

A WORKING DEMONSTRATION

An AI agent that awards a contract — accountably

How SIGNET makes autonomous procurement decisions governed, auditable, and safe to adopt.

THE QUESTION THAT ACTUALLY MATTERS

AI agents are arriving in procurement. They can read a tender, score bids, and draft an award in seconds. But the hard question was never whether an agent *can* make a decision — it is whether you can trust, audit, and stand behind the decision it makes. Who authorised it? What rule did it apply? Was the person who approved it actually empowered to? Can you prove the record wasn't altered afterwards?

Without convincing answers, autonomous procurement is a liability. SIGNET is built so the answers are **yes by construction**. This demonstration shows that working, end to end — and you can run it yourself at concert.foundation/standard/agent.

THE SCENARIO

A buyer runs a competitive procurement for managed network services, valued at €12 million. Two suppliers submit compliant bids:

- **Supplier A** bids €11.4M — the stronger quality submission.
- **Supplier B** bids €10.8M — the lower price, with strong social value.

The buyer's published evaluation rule is a Most Advantageous Tender model weighted **price 20%, quality 55%, social value 25%**. An evaluation agent must score the bids against that rule and propose the award.

How it unfolds — and what each step proves

1 The agent reads the opportunity

What happens. The agent reads the sourcing event and the evaluation policy attached to it.

What it proves. It works from the buyer's own published rule, held as data. The weights it applies are read straight from the policy, not hard-coded — change the policy, and the agent's behaviour changes with it.

2 It checks its mandate — and hits a limit

What happens. Its Mandate permits it to evaluate and propose awards, but only to an autonomous ceiling of €10M. This procurement is €12M, so the system requires a named human approval before the award stands — and that approval is itself verifiable: a category director acting under a delegation-of-authority credential whose €25M ceiling demonstrably covers the value.

What it proves. Human oversight is enforced by the system, and the human's authority is checkable, not assumed. The leash is structural — for the agent and for the approver.

3 It evaluates the bids — and the cheaper bid loses

What happens. It derives a price score from each bid and combines it with quality and social value under the published weights. Supplier A wins, 0.859 to 0.842 — the more expensive bid, €600K above its rival.

What it proves. The agent is not a cheapest-bid picker. The buyer's rule weights quality at 55%, so the agent awarded the higher-quality supplier — explainably, reproducibly, with the arithmetic on the record. Anyone asking 'is it justifying overspend?' finds the answer in the rationale: the published policy decided, not a preference.

4 It records the decision — with its reasoning

What happens. It issues an award Decision recording what was decided, under which mandate, the bids considered, the policy applied, a plain-language rationale, the verifiable human approval, and cryptographic provenance.

What it proves. The 'why' travels with the 'what'. An auditor, a regulator, or an unsuccessful bidder can see the complete basis months later — including who approved it and on what authority.

5 The award is made — and the whole sequence is sealed

What happens. An Award is issued with a regulatory standstill, and every action is written as an append-only, cryptographically hash-chained record.

What it proves. Alter any record after the fact and the chain breaks, detectably. The demonstration shows this live: tamper with a step, and verification turns red.

WHAT THE DEMONSTRATION VERIFIES

- **Valid.** Every record conforms to the published standard.
- **Bounded.** The decision stayed within the agent's mandate, with human approval where the threshold required it.
- **Authorised.** The approver's delegation-of-authority credential covers the awarded value — the approval is verifiable, not nominal.
- **Tamper-evident.** The audit trail is an unbroken cryptographic chain; altering any step is detected.
- **Attributable.** Every assertion carries provenance.

Why this matters for your organisation

- **It de-risks AI adoption.** Trust, control, and auditability are built-in properties, not policy documents.
- **It is not a one-trick demo.** The same harness governs supplier qualification (conditional outcomes with value caps) and multi-party reverse auctions (deterministic, provably fair closes) — three governed actions, one architecture. See them all at concert.foundation/standard.
- **It works with your stack and any model.** SIGNET orchestrates over existing systems, and the governance is identical whatever model does the reasoning — no AI-vendor lock-in.
- **It is an open standard, not a product.** Stewarded by the neutral Concert Foundation, verified by a public conformance suite every implementer passes identically.

WHAT IT DOES — AND DOESN'T — SHOW

The reasoning here is deterministic, so the result is identical every run — which is what makes it a *proof*, not a one-off performance. A live AI model drives the identical flow with every governance guarantee intact; that version is for live demonstration. We are deliberate about the line between the two.

See it for yourself

Run the demonstration in your browser at concert.foundation/standard/agent — including tampering with the record and watching verification fail. These claims are re-verified in automated testing on every change to the standard.

Concert Foundation · concert.foundation · hello@concert.foundation