graphiant

Governing AI Starts with Visibility

A Practical Framework for Building an AI-Ready Enterprise

Executive Summary

Artificial intelligence is moving into the enterprise faster than many governance programs can keep pace.

Employees are adopting AI assistants, copilots, generative AI platforms, and AI-enabled applications to summarize information, write code, analyze data, automate tasks, and make decisions. Some of these tools are sanctioned by IT. Others are adopted independently by employees or embedded within applications the organization already uses.

At the same time, AI is changing something more fundamental: how enterprise data is consumed, transformed, copied, and moved.

This creates a governance problem that cannot be addressed through acceptable-use policies, perimeter security, or static compliance checklists alone.

The question is no longer simply:

Which AI applications should our employees be allowed to use?

Organizations increasingly need to answer:

- Which AI applications are actually being used?
- What data is interacting with them?
- Where is that data moving?
- What policies should apply to it?
- Can those policies follow the data as it moves across increasingly distributed infrastructure?
- And can the organization act when AI activity violates those policies?

AI governance is therefore becoming inseparable from data governance.

And effective data governance begins with visibility.

This paper introduces a practical framework for evaluating AI governance across five capabilities:

Visibility → Control → Data Assurance → Compliance → Measurement

The objective is not to stop AI adoption. It is to establish the visibility, assurance, and control organizations need to adopt AI with confidence.

1. AI Has Changed the Governance Problem

AI adoption rarely happens through a single enterprise-wide decision.

It happens incrementally.

A developer begins using an AI coding assistant. A marketing team experiments with a generative AI platform. An analyst uploads information to an AI service to summarize it. A department purchases an AI-enabled SaaS application without involving the security team.

Meanwhile, established enterprise platforms continue adding their own AI capabilities.

The result is an expanding ecosystem of AI applications, models, APIs, integrations, users, data flows, and increasingly, autonomous software capable of acting on behalf of users and other systems.

This creates two interconnected governance challenges.

Application visibility

The organization's official AI environment may be very different from its actual AI environment. This gap contributes to Shadow AI: the use of AI applications, models, or capabilities outside an organization's sanctioned or governed environment.

> But underneath the application problem is a more foundational challenge: data.

AI systems consume, transform, generate, and redistribute data.

Information may move into an AI application, across a third-party provider, through an API, into a model or processing environment, and potentially across cloud regions and jurisdictions.

That means organizations cannot think about AI governance solely in terms of which applications are approved.

They must also understand how data is moving through the AI ecosystem and how the infrastructure underneath it must evolve.

> *"The current Internet was designed to connect users to applications. The AI era requires a new network, one designed to connect distributed intelligence and autonomous software across machines."*

— Khalid Raza, Founder & President, Graphiant

2. AI Governance Is a Data Governance Problem

For years, many organizations have approached governance through a combination of perimeter controls, vendor reviews, security policies, contractual requirements, and periodic compliance assessments.

AI challenges that model.

Data may now move between users, SaaS platforms, AI applications, cloud environments, third-party services, APIs, and AI model pipelines. It may be copied, transformed, summarized, or incorporated into new content.

Yet many organizations are not confident that their data-management foundations are ready for this shift.

The consequences extend beyond governance and security. The readiness of enterprise data can directly affect whether AI initiatives deliver their intended business value.

63%

OF ORGANIZATIONS EITHER DON'T HAVE, OR AREN'T SURE THEY HAVE, THE RIGHT DATA-MANAGEMENT PRACTICES FOR AI

60%

GARTNER PREDICTS THAT THROUGH 2026, ORGANIZATIONS WILL ABANDON 60% OF AI PROJECTS UNSUPPORTED BY AI-READY DATA

Together, these findings point to a broader issue: organizations are under pressure to accelerate AI adoption while many are still building the data foundations required to support it.

AI-ready data is not simply about having enough information available to an AI model. Organizations must also consider how data is governed, observed, classified, accessed, and moved across increasingly distributed environments.

This makes a foundational principle increasingly important:

> **You cannot separate AI governance from the data that powers AI.**

Organizations need to understand not simply whether an AI application is permitted, but what happens when enterprise data interacts with it.

That requires asking:

- What data moved?
- Where did it move?
- Who or what accessed it?
- Which application, model, or agent processed it?
- Which governance policies applied?
- Did the data cross a geographic or jurisdictional boundary?
- Can the organization reconstruct that activity later if necessary?

A governance failure is damaging enough when discovered immediately.

It becomes significantly more difficult when an organization discovers an exposure weeks or months later and cannot determine where affected data traveled, which systems interacted with it, or what policies were in effect.

AI governance therefore requires more than policy intent.

It requires technical assurance.

3. Visibility Is Necessary. But Visibility Alone Is Not Governance.

Consider an organization that has officially approved three AI applications.

Its AI governance policy addresses those three tools.

Employees, however, are actually accessing twelve.

The organization may have a policy.

It does not yet have complete governance.

The same problem exists at the data layer.

An organization may have policies defining where regulated, proprietary, or sensitive information is permitted to travel. But without sufficient visibility into data in motion, it may be difficult to determine whether those policies are being followed as data moves between networks, SaaS APIs, clouds, AI services, and other distributed environments.

This establishes the first principle:

> **Visibility is the prerequisite for control.**

Organizations cannot effectively govern AI applications they cannot identify.

And they cannot effectively govern data movement they cannot observe.

But discovery is only the beginning.

Knowing that an employee, application, or autonomous agent is interacting with an AI service does not by itself prevent inappropriate data access or stop an action that violates organizational policy.

Effective governance requires organizations to move from seeing → understanding → controlling → enforcing.

> *"AI governance can't stop at discovering Shadow AI. Enterprises need to know every AI being used, understand what data it can access, control what it can do, and have a kill switch when it crosses the line. In an agentic world, governance without enforcement isn't governance."*

— Vinay Prabhu, Chief Product Officer, Graphiant

This becomes increasingly important as enterprises move toward agentic AI.

Traditional AI tools largely respond to human prompts. AI agents can potentially interact with systems, call APIs, access data, initiate workflows, and perform actions.

The governance question therefore evolves from:

> **From.** What AI are employees using?
>
> **To.** What AI is operating in our environment, what can it access, what is it allowed to do, and can we stop it when necessary?

That requires more than observability.

It requires enforceable governance.

4. Not All Data Carries the Same Risk

Improving visibility creates another important realization:

Not all enterprise data should be governed the same way.

A publicly available marketing asset does not carry the same risk as intellectual property.

Anonymous operational information does not carry the same requirements as regulated personal data.

An internal presentation may have different requirements from financial records, customer information, source code, or proprietary product designs.

AI governance therefore requires organizations to develop a meaningful data taxonomy.

A meaningful data taxonomy

That taxonomy might distinguish between categories such as:

- Public information
- Internal business information
- Confidential information
- Intellectual property
- Personally identifiable information
- Financial information
- Regulated or jurisdiction-sensitive data
- Strategic operational data

Governance questions become more precise

Organizations can then begin answering more sophisticated governance questions:

- Which information can interact with public AI applications?
- Which information may only interact with sanctioned enterprise AI?
- Which datasets should never leave a particular geographic region?
- Which information should never be submitted to an external AI service?
- What information can an autonomous AI agent access?
- What actions should an agent be permitted to take using that information?
- And what happens when an application or agent violates those policies?

> **AI governance therefore becomes a question not simply of access, but of authority.**

5. A Practical AI Governance Framework

Organizations can approach AI governance through five progressive capabilities.

1. **Visibility.** What AI is being used, and where is data moving?
2. **Control.** What should AI be allowed to access and do?
3. **Data Assurance.** Is data moving as intended?
4. **Compliance and Accountability.** Can the organization prove its policies are working?
5. **Measurement and Optimization.** How is AI actually being consumed?

1. Visibility

What AI is being used, and where is data moving?

Organizations first need an accurate understanding of their AI environment.

Questions to ask:

- Which AI applications are employees accessing?
- Which are sanctioned?

- Where does Shadow AI exist?
- Which teams or locations are driving adoption?
- Which autonomous agents or AI-enabled applications are interacting with enterprise systems?
- How is usage changing?
- Can relevant data flows be observed across distributed infrastructure?

> **Governance principle.** You cannot govern what you cannot see.

2. Control

What should AI be allowed to access and do?

Visibility should inform policy.

Organizations need to establish boundaries around both applications and increasingly autonomous AI systems.

Questions to ask:

- Which AI applications should employees use?
- Which require additional review?
- What categories of data can interact with each?
- What data can an AI agent access?
- What actions can it perform?
- Can risky activity be restricted?
- Can access or activity be terminated when it violates policy?

> **Governance principle.** Governance without enforceable control is incomplete.

3. Data Assurance

Is data moving as intended?

Organizations need confidence that enterprise data is moving according to security, privacy, sovereignty, compliance, and business requirements.

That means moving beyond simply observing traffic toward providing assurance that data is following intended and sanctioned paths.

Questions to ask:

- Where is enterprise data moving?
- Which AI services are interacting with it?
- Are sensitive flows appropriately protected?
- Are geographic or sovereignty requirements being maintained?
- Can policies remain associated with data movement across distributed environments?
- Can the organization verify that data traveled as intended?

> **Governance principle.** Governance should extend to the data itself.

4. Compliance and Accountability

Can the organization prove its policies are working?

Static compliance checklists document intent.

Operational governance requires evidence.

Organizations should be capable of answering:

What data moved where, who or what accessed it, how was it processed, what actions were taken, and under what policy or regulatory constraints?

> **Governance principle.** A policy says what should happen. Assurance demonstrates what did happen.

5. Measurement and Optimization

How is AI actually being consumed?

Once organizations can see and govern AI activity, governance can evolve beyond risk management.

Organizations can begin understanding AI consumption as an operational and business metric.

This may include:

- Application adoption
- Usage trends
- Sanctioned versus unsanctioned AI
- AI consumption patterns

- Token-related usage
- Redundant AI services
- Changes in usage over time

> **Governance principle.** What can be measured can be better managed.

From Visibility to Enforceable Governance

The progression matters.

> **Visibility.** tells the organization what exists.
>
> **Control.** establishes what should be allowed.
>
> **Data Assurance.** helps determine whether data is moving according to those requirements.
>
> **Compliance.** creates accountability and evidence.
>
> **Measurement.** helps the organization continuously improve.

Together, these capabilities move AI governance beyond policies and dashboards toward something more operational:

Observe → Understand → Govern → Enforce → Optimize

That distinction becomes increasingly important as AI moves from tools that employees use to software capable of acting autonomously across enterprise environments.

The AI-ready enterprise will need to govern both.

Graphiant's Role

Graphiant approaches this challenge from the infrastructure layer.

Graphiant's Data Assurance capabilities are designed to bring visibility, policy-driven control, and assurance closer to the network carrying enterprise data.

As AI adoption expands, the Graphiant AI Dashboard extends this approach by helping organizations understand AI application usage across their environment, including sanctioned and unsanctioned applications, usage patterns, and token-related consumption.

The Dashboard is not an AI governance strategy by itself.

It provides a critical starting point:

> **Knowing what AI is actually being used.**

From there, organizations can begin making informed decisions about access, data movement, policy, control, and enforcement.

The ultimate goal is not to prevent AI adoption.

It is to create an infrastructure foundation capable of supporting AI innovation without losing control of the data that powers it.

Sources

1. Gartner, “Lack of AI-Ready Data Puts AI Projects at Risk,” February 26, 2025. Gartner predicts that through 2026, organizations will abandon 60% of AI projects unsupported by AI-ready data.

2. Gartner, “Lack of AI-Ready Data Puts AI Projects at Risk,” February 26, 2025. Gartner reported that 63% of organizations either do not have—or are unsure whether they have—the right data-management practices for AI and predicted that through 2026, organizations will abandon 60% of AI projects unsupported by AI-ready data.

Assured, Agile, Awesome.

Graphiant delivers a single, unified Network-as-a-Service platform built for the AI era. By combining a Stateless Core architecture with end-to-end encryption, real-time observability, and policy-driven routing, Graphiant helps enterprises and service providers connect users, sites, clouds, and AI workloads with the speed, security, and agility modern business demands.

Learn more at graphiant.com