



HERDING CATS
THE CONTROL PLANE FOR ENTERPRISE AI AGENTS

The Control Plane for Enterprise AI Agents

Autonomy without the chaos.



AI agents are becoming enterprise actors. They can access systems, make decisions and take actions. Herding Cats gives enterprises a vendor-neutral governance layer that determines what each agent can access, what it is permitted to do, and when a human must intervene, before consequential actions execute.

Discover every agent. Govern every action. Trace every decision.

Website

www.herdcats.ai

Interactive Demo

herdingcats.cyberly.com.au/demo

AI is moving from generating information to taking action.

Gartner predicted that **40% of enterprise applications would include task-specific AI agents by the end of 2026, up from less than 5% in 2025** (Gartner, 2025). Agents change the security model. They can now invoke tools, interact with applications, use APIs, access data and execute multi-step workflows with decreasing human involvement.

Three forces are converging:

MORE AUTONOMY

Agents are gaining access to real enterprise systems.

MORE CONSEQUENCE


Agent actions increasingly affect money, data, operations and people.

MORE ACCOUNTABILITY

Enterprises must demonstrate who authorized consequential automated decisions and why.

The missing layer is control over agent actions.


Herding Cats is a control plane sitting between AI agents and the actions they want to take.



AI agents are becoming privileged enterprise actors without control.

An AI agent can act across multiple systems. Yet its identity, authority and accountability are often fragmented across the applications, models and frameworks from which it operates.

An agent might:



Each action may be individually permitted. In sequence, they can create significant financial, security or regulatory consequences. Enterprises increasingly need to answer:

- Should this agent be allowed to take this action, in this context, right now?
- And if not, should it be blocked or should an accountable human decide?

That is the control gap Herding Cats is built to close. The platform enables the following:

Know the agent

- Agent registry tracks agent identity, owner, status and accountability.
- Governance graph maps the models, data, tools, credentials and systems each agent can access.

Control the action

- Governance requirements are converted from documents into machine-enforced executable policies.
- Decisions requiring human oversight are routed to the accountable person.
- For complex decisions, Herding Cats uses a non-deterministic policy engine that evaluates risk and business context to determine whether an action requires governance.

Prove the decision

The complete authorization and decision chain is recorded automatically.

Hold consequential actions at the boundary.

When policy requires human judgement, the action pauses and resumes only after an accountable decision.

HERDING CATS | Interactive product demo | Illustrative environment - Acme Financial Group

HUMAN APPROVALS

Human decisions. Accountable outcomes.

When policy requires human judgement, the action pauses and resumes only after an accountable decision.

APPROVAL REQUEST APR-4417

Enterprise Sales Assistant

Requested by Alex Chen - Account Executive

Within SLA
Waiting 12m of 4h

Customer	Acme Bank	Requested action	22% Discount
Contract value	\$2.3 Million	Reason	Competitive Renewal
Policy triggered	Discount >15%	Decision required from	Marcus Vance - Sales Director

Risk: Medium

Add a comment for the record (optional)

Demo controls - nothing is saved or transmitted

Approval timeline Live

- Agent requested 22% discount
10:42 - Enterprise Sales Assistant
- Discount Threshold policy triggered
10:42 - Policy engine
- Action held at boundary
10:42 - Runtime enforcement
- Routed to Sales Director
10:43 - Approval routing
- Awaiting human decision
— Marcus Vance
- Quote generation resumes on approval
— Enterprise Sales Assistant

Approver comments

Incumbent has tabled 19% on a like-for-like scope. 22% holds the renewal at three years rather than one.

Alex Chen - Account Executive - 10:44

Margin remains above floor at 22% given the multi-year term. No services


Previous decisions Acme Bank

18% discount - \$1.9M 14 Mar 2026 - M. Vance	Approved
25% discount - \$2.1M 02 Nov 2025 - M. Vance	Rejected
15% discount - \$840K	Approved

Illustrative demo data. Request APR-4417, Acme Bank and all decisions are fictional, temporary and never saved or transmitted.

Scenario

Consider an enterprise sales agent negotiating a customer renewal:



Without Herding Cats

The business either blocks a legitimate exception or relies on informal approval outside the workflow, slowing the sales cycle, creating inconsistent decisions, and weakening the audit trail.

With Herding Cats

The action is evaluated before execution. Because the 22% discount exceeds the policy threshold, Herding Cats automatically routes it to the Sales Director for approval before the offer can proceed.



Adaptive Governance

Routine decisions can follow **deterministic rules**, such as a fixed approval threshold. More complex decisions require context. In these cases, Herding Cats uses a **non-deterministic policy engine** to evaluate multiple business, risk and regulatory factors together and determine the appropriate governance path - including when human oversight is required.

Context-aware decision

Scenario: AI sales agent proposes a 22% discount on a five-year renewal.

Context Considered: contract value, number of users, term length, expected margin, customer history, competitive pressure and existing commercial commitments.

Outcome: Herding Cats evaluates the context to allow, block or escalate the discount for approval.

High-risk decision


Scenario: An AI hiring agent recommends rejecting a candidate because of an employment gap caused by serious illness.

Context Considered: hiring policy, role requirements, anti-discrimination obligations and whether the gap is relevant to job performance.

Outcome: Because the decision may affect the candidate's rights and involves sensitive health information, Herding Cats requires human review before it can proceed.

Know the blast radius before something goes wrong.


Herding Cats will initially target enterprises deploying agents into environments where actions carry meaningful security, financial or regulatory consequences.



When a security team selects an agent in Herding Cats, the Governance Graph immediately surfaces the full scope of that agent's operating environment.


This includes the models it calls, the data sources it accesses, the tools it invokes, the credentials it holds, the policies that govern it, and the enterprise systems it can affect. Each relationship is mapped and visible, not buried in configuration files or scattered across access control lists.

Security teams can immediately see an agent's potential blast radius and investigate what would be exposed if it were compromised, misconfigured or operating outside its intended scope.



Help us build the control layer for enterprise AI agents

As enterprise AI agents proliferate across regulated workflows, agent governance has the potential to become a horizontal control layer across security, risk and compliance. Herding Cats is built to provide that control layer, initially targeting organisations where AI agents are being given authority over consequential enterprise actions. For example:



Our differentiation

Existing security and AI-governance tools address parts of the problem, including identity, credentials, model governance and activity monitoring. Herding Cats brings these controls to the point of action - evaluating a proposed agent action against organisational policy before execution, determining whether it should proceed, be blocked or require human approval, and preserving the resulting authorization trail as audit evidence.

While emerging agent-security platforms such as Noma and Zenity focus on runtime security, access control and agent behaviour, Herding Cats is built around **governance of consequential enterprise decisions** - translating organisational policy into executable controls, applying business and risk context, routing decisions to accountable humans, and preserving the authorization trail as audit evidence.

Help us build the control layer for enterprise AI agents



For pilot and investor enquiries contact us at info@cyberly.com.au

We understand both the rules and the systems.

Herding Cats was founded by a team with deep expertise spanning cybersecurity engineering, legal governance, data architecture, and government operations. This combination - technical depth paired with regulatory fluency - is rare in the AI governance space, and it is the foundation of everything we build.

Anna Ko, Founder

Cybersecurity + Law + Governance

- Government cybersecurity experience across cloud security, engineering, IRAP assessments, governance, compliance, and policy
- Expertise in AI governance and regulations
- Master of Information and Cybersecurity, UC Berkeley
- Master of Laws

James Marshall, Co-Founder

Data + Architecture

- Experience across government in building data platforms, analytics environments, workflow systems, and operational dashboards at scale
- Background in actuarial science, finance and enterprise architecture
- Master of Actuarial Studies
- Master of Applied Finance

Progress

- Conducted research as part of UC Berkeley capstone project. Developed ATLAS (Adaptive Trust Layer for Autonomous Systems) as a precursor to Herding Cats.
- Working Herding Cats interactive prototype.
- Production architecture under development.
- Existing government relationships supporting early customer discovery.
- Accepted into leading AI and cloud startup programs including:
 - NVIDIA Inception,
 - Google for Startups,
 - AWS activate, and
 - Microsoft for Startups Founders Hub.

