Pyrana

PYRANA OVERVIEW

Production-grade AI Applications, not chat boxes.

Built around your processes, informed by your knowledge, controlled by the platform.

November 2026

pyrana.ai

PYRANA · THREE POINTS OF VIEW

Unlock AI transformation on a platform purpose-built for enterprise.
The speed and power of AI on a foundation you can trust.

01

The chat box is the wrong interface for most enterprise work.

Companies run on processes. Outcomes have to be consistent and repeatable whoever the user is, and user sophistication varies widely across the workforce. A process needs an application built around it, not a prompt.

→ The example · Where it fits

02

Build AI-native applications on one platform.

Not hundreds of SaaS tools traded for as many vertical AI tools. Not every solution vibe-coded or assembled from components. One platform where guardrails, reliability, supportability, scalability, explainability and auditability are already built, so no team rebuilds the foundation.

→ Decision frame · Platform · Checks

03

Output you can trust needs knowledge you manage.

Good AI output depends on good, governed context. Every decision, action and response traces back to the data and knowledge it considered, and to when that knowledge was true.

→ cortIQ

THE SHIFT

Software can now do the work.
The economics of every function change.

Systems of record

Keep track. People do the work.

Copilots

Help one person at a time. Results depend on who asks.

Agentic applications

Do the work, in an experience designed for your users.



THE PRIZE

10 to 25%

EBITDA gains from scaling single-task AI

Bain, 2025

30 to 50%

faster business processes

BCG, 2025

79%

lower cost per invoice at the best accounts payable teams

Ardent Partners, 2026

WHY IT IS STUCK

Most AI never reaches production.
The demo hides the hard part.

95%

of AI pilots show no measurable P&L impact

MIT NANDA, 2025

40%+

of agentic AI projects canceled by 2027

Gartner, 2025

21%

have mature governance for AI agents

Deloitte, 2026

An agent, a prompt, a happy path

WHAT THE DEMO SHOWS	WHAT PRODUCTION NEEDS
	Permissions, to the row — An agent sees only what the person can.
	Approval on the exact action — The person signs off on what will run.
	A record before every action — Who asked, who approved, which rule.
	Runs that survive failure — A restart resumes. It does not repeat.
	Knowledge that stays true — Policies change, and answers follow.
	Versions you can roll back — One step back when something breaks.

That part is rebuilt for every project, or never built. And when AI is a chat box, results depend on who is typing.

THE EXAMPLE

People should make decisions, not write prompts.
Put the AI behind the scenes in an interface designed for your users.

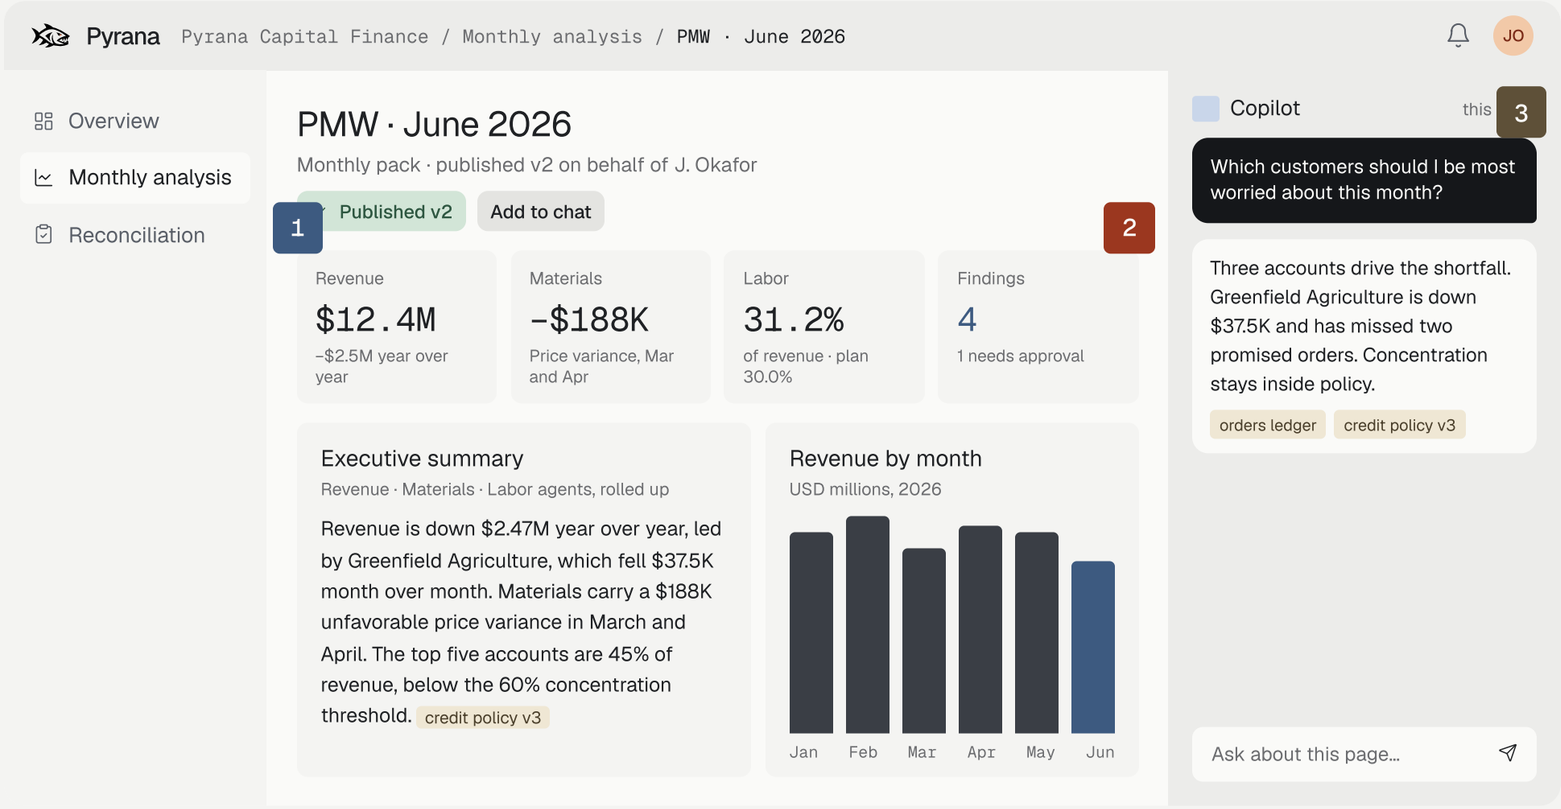

1 The work, already done

Revenue, materials, labor and findings for one entity, analyzed overnight.

2 One decision waiting

The finding that needs approval goes to the controller, with the evidence.

3 A copilot beside it

Questions about what is on the page. Every figure shows where it came from.

Illustrative · Pyrana Capital is a synthetic portfolio

PROOF IN A DEMO

Three questions to ask any vendor. Ask us too.

1 **What happens when it tries something your policy does not allow?**

LOOK FOR It stops and waits for the approver your policy names.

2 **Whose permissions did it use?**

LOOK FOR A record of who asked, who approved and which rule applied.

3 **What knowledge was used to make this decision?**

LOOK FOR Time-stamped knowledge attribution: the exact rule and source, as they stood when it ran.

Governance is not a dashboard. It is what happens on the next action.

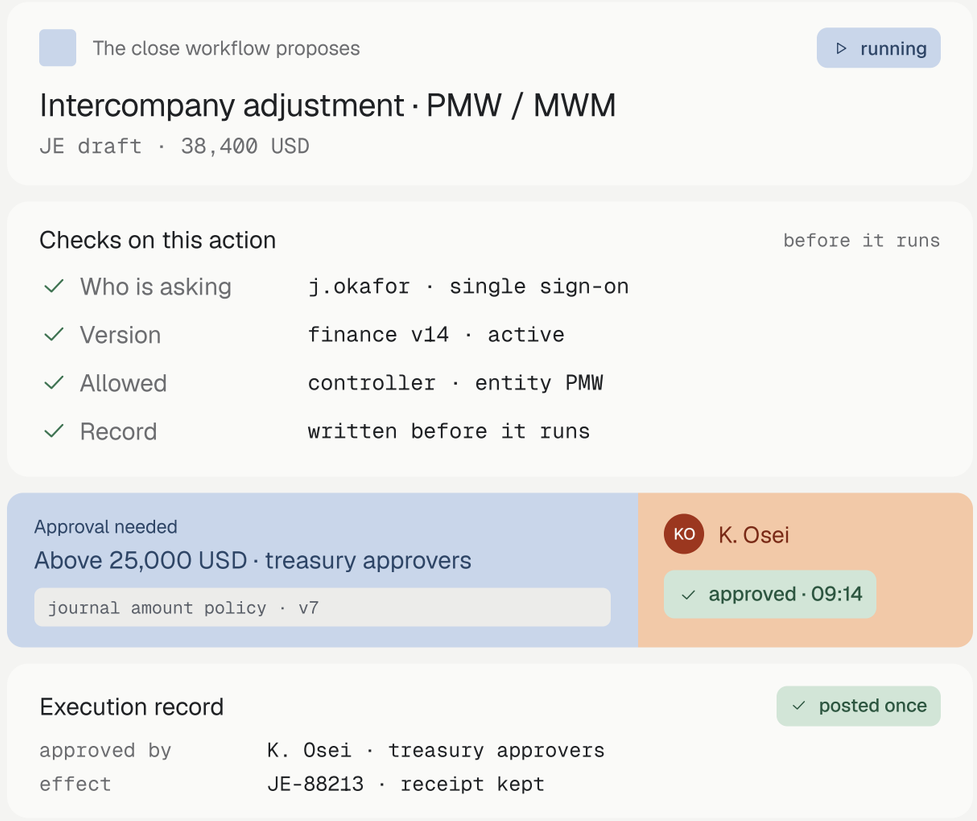

WHERE IT FITS

Copilots make people faster. They don't change how the business runs.
The processes that matter most need AI-native applications.

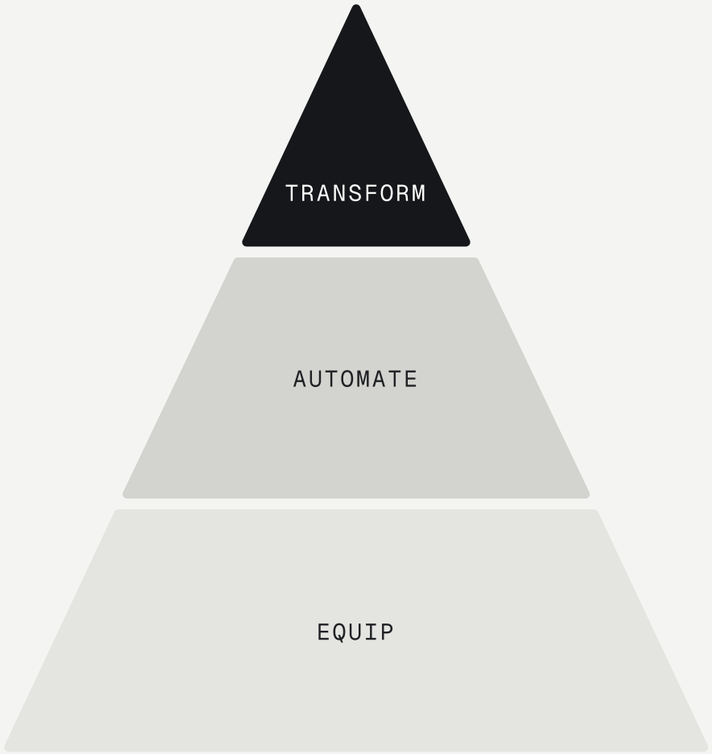

Tier 3 · Transform

AI-native applications for the work that matters most. This is Pyrana.

Regulated | Differentiating | Cross-system | High-consequence

Tier 2 · Automate

Copilot Studio, n8n and in-app agents take steps out of existing processes. Keep them.

Tier 1 · Equip

Copilots for every knowledge worker. Everyone should have one.

- Keep tiers 1 and 2 on the tools you have. Put tier 3 on one governed platform.

THE MARKET IN ONE PICTURE

A new vehicle for every trip does not scale. Enterprises need a fleet.



CHAT AGENTS ON YOUR SYSTEMS

A hired driver

For each car you already own. The route depends on whoever gives directions.

Process unchanged



BUILD FROM COMPONENTS

A kit car

Your engineers assemble every vehicle, and redo the foundation each time.

You have to assemble



VERTICAL AI TOOLS

A train

Fast between two fixed stations. Every new route needs its own track.

A different tool for each job



BUILD ON A PLATFORM

A fleet with dispatch

Vehicles, routes, licensing and dispatch, already in place. A new route is a schedule, not a new vehicle.

One foundation, many solutions

Pyrana is the fleet. Each new application is a new route, not a new vehicle.

THE DECISION FRAME

Don't trade SaaS sprawl for AI sprawl.
Build what your business needs on one platform.

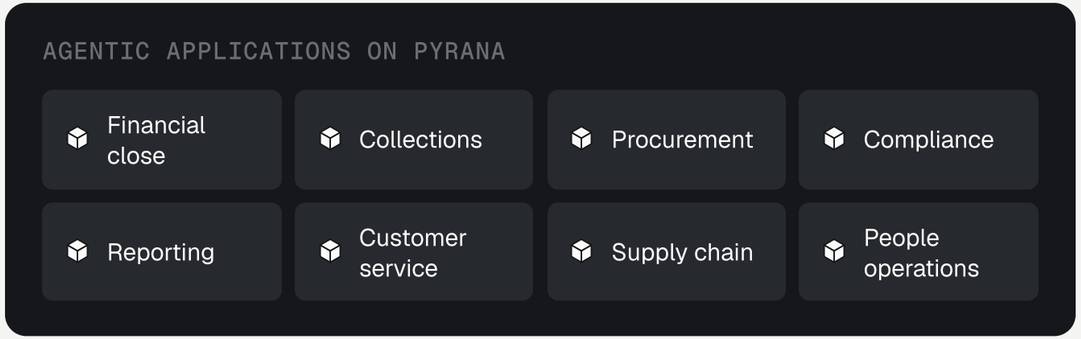

A FEW VERTICAL TOOLS

Only where a vendor owns knowledge you cannot reproduce.

DATA LAYER · BOUGHT	Fabric · Databricks · Snowflake	Read live, as the user
SYSTEMS OF RECORD · BOUGHT	ERP · HR · CRM	A few, with AI built in

WHY THE OLD ANSWER STOPPED WORKING

305

SaaS applications at the average company, $55.7M a year

Zylo, 2026

2,191

applications at the average large enterprise, 61% outside IT

Torii, 2026

45%

of AI code-generation tests produced a security flaw

Veracode, 2025

- Software is not dead. Buying a separate application for every workflow is.

THE ANSWER

Your teams should build business applications, not AI infrastructure.
The foundation is the hard part, and it is already done.

BUILDING IT YOURSELF

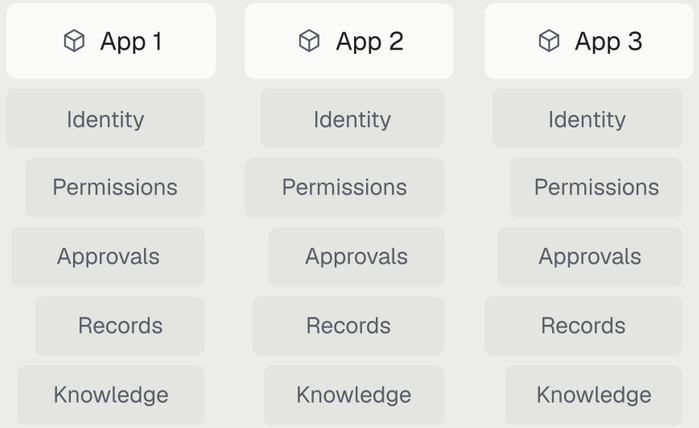

Every application rebuilds the foundation, each one a little differently.

WITH PYRANA

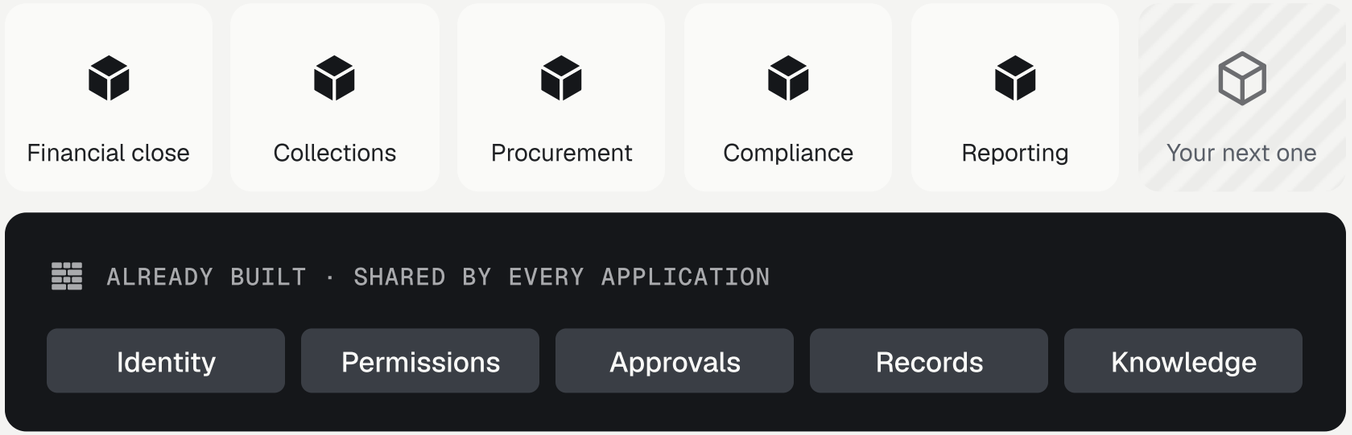

Each application is configuration on a foundation that already exists. None of them starts from zero.

MODULE A

Platform depth

How Pyrana works, how it compares, and what it works alongside.

ARCHITECTURE

Production-grade AI-native apps built for your business.

Pyrana delivers interfaces, workflows, agents, context, storage, and controls via configuration, not code.

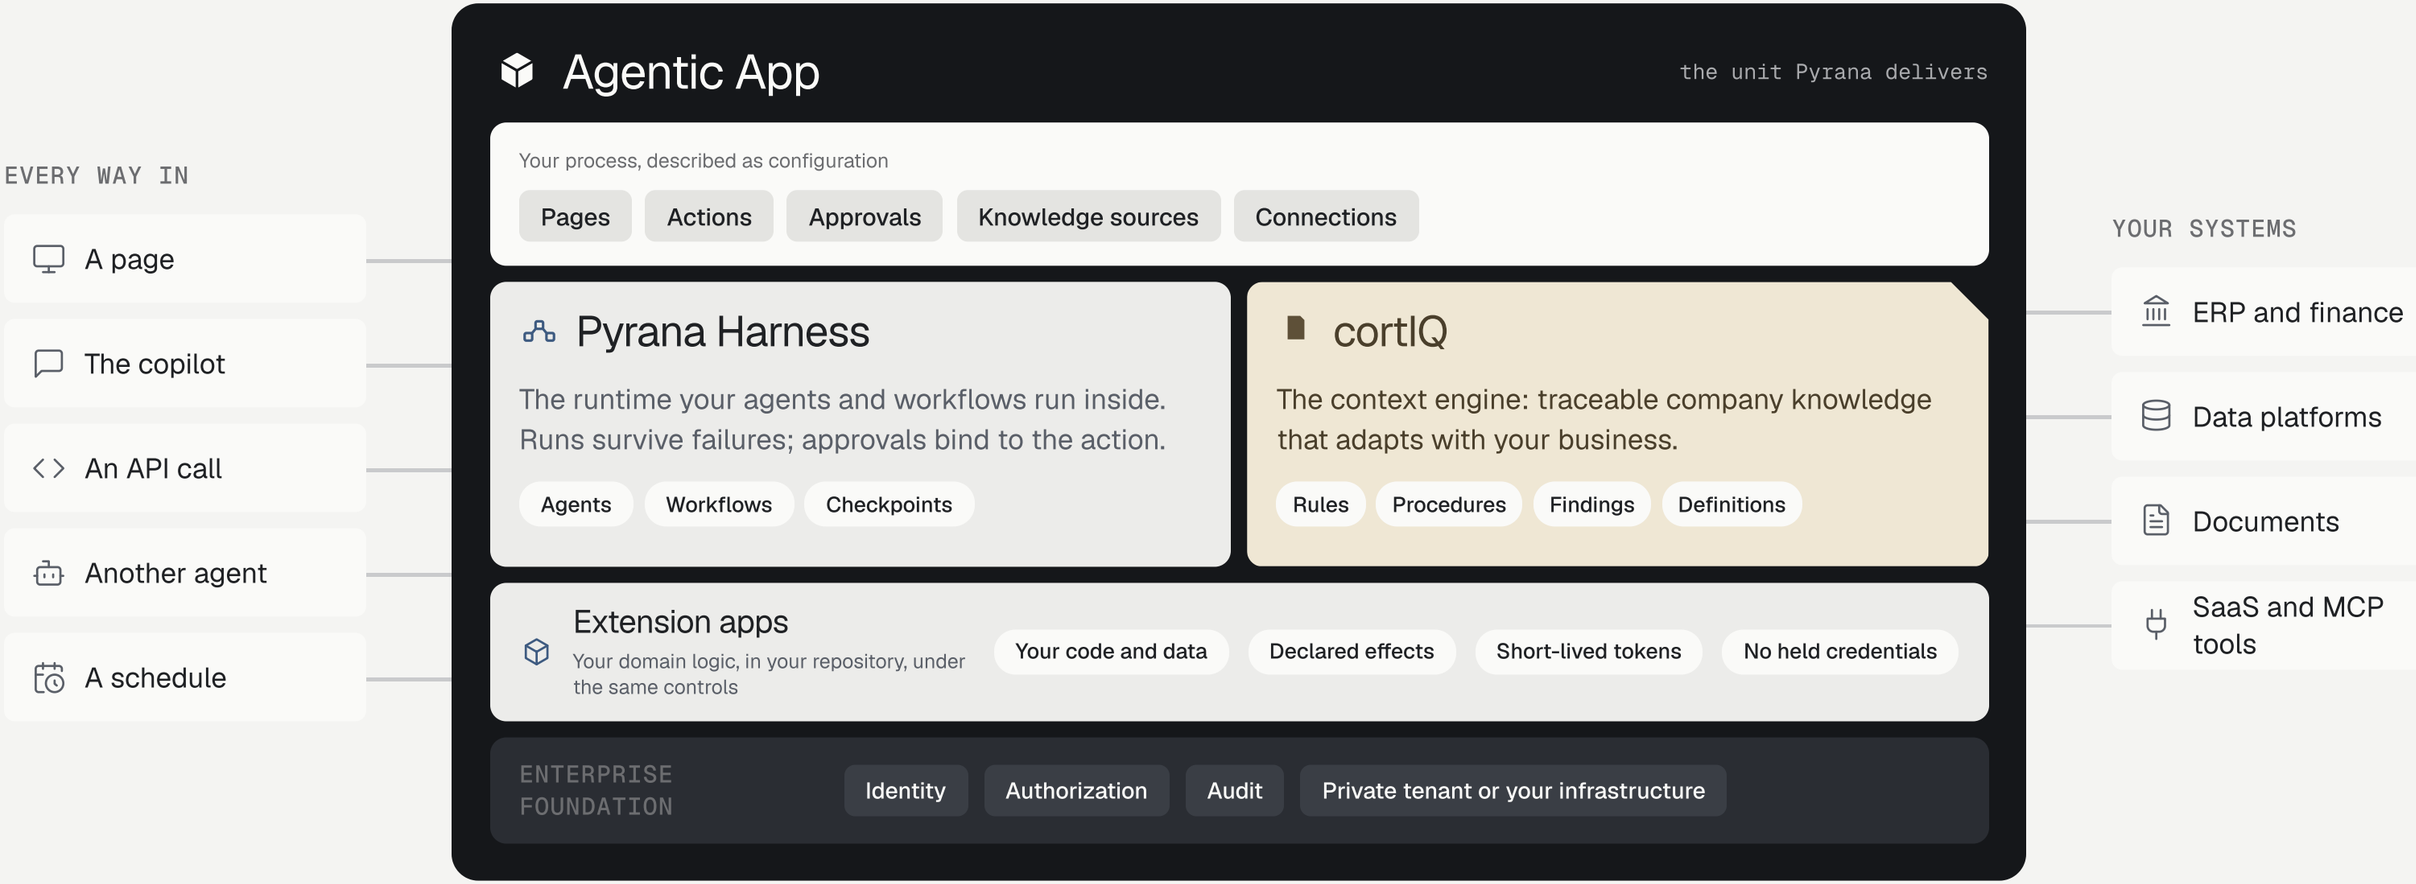

Already built in every application: one identity, one authorization model, one record, one knowledge layer.

GOVERNANCE IN EVERY ACTION

Governance is enforced by the platform, not requested in a prompt. Every action is checked, approved and recorded before it runs.

01

Who is asking?

An identity for every person and agent

02

Is it allowed?

Checked against the version and the person’s access

03

Who signs off?

The exact action waits for the named approver

04

Record it first

Written before anything runs

05

Then act, once

The approved action runs exactly once

Agents belong to applications.

They inherit the application’s rules and the authority of whoever started the work, and nothing more. What a person cannot see is simply absent.

One version at a time

Switch or roll back in one step.

Your code stays yours

In your repository, under the same rules.

PYRANA HARNESS

Enterprise workflows need a harness built for enterprise reality.

Pyrana Harness brings the guardrails, resilience and flexibility that make AI work you can rely on.

ONE RUN, START TO FINISH

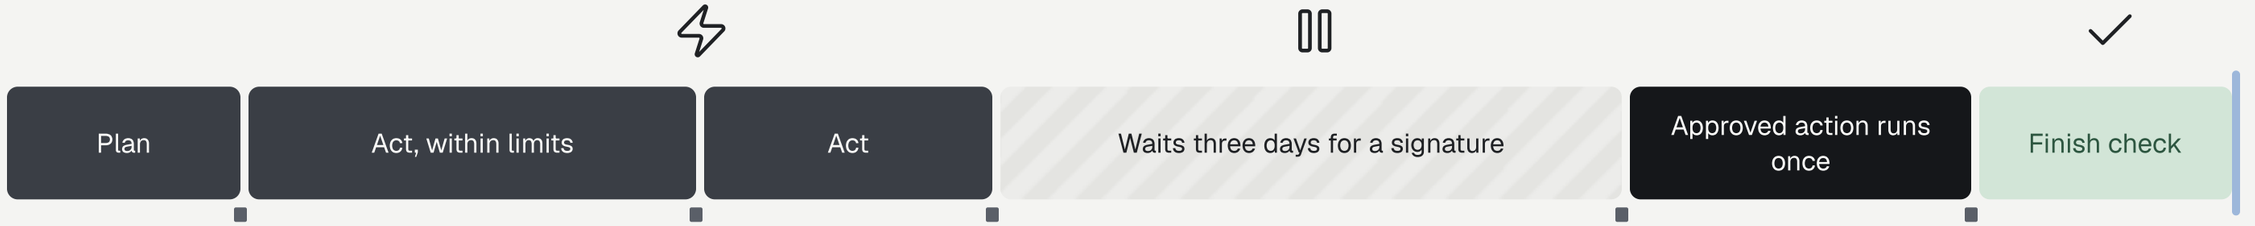

CHECKPOINT

Worker restarts
Resumes from the last checkpoint

No process held open
Pauses, then resumes on the decision

Done on evidence
Not when the model stops

Actions stay bounded

Every tool call checked against permissions and policy, with limits on steps and writes.

A model per agent

Any major provider, your own keys, never switched mid-run.

Five ways to run

On demand, by an agent, as a workflow, in batch or on an event.

Fan-out

Work handed to specialized agents, their results collected and checked.

CORTIQ · THE CONTEXT ENGINE

AI is only as trustworthy as the knowledge behind it.
Manage the knowledge, and every answer can be traced and defended.

ONE CONTEXT UNIT · ILLUSTRATIVE

Each manufacturing deviation gets an initial classification within 3 business days of discovery.

> "The Quality Unit shall assign an initial classification to each deviation within three (3) business days of discovery."
>
> Deviation Management Procedure · rev 4 · effective 17 Aug 2026

Scope Owned sites and contract manufacturing sites

REVISED, NEVER OVERWRITTEN

rev 2 · 5 days → rev 3 · 3 days → **rev 4 · current**

A decision made in February is judged by February's rule.

WHAT ONE RUN USED

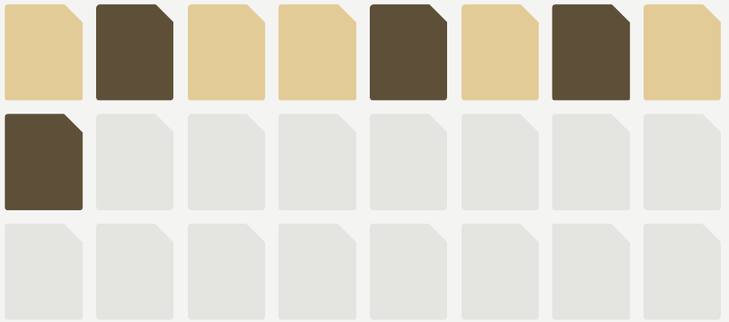

24 found

9 given to the agent

4 cited in the answer

Recorded on every run. A citation the agent was never given is caught by the record.

Better context beats more context. Provenance is a record, not a claim.

ARCHETYPES FOR AGENTIC APPLICATIONS

Faster and cheaper to build, cheaper and safer to run.
The gap widens with each additional application.

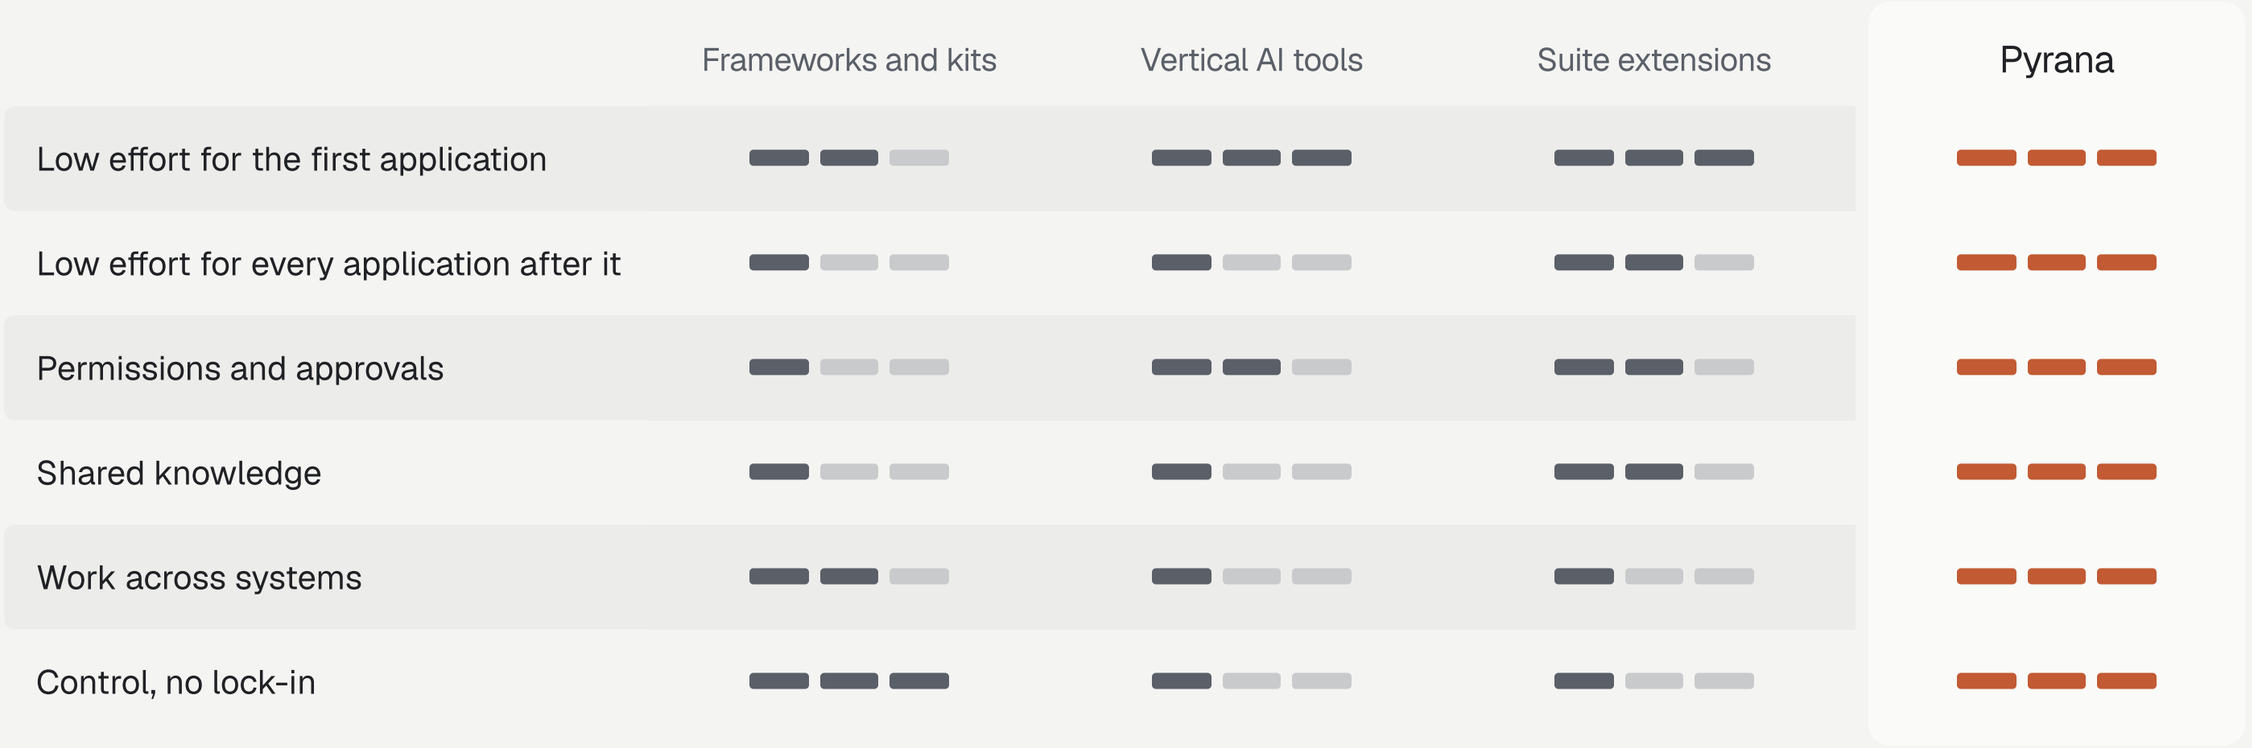

Frameworks and kits: LangGraph, AgentCore, Foundry, built in-house or by a consultancy · Vertical AI tools: Harvey, Sierra, Jasper · Suite extensions: Copilot Studio, Agentforce, ServiceNow

AND, NOT OR

Your data and AI investments keep paying off with Pyrana in the stack.

It connects to your data, runs on your LLM contracts, and reports into the control planes you already use.

Data platforms

Fabric, Databricks, Snowflake. Read live, as the user.

Agent control planes

Agent 365 and AI Control Tower keep the inventory.

Copilots

Your choice: Microsoft 365 Copilot, Claude Cowork, ChatGPT, Gemini.



Any model

OpenAI, Anthropic, Google, Azure, on your keys.

Temporal

Durable execution underneath the Harness.

MCP, both ways

Every application is an MCP server and uses MCP tools.

- Pyrana adds the application layer that turns what you have into governed work.

THE SHORT ANSWERS

The CIO's questions have short answers.
Hosting, data, security, models and lock-in are already settled.

Where does it run?
SaaS on a dedicated private tenant, or deployed into your own infrastructure. The same platform either way.

Who builds it?
Pyrana, third party service providers, or your own team.

Where does our data go?
Nowhere new. Agents query it live, as the user, and it never trains a model.

How fast?
About a week to configure with clean data. About six weeks to the first application in production.

Are we locked in?
Applications are readable configuration, the knowledge is yours, and your code stays in your repository.

What about security review?
SOC 2 aligned. Policies, test results and subprocessors shared on request.

Which models?
Any major provider, on your own keys, assigned per agent.

What happens after go-live?
Your team changes rules and thresholds by configuration. Versions roll back in one step.

MODULE B

Audience modules

One slide each for finance leaders, delivery partners, investors and analysts.

FOR FINANCE LEADERS

The monthly review becomes a proof.
Every run checked against your rules. You decide only the exceptions.

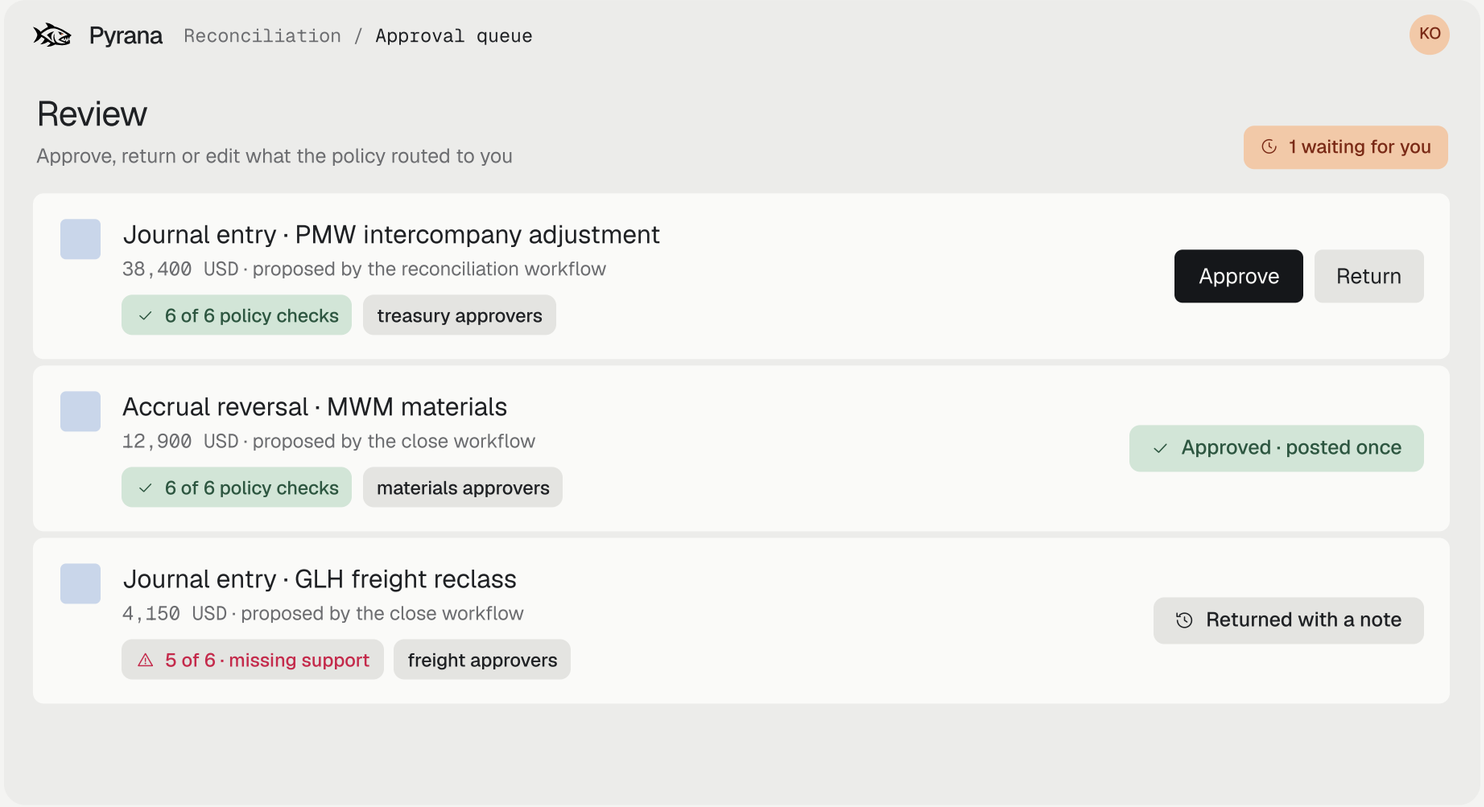

1 **Proposed**

A 38,400 USD entry, drafted by the close workflow

2 **Held**

Above 25,000 USD, it waits for the approver the policy names

3 **Approved**

K. Osei approves the exact entry, unchanged

4 **Posted once**

The approved entry runs once, and the record is kept

The record shows every run stayed inside your rules. You review only what waits for you.

FOR DELIVERY PARTNERS

The foundation is already built. You deliver what sits on top.

BUILT BY

Pyrana

Pyrana’s own delivery team, on the platform.

BUILT BY

Delivery partners

Third party service providers like you, enabled and supported by Pyrana.

BUILT BY

Your client’s team

The client’s own people describe an application and receive it.

WHAT YOU SKIP · ALREADY IN THE PLATFORM

- ✓ Identity and single sign-on
- ✓ Approvals on the exact action
- ✓ Runs that survive failure
- ✓ Permissions to the row
- ✓ A record of every action
- ✓ A shared knowledge layer

WHAT YOU SELL

Governed applications, delivered in weeks.

HOW YOU WIN

Bid against larger firms at a lower price, with a far lower cost of delivery.

No program fee | No referral fee | No exclusivity

FOR INVESTORS

Any lab can build a model.
Pyrana turns it into applications customers cannot switch off.

Models commoditize. The durable value is the application layer and the knowledge under it, already built and reused by every application a customer runs.

THE MARKET

Mid-market, $100M to $2B revenue, multi-site, often PE-backed. Under audit and margin pressure, short on talent to build.

WHO BUILDS

Pyrana, third party service providers, or the customer's own team.

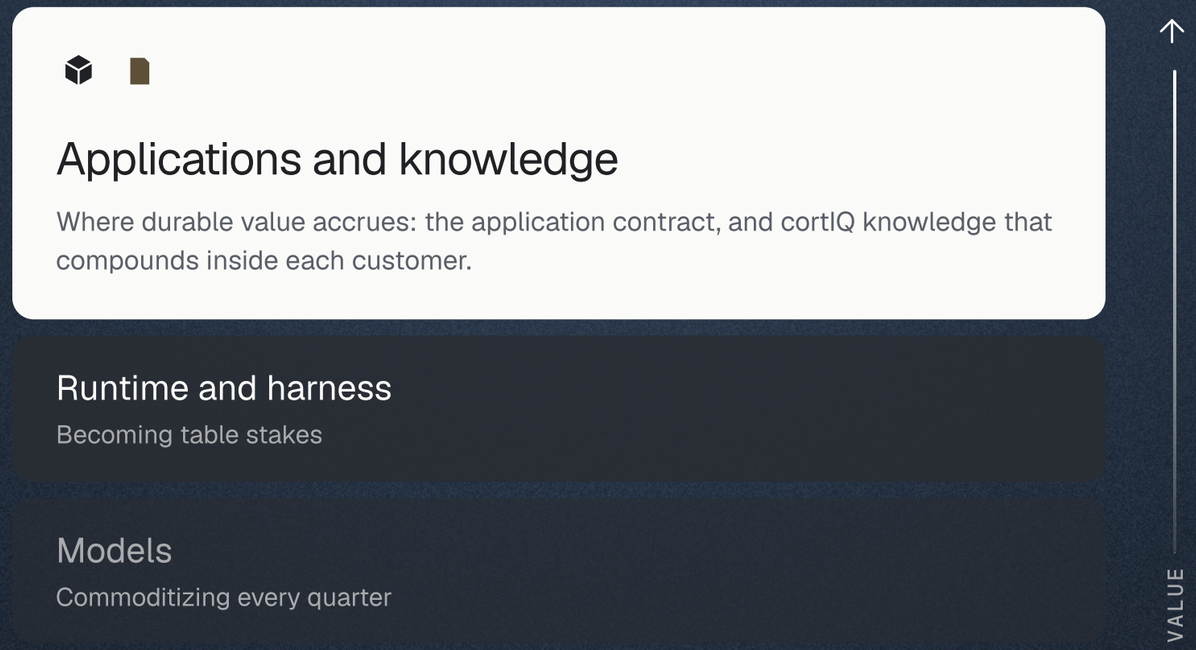

FOR ANALYSTS

Build and run are table stakes. Govern is the proof.

Application platform — Pyrana
Who may use it, on what, with whose approval, and what is recorded

Control plane — Agent 365, AI Control Tower
The inventory and policy over agents

Runtime — AgentCore, managed agents
Where the agent runs

Framework — LangGraph, CrewAI, Agents SDK
How the agent thinks

WHERE TO LOOK

`pyrana.ai/learn`
Category definitions and the layer taxonomy

`pyrana.ai/compare`
Twenty vendor comparisons with sourced detail

`pyrana.ai/glossary`
The vocabulary, in plain language

`pyrana.ai/resources`
Research on harnesses, approvals, knowledge and build vs buy

- Governance is in the application, not in a dashboard.

NEXT STEP

Start with one process. Prove it in production, then scale.

Six weeks to the first application; the second starts on the same foundation.

Day 1

A 45-minute working session on the process

Day 2

A written plan for the first application

Week 1

The application configured, with clean data

Week 6

The first application in production

Next

The second starts with the foundation in place

START WITH A PROCESS THAT IS

Regulated · Differentiating · Cross-system · Living in spreadsheets

Have a stalled pilot?

Bring it. Pyrana rebuilds it on the platform and shows the difference.