



LSAS Platform

Governed Workflows in Practice

Representative high-stakes
scenarios,
evaluated under LSAS governance.

CLINICAL, LEGAL, AND FINTECH SCENARIOS FROM THE LSAS SANDBOX

Audience: product, platform, security, and compliance stakeholders.

LSAS Platform

Enterprise AI Execution Authority™

ADDRESS

5365 W. 73rd Pl.
Arvada, CO 80003

CONTACT

+1 720.334.3838
www.lsasplatform.com
www.lsasplatform.com/contact

Governed Workflows in Practice

What this pack shows

This pack walks through three representative workflows the way LSAS Platform sees them: what the AI is asked to do, where the risk concentrates, which controls apply, and what the decision record looks like afterward. Clinical documentation, legal drafting, and payment-adjacent operations were chosen because they combine high-consequence outputs with realistic AI touchpoints.

How to read each scenario

- The workflow — what the AI touchpoint does.
- Where risk concentrates — the failure modes that matter most.
- Controls applied — the LSAS primitives in play: policy packs, evidence thresholds, deterministic outcomes (ALLOW, REDACTED, BLOCKED, ESCALATE_HITL), governed overrides, and decision telemetry.
- Representative decision envelope — the outcome and findings LSAS would emit for the illustrated pass.

Illustrative scenarios, not case studies. These scenarios are representative sandbox walkthroughs. They do not describe customers, deployments, or measured results. The LSAS sandbox is a controlled evaluation environment — not a claim of universal workflow coverage. Scenarios show representative behavior so teams can evaluate the control model before scoping a pilot.

Governed Workflows in Practice

Scenario 1 — Clinical note handling

The workflow

An AI assistant drafts clinical documentation from encounter notes: summarizing history, extracting medication lists, and proposing structured fields for the record.

Where risk concentrates

- Unsupported clinical assertions — statements in the draft with no grounding in the source notes.
- Identifier exposure — direct identifiers carried into summaries, downstream systems, or telemetry.
- Low-confidence extractions presented as fact — uncertain values rendered without qualification.

Controls applied

- Evidence thresholds: clinical assertions must trace to approved source content before release; ungrounded statements are recorded as findings.
- Redaction: direct identifiers are transformed or removed before the draft leaves the governed boundary — a REDACTED outcome, so the workflow continues instead of stalling.
- Escalation: low-confidence extractions route to a qualified reviewer rather than releasing silently — ESCALATE_HITL.
- Decision telemetry: every decision carries findings, evidence references, and the policy version for later review.

Representative decision envelope

Field	Representative value
decision	REDACTED
findings	1 unsupported assertion flagged; 2 direct identifiers redacted
evidence_status	partial grounding — 1 claim below evidence threshold
remediation	redacted draft released; flagged claim queued for reviewer
audit_metadata	tenant, app, environment, actor, timestamp, policy-pack version

Governed Workflows in Practice

Scenario 2 — Legal citation-sensitive drafting

The workflow

An AI assistant drafts legal memoranda and client-facing summaries, inserting case citations and statutory references.

Where risk concentrates

- Fabricated authority — citations that look real but resolve to nothing.
- Unsupported legal claims — conclusions stated without cited support.
- Superseded references — citations outside the approved authority set for the governing policy tier.

Controls applied

- Citation validators: every citation must resolve against an approved authority set before the draft can proceed.
- Deterministic outcomes: a draft containing fabricated authority is BLOCKED — it cannot be released under the governing policy tier.
- Audit trail of findings: each failed check is recorded with the offending passage and the rule that fired.
- Governed overrides: a licensed reviewer may authorize release with explicit exception capture — the override, the authority, and the reason are all recorded.

Representative decision envelope

Field	Representative value
decision	BLOCKED
findings	2 citations failed validation; 1 unsupported conclusion
evidence_status	citation check failed against approved authority set
remediation	draft returned to author with findings; override path available to licensed reviewer
audit_metadata	tenant, app, environment, actor, timestamp, policy-pack version

Governed Workflows in Practice

Scenario 3 — Payment-adjacent data controls

The workflow

An AI assistant helps operations staff reconcile payment-adjacent records: matching invoices, flagging anomalies, and preparing payout summaries.

Where risk concentrates

- Identifier exposure — account numbers and personal identifiers leaking into prompts, outputs, or logs.
- Unapproved data egress — record content crossing a connector boundary it was never cleared for.
- Consequential actions without approval — a summary that triggers a payout step nobody reviewed.

Controls applied

- Connector-boundary policies: data crossing an ingress or egress boundary is evaluated against explicit allow-rules; unapproved crossings are blocked.
- Redaction by default: identifiers are transformed before AI processing wherever the workflow allows it.
- Human escalation: any action with payout consequences requires explicit human approval — ESCALATE_HITL.
- One review artifact: the decision envelope gives reviewers the outcome, findings, evidence, and remediation in a single record.

Representative decision envelope

Field	Representative value
decision	ESCALATE_HITL
findings	payout-impacting action proposed; 1 identifier redacted at boundary
evidence_status	anomaly evidence present; human approval pending
remediation	held for human approval with full decision context
audit_metadata	tenant, app, environment, actor, timestamp, policy-pack version

Evaluate it in your environment.

Run these scenarios in the LSAS Platform sandbox, or book a scoping call to define one workflow, one boundary, and one definition of done for a 90-day pilot.

www.lsasplatform.com/playground · www.lsasplatform.com/contact

