FRONTIER.

FRONTIER RESEARCH / DECISION-GRADE ANALYSIS

Your AI Policy Can't Stop an Agent: The State of Enterprise AI Governance

Governance is becoming the infrastructure of delegated authority. A CIO field guide to cross-platform AI controls, 23 comparative assessments and 46 vendor families and offers—across models, applications and agents.

FR-032 / Dion Hinchcliffe / Frontier Research

REPORT PROFILE

Built to change a decision.

Governance is becoming the infrastructure of delegated authority. A CIO field guide to cross-platform AI controls, 23 comparative assessments and 46 vendor families and offers—across models, applications and agents.

REPORT	**FR-032**
AUTHOR	**Dion Hinchcliffe**
PUBLISHED	**Frontier Research**
PROFILE	**frontier-ro-0.1**

This PDF is a publication rendering of the canonical Frontier Research Object. The web report remains the live source of truth; citations, figures, methodology, and rights metadata travel with this edition.

The Frontier verdict

The enterprise has spent years teaching software to answer. It is now teaching software to act. Too much AI governance still assumes the answer is the dangerous part.

An approved model can power an unauthorized decision. A compliant chatbot can hand a tool an excessive permission. A perfectly logged agent can spend the budget before anybody reads the log. The market's central problem is no longer the absence of governance language. It is the distance between a business rule and the place where that rule must take effect.

Our thesis: the next durable AI governance platform will connect enterprise accountability to enforceable authority across heterogeneous AI. It will know which system is acting, for whom, on what data, toward which outcome, under which policy, at what cost—and have a dependable way to constrain that action. A registry without enforcement is an incomplete control system. Enforcement without business context is an incomplete governance system.

This is an expansion of the argument in Frontier's The Autonomous Enterprise: governance should increase the amount of responsibility a business can safely delegate. The objective is neither unrestricted automation nor an approval committee that becomes the enterprise's slowest API. It is more useful autonomy within explicit boundaries.[1]

Three findings should shape the CIO's buying decision:

1. Buy a control architecture before buying a console. GRC, identity, data governance, AI evaluation, gateways, security and workflow platforms are converging. A common screen does not prove a common policy model or a working enforcement path.
2. Separate enterprise reach from enforceable authority. A product can inventory AI across the company yet control only a subset. Another can reliably block a tool call but know little about the business obligation behind it. Both can be valuable. Neither alone establishes enterprise-wide governance.
3. Treat governance as part of production economics. Its return includes faster approval, reusable controls, less rework, bounded incidents and more accepted business outcomes. A platform that makes every legitimate action expensive or slow can destroy the value it is meant to protect.

   This report evaluates a comparative field of 23 offers and provides a wider directory of 46 vendor families and offers considered. It is not a ranked shortlist, a market-share study or a certification. Our evaluation is based on public primary-source evidence available on November 9, 2026. The map deliberately separates documented control paths from claims, launches and previews. The absence of public evidence is an uncertainty to resolve, not proof that a capability does not exist.

1. Define the thing being governed

The purchasing category is confusing because “AI governance” describes at least four different jobs: governing the organization's AI program; assuring the behavior of particular models and applications; authorizing production decisions; and detecting or responding to attacks. A fifth, AI used to automate GRC work, is often marketed beside them. These jobs overlap. They are not interchangeable.

The right unit of governance is the AI-enabled business system. That includes its purpose, accountable owner, models, prompts, retrieval sources, tools, identities, suppliers, deployment, affected people and downstream actions. A model inventory is one component of that system. An agent inventory is another. Neither is a complete inventory of responsibility.

The same foundation must cover a credit-risk model, a document classifier, a customer-service assistant and a purchasing agent. Their assurance requirements differ, but they should not require four incompatible definitions of owner, approval, exception, incident or evidence. NIST's AI Risk Management Framework provides a useful organizational foundation; its generative-AI profile adds relevant risk considerations. These frameworks guide a management system. They do not establish that a particular product enforces it.[2][3]

Non-agentic AI does not become yesterday's problem

Predictive and conventional machine-learning systems still need data lineage, independent validation where appropriate, bias and performance assessment, drift monitoring, version control, deployment approval and retirement. For consequential decisions, the enterprise also needs reason codes, contestability and an accountable process for correcting outcomes. An agent-native product that ignores these requirements is not an enterprise-wide replacement.

Generative applications add groundedness, prompt and retrieval integrity, sensitive-data disclosure, content safety, intellectual-property obligations and changing provider behavior. Evaluate the application and its evidence sources, not merely the model name. A model substitution can change quality, cost, region and contractual exposure without changing the application's business purpose.

Agents add delegated authority over a sequence of actions. Their plans can evolve, they can invoke tools, create subagents, consume resources and persist beyond one conversation. The relevant control must travel through the chain of delegation. OWASP's agentic risk work is useful here because it addresses problems such as goal manipulation and tool misuse; a conventional output filter does not cover that entire threat surface.[4]

The capability contract

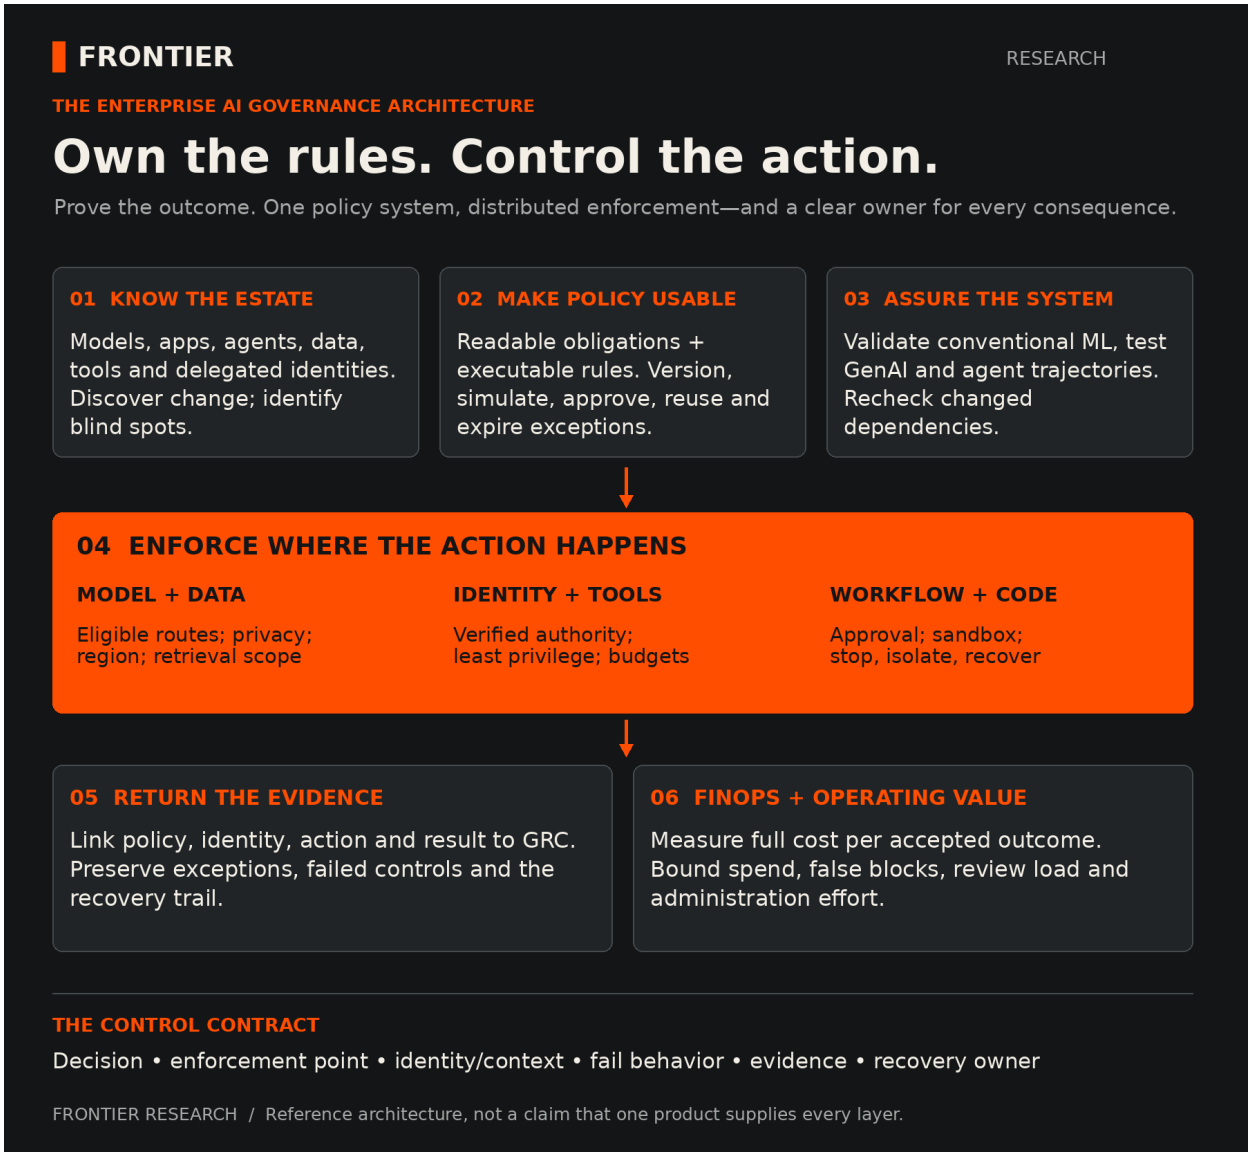
Capability	Non-agentic requirement	Additional agentic requirement	Evidence to demand
Discovery and inventory	Models, applications, embedded SaaS AI, owners and deployment versions	Agents, subagents, tools, MCP servers, delegated identities and execution routes	Reconcile discovered assets against cloud, SaaS, identity and development records; show what is invisible
Business accountability	Purpose, affected people, risk tier, approvals and exceptions	Named sponsor, delegated decision rights and limits on autonomy	An owner and time-limited authorization for a real production use case
Policy management	Versioned obligations, reusable controls, jurisdiction and domain overlays	Executable rules for tools, resources, budgets, approvals and delegation depth	Change, simulate, approve, distribute and revert one policy without editing every application

Validation and release	Quality, fairness where relevant, privacy, security, drift and release evidence	Trajectory tests, malicious tool results, authority escalation and multi-step failure tests	A release gate tied to the tested model, prompt, tool schema and policy versions
Identity and access	User/service identities, least privilege and data permissions	On-behalf-of context, short-lived authority, subagent restrictions and revocation	Revoke the human sponsor and test existing sessions, cached credentials and queued work
Runtime control	Data filtering, endpoint eligibility and application-specific decision constraints	Pre-action authorization, spending limits, sandbox boundaries and human escalation	A prohibited action prevented at the system that executes it—not merely an alert
Evidence and GRC	Traceable approvals, control tests, incidents and audit exports	Causal chain from intent through delegation, policy decision, action and business result	Reconstruct a disputed outcome without relying on an unverifiable AI narrative
Economics	Allocation by use case, provider and business owner	Token/tool/compute budgets, loop limits, retries and delegated spend	A budget enforced under concurrency, with a defined overrun tolerance
Recovery	Model rollback, corrected records, customer remediation and retirement	Stop new actions, isolate in-flight work, compensate completed actions and revoke descendants	A rehearsed incident involving more than one vendor and a non-reversible action
Administration	APIs, access control, useful defaults and manageable workflow	Fast policy propagation, operator separation and resilient enforcement	Measured onboarding effort, policy maintenance work and failure behavior

This is a requirements framework, not a claim that every row belongs in one SKU. The governance platform should orchestrate and prove the control chain. IAM, application authorization, security tools, FinOps systems and the business workflow will often perform parts of it.

2. What a strong platform actually looks like

A strong platform begins with a graph of responsibility. It links the use case to its owner and obligations, the application to its models and data, and the agent to its identity and tools. When a model, dataset, supplier term or tool permission changes, the platform should identify which approvals and tests have become stale. A static list of assets cannot answer that question reliably.



Frontier enterprise AI governance architecture connecting policy, runtime controls, evidence and outcome economics

Figure 1. Own the rules. Control the action. Prove the outcome. Frontier's reference architecture connects business policy to distributed enforcement and returns evidence to the accountable owner. The architecture is a design requirement, not a representation of any single vendor's product.

Next comes a policy system that both technical and nontechnical owners can manage. Legal and risk teams need readable obligations and exceptions. Engineers need testable rules and APIs. Business leaders need understandable decision rights. The translation between these views must itself be versioned and reviewed. Let an AI assistant draft a policy; do not let the draft silently become production authority.

The enforcement layer should be distributed. Model gateways constrain inference routes and data movement. Identity systems constrain credentials. Tool gateways and application APIs constrain actions. Sandboxes constrain code and network access. Workflow systems constrain approvals and commitments. A central governance service coordinates the rules and evidence; it should not become the only synchronous dependency for every enterprise action.

This is why cross-model support is necessary but insufficient. Supporting ten LLM APIs does not establish visibility into embedded SaaS AI, a legacy risk model or an agent executing a local shell command. Equally, listing a SaaS integration may mean an inventory import, not an ability to intercept its decisions. Procurement must ask what each connector can discover, inspect, authorize, stop and recover.

The control contract: six questions that cut through the pitch

For each important control, record: the decision being governed; the enforcement location; the identity and context available there; the fail behavior; the evidence retained; and the recovery owner. If any field is missing, the control is not ready for an unconditional production assurance claim.

Consider an agent permitted to issue a customer refund up to a defined limit. A prompt that tells it to stay under the limit is guidance. A classifier that spots suspicious language is a useful detection layer. The dependable boundary is an authorization check in the refund path using verified identity, order state, cumulative limits and the approved policy version. It must prevent repeated small refunds from evading the aggregate limit. If human approval is required, the approved transaction must be bound to the actual amount and recipient so it cannot change between review and execution.

The policy can be deterministic even when the model's reasoning is probabilistic. That does not make every policy input certain: identity can be stale, classification can be wrong and context can be incomplete. Define what happens under uncertainty. A harmless search may continue with reduced capability; a financial commitment may stop. “Fail closed” is not a substitute for a business continuity design.

Recovery is a business capability, not an undo button

The ability to deactivate an agent is valuable. It does not recall an email, reverse a disclosure or automatically restore every downstream record. Distinguish stopping future execution, cancelling queued work, revoking authority, reversing reversible transactions and compensating for irreversible consequences. Each needs a named owner and a tested path.

Frontier's earlier AI Passport analysis made a related point: a portable credential does not automatically create portable business authority or consistent enforcement in another platform. The receiving system must recognize the scope and actually apply it. The same principle applies to a policy exported from a governance console.[5]

GRC integration must carry evidence in both directions

A link to a GRC ticket is the beginning. The useful integration maps obligations to controls, controls to implementations, implementations to tests and operational evidence, and exceptions to accountable acceptances with expiry dates. An incident should invalidate the appropriate assurance claim and trigger reassessment. A revised obligation should identify affected systems and policies.

Avoid constructing a second, conflicting enterprise risk register inside the AI product. Designate authoritative records for obligations, identities, assets, decisions and evidence. Give internal audit an exportable view that includes failed controls and accepted exceptions, not just successful checks. Preserve the difference between first-line operation, second-line oversight and independent assurance.

FinOps belongs inside the boundary

An agent's cost is not just inference. Include tool fees, retrieval, storage, observability, guardrails, evaluations, review labor, retries, remediation and the cost of waiting for approval. A budget alert after a runaway loop is accounting evidence, not spending control. Enforce per-task and aggregate limits with a defined concurrency allowance; make delegated agents inherit the parent budget rather than minting new spending authority.

FinOps guidance distinguishes resource measures from business units and recognizes the fragmented economics of AI. Frontier's token-routing research takes that further: the useful denominator is an accepted business outcome. Governance should report both its own operating cost and its effect on the quality and cost of those outcomes.[6][7][8]

Do not promise a universal ROI percentage. Compare the full workflow before and after: approval lead time, exception load, accepted outcomes, adverse outcomes and total cost. A lower token bill can coexist with a more expensive operation.

Ease of administration is a control property

If every onboarding requires bespoke professional services, teams will route around the platform. If every alert needs an expert, the queue will become the policy. Ask for elapsed time and staff hours to onboard three different environments, establish one shared policy, add a jurisdictional exception and upgrade an agent framework without losing enforcement.

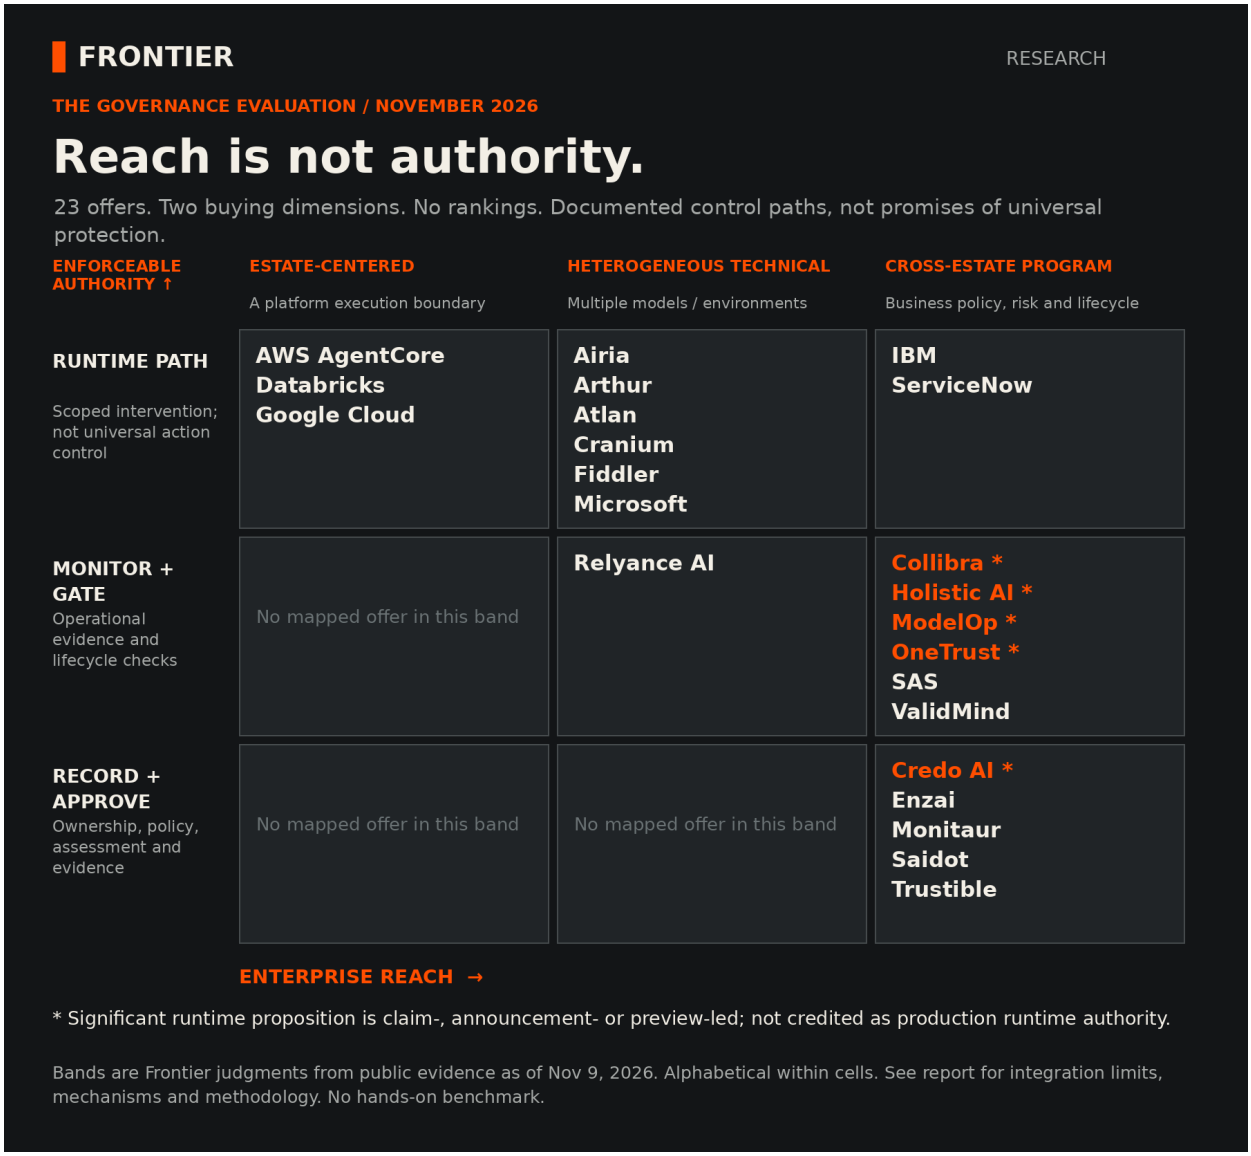
Measure policy reuse, stale approvals, connector freshness, false blocks, review capacity and time to revoke authority. Separate one-time implementation from the continuing cost of integrations, evaluations and policy maintenance. A product that is easy to demonstrate can still be difficult to operate.

3. Two dimensions that matter more than a feature count

Our evaluation uses enterprise reach and enforceable authority. These are independent dimensions. More breadth does not establish deeper control, and deeper local control does not establish enterprise coverage.

Enterprise reach asks how far the governance context extends. The map uses three bands: an estate-centered control service; heterogeneous technical coverage; and a cross-estate business governance program. The last band requires a substantive proposition for lifecycle ownership, policy, risk and evidence beyond one execution environment. It does not mean universal discovery or protection. A registry can have broad reach while its runtime controls remain narrow.

Enforceable authority asks what public evidence supports today. Its three bands are record and approve; monitor and gate; and a documented runtime intervention path. The upper band includes scoped blocking, externally applied guardrail decisions or supported agent deactivation. These are materially different mechanisms, identified in the profiles. It is not a claim of comprehensive pre-action enforcement, automatic rollback or identical coverage across products.



Frontier qualitative evaluation of 23 AI governance offers by enterprise reach and enforceable authority

Figure 2. Reach is not authority. Twenty-three offers positioned in qualitative evidence bands. Names within a cell are alphabetical, not ranked. An orange asterisk identifies a significant runtime proposition supported principally by a claim, announcement or preview; it does not lift the vendor into the runtime band. No distance between names implies a measured performance difference.

How to read—and challenge—the map

The map evaluates product propositions, not company quality or financial strength. It is deliberately conservative about runtime maturity. A vendor may sell additional capabilities not sufficiently documented in our reviewed evidence. Conversely, a documented capability may require a particular edition, deployment, connector, configuration or separately licensed component.

We mapped 23 offers to expose the architectural tradeoffs, including three estate-centered services as boundary cases. The broader directory includes security, observability, data, GRC and application platforms that can supply parts of the architecture. Inclusion in the map is not a purchasing recommendation; exclusion from it is not a negative rating.

Our evidence order is implementation documentation and release notes first, product documentation second, and announcements or marketing claims third. A research/private preview does not establish general availability. An acquisition agreement does not establish a completed acquisition, much less integrated products. A protocol logo does not establish policy portability. We did not conduct vendor interviews, hands-on comparative testing, security audits, customer-reference calls or pricing negotiations. These are architectural assessments from public evidence, not measured production results.

4. The comparative field: where the control boundary sits

IBM — broad governance context; inspect the runtime destination

IBM's watsonx.governance combines a substantial enterprise governance proposition with expanding agent discovery. Its October update connects agent governance to development and broadens external asset visibility. The important distinction is that runtime control synchronization is described for watsonx Orchestrate, while external monitoring integrations can supply metrics and trigger governance workflows. Those are different control paths.[9][10]

Frontier assessment: a credible candidate for organizations seeking continuity between established model governance and agent accountability. Buying test: take a policy across an IBM agent and two non-IBM runtimes; show which actions are prevented, which generate a workflow, and which remain outside control. The map credits a bounded runtime path, not uniform enforcement across every discovered asset.

ServiceNow — governance attached to the operating workflow

AI Control Tower's attraction is its proximity to enterprise workflow, ownership and remediation. Release notes describe deactivation and reinstatement for supported external agent environments as well as ServiceNow agents, with scope limitations. Edition matters: the broader enterprise proposition should not be confused with a Now Assist-focused configuration. The Control Tower Alliance expands the integration agenda; partner participation is not evidence that every integration is qualified and deployed.[11][12][13]

Frontier assessment: strategically strong where ServiceNow already coordinates operational responsibility. Buying test: deactivate a supported external agent, then inspect queued work, credentials and downstream effects. Platform deactivation is not the same as authorizing each tool call or reversing completed actions. Require an exportable decision trail so governance does not become inseparable from one workflow platform.

OneTrust — privacy and risk context moving toward runtime

OneTrust's proposition joins AI governance to established privacy and risk context. Its CoRIE announcement describes runtime decisions using organizational context, including allowing, stopping or escalating agent tool activity. That is strategically relevant: a tool call may be technically permitted but inconsistent with consent or purpose. The reviewed launch material does not establish a universal, generally available connector and deployment matrix.[14][15]

Frontier assessment: compelling where existing privacy obligations should become operational constraints. The map credits governance and monitoring, with the runtime expansion marked as announcement-led. Buying test: demonstrate a consent withdrawal propagating into a running agent, identify the actual enforcement point, and document tenant availability and behavior when context cannot be retrieved.

Credo AI — policy discipline, with a distinct runtime preview

Credo AI emphasizes inventory, policy management, risk assessment and governance workflows. Its Agent Governor is explicitly a research preview. The reviewed page packages support for Claude Code and identifies additional policy and audit capabilities as forthcoming. The runtime proposition is therefore a direction to test, not evidence that Credo's entire governance portfolio already controls arbitrary agent execution.[16][17]

Frontier assessment: relevant for formalizing a reusable AI governance program across business units and suppliers. Buying test: show how an approved policy becomes a tested technical rule, how exceptions propagate, and which production runtimes are contractually supported. Evaluate the governance platform and preview separately; do not let enthusiasm for the latter erase the value or limitations of the former.

Holistic AI — ambitious lifecycle breadth; verify the intervention claims

Holistic AI describes discovery, inventory, testing, monitoring and agent-oriented control capabilities, including graph-based visibility and runtime interventions. This breadth is attractive because assurance before deployment and observation after deployment should inform one another. However, the reviewed product-level claims alone do not establish consistent blocking, revocation or recovery across an enterprise estate.[18]

Frontier assessment: a broad platform proposition whose enforcement coverage warrants unusually concrete diligence. The map places it in monitor-and-gate, with a runtime-claim marker. Buying test: choose an unsupported-looking edge case—an external SaaS agent or a locally executed tool—and require a live demonstration, supported-connector list and failure contract. Do not substitute an attractive agent graph for a tested intervention.

ModelOp — lifecycle controls meeting the execution environment

ModelOp's Enterprise AI Command Center emphasizes portfolio visibility, lifecycle controls and operating evidence. Its November 4 NVIDIA OpenShell integration announcement addresses a meaningful gap: connecting enterprise policy to execution permissions rather than leaving governance at the model record. The announcement is current evidence of product direction, not an independent production validation.[19][20]

Frontier assessment: relevant for enterprises that want governance to coordinate heterogeneous technical controls without replacing every development platform. The map credits lifecycle monitoring and gates, while flagging the new runtime proposition. Buying test: change a risk classification or approval and observe its propagation into the sandbox and tool gateway. Identify who owns the mapping, supported versions and fail behavior when the governance service is unavailable.

Collibra — data context is an advantage; preview semantics matter

Collibra's AI Command Center brings AI governance close to data ownership and lineage. Its Guardian Agent documentation is particularly instructive: the capability is a private preview; enforcement requires deployment instrumentation; and an escalation decision blocks without a completed approval flow in the reviewed documentation. Checking a tool result after execution cannot undo the tool's side effect.[21][22]

Frontier assessment: strong strategic relevance where trusted data context is the missing foundation. The map does not promote private-preview controls into a production runtime rating. Buying test: distinguish pre-tool authorization from post-tool inspection, test the actual approval path, and require evidence that intended agents are instrumented. Enabling a feature globally is not proof that every execution route uses it.

Microsoft — formidable native controls, uneven external protection

Microsoft combines Agent 365, identity, security and Purview capabilities across a large enterprise footprint. Its own Purview documentation makes a crucial distinction: visibility into AI agents and protection/compliance support are different columns, with coverage depending on the parent application. External platforms do not uniformly inherit the controls available in Microsoft's native environment.[23][24]

Frontier assessment: an important architectural anchor for Microsoft-centered enterprises, but not evidence of automatic control over every discovered external agent. Buying test: name the exact host, license, identity and data path for each required control. Use the coverage matrix rather than a suite-level promise. The map credits heterogeneous reach and bounded runtime paths, not a single universal Microsoft governance product.

SAS — retain the non-agentic governance discipline

SAS AI Navigator addresses AI inventory and governance foundations; Model Manager supplies established model lifecycle capabilities. This matters because an enterprise-wide program must govern conventional analytical systems alongside newer agents. These are distinct products and responsibilities, not automatically one entitlement or one enforcement architecture.[25][26]

Frontier assessment: relevant where regulated analytics and model operations remain material. Buying test: reconcile the model inventory with agent and third-party application inventories, then show which approval and monitoring controls carry across them. Public evidence reviewed here does not establish a general cross-vendor agent tool-authorization layer. That gap should drive integration requirements rather than dismissal of mature model governance.

Saidot — make dependencies and responsibility visible

Saidot emphasizes relationships among AI systems, models, datasets, risks and governance responsibilities. Its agent catalogue documentation provides a concrete integration example with Azure AI Foundry. This is useful for understanding how a change to a dependency affects several AI systems.[27][28]

Frontier assessment: a relevant program-governance option for organizations that need a connected inventory and practical accountability. Buying test: discover a changed agent configuration, identify affected approvals and trigger the appropriate operational response. APIs and MCP connectivity can improve access to governance context; they do not by themselves prove interception of production actions. The map therefore credits record-and-approve authority.

Monitaur — evidence continuity across the AI lifecycle

Monitaur presents a lifecycle governance platform spanning policies, models and AI systems, including newer generative and agentic uses. Its value proposition is continuity of governance evidence and responsibility across a portfolio.[29]

Frontier assessment: relevant where the principal failure is inconsistent ownership, review and auditability rather than lack of another runtime filter. Buying test: connect a failed production control to the corresponding approval, exception and remediation record. We did not establish a general cross-platform tool-interception layer from the reviewed public material. Treat runtime integration as a procurement requirement to demonstrate, not as an inferred feature of lifecycle coverage.

ValidMind — independent validation remains consequential

ValidMind focuses on model-risk and AI governance workflows, validation evidence and documentation. Its relevance is especially clear where models affect regulated or consequential decisions and the organization must explain the basis for approval.[30]

Frontier assessment: a serious part of the architecture when independent challenge and reproducible validation matter. Buying test: extend validation from a static model to a changing agent workflow with tool and policy dependencies. Determine which changes invalidate prior approval. The map credits evidence, monitoring and release discipline; it does not infer real-time action control from strong validation tooling.

Enzai — governance of agent permissions needs an execution partner

Enzai's agentic governance material explicitly addresses autonomy, permitted actions, escalation and coordination. These are the right organizational questions: what authority has been delegated, and under which conditions?[31]

Frontier assessment: relevant for translating policy into an intelligible governance program. Buying test: take an action whitelist from the governance record into a production tool and demonstrate enforcement, version synchronization and exception handling. We place the reviewed proposition in record-and-approve because documenting or assessing an action boundary is not sufficient evidence that every relevant runtime enforces it.

Trustible — policy and evidence as reusable enterprise assets

Trustible focuses on operationalizing AI policy through inventories, assessments, workflows and evidence. Its strategic opportunity is to make governance repeatable across teams instead of forcing every project to invent a new approval process.[32]

Frontier assessment: relevant where fragmented review is delaying useful AI adoption. Buying test: reuse one control across several systems, change it once and show which systems require reassessment. Separate automated collection of evidence from automated prevention of an action. The reviewed public material supports program governance more clearly than a universal runtime mediation claim.

Cranium — discover the technical estate and test its defenses

CraniumONE emphasizes AI discovery, security testing and governance across infrastructure and model environments. Guardian documentation provides a runtime integration path and makes initial listening behavior explicit. Deployment mode matters: observation should not be mistaken for active prevention.[33][34]

Frontier assessment: relevant when shadow AI and technical exposure are the immediate problem. The map credits a bounded runtime path. Buying test: move a deployment from listening to enforcement, measure false blocks, and show how technical findings attach to business owners and risk acceptances in the enterprise GRC system. Do not assume security inventory alone captures all business use cases.

Relyance AI — connect AI behavior to data exposure

Relyance brings data context, discovery and AI evaluation into the governance discussion. Its materials distinguish read-only discovery mechanisms from evaluation and runtime capabilities. That distinction is useful: understanding where sensitive information can flow is not equivalent to blocking the flow.[35][36]

Frontier assessment: relevant where AI governance is inseparable from data usage and privacy. The map conservatively credits technical visibility and monitoring. Buying test: trace a prohibited data use from discovery to a concrete enforcement point, then identify the product, deployment and integration required. Require a clear answer about which content is inspected, retained or excluded; governance telemetry can itself create a sensitive data store.

Airia — governance close to orchestration

Airia combines governance with an orchestration and security proposition. Its September announcement specifically describes generally available real-time enforcement for Claude; its runtime documentation also identifies policy-scope exclusions for certain system agents. Those details are more useful than an undifferentiated “all agents” claim.[37][38][39]

Frontier assessment: relevant for organizations willing to govern supported execution paths through the platform's integration model. Buying test: enumerate externally hosted agents, included system activity and bypass routes. Establish how traditional model governance and enterprise GRC evidence are maintained alongside agent controls. The map credits documented runtime scope without assuming that all enterprise activity passes through it.

Fiddler — behavior evidence and runtime guardrails

Fiddler connects model and agent observability with guardrails and governance-oriented evidence. Its runtime proposition includes inspecting inputs and outputs; that is a useful technical control, but it should not be confused with business authorization for every downstream action.[40][41]

CrowdStrike announced an agreement to acquire Fiddler on November 3. The reviewed notice describes a proposed transaction and future product direction, not completed integration. This report therefore evaluates Fiddler's documented offer separately.[42]

Frontier assessment: relevant where behavior monitoring, evaluation and runtime safeguards are central. Buying test: connect a detected failure to the business decision and remediation workflow, and obtain current support and roadmap commitments during the corporate transition. Do not buy a future combined product as if it already exists.

Arthur — the application still has to enforce the verdict

Arthur describes discovery across major cloud environments, model and agent evaluation, and guardrails. Its explanation of the control path is especially important: the guardrail can return a decision that the application enforces inline. The decision service and the enforcement implementation are separate responsibilities.[43][44]

Frontier assessment: relevant for engineering-led teams seeking behavior assurance across model and agent deployments. Buying test: deliberately ignore or time out the guardrail response and show what the application does. Secure and test the caller, not just the classifier. The map credits the documented mediated runtime path, with the application integration explicitly inside the buyer's responsibility boundary.

Atlan — data context plus an agent execution boundary

Atlan connects AI governance to data ownership, lifecycle context and lineage. Its product documentation distinguishes the core context layer from Trident, an agent/skill registry with traces, evaluations and cost records. Trident routes model and tool calls through an AI Gateway to control access.[45][46]

Frontier assessment: relevant where the enterprise needs both governed context and a registry for technical agent activity. The map credits a documented routed runtime path and heterogeneous technical reach, not universal application coverage. Buying test: establish which product owns each policy and prove that the agent cannot use an alternative model or tool route. Reconcile the technical registry with business-purpose approvals and the enterprise GRC record.

Three estate-centered boundary cases: AWS, Google Cloud and Databricks

AWS made AgentCore Policy generally available in March 2026. Its Cedar-based policy enforcement at the AgentCore Gateway is a concrete example of externalizing authorization from agent reasoning. It governs the path that uses that gateway; model or framework flexibility does not make it an inventory and GRC system for all enterprise AI.[47]

Google Cloud documents Agent Gateway coverage with explicit availability distinctions. The reviewed table lists available Gemini Enterprise Agent Platform paths and other environments as coming soon. Evaluate the available deployment, not the eventual cross-cloud ambition.[48]

Databricks connects AI Gateway governance to its data and identity environment. This can be valuable where data, models and agent activity already converge. Confirm the actual endpoints, workspace scope and policies covered rather than extrapolating to every external application.[49]

These three appear in the map to make a purchasing point: deep local controls are valuable building blocks. They become enterprise governance only when linked to portfolio accountability, consistent policy and evidence outside their own estate.

5. The industry's maturity: useful parts, unfinished system

The market is more mature at documenting responsibility than at maintaining it through a changing chain of autonomous execution. That is our architectural judgment from the reviewed evidence, not a statistical finding about customer deployments.

Capability area	Frontier maturity judgment	What remains unresolved
Inventory, assessment and policy workflows	Established product category; implementation quality varies	Completeness of discovery, embedded SaaS AI and stale ownership
Traditional model validation and lifecycle evidence	Established in specialist and platform offerings	Consistent extension into changing agent trajectories
GenAI evaluation and observability	Substantial tools and documented deployments; workload-specific validity	Reliable labels, evaluator bias, real-world task coverage and evidence cost
Runtime model and tool safeguards	Real, but fragmented by route, protocol and integration	Bypass resistance, fail behavior, side effects and configuration drift
Cross-vendor delegated authority	Emerging	Shared scope semantics, identity propagation, revocation and policy portability
Business recovery and outcome economics	Often distributed across other systems	End-to-end accountability, compensation and accepted-outcome measurement

The acquisition wave reinforces convergence, not completion. Security vendors want behavior context; observability vendors want a place in the control decision; GRC vendors want continuous evidence. Buyers should resist the inference that common ownership instantly produces a common architecture. Product identities, supported integrations, data terms and pricing can remain separate after a deal closes.

“Enterprise-wide” should therefore be treated as an acceptance test. Ask vendors to show coverage across a predictive model, an externally hosted assistant, an internally built agent and an agent embedded in SaaS. Require the same business policy and traceable identity across them. The right answer may involve multiple enforcement products and one governance record. That is acceptable if responsibilities and evidence are coherent.

Regulatory mapping is also a maintenance problem. Obligations differ by jurisdiction, role, system use and implementation date. The European Commission's live AI Act timeline illustrates why a static checklist is inadequate. Require versioned obligations, effective dates and legal review of applicability; a platform's mapping does not certify the enterprise's compliance.[50]

6. Buy the missing control, not another inventory

There are three sensible starting architectures. None is universally best.

Program-first: use an AI governance or GRC-oriented platform as the authoritative record for ownership, obligations and approval, then connect existing identity, data, security and runtime controls. This fits a heterogeneous company with fragmented accountability. The risk is spending a year improving documentation while dangerous execution paths remain open. Tie the first release to working enforcement in one consequential workflow.

Runtime-first: start with gateways, agent security and platform controls where agents already operate, then connect decisions and evidence to the enterprise governance program. This fits an organization with active agent deployments and immediate technical exposure. The risk is accumulating several incompatible definitions of policy and exception. Establish common identifiers and evidence contracts before expanding.

Estate-first, federated outward: use the strongest controls inside a dominant cloud or application platform, with a separate enterprise record and explicit boundaries for other estates. This can reduce initial integration effort. The risk is letting the dominant platform define what counts as governed. Maintain an independent coverage register and test exits before dependency grows.

The CIO should own the architecture and accountable integration, not every business decision. Business owners authorize outcomes and risk within their remit. Security owns technical protection and response. Data leaders own data controls. Legal and risk define obligations and challenge assumptions. Finance validates economics. Audit needs independent evidence. Frontier's CIO 2030 frames this as allocation of intelligence and responsibility across the enterprise; AI governance is the operating mechanism that makes that allocation credible.[51]

A procurement proof suite

Run the same tests on the proposed configuration, including partner components. Record the result, latency, operator effort and evidence—not just a pass/fail demo.

1. Find the unknown: introduce an unregistered model, SaaS assistant or agent. Measure discovery delay and explain the blind spots.
2. Change the system: swap a model, prompt, retrieval source or tool schema. Show which approval and tests become invalid.
3. Break the authority chain: revoke the sponsoring user and test active sessions, delegated agents, cached credentials and queued actions.

4. Attack through context: place hostile instructions in a retrieved document or tool result. Test whether the action boundary holds even if the model follows them.
5. Exceed the budget: run concurrent agents, retries and subagents against a shared limit. Measure the maximum possible overshoot.
6. Remove the governance service: simulate policy, network and telemetry failures. Confirm the specified reduced mode, stop behavior and recovery.
7. Dispute a decision: reconstruct the actual identity, policy, evidence, tool request and business effect. An AI-generated explanation is not a substitute for the underlying record.
8. Contain and recover: stop further action, handle queued work and compensate for one irreversible effect. Measure business recovery, not only alert delivery.
9. Move the workload: export policies, identifiers, evidence and evaluations. Re-create the important control in a second environment.
10. Operate it: have the intended administrator update a policy and handle exceptions without the vendor's demonstration team.

    Define thresholds before testing. Discovery and revocation times should reflect the business exposure. False-block and latency tolerances should reflect the workflow. Do not accept an aggregate score that lets excellent documentation compensate for a failed mandatory authorization control.

The first 90 days

Days 1–30: establish the truth. Select three representative systems: a conventional model, a generative assistant and an agent with limited write authority. Reconcile inventories, name owners, map data and actions, and choose a single consequential policy. Establish baseline cost and accepted outcomes. Document where the organization cannot currently observe or intervene.

Days 31–60: connect the rule to the action. Implement the policy at the real enforcement points. Integrate evidence into the existing risk and incident workflow. Run the proof suite, including outages and bypass attempts. Measure both blocked bad actions and disrupted legitimate work. Retire duplicate records rather than adding another permanent spreadsheet.

Days 61–90: expand only what has earned expansion. Review results with the business owner, security, risk and finance. Extend authority where controls and recovery work. Keep other workflows read-only or human-approved. Negotiate pricing, evidence export, policy portability, support responsibilities and connector maintenance before broad deployment.

Frontier's Future of Work research argues for designing work around intent, evidence and accountable outcomes. The same discipline applies here. A governance purchase should change how the enterprise delegates and verifies work—not merely give the existing approval process a more impressive dashboard.[52]

7. The decision the CIO cannot outsource

The strategic prize is not an AI system that never makes a mistake. It is an enterprise that can grant useful authority, detect when the conditions for that authority change, and contain the consequences without losing control of its customers, finances or obligations.

Choose the governance platform that can connect your business rules to your actual execution estate. Make every broad claim resolve into a supported control path. Keep ownership of policy, evidence and accepted outcomes even when you rent the models, workflows and enforcement technology.

An audit trail can explain a failure. The next generation of AI governance must also help prevent the unauthorized action—and make the authorized one easier.

Methodology and limitations

Evidence cutoff: November 9, 2026. Method: broad public-web discovery followed by primary-source review of product documentation, release notes, corporate notices and standards guidance; architectural synthesis; qualitative positioning against enterprise reach and enforceable authority. Prior Frontier research supplies the strategic framework, not independent verification of vendor claims. The research covers public offers visible in this review and is not a guarantee that every provider worldwide has been identified.

We prioritize enterprise relevance across conventional ML, generative AI and agents, heterogeneous environments, policy administration, lifecycle evidence, runtime controls, GRC integration and operating economics. Adjacent security, observability and estate-specific services are included to expose integration choices. We do not use market share, customer counts, analyst rankings or vendor-reported savings to establish comparative superiority.

No original survey, vendor interview, lab benchmark, production security audit, customer-reference study or commercial quote comparison was conducted. Public documentation varies in depth; a lower evidence band can reflect incomplete public information. Availability is feature-, edition-, region- and connector-specific. Descriptions of launches and previews retain those limitations. Vendor marketing is evidence of a proposition, not proof of its effectiveness. Positioning is Frontier's judgment and should be challenged through the common proof suite.

This draft was developed with AI-assisted source collection, synthesis and visual construction for named-author editorial review. Source links support factual claims; strategic interpretations and the evaluation framework are Frontier analysis. Regulatory discussion is an operational purchasing lens, not a legal opinion. Product and ownership status can change after the cutoff.

Copyright 2026 Frontier Research. Research inquiries: inquiries@thisisfrontier.ai. Reprints and visual licensing: reprints@thisisfrontier.ai.

Addendum: solutions considered

The following directory records the wider research universe. It includes full governance propositions and adjacent components because enterprise coverage usually requires both. Entries are not endorsements or a ranking. Product families and acquired businesses are consolidated where practical; the number of entries is not a count of independent corporations. Descriptions identify the role evaluated, not an assertion that each offer supplies every capability in this report. Links point to vendors or official product documentation.

1. IBM — Enterprise governance. Lifecycle risk, AI inventory and agent governance; external discovery and runtime enforcement have different boundaries.[9]

2. ServiceNow — Enterprise governance. AI Control Tower connects accountability, workflow and supported agent interventions; editions and connector scope matter.[11]

3. OneTrust — Enterprise governance. Privacy and risk context for AI governance, with an announced CoRIE runtime expansion; establish production availability.[14]

4. Credo AI — Enterprise governance. Policy-centered governance and evidence workflows; Agent Governor is a separate research preview.[17]

5. Holistic AI — Enterprise governance. Discovery, assurance and lifecycle governance with broad agent-control claims requiring deployment-level proof.[18]

6. ModelOp — Enterprise governance. Portfolio and lifecycle control coordination, with newly announced execution-policy integrations.[20]

7. Collibra — Enterprise governance. Data-connected AI governance; Guardian Agent runtime controls remain a private-preview diligence item.[22]

8. Microsoft — Enterprise platform. Agent 365, Purview and related identity/security controls; visibility and protection differ by host application.[23]

9. SAS — Enterprise governance. AI Navigator and model-management products connect program inventory and conventional model lifecycle disciplines.[25]

10. Saidot — Enterprise governance. Connected AI inventory, risk and dependency governance, with agent catalogue integrations.[28]

11. Monitaur — Enterprise governance. Lifecycle evidence and policy management across an organization's AI portfolio.[29]

12. ValidMind — Model risk and governance. Validation, model-risk documentation and governance workflows; especially relevant to consequential decision systems.[30]

13. Enzai — Enterprise governance. AI policy and risk workflows addressing agent autonomy, action boundaries and escalation.[31]

14. Trustible — Enterprise governance. Reusable policy, assessment and evidence workflows for an enterprise AI program.[32]

15. Cranium — Technical governance and security. AI discovery, security testing and runtime integration; listening and enforcement modes must be distinguished.[34]

16. Relyance AI — Data and AI governance. Data context, AI discovery and evaluation; establish the enforcement path separately from read-only discovery.[35]

17. Airia — Orchestration and governance. Agent orchestration, governance and supported runtime controls, including a specific Claude GA announcement.[38]

18. Fiddler — Observability and guardrails. Model/agent behavior evidence and guardrails; CrowdStrike acquisition agreement announced, not treated here as completed integration.[42]

19. Arthur — Observability and guardrails. Evaluation, discovery and guardrail verdicts; application enforcement remains an explicit integration responsibility.[43]

20. Atlan — Data context and AI governance. Data-governance context plus Trident agent registry and gateway-routed model/tool access; verify product and route boundaries.[46]

21. Amazon Web Services — Estate-centered runtime. AgentCore Policy enforces authorization through its gateway; a building block rather than a complete enterprise GRC program.[47]

22. Google Cloud — Estate-centered runtime. Agent Gateway and related controls; available and coming-soon deployment coverage must remain separate.[48]

23. Databricks — Estate-centered runtime. AI Gateway governance close to the data, model and identity environment; verify endpoint and workspace coverage.[49]

24. SAP — Enterprise platform. AI Gateway and governance announcements connect agents with enterprise process context; confirm orderability and integration scope.[53]

25. Dataiku — AI lifecycle platform. Governance workflows and model/agent lifecycle records; the Govern deployment is a distinct architectural component.[54]

26. DataRobot — AI lifecycle platform. Model and AI lifecycle governance within a broader development and operations platform; validate external-estate depth.[55]

27. Archer — GRC platform. AI use-case, risk and control governance linked to enterprise risk management; runtime enforcement needs technical integrations.[56]

28. MetricStream — Adjacent GRC platform. AI-assisted GRC and risk workflows; AI used for GRC should not be mistaken for comprehensive governance of all AI execution.[57]

29. Diligent — Adjacent GRC platform. Governance and risk workflows with AI capabilities; verify the distinction between GRC automation and runtime AI control.[58]

30. Optro (formerly AuditBoard) — GRC platform. Risk and assurance platform under the new Optro identity; evaluate AI governance and inherited FairNow capabilities as one vendor family.[59]

31. TrustArc — Privacy and governance. AI governance connected to privacy obligations and assessment processes; not assumed to mediate all runtime traffic.[60]

32. Drata — Trust and compliance. Announced continuous agent inventory and monitoring expansion; reviewed announcement specifies limited availability.[61]

33. Cisco — AI security. AI Defense covers discovery, validation and runtime security; link security controls to business governance and GRC.[62]

34. Palo Alto Networks — AI security. Agent and AI security, including acquired Protect AI capabilities; evaluate the current combined portfolio rather than duplicate historical vendors.[63]

35. Check Point / Lakera — AI security. Lakera-based AI security and runtime protections within Check Point's broader security proposition.[64]

36. CrowdStrike — AI security. Security platform pursuing agent behavior and AI trust capabilities; announced Fiddler and Dreadnode deals are not proof of delivered integration.[65]

37. Dynatrace / Arize — Observability and evaluation. AI tracing, evaluation and observability; Dynatrace reports the Arize acquisition completed. Evidence collection is not a full GRC program.[66]

38. Snowflake / TruEra — Data and AI assurance. TruEra's move into Snowflake places model/AI assurance near the data platform; verify the current product and entitlement.[67]

39. Salesforce — Application platform. Agentforce's integrated trust architecture governs supported Salesforce execution paths; not assumed to govern unrelated enterprise agents.[68]

40. UiPath — Automation platform. AI Trust Layer policies manage supported third-party model usage across UiPath organizations, tenants and groups.[69]

41. NVIDIA — Execution building block. OpenShell supplies an execution-environment component; the buyer still needs enterprise policy, accountability and evidence integration.[70]

42. WitnessAI — AI security and runtime. Agent/model traffic controls; the vendor explicitly scopes enforcement to traffic traversing its platform.[71]

43. Zenity — Agent security. Discovery, governance and defense for agents across supported SaaS, cloud and endpoint environments.[72]

44. HiddenLayer — AI security. AI discovery, attack testing and runtime defense; establish how security findings become accountable business controls.[73]

45. F5 / CalypsoAI — AI security and gateway. F5 completed the CalypsoAI acquisition and introduced guardrails/red-team capabilities; route and deployment scope remain purchasing questions.[74]

46. Kong — Gateway and policy enforcement. Model, MCP and agent traffic controls; policies differ by traffic type and govern the routes that use the gateway.[75]

SOURCES / CITATIONS

Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

01 **Frontier: The Autonomous Enterprise**
https://www.thisisfrontier.ai/research/the-autonomous-enterprise

02 **NIST: AI Risk Management Framework**
https://www.nist.gov/itl/ai-risk-management-framework

03 **NIST: Generative AI Profile**
https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

04 **OWASP: Top 10 for Agentic Applications 2026**
https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/

05 **Frontier: ServiceNow AI Passport analysis**
https://www.thisisfrontier.ai/news/servicenow-s-ai-passport-turns-agent-governance-into-a-platform-bet-32f19955

06 **FinOps Foundation: Unit economics**
https://www.finops.org/framework/capabilities/unit-economics/

07 **FinOps Foundation: AI technology category**
https://www.finops.org/framework/technology-categories/ai/

08 **Frontier: Rent the Models. Own the Rules**
https://www.thisisfrontier.ai/research/rent-the-models-own-the-rules-ai-token-routing

09 **IBM: October 2026 watsonx.governance capabilities**
https://www.ibm.com/new/announcements/new-in-ibm-watsonx-governance-govern-ai-agents-as-you-build-expand-accountability-and-streamline-compliance

10 **IBM: AI asset discovery expansion**
https://www.ibm.com/new/announcements/expanding-ai-asset-discovery-for-broader-visibility

11 **ServiceNow: AI Control Tower release notes**
https://www.servicenow.com/docs/r/release-notes/ai-control-tower-rn.html

12 **ServiceNow: AI Control Tower editions and workflows**
https://www.servicenow.com/docs/r/intelligent-experiences/ai-control-tower/configuring-ai-control-tower-workflows.html

13 **ServiceNow: Control Tower Alliance announcement**
https://newsroom.servicenow.com/press-releases/details/2026/ServiceNow-launches-Control-Tower-Alliance-to-help-enterprises-adopt-AI-safely-at-scale/default.aspx

14 **OneTrust: CoRIE announcement**
https://www.onetrust.com/news/onetrust-introduces-corie-to-govern-ai-agents-at-machine-speed/

15 **OneTrust: AI governance platform expansion**
https://www.onetrust.com/news/onetrust-unveils-platform-and-product-innovations-to-govern-ai-at-enterprise-scale/

16 **Credo AI: Governance platform**
https://www.credo.ai/product

17 **Credo AI: Agent Governor research preview**

https://www.credo.ai/agent-governor

18 **Holistic AI: Governance platform and runtime claims**

https://www.holisticai.com/ai-governance-platform

19 **ModelOp: Enterprise AI Command Center**

https://www.modelop.com/enterprise-ai-command-center

20 **ModelOp: November 4 OpenShell integration announcement**

https://www.globenewswire.com/news-release/2026/11/04/3395334/0/en/modelop-unveils-industry-first-full-stack-ai-policy-enforcement-with-nvidia-openshell.html/

21 **Collibra: AI Command Center**

https://www.collibra.com/products/ai-command-center

22 **Collibra: Guardian Agent private-preview limitations**

https://productresources.collibra.com/docs/release-notes/Content/AICommandCenter/GuardianAgents/co_guardian-agent.htm

23 **Microsoft: Purview agent coverage matrix**

https://learn.microsoft.com/en-us/purview/ai-agents

24 **Microsoft: Agent 365 May 2026 update**

https://techcommunity.microsoft.com/blog/agent-365-blog/what%E2%80%99s-new-in-agent-365-may-2026/4516340

25 **SAS: AI Navigator support**

https://support.sas.com/en/software/ai-navigator-support.html

26 **SAS: Model Manager support**

https://support.sas.com/en/software/model-manager-support.html

27 **Saidot: AI governance platform**

https://www.saidot.ai/

28 **Saidot: Agent catalogue documentation**

https://help.saidot.ai/knowledge-base/using-agent-catalogue

29 **Monitaur: AI governance platform**

https://www.monitaur.ai/platform

30 **ValidMind: Platform**

https://validmind.com/platform/

31 **Enzai: Agentic AI governance**

https://www.enz.ai/product/agentic-ai-governance

32 **Trustible: Platform overview**

https://trustible.ai/platform-overview/

33 **Cranium: CraniumONE platform**

https://cranium.ai/platform/

34 **Cranium: Guardian quickstart**

https://docs.cranium.ai/en/articles/24872-guardian-quickstart-guide

35 **Relyance AI: Platform capabilities**

https://docs.relyance.ai/docs/introduction-to-relyance-ai/relyance-ai-capabilities/

36 **Relyance AI: AI discovery and governance**

https://www.relyance.ai/use-cases/ai-governance

37 **Airia: Governance documentation**

https://airia.ai/docs/governance/overview

38 **Airia: Claude enforcement general availability**

https://www.globenewswire.com/news-release/2026/9/29/3370769/0/en/airia-makes-real-time-policy-enforcement-for-claude-generally-available.html

39 **Airia: Runtime policy scope**

https://airia.ai/docs/secure/runtime-security/managing-policies

40 **Fiddler: Guardrails**

https://www.fiddler.ai/guardrails

41 **Fiddler: Governance and risk capabilities**

https://www.fiddler.ai/ai-governance-risk-compliance

42 **CrowdStrike: Agreement to acquire Fiddler**

https://crowdstrike.gcs-web.com/news-releases/news-release-details/crowdstrike-acquire-fiddler-deliver-agent-observability-ai-trust

43 **Arthur: Agent discovery and governance**

https://www.arthur.ai/discover-and-govern-agents

44 **Arthur: Guardrails documentation**

https://docs.arthur.ai/docs/guardrails-overview

45 **Atlan: AI governance**

https://docs.atlan.com/product/capabilities/governance/ai-governance

46 **Atlan: Trident registry and AI Gateway**

https://docs.atlan.com/get-started/concepts/atlan-products

47 **AWS: AgentCore Policy general availability**

https://aws.amazon.com/about-aws/whats-new/2026/03/policy-amazon-bedrock-agentcore-generally-available/

48 **Google Cloud: Agent Gateway coverage and availability**

https://docs.cloud.google.com/gemini-enterprise-agent-platform/govern/gateways/agent-gateway-overview

49 **Databricks: AI Gateway governance**

https://docs.databricks.com/aws/en/ai-gateway/ai-governance

50 **European Commission: AI Act implementation timeline**

https://ai-act-service-desk.ec.europa.eu/en/ai-act/eu-ai-act-implementation-timeline

51 **Frontier: CIO 2030**

https://www.thisisfrontier.ai/research/cio-2030

52 **Frontier: The Future of Work in the AI Era**

https://www.thisisfrontier.ai/research/the-future-of-work-in-the-ai-era

53 **SAP: AI governance announcements at TechEd**

https://news.sap.com/2026/10/sap-teched-shows-how-to-lead-with-ai-governance/

54 **Dataiku: Governance documentation**

https://doc.dataiku.com/dss/latest/governance/index.html

55 **DataRobot: AI governance**

https://www.datarobot.com/product/ai-governance/

56 **Archer: AI governance**
https://www.archerirm.com/ai-governance

57 **MetricStream: AI for GRC**
https://www.metricstream.com/metricstream-ai-for-grc.htm

58 **Diligent: AI governance and risk**
https://www.diligent.com/ai-governance-risk

59 **Optro: AuditBoard rebrand and platform direction**
https://optro.ai/blog/meet-optro-auditboard-unveils-new-identity-as-ai-transforms-grc

60 **TrustArc: AI governance**
https://trustarc.com/solutions/ai-governance/

61 **Drata: Agent governance limited availability**
https://drata.com/about/news/drata-extends-trust-management-platform-to-continuously-monitor-and-govern-ai-agents

62 **Cisco: AI Defense data sheet**
https://www.cisco.com/c/en/us/products/collateral/security/ai-defense/ai-defense-ds.html

63 **Palo Alto Networks: Agent security**
https://www.paloaltonetworks.com/ai-security/agent-security

64 **Check Point: Lakera acquisition**
https://www.checkpoint.com/press-releases/check-point-acquires-lakera-to-deliver-end-to-end-ai-security-for-enterprises/

65 **CrowdStrike: Dreadnode acquisition agreement**
https://ir.crowdstrike.com/news-releases/news-release-details/crowdstrike-acquire-dreadnode-advance-ai-trust-and-deliver/

66 **Dynatrace: Arize acquisition completed**
https://www.dynatrace.com/news/blog/dynatrace-completes-acquisition-of-arize/

67 **TruEra: Snowflake ownership notice**
https://truera.com/

68 **Salesforce: Trust and Agentforce**
https://help.salesforce.com/s/articleView?id=ai.copilot_trust.htm&language=en_US&type=5

69 **UiPath: AI Trust Layer policies**
https://docs.uipath.com/automation-cloud/automation-cloud/latest/admin-guide/managing-ai-trust-layer-policies

70 **NVIDIA: OpenShell documentation**
https://docs.nvidia.com/openshell/latest/about/installation

71 **WitnessAI: Agent controls and traffic boundary**
https://witness.ai/for-agents/

72 **Zenity: Agent security platform**
https://zenity.io/platform

73 **HiddenLayer: AI security platform**
https://www.hiddenlayer.com/

74 **F5: CalypsoAI acquisition completed and guardrails**
https://www.f5.com/company/blog/what-are-ai-guardrails

75 **Kong: MCP gateway policy boundaries**
https://developer.konghq.com/ai-gateway/entities/ai-mcp-server/

FRONTIER.

TAKE THE NEXT STEP

Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

Briefings, advisory work, executive discussions, and research inquiries.

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.