

— ENTERPRISE AI · CONTROL PLANE

The **AI control plane** for governed enterprise execution.

Operationalize AI across the systems that already run your business — orchestrating agents, knowledge, approvals, and observability, without rip-and-replace.

● No rip-and-replace

● API-first

● Human-governed

● Observable by design

— THE REAL CONSTRAINT

AI pilots are everywhere. Governed execution is rare.

The model layer is no longer the constraint. The hard part is connecting AI to real workflows, system context, approvals, and controls.

CONTEXT

Fragmented context

Knowledge lives across apps, documents, tickets, data platforms, and teams.

WORKFLOW

Manual workflow debt

High-value work still depends on emails, spreadsheets, handoffs, and approvals.

LATENCY

Decision latency

Insights arrive late, lack context, or never connect to a clear action.

GOVERNANCE

Governance friction

Security, compliance, and audit needs keep pilots from scaling to production.

Envision closes the gap between AI capability and governed enterprise execution.

Modernize workflows without replacing core systems.

Core systems are mission-critical. Envision adds intelligence above them — preserving your architecture, security, and controls.

AI-assisted workflows

Triage · Review · Approve · Execute · Monitor

Envision governed execution layer

Connect · Orchestrate · Govern · Observe · Learn

Existing systems of record

ERP · CRM · ITSM · EHR · CLM · DMS · Data Warehouse · Custom Apps

Connect

APIs, events, documents, data.

Orchestrate

Agents coordinate across systems.

Govern

Identity, policy, approvals, audit.

Learn

Feedback improves reusable patterns.

Three product engines. One enterprise AI platform.

Agent Studio

Governed multi-agent orchestration

Design, deploy, and govern agents that coordinate tasks, approvals, and system actions.

Automate

Knowledge AI

Trusted enterprise knowledge layer

Turn documents, records, policies, and tickets into access-aware contextual intelligence.

Inform

PredictIQ

Operational decision intelligence

Score risk, urgency, SLA exposure, priority, and outcome likelihood from workflow context.

Predict

Agents know what to do. Knowledge AI knows where to look. PredictIQ knows what to prioritize.

— HOW IT WORKS

From signal to governed action.

01

Signal

Work enters

02

Context

Evidence retrieved

03

Reason

Recommendation

04

Approve

Human gate

05

Execute

Bounded action

06

Observe

Trace captured

Runtime controls: Identity · Policy enforcement · Tool permissions · Approval gates · Audit logging · Telemetry · Model routing

Workflow-safe

Recommendations stay tied to business processes and approval rules.


System-aware

Actions map to APIs, records, permissions, and operational boundaries.

Evidence-backed

Every recommendation cites retrieved context, confidence, and trace.

Governed AI execution, understandable in 30 seconds.



Runtime control makes AI deployable.



Governance is the production safety layer that lets AI operate inside enterprise workflows.

Repeatable patterns for real enterprise workflows.

IT incident management

- Retrieve logs + runbooks
- Classify severity + owner
- Score SLA exposure
- Human approves remediation

MTTR · triage time · first-touch

Claims & revenue cycle

- Check EHR + payer rules
- Score denial risk
- Recommend evidence
- Human approves exceptions

clean-claim rate · days in AR · manual review

Compliance & contract review

- Retrieve clauses + policy
- Flag deviations + risk
- Score exceptions
- Legal approves

review cycle · policy consistency · audit

Illustrative ranges, validated against your baseline during the pilot.

— OBSERVABILITY

Observable by design.

AI-assisted workflows must be measurable, monitorable, auditable, and reliable.

Throughput

workflow volume

Approval latency

time to gate

Override rate

trend

Audit events

captured

Runtime trace: Request → Retrieval → Model route → Agent plan → Approval → Write-back → Audit

Signals monitored: drift · confidence decay · latency spikes · approval bottlenecks · exception growth · adoption · SLA risk

— DESIGNED TO FIT

API-first, deployment-flexible, security-boundary-aware.

Integration-first

APIs, events, documents, approved connectors.

Deployment flexibility

Cloud, private cloud, hybrid, controlled environments.

Security boundaries

Identity, access, encryption, approved data movement.

Fail-safe operations

Read-only → approved → low-risk automation only.

Operational reliability

Monitoring, fallbacks, versioning, runbooks.

Reusable patterns

Workflows, action libraries, governance templates.

STAGE 1

Read-only assist

STAGE 2

Human-approved action

STAGE 3

Controlled execution

STAGE 4

Scale patterns

— WHY ENVISION IS DIFFERENT

Beyond copilots and model wrappers.

Generic copilots	Envision AI
Help inside one application	Orchestrates across systems
Limited workflow execution	Connects context, agents, approvals & actions
Weak cross-system context	Uses governed enterprise knowledge
Limited observability	Workflow telemetry & audit evidence
Limited approval architecture	Embedded human approval gates
A better prompt is easy to copy. Governed execution across systems, policies, approvals, and telemetry is not.	

We prove it in your environment — before you scale.

Envision is early by design. Instead of slideware, we run a governed 90-day value proof on 2–3 of your real workflows and measure against your own baseline. The proof is the pitch.

Founding design partners

A small first cohort co-builds the playbooks — roadmap influence + preferred terms.

Measured, not claimed

Outcomes validated against your pre-pilot baseline; ranges are illustrative until proven.

Low-risk by construction

Read-only first, approved actions next, scale only what evidence justifies.

Built for the controls enterprise security teams require.

Identity & access

SSO, RBAC, least-privilege tool permissions, scoped credentials.

Data boundaries

Your data stays in your boundary. Not used to train models. Encrypted.

Compliance posture

Designed toward SOC 2, GDPR, and HIPAA controls; audit-ready evidence.

Human-in-the-loop

Material actions require explicit approval; nothing high-risk runs alone.

Full auditability

Every retrieval, decision, approval, and action is traced and logged.

Deployment choice

Cloud, private cloud, hybrid, or on-prem for regulated environments.

— PROVE VALUE IN 90 DAYS

A practical, low-disruption pilot model.

WEEKS 1–2

Select & baseline

Choose 2–3 workflows, owners, data sources, approval rules, baseline KPIs.

WEEKS 3–5

Connect & configure

Integrate systems + knowledge; configure retrieval, agents, policies, telemetry.

WEEKS 6–9

Run AI-assisted

Operate with human oversight; capture approvals, overrides, latency, outcomes.

WEEKS 10–12

Validate & scale

Compare to baseline; confirm production-readiness; build the scale roadmap.

Start with the workflows slowing execution today.

Plan a 90-day value proof → suman@quadframe.work