

Agent Service Assurance Starter Kit

A practical preflight, evidence record, and postflight checklist for machine-purchased services.

Before payment

Decide whether this service transaction should proceed for this buyer, task, region, and cost envelope.

At execution

Bind the selected offer, policy decision, payment authorization, service version, and actual route.

After delivery

Verify delivered cost, latency, result shape, side effect, and buyer-defined outcome.

After failure

Preserve evidence for retry, reroute, refund, dispute, and source-reliability calibration.

What this kit does not claim

It is a design and discovery instrument—not a certification, legal opinion, production control, or guarantee. Replace every blank with evidence from your own environment.

Seven checks before an agent pays

- **Economic authority:** identify the principal, delegated agent, wallet/account, task, and maximum authorized spend.
- **Service identity:** resolve seller, endpoint, marketplace aliases, ownership, version, protocol revision, and supported capabilities.
- **Offer binding:** preserve price, asset, chain/rail, recipient, SLA, refund/recourse terms, expiry, and a signed offer hash.
- **Policy fit:** test data boundary, jurisdiction, buyer allow/deny rules, approval state, cost cap, latency ceiling, and fallback policy.
- **Current evidence:** compare seller claims, marketplace state, native telemetry, status feeds, independent probes, and buyer history.
- **Alternative routes:** record at least one substitutable service or a deliberate no-purchase option.
- **Decision:** issue ALLOW, CHALLENGE, REROUTE, or DENY with a short expiry, confidence, reason codes, and evidence references.

Minimum decision record

Field	Example	Evidence
decision_id / expires_at	dec_... / +120s	Signed decision
buyer / agent / task	Principal + delegated runtime + purpose	Customer-private policy
service / version / route	Resolved entity and selected endpoint	Registry + live probe
offer_hash / cap / SLA	Terms frozen at selection time	Signed offer
verdict / confidence	ALLOW · 0.82	Reason codes + sources

Know what predicted delivery

Uptime, price, identity, and risk signals are increasingly available from native platforms. The useful graph preserves which evidence source actually predicted a successful economic outcome.

Source class	Capture	Question
Seller claim	Claim, version, signature, observed time	Was the promise current and accurate?
Marketplace / registry	Listing, aliases, categories, price, commit	What existed at selection time?
Native gateway / monitor	Route, tokens, cost, drift, approval, trace	What did the platform observe?
Independent probe	Region, protocol, latency, fault, timestamp	Did outside evidence disagree?
Buyer trace	Task, internal policy, execution, final outcome	Did the buyer receive economic value?

Legal-access provenance

For each source, preserve whether access came from an explicit API, licensed feed, public registry, status page, customer-authorized session, user-delegated browser, or server-side crawler—plus authentication basis, applicable terms, rate limits, retention rule, and legal-review state.

Data boundary

Customer-private Requests, policies, contract prices, caps, identities, approvals, traces, incidents, disputes, and transaction outcomes.	Assurance-proprietary Raw probes, fault injection, entity resolution, source calibration, drift, counterfactual routing, and cross-customer outcome models.
---	---

Verify the actual economic outcome

- ❑ Did the executed seller, endpoint, version, gateway, facilitator, network, recipient, and asset match the decision?
- ❑ Did the service start only once? Check replay, duplicate execution, decision-nonce reuse, retries, and idempotency.
- ❑ Did the charged amount match the signed offer and buyer cap, including fees and delivered cost?
- ❑ Did the response meet the promised latency, schema, freshness, capability, and error budget?
- ❑ Did the external effect occur—record created, identity verified, shipment booked, dispute updated, or other buyer-defined side effect?
- ❑ Was payment settled but service not delivered, service delivered but payment mismatched, or outcome still unknown?
- ❑ Which intervention—allow, challenge, reroute, retry, refund, dispute, or human approval—improved the result?

Verdict semantics

Verdict	Meaning	Next action
PASS	Bound route, payment, delivery, and outcome match.	Release / close.
PARTIAL	Useful delivery with a measurable shortfall.	Adjust, retry, or claim credit.
FAIL	Material mismatch or non-delivery.	Hold, reroute, refund, or dispute.
UNKNOWN	Evidence cannot establish the outcome.	Challenge or require human review.

Use the kit on one real transaction

1. **5 min:** name the buyer, agent, paid external service, and economic outcome.
2. **5 min:** freeze the current offer, identity, price, policy, route, and protocol versions.
3. **8 min:** map seller, marketplace, native-platform, independent, and buyer evidence.
4. **7 min:** select the smallest inline hook: gateway policy, tool proxy, MCP challenge, workflow gate, or post-settlement monitor.
5. **5 min:** define PASS/FAIL/UNKNOWN, one fault injection, one reroute, and the receipt that would prove value.

Exit criteria

- ☐ A named transaction and owner.
- ☐ A short-lived decision point the workflow can enforce.
- ☐ A postflight outcome the buyer can verify.
- ☐ A customer-private / public / proprietary data boundary.
- ☐ A kill rule if native controls make independent assurance redundant.

Next step: bring one consequential external-service transaction to InferenceView. We will tell you whether there is a credible independent-assurance case—or whether the native gateway already covers it.

inferenceview.com · hello@inferenceview.com