

# How to Mature Generative and Agentic AI Governance for Enterprise Applications

29 October 2025

By: Max Goss, Stephen Emmott, Tristan Iles

As megavendors such as Microsoft, Google and Salesforce embed more GenAI and agentic AI capabilities into their solutions, IT increasingly struggles to keep up with the pace of change. Application leaders must establish operational AI governance at the application level to proactively mitigate risk and drive value.

## Overview

### Key Findings

- 70% of IT leaders say they have a centralized AI strategy, but only 34% can actually apply it to their enterprise applications. The result is a growing disconnect between the organization's central AI policies and how they are applied across their application portfolio.
- Most IT leaders (68%) say they struggle to keep up with the pace of change as application vendors continually add new embedded generative and agentic AI features, with 38% admitting that their application vendors drive their GenAI strategy.
- IT is constantly playing defense and is often unable to vet new AI features before the vendors make them available to employees.
- Only 15% of IT leaders strongly believe they have the right governance models in place to manage AI agents in their enterprise applications, while 84% believe they need additional technical controls to mitigate risk. Despite this, only 16% are investing in dedicated AI security and governance products. Most are focusing on developing policies and training.

## Recommendations

- Unify AI initiatives by developing a centralized AI governance committee that is responsible for building and coordinating the organization's AI strategy. Ensure representation from key enterprise application leaders (ERP, CRM, digital workplace) working alongside senior IT and business roles.
- Build AI operational governance at the application or domain level. Develop the roles and responsibilities to monitor and manage vendor releases, empowering application teams to evaluate embedded AI tools and block or restrict those that do not align with the organization's AI strategy.
- Enforce your AI governance policies by integrating AI-specific security and governance controls into your application portfolio. Consider third-party AI trust, risk and security management (AI TRiSM) vendors alongside first-party controls provided by the application vendors.

## Strategic Planning Assumption

By 2029, more than 50% of user interactions linked to enterprise business processes will leverage large language models to bypass traditional enterprise applications, up from less than 5% in 2025.

## Introduction

According to Gartner survey data, 68% of IT leaders struggle to keep up with the rapid inclusion of generative AI (GenAI) and agentic AI in their major enterprise applications — ERP, CRM and digital workplace (DW).

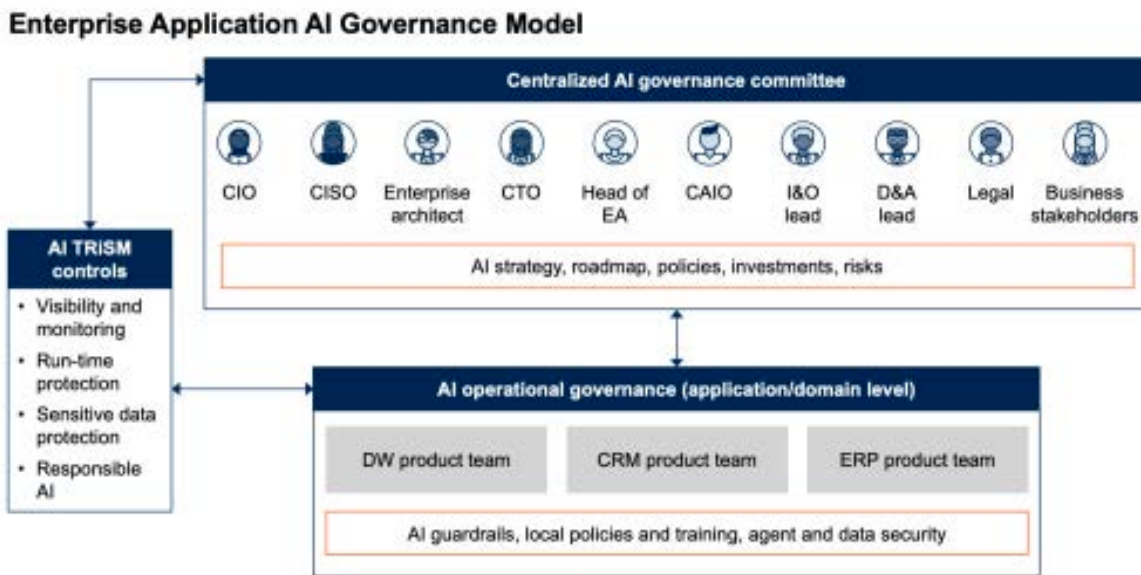
A key reason for this is that while many organizations say they have a centralized GenAI strategy, most are unable to actually apply it to their applications and enforce policies.

*70% of organizations say they have a centralized AI strategy, but only 34% can apply it consistently to their enterprise applications. This highlights a clear disconnect between the organization's AI strategy and the reality of what is happening in their critical enterprise applications. Worryingly, 30% report not having a centralized AI strategy at all.*

Ineffective governance is one of the biggest blockers to AI deployments, with only 15% of application leaders reporting high confidence in their ability to manage and govern AI agents in their enterprise applications. As a result, over 50% of organizations limit their AI rollouts to low-risk or trusted users only. However, this also limits the value they get from their AI investments. Gartner survey data highlights that these organizations are over three times less likely to report high value from their GenAI tools.

Governance should be an enabler for GenAI value; however, it is more often a blocker. Application leaders often ask, "How can we mature our AI governance for our major enterprise applications?"

Gartner has created an AI governance model to help leaders to understand the roles and responsibilities as well as the technical controls required to mature their AI governance and proactively apply it to their applications (see Figure 1).



Source: Gartner  
826141

Gartner

Gartner’s enterprise application AI governance model consists of three components:

- **Centralized AI governance committee:** The AI governance committee builds the overall AI roadmap for the organization and makes the ultimate decisions on tool investment and risk. It also defines centralized policies for AI usage.
- **AI operational governance (application/domain level):** These groups apply the AI strategy to key enterprise applications (e.g., ERP, CRM, DW). They align application roadmaps with business needs and adapt centralized AI policies to specific contexts — such as managing agents and data use in tools like Microsoft 365 Copilot.
- **AI TRiSM controls:** To enforce AI governance, organizations must go beyond traditional security measures and apply AI TRiSM capabilities centrally across AI models and services and locally within applications. These controls include AI and agent visibility, runtime protection, and safeguarding sensitive data.



## Analysis

### Develop a Cross-Functional AI Governance Committee

Gartner defines AI governance as the process of creating policies, assigning decision rights and ensuring organizational accountability for risks and investment decisions in the application and use of AI techniques (see Reference Guide for AI Governance).

A centralized and cross-functional AI governance function or committee is required for developing an organization wide AI strategy and the policies and tooling to enforce safe usage. While AI governance should be its own initiative, it should follow a functioning governance operating model, such as IT or data governance.

*Organizations with a centralized GenAI strategy that is applied consistently across their applications are over 3x more likely to report high confidence in their ability to govern AI, deliver effective change management, drive skills and assess AI's value. They are also 2x more likely to report significant value from their AI investments.*

Figure 2 shows example roles that are fundamental to the success of the AI governance committee. This committee should be led by the organization's AI lead, usually a senior head of AI or chief AI officer. However, it must also contain cross-functional representation from security, compliance, business leaders and, crucially, application leaders who represent the core enterprise applications that the organization owns (e.g., Microsoft 365, Salesforce, SAP).

## Key Roles Making Up the AI Governance Committee



Source: Gartner  
826141

Gartner

If your organization is one of the 30% that does not have a centralized or formalized AI strategy, start by implementing the central AI governance committee. Where one already exists, ensure that it is suitably empowered. Application leaders will need to work with IT senior leadership and the CIO office to ensure that the group has the right level of support and backing to succeed. If you are struggling to secure senior leadership support, highlight that organizations that can apply their Gen AI strategy consistently are two times more likely to see significant value from their GenAI investments.

For the 70% of organizations that already have an AI governance strategy but can't consistently apply it across their enterprise applications, the first task is to ensure that the head of enterprise applications role, or nominated deputy, attends the AI governance committee. They must be empowered to bridge the gap between the organization's stated AI strategy and the reality of what is happening in the major DW, ERP and CRM applications. This will enable application owners to provide key feedback, and will help to ensure that embedded AI tools adhere to the organization's AI strategy.

## Build Proactive AI Operational Governance at the Application Level

While the AI governance committee is concerned with the overall AI roadmap for the organization, companywide policies and risk mitigation, operational governance is focused on the implementation of these policies at the application or domain level — something that is so often missed.

---

*In the context of enterprise applications, Gartner defines AI operational governance as a framework for defining and implementing proportionate policies and controls at the domain or application level to mitigate AI risk, guide user behavior and enable the organization to succeed with AI.*


---

### Establishing AI Operational Governance

To enable effective operational governance and to help align application and business strategy, Gartner recommends that the major enterprise applications (ERP, CRM, DW) be managed through proactive product teams. Each product team reports to an operational governance group that manages the applications in that domain. These groups provide the perfect vehicle to establish domain-specific AI policies and controls, aligned with the organization's AI governance strategy. For more information on the product team approach, see Drive Value With Enterprise Application Product Teams.

Under the product team model, AI operational governance becomes an extension of the application or domain governance group. The group is chaired by the application/domain product owner and contains key representatives from security, compliance and identity who work with the business owners as part of a multidisciplinary team (see Figure 3).

## Operational Governance Requires a Multidisciplinary Approach



Source: Gartner  
826141

Gartner

Together, these roles work to ensure that the organization’s AI roadmap and policies are implemented in the applications under their purview. For example, if the organization’s AI policy is to restrict end-user agent creation, this group should ensure that the applications they own follow this policy, but address any key nuances. For example, the application may include access to low-risk agents that could be safely enabled for specific use cases. This requires working closely with the central AI governance committee and also proactively monitoring the application vendor’s roadmap to ensure the organization stays ahead of any changes.

Each product owner (ERP, CRM, DW) reports into the head of enterprise applications role. This role escalates key AI risks or issues to the central AI governance committee, ensuring connectivity between organization wide and application-level AI policies.

### Enforce AI Governance With TRiSM Controls

For AI operational governance to be effective, application leaders need to be able to enforce the policies they create and monitor for signs of abuse or noncompliance. To do this, they need AI-specific trust, security and risk management controls, or AI TRiSM.

## AI Presents New Organizational Risks

AI presents several new risks that can not be adequately mitigated through existing application and data security controls. For example:

- The prompt/response architecture that AI chatbots use can be exploited by bad actors or malicious insiders through techniques such as prompt injections.
- AI tools can be used either deliberately or by accident to spread misinformation or to exfiltrate sensitive data.
- The proliferation of AI agents presents a new attack vector, and many organizations are concerned with managing AI agent sprawl.

Figure 4 highlights the AI risks and the security and risk mitigations needed.

**AI Risk Factors and Mitigations**

<b>Compromise and threat vectors</b>	<b>Compromises or attacks against AI entities</b>	<b>Malicious activities</b>	<b>AI hallucinations or user mistakes</b>
<b>Types of compromise</b>	Altered AI behavior, AI theft	Data compromise, poisoning or drift	Misinformation, faulty decision making
<b>Security and risk management approach</b>	<b>AI entity</b> Integrity (model, app, agent)	<b>AI data</b> Protection and privacy	<b>Responsible AI</b> Education and training

Source: Gartner  
826141

Gartner

Embedded AI tools can increase these if compromised, as a result of their tight integration with core business applications and access to sensitive data. For more information on embedded AI risks see *Exposing and Managing Embedded AI: Tools for Transparency and Vendor Oversight*.

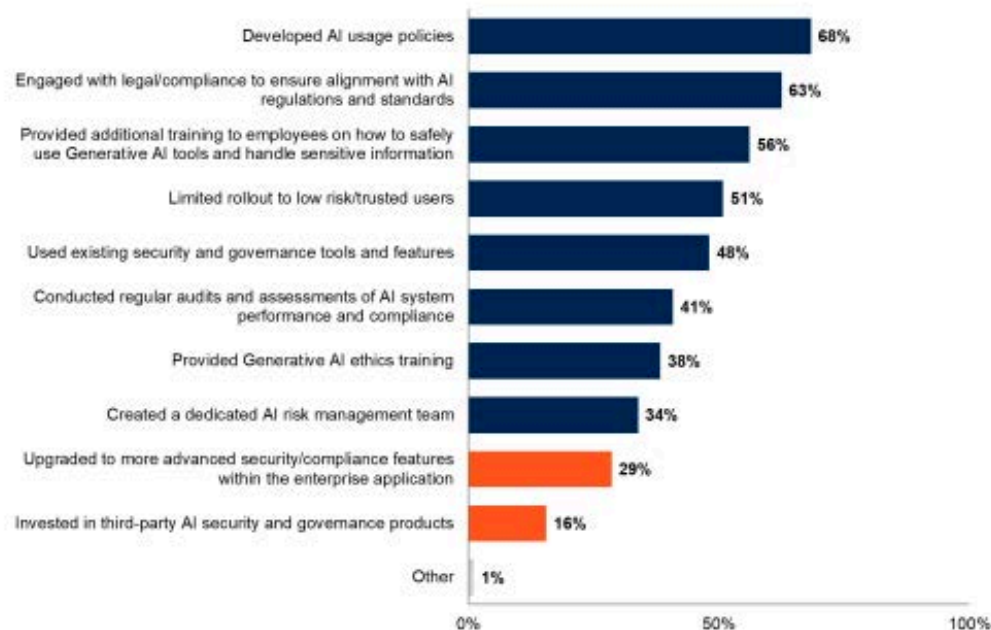
Many application leaders do not trust that their vendors have the necessary AI security and governance capabilities to mitigate the risks introduced by AI. Only 26% have high trust in the vendor’s ability to provide adequate AI security and governance controls, and only 16% have high trust in the vendor’s ability to mitigate hallucination risks. Organizations should look beyond their existing application vendors to provide AI security controls. However, most are not yet doing this.

## Most Rely on User Policies and Training Rather Than on Dedicated Security Controls

Policies and standards are a vital component of AI governance, but without controls to enforce them, they are ineffective. While many IT leaders have created AI usage policies and worked with legal and compliance teams to build regulations and standards, less than a third have upgraded security and compliance features in their existing applications, and just 16% have invested in third-party AI governance tools (see Figure 5). IT leaders need to complement policies and training with an equal focus on implementing AI-specific security controls in their applications.

### Approaches Taken to Mitigate GenAI Risks

Multiple responses allowed



n = 360; IT Leaders, excluding unsure  
H2\_Q08. Which of the following approaches has your organization taken to mitigate potential risks associated with the rollout of Generative AI tools?  
Source: 2025 Gartner Generative and Agentic AI in Enterprise Applications Survey  
838963

## Invest in AI TRiSM Controls Together With Responsible AI Training

To enforce your AI policies, augment the security and governance capabilities provided by your existing enterprise applications with capabilities offered by vendors in the AI TRiSM market. Many of the vendors in this market offer runtime protection to secure prompts and responses, together with monitoring and remediation capabilities to detect anomalies, such as when an agent's behavior changes.

*Organizations that invest in third-party AI governance products are almost two times more likely to report higher levels of value from their AI tools, highlighting a clear link between effective and enforceable governance and the value that organizations get from their AI investments.*

Look for AI TRiSM controls that are model- or vendor-agnostic and can be applied to different embedded AI tools across the enterprise application portfolio.

For embedded AI, it is critical to have the following capabilities, which can be obtained by a combination of AI TRiSM vendors, AI management and governance features offered by the application vendors (often as an add-on license), and the organization providing responsible AI training and education:

- **AI visibility:** Catalog embedded AI tools across enterprise applications, who owns the tools and what risks they present.
- **Runtime inspection and enforcement:** Intercept prompts/responses in real time (e.g., via a proxy) to ensure that users' AI interactions are safe and conform to organizational policies.
- **AI data protection:** Ensure that embedded AI does not leak sensitive data by using appropriate data classification, life cycle and access controls.

- **Anomaly detection:** Determine whether an AI entity (such as an AI agent) is behaving as expected, and monitor for deviations or anomalies. For example: An AI agent is compromised, and in addition to its original actions, it now sends an email to an external domain.
- **AI agent access and life cycle controls:** Enforce who can create and share agents and ensure that when built, AI agents conform to least-privilege architecture and have appropriate life cycle policies to mitigate the risks of agent sprawl.
- **Cost management:** Forecast and track consumption costs when using GenAI tools and agents, particularly as many vendors offer pay-as-you-go billing options.
- **Responsible AI:** Develop a culture of responsible AI, working with security teams to educate and empower end users to utilize embedded AI tools safely. End users must know how to validate AI output and identify misinformation. Gartner believes that responsible AI education will become just as important as cybersecurity training in the years to come and should form part of your organization's mandatory security training.

## Evidence

2025 Gartner Generative and Agentic AI in Enterprise Applications Survey. This study was conducted to understand the key challenges and opportunities when deploying generative AI (GenAI) tools, and where organizations should focus their AI investments. This research also aims to understand what stage organizations are at on their AI agent journey and their thoughts on AI agents. The research was conducted online from May through June 2025 among 360 respondents from organizations with at least 250 full-time employees across all industries (except IT software) in North America (n = 149), Europe (n = 140) and Asia/Pacific (n = 71). Soft quotas were established for country, company size, and respondent's function type and job level to ensure a good representation across the sample. Organizations were required to have deployed or plan to deploy in less than one year at least one generative AI tool in at least one core enterprise application domain: digital workplace applications, customer relationship management applications, or enterprise resource planning applications. Respondents were team leaders or above, excluding C-level, and involved in the rollout of generative AI tools; they were required to have certain responsibilities regarding these generative AI tools. Disclaimer: The results of this survey do not represent global findings or the market as a whole, but reflect the sentiments of the respondents and companies surveyed.

Gartner, How to Mature Generative and Agentic AI Governance for Enterprise Applications, by Max Goss, Stephen Emmott and Tristan Iles, 29 October 2025. This graphic was published by Gartner, Inc. as part of a larger research document and should be evaluated in the context of the entire document. The Gartner document is available upon request from NLP Logix. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.