The agent said the refund succeeded.
The system of record said otherwise.

Pruvz creates a business evidence record for every AI agent action it covers. It brings together what the agent saw, the policy in force, its decision and reported result, then independently checks the outcome in your business systems. If the agent's report differs from the verified outcome, Pruvz flags the mismatch quickly and shows the evidence behind it.

One action. One evidence record.

A customer asks for a $240 refund. The agent approves it and reports success. Illustrative example.

What the agent and tools reported	LOOKS CLEAN
Policy	Refund Policy v4.2
Refund tool	Returned SUCCESS
CRM	Case CS 1198 closed
Agent report	"Refund complete."

What Pruvz verified	OUTCOME MISMATCH
Billing	No refund for invoice 2291
CRM	Case closed
Verdict	Mismatch recorded
Review	Sent for human review

01 Context · 02 Policy · 03 Decision · 04 Action · 05 Pruvz check · 06 Result and review

All six parts are kept together in one record you can open later.

A trace shows what the agent attempted.

Pruvz reads the business systems, verifies the result and keeps it as evidence.

When the report matches reality

The evidence stays.

Pruvz keeps the context, policy, decision, reported result and verified outcome together. If a customer questions the result, a manager reviews the decision, or risk, compliance, or audit teams need proof, the answer is already there. No rebuilding the story from logs spread across systems.

When the report differs from reality

Pruvz shows the gap.

- **The agent reported success, but the action never happened**
- **The amount or account was wrong**
- **One system was updated, but another was not**
- **The decision did not follow the policy in force**
- **Conflicting system results, routed for mismatch review**

Across refunds, subscriptions, orders and more.

Every covered action rolls up into one view.

ILLUSTRATIVE PRODUCT VIEW, SAMPLE DATA

1,284	98.2%	23	11
Business actions, last 30 days	Verified in systems of record	Outcome mismatches	Awaiting human review

Agent	Actions	Verified	Mismatches
Refund agent	412	99.3%	3
Billing update agent	508	97.0%	15
Order fulfillment agent	364	98.6%	5

Illustrative sample data. The 11 actions awaiting human review are a subset of the 23 outcome mismatches.

Managers see the overall picture. Operations, risk and compliance can open any metric and inspect the actions and evidence behind it.

Fits alongside your agents.

Read-only.

Pruvz verifies by reading your systems. It does not execute business actions.

Non-blocking.

Verification runs after the action, outside the agent's execution path.

API integration.

Your agent or orchestrator sends authenticated action and completion events through the API. Designed to fit the workflow you already run.

Security: tenant-isolated evidence, encryption in transit and at rest, append-only evidence with cryptographic commitments that support independent tamper-evident verification, and role-based review.

Working product · Customer pilots open · Founding design-partner program. Ready to connect to a customer's systems for a proof of value.

Watch the product demo
pruvz.ai/demo