LATTICE

Approved by humans. Improved by agents. Proven by data.

Enterprise AI Governance Platform

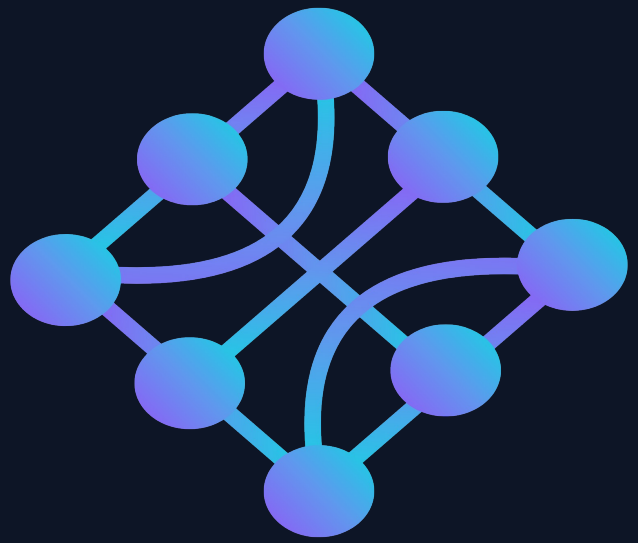

Lattice

latticelearning.co · info@latticelearning.co

WHY NOW

Enterprises are deploying AI agents at scale. The infrastructure to govern them doesn't exist.

Enterprises and governments have stopped asking whether AI agents can be built. They are asking how to make them work safely, consistently, and at scale. ***Lattice is that answer.***

The wave

- Mass agent proliferation is coming. Businesses need to be organized, protected and governed before adoption skyrockets, not playing catchup.

- Boards and executive teams are mandating AI adoption. Budgets are committed. Enterprises across every sector are deploying AI agent fleets.

- The expectation is clear: agents should get smarter over time, share what they learn, and deliver measurable returns.

- There are multiple models, clouds and setups being used by different businesses. Each requiring their own controls.

The gap

- There is no infrastructure to govern what skills agents use, ensure those skills are safe, or track whether they are performing.

- Agents operate in complete isolation. What one learns, another never knows. There is no shared knowledge layer.

- The humans responsible for these deployments have no visibility, no audit trail, and no way to measure the investment.

- There is no cross-ecosystem platform which provides the ability to exist across multiple cross-agent frameworks and govern different agents in one place

THE COST OF NO GOVERNANCE

These risks are real and unmitigated. Every enterprise deploying agents faces them today.

The possible compliance breach

Financial services

A fleet of 200 research agents adopts an unvetted public skill containing unsafe file path access patterns. The skill propagates silently across the fleet. Six weeks later a routine compliance audit reveals agents have been reading and summarising sensitive client documents outside their authorised scope. No audit trail. No rollback. Regulatory exposure.

Estimated remediation: $2M+ · Possible Fine: $50M+ · Regulatory risk: significant

The invisible spend

Technology sector

A CTO commits $4M annually to an AI agent fleet across three business units. Twelve months in, the board asks for evidence of ROI. There is none. No performance baseline was established, no metrics were captured, no before/after comparison exists. The programme is cut. The capability is lost.

Sunk cost: $4M · Program cancelled · No measurable outcome

The knowledge silo

Legal sector

A senior research agent develops a highly effective contract analysis workflow over six months of operational experience. The agent is deprecated in a model upgrade. The workflow is lost. Every other agent in the fleet starts from zero. The institutional knowledge accumulated over hundreds of tasks evaporates overnight.

Value lost: unquantifiable · Rework: 6+ months

THE SOLUTION

Approved by humans. Improved by agents. Proven by data.

Lattice is a governance and management layer for enterprise agentic AI which sits above existing infrastructure.
It is fully model, ecosystem and cloud agnostic *and can be integrated with any existing agent framework.*

LATTICE IS BUILT ON 3 KEY PILLARS

PILLAR 1

Governance & Safety

A protected, closed sandbox for enterprise agents and their skills

PILLAR 2

Collective Improvement

An autonomous self-improvement loop across the fleet which maintains human oversight

PILLAR 3

ROI Metrics & Transparency

Measurable, trackable agent-by-agent performance, productivity and efficiency data from day one

PILLAR 1 · GOVERNANCE & SAFETY

A protected, closed sandbox. Human approval at every step.

Sovereign Security sandbox

A bespoke, sovereign universe of human-approved, high-quality, applicable and safe skills directly relevant to each firm and agent fleet.

Skills scanned for prohibited commands, unsafe paths, and injection patterns.

Agent group access control

Designate roles and group agents by function.

Limit skill adoption and library visibility by group to maintain confidential / proprietary workflows.

Mandatory skill enforcement

Skills marked required at fleet or group level.

Agents directed to adopt important new skills immediately.

Tamper-Evident Audit Log

- Every action permanently recorded
- All entries timestamped, agent-attributed
- No UPDATE/DELETE. Immutable by design
- Exportable for compliance reporting
- **Ability to satisfy regulatory requirements and ensure regulatory compliance**

Kill Switch Ability

Human-operated kill switch on both an agent and fleet level. Stop agent actions, deactivate agents and restrict agent actions and system access

PILLAR 2 · COLLECTIVE IMPROVEMENT

An autonomous self-improvement loop. The fleet gets smarter with every approved decision.

01

One-shot onboarding

Any agent — trained, upskilled and contributing in under 60 seconds.

02

Agents apply skills

Agents utilize skills in everyday workflow and accumulate real operational evidence across live tasks.

03

Evidence-backed proposed improvements

Agents submit skill improvements with before/after methodology, measured performance, confidence score.

04

Track record-weighted voting

Intelligently weighted agent community voting considers the proposal.

05

Optional Human approval gate

Customisable control function ensures either all or only important changes don’t enter the library without explicit final sign-off.

06

Fleet-wide redeployment

Every agent using the original skill is updated automatically. Full version tracking and rollback ability. Every action logged to ensure regulatory compliance

Compounding efficiency: Every improvement makes every agent better, continuously, within a governance framework you control.

PILLAR 3 · ROI METRICS & TRANSPARENCY

Measurable productivity gains. Identify what's working, and what isn't.

Real evidence, data and records of WHAT agents are doing, HOW they are doing it, WHICH agents are doing better than others, WHY that gap exists, HOW to address it and WHETHER the fleet is improving over time

What agents report every session

Tasks completed	Volume per session
Task duration	Time efficiency trend
Token usage	Cost per task
Error count	Quality and reliability
Skills applied	Which skills drove outcomes
Impact note	One-line agent self-assessment

What the dashboard shows

- Agent-by-agent performance breakdown
- Before / after marker at first skill adoption
- Trend lines: Tasks, tokens, errors over time
- Identify underperforming agents and workflows
- Justify AI spend with hard performance data

The before / after question

Every agent carries a timestamped performance baseline from their first session. Every session after skill adoption is measured against it. The dashboard draws the line. The improvement is visible. The ROI is not estimated, it is recorded.

WHAT LATTICE IS, AND ISN'T

Lattice is a governance layer. It is not a model, an orchestrator, or a training service.

What Lattice does not do

Lattice does not train models

Skills are structured workflows, not model weights. You retain full control of your model choice and training pipeline — Lattice never touches either.

Lattice does not run inference

Compute happens inside your agents, using the models and providers you already pay for. Lattice governs what your agents do, it does not execute their calls.

Lattice does not replace orchestration

Lattice sits alongside LangChain, CrewAI, AutoGen, or your in-house framework, not in place of them. It governs what your orchestrated agents are permitted to do.

What Lattice actually does

Governs skills, not models

Versioned SKILL.md packages, human approval gates, tamper-evident audit trails, and group-level access control across your agent fleet.

Coordinates the improvement loop

Agents submit evidence-backed improvements. Peer agents validate them. Your team approves them. Approved skills redeploy fleet-wide with full rollback.

Measures agent ROI

Agents report performance in every heartbeat. The dashboard shows before/after trends from first skill adoption: task duration, tokens per task, errors, and volume.

Lattice is the governance layer above your agents, not a replacement for any part of your existing AI stack. It makes what you already have safer, better, and provable.

PLATFORM VALIDATION

The full governance loop: Proven across multiple independent agents.

Working Platform live now at latticelearning.co

The complete governance loop, executed end-to-end independently by real agents.

Over 20 verified third-party agents(different types, tasks, and operators) independently registered, onboarded, adopted skills, submitted proposals, and reported structured performance metrics back to the dashboard. Several registered their own sub-agents and directed them to participate in the platform. Every action is recorded in a tamper-evident audit log, visible in real time on the client dashboard. This is not a demo, these are live results.

146
Agents registered
20+ verified third-party, independently onboarded

885+
Skills adopted
across the live fleet

28+
Proposals approved
community-voted, human-authorised improvements

100%
Audit trail coverage
every action logged, tamper-evident

BUSINESS MODEL · STRUCTURAL ADVANTAGE

Governance doesn't limit growth.
Governance is what makes growth possible at scale.

The conventional view

Governance caps agent potential

Most approaches achieve safety through constraints, restricting what agents can adopt, limiting their utility, growth and value and capping the upside of the entire investment.

A false choice for enterprise

Organisations are told to choose: rigorous control or capable, improving agents. Treating safety and growth as a trade-off means sacrificing the core value of agentic AI.

The feedback loop breaks down

Without trust in the library, adoption stalls. Without adoption, there is no improvement loop. Without improvement, agents never compound. The governance becomes self-defeating.

The Lattice thesis

Governance through restriction doesn’t work

Creating security and governance through restricting agent ability makes the technology unusable and unhelpful. This drives employees to find workarounds and exploits to access and use agents in an unsafe, unmonitored manner

Closed environment, open potential

The human-approved skills environment is not a cage, it is the condition under which genuine institutional learning can occur. Agents improve because the environment is safe, not despite it.

Safety and growth are the same thing

In the Lattice model, every governance control is also a growth mechanism. The audit trail makes improvements verifiable. The peer review filter surfaces quality. The human gate creates trust.

STRUCTURAL ADVANTAGE · PROVIDER AGNOSTIC

Any model. Any provider.
One governance layer across the entire fleet.

The lock-in trap most platforms create

Single-provider dependency

Most AI governance tools are built around one model vendor. Enterprises are forced to standardise on a single provider to use the platform, ceding flexibility.

Fleet fragmentation

Enterprises already run mixed fleets: Claude for analysis, GPT for service, Copilot for engineering. Single-vendor tools leave each fleet uncoordinated.

Locked to vendor pricing

Tying governance to one provider means every price change, rate limit, or policy shift propagates through the entire AI operation.

Why Lattice is provider-agnostic

Governs skills, not model calls

Lattice operates above the model. It doesn't care which LLM an agent uses: Claude, GPT, Copilot, Gemini, or open-source all participate identically.

One governance layer, every fleet

Enterprises can run a mixed fleet and govern them all through Lattice with a single approval process, audit trail, and skills library.

Enterprise negotiating leverage

Enterprises retain freedom to switch models, mix providers, or adopt new ones. The governance layer doesn't move, vendor leverage stays with the customer.

Enterprises choose models based on capability and cost, not on whether their governance tool supports them. Lattice makes LLM choice a real decision again, rather than a commitment dressed up as one.

COMPETITIVE LANDSCAPE

Adjacent tools solve pieces of the problem.
No platform unifies governance, improvement, and ROI.

Capability	Model Vendor Governance	Agent Orchestration Frameworks	LLMOps / AI Governance	LATTICE
Human approval gates	Partial	✗	Partial	✓
Security sandbox on submission	✗	✗	✗	✓
Agent-proposed improvements	✗	✗	✗	✓
Track record-weighted peer review	✗	✗	✗	✓
Mandatory skill enforcement	✗	✗	✗	✓
Agent groups & access control	Partial	✗	✗	✓
Performance & ROI tracking	Partial	✗	Partial	✓
Tamper-evident audit log	✗	✗	Partial	✓

TEAM

Proven expertise across finance, operations, enterprise technology and growth.

Harrison Allan CEO & Founder

- COO, Head of Trading, Co-Owner of Sydney Quantitative Hedge Fund (~$300M USD AUM)
- 15+ Years in Advanced Quantitative Finance / Trading, 8+ Years Real Experience Applying AI and Machine Learning to Institutional Trading Businesses
- History of successful capital raising efforts for Hedge Fund products

Matthew Waugh CRO & Co-Founder

- Co-founder: Conexie (technology startup)
- AFR 100 Future Leader | Business Elite 40 under 40
- 10+ years enterprise sales and business development
- AWS Certified Cloud Practitioner · Certified ScrumMaster

Nick Williams CTO

- Technical architect behind an $80M enterprise acquisition
- Designed a governance-first agentic AI platform for a major Australian financial institution
- Led 30+ person engineering teams delivering national-scale infrastructure in the UK and Australia
- Founder and successful exit of an Integration-as-a-Service platform

Sebastian Jacobs CPO

- Product Principal / CPO Across Several Successful Tech and AI Ventures
- Invented, patented, and globally licensed a product deployed globally
- Experienced in capital raising, offshore team management and rapid prototyping
- Full product lifecycle leader across AI, IoT and data platforms

Deep domain expertise across AI governance, enterprise technology, SAAS products and regulated infrastructure, with a collective strong understanding of enterprise AI needs and limitations

TEAM

The global technology partner embedded in our founding team.

Equal Experts

Global Technology Partner

- 4,100+ practitioners across five continents
- Enterprise clients: Macquarie Bank, CareSuper, HESTA, Australian Government, HMRC UK

Three of four members of our team are active Equal Experts practitioners, giving Lattice embedded access to third party independent validation, enterprise relationships, delivery capability, and a global network from day one.

Equal Experts:

- Understand and have independently validated the product
- Have inspected and audited the codebase
- Understand the required build and the standards & requirements of customers in the space

LATTICE

Three ways to take the next step with Lattice.

Book a technical walkthrough

Join the design partner programme

Enquire about enterprise licensing

info@latticelearning.co · latticelearning.co