Magic Quadrant for AI Governance Platforms

16 June 2026 - ID G00841467 - 52 min read

By Lauren Kornutick, Sumit Agarwal, and 3 more

AI governance platforms are designed to centrally define, approve and enforce responsible AI policies across comprehensive AI use cases, applications and agents. As an emerging market, these platforms are essential for operationalizing responsible AI throughout an organization's entire AI ecosystem.

Strategic Planning Assumptions

- By 2027, AI governance will become a requirement of all sovereign AI laws and regulations worldwide.
- By 2027, AI governance and responsible AI capabilities will be part of 75% of AI platforms, making them the main area of competition.

Market Definition/Description

Gartner defines AI governance platforms as tools designed to ensure organizations comply with their responsible AI practices, organization policy, regulations, and other risk management frameworks/industry standards. They enable AI leaders and other leaders to streamline AI governance processes organizationwide and are a central repository that links trust, risk and security runtime controls for AI systems and third-party AI usage. They automate workflow approvals for new AI use cases, applications and agents, and support risk-based, real-time execution of responsible AI guardrails.

AI governance platforms (AIGPs) are tailored to the organization's AI governance leader. This leader is responsible for setting internal governance policy across common

responsible AI principles (RAI) and accountable for providing corporate assurance that policy rules that can be translated to technical controls are enforced at runtime. AIGPs serve a wide range of assets built using multiple AI techniques and must be able to support any AI use case. AIGPs must be interoperable across the organization's technology and data stack as well as domain-specific tools addressing operational execution of governance policy.

AIGPs tie the corporate oversight and application of AI governance policy to real-time execution of these requirements for responsible AI practices across AI systems and third-party AI usage in the organization. These platforms automate governance policy, and manage and report on AI risks and acceptable use adherence in the enterprise across all forms of AI. They serve as oversight systems that continuously manage, implement and enforce the necessary trust, risk and security controls (e.g., data and model guardrails). This aligns with requirements to demonstrate that the organization has implemented and is governing all AI use cases, including agents and third-party applications or models.

AIGP tools facilitate the ongoing AI use-case risk assessment and approval process for AI systems, such as models, applications or agents, and streamline information exchange with AI governance stakeholders. They incorporate real-time observability and responsible AI policy guardrail enforcement along with audit trails.

An AIGP serves as a central repository for continuous monitoring and policy enforcement from AI governance rules that cover corporate responsible AI policy, regulations, frameworks and standards. It also has the ability to capture data and/or metadata from more than one of the following operational governance categories: acceptable use, identity and other organization-level security policies; observability; and data governance. AIGPs must have policy engines (e.g., prepackaged rules and/or models) to adhere to common regulations (e.g., EU AI Act), frameworks and standards, such as NIST AI RMF and ISO 420001, with the option to customize rules for corporate policy and apply enforcement at runtime.

Mandatory Features

- **AI discovery and registry:** Provides a centralized, discoverable registry of all AI use cases and AI services like SaaS with embedded AI that can support multiple use cases. These use cases include applications, agents and models within the organization, including version history, metadata (purpose, data sources, algorithms),

documentation (e.g., model cards or systems cards, agentic decision logic), ownership, development stage and deployment status.

- **Compliance risk management:** Catalogs the risks applied to AI applications, agents, models or use cases. The ability to classify, assess and mitigate AI-specific risks (bias, fairness, robustness, etc.), including content libraries that address unique laws for AI and data protection (e.g., EU AI Act, GDPR), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO 42001), as well as address compliance and decision making with organizational AI policy, such as acceptable use and common RAI principles of accountability, fairness, explainability, transparency, and security and safety.
- **Policy management and enforcement:** Provides centralized, automated management and enforcement of AI-specific policies via multiple guardrails, including control validation for AI-specific risks (e.g., bias, data leakage, trust, privacy, security), access controls, use-case alignment and other enterprise policies, remediation recommendations, and compliance reporting.
- **Dynamic risk scoring:** Enables model builders or AI agent, application or use-case product owners to continuously monitor, understand and diagnose the performance and behavior of AI models, applications and agents in production. This enables stakeholders to understand why an AI system behaves a certain way through explainability techniques. This includes providing alerts, dashboards and historical trend analysis to ensure that AI systems remain accurate, reliable and compliant over time.
- **Evidence collection:** Provides documentation for trust, risk and security assessments, testing and validation results (such as security, bias detection, and model), and risk and compliance remediation evidence.
- **Interoperability:** Enables different systems, devices, applications or agents to exchange and utilize information effectively.
- **Workflow and approvals:** Enables the automation of routine governance tasks such as a new AI use-case model, application or agent approval by a governing body; risk and security assessments internal to the organization or with third parties; approvals; testing procedures; and documentation generation and feedback loop from detection to remediation. Facilitates communication and coordinated action among diverse stakeholders. Includes structured signoff, attestation and approval requirements.

- **Audit trail:** Provides comprehensive audit trails of actions taken in the platform and, where applicable, automatically logs all activities related to the AI life cycle.

Optional Features

- **AI usage reporting:** Automates the generation of standardized documentation for AI models and use (e.g., model cards, datasheets) for auditors, regulators, and so on. Reporting should include the ability to customize dashboards for end users in nontechnical roles and the ability to observe monitoring in real time.

- **Data usage mapping:** Captures data mapping used by various AI entities and tracks use and misuse over time. This may also include the ability to track the provenance of training data and interface with data governance platforms to include data lineage, classification, ownership and data observability information. This feature can be captured via the AI governance platform directly or via interoperability with a D&A governance platform or similar technology solution.

- **AI value tracking:** AIGPs capture the use-case requirements and the expected business value or other nonfinancial KPIs (such as usage or hours saved). AIGPs may extend their observability capabilities to keep track of the outcomes and the associated value to enable a single organizational view of the value achievement. AIGPs may be interoperable with tools that govern spend.

- **Business-friendly user experience:** Enables targeted users to easily navigate and use the tool to complete their tasks without the need to consult with product subject matter experts (SMEs) or technical staff. This could be interpreted at minimum to mean that the majority of users will not revert back to tools such as spreadsheets after using the tool and/or including visualization capabilities in the user experience (UX).

- **Ease of implementation:** Allows users to quickly adopt a new instance of the tool to support AI governance activities without the need to heavily customize off-the-shelf templates/prebuilt workflows or make changes to the underlying data model. This could include how-to guides, AI agents or easy plug-ins to implement the tool.

Magic Quadrant

Figure 1: Magic Quadrant for AI Governance Platforms

Gartner

Vendor Strengths and Cautions

Airia

Airia is a Visionary in this Magic Quadrant. Its AI governance platform (AIGP) is focused on providing centralized visibility into AI agents, models and unmanaged AI, enabling real-time, in-line policy enforcement through an AI/Model Context Protocol (MCP) gateway. It also includes automated controls for compliance and risk management aligned with regulations (e.g., EU AI Act), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO/IEC 42001).

Airia's primary market presence is in North America (NA), with early expansion into Europe to support the EU AI Act, and presence in Asia/Pacific (APAC) and Latin America

(LATAM). Its customer base consists of midsize to large enterprises in the technology, telecom, banking and financial services, and healthcare industries.

Its roadmap includes guardian agents for autonomous monitoring and adaptive policy enforcement for agent-based AI systems, living disclosures for continuously updated AI transparency artifacts, and automated AI change management features that support versioning, evaluation and validation across distributed AI assets.

Strengths

- **Product strategy:** Airia's embedded active AI governance control plane enforces policy at runtime rather than point-in-time assessments. A unified life cycle covers discovery, assessment, approval and monitoring with governance regression on models or agent changes to ensure continuous risk mitigation.
- **Regulatory alignment and compliance automation:** The platform aligns with prepackaged regulations, frameworks and standards to build a defensible compliance posture. It achieves this through AI-assisted assessments, versioned policy packs, automated evidence capture and continuous monitoring.
- **Runtime security and agent governance depth:** Airia governs agentic AI using multilayer prompt injection, jailbreak detection, parameter-level agent constraints, data loss prevention guardrails, adversarial red teaming and centralized agent inventories with real-time observability. The depth of these agent governance features helps clients prevent attacks, protect sensitive data and ensure compliance.

Cautions

- **Delivery ecosystem maturity:** As Airia continues to scale, its expanding partner ecosystem increasingly supports enterprise implementation needs. However, global system integrator engagement remains limited compared to established vendors.
- **Limited global coverage:** Despite its strength in U.S. and EU frameworks, Airia offers fewer prebuilt regulatory packs for APAC, the Middle East, and other regions. Enterprises with broad geographic exposure may require custom frameworks or depend on roadmap expansions to achieve coverage parity globally..
- **Support for edge environments:** Airia's edge and on-device AI capabilities require moderate local runtime resources when deploying local inference with an on-device SLM. This requirement may limit deployment on lightweight, low-powered clients,

restricting adoption for organizations that operate in desktop-constrained, regulated environments.

Cranium AI

Cranium AI is a Niche Player in this Magic Quadrant. Its Cranium AI Security & Governance Platform focuses on the convergence of AI governance and security into a single continuous platform. Its product enables organizations to create a unified AI asset inventory via discovery across code, cloud and third parties, and operationalizes security risk and compliance guardrails from a single platform.

Cranium AI's operations are primarily focused in NA, and its clients tend to be large global enterprises in the technology, telecom, financial services and healthcare industries.

Its roadmap includes adding an autonomous guardian agent that continuously monitors, assesses, and acts on AI governance signals across the enterprise. It also includes a personalizable UI that will offer a role-based, configurable user experience (UX) to adapt Cranium to different personas, including chief information security officers (CISOs), compliance officers, data scientists, third-party risk managers and developers.

Note: Cranium AI closed its acquisition of Aiceberg, an agentic AI security and risk management company, on 21 May 2026, after the process of the research period for this Magic Quadrant had closed. As such, Gartner's evaluation in this Magic Quadrant does not cover the inclusion of Aiceberg. The acquisition extends Cranium's platform from model- and application-layer governance into the monitoring and control of autonomous AI agents; broadens Cranium's coverage of the emerging agentic AI use case; and adds Aiceberg's engineering team, with its founder appointed Cranium's CTO. Gartner will provide additional insight and research to clients as more detail becomes available.

Strengths

- **Third-party AI risk:** Cranium extends its dynamic risk scoring to evaluate third-party risk. By combining continuous scoring with automated red teaming, it unifies compliance and AI security risk, enabling organizations to manage their broader risk ecosystems.

- **Policy runtime enforcement:** Cranium captures full spectrum policy management and helps end users build policies to onboard AI usage. It automates predeployment guardrails to provide in-line runtime enforcement with a five-layer risk framework covering context, information, content, instruction and alignment. The platform integrates with security information and event management (SIEM) tools and IT service management (ITSM) tools, enabling policy enforcement through existing enterprise GRC workflows.
- **Comprehensive AI inventory:** Cranium's multilayer, native AI discovery tool enables organizations to map and govern all AI assets by identifying AI-containing repositories at scale, and extracts models, dependencies and agent definitions. Its Agent Discovery generates detailed bills of materials, including instructions, roles and relationships. Shadow AI discovery identifies AI services across the enterprise, providing visibility into unauthorized and unmanaged AI use alongside code-level and cloud-level discovery.

Cautions

- **Growth expansion phase:** Cranium is a growth-focused company actively investing in product development and market expansion, and has demonstrated enterprise customer traction across multiple verticals. Given that the company is expanding, enterprise buyers should evaluate against enterprise third-party due-diligence processes.
- **Compliance management:** Cranium's compliance capabilities automatically map controls across multiple regulatory frameworks, with automated scoring and reminders and support for custom frameworks. Fully automated cascading of regulatory changes across all use cases is maturing on the roadmap.
- **Customer segment alignment:** Cranium's product best addresses AI security and AI agent governance use cases via runtime AI capabilities delivered through its Arena Shield offering, since it is optimized for the cybersecurity buyer. Buyer-specific messaging and capabilities across additional personas responsible for AI governance continues to mature.

Credo AI

Credo AI is a Visionary in this Magic Quadrant. Its Credo AI governance platform focuses on becoming a central layer that integrates with the enterprise tech stack to provide AI-specific risk discovery and visibility. Credo AI structures AI governance as an operational workflow, starting with intake questionnaires that leverage a Triggers & Actions engine to automatically start a risk-scoring process. This in turn assigns policy packs and manages reviewer routing.

Credo AI's operations are mostly focused on NA, and its clients tend to be midsize to large organizations in the financial services, healthcare and technology sectors.

Its roadmap includes a focus on multientity governance, including platforms, tools and MCP servers, to seamlessly expose governance operations to AI agents.

Strengths

- **Discovery and visibility:** Credo AI enables organizations to monitor the entire AI stack to address compliance risk through comprehensive AI asset management and a centralized inventory. Organizations track first-party models and third-party applications using automated discovery of unmanaged "shadow AI," which is achieved through cloud access security broker and proxy log ingestion, enabling customers to uncover hidden vulnerabilities and securely govern all first- and third-party AI assets.

- **Dynamic risk scoring:** Credo AI provides a highly contextual view of risk that adapts to different users and risk postures, enabling better risk communication for stakeholders. The platform enables configurable risk quantification by aggregating likelihood and impact scores across diverse risk types, delivering inherent and residual views through visual heat maps.

- **Vision:** Credo AI's product strategy demonstrates a well-developed understanding of the market, grounded in customer feedback. Its multientity governance roadmap focuses on guardian agents for automated oversight, an MCP server for seamless agentic-based integration, and specialized governance models for regulated verticals like healthcare and insurance.

Cautions

- **Data provenance limitation:** Credo AI's limited native data usage mapping and data provenance capabilities may create visibility gaps for organizations seeking a fully holistic view of AI risk across data, models and agents. While Credo AI integrates with external data governance tools, buyers evaluating the platform as a central

governance layer should assess whether these integrations provide sufficient end-to-end visibility and control for their needs.

- **Customer segment alignment:** Credo AI's marketing targets AI governance leaders, chief data and AI officers, and CISOs, but has made a deliberate choice to not specialize in messaging or product capabilities for all buyers' needs. Consequently, prospective customers looking to fully address AI security and AI agent governance use cases must consider additional technology tools to integrate with Credo AI's central governance plane.
- **Growth-stage startup:** Credo AI is in an active growth stage operating in a rapidly evolving and competitive market. While it demonstrates strong growth and customer traction, it is still scaling its operations, ecosystem and delivery capabilities.

Holistic AI

Holistic AI is a Challenger in this Magic Quadrant. Its Holistic AIGP is focused on providing a centralized hub for managing the full AI life cycle, offering automated policy orchestration and role-based workflows tracked through configurable dashboards for portfolio-level visibility.

Holistic AI's operations are presently focused in NA and Europe, but are expanding globally, and its clients tend to be large organizations operating in the healthcare, professional services (e.g., banking and consulting) and telecom sectors.

Its roadmap includes the planned launch of operative guardian agents (September 2026), which will provide autonomous, agentic-based oversight and real-time enforcement across an organization's entire AI estate. Additionally, Holistic AI plans to deliver an AI configuration management database (October 2026) and an "AI exchange" for the reuse of AI assets (November 2026) to enhance enterprisewide visibility and AI capability adoption.

Strengths

- **Fast implementation:** Holistic AI gets its customers up and running quickly with an average setup time of six weeks. To improve collaboration across roles, it features a configurable, role-based UI with embedded assistive capabilities like no-code policy builders and an AI copilot for administrators and users.

- **Workflow and auditability:** The product provides automated, multistage workflow engines for AI intake and risk assessment that integrate with continuous integration/continuous delivery (CI/CD) pipelines to enforce governance gates. These features run in tandem while maintaining an immutable, tamper-proof audit trail that captures every action and decision for regulator-ready compliance, which is unique.
- **Product balance and strategy:** Holistic AI's features and product strategy provide clients with a future roadmap that drives depth and builds more automation into a balanced set of capabilities. This fits well within the rapidly evolving AIGP market.

Cautions

- **AI value and risk quantification:** The solution relies on a semiquantitative approach to risk and value quantification and does not natively address stochastic quantitative models — neither proprietary nor using monte carlo or FAIR models. This limits the ability to translate technical AI risk scores into hard quantitative numbers tied to business impact language.
- **Broad industry focus:** Holistic AI has a broader customer base in critical infrastructure, such as technology and telecom, manufacturing, and service industry verticals While this reflects a diversified customer base when compared with traditional regulated industries, this focus may not be sufficiently broad for all regulated industries.
- **Vision:** Holistic AI's product vision lacks native integration depth for AgentOps specifically related to data governance, lineage and security. This may create adoption barriers for AI-forward prospects requiring out-of-the-box connectors for their specific technical stacks.

IBM

IBM is a Leader in this Magic Quadrant. Its watsonx.governance platform is focused on enterprise AI governance and assurance. It supports AI life cycle oversight and includes policy enforcement, risk controls and validation. It integrates AI governance with AI security ecosystem partners and embeds IBM OpenPages GRC to power its watsonx.governance console. It provides prebuilt policy and control mappings for regulations (e.g., EU AI Act), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO/IEC 42001) to support compliance and runtime policy enforcement.

IBM has global operations, with growth concentrated in the Americas, EMEA and APAC. Its customer base is predominantly large enterprises and public sector organizations in regulated industries, with strong representation in the financial services, technology and telecom sectors.

Its roadmap includes a governance graph to centralize AI asset inventories and lineage, AI horizon scanning for regulatory monitoring, and risk advisor agents to automate risk and compliance use cases.

Strengths

- **Overall viability:** IBM's long-term viability is critical for enterprise customers with complex needs. Its global presence, brand awareness, diversified portfolio and proven record as a trusted enterprise technology partner support client confidence, even in an emerging market.
- **Geographic strategy:** IBM's sales and support teams and large partner network provide localized customer support in nearly all regions. Its flexible deployment options, such as on-premises and hybrid approaches, and expanded data center support make it a strong fit for global organizations with distributed global teams.
- **Industry expertise:** IBM watsonx.governance excels at serving highly regulated global organizations operating in the financial services, healthcare and energy sectors. It supports complex organizational structures by enabling AI governance teams to aggregate, analyze and report on risk exposures across business units and geographies.

Cautions

- **Modular coverage:** AI-native experiences are still maturing and are largely modular, which further drives governance silos. This may slow adoption for organizations prioritizing lightweight, highly automated governance tools that are a single source of truth for the AI governance life cycle until the capability is delivered in June 2026.
- **Interoperability and data coverage:** Achieving a unified AI governance view often depends on partner data sources and additional configuration. This reliance can increase deployment effort for customers seeking comprehensive, out-of-the-box visibility across AI risks, assets and controls.

- **Time to implementation:** IBM's governance model requires significant implementation effort, which may overwhelm organizations with decentralized teams or early-stage governance maturity. Consequently, these buyers may experience a longer time to value.

ModelOp

ModelOp is a Visionary in this Magic Quadrant. Its ModelOp Center is focused on being the central AI system of record for enterprises. ModelOp's AIGP provides centralized inventory, policy oversight, dynamic risk, traceability and approvals across the AI life cycle. It extends beyond risk management to address runtime behavior, policy adherence, risk, performance, cost and ROI change detection, and remediation coordination through interoperable workflows.

ModelOp's operations are mostly focused in the U.S., and its clients tend to be large multinational organizations in heavily regulated industries.

Its roadmap includes an agentic workflow development framework, enabling the creation of custom agents to automate tasks in the governance life cycle, as well as an advanced FinOps UX to provide proactive cost-benefit analysis across the portfolio of AI solutions.

Strengths

- **Regulatory framework support:** ModelOp has out-of-the-box support for regulations (e.g., EU AI Act, General Data Protection Regulation), frameworks (e.g., NIST AI RMF), as well as vertical-specific regulations. Clients are also able to build their own standards and customize existing frameworks. This support is operationalized through ModelOp's life cycle automation and workflow engine.
- **Detailed AI documentation:** The information cards that ModelOp dynamically generates are notably detailed, including risk tiering and prompt templates. This creates visibility into model inventories, and their utility extends beyond compliance audit use cases.
- **Emerging customer-centric roadmap**: ModelOp's product balances pragmatic improvements with forward-looking innovations designed for the emerging AI governance leader, making it well-suited for customers seeking a central policy and governance control plane that can also address cost governance and agentic development. The Completeness of Vision will equip buyers with a long-term

approach that they can scale from risk management into the actionable data needed to track AI value.

Cautions

- **Delivery ecosystem maturity:** As ModelOp continues to scale, its partner and delivery ecosystem is still developing. Buyers should factor in organizational readiness as part of their implementation planning, as people and process decisions rather than technical configuration typically account for the bulk of deployment time.
- **Enterprise deployment-focused:** ModelOp does not offer a public SaaS solution, so clients retain responsibility for the infrastructure to operate ModelOp. The platform itself is managed by ModelOp, but it does not offer hosting as part of the solution.
- **Pricing:** ModelOp's consumption-based pricing is tied to AI solutions under management, leading to variable costs between contracts. While the average annual cost is higher than many solutions in this market, the model includes unlimited users and does not require the additional professional services investment that comparable solutions often carry.

Monitaur

Monitaur is a Visionary in this Magic Quadrant. Its AI governance platform is focused on risk management and compliance. Its product serves consequential, high-risk and all other use cases with a control library that links controls to key AI projects or system risks. These tiered controls are mapped to regulations and standards, as well as Monitaur's policy templates, creating a scalable policy and control library that is transferable across AI types, geographies and use cases.

Its operations are mostly focused in the U.S., and its clients tend to be large, heavily regulated organizations in the insurance and financial services verticals.

Its roadmap includes further integrations with AI development platforms, enhancements to risk quantification and business alignment with native KPIs and metric management.

Strengths

- **Compliance risk management:** Monitaur utilizes a structured control library that is directly aligned with relevant regulations and standards to address compliance requirements systematically throughout the AI life cycle. By automating key

compliance tasks, such as testing and monitoring, the platform reduces reliance on manual processes, minimizes the risk of oversight and enforces governance policies at runtime that extend beyond traditional GRC tools.

- **Risk quantification:** Monitaur calculates expected risk using probability and conditional loss formulas, which are then reduced by control-level mitigations. This approach allows buyers to understand their exposure and effectively prioritize their AI remediation efforts.
- **Vertical strategy:** Monitaur's proven experience serving large global insurers demonstrates its ability to navigate highly complex environments. Enterprise buyers in equally complex, regulated sectors (e.g., banking, healthcare, pharmaceuticals) can use the platform to manage their own strict compliance needs.

Cautions

- **Partner programs:** Monitaur is in an active growth stage in an evolving market. Therefore, its partner programs and initial co-selling are currently limited. This potentially impacts enterprises requiring mature delivery ecosystems.
- **Product strategy:** Monitaur's product plan is focused on delivering governance for high-risk use cases. This could be construed as too limited a view for enterprises looking for solutions with broad AI usage control and operational governance requirements.
- **Geographic strategy:** While Monitaur has the option to deploy its product using Amazon Web Services or other cloud location of preference, it has limited global technical support and no options for on-premises deployments. This may restrict adoption for organizations with more rigid definitions of regional sovereignty or where end-user preference is to deploy on-premises for added security.

OneTrust

OneTrust is a Visionary in this Magic Quadrant. Its AI governance platform provides a centralized system of record for AI risk and compliance, automated AI inventory, risk classification, and assessments aligned to regulations (e.g., EU AI Act), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO/IEC 42001).

OneTrust operates with broad geographic coverage, with an established customer base in NA, Europe and APAC. Its clients are primarily midsize to large enterprises in regulated industries, including technology, telecom, manufacturing, banking and financial services, and investment services.

Its roadmap includes agent protocol policy contracts, intended to apply policy-as-code controls at the agent runtime layer; third-party AI intelligence to support extended AI supply chain risk management; and an AI governance agent designed to automate use-case intake, risk tiering and evidence collection.

Strengths

- **Regulatory alignment:** OneTrust provides native out-of-the-box support for standards-based frameworks. Continuously updated regulatory content, combined with customizable frameworks and audit-ready evidence, supports large enterprises managing complex compliance obligations.
- **Comprehensive governance workflows:** The platform delivers end-to-end AI governance workflows spanning intake, risk assessment, approvals and runtime enforcement. Risk-adaptive controls dynamically adjust approval rigor and monitoring intensity, integrating directly with CI/CD, machine learning operationalization (MLOps) and enterprise workflows to embed governance into AI delivery processes.
- **Comprehensive inventory and runtime controls:** OneTrust offers broad AI coverage across systems, models, datasets and agent-based AI. By combining a centralized inventory with dynamic risk scoring and extensive runtime controls, the platform enables clients to safely scale their AI initiatives while actively preventing bias and sensitive data leaks.

Cautions

- **Manual implementation burden:** OneTrust's AI governance solution looks standardized and simple on the surface, but users are responsible for most of the detailed implementation. What's possible is not the same as what's provided or recommended.
- **Attribute-based access controls:** OneTrust's attribute-based access control capabilities (ABAC) are not yet developed with the level of granularity needed to be successful across all AI assets and workflows. Organizations with complex authorization requirements may experience limitations in fine-grained, policy-driven

access control enforcement at large scale until this feature is released in September 2026.

- **Evolving agent identity governance:** Agent and nonhuman identity governance is still maturing. While OneTrust integrates with enterprise identity and access management systems, deeper support for agent credentialing, delegated actions and workload identity is under development, which may limit advanced control over autonomous agent ecosystems today.

Relyance AI

Relyance AI is a Niche Player in this Magic Quadrant. Its AIGP is focused on unifying AI security, data security, and privacy into one platform via Data Journeys.

Relyance's operations are primarily focused in the U.S., and its clients tend to be in critical infrastructure sectors such as telecommunications, healthcare and financial services.

Its roadmap includes real-time data flow governance to enable continuous runtime visibility into how data flows through AI systems, services and agents in production. Relyance also plans to extend AI agent governance and security across code and cloud platforms, including mapping agent identity, data exposure, privilege levels and tool access.

Strengths

- **Data-focused AI governance:** Relyance provides broad native discovery, integrations (e.g., models, agents, datasets, shadow AI), automated metadata and lineage capture across major cloud AI platforms, identity providers and data stores for AI use cases. This provides customers with high inventory completeness and contextual metadata with minimal manual tagging, accelerating time to value and improving accuracy of risk assessments.

- **Continuous, audit-ready evidence:** The platform autogenerates live evidence and continuously maps AI inventory to regulations (e.g., EU AI Act), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO/IEC 42001), enabling customers to produce current compliance artifacts without manual assembly.

- **Flexible deployment choices:** Relyance offers SaaS, private cloud, on-premises subscriptions, tiered packaging, volume discounts and modular add-ons. This

flexibility enables customers to align procurement and deployment to compliance and architecture constraints without needing custom licensing models.

Cautions

- **Runtime governance gaps:** Several runtime controls that are essential for holistic program-level AI governance are unsupported, including bias monitoring, prompt injection and jailbreak prevention, hallucination checks, model/agent observability, rate limiting, and access-control enforcement. This leaves customers exposed to operational risks and forces them to deploy complementary tools or accept residual runtime governance gaps.
- **Limited workflows:** Relyance does not provide end-to-end intake/deployment approval workflows and reports no telemetry for drift, bias, toxicity, performance or cost. This limits customers' ability to operationalize approvals, detect runtime agent and model issues, or measure governance effectiveness specific to top-down, high-risk AI governance deployments.
- **Integration drag:** Customers in highly distributed or complex environments may require upfront configuration, as coverage of specific third-party AI vendors depends on integration and API availability. For SaaS sources, OAuth-based setup enables integration, while cloud and noncloud environments use an infrastructure-as-code pattern that automates integration setup and accelerates time to value.

Saidot

Saidot is a Niche Player in this Magic Quadrant. Its AIGP is graph-based and focused on enabling scalable governance. Saidot's knowledge graph architecture means information is entered once and reused across all relevant contexts — risks, controls, evidence and policies propagate automatically through connected systems, models, agents, tools and datasets. It is grounded in regulations (e.g., EU AI Act), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO/IEC 42001), with curated libraries and MCP-based agentic integration for developer workflows.

Saidot's operations are primarily focused in Europe, and its customers tend to be in financial services, large consultancies and the public sector.

Its roadmap includes policy-as-code and automated evidence collection. This will enable replacing static documents as evidence with continuous evidence collection and

programmatic compliance verification across AI pipelines, including AI agents.

Strengths

- **Policy packs:** Customers in heavily regulated sectors (e.g., financial services and government) benefit from Saidot's industry-specific policy monitoring and tailored policy templates.
- **Marketing execution:** Saidot integrates a responsible AI focus into its product strategy, aligning marketing efforts with these core product themes. Its approach is selective and aimed at AI governance and responsible AI leaders.
- **EU focus:** Organizations operating in the EU benefit from Saidot's regional focus, which helps them navigate and comply with the EU AI Act and ISO/IEC 42001 standard implementation.

Cautions

- **Partner ecosystem:** Saidot's partner programs and initial co-selling are limited, potentially impacting enterprises requiring mature delivery ecosystems.
- **Limited full-cycle governance:** The product functionality is focused on policy and control mappings with AI models, agents and systems. It provides triggers based on model provider changes, then automatically triggers actions from runtime observability events through the observability ingestion module. Binding these triggers to specific Saidot controls is on the roadmap, and does not extend to value and cost management.
- **Product scope:** Saidot's product has a limited scope of policy implementation. Its roadmap does not include related AI usage with security and privacy, and its policy runtime enforcement consists of observability-ingestion-module, webhook-based actions and internal actions. Customers seeking advanced features outside of frameworks to translate regulations to policies and controls outside of the Microsoft environment must rely on third-party tools and integrations to fill these gaps.

SAP

SAP is a Niche Player in this Magic Quadrant. Its SAP AI Agent Hub is focused on building and maintaining a continuously updated, organizationwide AI inventory complete with a

view of risk exposure across embedded AI, models in use and all agents connected to their business context.

SAP's operations are geographically distributed, and its customers tend to be midsize to large enterprises in the utilities, energy and technology verticals.

Its roadmap includes extending guardian agent capabilities with AI-powered remediation recommendations through an AI agent embedded in the governance workflow. This agent performs conformance checks via mining, derives guardrails, suggests specific remediation steps, drafts policy language, prioritizes actions by risk impact and addresses multiagent orchestration via the AI Agent Hub.

Strengths

- **Overall viability:** SAP's established global footprint, diversified portfolio and corporate strategy offer a low-risk option for enterprise buyers. SAP's long-term market stability is a differentiator for organizations that want to extend their existing platforms into emerging AI governance capabilities.
- **SAP integration:** AI Agent Hub provides runtime enforcement for AI agents and other use cases built within the SAP environment and tightly integrates with solutions across the SAP portfolio. It is a bundled product with SAP LeanIX, Joule Studio (formerly SAP Build), SAP Signavio, SAP Integration Suite and other components of the SAP portfolio.
- **Customer experience:** SAP provides resources and support for customers aiming to advance their automation efforts. SAP partners include technology vendors, cloud providers and consulting firms, ranging from Global Strategic Service Partners that manage massive enterprise transformations to specialized independent software vendors selling apps on the SAP Store.

Cautions

- **Non-SAP governance gaps:** Although SAP markets its AIGP to govern any AI use case, it lacks automated enforcement for non-SAP solutions. Consequently, organizations governing third-party AI use cases will face increased manual governance burdens when operating outside of the native SAP environment.
- **Product innovation:** SAP continues to invest in AI Agent Hub governance, with a unique "Verified" status feature that provides runtime enforcement for AI asset reuse

by developers (e.g., in Joule Studio and third-party agent access to SAP systems). However, SAP lacks critical features for comprehensive AI governance across the organization, particularly in risk management, value tracking and cost optimization outside of SAP systems.

- **Dynamic risk scoring:** The product does not continuously monitor AI risk in real time or provide alerts at runtime. End users must update risks manually or rely on a static risk calculation, delaying threat detection and increasing manual workloads when managing enterprise AI.

ServiceNow

ServiceNow is a Leader in this Magic Quadrant. Its AI Control Tower is focused on centralized risk and compliance management across all AI assets. Impact assessments identify, score and mitigate risks, with dynamic scoring that automatically recalculates as controls are applied.

ServiceNow has global operations, with growth concentrated in the Americas, EMEA and APAC. Its customer base is predominantly large enterprises and upper-midmarket and public sector organizations in the financial services, banking and investment sectors, followed by the technology and telecom sectors.

Its roadmap includes expanding multicloud automated AI discovery, including detection of shadow AI in SaaS applications and code repositories, where discovered assets are autoenrolled in governance workflows with risk classification. ServiceNow also plans to extend AI agent identity governance from human users to AI agents operating autonomously across the enterprise via a Veza and Moveworks integration.

Strengths

- **Overall viability:** ServiceNow's long-term viability is critical for enterprise customers with complex needs. Its global presence, brand awareness, large enterprise footprint in ITSM, growing presence in CRM and HR, and proven record as a trusted enterprise technology partner support client confidence in an emerging market.

- **Policy packs with GRC integration:** ServiceNow provides out-of-the-box content packs for the EU AI Act and NIST AI RMF to automate compliance mapping, while allowing customers to customize controls and guardrails. The product also extends GRC workflows to enforce controls and features strong audit trails and approval flows.

- **Security and compliance certifications:** ServiceNow maintains more than 20 compliance, privacy and cybersecurity certifications, including ISO/IEC 42001 for AI governance and FedRAMP certification, making it a strong option for customers who prioritize those certifications during supplier onboarding.

Cautions

- **Cost-prohibitive for all AI:** ServiceNow's AI governance is chargeable for governing AI use cases outside of ServiceNow. Despite being purchased and paid for up front, costs will be based on the deployment and use of the AI use cases in production.
- **Integration complexity:** Integrating ServiceNow's AI governance features with diverse data sources, external AI systems, or specialized compliance and security tools for enhanced runtime inspection and enforcement can be complex due to organization requirements, and therefore may require custom configuration or additional middleware.
- **Limited AI-powered governance automation:** At the time of evaluation, ServiceNow's product lacked key features to improve the end-user experience for AI-based risk assessment and automated regulatory change impact analysis. The lack of native guardian agent capabilities could lead to end users spending too much time on inefficient GRC processes rather than more impactful AI governance activities.

Truyo

Truyo is a Leader in this Magic Quadrant. Its Truyo AIGP is focused on providing centralized AI risk and compliance management by maintaining an enterprise AI inventory, conducting AI impact and risk assessments, and generating audit-ready evidence. It supports ongoing monitoring, regulatory change tracking and governance workflows to demonstrate compliance and risk mitigation across the AI life cycle.

Truyo's operations are mostly focused in NA, and it has global expansion plans. Its clients tend to be regulated entities in the healthcare, financial services and government sectors.

Its roadmap includes native guardian agent capabilities via AI-powered vendor and risk assessment automation, with website and document scanning to detect AI usage and privacy risks;

intelligent prepopulation of assessments; and dynamic risk scoring and continuous monitoring.

Strengths

- **AI discovery:** Truyo provides native shadow AI and AI agent discovery that automatically links discovered agents to govern use-case records, risk assessments and approval workflows. This gives customers visibility into emerging agentic risks and enables faster assessment and oversight of previously unknown AI activity.
- **Adaptable workflow:** Truyo offers configurable, role-based workflows, connectors and sector-focused messaging to operationalize vertical requirements without custom code. Therefore, customers can adapt intakes, assessments and approvals to specific industry rules with lower implementation complexity and clearer governance mapping.
- **Customer focus:** Truyo added agent discovery in 2026 and actively incorporates customer feedback into its roadmap. It ships quarterly major releases, enabling customers to address emerging risks faster and benefit from frequent releases.

Cautions

- **Sales and channel strategy:** Truyo's heavy reliance on channel sales (e.g., 70% in NA and 100% in several regions) and a small direct sales team of four full-time equivalents can affect direct account engagement, escalation responsiveness, and continuity of pre- and postsales relationships for some customers.
- **Runtime enforcement limitations:** Runtime enforcement is governance-centric and does not perform in-line blocking or autonomous model intervention. Customers that require real-time, in-line execution protections must rely on provider/platform controls or additional tooling, increasing integration complexity and residual operational risk.
- **Limited executive functionality:** There is limited support for executive-level summaries or out-of-the-box, board-ready views tailored for AI governance boards and senior leadership, which may require manual report assembly or additional services to inform strategic decisions. Customers who require formal committee voting, quorum tracking and documented dissent will need workarounds or professional services until the planned enhancements are delivered, which can slow governance committee adoption.

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant may change over time. A vendor's appearance in a Magic Quadrant one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. It may be a reflection of a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

This is the inaugural publication of this research. Therefore, no vendors have been dropped.

Inclusion and Exclusion Criteria

In addition to Gartner client relevance, as determined by analyst expertise and opinion, providers needed to meet the following criteria to qualify for inclusion:

- Vendors must meet all the mandatory features as described in the market definition; all features must be generally available on or before 1 April 2026 and form a stand-alone solution.
- Vendors must have more than 10 paid deployments of the stand-alone, independent, commercially off-the-shelf AIGP product as of 25 February 2026.
- Must support governance of all types of AI, as defined below:
    - AI that is built by the organization’s data science, AI engineering or software engineering team using a platform specific for AI development
    - AI that is custom-built on top of an enterprise app to solve a business need
    - Third-party app or model (any AI that is procured either by a central IT team or business user)
    - Agentic AI that is built by the organization’s data science or software engineering team using a platform specific for AI development
    - Agentic AI that is custom-built on top of an enterprise app to solve a business need

- Any third-party agent (any AI that is procured either by a central IT team or business user)

Geographic Presence: Vendor must have products deployed in **more than two** of the following regions: U.S., Canada, LATAM, U.K., Europe, Middle East and Africa, APAC

Magic Quadrant Exclusion Criteria

- Vendors with AIGP capabilities embedded within a broader enterprise or cyber GRC tool, data management, data security, data and analytics governance, or data science/machine learning platform, but that do not offer a stand-alone, independent, commercially off-the-shelf AIGP product.
- Vendors whose AIGP offering is predominantly focused on governing the use of AI within their own platforms.
- Products not aimed toward AI leaders and AI governance leaders.

Honorable Mentions

The AIGP market has many vendors. Of these, 13 were found relevant to our clients and were selected for evaluation in this Magic Quadrant. However, the exclusion of a provider does not mean that the vendor lacks viability. The following are noteworthy vendors not included in the formal analysis. These six vendors could be appropriate for clients, contingent on their requirements.

The vendors are displayed alphabetically.

Enzai

Enzai provides a purpose-built AIGP designed to help enterprises manage AI use across the life cycle. The platform maintains a centralized record of AI use cases, systems, models, agents and vendors.

The platform supports configurable workflows for intake, review, approval and ongoing oversight. Enzai places strong emphasis on regulations (e.g., EU AI Act), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO/IEC 42001). Its roadmap highlights expanded support for agentic AI governance and greater automation in policy enforcement and risk scoring.

LatticeFlow AI

LatticeFlow AI is an emerging vendor in AI governance. It offers a platform designed to transition organizations from superficial check-the-box AI risk management to evidence-based AI governance. By utilizing Swiss engineering and scientific research, the company offers deep technical evaluations and continuous monitoring for autonomous, agentic AI systems.

Its platform capabilities include AI asset discovery, ensuring compliance with international standards and regulations like NIST RMF, ISO42001 and the EU AI Act, and performing automated security testing. Real-world applications of its technology span across critical infrastructure, global finance, frontier AI labs and defense operations to ensure system reliability and performance. Through strategic partnerships and a public registry of AI frameworks (AI Atlas), the platform maps AI risks to technical controls and interprets complex technical evidence into actionable insights for secure innovation.

Modulos

Modulos AG offers the Modulos Risk-centric AI Governance Platform. This end-to-end platform automates regulatory gap assessments, risk monitoring and remediation tracking for AI systems. The platform provides a unified Governance Graph that connects risks, controls, evidence, requirements and frameworks. AI agents automate routine GRC tasks such as evidence collection and control assessments.

A differentiating feature is quantifying AI risk in monetary terms. It supports regulations (e.g., EU AI Act), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO/IEC 42001). Real-time compliance posture is maintained through automated monitoring and integration with existing tech stacks via open APIs, which also surface shadow AI by reconciling discovered usage against the governed inventory.The platform also features structured review workflows and generates audit-ready reports.

.

Singulr

Singulr delivers a full life cycle AI governance control plane that combines continuous AI discovery, contextual risk scoring, automated onboarding, runtime guardrails and audit-ready reporting. It spans employee AI use, custom models, third-party tools, embedded AI and agentic systems.

The platform takes an integration-first approach, working with existing security, IT, cloud, MCP, and SaaS agentic platform controls rather than forcing proxy-based deployments. The platform also provides governance for endpoint AI agents like Claude, Cursor and ChatGPT via native OS sensors and browser extensions. Its broad coverage of emerging agentic platforms appeals to large organizations looking for an independent governance layer on top of existing security and IT stacks.

Trustible

Trustible offers the Trustible Responsible AI Governance Platform, an AI-powered SaaS solution designed to orchestrate enterprise AI programs for legal, risk, compliance and technology teams. The platform centralizes AI use cases, agents, models, datasets and vendors into a unified inventory for visibility and risk prioritization, supported by customizable workflows for reviewing proposals, assigning risk levels and tracking approvals.

Trustible provides governance blueprints, policy templates, risk taxonomies and audit preparation tools that align with regulations (e.g., EU AI Act), frameworks (e.g., NIST AI RMF) and standards (e.g., ISO/IEC 42001). It does not natively provide centralized runtime enforcement, opting instead to orchestrate via API and MCP controls.

WitnessAI

WitnessAI offers the WitnessAI Unified AI Security and Governance Platform, a unified AI security and governance solution securing employees, models, applications and agents. It classifies AI activity by intent, not keywords, and sees every AI surface, including 4,500+ apps, thick clients, IDEs, agents and API traffic.

Governance spans three layers: Observe (discover all AI, shadow AI, agentic plugins, MCP servers), Control (inline block, reroute, redact, workflow actions based on identity, intention, model and app, and agent approved-tool enforcement) and Protect (bidirectional runtime AI defense blocking prompt injection and filtering harmful output across 100+ models). It integrates through the option that best fits each AI surface and works with a wide variety of security, identity management and cloud platforms to enforce controls where necessary.

Evaluation Criteria

Ability to Execute

Gartner evaluates a vendor's Ability to Execute by meticulously analyzing its products, services, viability and the overall customer experience it delivers. The ultimate measure of a vendor's Ability to Execute lies in its capacity to fulfill its commitments and its track record of success in doing so.

In alignment with this, Gartner's Magic Quadrant for AI Governance Platforms assigns "high" priority to the criteria of product or service and overall viability These elements are crucial indicators of a vendor's capability to deliver on its promises effectively and resist turbulent conditions in an emerging market.

The criteria for marketing execution and customer experience are assigned a "medium" weighting. This weighting underscores the necessity for vendors to target the right organization leaders responsible for AI governance and support their products. Sales execution/pricing is assigned a "low" rating consistent with the nascent nature of the emerging market.

Market responsiveness/record and operations are not evaluated in this first Magic Quadrant iteration. As this is the inaugural release of this Magic Quadrant, historical performance and business operations are not a distinguishing factor. However, this aspect is anticipated to gain importance in future iterations.

Ability to Execute Evaluation Criteria

Evaluation Criteria	*Weighting*
Product or Service	High
Overall Viability	High
Sales Execution/Pricing	Low
Market Responsiveness/Record	NotRated
Marketing Execution	Medium

Evaluation Criteria	*Weighting*
Customer Experience	Medium
Operations	NotRated

Source: Gartner (May 2026)

Completeness of Vision

Gartner conducts a thorough evaluation of vendors' ability to comprehend both current and future market and technology trends, customer needs and competitive dynamics, collectively known as their Completeness of Vision.

This assessment ultimately hinges on the vendors' understanding of how market forces can be leveraged to generate growth opportunities. This qualitative evaluation is informed by Gartner's extensive interactions with end users and its comprehensive market insights.

Four criteria are assigned "high" weightings. As the AIGP market continues to evolve, a deep market understanding, robust offering (product) strategy, innovation and geographic strategy are all critical components for vendors to consistently deliver value to customers amid expanding customer demands. Moreover, global enterprises need contingency plans due to geopolitical tensions that could change the delicate balance of how they conduct business across boundaries. Therefore, we have prioritized sovereign AI now, which is a key consideration in assigning geographic strategy a "high" weighting.

Vertical strategy is assigned a "low" weighting. Although this factor is important, Gartner considers this to be less important than other factors in a nascent market where most buying behaviors are driven by regulatory compliance obligations. We believe a more accurate indicator of a vendor's vision is a robust product strategy that plans to embed AI security, AI risk governance, AI operations and AI agents and innovation to incorporate guardian agents.

The marketing strategy and business model were not evaluated. The majority of AIGPs employ similar business model approaches, making it a nondifferentiating factor in

assessing a vendor's vision. Furthermore, marketing strategy is closely tied to sales strategy, and Gartner considers it too soon to assess the market on these criteria since it has just gained its grounding.

Completeness of Vision Evaluation Criteria

Evaluation Criteria	*Weighting*
Market Understanding	High
Marketing Strategy	NotRated
Sales Strategy	NotRated
Offering (Product) Strategy	High
Business Model	NotRated
Vertical/Industry Strategy	Low
Innovation	High
Geographic Strategy	High

Source: Gartner (May 2026)

Quadrant Descriptions

Leaders

Leaders are in the strongest position to influence the market's growth and direction. They demonstrate a market-defining vision for how AIGPs can help organizations support an enterprisewide AI governance program as a first line of defense, encompassing a holistic AI governance process that includes compliance risk management, security and operations of all AI use in the organization.

Leaders can execute against that vision through products and services and have demonstrated business results in the form of revenue and earnings. They excel in their combination of market understanding, innovation, product features and functions, and overall viability.

While maintaining a well-established base of long-term customers, Leaders show a consistent ability to win new deals. They have customers in many geographic regions, cover a wide variety of industries and serve customer organizations of a range of sizes. Leaders are often the vendors that other providers measure themselves against.

Challengers

Challengers have established presence, credibility and viability, and have demonstrated the ability to meet customers' expectations in terms of functionality and customer experience. Challengers may have a good vision for technology, but may not have fully won over business stakeholders and IT executives.

Challengers are positioned well to succeed. However, they may not demonstrate thought leadership or innovation to the same degree as Leaders. They may be a good choice for organizations that value execution over vision and leading-edge functionality.

Visionaries

Visionaries are ahead of most competitors in terms of delivering innovative products and product strategy. They are sometimes smaller vendors or newer entrants that embody trends that are shaping, or will shape, the AIGP market. Visionaries have a strong vision and roadmap, which brings innovation and strong functionality to their platforms.

Visionaries may be a good choice for organizations that want an opportunity to skip a generation of technology. They may offer a competitive advantage or a chance to influence their product roadmap. They might be acquired or face a challenge to increase their market share. However, as these vendors mature and prove their Ability to Execute, they may become Leaders.

Niche Players

Niche Players may offer viable AIGPs but they frequently lack a complete execution of product vision that focuses on all four use cases and/or are early-stage startups struggling with driving enterprise AI adoption. Some Niche Players focus on specific

regulated industries like banking or insurance, while others prioritize capabilities for cybersecurity leaders who co-own AI governance.

Niche Players can often offer the best solutions to meet the needs of particular organizations, considering role priority or the price-to-value ratio of their solutions. These vendors may win deals in specific regions or industries. However, they are typically not winning new business across multiple regions or industries at the same pace as vendors in the other quadrants.

Some Niche Players demonstrate a degree of vision that suggests they might become Visionaries, but they may struggle to make this vision compelling. They may also struggle to develop a track record of continual innovation. Other Niche Players may have the opportunity to become Challengers if they continue to develop products with a view to improving their overall execution.

Context

The AIGP market is a distinct market consisting of legacy enterprise vendors who've built a stand-alone, cohesive AI governance product into their existing portfolio, as well as early- to growth-stage startups who pioneered the category. We identified more than 100 vendors marketing AI governance capabilities.

While most vendors position their solutions as AIGPs, given that AI is moving so quickly, the reality is that few products meet the comprehensive needs of enterprise-level AI governance leaders. This leader is not necessarily a role in the organization, but is responsible for key governance activities that ensure the appropriate oversight of AI. These activities are frequently conflated with a governance, risk and compliance (GRC) process because they address risk management; however, AI governance moves beyond enterprise and cyber GRC and cuts into operations.

Therefore, AIGPs should be viewed not as extension of GRC processes, but as a single source of truth for AI governance operations. Put simply, an AIGP serves as the organization's central governance plane to:

- Approve AI use
- Define the policies and guardrails for AI systems (models, apps, agents)

- Monitor AI system behavior behavior

Enforce critical enterprise guardrails wherever they are deployed (this is particularly important for agents) These objects could be deployed in different platforms, but the control plane (and accordingly inventory) is centralized. AIGPs must be able to integrate with other technology solutions to apply these objects, monitor them continuously and interfere to take action if there is a violation (referred to as “runtime enforcement”).

As a result, this Magic Quadrant focuses on the limited subset of the broader AI governance and security market that addresses the needs of the leader responsible for AI governance who is responsible for all AI use in the enterprise. We include only vendors that offer qualifying tools designed to support holistic AI governance processes across AI risk and compliance, AI operations, AI security and AI agent governance.

Recommendations for selecting an AIGP vendor:

- Assemble a cross-functional team, including your CIO and AI leader, and assurance partners in legal, compliance and cybersecurity, to evaluate vendor options based on required functionality, innovation and integration needs. This increases the likelihood of successful selection and user adoption.
- Consider whether existing enterprise vendors are sufficient for your organization’s planned AI journey. Ensure that your chosen solution can serve as a central AI inventory and scale with your growth plans and future needs, such as expanding to agent deployment for complex use cases.
- When evaluating early-stage and growth-stage startups, carefully consider overall long-term viability and evaluate whether the product’s capabilities outweigh the uncertainty of inevitable market consolidation.
- Assess vendor performance against your expectations. Given the variability in performance, avoid long-term lock-in for both startups and enterprise solutions offering modules. Few startups lay out a clear vision for enterprise players.
- This market is ripe for acquisition and other M&A activity. Consult Gartner if an AI governance solution you are using or considering is being acquired by another vendor or if your existing enterprise vendor acquires an AIGP.
- If procuring at this early stage, negotiate and lock in multiyear preferential pricing. Vendors are likely to significantly increase prices as they grow and demonstrate a

consistent ability to sell into large, complex enterprises with big budget centers.

Market Overview

The AI governance platform market comprises products designed to centrally define, approve and enforce responsible AI (RAI) policies across comprehensive AI use cases, applications and agents. This market is projected to grow rapidly at a 67.5% CAGR, expanding from $65 million in 2024 to $1,434 billion by 2030.

AIGPs act as a centralized governance plane, uniquely providing real-time observability, dynamic risk scoring and runtime enforcement to monitor and intervene if policy violations occur. Unlike traditional GRC tools — which are largely static repositories enabling point-in-time risk assessments, lack real-time enforcement capabilities and function more like cloud-based spreadsheets — AIGPs are specifically built to manage the operational layer of AI risk as a first line of defense.

Key Market Trends

- **Market evolution and consolidation:** In the next two years, the AIGP market is expected to undergo consolidation as startups are acquired by vendors in adjacent markets such as AI security and GRC to close product gaps, as evidenced by acquisitions in this space that have already occurred. Furthermore, AIGPs will need to evolve to comprehensively manage AI cost and agent sprawl, which will become primary drivers for platform adoption.

- **AI agent governance will force continued consolidation with AI security:** Agents operating without policy enforcement and in-line blocking of unwanted activity can cause irreversible harm for organizations. For AI governance to be effective, security and governance controls that address behavioral anomaly detection and in-line blocking, dynamic least-privilege access enforcement and multiagent trust chain verification will fully converge in the next two years for agents.

- **Rise of guardian agents:** AIGPs will increasingly focus on decision governance and situational awareness to support complex agentic use cases. This includes introducing guardian agents — AI-based governance workflows that will continuously monitor baseline context and execute real-time insights via multitiered evaluations to safely oversee autonomous decision making.

- **Decision governance for autonomous agents:** As AI agent use cases become more integrated into enterprise ecosystems, aligning AI intent with outcomes through "decision governance" will be critical.
- **Regulatory requirements:** Navigating new and complex regulations like the EU AI Act and emerging U.S. state regulations like the Colorado AI Act is a major driver for AIGP investment. Furthermore, as AI incidents become mainstream and tracked in databases, organizations are seeking AIGPs to prevent reputational damage and satisfy underwriters who increasingly require robust AI controls to provide cyber and AI liability insurance.
- **Distinction from cybersecurity and data tools:** While cybersecurity tools will continue to focus broadly on all AI use, in the future, AIGPs will clearly distinguish themselves by targeting high-risk deployments, enforcing responsible AI policies that include security use cases, and managing high risk, high-value or consequential agents. Additionally, data-centric AIGPs (DAGPs) are expected to emerge to prioritize controlling the provenance of data throughout an organization.
- **Sovereign AI:** Sovereign AI is becoming a top priority for executives and is influencing technology decisions amid ongoing regulatory uncertainty in various regions. AIGPs will evolve to apply additional regulatory resiliency to the various AI technology decisions, including data, to ensure adherence to the necessary standards.

⊕ Evaluation Criteria Definitions