

Agentic Systems

SPECIMEN INTERVENTION

Design how humans, agents and tools should actually work together - and prove the operating architecture on one workflow before it is rolled anywhere.

Specimen intervention. An I.O.V worked example of the method - no client is named and no client material is reproduced. Figures are illustrative thresholds, not measured outcomes.

PLATE 01 - WORKFLOW AND DECISION MAP

A high-consequence back-office workflow: an inbound customer obligation is interpreted, priced against policy, and either settled or escalated.

01 · INTAKE · AGENT-HELD

The obligation arrives through mail, portal or partner feed. Normalised into one record with its source and timestamp.

02 · INTERPRETATION · AGENT-HELD

Documents read, entities resolved, obligation classified against the policy taxonomy. Confidence recorded per field.

03 · POLICY MATCH · CONTESTED

The classified obligation is priced against written policy. Ambiguity is surfaced rather than averaged away.

04 · DECISION · HUMAN-HELD · DECISION

Settle, adjust or refuse. This is where consequence sits: money moves, or a customer is told no.

05 · EXECUTION · AGENT-HELD

Settlement instructed, ledger written, customer notified. Every action carries the mandate it acted under.

06 · REVIEW AND LEARNING · HUMAN-HELD

Sampled review of decided cases. Disagreements feed the policy taxonomy, not the model alone.

PLATE 02 - HUMAN AUTHORITY MATRIX

Who may act, under whose mandate, inside which limits - and where a human must remain in the loop.

INTAKE AGENT

Standing mandate · read-only outside its own record

READ CASE RECORD	May act
CLASSIFY OBLIGATION	May act
PRICE AGAINST POLICY	Must escalate
SETTLE WITHIN THRESHOLD	May not act
REFUSE OR ADJUST	May not act
CHANGE POLICY RULE	May not act

PRICING AGENT

Standing mandate · bounded by the written threshold

READ CASE RECORD	May act
CLASSIFY OBLIGATION	May act
PRICE AGAINST POLICY	Inside limits
SETTLE WITHIN THRESHOLD	Inside limits
REFUSE OR ADJUST	Must escalate
CHANGE POLICY RULE	May not act

CASE OPERATOR

Role mandate · accountable named human

READ CASE RECORD	May act
CLASSIFY OBLIGATION	May act
PRICE AGAINST POLICY	May act
SETTLE WITHIN THRESHOLD	May act
REFUSE OR ADJUST	May act
CHANGE POLICY RULE	Must escalate

REVIEWER

Sampling mandate · may reverse a decided case

READ CASE RECORD	May act
CLASSIFY OBLIGATION	May act
PRICE AGAINST POLICY	May act
SETTLE WITHIN THRESHOLD	Must escalate
REFUSE OR ADJUST	May act
CHANGE POLICY RULE	Must escalate

SYSTEM OWNER

Change mandate · cannot decide an individual case

READ CASE RECORD	May act
CLASSIFY OBLIGATION	May not act
PRICE AGAINST POLICY	May not act
SETTLE WITHIN THRESHOLD	May not act
REFUSE OR ADJUST	May not act
CHANGE POLICY RULE	May act

WRITTEN LIMITS

- Settlement above the written threshold is never agent-held, whatever the model's confidence.
- No actor may both decide a case and change the rule that decided it.
- Every agent action is written with the mandate it acted under; an action without a mandate is void.

ESCALATION, REVERSAL AND REVOCATION**E1 · DETECT**

Threshold breach, low confidence, policy conflict or reviewer disagreement raises the case.

E2 · HOLD

The action is stopped before effect. Nothing settles while a case is held.

E3 · ESCALATE

Routed to the named operator with the record, the confidence and the mandate that was claimed.

E4 · REVERSE

A decided case is unwound: ledger corrected, customer re-notified, reversal filed against the original mandate.

E5 · REVOKE

The agent's mandate is withdrawn immediately and its queue reverts to human hands. Revocation needs one action, not a release.

PLATE 03 - FIRST-PROOF SPECIFICATION

The one workflow proved first, written so it can pass or fail on evidence rather than on enthusiasm.

WORKFLOW UNDER TEST

One obligation class, one region, one queue - the narrowest slice where the consequence is real and the exposure is small.

ASSUMPTION UNDER TEST

That interpretation and policy matching can be agent-held at production quality while the settle-or-refuse decision stays human-held.

EVIDENCE IT MUST PRODUCE

Per-case: classification agreement against reviewer judgement, cost per run, time to decision, escalation rate and every reversal with its cause.

PASS LINE

Agreement at or above the written threshold, no unmandated action, and every escalation caught before effect - held across a full review cycle, not a good week.

FAIL LINE

One settlement without a mandate, or a reversal that could not be traced to its decision, ends the proof regardless of the aggregate numbers.

CONTROLS IN PLACE BEFORE DAY ONE

Written thresholds, hold-before-effect on every agent action, sampled reviewer queue, and one-action revocation.

THE 90-DAY SEQUENCE

DAYS 0-30 · INSTRUMENT AND SHADOW

The workflow is mapped as it actually runs, controls are stood up, and agents run in shadow with no effect on live cases.

DAYS 30-60 · GOVERNED PRODUCTION

Agent-held steps go live inside written limits. Every decision stays human. Evidence accumulates against the pass line.

DAYS 60-90 · ARGUE THE LEVERAGE

Quality and cost per run are read against the pass line, and the next workflow is scoped - or the proof is stopped and said so.

PROVENANCE

Specimen intervention. An I.O.V worked example of the method - no client is named and no client material is reproduced. Figures are illustrative thresholds, not measured outcomes.

<https://iov.agency/agnostic-systems>