AI Control Plane for IT and Security

Why Enterprises must govern AI with a control plane, not point security tools

Executive Summary

Enterprises are entering a new phase of AI adoption. What began as experimentation with chat interfaces has evolved into agentic AI systems that reason, plan, access enterprise data, and execute workflows autonomously. As AI capabilities are embedded into enterprise software, applications increasingly behave like agents.

To support this shift, enterprises are deploying a new layer of AI infrastructure such as model gateways, agent platforms, orchestration layers, and MCP servers, while policy, risk, and accountability remain anchored in legacy Systems of Record.

An AI Control Plane addresses this gap. Instead of blocking AI by default, it governs AI systems by connecting new AI infrastructure to enterprise ground truth, applying policy dynamically, and providing continuous visibility and evidence. This enables IT and Security to accelerate AI adoption safely, rather than slow it down.

FORRESTER®

An agent control plane is one that inventories, governs, orchestrates, and assures heterogeneous AI agents across vendors and domains.

LangGuard is an example of a purpose-built AI Control Plane for IT and Security.

1. What Is an AI Control Plane?

An **AI Control Plane** is a centralized governance layer that operates across AI infrastructure to provide visibility, policy enforcement, and auditability for AI systems, without rewriting applications or intercepting execution by default.

An AI Control Plane governs:

- **Who** AI acts as (identity and credentials)
- **What** AI can access (data, tools, models)
- **Why** AI is acting (intent and task context)
- **How** AI behavior aligns with enterprise policy

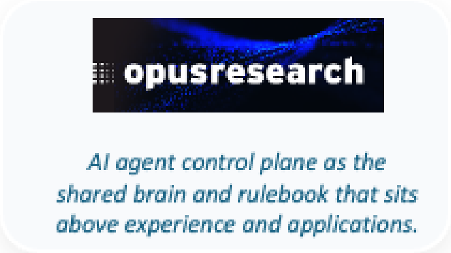

opusresearch

AI agent control plane as the shared brain and rulebook that sits above experience and applications.

Unlike point security products that inspect packets, prompts, or endpoints in isolation, an AI Control Plane enriches those signals with deep context about AI identity, intent, data access, and behavior, so that existing enterprise security controls can detect and respond to AI-centric threats.

2. Why Enterprises Need an AI Control Plane

Enterprise AI introduces a **structural fragmentation problem** that cannot be solved by existing point security and IT tools.

- **AI is fragmented** across models, agent platforms, embedded AI features, and external services, each with different execution paths and behaviors at runtime.
- **Data is fragmented** across data lakes, SaaS applications, APIs, MCP servers, and user environments, and is pulled dynamically at runtime.

An AI Control Plane is designed to connect fragmented AI systems and data into a single, governable system.

2.1 The Rise of New AI Infrastructure

To operationalize AI at scale, enterprises are deploying a new layer of AI infrastructure, including:

- Multiple foundation and fine-tuned models
- Agent frameworks and orchestration platforms
- AI gateways for model access, routing, and cost control
- MCP gateways exposing enterprise tools and data
- Embedded AI across analytics, collaboration, and SaaS platforms

This infrastructure is **dynamic and autonomous by design**. Agents select models, tools, and data at runtime.

2.2 Enterprise Policy and the Governance Gap

While AI infrastructure is evolving rapidly, enterprise policy remains anchored in legacy **Systems of Record**, including:

- Identity and access management systems
- Service catalogs and CMDBs
- Cloud and endpoint security platforms
- Data classification and governance tools
- Incident response and audit systems

These systems were designed to govern **human-driven workflows**, not autonomous AI systems operating across models, tools, and data.

As a result, when AI infrastructure operates independently of enterprise policy systems, a **governance gap** emerges. The seams between AI, data, identity, and policy are left ungoverned.

Without a unifying control layer:

- Visibility becomes partial and inconsistent
- Enforcement becomes reactive or blunt
- Proof and accountability are difficult to establish

AI Control Planes exist to close this gap. They provide the missing layer that connects new AI infrastructure with enterprise Systems of Record, restores end-to-end visibility, and enables **governed autonomy at scale**.

3. What an AI Control Plane Does

An AI Control Plane provides a consistent lifecycle management for AI systems:

- **Discover** AI agents, models, tools, and data access
- **Screen** AI activity against enterprise policy and systems of record
- **Approve** AI usage using evidence-based decisions
- **Provision** scoped identities, credentials, and permissions
- **Monitor** AI behavior and intent continuously
- **Enforce** AI dynamically without breaking workflows
- **Audit** AI actions with enterprise-grade evidence

It is a **control layer for AI as infrastructure**.

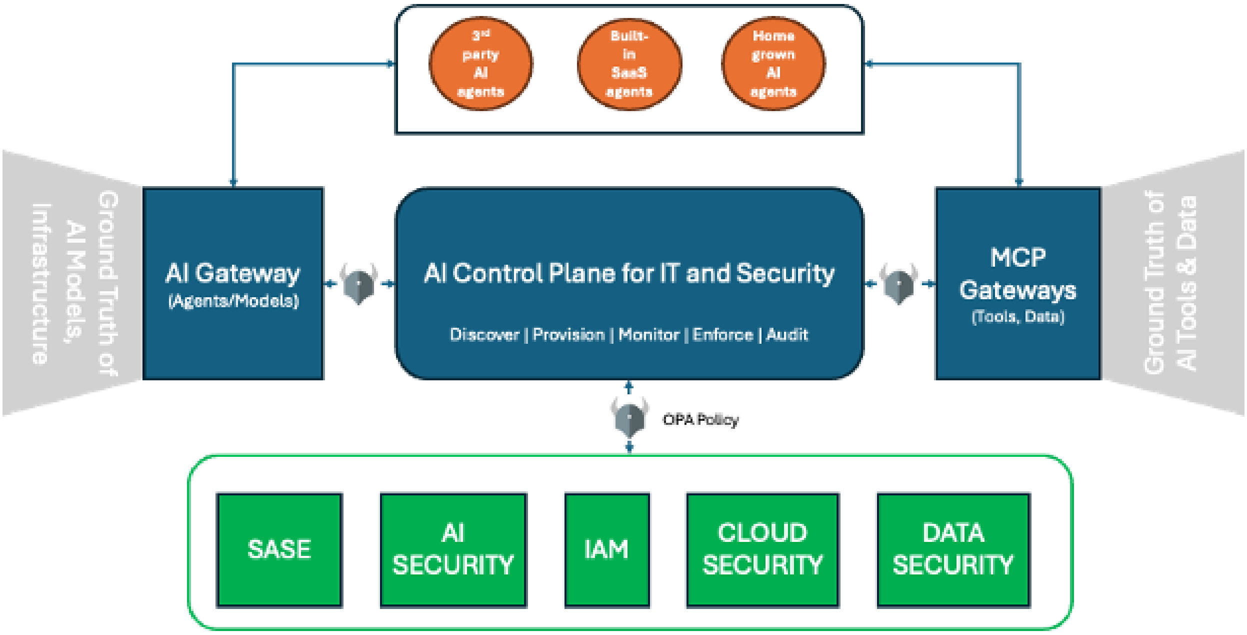

4. Governing AI Without Killing Productivity

AI safety is non-negotiable but **is not the same as blocking**. As enterprises adopt agentic AI, IT and Security face a new challenge: AI systems act autonomously, adapt dynamically, and operate across models, tools, and data.

4.1 Why Block First Breaks Agentic AI

Many security controls were designed for deterministic systems, where intent is explicit and execution paths are predictable. In those environments, block by default is often acceptable. But, Agentic AI behaves differently.

Agentic AI:

- Select tools dynamically
- Pull data contextually
- Chain actions across systems
- Adapt behavior based on goals and outcomes

When controls lack context about AI intent, identity, and authority, they default to blunt enforcement. Blocking an agent often disables the very workflow it was designed to automate, without explaining why or how to remediate the issue.

This does not improve safety. It erodes trust and drives workarounds.

4.2 What IT, Security, and Leadership Actually Need

Boards and executives are not debating whether to deploy AI. They expect IT and Security to:

- Enable safe acceleration
- Make risk visible and manageable
- Explain AI behavior with evidence
- Demonstrate control at scale

AI Control Planes provide the missing layer to do this credibly. They give IT and Security continuous visibility into AI behavior, the ability to adjust policy dynamically, and the evidence needed to balance **AI safety and AI productivity**.

5. Policy as the Control Fabric for Enterprise AI

The core challenge in AI governance is not defining policy but it is the **consistent application of policy across fragmented AI systems**.

Enterprise policies already exist. Identity, access, data handling, risk, and compliance rules are well established across Systems of Record. Agentic AI introduces new actors, behaviors, and execution paths that span models, tools, data, and infrastructure, often outside the direct control of those systems.

An AI Control Plane addresses this challenge by acting as the control fabric that applies enterprise policy **consistently and contextually** across AI infrastructure. For example, LangGuard uses **Open Policy Agent (OPA)** as its policy foundation, enabling:

- A common, open standard for policy definition
- Separation of policy from AI execution
- Extensibility across AI infrastructure and enterprise systems

LangGuard operationalizes OPA specifically for AI gateways, agent platforms, MCP servers, and integrations with existing Systems of Record. At a high level, this policy fabric supports several **types of AI governance controls**:

- **Discovery and Classification Controls**
  Identify and classify AI systems, agents, models, tools, and data sources and establish visibility into sanctioned and shadow AI.

- **Identity and Authority Controls**
  Govern how AI systems act on behalf of users or services, including identity assignment, credential scoping, delegation, and least-privilege enforcement.

- **Data and Context Controls**
  Define what data AI systems may access, under what conditions, and in what context thereby ensuring exposure is intentional, scoped, and auditable.

- **Behavior and Safety Controls**
  Evaluate AI inputs, outputs, and actions for safety risks such as prompt manipulation or unintended behavior.

- **Operational and Cost Controls**
  Manage AI usage, scope, and cost to prevent runaway execution while enabling teams to scale responsibly.

- **Evidence and Integration Controls**
  Export AI behavior, decisions, and evidence into existing security, compliance, audit, and incident response systems.

Conclusion

An AI Control Plane gives enterprises the ability to govern autonomous AI without rewriting systems or sacrificing productivity. It provides continuous visibility into AI behavior and intent, enables Security to manage risk without indiscriminate blocking, and gives leadership evidence they can trust as AI adoption scales. As enterprises move from experimentation to agentic AI operating at scale, governance must evolve beyond detection and shutdowns.

To learn more about LangGuard AI Control Plane, visit https://www.langguard.ai or follow us on LinkedIn