


Agentic Architecture Framework - UAAF

Scaling Agentic Work on Pega: Where UAAF Adds the Governance Lay



The Value Proposition

You invested in Pega. Pega executes and audits the work beautifully. Now scale agentic programmes without governance stalls. UAAF adds the one layer Pega cannot: **proof that your agents stay inside their boundary**. Infinity '26 makes the integration native.

That is the entire value proposition. What follows is evidence — verified against what Pega Infinity '26 actually ships, and what it leaves to you.

Pega owns orchestration and audit. UAAF owns governance and proof. Together they answer the eight questions your risk committee will ask before approving scale.

The Core Problem

Your first Pega agentic pilot worked. The demo was solid. The case lifecycle held. The audit trail was perfect.

Then governance asked questions Pega cannot answer — because they sit outside orchestration.

Your current stall is not Pega's fault. It is a governance gap. UAAF fills it.

Pega Provides


Orchestration, audit trail, approval gates, fairness screening, outcome measurement

UAAF Adds

Governance charter, bounded error verification, cryptographic proof, independent monitoring

Together

Evidence-based answers to every risk committee question before approving agentic scale



1. What You Get from Pega, and What UAAF Adds

Pega is strong at:

Deterministic and AI case lifecycle

Stages, SLAs, escalation – this is your primary guardrail

Immutable audit trail

Strong traceability, already accepted by auditors

Approval gates and HITL

Maker-Checker patterns, escalation, safe defaults

Tool discovery and invocation

Pega as an MCP server, A2A on agents; third-party agents can invoke your processes

Fairness screening

Ethical Bias Check, T-Switch transparency

Outcome measurement

Time-to-resolution, HITL reduction, SLA evasion – all observable from your Pega cases

UAAF adds the governance layer that sits beside Pega:

Governance charter

Versioned, signed, enforced as Pega rules but authored outside Pega

Non-determinism risk register

Tracking per-domain hallucination and model fragility

Bounded error verification

Proving agent decisions are sound at a measured error rate you control

Cryptographic authority binding

Proving what executed in downstream systems matches what was approved in Pega

Behavioral evaluation

Measuring whether agents still behave correctly as case mix drifts

Independent monitoring

Watching agents from outside Pega, in case Pega itself is compromised

Evidence repository

Holding the proof that regulators and auditors will demand

2. The Capability Map

This is not "Pega vs UAAF." It is "what you bought from Pega" and "what UAAF adds on top."

PEGA FOUNDATION

Pega gives you — use these as your foundation

- **Deterministic & AI case lifecycle** — Stages, transitions, SLAs, escalation rules. The agent is free inside a stage; everything else is locked.
- **Immutable case history** — Full audit trail, auditor-accepted, retrievable years later.
- **Approval gates** — Maker-Checker, safe defaults on timeout, agentic assignment agent (proactive outreach).
- **Boundary Guardian pattern** — When rules before effectful connectors. Pattern exists; discipline is yours to enforce.
- **Tool exposure via MCP** — Pega as server, A2A on agents. External agents (Claude, Gemini, OpenAI) can invoke your processes.
- **Agentic Process Fabric** — Registry of agents and workflows, with per-workflow discoverability flags.
- **Fairness screening** — Ethical Bias Check pre-deployment, T-Switch transparency. Shipped with Pega.
- **Outcome measurement** — Observable directly from Pega case data — no extra instrumentation needed.

UAAF LAYER

UAAF gives you — this sits beside Pega

- **Governance charter** — Signed, versioned declaration of bounded freedoms and Never-Do constraints. Lives outside Pega; generates Pega rules.
- **Bounded error verification** — Domain reasoner trained on your Pega case history, verified against your Pega rules. Calls Pega over MCP. When confident, commits; when unsure, routes to your Pega approval queue.
- **Cryptographic authority binding** — Proof that what executed downstream matches what was approved in Pega. Independent notary, cryptographic signature.
- **Behavioral evaluation** — Measuring whether your agents still behave on live Pega work. Golden dataset, judge models, calibration against rubric.
- **Agent identity lifecycle** — Issuance, scoping, rotation, revocation of agent credentials. Sits in your secrets platform; governs which agents can invoke which Pega processes.
- **Independent monitoring** — Watches your agents from outside Pega. Cannot be suppressed if Pega is compromised.
- **Non-Determinism Risk Register** — Tracking per-domain hallucination risk, model fragility, drift impact. Governance team maintains.
- **Evidence repository** — Holds decision provenance, evaluation records, incident history. What regulators will ask for.

☐ *The pattern: Pega covers orchestration and audit. UAAF covers governance and proof. Together they answer the eight questions your risk committee will ask before approving scale.*

3. Why UAAF Solves Your Governance Stall

Your first Pega agentic pilot worked. The demo was solid. The case lifecycle held. The audit trail was perfect. Then governance asked:

"How do we know the agent's decision is correct?"

"Can you prove what executed matches what was approved?"

"What evidence will we show auditors when a decision is disputed?"

Pega cannot answer these because they sit outside orchestration. Your current stall is not Pega's fault — it is a governance gap. UAAF fills it.

With UAAF + Pega together:

1 Risk committee approval moves faster

Answers to the eight governance questions are evidence-based, not asserted.

2 Blueprint Delivered becomes more precise

You enter Blueprinting with a signed governance charter and autonomy envelope, not a blank slate for governance to argue about later.

3 Integration spend is visible earlier

UAAF assessment measures API maturity before Pega build, not discovered during Authoring.

4 Fewer failed pilots

Processes that need no reasoning are explicitly excluded from agentic scope. You do not waste budget on an agent doing work a Pega decision table already did.

5 Expansion mapping works

Scoring across a function shows where the next three Pega missions are. One case type becomes a programme.

6 Renewal conversations have a narrative

Re-run the assessment after each phase and you have movement evidence — which is what CFO renewal conversations are built on.

4. How UAAF Plugs into Your Pega Delivery

You are using Blueprint Delivered. Good. UAAF aligns to that lifecycle.



Blueprinting

Authoring

Value
Activation

Blueprinting

Architect: process alignment, integration readiness.

UAAF provides: Signed governance charter (bounded freedoms, Never-Do constraints), autonomy envelope, risk classification. This must sign before Blueprint generation.

Authoring

Developer: case build, stages, connectors.

UAAF provides: Guardrail overlay: approval gates, exception taxonomy, tool tiering, Boundary Guardian placement, which stages stay deterministic.

Value Activation

Full team: testing, deployment, monitoring.


UAAF provides: Evaluation harness defined before build, acceptance criteria, production monitoring plan, evidence pack for your risk committee.



Critical sequencing rule: Do not complete Pega Blueprint design before your UAAF governance charter is signed. Pega's Blueprint is fast and faithful – it encodes whatever assumptions are in the room and produces correct case types in days. Unwinding assumptions after Blueprint costs more than the Blueprint saved.

5. How Your Teams Use Both Together

UAAF is not a Pega credential and not a Pega role. It is a governance discipline spanning your architecture, platform, and risk teams.



Your Pega Academy training is separate and essential. UAAF enablement is separate and orthogonal. The overlap is architectural — knowing where Pega's deterministic guarantees end and UAAF's empirical guarantees begin.

6. What Regulators and Auditors Will Ask


Requirement	Pega Provides	UAAF Adds
EU AI Act Art. 14 Human oversight	Approval workflows, escalation, audit history	Proof the approver looked; engagement telemetry in the Pega record
Art. 17 Quality management	Deployment gating, observability	Golden dataset, judge models, calibration on live Pega cases
Art. 12 / 72 Logging	Case history	Cross-plane trace correlation; independent behavioral monitoring
ECOA / Fair lending	Ethical Bias Check (pre-deployment)	Adverse-action notice as mandatory Pega checkpoint; ongoing fairness measurement
NIST AI RMF	AI Governance Framework, decision governance	Continuous measurement; behavioral evaluation against rubric



Timing note: The EU AI Act's high-risk obligations defer to December 2027 and August 2028. The deferral removes the deadline, not the requirement. You cannot manufacture a year of behavioral evidence in the final quarter of 2027. The deferral is your window to build UAAF governance infrastructure.

7. The First Ninety Days – Proof Before Scale

Do not commit to scaling your Pega agentic programme until governance approves it. This is how UAAF proves governance:



8. What You Actually Get, Stated Honestly

PEGA INFINITY '26

Where Pega is stronger than you may have credited it:

The case lifecycle as an auditable, business-readable state machine. Maker-Checker and HITL as native patterns. Outcome measurement directly from case data. Ethical Bias Check as a shipped control.

And in Infinity '26: MCP as a first-class integration surface, A2A for inter-agent delegation, Process Fabric with exposure controls, design-time MCP for your development tools, and 50+ Agent Skills covering most of what you would otherwise build.

UAAF ADDITIVE LAYER

Where UAAF is genuinely additive:

The governance charter with a signature and version history. The non-determinism risk register. Cryptographic proof that what executed matched what was approved. Agent identity lifecycle (outside Pega's scope). Independent behavioral monitoring from outside Pega. Ongoing capability evaluation. The incident taxonomy with regulator timelines. Cross-organisation agent trust. And the evidence repository.

None of these are Pega workflow features. That is not a weakness. It is a boundary. Pega is excellent at what Pega owns. UAAF is what you own.

8. What You Actually Get, Stated Honestly

PEGA INFINITY '26

WHERE PEGA IS STRONGER THAN YOU MAY HAVE CREDITED IT:

IN BUILT-IN FEATURES:

- Case lifecycle as an auditable, business-readable state machine.
- Maker-Checker and HITL as native patterns.
- Outcome measurement directly from case data.
- Ethical Bias Check as a shipped control.

AND IN INFINITY '26 ADDITIONS:

- MCP as a first-class integration surface.
- A2A for inter-agent delegation.
- Process Fabric with exposure controls, and development tools.
- 50+ Agent Skills covering most of what you would otherwise build.

UAAF ADDITIVE LAYER


WHERE UAAF IS GENUINELY ADDITIVE:

- The governance charter with a signature and version history.
- The non-determinism risk register.
- Cryptographic proof that what executed matched what was approved.
- Agent identity lifecycle (outside Pega's scope+).
- Agent identity lifecycle (outside Pega's scope).
- Independent behavioral monitoring from outside Pega.
- Ongoing capability evaluation.
- The incident taxonomy with regulator timelines.
- Cross-organisation agent trust.
- And the evidence repository.

None of these are Pega workflow features. That is not a weakness. It is a boundary. Pega is excellent at what Pega owns. UAAF is what you own.

9. Five Things to Verify About Your Pega '26

Before committing to agentic scale:



The Next Step



One regulated Pega process. One Pega approval gate. Ninety minutes.

Pick a Pega process where agentic roadmap has stalled in governance (not in Pega implementation). We run a governance readiness review with three stakeholders, verify your Pega '26 against governance requirements, and produce a mission map your Pega team consumes directly for Blueprint and Authoring.

If governance says: stay deterministic

We will say so – and you avoid a Pega pilot that would have stalled in risk. No wasted build budget.

If governance approves: agentic deployment

UAAF governance infrastructure becomes the second work-stream alongside your Pega build. You ship both together.