

=====

Tenant: synthetic-governance-sample ? SYNTHETIC ? not a live customer
As of: 2026-08-08T12:00:00.000Z
Window: 30 days
Engine: 1.0.0
Clinicians: 12

1. What fires

Accountable justification asks the ordering clinician for one free-text clinical justification on a low-appropriateness advanced-imaging order (CT, MR, PET, or NM) at order-sign, without blocking the signature.

2. Fire-rate ceilings

Per clinician per rolling day: 4 [modeled]

The interruption budget. ARKA asks for justification only on low-appropriateness advanced-imaging orders, so the honest unit of burden is prompts per clinician per working day, not a share of all orders. Four accountable-justification prompts in a day sits well under published alert-fatigue thresholds and below the volume a busy orderer already sees from safety alerts. On breach: suppress the card and record 'ceiling-day'.

Per clinician per rolling 30 days: 40 [modeled]

A secondary guardrail at ten times the daily ceiling. It prevents a single high-volume cohort or a bad month from exhausting clinicians even when the daily cap is respected most days. On breach: suppress and raise a governance alert ('ceiling-30d').

Organisation share of eligible advanced-imaging orders: 8.0% [modeled]

A hard organisational stop. If AJ would fire on more than 8 % of eligible advanced-imaging orders org-wide over a rolling 30 days, something is wrong with the engine threshold or the indication cohort ? not with the clinicians. On breach: tighten the score gate to ? 2 and raise a governance event ('ceiling-org-share'), rather than silently accepting an 8 %+ interrupt rate.

3. Measured fire rate (this period)

Per-clinician median: 1 [measured]
Per-clinician p95: 1 [measured]
Per-clinician max: 1 [measured]
Org share: 6.0% [measured]
Fired / eligible: 12 / 200 [measured]
Suppressed (day): 0 [measured]
Suppressed (30d): 0 [measured]
Suppressed (org): 0 [measured]

4. Response and completion

Response share (prominent): 100.0% [measured]
Median added interaction: 8000 ms [measured]
p95 added interaction: 8000 ms [measured]
Server contribution p95: 0 ms [measured] (budget 800 ms)
Interaction p95 vs budget: 8000 ms [measured] (budget 10000 ms)

5. Mechanism status

State: ACTIVE ? written share 100.0% [measured] ? lane I1
Accountable justification: ACTIVE ? verified chart writes for clinician responses meet the mechanism floor; the published effect may be expected for this deployment.

6. Commitments (firewall.json forbidden[])

- storing, logging, transmitting to a third party, or reading justification text
- joining justification text or its hash to any performance measure, peer packet, ledger cell, incentive computation or clinician-level aggregate
- blocking, delaying or conditioning an order signature
- requiring a minimum length, a structured taxonomy, or a selection in place of free text
- firing above the published ceilings
- claiming an effect for a deployment where the record write is unavailable

Every number above carries a provenance label. Small cells are suppressed at $n < 11$.

Appendix ? Schema (J6 no-text)

=====

ARKA does not store what your clinicians write. The justification is composed in your EHR, written with your clinician's own credentials, and lives in your chart. We keep a hash, a length and a timestamp, so we can tell you the mechanism was present. There is nothing for us to hand over, mine, or lose.

You can verify this. Run 'npm run aj:verify-no-text' against your own instance, or read the schema ? aj_justification_receipts has no text column, and a CI check fails the build if anyone adds one.

Migration: supabase/migrations/069_aj_core.sql

aj_justification_receipts columns include text_sha256 and char_length ? never a text, justification, narrative, comment, or body column.
Verify: npm run aj:verify-no-text ? CI fails the build if a forbidden column is added.