Make the sandbox case easy to defend.

A facilitator-ready guide for taking a FlowAI demonstration from workflow diagnosis through synthetic case review, extraction, screening, risk reasoning, audit evidence, and a scoped pilot close.

Open the sandbox ->: /sandbox · Download PDF: /resources/sandbox-demo-guide.pdf · Open homepage demo: /#demo · Compare live guide: /resources/demo-guide

The outcome · 25 MINUTES Prospect sees a repeatable synthetic case path. • Presenter explains fields, checks, reasoning, and routing. • Conversation ends with evidence and a pilot hypothesis.

Safety note No real PII. Use only the seeded sandbox cases. Never paste names, documents, emails, CPF numbers, or other real applicant identifiers.

BEFORE YOU RUN

Set the boundary, then set the question.

This is a sandbox demonstration, not a production decision. Say that plainly, agree on the workflow question you will test, and keep the prospect's compliance policy in the driver's seat.

Use seeded inputs: Stay with the named clean, PEP, and OFAC/UBO examples so the walkthrough is repeatable.

Say what is synthetic: Every person, document, hit, score, and decision shown is a demonstration fixture.

Keep policy visible: Ask where their team would approve, escalate, reject, or request more evidence.

Print this page: Use the browser print dialog or save as PDF; the interactive chrome is intentionally omitted.

Presenter reset Presenter reset: if someone offers to share a real applicant packet, stop and redirect to the seeded sandbox case. The demo should never become a data-transfer channel.

RUN OF SHOW

Eight beats from question to next step.

Keep the evidence narrative ahead of feature tourism. The links below are useful during the live session; their full text remains visible when printed.

TIME	BEAT	PRESENTER OUTCOME
0:00–2:00	**Frame the synthetic rules**	Confirm no-real-PII boundaries and the workflow question to test.
2:00–5:00	**Diagnose the current workflow**	Map intake, queues, checks, routing, evidence, and success metrics.
5:00–8:00	**Run a seeded sandbox case**	Establish the clean baseline and narrate intake through decision.
8:00–11:00	**Review extracted fields**	Connect structured output to analyst review and downstream routing.
11:00–14:00	**Explain sanctions and PEP screening**	Show how possible matches become policy questions, not automatic claims.

TIME	BEAT	PRESENTER OUTCOME
14:00–18:00	**Inspect reasoning and routing**	Compare clean, PEP, and OFAC/UBO outcomes in the playground.
18:00–21:00	**Separate audit evidence**	Distinguish playground output from a persisted case_id record.
21:00–25:00	**Close with a scoped pilot**	Agree on volume, owners, policy boundaries, measures, and a date.

FRAME THE SANDBOX

01 Make the rules part of the demonstratio

The safest demo begins with an explicit boundary: synthetic cases show the sh
a real applicant or replace a compliance decision.

Say

“Everything we show is seeded synthetic data. We will not enter real names, d
identifiers. We’ll use the case to inspect the workflow, not to make a real decisi

Do

1. Ask what workflow the prospect wants to pressure-test.
2. Confirm the audience and their policy owner.
3. Keep this guide and the sandbox playground (/sandbox) open.

Show

Point to the seeded case tabs and the visible decision, reasoning, and raw JS

> Expected result The prospect agrees that the walkthrough is a safe, repeatable inspection of a KYC/KYB flow—not a production test or a promise that synthetic outcomes match their final policy.

DIAGNOSE

02 Start with their current workflow.

Use the prospect’s own operating language to decide which sandbox result will be useful.

Say

“Before we run the case, where does an applicant slow down, become hard to defend, or move between teams today?”

Do

- Map document intake and missing-data handling.
- Ask who owns manual review and escalations.
- Ask what evidence an auditor or regulator needs.
- Capture the metric they want to improve.

Show

Relate their answer to the path: intake -> extraction -> screening -> risk score -> APPROVE / ESCALATE / REJECT -> evidence.

Listen for the bottleneck.

Ask about document quality, list-match review, analyst queues, exception ownership, SLA, and whether the current evidence can be replayed later.

Transition: “Let’s use a clean seeded case first, then introduce the exact risk signals your team would route differently.”

SEEDED CASE

03 Run the clean baseline in the sandbox.

Open the API playground (/sandbox) or homepage demo (/#demo) and establish the intake-to-decision sequence before showing exceptions.

Say

“We’ll start clean so the shape of a normal case is clear. Then we’ll deliberately add a PEP and a sanctions scenario to test routing.”

Do

1. Select Individual KYC.
2. Choose Clean applicant.
3. Run the case once; do not paste custom data.
4. Explain any work-email gate if the homepage demo shows it.

Show

Narrate intake, extraction, checks, score, decision, latency, and the structured response. Link to the API quickstart (/docs/quickstart) if implementation detail is requested.

Expected result / recovery The clean preset returns a structured decision and reasoning view. If the gate or rate limit interrupts the run, explain the expected access boundary, stop retrying, and continue from the captured result rather than improvising a real case.

EXTRACTION

04 Read the returned fields like an operator.

Move from a decision badge to the structured data that makes the next action understandable.

Say

“I’m going to read this as an analyst: what was available, what was checked, what changed the risk, and what should the next person or system do?”

Do

Open the response and connect each field to the workflow diagnosis. Call out missing or ambiguous data as a reason to review, not as a failure hidden below the fold.

Show

extracted_data · sanctions_hits · aml_flags · risk_factors

Useful field language

- extracted_data is the structured identity and document information available to the decision.
- sanctions_hits identifies a possible match returned by configured screening.
- aml_flags and risk_factors explain which signals shape review and routing.

Expected result The prospect can explain how a structured response would move into their case-management or analyst workflow. Keep the reminder visible: this is a synthetic result, not proof about any real person.

SCREENING

05 Explain sanctions and PEP signals without overclaiming.

Use seeded risk examples to show why a possible match changes the next action, while leaving the final policy judgment with the prospect.

Say

“A possible PEP or sanctions match is a signal for investigation and routing. It is not, by itself, a final finding about a real customer.”

Do

- Run the PEP hit example.
- Ask whether policy escalates or rejects it.
- Switch to Business KYB.
- Run OFAC hit on UBO.

Show

Point to the match signal, risk factors, affected party, and the decision route. Use the API reference (/api-docs) for the documented response shape.

> Recovery note: if the prospect asks whether the seeded name is a live list match, answer that it is a demonstration fixture. The useful discussion is what verification, enhanced due diligence, or human approval their policy requires next.

RISK REASONING

06 Inspect the chain, then map it to routing.

Show how the visible reasoning and score support an operational route without turning a demo score into a universal policy threshold.

Say

“The point is not only the final badge. It is whether an operator can see which checks led there and what should happen next.”

Do

1. Compare clean, PEP, and OFAC/UBO presets.
2. Open the reasoning-chain rows.
3. Show the raw JSON toggle and latency.
4. Ask which outcomes stay human-reviewed.

Show

Use Decision Theater (/live/theater) only as an optional visual deep dive into synthetic reasoning—not as an audit record or production case.

> Expected result The prospect can distinguish a signal, a score, a decision, and a routing rule. Their policy—not the seeded example—defines the production threshold and escalation owner.

AUDIT EVIDENCE

07 Separate playground output from persisted evidence.

This is the trust-building caveat. The browser playground is for inspection; the authenticated triage path is where a case

can be retrieved by an audit identifier.

Say

“The playground run is not automatically the audit record. A persisted triage event is the record we retrieve with a case_id when audit_logged is true.”

Do

Use the documented POST /api/triage path when you have an appropriate sandbox key. Capture its case_id and audit_logged response fields.

Show

Open the triage docs (/docs/api/triage) and retrieve the record with the documented GET /api/triage/:caseId lookup.

> Request path GET /api/triage/<case_id>

Point to the stored evidence

- Decision, score, reasons, and model metadata.
- Extracted data, sanctions hits, and AML flags.
- Processing time, timestamp, and case identifier.

> Evidence language Say “this is the persisted triage record when audit_logged is true,” not “every playground run is audited.” If no case ID is returned, explain that the decision completed but persistence was not confirmed in that environment.

08 · CLOSE THE PILOT

Turn the demo into a measurable operating hypothesis.

Close with the prospect’s own workflow, policy boundaries, and evidence requirements—not a generic invitation to try the product.

Say

“Based on what you shared, I’d recommend a focused pilot. Let’s define the first queue, the cases that stay human-reviewed, and the evidence we need to call it successful.”

Do

Confirm monthly volume, jurisdictions, document mix, owner, escalation policy, baseline, success metrics, and the next review date.

Show

Share the pilot scope template (/docs/pilot_scope_template.md), then offer pilot access (/founding) or a working session (/book).

> Pilot recommendation “I recommend a [duration] pilot for [volume] cases per month across [jurisdictions], starting with [document mix]. We’ll keep [PEP / sanctions / missing-data] cases under [your escalation policy], with [owner] accountable. We’ll measure [decision time], [straight-through or escalation rate], [false-positive rate], and [evidence completeness]. We’ll review the first results on [date].”

Closing checklist

- First workflow scope is named.
- Operational owner is accountable.

- Jurisdictions and document mix are agreed.
- Escalation and evidence policy are explicit.
- Success metrics have a baseline.
- Next meeting has an owner, date, and purpose.

Open pilot scope template ->: /docs/pilot_scope_template.md · View pilot offer: /founding · Book next session: /book

GUIDE FOOTER

FlowAI resources

FlowAI: / · Resources: /resources · Sandbox: /sandbox · Docs: /docs · Privacy: /privacy

© 2026 FlowAI LLC. All rights reserved. Formed in Kansas, USA. FlowAI® is a registered trademark of FlowAI LLC.