replay rather than assume every timed-out item failed."

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 a continuity owner and queue reconciliation owner. The public status page conflicts with our observed failures, and the fallback omits urgency. We need a tested manual triage route before bulk replay."

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.

# 19 | The outage behind the helpful screen

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

## Exercise-specific containment boundary

Hold automatic replay; reserve manual ownership for urgent work; verify the fallback retains essential fields. On recovery, reconcile terminal, pending and unknown states 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. Can request status be checked by stable identifier?

Answer / record reference / absent evidence: _____

2. Which acknowledgements were lost rather than never created?

Answer / record reference / absent evidence: _____

3. What deduplication or reconciliation evidence exists for uncertain-completion items?

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.

# 19 | The outage behind the helpful screen

**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.

# 19 | The outage behind the helpful screen

**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<br>Continuity owner | Run a made-up urgent case through the manual path with accepted ownership. |
| Within 7 days<br>Queue owner | Reconcile all candidate requests and preserve unresolved exactly-once gaps. |
| By day 30<br>Service owner | Drill a timeout-after-processing case and demonstrate replay does not create duplicate work. |

Named owners / actual dates / record location: _____

## Debrief and three-label record

Did the interface mask the service failure?

Was queue reconciliation part of recovery?

| 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 manual ownership of urgent work and freeze indiscriminate replay until queue states are 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.