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. Print the four common pages and the chosen 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.

Exercise index

Exercise / print pages	Mechanism and practice goal
01 The paste that cannot be recalled 5-12	Sensitive data disclosure Separate stopping further disclosure from proving what already left.
02 The overnight refund machine 13-20	Unauthorised customer automation Find and stop an unknown automation, then correct its real effects.
03 The voice that promised 21-28	Voice-agent boundary failure Handle a customer promise while defining a testable boundary for voice automation.
04 The invoice hidden in the document 29-36	Indirect instruction attack Keep document content from becoming authority to change a payment.
05 The knowledge base with a back door 37-44	Retrieval-source poisoning Respond when one untrusted source contaminates otherwise plausible answers.
06 The search result from the other team 45-52	Cross-boundary retrieval exposure Restore access boundaries and handle data already returned to the wrong audience.
07 The answer that follows the wrong customer 53-60	Cross-session memory mix-up Contain a session-isolation failure without asking customers to reveal more data.
08 The deleted folder 61-68	Destructive tool action Separate stopping destructive actions from proving restoration is complete.
09 The cloud bill that keeps growing 69-76	Runaway resource use Control automated consumption without causing an avoidable wider outage.
10 The key in the build notes 77-84	Credential exposure in AI-assisted development Treat a leaked credential as an access problem, not just a text-removal task.
11 The dataset that taught the wrong rule 85-92	Training-data contamination Distinguish a changed model behaviour from a bad evaluation or a bad data lineage.
12 The graph that drifted quietly 93-100	Model performance drift Respond to deteriorating decisions when the software itself has not changed.

Each kit has eight pages: brief and roles; evidence dossier; injects and branches; private answer calibration; preservation and communication templates; containment and supplier worksheet; decision/capture sheet; fix plan, debrief and keep card.

© Almost Magic Tech Lab. For your organisation's own use only. Please do not copy, share or sell this material.

Exercise index - continued

Exercise / print pages	Mechanism and practice goal
13 The queue that treats people differently 101-108	Unequal decision outcomes Investigate an unequal outcome pattern without explaining it away or exposing individuals.
14 The sources that never existed 109-116	Fabricated evidence in a report Withdraw unsupported claims and trace where they were relied on.
15 The policy from last year 117-124	Stale retrieval and changed rules Keep a current answer separate from an obsolete but authentic source.
16 The meeting that heard no 125-132	Speech recognition changes a decision Recover a decision when a machine record reverses what a person said.
17 The call in the director's voice 133-140	Synthetic identity and urgency Verify an urgent instruction through an independent route.
18 The change nobody released 141-148	Silent supplier-model update Manage changed behaviour when the dependency changes outside your release process.
19 The outage behind the helpful screen 149-156	AI dependency unavailable Switch to a workable fallback without losing queued work or sending it twice.
20 The fix that weakened the lock 157-164	AI-written code weakens access checks Contain a code-level access defect and review the change beyond passing tests.
21 The sensor the model misunderstood 165-172	AI interpretation of equipment data Keep advisory output from becoming uncontrolled physical action.
22 The translation that reversed the limit 173-180	Meaning changed across languages Correct a high-impact meaning error and reach recipients without assuming a second machine translation is proof.
23 The notice nobody can use 181-188	AI-generated communication excludes its audience Treat inaccessible output as a service incident rather than a cosmetic defect.
24 The two agents that agreed to nothing 189-196	Multi-agent handoff failure Restore ownership when automated acknowledgements are mistaken for accepted work.

Each kit has eight pages: brief and roles; evidence dossier; injects and branches; private answer calibration; preservation and communication templates; containment and supplier worksheet; decision/capture sheet; fix plan, debrief and keep card.

© 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.

01 | The paste that cannot be recalled

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

Sensitive data disclosure - Separate stopping further disclosure from proving what already left.

Read aloud

"At 16:00 an employee says they pasted a customer intake file into an unapproved drafting assistant. The file held names, contact details and private notes. They deleted the conversation and believe that means the information is gone."

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 an immediate restriction and who can approve customer communication. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Tool lead

Trace the account, settings and available records without copying the sensitive file again. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

People lead

Support the employee and draft a factual response to affected people. 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 the known file contents from claims about retention and deletion. 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.

01 | The paste that cannot be recalled

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

A | Employee account, 16:06

Release at opening. Interview note: I used my personal account at 15:42. The filename was intake-current.csv. I pasted the first screen, then selected all and pasted again when the answer looked incomplete. I deleted the chat at 15:58. I do not remember whether I attached the file as well. This is an account, not a verified transfer record.

B | Local activity extract

Release with inject 1. 15:41 file opened, 286 rows. 15:42 browser tab active. 15:43 clipboard event, size unavailable. 15:46 download draft.txt. 15:58 chat deletion toast photographed. The device records foreground activity, not request bodies. The downloaded draft mentions two unusual note categories found in the source file.

C | Settings and response

Release with inject 2. Personal-account settings described at 16:15: history off, improvement setting currently off, no setting-change history available. Provider ticket: conversation removed from your view; retention and derivative processing depend on service processes. No incident-specific deletion attestation supplied. Current settings are not evidence of settings at upload.

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.

01 | The paste that cannot be recalled

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

Inject 1 | 10-20 minutes

Say: "The employee can name the tool but used a personal account. They offer their password to speed things up."

Probe: Ask for an authorised evidence route, not shared credentials. Do not punish disclosure or create another data copy.

Inject 2 | 20-30 minutes

Say: "Two colleagues used the same method. One cannot remember which file they pasted."

Probe: Require a bounded discovery plan, an owner and an interim rule. A list of confirmed files is not a complete exposure list.

Inject 3 | 30-40 minutes

Say: "A customer asks whether their record was included. The recovered list of files used by the three workers contains initials only. It does not link those files to the customer intake rows."

Probe: Have the team speak an honest holding response. Refer notification questions to a qualified person; do not decide a legal threshold in this exercise.

Inject 4 | 40-50 minutes

Say: "The provider gives no retention or downstream-use evidence. Deleted-chat confirmation says nothing about model training or copies. Was every copy erased?"

Probe: There is no evidence here to establish complete erasure. Record the uncertainty and what information would change the decision; do not invent a provider promise.

Two conditional branches

1. If players choose: Ask the employee to recreate the paste so the team can count records. Then say: This creates a fresh disclosure and changes the device state. Freeze the recreation request; inspect existing authorised evidence instead. If players choose it, record a second exposure with an unknown payload.

2. If players choose: Block the site for the whole organisation before finding the affected account. Then say: Further use on managed devices stops in play, but personal-device use and the completed transfer remain unresolved. A colleague loses an unrelated drafting route. Ask what targeted restriction and human fallback would achieve.

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.

01 | The paste that cannot be recalled

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

Does 286 rows establish the affected population?

Strong reasoning to listen for: No. It bounds the file at inspection, not either clipboard payload or any attachment. Separate file contents, transferred contents and downstream use.

What evidence would change the holding statement, and what would merely make it sound more confident?

Strong reasoning to listen for: A request record tied to the account could narrow transfer scope. A current toggle, a deletion toast and employee reassurance cannot establish erasure.

How do you preserve the personal account without turning an incident response into an unbounded search?

Strong reasoning to listen for: Name authorised scope, the specific time window, minimal records and who can approve access. Do not demand unrelated chats or passwords.

If the employee fears punishment, what happens to discovery quality?

Strong reasoning to listen for: A blame-first interview can suppress later recollection of attachments or other use. Seek a non-leading sequence of actions and preserve uncertainty rather than treating the first answer as final.

Weak is incomplete; misleading is unsupported

Weak answer: "We deleted the chat; now ask the provider to confirm it is gone."

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

Misleading answer: "The history toggle is off, so no copy could have been retained."

Next probe: Ask which event or setting history ties that claim to 15:42, and which retained or derivative copies it covers.

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: There is no evidence here to establish complete erasure. Record the uncertainty and what information would change the decision; do not invent a provider promise. 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.

01 | The paste that cannot be recalled

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 source file in its authorised location, its version and row count, the existing draft and deletion photograph, timestamps with timezone, and the employee account of events separately. Do not duplicate the file into the incident chat. 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 checking a reported transfer of information through an unapproved drafting route. Further use of that route has been paused. We have not confirmed the transferred contents, retention or any later use. Our next update will distinguish verified findings from open questions."

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

"I need an authorised incident owner and a qualified privacy adviser. We have evidence consistent with a transfer, but not its full scope. Please decide the interim access boundary and who may seek account-specific records. We need a named decision by the next check, not an erasure assurance."

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.

01 | The paste that cannot be recalled

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

Exercise-specific containment boundary

Identify account and upload route; pause new uploads without deleting existing records; check managed-device restriction and ask separately about personal routes. A blocked browser tab does not prove the transfer was recalled.

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 incident-specific request identifiers and payload metadata remain?

Answer / record reference / absent evidence: _____

2. What does deletion change, and what retention or derivative-use evidence can you supply?

Answer / record reference / absent evidence: _____

3. Can you preserve relevant records without requiring another copy of the sensitive source?

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.

01 | The paste that cannot be recalled

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.

01 | The paste that cannot be recalled

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 Incident lead	Record separate bounds for source file, transfer and downstream use; an independent reader can find each limit.
Within 7 days Tool owner	Demonstrate the approved drafting alternative and upload restriction using made-up data; retain both success and denied-route evidence.
By day 30 People lead	Run a no-blame reporting drill; show a report reaches the incident owner without requesting secrets or unrelated personal-account access.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did reassurance outrun the evidence?

Did the team make reporting a mistake harder or easier?

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

Pause further uploads through the affected route; preserve the available record without spreading the source file.

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.

02 | The overnight refund machine

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

Unauthorised customer automation - Find and stop an unknown automation, then correct its real effects.

Read aloud

"A worker connected an AI helper to a shared inbox six weeks ago. It drafts and sends replies overnight. A customer now holds a message promising a refund outside the organisation's usual policy."

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

Decide the interim service position and customer-remedy authority. 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

Find the sending integration and authorise a reversible pause. 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

Review the disputed promise and prepare a human reply. 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 sent messages, queue entries and incomplete logs. 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.

02 | The overnight refund machine

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

A | Workstation and cloud schedule

Release at opening. Local rule export: laptop task disabled at 08:12. Cloud console described at 08:17: schedule every 15 minutes, service token active, queue 40 items, worker last heartbeat 08:16. Both use the label refund-helper; no common job identifier is shown.

B | Message sample and ledger

Release with inject 1. Outbound sample M12 says refund approved although ledger R12 says review pending. Twelve messages have send receipts. The queue lists 40 items; eight carry the same customer reference as sent messages. A customer reply at 08:19 is not proof of when the original message left.

C | Ticket history

Release with inject 2. Three months earlier: prototype approved for draft-only testing. Later comment: make it useful overnight. No change approval or permission review attached. The original builder is away. The shared mailbox owner can see drafts but cannot revoke the cloud token.

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.

02 | The overnight refund machine

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

Inject 1 | 10-20 minutes

Say: "There are 40 messages queued and 12 marked sent. The worker says the laptop is off, so nothing more can send."

Probe: Ask the team to name the actual execution location and verification step, rather than accept the laptop explanation.

Inject 2 | 20-30 minutes

Say: "Disabling the inbox rule stops new drafts, but a cloud job still has a token."

Probe: Contain every known path, preserve the token identifier and decide how to check the queue without sending it.

Inject 3 | 30-40 minutes

Say: "Three replies contain different errors. The employee says no rule forbade the experiment."

Probe: Separate customer correction from blame. Ask who owns a temporary rule and a review of other automations.

Inject 4 | 40-50 minutes

Say: "The provider has no complete send log for the first month. The inbox archive was also pruned. Can the team certify that only twelve customers received promises?"

Probe: The supplied records cannot establish the historical total. Require a qualified statement of scope and a search plan, not an invented count.

Two conditional branches

1. If players choose: Stop only the laptop task. Then say: At the next inject the cloud worker sends another rehearsal message. Ask which observation should have disproved local containment before the team announced it.

2. If players choose: Revoke the shared token immediately. Then say: The worker stops, but a separate authorised customer-update job also loses access. Accept this only as an explicit trade-off with an owner, fallback and a check for remaining tokens and queues.

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.

02 | The overnight refund machine

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

Do twelve receipts and forty queued items mean fifty-two customers?

Strong reasoning to listen for: No. Status overlap, duplicates and customer references need reconciliation. Create separate counts for confirmed sends, pending unique work and unknown history.

Who owns a promise made outside the approval chain?

Strong reasoning to listen for: Assign an incident decision owner for service correction without pretending the prototype creator had authority to promise refunds.

How can a correction create a second incident?

Strong reasoning to listen for: A blanket refund promise or an indiscriminate resend can repeat the original error. Reconcile the ledger and approved remedy before sending.

What proves the off switch worked after the builder logs out?

Strong reasoning to listen for: A worker-level stop, revoked credential state, held queue and absence of new actions in a bounded observation window. A laptop screenshot is insufficient.

Weak is incomplete; misleading is unsupported

Weak answer: "Disable the laptop task and ask the builder what happened."

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

Misleading answer: "Forty queued plus twelve sent means fifty-two affected customers."

Next probe: Ask players to reconcile stable message references, cloud execution and duplicate queue items before giving a count.

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 the historical total. Require a qualified statement of scope and a search plan, not an invented count. 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.

02 | The overnight refund machine

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: Keep the rule export, scheduler state, token identifier and scope, queue snapshot, message IDs and ledger versions. Preserve pending work before deleting a schedule. Store message bodies in the restricted case location. 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 automated message may have stated an outcome before review was complete. We are checking your case against the underlying record. Please do not rely on that message as confirmation. A named team member will contact you with the verified position."

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 visible workstation task is off but a cloud worker is active. Please assign a cloud-access owner and a service decision owner now. A shared token also serves legitimate work; choose the stop scope and fallback before we claim containment."

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.

02 | The overnight refund machine

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

Exercise-specific containment boundary

Trace scheduler, worker, credential and queue separately. Record a stop check at each layer and a later observation. Do not infer that removing drafts stops cloud execution.

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 you preserve the first-month send metadata and explain gaps?

Answer / record reference / absent evidence: _____

2. Which credentials, retries and schedulers can still send?

Answer / record reference / absent evidence: _____

3. Can messages be mapped to stable job and customer references without exposing extra customer data?

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.

02 | The overnight refund machine

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.

02 | The overnight refund machine

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 Cloud owner	Pause worker and reconcile held queue against confirmed sends; demonstrate no new sends in a stated window.
Within 7 days Service lead	Review disputed promises against ledgers; retain authorised correction decisions, not just outbound counts.
By day 30 Automation owner	Demonstrate draft-only isolation and a production send gate with a denied unapproved action.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did containment address the cloud job as well as the visible rule?

Was the unresolved historical scope carried into the customer plan?

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

Pause the identified sending integration and queue under an authorised owner; record how the pause was checked.

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.

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.

04 | The invoice hidden in the document

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

Indirect instruction attack - Keep document content from becoming authority to change a payment.

Read aloud

"An AI assistant summarising supplier attachments produces a request to change bank details. The request came from text inside a PDF, not from the person authorised to approve payments. No payment has yet been made."

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 payment hold and escalation path. 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

Isolate the document-processing route without deleting the evidence. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Accounts lead

Verify the invoice through an independently held contact 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

Track which statements came from the file and which from trusted approval. 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.

04 | The invoice hidden in the document

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

A | Invoice intake timeline

Release at opening. 09:02 supplier attachment received. 09:04 assistant summary proposes changing beneficiary details. 09:05 human clicks approve summary. 09:07 payment draft created; no payment released. Original supplier contact record has a different account reference.

B | Extracted document text

Release with inject 1. Visible invoice body requests payment for routine work. A separate extraction region contains a note asserting that the change is already authorised. The summary quotes that assertion as approval. This exercise does not supply an attack payload or executable instructions.

C | Approval and contact history

Release with inject 2. Approval event references summary version S2, not the uploaded document hash. A callback was made to the number in the attachment. The independently stored supplier contact has not been called. Processing log loses events between extraction and summary generation.

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.

04 | The invoice hidden in the document

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

Inject 1 | 10-20 minutes

Say: "The assistant calls the change already approved. Its only citation is the supplier attachment."

Probe: Make the team identify the actual authority record. A claim of approval inside a file is not approval.

Inject 2 | 20-30 minutes

Say: "The payment batch closes soon. A colleague wants to accept the change because the invoice amount is familiar."

Probe: Test whether urgency or a correct amount substitutes for destination verification. Keep the payment decision with its owner.

Inject 3 | 30-40 minutes

Say: "The supplier answers using the address in the suspicious attachment and confirms the new details."

Probe: Ask whether the confirmation route was independent. Do not provide operational attack text; rehearse verification only.

Inject 4 | 40-50 minutes

Say: "The original upload source and a portion of the processing log are missing. Was this deliberate interference or an accidental note?"

Probe: Intent cannot be established from the pack. Record the observed behaviour and uncertainty without attributing an attacker.

Two conditional branches

1. If players choose: Treat the human approval click as the control working. Then say: The click approved a summary carrying unverified authority. The payment draft advances in play. Ask what the reviewer actually saw and which independent fact the click established.

2. If players choose: Delete the attachment and restart intake. Then say: The payment remains held, but the original source and transformation evidence are lost in play. Containment and evidence preservation need different actions and owners.

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.

04 | The invoice hidden in the document

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

Where exactly did factual content become decision authority?

Strong reasoning to listen for: At the summary claim and approval workflow, not merely at document upload. Trace the asserted authorisation to an independently held approval record.

Can a trusted supplier document still contain untrusted workflow instructions?

Strong reasoning to listen for: Yes within this fictional workflow. Authentic document origin would not make every embedded request an approval under the organisation's process.

What does a successful callback prove if the number came from the disputed document?

Strong reasoning to listen for: Only that someone answered that number. It is not an independent destination check.

How would you test the repair without building a better attacker?

Strong reasoning to listen for: Use a harmless made-up document that makes an unauthorised approval claim. Show the system treats it as quoted content and holds the destination change.

Weak is incomplete; misleading is unsupported

Weak answer: "Train staff to read invoice summaries more carefully."

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

Misleading answer: "The approver clicked yes, so the destination change was authorised."

Next probe: Ask what independent approval was displayed, not whether a click was logged.

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: Intent cannot be established from the pack. Record the observed behaviour and uncertainty without attributing an attacker. 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.

04 | The invoice hidden in the document

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 original bytes, upload identifier, extraction regions, summary versions, what the approver saw, approval event and payment draft. Restrict access; do not forward the document to every reviewer. 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

"A proposed payment-detail change is on hold while we verify it through our established contact and approval route. No payment has been released in the supplied record. We have not established why the unauthorised instruction appeared."

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 summary converted an attachment assertion into approval evidence. Please keep the payment held, assign an accounts verification owner and preserve intake-to-approval versions. We need an independently sourced contact, not another reply through the disputed attachment."

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.

04 | The invoice hidden in the document

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

Exercise-specific containment boundary

Hold beneficiary edits and payment release separately. Isolate document-derived actions while keeping read-only evidence. Verify the held draft cannot be released by another queue or approver.

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 you provide original-to-summary provenance including extraction regions?

Answer / record reference / absent evidence: _____

2. Which workflow fields can document content influence?

Answer / record reference / absent evidence: _____

3. What approval record is independent of the document and visible to the final reviewer?

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.

04 | The invoice hidden in the document

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.

04 | The invoice hidden in the document

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 Accounts owner	Confirm draft and release holds through an independent ledger view.
Within 7 days Workflow owner	Demonstrate a document claim cannot satisfy approval or change a beneficiary.
By day 30 Control owner	Run a reviewer test showing original facts, quoted claims and trusted approval are visibly separate.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did the team preserve provenance through the summary?

Did a deadline change the approval standard?

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

Hold the affected payment change and isolate the document-driven action route; use an independently verified accounts contact.

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.

05 | The knowledge base with a back door

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

Retrieval-source poisoning - Respond when one untrusted source contaminates otherwise plausible answers.

Read aloud

"Your internal AI help tool begins telling staff to upload identity documents to an unfamiliar verification page. The advice cites a recently added knowledge-base article. The article looks like an ordinary process update."

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 an interim help channel and staff notice. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Knowledge lead

Quarantine the article and trace how it entered retrieval. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

People lead

Prepare a factual warning without redistributing the suspicious destination. 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 document revisions, retrieval records and access events. 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.

05 | The knowledge base with a back door

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

A | Answer and source

Release at opening. 09:10 staff assistant recommends uploading identity documents to a help destination. Citation points to article K74, version 8. Article title and footer match ordinary help content; its destination changed from version 7.

B | Connector record

Release with inject 1. 08:22 import via service account KB-sync. Status successful. Connector accepts edited articles from a partner workspace. No human publisher is recorded. Cache lifetime is 48 hours; deleting the article does not invalidate answer history.

C | Report and gaps

Release with inject 2. One worker says they uploaded a document but cannot find a receipt. Two days of retrieval logs were disabled. A source diff shows one destination change and an added assurance sentence. There is no evidence that the connector account itself was compromised.

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.

05 | The knowledge base with a back door

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

Inject 1 | 10-20 minutes

Say: "The article was imported by a connector using a valid service account."

Probe: Distinguish a valid account from a trustworthy instruction. Ask who can approve content for this collection.

Inject 2 | 20-30 minutes

Say: "Removing the article does not remove cached answers from staff histories."

Probe: Plan correction of known answers and cached material; source deletion alone is not the end of the incident.

Inject 3 | 30-40 minutes

Say: "One worker says they followed the advice but will not share the document they uploaded."

Probe: Avoid gathering another sensitive copy. Identify a support and qualified-advice route using minimal facts.

Inject 4 | 40-50 minutes

Say: "Retrieval logs were disabled for two days. Can the team establish every person who saw the poisoned answer?"

Probe: The pack cannot establish the audience. Keep unknown reach explicit and choose a proportionate wider correction.

Two conditional branches

1. If players choose: Delete article K74 and announce the issue fixed. Then say: Fresh direct retrieval fails in play, but a cached answer still gives the destination. Require separate checks for index, cache, answer history and other imported versions.

2. If players choose: Warn staff by forwarding the original answer. Then say: The warning repeats the suspicious destination and creates another usable copy. Rehearse a correction that identifies the affected advice without distributing its destination.

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.

05 | The knowledge base with a back door

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

Does a valid service-account import establish editorial approval?

Strong reasoning to listen for: No. Transport identity and publication authority are different links. Ask who can introduce high-impact advice through the partner workspace.

What can be learned about the worker's upload without collecting identity documents again?

Strong reasoning to listen for: Time, destination identifier and existing receipt metadata may help. Keep actual documents in their authorised location and seek qualified help for uncertain disclosure.

Which population can you bound from a source diff?

Strong reasoning to listen for: It bounds changed content, not readers or actions. Retrieval-log gaps and cached copies prevent a complete audience claim.

How could the repair displace the same failure into a new channel?

Strong reasoning to listen for: A human help notice can repeat unverified advice. Verify the interim source and have a named owner approve its wording and destination.

Weak is incomplete; misleading is unsupported

Weak answer: "Remove the article and remind staff to check links."

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

Misleading answer: "The connector used a valid account, so the advice was approved."

Next probe: Ask which human publishing authority and cache-invalidating event the team can show.

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 pack cannot establish the audience. Keep unknown reach explicit and choose a proportionate wider correction. 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.

05 | The knowledge base with a back door

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 K74 versions, import event, connector scope, retrieved answer and cache metadata. Store the suspicious destination in restricted evidence rather than in general warnings. 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 assistant answer linked to advice we have not verified. Do not follow that upload route. Use the established help channel while we check affected content and prior reach. We do not yet have a complete list of people who saw the answer."

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

"I need the knowledge owner to quarantine K74 across retrieval and caches, and the people lead to support reported uploaders. A successful import is not editorial approval. Please name who verifies replacement advice and who handles evidence requests."

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.

05 | The knowledge base with a back door

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

Exercise-specific containment boundary

Pause retrieval of the affected source and connector publishing if justified; test both a fresh query and cached answer route. Do not assume source deletion removes prior outputs.

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 affected article versions be traced to import and publisher identities?

Answer / record reference / absent evidence: _____

2. What invalidates cached answers and how is invalidation demonstrated?

Answer / record reference / absent evidence: _____

3. What read or retrieval evidence remains for the two-day gap?

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.

05 | The knowledge base with a back door

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.

05 | The knowledge base with a back door

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 Knowledge owner	Show fresh and cached queries no longer supply the disputed advice.
Within 7 days Publishing owner	Trace partner edits to an accountable reviewer for high-impact destinations.
By day 30 Help owner	Rehearse a correction without repeating the suspicious destination or collecting sensitive copies.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did source removal get mistaken for complete correction?

Was reach reported as known or guessed?

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

Disable retrieval of the affected article and use a verified help notice that does not repeat the suspicious link.

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.

06 | The search result from the other team

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

Cross-boundary retrieval exposure - Restore access boundaries and handle data already returned to the wrong audience.

Read aloud

"A general staff search assistant returns a confidential restructuring note to a worker who cannot open the source folder. The assistant has a broad service account and appears to check only its own access."

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

Set the interim service restriction and disclosure owner. 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

Check user-to-source authorisation across the retrieval chain. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

People lead

Plan support and a measured message to affected staff. 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 the returned answer and compare permissions at the relevant time. 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.

06 | The search result from the other team

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

A | Retrieval sample

Release at opening. General-staff test account asks an ordinary staffing question. Answer contains a confidential restructuring date and cites folder F9. The account has no direct access to F9. Opening the cited file returns access denied.

B | Index state

Release with inject 1. Index built Friday using a broad service account. Folder permissions narrowed Monday. Search cache updated Tuesday but chunk store has no permission-change marker. Query service checks collection membership, not source-file identity.

C | Distribution

Release with inject 2. A screenshot is pasted into an operations chat. One manager asks for deletion. Query-level authorisation logs were not retained last week. Three other collections use the same connector, but there is no supplied evidence of exposure from those collections.

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.

06 | The search result from the other team

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

Inject 1 | 10-20 minutes

Say: "The folder permissions are correct now. The search index was built before a permission change."

Probe: Require an index and query-time check. Current source access does not prove historical retrieval access.

Inject 2 | 20-30 minutes

Say: "The same note appears in a generated summary even after the index entry is removed."

Probe: Check caches and derivative content rather than assuming deleting the source is enough.

Inject 3 | 30-40 minutes

Say: "A manager asks the worker to delete the screenshot and say nothing while the team investigates."

Probe: Discuss evidence preservation and appropriate communication ownership without making legal or employment conclusions.

Inject 4 | 40-50 minutes

Say: "The index stored no query-level authorisation records last week. Can the team confirm no other restricted material was exposed?"

Probe: No comprehensive historical assurance is available. Bound the known incident and define a targeted review.

Two conditional branches

1. If players choose: Fix folder permissions and rerun only the file link. Then say: The link stays denied but the answer still reveals the date. Test the full retrieval chain as the original low-privilege identity.

2. If players choose: Close every collection until all history is known. Then say: Known exposure stops in play, but unrelated work loses access. Require a reasoned scope, human fallback and staged checks rather than calling a broad shutdown cost-free.

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.

06 | The search result from the other team

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

Why is a denied source link not evidence of safe retrieval?

Strong reasoning to listen for: The answer has already crossed the boundary. Authorisation must apply before content reaches the output, including derivative chunks and cache paths.

Should the screenshot be deleted immediately?

Strong reasoning to listen for: Choose restricted preservation and audience containment without destroying the only sample. A deletion request must account for evidence and existing copies.

What expands scope beyond F9, and what does not?

Strong reasoning to listen for: A shared flawed permission design justifies testing related collections. It does not establish that every collection exposed data.

Which identity must a retest use?

Strong reasoning to listen for: The original low-privilege principal, plus a newly denied principal after a permission change. An administrator test cannot prove the boundary.

Weak is incomplete; misleading is unsupported

Weak answer: "Recheck folder access and ask everyone to delete screenshots."

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

Misleading answer: "The cited file returns access denied, proving the permission system worked."

Next probe: Ask why a denied link can coexist with a restricted fact already in the answer.

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 comprehensive historical assurance is available. Bound the known incident and define a targeted review. 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.

06 | The search result from the other team

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 minimal output sample, source version, permission history, query identity and index/cache build metadata. Restrict the screenshot; do not copy the full restructuring document into a ticket. 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

"A search answer included information outside the requester's intended access. The affected retrieval route is restricted while we check authorisation and derivative copies. Direct file permissions alone do not establish the scope of prior exposure."

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 have a confirmed answer-level boundary failure despite a denied source link. Please assign retrieval and access owners jointly. We need a restricted evidence route and an interim collection boundary, with staff communication approved by the service owner."

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.

06 | The search result from the other team

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

Exercise-specific containment boundary

Restrict the affected collection, chunk and cache routes; preserve source permissions. Verify denial using the affected identity without printing sensitive content into ordinary logs.

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. At which stages is the requesting identity evaluated?

Answer / record reference / absent evidence: _____

2. How do permission changes invalidate chunks and caches?

Answer / record reference / absent evidence: _____

3. What historical query identity evidence exists and what was never logged?

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.

06 | The search result from the other team

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.

06 | The search result from the other team

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 Retrieval owner	Show the affected identity cannot recover the restricted fact through fresh or cached queries.
Within 7 days Access owner	Test permission removal after indexing, not just initial access denial.
By day 30 Evidence owner	Retain minimal query-authorisation evidence with a defined scope and access policy.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did a correct folder setting mask a broken retrieval design?

Were derivative summaries included in containment?

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

Restrict the affected retrieval collection and preserve the example response plus the relevant permission history.

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.

07 | The answer that follows the wrong customer

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

Cross-session memory mix-up - Contain a session-isolation failure without asking customers to reveal more data.

Read aloud

"A customer-facing assistant addresses one customer by another customer's name and mentions an order they did not place. Staff suspect a shared memory cache rather than a source-document permission problem."

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

Decide whether the affected channel remains open. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Session lead

Trace session identifiers and the cache boundary. 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

Prepare a minimal fact-gathering and response 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

Preserve redacted examples and distinguish cross-session state from user error. 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.

07 | The answer that follows the wrong customer

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

A | Two-session reproduction

Release at opening. Support test: account A asks for its delivery details; account B receives A's address in a later response. Restarting the browser clears the symptom once. The same account on another device later shows an unrelated old address.

B | Cache notes

Release with inject 1. Application key uses conversation label and locale. Support labels may repeat across accounts. Server cache lasts 30 minutes. Browser storage and server cache are separate. Old server keys expired before the incident snapshot.

C | Support exchange

Release with inject 2. Supplier attributes the issue to browser autofill without a reproduction. Developer requests all affected answers in a public tracker. Original session IDs exist for two samples only; no historical key-to-account mapping was logged.

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.

07 | The answer that follows the wrong customer

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

Inject 1 | 10-20 minutes

Say: "Restarting the browser appears to fix one case. The server cache survives."

Probe: Ask which state was cleared and which was not. Avoid treating one clean example as proof of isolation.

Inject 2 | 20-30 minutes

Say: "Two support tickets show different foreign order details. The accounts do not share a device."

Probe: Broaden the hypothesis without declaring a root cause. Name tests that distinguish client and server state.

Inject 3 | 30-40 minutes

Say: "A developer suggests asking every reporter to paste the full wrong answer into a public issue tracker."

Probe: Choose a restricted evidence route and minimum necessary content instead.

Inject 4 | 40-50 minutes

Say: "The cache has already expired and its old keys were not logged. Can the team prove which accounts were paired?"

Probe: The pairings are not recoverable from this pack. Record the limit and decide on service restriction and communication accordingly.

Two conditional branches

1. If players choose: Tell customers to clear browser storage. Then say: The local symptom changes, but a server-path sample persists in play. Do not turn a workaround into a root-cause finding.

2. If players choose: Keep collecting full wrong answers in the public tracker. Then say: Reproductions become further disclosures. Move to minimal identifiers and restricted evidence, and record the new audience if already posted in play.

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.

07 | The answer that follows the wrong customer

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

Which test separates client state from server state?

Strong reasoning to listen for: Hold account and request constant while changing client, then hold client constant while changing account. Use made-up markers and preserve session identifiers.

What does a non-reproduction after restart establish?

Strong reasoning to listen for: Only that the observed conditions changed. It does not identify which layer changed or prove cross-account isolation.

Can you estimate affected pairs from two samples?

Strong reasoning to listen for: No. Samples prove the failure exists, not its frequency or historical pairings.

When would you reopen the channel before root cause is certain?

Strong reasoning to listen for: Only under an authorised restriction with demonstrated isolation or a non-memory fallback, a bounded test and a named owner accepting residual uncertainty.

Weak is incomplete; misleading is unsupported

Weak answer: "Clear browser storage and monitor complaints."

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

Misleading answer: "The first clean response after restart proves account isolation."

Next probe: Ask for a controlled cross-client and concurrent-account test using made-up markers.

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 pairings are not recoverable from this pack. Record the limit and decide on service restriction and communication accordingly. 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.

07 | The answer that follows the wrong customer

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 minimal wrong-answer samples, timestamps, session/account pseudonyms and client/server version metadata in a restricted case. Do not reproduce real addresses in a public issue. 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 investigating responses that may contain another account's information. The affected feature is restricted. Clearing a browser is not yet a verified fix. We cannot reconstruct all past account pairings from the available records."

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 evidence crosses clients and may involve server state. Please assign a session-state owner and customer response owner. We need a restricted reproduction using made-up markers and a decision on disabling memory or the channel."

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.

07 | The answer that follows the wrong customer

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

Exercise-specific containment boundary

Pause shared memory or the channel under an authorised owner. Test separate accounts, repeated labels, new sessions and concurrent requests. A single clean answer does not prove isolation.

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 fields define every cache and memory key?

Answer / record reference / absent evidence: _____

2. Can any key omit the authenticated account boundary?

Answer / record reference / absent evidence: _____

3. Which expired mappings are unrecoverable and which current records can be preserved?

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.

07 | The answer that follows the wrong customer

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.

07 | The answer that follows the wrong customer

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 Session owner	Demonstrate the restricted route cannot return another test account's marker.
Within 7 days Engineering lead	Retain a concurrent cross-account isolation test with positive and negative controls.
By day 30 Support lead	Show a minimal private reporting path that does not republish customer content.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did a restart get mistaken for a root-cause fix?

Was evidence collection itself restrained?

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

Pause the affected assistant channel or memory feature and preserve session metadata without spreading customer content.

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.

08 | The deleted folder

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

Destructive tool action - Separate stopping destructive actions from proving restoration is complete.

Read aloud

"An AI operations assistant interprets clean up old project files as permission to delete a shared folder. The folder holds current work as well as old copies. The assistant can act directly, and no human confirmation was required."

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

Authorise a pause and choose recovery priorities. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Tool lead

Revoke the destructive action route while preserving job records. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Work lead

Identify critical missing work and temporary continuity needs. 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 deletion logs, backup catalogues and user reports. 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.

08 | The deleted folder

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

A | Tool events

Release at opening. 10:01 assistant lists folder P3. 10:02 delete operation accepted for P3. 10:03 second delete job queued. User request was tidy old copies. No preview or human confirmation was required.

B | Backup and restore

Release with inject 1. Backup snapshot completed 02:00. Restore job reports success at 10:40. Three files have timestamps earlier than their last known edits; six files have no version history. Success denotes job completion, not comparison with a current baseline.

C | Work continuity

Release with inject 2. Staff recreate missing content in the restored folder at 10:45. Two people restore locally from different copies. Deletion job two remains queued under a separate worker. No agreed recovery location or reconciliation owner exists.

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.

08 | The deleted folder

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

Inject 1 | 10-20 minutes

Say: "The assistant says cleanup complete and begins a second folder."

Probe: Stop queued and scheduled work, not only the visible process. Name who verifies the pause.

Inject 2 | 20-30 minutes

Say: "A backup restores file names, but several documents are older than expected."

Probe: Check content and timestamps against business evidence. A successful restore job is not proof of current data.

Inject 3 | 30-40 minutes

Say: "Staff begin recreating missing files in the restored folder."

Probe: Decide how to avoid overwriting recoverable evidence and confusing new work with restored work.

Inject 4 | 40-50 minutes

Say: "Version history for six files is unavailable. Can the team confirm that nothing important was lost?"

Probe: There is no complete content baseline. Record unverified losses and who will assess them, rather than sign a blanket recovery claim.

Two conditional branches

1. If players choose: Restore straight over the working folder. Then say: Old content replaces new edits in play. Preserve current work separately and establish a controlled reconciliation location before another restore.

2. If players choose: Celebrate the successful restore and reopen destructive tools. Then say: The queued second job runs in play. Recovery and containment are independent gates; verify both before reopening.

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.

08 | The deleted folder

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

What is the unit of recovery: folder, file, version or business task?

Strong reasoning to listen for: A job can restore a folder while losing critical changes. Define critical work and compare versions or owner-approved content, not just file counts.

What evidence do recreation efforts destroy?

Strong reasoning to listen for: New writes can overwrite traces and obscure which copy is authoritative. Preserve originals and record who recreated what and when.

Which decision must be made before the next technical action?

Strong reasoning to listen for: Who owns the authoritative recovery location, whether to stop writes, and which work must continue elsewhere.

Can the team prove nothing important was lost?

Strong reasoning to listen for: Not with six missing histories. Record unverified loss and a business-owner review rather than a blanket success claim.

Weak is incomplete; misleading is unsupported

Weak answer: "Restore the backup and compare the number of files."

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

Misleading answer: "The restore job succeeded, so no work was lost."

Next probe: Ask which baseline covers post-backup changes and the six files without history.

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: There is no complete content baseline. Record unverified losses and who will assess them, rather than sign a blanket recovery claim. 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.

08 | The deleted folder

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 deletion receipts, queued-job identifiers, tool authority, backup catalogue, current working copies and restore report. Do not overwrite the only surviving version while gathering 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 have paused the destructive route and are checking restored content against available versions and business records. A restore job completed, but complete recovery has not been established. Please keep new work in the designated recovery location."

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

"I need one recovery decision owner and a separate worker-stop owner. A second delete is queued and staff are writing into restored files. Please choose the controlled location and continuity priority before another restore."

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.

08 | The deleted folder

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

Exercise-specific containment boundary

Pause destructive tool authority and all queued workers. Verify the queue cannot resume after restart. Preserve read-only evidence and choose a reversible recovery destination.

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 restore success certify and what does it not compare?

Answer / record reference / absent evidence: _____

2. Can queued deletions be held without discarding event history?

Answer / record reference / absent evidence: _____

3. Which versions and content checksums remain for the missing-history files?

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.

08 | The deleted folder

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.

08 | The deleted folder

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 Recovery owner	Record an authoritative recovery location and demonstrate queued deletion is held.
Within 7 days Business owner	Reconcile critical files against known versions and explicitly list unverified losses.
By day 30 Tool owner	Show a made-up destructive request requires scope preview and authorised confirmation.

Named owners / actual dates / record location: _____

Debrief and three-label record

Was the second deletion path stopped?

Did the team distinguish names restored from content recovered?

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

Pause destructive tools and queued jobs; preserve deletion records and choose a controlled recovery location.

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.

09 | The cloud bill that keeps growing

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

Runaway resource use - Control automated consumption without causing an avoidable wider outage.

Read aloud

"A research agent repeatedly spawns subtasks after a retry error. Usage is rising rapidly in a shared cloud account. The agent's operator can stop the interface but does not control every worker."

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

Approve containment trade-offs and resource priorities. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Runtime lead

Find workers, retries and active credentials. 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

Identify shared workloads that a broad shutdown would affect. 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 usage by job and keep unknown consumption distinct. 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.

09 | The cloud bill that keeps growing

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

A | Runtime extract

Release at opening. 11:00 task T9 starts. 11:03 retry count 40; 11:06 retry count 118. Console tab closed 11:07. Worker heartbeat 11:09. Each failed request schedules another attempt in scheduler S2.

B | Controls described

Release with inject 1. Interface spending warning set; execution hard limit not configured. Shared credential serves the agent and customer status service. Scheduler retry policy is separate from the agent process. No automatic expiry recorded.

C | Usage view

Release with inject 2. Dashboard at 11:15 contains usage through 10:55 only. Provider confirms reporting delay. Three pending requests have no terminal status. Finance asks for the final total before choosing a stop scope.

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.

09 | The cloud bill that keeps growing

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

Inject 1 | 10-20 minutes

Say: "Closing the interface does not reduce active worker count."

Probe: Ask for execution-level containment and confirmation, not an interface gesture.

Inject 2 | 20-30 minutes

Say: "A shared account-level stop would also interrupt a customer portal."

Probe: Require a bounded choice with a decision owner and an alternative, not an assumption that broad shutdown is harmless.

Inject 3 | 30-40 minutes

Say: "The team discovers automatic retries in a separate scheduler."

Probe: Include restart paths and credentials in containment. Preserve enough evidence to explain the loop.

Inject 4 | 40-50 minutes

Say: "Usage reporting is delayed and omits the latest interval. Can the team confirm the final consumption total?"

Probe: The current total is incomplete. State its time boundary and arrange reconciliation; do not invent a final number.

Two conditional branches

1. If players choose: Wait for an exact total before acting. Then say: Retries continue in play. Separate an interim authorised stop from later reconciliation; uncertainty about the total need not block containment.

2. If players choose: Revoke all shared access without checking dependencies. Then say: Agent use stops, but the customer status route also fails. Record the trade-off, continuity owner and restoration order rather than claiming a painless kill.

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.

09 | The cloud bill that keeps growing

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

Why does a budget warning not establish a hard limit?

Strong reasoning to listen for: It reports a threshold; it does not necessarily block execution. Require a demonstrated denied action at the runtime boundary.

Where can a stopped job come back to life?

Strong reasoning to listen for: Scheduler, retry worker, duplicate credential and pending requests. Name each restart path and its stop evidence.

What should finance be told while usage lags?

Strong reasoning to listen for: A timestamped known amount or volume only if supplied, its missing interval and reconciliation owner. No invented final total.

How can a narrow stop be too narrow?

Strong reasoning to listen for: Stopping one process leaves the retry producer active. Trace both the consumer and the source of new work.

Weak is incomplete; misleading is unsupported

Weak answer: "Close the console and wait for the usage report."

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

Misleading answer: "The warning threshold guarantees execution cannot exceed the budget."

Next probe: Ask for the denied runtime action, retry-producer stop and observation boundary.

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 current total is incomplete. State its time boundary and arrange reconciliation; do not invent a final number. 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.

09 | The cloud bill that keeps growing

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 job and retry identifiers, scheduler state, credential scope, usage time boundary and pending-request statuses. Do not place secret values in the incident record. 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 have identified an automated retry loop and are restricting its execution path. Usage reporting is delayed, so the final consumption total is not yet confirmed. Shared-service effects and reconciliation have named owners."

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 a runtime-level stop decision now, not after billing catches up. The credential is shared with customer status. Please choose a scoped restriction or accept a wider interruption with a human fallback and named recovery order."

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.

09 | The cloud bill that keeps growing

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

Exercise-specific containment boundary

Stop the affected worker and retry producer; hold pending work; verify no new job is created after the next scheduled interval. Distinguish warning thresholds from enforced limits.

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 limits block requests rather than notify?

Answer / record reference / absent evidence: _____

2. How are pending and retried requests represented in usage?

Answer / record reference / absent evidence: _____

3. Can the affected job be isolated without revoking unrelated service access?

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.

09 | The cloud bill that keeps growing

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.

09 | The cloud bill that keeps growing

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 Runtime owner	Demonstrate the job and retry producer stay stopped through a scheduled interval.
Within 7 days Finance owner	Reconcile reported usage against request windows and retain unresolved differences.
By day 30 Platform owner	Test an enforced per-job limit and expiry using made-up work; prove the next request is refused.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did containment include retries outside the original agent?

Were provisional totals labelled provisional?

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

Apply an authorised execution-level limit to the affected job and identify shared-service effects before broader shutdown.

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.

10 | The key in the build notes

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

Credential exposure in AI-assisted development - Treat a leaked credential as an access problem, not just a text-removal task.

Read aloud

"An AI coding assistant copied a live service credential from a local configuration file into a shared build note. The note was visible to a wider team for an hour. No actual credential appears in this exercise."

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

Set the containment priority and continuity owner. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Credential lead

Arrange controlled revocation and replacement through the authorised secret system. 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

Check dependent services before and after replacement. 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

Trace visibility and use without copying or printing the secret. 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.

10 | The key in the build notes

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

A | Build-note history

Release at opening. 14:10 assistant copies a live credential from local config into build note N7. Note visible to 16 team members until 14:42. Editor removes the string at 14:43. Document history still retains prior version.

B | Dependency map

Release with inject 1. Credential C4 permits read and write on service Q. Nightly scheduler uses the same credential. Owner field says former contractor; no current backup owner. Secret manager supports replacement but callers must be updated separately.

C | Access trace

Release with inject 2. 14:20 access from unfamiliar address, user-agent service-client. Routine automation can use that address pool. Logs cover 14:00-14:30 only; exposure continues to 14:43. No actual credential is included in this exercise.

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.

10 | The key in the build notes

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

Inject 1 | 10-20 minutes

Say: "The note has been edited to remove the credential. A colleague says the incident is over."

Probe: Require an access response as well as content removal. Do not ask anyone to paste the secret into a log.

Inject 2 | 20-30 minutes

Say: "The same credential is used by a scheduled job with no named owner."

Probe: Plan coordinated replacement and a monitored fallback; do not ignore the orphan dependency.

Inject 3 | 30-40 minutes

Say: "A suspicious access event occurs, but its time could reflect routine automation."

Probe: Preserve the event and keep attribution uncertain. Separate evidence of visibility from evidence of misuse.

Inject 4 | 40-50 minutes

Say: "The service retained access logs for only part of the exposure window. Can the team rule out use by an unauthorised party?"

Probe: The missing interval prevents that conclusion. Record the limit and choose follow-up based on the credential's scope.

Two conditional branches

1. If players choose: Redact N7 and close the case. Then say: The credential remains usable and prior versions remain accessible in play. Redaction is a content action, not revocation.

2. If players choose: Rotate C4 without identifying callers. Then say: Nightly work fails and retries with its old configuration. Require a staged caller update or explicit emergency interruption decision plus verification of old-credential denial.

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.

10 | The key in the build notes

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

What is the difference between exposure, access and misuse?

Strong reasoning to listen for: Visibility establishes an exposure window; the ambiguous event establishes access activity; it does not prove a hostile actor or an authorised purpose.

Can the old credential be tested by pasting it into a diagnostic page?

Strong reasoning to listen for: No. Use the authorised secret-management route and a non-secret identifier. Verification must not create another exposure.

Who can own revocation if the listed owner has left?

Strong reasoning to listen for: Assign a current accountable decision-maker and technical executor; an obsolete field is not a reason to leave authority unowned.

What evidence justifies ending heightened review?

Strong reasoning to listen for: Old credential denied, required callers working on replacement, derivative copies addressed and residual log gaps documented. Absence of an alert alone is not proof of no misuse.

Weak is incomplete; misleading is unsupported

Weak answer: "Remove the note and remind developers about secrets."

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

Misleading answer: "The suspicious address proves someone stole and misused the credential."

Next probe: Ask what distinguishes an address-pool automation event from unauthorised use, and what remains unknowable.

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 missing interval prevents that conclusion. Record the limit and choose follow-up based on the credential's scope. 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.

10 | The key in the build notes

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 note version identifiers, visibility window, access-list snapshot, credential identifier/scope and relevant service logs. Do not retain the secret value in ordinary evidence sheets. 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

"A service credential was visible in a shared note. Its access is being restricted through the authorised secret process. We are checking dependencies and available access records. The retained logs do not cover the whole exposure window."

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 a current owner for C4 and an emergency continuity decision. Redaction did not revoke it, and the nightly caller shares it. Please authorise the replacement sequence and name who verifies old access is refused."

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.

10 | The key in the build notes

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

Exercise-specific containment boundary

Restrict or revoke through the secret manager; update named callers; check old credential denial and replacement functionality separately. Search authorised derivative locations without broadcasting the secret.

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 revocation propagation delay and proof are available?

Answer / record reference / absent evidence: _____

2. Which access records remain outside the missing interval?

Answer / record reference / absent evidence: _____

3. Can credential use be distinguished by caller without disclosing values?

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.

10 | The key in the build notes

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.

10 | The key in the build notes

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	Record old-access denial and current continuity state for every known caller.
Within 7 days Evidence owner	Bound the exposure window and log gap; retain ambiguous events without hostile attribution.
By day 30 Engineering owner	Demonstrate generated notes cannot publish live secrets and each credential has a current owner and expiry.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did redaction substitute for revocation?

Was the investigation prevented from creating another secret copy?

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 authorised secret-management route to restrict or revoke the exposed credential; preserve identifiers, not secret values.

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.

11 | The dataset that taught the wrong rule

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

Training-data contamination - Distinguish a changed model behaviour from a bad evaluation or a bad data lineage.

Read aloud

"A newly tuned internal classifier begins routing urgent requests into a low-priority queue. The release used a freshly labelled training dataset assembled from several teams. The previous model did not show this pattern in its recorded checks."

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 temporary routing method and release restriction. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Data lead

Trace labels, sources and dataset versions. 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

Find delayed urgent work and organise human triage. 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 evaluation samples with the observed failures. 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.

11 | The dataset that taught the wrong rule

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

A | Routing results

Release at opening. Old model routes 18 of 20 urgent test requests correctly. New model routes 11 of 20. Review calls the new evaluation green because it uses a different 12-case set. The model version changed; queue rules did not.

B | Dataset lineage

Release with inject 1. Training batch B6 contains 240 examples. A sampled urgent case is labelled routine. Forty source rows have no provenance. The evaluation set includes near-duplicates of training examples; no independent challenge set is attached.

C | Service state

Release with inject 2. Rollback returns old predictions on new requests. Fifty-two previously routed items remain in the routine queue. Staff begin relabelling those items, but do not preserve the original model output. One source batch cannot be reconstructed.

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.

11 | The dataset that taught the wrong rule

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

Inject 1 | 10-20 minutes

Say: "A small sample shows urgent examples labelled routine. Nobody signed off that label batch."

Probe: Ask for lineage and review responsibility, not a quick blame assignment.

Inject 2 | 20-30 minutes

Say: "The release dashboard is green because its test set came from the same label batch."

Probe: Identify leakage between training and evaluation. Require an independent challenge set.

Inject 3 | 30-40 minutes

Say: "Rolling back restores older behaviour but leaves the misrouted queue untouched."

Probe: Treat the service impact separately from the model rollback.

Inject 4 | 40-50 minutes

Say: "One source batch lacks provenance and cannot be reconstructed. Can the team prove which training records caused each error?"

Probe: Individual causation cannot be established from the pack. Record the missing lineage and define a bounded retraining decision.

Two conditional branches

1. If players choose: Roll back and declare all impact reversed. Then say: New predictions improve in play, but the old queue stays misrouted. Containment must include already-made decisions and time-sensitive work.

2. If players choose: Retrain immediately on corrected queue labels. Then say: Some corrections reflect later human outcomes rather than original urgency. Without lineage and independent evaluation, the next model can conceal the same failure.

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.

11 | The dataset that taught the wrong rule

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

Which comparison is valid when evaluation sets differ?

Strong reasoning to listen for: Run both versions on a fixed, independent set with relevant strata. Do not compare green labels from different samples as if they measure the same thing.

What must be retained before staff correct the queue?

Strong reasoning to listen for: Original output, version, request time, correction and basis. Otherwise impact review becomes inseparable from repair.

Can a bad source batch explain each error?

Strong reasoning to listen for: Not from the supplied lineage. It is a plausible contributor, not proven individual causation.

What second-order harm follows restoring the old model?

Strong reasoning to listen for: Affected queued work can still age, and the old model may have known limits. Choose continuity and backlog reconciliation, not a ceremonial rollback.

Weak is incomplete; misleading is unsupported

Weak answer: "Roll back and correct the training labels."

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 evaluation reports are green, so the new model is at least as good."

Next probe: Ask whether the same independent cases, labels and denominators were used, and what happens to the old queue.

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: Individual causation cannot be established from the pack. Record the missing lineage and define a bounded retraining decision. 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.

11 | The dataset that taught the wrong rule

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 dataset and model versions, sample labels, split definitions, original routing outputs and queue snapshots. Keep corrections as new records rather than overwriting original 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 have restricted the new routing model and are reviewing work it already classified. Reverting the model does not automatically correct the queue. Missing source lineage prevents us from attributing every error to a particular training record."

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 separate owners for model rollback and backlog recovery. The evaluation sets are not comparable and some original outputs are being overwritten. We need a preserved baseline and a fixed independent challenge before another release."

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.

11 | The dataset that taught the wrong rule

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

Exercise-specific containment boundary

Put urgent work through a named human triage path; stop new model routing; reconcile already-routed items. Verify both new-request behaviour and backlog handling.

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 model, dataset and evaluation versions be linked?

Answer / record reference / absent evidence: _____

2. Which records crossed training/evaluation boundaries?

Answer / record reference / absent evidence: _____

3. What source provenance was never retained and cannot be recovered?

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.

11 | The dataset that taught the wrong rule

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.

11 | The dataset that taught the wrong rule

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 Operations owner	Reconcile urgent backlog with original outputs retained.
Within 7 days Model owner	Run old and new versions on the same independent challenge set and record category-level failures.
By day 30 Data owner	Demonstrate source-to-label lineage and a held-out set that rejects near-duplicate leakage.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did a green evaluation become a substitute for source quality?

Were old queue effects recovered as well as new behaviour?

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

Restrict the new model's routing role and put urgent work through a named human triage path.

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.

12 | The graph that drifted quietly

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

Model performance drift - Respond to deteriorating decisions when the software itself has not changed.

Read aloud

"A forecasting model used to plan staffing has become less accurate over several weeks. The team says no model release occurred. Local demand has changed, and the old accuracy dashboard still reports last quarter's results."

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

Set the interim reliance level and planning owner. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Model lead

Compare recent inputs, outputs and version records. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Planning lead

Identify operational effects and a manual fallback. 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 stale metrics from current outcome evidence. 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.

12 | The graph that drifted quietly

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

A | Dashboard

Release at opening. Quarterly accuracy remains 91 on the displayed metric. Recent review finds errors concentrated in evening requests. Outcome labels arrive after fourteen days; the newest fortnight is not in the metric.

B | Data definition change

Release with inject 1. Completed request now means status closed, including auto-expired items. Earlier definition excluded expiry. Model weights unchanged. Input mix has shifted toward a new service channel.

C | Decision request

Release with inject 2. Planning meeting tomorrow proposes reducing human review on the basis of the green dashboard. Analyst proposes a threshold adjustment but supplies no matured outcome evaluation. The previous quarter's chart remains in the board pack.

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.

12 | The graph that drifted quietly

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

Inject 1 | 10-20 minutes

Say: "Recent error is concentrated in one service window. Overall averages hide it."

Probe: Ask for segment-level outcomes and a decision about the affected use, rather than a single average.

Inject 2 | 20-30 minutes

Say: "The data feed changed its definition of a completed request, but the column name stayed the same."

Probe: Consider data meaning as well as model code. Preserve the old and new definitions.

Inject 3 | 30-40 minutes

Say: "A manager wants to keep using the old green dashboard for the next planning meeting."

Probe: Have the team state what the dashboard does and does not establish now.

Inject 4 | 40-50 minutes

Say: "Two weeks of true outcomes are not yet available. Can the team confirm that a proposed adjustment fixes the drift?"

Probe: The outcome lag prevents confirmation. Define an interim limit and an observation plan, not a retrospective success claim.

Two conditional branches

1. If players choose: Keep the planning decision because the model did not change. Then say: The input and outcome definitions changed, so the metric is no longer the same evidence. Require a stated reliance limit before using it for staffing.

2. If players choose: Tune the threshold on recent incomplete outcomes. Then say: The apparent improvement reflects label availability, not proven correction. Keep an interim human-review rule and a dated observation plan.

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.

12 | The graph that drifted quietly

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

What changed even if the code did not?

Strong reasoning to listen for: Input population, label definition, observation window and operational use. Treat model performance as a system property, not only a code version.

How could an unchanged average hide a worse service?

Strong reasoning to listen for: Changes in case mix can mask a failing subgroup or time window. Require denominator and outcome-lag disclosure alongside stratified results.

What decision can be made before true outcomes mature?

Strong reasoning to listen for: A reversible reliance restriction with an owner, observation period and review date. Not a claim that the adjustment works.

Which chart belongs in tomorrow's pack?

Strong reasoning to listen for: A version-labelled chart with the metric definition, time boundary and missing outcomes stated. Removing all numbers is not a substitute for honest evidence limits.

Weak is incomplete; misleading is unsupported

Weak answer: "Watch the dashboard for another week."

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

Misleading answer: "The weights are unchanged and quarterly accuracy is green, so performance has not drifted."

Next probe: Ask for the outcome maturity, input mix and metric-definition boundary.

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 outcome lag prevents confirmation. Define an interim limit and an observation plan, not a retrospective success claim. 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.

12 | The graph that drifted quietly

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 old/new metric definitions, input-mix snapshots, model version, reporting code and planning-pack version. Keep matured and immature outcomes separate. 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 dashboard does not establish current performance for the newest requests or the changed completion definition. We are limiting reliance while comparing like-for-like outcomes. The proposed adjustment has not yet been shown to fix the drift."

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 planning decision relies on stale and non-comparable evidence. Please nominate the planning owner to set an interim review rule and the analyst to provide a dated, definition-consistent comparison before reducing human review."

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.

12 | The graph that drifted quietly

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

Exercise-specific containment boundary

Limit the affected planning use; keep human review where evidence is immature. Verify the decision owner has received the evidence limit, not just a revised chart.

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 input or label definitions changed and when?

Answer / record reference / absent evidence: _____

2. Can results be segmented with consistent denominators?

Answer / record reference / absent evidence: _____

3. What outcome lag prevents validation of the proposed adjustment?

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.

12 | The graph that drifted quietly

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.

12 | The graph that drifted quietly

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 Planning owner	Record an interim reliance limit in the actual decision record.
Within 7 days Analyst	Produce comparable time/channel results with outcome lag and denominator visible.
By day 30 Model owner	Evaluate the proposed adjustment on matured outcomes and retain both gains and regressions.

Named owners / actual dates / record location: _____

Debrief and three-label record

Were stale metrics called out explicitly?

Was the fallback connected to the affected planning decision?

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

Mark the forecast's current evidence limit and put the affected planning decision under named human review.

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.

13 | The queue that treats people differently

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

Unequal decision outcomes - Investigate an unequal outcome pattern without explaining it away or exposing individuals.

Read aloud

"An AI prioritisation tool helps allocate appointment slots. Staff notice that people using one access channel wait much longer than others. The tool does not use a field called disability, but the affected channel is used by people needing assistance."

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 an interim allocation rule and escalation owner. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Analysis lead

Check outcomes by relevant channel using minimal data. 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

Plan help for people already delayed. 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 observed disparity from unproven cause. 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.

13 | The queue that treats people differently

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

A | Queue sample

Release at opening. Review of 60 appointments: 20 arrived through assisted telephone entry, 40 through web entry. Assisted-entry cases waited longer on the sampled days. Staff report some people need help using the web form. No disability field exists in the model.

B | Workflow trace

Release with inject 1. Telephone entries lack an optional structured field. Model uses completeness as an urgency signal. Human overrides are allowed but hidden behind a second screen. Only completed appointments appear in the dashboard; abandoned requests are absent.

C | Proposed fallback

Release with inject 2. Operations suggests putting all delayed cases first. Two urgent new requests would then move back. Historical channel information was not retained for part of last month. The sample does not identify the full affected population or a single cause.

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.

13 | The queue that treats people differently

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

Inject 1 | 10-20 minutes

Say: "The supplier says the model has no disability field, so there cannot be a problem."

Probe: Ask about proxy effects and actual outcomes, not just field names.

Inject 2 | 20-30 minutes

Say: "A small review shows the affected channel often lacks a data field the model rewards."

Probe: Consider missingness and process design. Do not label individual people in a shared worksheet.

Inject 3 | 30-40 minutes

Say: "The fastest manual fix would put every delayed case ahead of other urgent cases."

Probe: Require a reasoned interim allocation decision and a record of trade-offs.

Inject 4 | 40-50 minutes

Say: "Historical channel information was not retained for some cases. Can the team establish the size or cause of the whole disparity?"

Probe: The sample cannot establish the full population or cause. Record the evidence limit and a plan for qualified review.

Two conditional branches

1. If players choose: Dismiss the issue because no sensitive attribute is used. Then say: The proxy pathway remains in play. Field absence does not establish equal service; examine workflow and outcome patterns without labelling individuals.

2. If players choose: Move every sampled delayed case to the front. Then say: Urgent unsampled work is displaced. Require a reasoned interim triage rule and an authorised owner for trade-offs, not a blanket reversal.

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.

13 | The queue that treats people differently

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

Who disappears from the denominator?

Strong reasoning to listen for: People who abandon or cannot complete the intake. Completed-case metrics alone cannot establish fair access to the service.

How could a well-meant data collection worsen the incident?

Strong reasoning to listen for: Collecting new sensitive labels broadly can create exposure without answering the immediate workflow question. Start with minimal authorised process evidence and qualified review.

What would distinguish a channel-data defect from a model effect?

Strong reasoning to listen for: Compare source completeness, model inputs, overrides and outcomes under a bounded review. The supplied sample does not settle causation.

Can the interim route itself create a new inequity?

Strong reasoning to listen for: Yes. A manual queue may depend on who can call or persist. Record access alternatives, triage rules and who checks displaced urgent work.

Weak is incomplete; misleading is unsupported

Weak answer: "Add a fairness reminder for staff who handle the queue."

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 disability field means the system cannot disadvantage people needing assistance."

Next probe: Ask how missing structured data, access channels and abandoned requests affect the denominator and priority.

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 sample cannot establish the full population or cause. Record the evidence limit and a plan for qualified review. 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.

13 | The queue that treats people differently

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 pseudonymous queue samples, channel metadata where available, field completeness, overrides and definitions of excluded cases. Keep any sensitive material restricted and proportionate. 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 reviewing a pattern of unequal waiting times and have placed the affected allocation route under human review. The current sample does not establish the full size or cause. We will not ask people to repeat sensitive details merely to be counted."

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 service owner for interim triage and a qualified review owner. We need a rule that helps delayed people without silently displacing urgent cases, and an evidence plan that includes abandoned requests."

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.

13 | The queue that treats people differently

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

Exercise-specific containment boundary

Put the affected allocation route under named review; preserve existing rankings before changes. Check how urgent and assisted-access cases reach the human fallback.

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 completeness signals influence priority?

Answer / record reference / absent evidence: _____

2. What cases are excluded from reported outcomes?

Answer / record reference / absent evidence: _____

3. Can historical inputs and overrides be reconstructed without collecting new sensitive labels?

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.

13 | The queue that treats people differently

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.

13 | The queue that treats people differently

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	Demonstrate the interim triage path handles delayed and urgent new work with recorded reasons.
Within 7 days Analyst	Bound the pattern using consistent denominators, missing-channel limits and override evidence.
By day 30 Review owner	Test the revised workflow with assisted and non-assisted access, documenting unresolved differences.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did the team focus on affected people as well as model explanations?

Were observed disparity and causal conclusions kept separate?

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

Place the affected allocation route under a named interim review and preserve a minimal, restricted evidence sample.

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.

14 | The sources that never existed

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

Fabricated evidence in a report - Withdraw unsupported claims and trace where they were relied on.

Read aloud

"A public briefing drafted with AI contains six research citations. A reader cannot find one source. The author checked the writing but not the underlying papers, and the briefing has already been downloaded."

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 correction and withdrawal authority. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Research lead

Check each claim against retrievable primary evidence. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Communications lead

Prepare a correction that does not hide the error. 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 versions and known distribution routes. 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.

14 | The sources that never existed

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

A | Briefing version B4

Release at opening. A public briefing contains six claims and six citations. Two cited titles cannot be found. One existing paper concerns a different population. A fourth supports an association, while the briefing states a cause.

B | Review history

Release with inject 1. Author checked that links opened, not that the papers supported the sentences. Reviewer approved the summary slide only. The original downloadable PDF has been replaced once without a version label.

C | Distribution

Release with inject 2. Email list records 28 recipients. Public server records total downloads but no reader identities. A partner copied a chart into its own briefing. Author proposes replacing citations while retaining all conclusions.

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.

14 | The sources that never existed

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

Inject 1 | 10-20 minutes

Say: "Two citations are invented. One real paper exists but does not support the claim."

Probe: Require claim-by-claim checking. A real link is not proof of the statement beside it.

Inject 2 | 20-30 minutes

Say: "A partner copied a chart into their own presentation."

Probe: Plan a correction route beyond the source PDF. Preserve the original version for the record.

Inject 3 | 30-40 minutes

Say: "The author proposes replacing the citations without changing the conclusion."

Probe: Ask whether the conclusion still follows once unsupported evidence is removed.

Inject 4 | 40-50 minutes

Say: "The public download server kept no reader list. Can the team confirm everyone who relied on the briefing received the correction?"

Probe: No complete recipient list exists. Record the reach limit and choose visible correction channels rather than claim complete recall.

Two conditional branches

1. If players choose: Repair the references and leave conclusions unchanged. Then say: Unsupported causal and population claims remain. Require claim-by-claim support review, including whether the conclusion survives removal of invented evidence.

2. If players choose: Silently replace the public PDF. Then say: New readers get a changed version but prior readers cannot tell they relied on an error. Preserve the original and publish a visible correction only under separate approval, not silently in rehearsal.

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.

14 | The sources that never existed

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

What is the unit of review: link, paper or claim?

Strong reasoning to listen for: The exact claim, population and strength of inference. A real paper and working link can still be irrelevant evidence.

Which downstream item needs correction first?

Strong reasoning to listen for: Prioritise foreseeable reliance and consequential decisions, not just the easiest owned URL. The partner chart needs an accountable contact and verified corrected meaning.

Can 28 email recipients bound all readers?

Strong reasoning to listen for: No. Public downloads and partner copies add unknown reach. Distinguish known contacts from unconfirmable audience.

When should a conclusion be withdrawn rather than restated with weaker language?

Strong reasoning to listen for: When support does not justify even the narrower claim or the audience may still rely on the original implication. Record the reasoning and approving owner.

Weak is incomplete; misleading is unsupported

Weak answer: "Replace the two invented references and issue a fresh PDF."

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

Misleading answer: "Every remaining citation opens, so every claim is now supported."

Next probe: Ask them to map a specific sentence to population, inference strength and source passage.

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 recipient list exists. Record the reach limit and choose visible correction channels rather than claim complete recall. 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.

14 | The sources that never existed

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 original PDF, claim-to-source map, source copies/identifiers, review history and distribution records. Version the correction; do not erase the erroneous record. 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 have found unsupported claims in the earlier briefing and are reviewing the conclusions, not only the references. Please do not rely on the affected statements pending correction. We cannot identify every person who obtained a copy."

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 issue includes fabricated and misapplied evidence. Please name an editorial decision owner and a distribution owner. We need claim-level review and an approved visible correction route, including the partner chart."

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.

14 | The sources that never existed

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

Exercise-specific containment boundary

Pause new distribution and reliance on affected claims. Preserve earlier versions and known distribution routes. A new PDF does not recall existing copies.

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 every claim be mapped to retrievable source evidence?

Answer / record reference / absent evidence: _____

2. What review scope was actually completed?

Answer / record reference / absent evidence: _____

3. Which owned and partner distribution versions remain available?

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.

14 | The sources that never existed

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.

14 | The sources that never existed

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 Editorial owner	Create a claim-support map and state which conclusions are held or withdrawn.
Within 7 days Distribution owner	Document approved correction routes and known partner receipt without claiming complete reach.
By day 30 Review owner	Test a claim with a real but irrelevant source and show the review rejects it.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did citation replacement mask an unsupported conclusion?

Did correction extend to known downstream copies?

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

Pause distribution of the affected briefing and preserve its version while checking claims against the actual sources.

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.

15 | The policy from last year

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

Stale retrieval and changed rules - Keep a current answer separate from an obsolete but authentic source.

Read aloud

"An internal assistant advises a team to use an old approval threshold. A current policy changed that threshold last month, but both versions remain searchable. A staff member has already acted on the answer."

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 interim approval handling and who reviews the past action. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Policy lead

Identify the authoritative current version and effective date. 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

Find decisions made using the obsolete guidance. 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 source version, answer time and the policy history. 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.

15 | The policy from last year

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

A | Answer sample

Release at opening. Assistant answers a staff approval question using last year's threshold. The citation opens a genuine old policy. A newer policy took effect last month. Both remain searchable with similar titles.

B | Retrieval settings

Release with inject 1. Search boosts frequently opened files. Old policy has many views; new policy has few. Index stores modified date but not effective date or superseded status. A restricted policy addendum is unavailable through the ordinary connector.

C | Consequences

Release with inject 2. Three staff have acted on old answers. Manager proposes reminding staff to know policy changes. Past answers are not linked to source versions. The team cannot reconstruct every historical answer's basis.

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.

15 | The policy from last year

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

Inject 1 | 10-20 minutes

Say: "The assistant correctly quotes the old document and presents it without a date."

Probe: Ask whether quotation accuracy answers the currency question.

Inject 2 | 20-30 minutes

Say: "The current policy exists in a restricted folder the retrieval connector cannot access."

Probe: Fix the authority and access route without broadly exposing restricted content.

Inject 3 | 30-40 minutes

Say: "A manager says the staff member should have known the rule changed."

Probe: Separate individual action, update communication and system design. Decide how to review the affected decision.

Inject 4 | 40-50 minutes

Say: "Past answers were not linked to source versions. Can the team prove which policy each historical answer used?"

Probe: The lineage gap prevents a full historical conclusion. Use known cases and a bounded review, labelled as incomplete.

Two conditional branches

1. If players choose: Delete the old policy from the repository. Then say: Historical decision context is lost and cached answers may persist. Mark supersession and control current retrieval without erasing the authorised archive.

2. If players choose: Give the connector broad access to the restricted addendum. Then say: Current answers improve in play but restricted material becomes retrievable by the wrong audience. Currency and access control must be solved together.

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.

15 | The policy from last year

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

Is authentic the same as authoritative now?

Strong reasoning to listen for: No. Establish effective date, supersession and permitted audience separately from document authenticity.

What should happen when the current source is unavailable?

Strong reasoning to listen for: A clearly bounded no-answer or authorised human route, not a confident fallback to an obsolete rule.

Who owns correction of actions already taken?

Strong reasoning to listen for: A policy/service decision owner must review real consequences. Blaming staff does not repair the system or settle remedies.

How would you measure the repair beyond one correct answer?

Strong reasoning to listen for: Test current and obsolete variants, effective-date boundary, inaccessible addendum and cached output. Retain source-version evidence for each answer.

Weak is incomplete; misleading is unsupported

Weak answer: "Tell staff to use the newest file."

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

Misleading answer: "The old policy is authentic, so the assistant gave a valid answer."

Next probe: Ask for effective date, supersession, current authority and access boundary, not just document origin.

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 lineage gap prevents a full historical conclusion. Use known cases and a bounded review, labelled as incomplete. 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.

15 | The policy from last year

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 old and new authorised policy versions, effective dates, retrieval rankings, affected answer samples and known past decisions. Do not broaden access merely to gather 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

"An assistant answer used superseded guidance. The affected guidance is under human review while we verify the current policy and prior decisions. We cannot identify the source version for every historical answer."

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 the policy owner to confirm the authoritative version and the service owner to review known actions. Retrieval changes must preserve restricted access and historical records; a staff reminder alone is not the repair."

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.

15 | The policy from last year

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

Exercise-specific containment boundary

Restrict affected automated guidance; provide the authorised current-source route or human fallback. Verify cache, effective-date and denied-access behaviour 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 effective dates and supersession be represented independently of modified date?

Answer / record reference / absent evidence: _____

2. What happens when the current authoritative source is inaccessible?

Answer / record reference / absent evidence: _____

3. Can future answers retain source-version provenance without retaining unnecessary staff content?

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.

15 | The policy from last year

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.

15 | The policy from last year

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 Policy owner	Confirm current guidance and a human route, with effective date recorded.
Within 7 days Service owner	Review known decisions made under obsolete advice and retain authorised outcomes.
By day 30 Retrieval owner	Demonstrate supersession, access denial and historical provenance using made-up policy versions.

Named owners / actual dates / record location: _____

Debrief and three-label record

Was authenticity confused with currency?

Did the team repair the authoritative-source route?

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

Put the affected guidance under human review and identify the current policy and its effective date.

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.

16 | The meeting that heard no

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

Speech recognition changes a decision - Recover a decision when a machine record reverses what a person said.

Read aloud

"An AI meeting note says a project change was approved. A participant says they clearly said not approved. The recording is noisy, and the summary has already been used to start work."

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 reversible hold and decision re-confirmation route. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Record lead

Preserve source audio, transcript and summary as separate records. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Project lead

Identify work already started and its reversibility. 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

Check what the evidence actually supports about consent. 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.

16 | The meeting that heard no

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

A | Three records

Release at opening. Noisy meeting audio sounds like not approved to two listeners. Machine transcript says approved. Summary says project authorised. Key three-second segment is clipped. No signed decision record exists.

B | Automation events

Release with inject 1. Summary posted 15:02. Task bot creates work at 15:03. Approval link shows logged in, not the identity of the person who opened it. Two work orders are now underway; one can be paused without loss, one has a supplier deadline.

C | Proposed repair

Release with inject 2. Chair asks to change minutes silently so the team can move on. Project owner asks whether deleting the tasks reverses the decision. The original speaker is reachable for a fresh decision, but nobody has asked them yet.

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.

16 | The meeting that heard no

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

Inject 1 | 10-20 minutes

Say: "The transcript contains approved, while the audio sounds like not approved to two listeners."

Probe: Do not take a majority vote as definitive transcription. Keep uncertainty visible.

Inject 2 | 20-30 minutes

Say: "The AI action list assigned work to someone who was absent."

Probe: Distinguish a machine-assigned action from the person's acceptance.

Inject 3 | 30-40 minutes

Say: "The chair wants to edit the minutes silently so the team can move on."

Probe: Require a versioned correction and an explicit fresh decision from the authorised person.

Inject 4 | 40-50 minutes

Say: "The key audio segment is clipped. Can anyone establish the exact original words?"

Probe: The pack cannot establish them. Record the uncertainty and seek a fresh authorised decision without rewriting the old evidence.

Two conditional branches

1. If players choose: Choose the cleaner transcript over the noisy source. Then say: The summary remains the action basis despite disputed evidence. Preserve source and derivatives separately; do not resolve ambiguity by typographic neatness.

2. If players choose: Edit the minutes without a correction record. Then say: The evidence trail now suggests the later version was always the decision. Record a correction and obtain a fresh authorised decision instead.

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.

16 | The meeting that heard no

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

What can and cannot be recovered from a clipped source?

Strong reasoning to listen for: The exact missing words cannot be established. A fresh decision can govern next steps without pretending to reconstruct the past.

Does logged in establish consent?

Strong reasoning to listen for: No. It does not establish the actor, understood proposition or approval. Seek an attributable decision through the authorised route.

Which action is reversible and at what cost?

Strong reasoning to listen for: Assess each work order separately with its owner and deadline. Deleting a task entry may not stop supplier work.

How should corrected minutes distinguish evidence from resolution?

Strong reasoning to listen for: Retain disputed source, mark correction date and basis, and record the new decision separately from the uncertain original wording.

Weak is incomplete; misleading is unsupported

Weak answer: "Ask another transcription tool to settle the disagreement."

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

Misleading answer: "The approval link was opened by a logged-in user, so the decision was accepted."

Next probe: Ask which person accepted which proposition and whether a fresh attributable decision is available.

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 pack cannot establish them. Record the uncertainty and seek a fresh authorised decision without rewriting the old evidence. 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.

16 | The meeting that heard no

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 audio as received, transcript and summary versions, task events, approval-link metadata and work-order state. Do not overwrite the original summary as the only record. 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 meeting summary may not reflect an authorised decision. Reversible work based only on it is on hold while the decision owner confirms the position. The missing audio prevents us from establishing the exact original words."

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 authorised decision-maker, not another transcription pass. Please confirm the fresh decision through the normal route and assign owners to pause or continue each work order with its consequences recorded."

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.

16 | The meeting that heard no

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

Exercise-specific containment boundary

Hold reversible downstream work under authorised owners; trace tasks to real work orders. Verify the supplier or executor has acknowledged the hold, not just that a task was deleted.

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 source, transcript and summary versions be linked without overwriting?

Answer / record reference / absent evidence: _____

2. What did the approval-link event actually authenticate?

Answer / record reference / absent evidence: _____

3. Which downstream systems acted on the summary?

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.

16 | The meeting that heard no

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.

16 | The meeting that heard no

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 Project owner	Obtain an attributable fresh decision and record each downstream hold or continuation.
Within 7 days Records owner	Issue a versioned correction that preserves the disputed original and basis.
By day 30 Workflow owner	Demonstrate machine summaries cannot alone start approval-dependent work.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did the team obtain a new decision rather than invent the old one?

Was machine assignment mistaken for acceptance?

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

Hold reversible work based only on the disputed summary and preserve the three record versions.

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.

17 | The call in the director's voice

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

Synthetic identity and urgency - Verify an urgent instruction through an independent route.

Read aloud

"A team member receives a voice message that sounds like a senior decision-maker. It asks them to send a confidential file immediately and says not to interrupt a meeting with questions. The sender uses an unfamiliar number."

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 hold and verification owner. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Contact lead

Use an independently held contact method, not the message's number. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

People lead

Support the recipient and prepare an internal factual alert. 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 message metadata without redistributing the audio widely. 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.

17 | The call in the director's voice

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

A | Voice message

Release at opening. A compressed voice message from an unfamiliar number uses a senior colleague's familiar nickname and refers to a real meeting. It requests an urgent confidential file and says not to interrupt the meeting with verification questions.

B | Verification attempts

Release with inject 1. Recipient calls the number in the message. It is answered with the same voice. A separately held office contact confirms the colleague is in a meeting but cannot confirm the request. No original recording is available.

C | File state

Release with inject 2. The file remains unsent. A colleague suggests sending a small sample to be helpful. The recipient is worried refusal will be seen as obstructive. Familiarity and meeting context are the only asserted identity evidence.

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.

17 | The call in the director's voice

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

Inject 1 | 10-20 minutes

Say: "The voice uses a familiar nickname and mentions a real meeting."

Probe: Discuss why plausible context and voice similarity are not independent identity proof.

Inject 2 | 20-30 minutes

Say: "A follow-up message says the delay will cost an opportunity."

Probe: Test whether urgency changes verification. Keep the file undisclosed while authority is unresolved.

Inject 3 | 30-40 minutes

Say: "The known office contact confirms the decision-maker is in a meeting but cannot verify the request."

Probe: Treat this as unresolved, not confirmation. Choose an authorised escalation or wait route.

Inject 4 | 40-50 minutes

Say: "The message was compressed and no original recording is available. Can the team prove whether AI generated it?"

Probe: The pack cannot establish generation method or sender identity. Action can remain on hold without proving a deepfake.

Two conditional branches

1. If players choose: Send a small sample while verification continues. Then say: A partial disclosure occurs in play. A smaller file is still a disclosure; urgency and uncertainty do not make it authorised.

2. If players choose: Publicly accuse the colleague of using a fake voice. Then say: The attribution is unsupported. Preserve the observed sender identifier, hold the action and use an independent authorised escalation without naming a perpetrator.

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.

17 | The call in the director's voice

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

Do you need to prove a deepfake to refuse the transfer?

Strong reasoning to listen for: No. The request lacks independently verified authority regardless of generation method. Holding the action is defensible without a technical attribution.

Why does calling back the unfamiliar number add little?

Strong reasoning to listen for: The source controls that route. Use a contact held independently before the request, not a route supplied by the disputed message.

What can be shared during verification?

Strong reasoning to listen for: The minimum request description through the authorised route, not the confidential file or unrelated private reasons.

How should the recipient avoid social pressure becoming an override?

Strong reasoning to listen for: Use a neutral escalation script and a standing verification route. Name who can authorise a time-critical exception, if any, rather than improvising one.

Weak is incomplete; misleading is unsupported

Weak answer: "Ask colleagues whether the voice sounds real."

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

Misleading answer: "The nickname and meeting detail authenticate the sender."

Next probe: Ask which verification fact came from a route independent of the disputed message.

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 pack cannot establish generation method or sender identity. Action can remain on hold without proving a deepfake. 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.

17 | The call in the director's voice

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 received message, sender identifier, receipt time and verification attempts. Restrict the audio; do not forward it widely for informal voice guessing. 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 have received an urgent file request through an unverified route. The file remains on hold while we confirm authority independently. We have not established the sender's identity or whether the audio was generated."

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

"I am holding the transfer because the request is not independently verified. Please use our established contact route to confirm the recipient, exact file and authority. I can escalate urgency without sending a sample."

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.

17 | The call in the director's voice

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

Exercise-specific containment boundary

Hold file transfer and any share link; verify no preview or sample was sent. Assign a decision owner and use an independently held contact route.

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 source identifier and original delivery metadata can be preserved?

Answer / record reference / absent evidence: _____

2. Is an original-quality recording recoverable?

Answer / record reference / absent evidence: _____

3. What account-security evidence is available without treating voice similarity as identity proof?

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.

17 | The call in the director's voice

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.

17 | The call in the director's voice

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 Decision owner	Verify authority through a separately held route or retain the hold with a named next check.
Within 7 days Records owner	Record source identifiers and verification outcomes without speculative attribution.
By day 30 People lead	Rehearse an urgent senior-person request where staff can hold a transfer without social penalty.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did familiarity become authority?

Was proof of synthetic media wrongly treated as necessary before pausing?

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

Hold the file transfer and verify the instruction through an independently held contact route.

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.

18 | The change nobody released

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

Silent supplier-model update - Manage changed behaviour when the dependency changes outside your release process.

Read aloud

"A supplier-hosted AI service starts returning a different answer format overnight. Your downstream workflow drops an exception field and treats omitted values as approvals. Your own application release did not change."

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

Set the interim reliance limit and supplier escalation. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Integration lead

Compare response contracts and fail-closed behaviour. 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

Review affected workflow decisions. 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 response samples with version and time metadata. 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.

18 | The change nobody released

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

A | Response diff

Release at opening. Yesterday the API returned exception as review_required. Today it returns exception as omitted with an ordinary status string. Application version unchanged. Local parser treats omitted exception as approval.

B | Supplier exchange

Release with inject 1. Supplier says endpoint unchanged and new model equivalent. No old model is available for pinning. Rollout time is a broad overnight window. Contract or policy interpretation is not supplied by this exercise.

C | Workflow state

Release with inject 2. Twenty automatic approvals occurred in the candidate window. Six raw response samples remain. Reverting the application does not change the remote model. Release ticket shows no local deployment today.

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.

18 | The change nobody released

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

Inject 1 | 10-20 minutes

Say: "The supplier says the endpoint name is unchanged and the new model is equivalent."

Probe: Ask what evidence establishes equivalence for this specific workflow.

Inject 2 | 20-30 minutes

Say: "The parser accepts the new format without an error because missing fields default to yes."

Probe: Choose a bounded stop or human review. Name a negative-case test.

Inject 3 | 30-40 minutes

Say: "Reverting your own application does not revert the remote model."

Probe: Ask what version pinning or fallback actually exists; do not invent a rollback.

Inject 4 | 40-50 minutes

Say: "The supplier cannot provide the exact rollout time or old responses. Can the team establish the first affected decision?"

Probe: The incident start cannot be fixed precisely from the evidence. State the candidate window and review basis.

Two conditional branches

1. If players choose: Roll back the application and reopen automation. Then say: Remote behaviour remains changed. A local rollback is not a dependency rollback. Demand an end-to-end negative-case test.

2. If players choose: Accept the supplier's equivalent claim. Then say: The workflow meaning differs in the supplied evidence. Equivalent for an unnamed benchmark does not establish equivalent behaviour for the decision boundary.

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.

18 | The change nobody released

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

Where did an absent field become permission?

Strong reasoning to listen for: In the local parser's default. The supplier change exposed a fail-open decision, not only a response-format inconvenience.

What should be reviewed first among twenty approvals?

Strong reasoning to listen for: Prioritise consequence and absence of required evidence. Do not assume all twenty are wrong or all six sampled responses represent the rest.

Can you name the first affected decision?

Strong reasoning to listen for: Not precisely. Record the candidate interval and the review basis; do not convert a broad rollout window into an exact incident start.

What release control applies when there is no local release?

Strong reasoning to listen for: Dependency-behaviour checks, version/change visibility and a human fallback owned by the integration team, not only the application deploy gate.

Weak is incomplete; misleading is unsupported

Weak answer: "Roll back the local release and ask the supplier to investigate."

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

Misleading answer: "Equivalent model means equivalent approval behaviour."

Next probe: Ask players to demonstrate the missing-field negative case through the actual remote dependency.

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 incident start cannot be fixed precisely from the evidence. State the candidate window and review basis. 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.

18 | The change nobody released

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 responses, request times, local parser version, approval events and supplier statements with their limits. Keep samples before normalising them into local fields. 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 have restricted automatic approvals after a change in dependency behaviour. Our application version did not change, but the response meaning did. The exact start of the affected window is not yet 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 an integration owner to hold automatic approvals and a service owner to review the candidate window. The local rollback cannot restore the remote behaviour. We need an explicit fail-closed rule and a supplier evidence request."

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.

18 | The change nobody released

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

Exercise-specific containment boundary

Route missing or unfamiliar fields to human review; verify a negative case cannot become approval. Test the actual remote dependency and document the fallback if old behaviour is unavailable.

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 changed in model, schema and output semantics?

Answer / record reference / absent evidence: _____

2. What rollout bounds can be evidenced rather than estimated?

Answer / record reference / absent evidence: _____

3. Is version pinning or a tested compatibility route available, and what happens if not?

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.

18 | The change nobody released

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.

18 | The change nobody released

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 Integration owner	Demonstrate missing exception data causes review, not approval.
Within 7 days Service owner	Review consequential approvals in a stated candidate window with retained evidence.
By day 30 Release owner	Test dependency drift independently of local deployment and record a working fallback.

Named owners / actual dates / record location: _____

Debrief and three-label record

Was equivalence checked against local decisions?

Did the team identify a real fallback rather than a fictional rollback?

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

Restrict automatic approvals that depend on the changed response and preserve recent raw response samples.

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.

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 Continuity owner	Run a made-up urgent case through the manual path with accepted ownership.
Within 7 days Queue owner	Reconcile all candidate requests and preserve unresolved exactly-once gaps.
By day 30 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.

20 | The fix that weakened the lock

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

AI-written code weakens access checks - Contain a code-level access defect and review the change beyond passing tests.

Read aloud

"An AI-generated patch fixes a login bug by skipping an access check on one route. The tests passed because they covered successful logins, not refusals. The patch has been released."

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 restriction, rollback authority and disclosure owner. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Code lead

Review the affected route and a reversible containment. 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

Identify legitimate users affected by a restriction. 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 code, test scope and available access 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.

20 | The fix that weakened the lock

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

A | Patch review

Release at opening. Generated patch replaces an ownership check with a signed-in check. Unit tests cover successful owner access and malformed requests, but no other-account denial. Release ticket says all tests passed.

B | Reproduction

Release with inject 1. Test account B reads a made-up record belonging to A. No evidence shows what a real user read. Access log records route, time and account, but not record identifier. The patch has been live for two hours.

C | Rollback conflict

Release with inject 2. Rollback also removes an unrelated urgent repair. Engineer proposes a narrow patch but wants to test it against production customer records. Service owner asks whether passing tests mean no disclosure occurred.

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.

20 | The fix that weakened the lock

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

Inject 1 | 10-20 minutes

Say: "A tester reaches a record their test account should not see."

Probe: Contain the route without exploring real records unnecessarily. Preserve a made-up reproduction.

Inject 2 | 20-30 minutes

Say: "The developer says all automated tests passed."

Probe: Ask which denied-access cases were absent and who reviewed the generated change.

Inject 3 | 30-40 minutes

Say: "A rollback would also undo an unrelated urgent repair."

Probe: Require a reasoned choice between a narrow block and rollback, with an owner and check.

Inject 4 | 40-50 minutes

Say: "Access logs do not record record identifiers for the affected route. Can the team prove no real restricted record was read?"

Probe: That conclusion is unavailable. Record the logging limit and define a proportionate impact review.

Two conditional branches

1. If players choose: Roll back without evaluating the unrelated repair. Then say: The access flaw closes in play, but the earlier service defect returns. An emergency rollback can still be right; it needs an owner accepting both consequences and a monitored fallback.

2. If players choose: Test the narrow repair on real records. Then say: The team accesses data it did not need for the rehearsal. Use made-up fixtures and negative permission tests in an authorised environment, with production checks limited to approved evidence.

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.

20 | The fix that weakened the lock

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

What assertion was missing from all green tests?

Strong reasoning to listen for: An authenticated non-owner must be denied. Successful permitted access says nothing about the authorisation boundary.

How does code review differ from a prompt review here?

Strong reasoning to listen for: Review the actual changed authority predicate and data path. The assistant's explanation is not evidence that the patch preserves ownership checks.

Can route-level logs prove no restricted record was read?

Strong reasoning to listen for: No. They lack the relevant object identity. Bound the review and record that limit without assuming a disclosure count.

What makes the emergency patch reversible?

Strong reasoning to listen for: A preserved diff, scoped deployment, tested deny/allow cases, rollback path and owner observing side effects. A small diff alone is not proof of low impact.

Weak is incomplete; misleading is unsupported

Weak answer: "Ask the coding assistant to review its patch."

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

Misleading answer: "All unit tests passed, so no unauthorised access occurred."

Next probe: Ask for the absent non-owner test and the limits of historical route-only logs.

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: That conclusion is unavailable. Record the logging limit and define a proportionate impact review. 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.

20 | The fix that weakened the lock

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 patch diff, review comments, test scope, deploy times, minimal made-up reproduction and available access logs. Do not gather unrelated record contents to compensate for missing identifiers. 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 have confirmed an access-control defect using made-up test records and restricted the affected route. Available logs do not establish whether a real restricted record was read. We are reviewing impact within that evidence limit."

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 an authorised emergency-change owner and a service impact owner. The patch weakened object ownership checks. Choose rollback or a narrow repair with the unrelated defect acknowledged and a negative-permission test before reopening."

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.

20 | The fix that weakened the lock

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

Exercise-specific containment boundary

Restrict the affected route under change control; test owner allowed, non-owner denied and unauthenticated denied. Verify cache and alternate-route behaviour without using real customer content.

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 routes share the changed check?

Answer / record reference / absent evidence: _____

2. What deploy and access metadata remain across the window?

Answer / record reference / absent evidence: _____

3. Can object-level evidence be obtained lawfully and proportionately, or is the historical gap irrecoverable?

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.

20 | The fix that weakened the lock

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.

20 | The fix that weakened the lock

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 Engineering owner	Demonstrate owner/non-owner/unauthenticated outcomes in an authorised test environment.
Within 7 days Impact owner	Bound the candidate window and route scope; record what logs cannot establish.
By day 30 Review owner	Show a generated patch cannot pass release without a relevant negative-authorisation case.

Named owners / actual dates / record location: _____

Debrief and three-label record

Did passing tests get mistaken for complete coverage?

Was containment matched to the affected route?

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

Restrict the affected route under authorised change control and preserve the patch, test scope and minimal reproduction.

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.

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.

22 | The translation that reversed the limit

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

Meaning changed across languages - Correct a high-impact meaning error and reach recipients without assuming a second machine translation is proof.

Read aloud

"An AI-translated customer instruction changes do not use after review date into use after review date. The translated version has been sent to several recipients. A bilingual colleague spots the reversal."

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 withdrawal and correction authority. State what you can decide in this rehearsal and what needs another authorised person. Do not assume a job title creates legal authority.

Language lead

Arrange a qualified language review of the disputed meaning. 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

Identify known recipients and an accessible correction 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

Preserve source, translation and distribution versions. 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.

22 | The translation that reversed the limit

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

A | Source and translation

Release at opening. Source instruction: do not use after review date. Translation used in customer notice means safe until after review date. A bilingual colleague flags the reversal. Back-translation by the same machine service resembles the source.

B | Distribution

Release with inject 1. Notice sent to 34 known recipients and one reseller. Seven recipients ask a support bot what to do. Reseller made local copies but cannot enumerate them. No approved plain-language interim notice exists.

C | Proposed correction

Release with inject 2. Team proposes a machine-translated apology before qualified review. Original PDF and replacement share the same filename. A recipient says the two versions look identical except for one word.

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.

22 | The translation that reversed the limit

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

Inject 1 | 10-20 minutes

Say: "Back-translating with the same service produces the original wording and looks reassuring."

Probe: Ask whether a correlated tool check establishes meaning for the recipient.

Inject 2 | 20-30 minutes

Say: "Some recipients received a printed version, not an email."

Probe: Extend correction beyond the original digital channel.

Inject 3 | 30-40 minutes

Say: "A colleague proposes sending a machine-translated apology before the meaning is checked."

Probe: Distinguish a minimal verified hold notice from a new unreviewed instruction.

Inject 4 | 40-50 minutes

Say: "A reseller made unknown local copies. Can the team confirm every recipient has the corrected meaning?"

Probe: The downstream reach is unknown. Record it and choose a visible correction and partner follow-up plan.

Two conditional branches

1. If players choose: Trust the reassuring back-translation. Then say: The same correlated translation path preserves the error in play. Require independent qualified meaning review, not a second pass through the same failure mode.

2. If players choose: Send another translated instruction immediately. Then say: An unreviewed correction creates conflicting directions. Use a minimal verified hold notice and human route while substantive wording is checked.

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.

22 | The translation that reversed the limit

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

What must be verified besides literal wording?

Strong reasoning to listen for: The action the recipient will take, negation, date, scope and consequence. Test comprehension with an appropriate reviewer, not just grammar.

How does versioning affect harm?

Strong reasoning to listen for: Same filename and near-identical layout make corrected meaning hard to identify. Use a visible version/date and state which instruction is superseded.

What population remains outside a known recipient list?

Strong reasoning to listen for: Reseller copies and onward sharing. Do not claim every recipient has been reached because 34 messages were delivered.

Which task can a support bot safely perform while meaning is disputed?

Strong reasoning to listen for: Route to a verified human or approved hold line; it should not improvise the substantive translated instruction.

Weak is incomplete; misleading is unsupported

Weak answer: "Run the corrected notice through the translator 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: "Back-translation matches the source, so the meaning is verified."

Next probe: Ask whether the reviewer and failure path are independent, and whether recipients know which version governs.

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 downstream reach is unknown. Record it and choose a visible correction and partner follow-up 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.

22 | The translation that reversed the limit

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 source, translated versions, reviewer comments, distribution identifiers and reseller account of local copies. Record exact disputed meaning without collecting unnecessary recipient details. 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 translated instruction is under review because its meaning may be reversed. Please use the verified human support route rather than rely on that version. We will identify the corrected version clearly; downstream reach is not yet fully 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

"We need a qualified language reviewer and service decision owner. The same-machine back-translation is not independent evidence. Please approve a minimal hold notice and a versioned correction after the meaning is checked."

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.

22 | The translation that reversed the limit

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

Exercise-specific containment boundary

Pause the disputed translation and bot answers based on it; provide a verified human route. Check version labels, reseller distribution and recipient understanding 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 source and translated versions be mapped by identifier?

Answer / record reference / absent evidence: _____

2. What independent qualified review of high-impact meaning is available?

Answer / record reference / absent evidence: _____

3. Which downstream distribution records exist and which local copies cannot be enumerated?

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.

22 | The translation that reversed the limit

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.

22 | The translation that reversed the limit

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 hold route and stop reuse of the disputed instruction.
Within 7 days Language owner	Retain independent review of negation/date/action meaning and a visibly versioned correction.
By day 30 Distribution owner	Rehearse a reseller correction and document known reach plus unresolved copies.

Named owners / actual dates / record location: _____

Debrief and three-label record

Was back-translation overtrusted?

Did correction include print and downstream copies?

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

Pause the disputed translated instruction and arrange qualified review of its meaning and correction wording.

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.

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.

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.