SimBI by Data Sense

Policy control at the point of execution

For AI agents in regulated financial services

The agent proposes. Your institution's policy decides.
An integrated boundary enforces. Evidence records.

November 2026

01. THE PROBLEM

Agents now act on regulated data and systems. Firms can’t yet show each action complied with policy.

Static permissions

Necessary, but incomplete

Say who may touch a resource, not why, for whom, with what data, or for what purpose.

Prompt guardrails

Advisory, not authority

Live inside the model they govern, and can be ignored or rewritten.

Agent self-attestation

Only as complete as what is declared

Verifying the sources and intent an agent reports is useful, but the agent chooses what to report.

Post-hoc observability

Too late to prevent

Explains an event after it has already happened.

> The missing piece is a decision made from **what the institution knows** (its policy, entitlements, data lineage and data meaning), applied to the exact operation before it runs, so it can still be stopped.

28% of financial services AI products already run fully autonomously.

Source: CCAF, 2026 Global AI in Financial Services Report, p. 52 (vendor survey).

02. HOW IT WORKS

The agent proposes. The institution's knowledge decides.

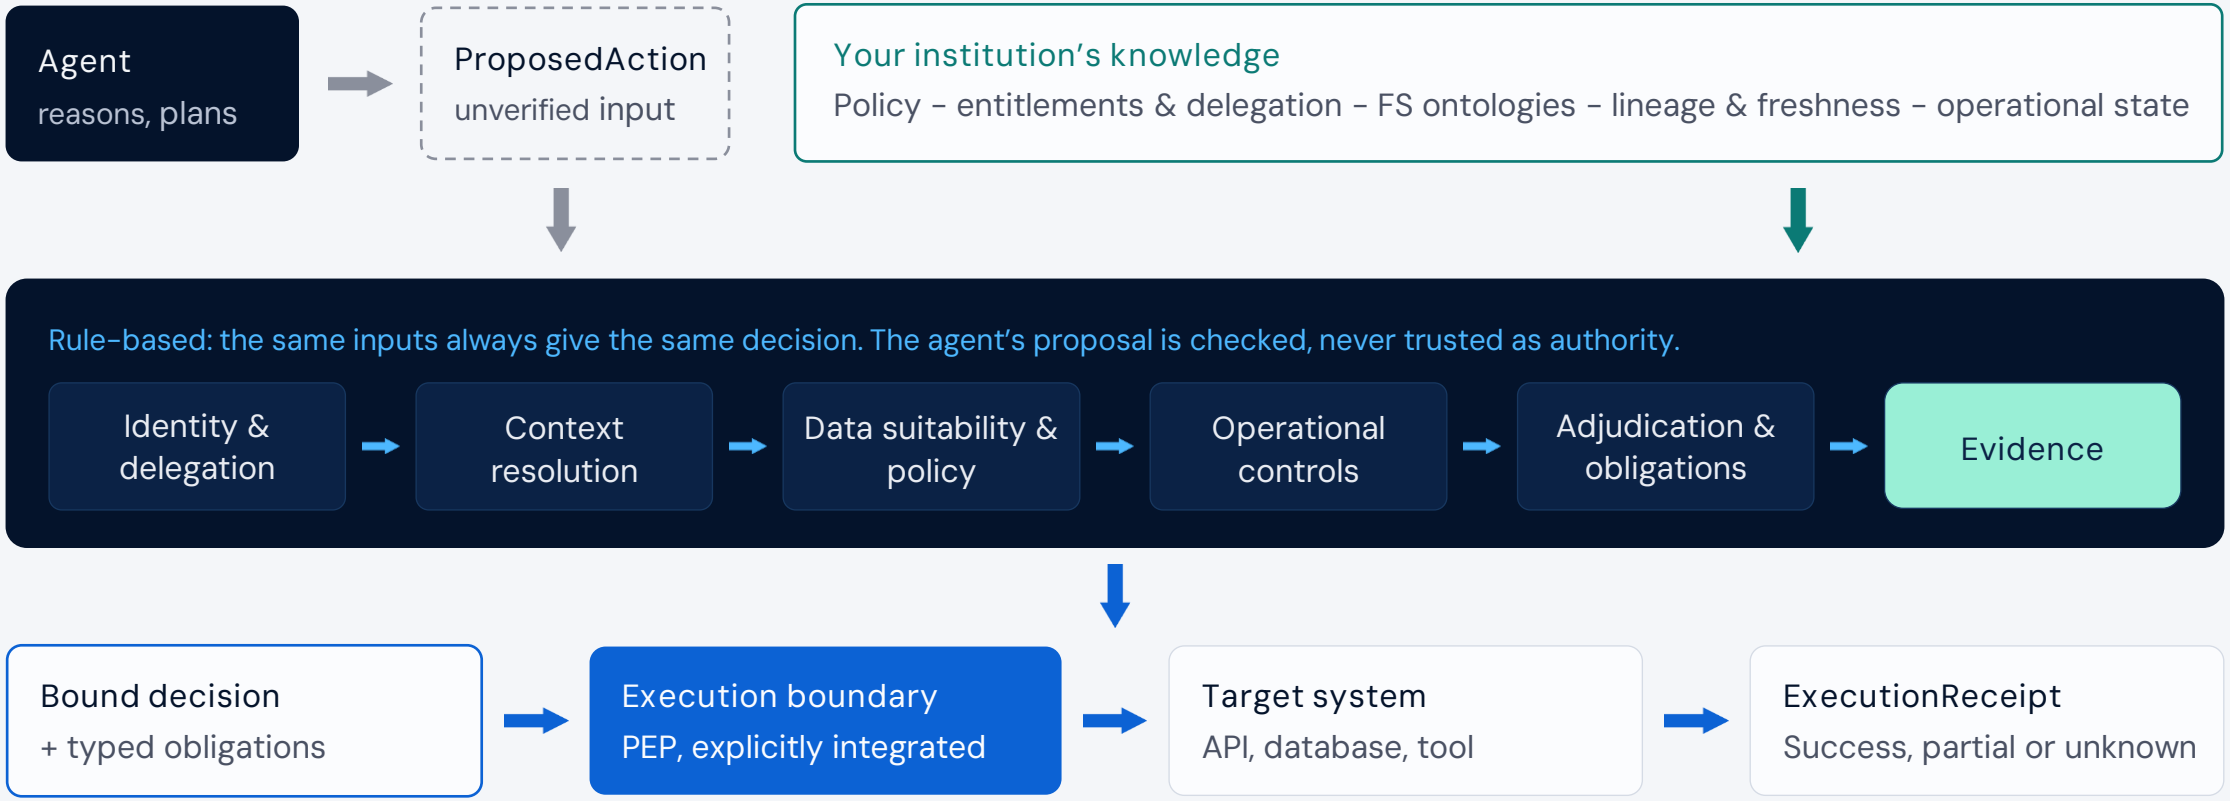

03. THE FINANCIAL-SERVICES PACKAGE

Domain knowledge ships with SimBI

Ontologies

Aligned to ontologies such as FIBO, so policy is written in business terms and applied to exact fields.

Lineage

Each input's source, version and freshness is pinned to the decision. Works with tools such as OpenLineage.

Data suitability

Is this data fit for this use? Assessed separately and recorded in every decision. A fit result is evidence, never permission.

Plug-and-play

Keep your identity provider, policy engines and catalog. Connectors and enforcement adapters bind them to the exact operation.

Open core

Decision contracts, a basic decision engine, lineage connectors, and MCP and framework adapters that fail closed.

Financial-services edition

Ontologies, lineage mappings, regulator-mapped policy packs, hardened adapters and signed supervisory records.

Status — Designed, not yet built. The FIBO subset is still being scoped.

04. FOR REGTECH AND SUPTECH

What a supervisor can ask of every decision

Question	Evidence in the record	Where
Who acted, for whom?	Verified agent identifier and its authority chain, scope and validity	context_refs
Which rule applied?	Policy identifier, version and effective time	policy_refs
On what information?	Each source's authority, freshness, lineage and version, pinned	context_refs
Was the data fit for this purpose?	Suitability result, recorded separately from authorization	*Planned*
Was a human involved?	Reviewer, approved scope, approval time and expiry	review_refs
What actually ran?	Executor status, attempt and obligations verified	ExecutionReceipt
What couldn't be verified?	Explicit NOT_EVALUATED and recorded unknowns	resolution_state

Decision and execution are separate records. Both are evidence for supervision, not proof of correctness or of complete coverage.

05. ARCHITECTURE DECISIONS

Decisions that shape the control plane

Decision	Why	Trade-off
Keep data suitability separate from authorization	Data being fit for a task isn't permission to use it. Our earlier POC let a passing check read as an allow.	A suitability assessment is a required input: if it's missing the operation is held; if the data isn't fit, policy denies it.
Govern a bounded operation, not a conversation	A decision must bind to one exact action, target and resource version	Ambiguous requests are clarified, escalated or rejected
Deterministic core; LLMs only advise	Decisions can't depend on model availability or model choice	Advisory classifiers (proposed) may suggest routes, never authorize actions
Three-valued authorization: ALLOW, DENY, NOT_EVALUATED	Missing or stale evidence must never look like permission	Outages or missing information hold work rather than let it through
Standard security access-control roles (PEP, PDP, PIP, PAP)	Security teams already know them, and existing policy engines can plug in	We adopt the roles only, not the XACML policy language they come from



06. EXAMPLE (ILLUSTRATIVE)

Policy shapes the operation, not just whether it runs

ProposedAction "Pull account balances and transactions from the treasury management system for today's liquidity report."

1. Resolve

Delegated agent, Treasury Ops.
Target: TMS, read. Purpose: liquidity reporting.
Ontology: fields are account-level.
Lineage: today's extract is current.

2. Decide

ALLOW + obligations:
balance and transaction fields only;
no customer-level data;
row limit.

3. Enforce

The PEP at the TMS connector applies the field list and limit *before* the query runs, with a scoped read credential.

4. Record

Decision and receipt are stored separately: policy version, context versions, lineage, authorizations, and what actually ran.

If the agent asks for customer-level fields
→ DENY

If the policy requires a human approver
→ ESCALATE, held

If the entitlement source is unreachable
→ INDETERMINATE, not run

07. AGENT IDENTITY

Pluggable identity, provable authority

Identity mechanisms are plug-ins: KERI in the proof of concept, vLEI and others through adapters. Delegated authority must be a verifiable ACDC chain. First spec: an AML agent reads transaction evidence for one case.

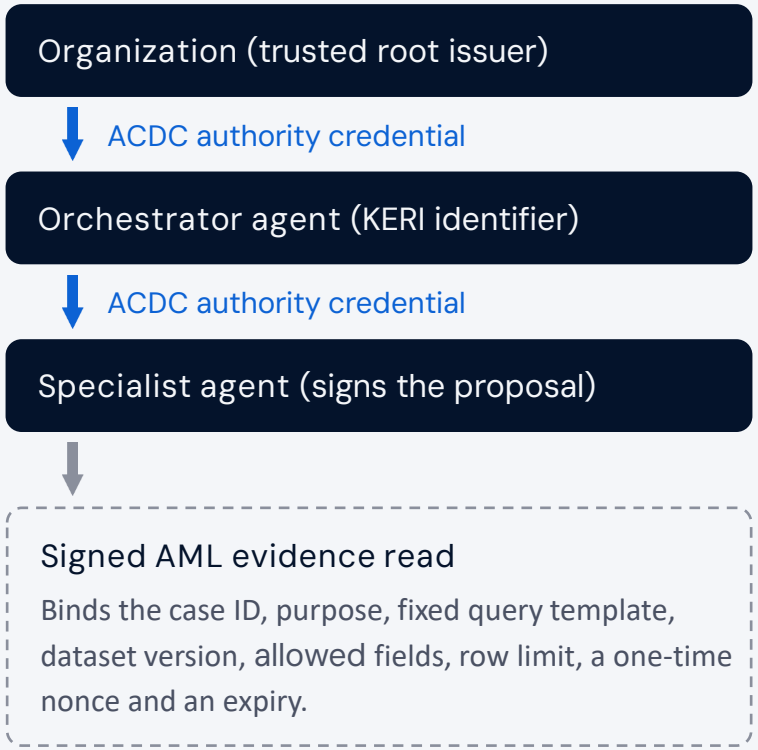

Test case	Result
Complete chain, exact in-scope request	ALLOW (evaluation only)
Asks for an excluded field or too many rows	DENY
Child credential claims more than its parent	DENY
Tampered payload or wrong signer	NOT_EVALUATED
Broken chain or untrusted root	NOT_EVALUATED
Key or credential status stale or unknown	NOT_EVALUATED

Status: evaluation-only spec on sample test inputs; enforcement is Phase 2. Identity is one input: policy, suitability and lineage still decide.

08. ENTERPRISE APPLICATION

Plug into the stack you already run

Topology & coverage

Gateway, tool adapter, sidecar, library or queue. Each has its own bypass routes; coverage is the integrated paths only.

IAM & policy fit

Keep your identity provider, identity standard and policy engines. SimBl binds them to the exact operation.

Latency & uptime

Decisions sit in the request path. Mandatory controls fail closed, so their uptime becomes the agent's uptime.

Scoped credentials

The PEP holds operation-scoped credentials, not the agent. Expiry, replay and idempotency are checked at dispatch.

Human review load

ESCALATE creates queues. Approvals bind to one exact scope and expire, so the cost of review must be designed for.

Audit evidence

Decision and receipt records need retention, redaction and access rules that satisfy audit without overclaiming proof.

Adoption path Open core → observe-only → one enforced, low-impact read → ontology and lineage providers → MCP and framework adapters

09. STATUS AND ROADMAP

Working today, designed, and the next proof point

Working today: pattern proofs

Power BI MCP and VaultGraph prove bounded resolution and gated queries in separate Data Sense tools. Earlier FastAPI/DSPy POC is retired.

Designed: ADI

Contracts & predicate (Phase 0 draft), decision states, typed obligations, KYA and AML specs, FIBO-aligned ontologies, lineage connectors, and the two editions.

Next: first proof

One low-impact read through one connector and one PEP, with decision, receipt, bypass and failure tests.

0	1	2	3	4	5
Contracts & threat model	Additive schemas & evidence	One enforced vertical slice	Context providers	MCP & framework adapters	Interop & hardening

Each phase closes with an evidence gate and no enforcement claim is made before the gate is proven.

WORKING TOGETHER

Prove enforcement on one bounded operation, then widen.

What we're asking partners for

Design partner

A read-only regulated workflow, such as AML evidence retrieval, to run the first enforced slice on.

Credential profile

Help deciding which KERI/ACDC credential formats, issuers and status sources to trust for agent authority.

Supervisory input

Guidance on what a decision record must contain to be useful to a supervisor.

Nicholas Lombardi
nlomb@datasenseco.com
datasenseco.com/simbi

SimBI by Data Sense

APPENDIX. IDENTITY DECISIONS

Decisions on identity and evidence

Decision	Why	Trade-off
Pluggable identity; KERI for the MVP *(vLEI and others plug in via adapters)*	Agent-identity standards are still emerging; KERI gives verifiable key state	One adapter per protocol; KERI specs and keripy pinned first
ACDC credentials for delegated authority	Authority chains are explicit, scoped and checkable hop by hop	Schemas, root issuers and status sources still need to agree
JSON records that carry their own content hash (SAID), signed in CESR format	One canonical, signable form native to KERI/ACDC	Envelope rules pinned in Phase 0; no interop until verified
Record IDs are not identity	UUIDs and hashes give idempotency and lineage joins, not authentication	Two ID systems, kept deliberately separate
Decision and receipt are separate records	Supervisors need to assess what was allowed and what actually ran	Two records to correlate, retain and protect