ai integrators
COUNCIL

FEBRUARY 2026

ai integrators
COUNCIL

WHITE
PAPER

Executive Summary

AI integrators are essential to enterprise AI adoption. Most technology providers of AI systems typically do not develop or commoditize foundation models, like Large Language Models (LLMs), nor do they control how AI systems are deployed across the public and private sectors. Instead, they configure, customize, and deliver AI systems that bring together the power of foundation models, insights from data provided by AI deployers, and product features that empower enterprises to harness the benefits of AI.

This growing group of companies, known as AI Integrators, serve as the bridge between developers of foundation models and deployers of AI systems. Their role has been misunderstood and often overlooked in emerging legislative and regulatory frameworks. Typically, these frameworks rely on binary distinctions between “developers” and “deployers,” but the AI ecosystem is more complex. AI Integrators do not fit neatly into either bucket, creating legal uncertainty and muddled compliance obligations. It is time to move beyond the developer-deployer binary.

The AI Integrators Council represents the voice of AI Integrators. Founded in 2025, we are a group of technology providers on the frontlines of enterprise AI adoption. The AI Integrators Council serves as the premier voice for companies operationalizing AI across sectors, bridging development and deployment through integration into systems, platforms, and applications. Collectively, we provide AI products deployed by millions of organizations globally, including the Fortune 500.

Our AI products are used by a range of civilian and defense agencies at all levels of government and a vast array of public and private companies. We are comprised of leading AI companies including **Alteryx, Atlassian, Box, Cognizant, Docusign, Peraton, SAIC, Salesforce, and Twilio**.

Our white paper offers three key insights for AI policymakers. Our breadth of engagement with AI use cases informs our depth on what’s needed in AI policy. Building on this experience, we have identified three insights that AI policymakers should keep top-of-mind as they develop AI policy initiatives to ensure that AI Integrators can continue to play a critical role in enterprise AI adoption.

1) Understanding the role and impact of AI Integrators in the AI value chain. Our paper provides a framework for the AI value chain and the role of AI Integrators in it, including examples from member companies of the AI Integrators Council.
2) Appreciating the limitations of the developer-deployer binary for AI policymaking. We explain the shortcomings of the developer-deployer distinction, especially as the AI value chain expands, and we offer legislative examples of misplaced reliance on this binary.
3) Practical approaches to addressing AI Integrators in AI policy. We offer draft definitions and similar tools to support AI policymakers in developing policy initiatives that strengthen and enable the full AI value chain, including AI Integrators.

AI Integration in Practice

AI integrators drive growth in AI adoption and deliver innovation in practical applications of artificial intelligence. Integrators represent a diverse set of companies offering a broad range of AI-power solutions, from improving workplace productivity and streamlining communications to better securing and storing data. Leveraging foundation models and their underlying infrastructure, AI integrators engineer new products and services to add value across industries.

Across the AI Integrators Council, member companies contribute to AI innovation through their commercial offerings, R&D investments, workforce skilling, and other initiatives that promote AI use. The following descriptions provide a thumbnail sketch of member companies' activities.

Alteryx integrates Generative AI and Machine Learning across its platform to enable users to harness AI throughout the data analytics lifecycle to enhance data preparation, analytics, and insights. For example, Alteryx Co-Pilot, an interactive AI assistant, assists users in building workflows with step-by-step guidance, analytics tool recommendations, and automated tool placement, while a suite of Generative AI tools helps users prepare data, integrate LLMs into their workflows, and easily transform analytic outputs into automated, tailored reports and documents.

Atlassian helps teams plan, build, and run the work that matters most. More than 350,000 customers worldwide, including over 80% of the Fortune 500 and 60% of the Forbes AI 50, rely on Atlassian as their AI-powered System of Work. With decades of rich context about how organizations actually operate—their goals, projects, incidents, decisions, and the knowledge spread across tools and teams—Atlassian builds AI into the fabric of how work already gets done. Rovo is Atlassian's AI built for teamwork, connecting people, work, and organizational knowledge across Atlassian and connected third-party tools. It transforms scattered work into shared context, trusted answers, and context-aware agents that help teams move faster, make smarter decisions, and keep complex initiatives moving forward. Rovo uses an intelligent mix of foundational models, including Atlassian-hosted and multiple third-party models, but its real strength lies in its deep integration with the data, permissions, and workflows customers already trust in Atlassian Cloud. These capabilities are delivered with enterprise-grade privacy, security, and data governance built in, so institutions can adopt AI where their most important work already lives. The response has been strong: Rovo has surpassed 5 million monthly active users, underscoring the demand for AI that solves human collaboration problems and is deeply and thoughtfully embedded in the tools teams rely on every day.

Box AI is an enterprise-grade AI layer built directly into Box that helps organizations securely understand, analyze, and act on their most critical content. Box AI harnesses the power of leading AI models—including those from OpenAI, Google, Anthropic, and others—and applies them to documents where they already live, enabling users to ask questions, generate summaries, extract key information, and automate insights across large volumes of unstructured data. All Box AI interactions are permissions-aware and governed by Box's industry-leading security, compliance, and privacy controls, ensuring it only accesses content a user is authorized to see and that customer data is never used to train AI models without explicit consent. The result is faster, more informed decision-making and AI-powered productivity, delivered in a way that meets the rigorous security and governance requirements of enterprises, regulated industries, and the public sector.

Cognizant acts as an integrator by connecting third-party AI models, proprietary tools, and client data to create tailored solutions for complex business challenges. For example, in healthcare, Cognizant brings together vendor AI platforms and client systems to automate claims processing, flag anomalies in patient records, and streamline communications—without fundamentally altering the underlying models. In financial services, Cognizant integrates fraud detection technologies with client workflows, ensuring regulatory compliance and enabling clients to set their own business rules. Across industries, Cognizant's integration work focuses on configuring, customizing, and documenting AI systems so clients can harness advanced analytics and automation while maintaining control over how these technologies are used.

Docusign Iris is an agreement-specific AI engine that is built on decades of agreement expertise enabling businesses to create agreements faster, negotiate more effectively, and sign contracts with greater confidence to deliver maximum business impact. Iris helps flag potential risks during contract review, extracts key insights so teams can quickly understand obligations, and easily connects into existing systems and workflows.

Peraton fuses foundational AI models with custom-built tools and specialized datasets to help America's government—including its armed forces, intelligence community, and other national security agencies—tackle their most critical challenges. Peraton tailors AI systems to analysts', operators', and decision-makers' unique contexts, adapts them to integrate with customers' existing technology, and fine-tunes them to make sense of datasets that would otherwise be disconnected and disparate. The result are platforms that help US investigators identify healthcare fraud costing taxpayers billions, enable US operators to understand the information environment in hostile territory, and position US cyberdefenders to counter incoming attacks—whether those platforms are deployed at the enterprise-level or at the mission's edge.

SAIC As a leading federal technology integrator, SAIC plays a critical role in deploying AI solutions across government missions, ensuring that AI technologies are effectively and securely integrated into operational environments. SAIC is solving our nation's most complex modernization and readiness challenges across the defense, space, federal and state & local civilian, and intelligence markets. Our robust portfolio of offerings includes high-end solutions in systems engineering and integration; enterprise IT and cloud services; cybersecurity; software; and advanced analytics, AI, ML, and data sciences. Our expertise in AI enables us to deliver innovative solutions that support full scale mission decision making, governmental efficiency, and mission success.

Salesforce is the leading customer relationship management technology, helping our customers build and improve their customer relationships. Salesforce helps businesses large and small to stay ahead of changing customer expectations with data tools, trusted AI, and best-in-class apps for sales, service, marketing, commerce, and IT — all on one integrated platform. With Salesforce, employees can be more productive with a single view of customer data that keeps every team in sync and by integrating apps, data, metadata, and agents on a deeply unified platform, and by infusing agentic capabilities into each of the apps, Salesforce has created a platform for the Agentic Enterprise.

Twilio uses AI to help businesses connect with their customers in smart, simple, and personal ways. AI is built into every part of our platform so companies can deliver the right message to the right person at the right time. We make personalized communication faster and easier while ensuring trust, transparency, and control.

Why It Matters: AI Integrators in the AI Value Chain

Snapshot View of AI Value Chain

Enablement	Intelligence	Integration	Deployment
Infrastructure	Model developers	AI integrators	AI deployers
Hardware Materials required to support AI models and systems, such as GPUs, memory, interconnects, solid state drives Cloud Public and private cloud services that enable AI models and systems	Generative models Models like LLMs that produce text, image, video, and other content Non-generative models Models that enhance ranking, prediction, relevancy, and other computational tasks	Copilots, assistants, agents, and other AI-powered tools and services that integrate foundation models	Public and private sector organizations, groups, and individuals putting AI systems into use

This framing is adapted from *Artificial intelligence: sizing and seizing the investment opportunity* (UBS, July 2024).

The AI value chain comprises several different and dynamic layers that provide critical inputs for AI products and services. At the infrastructure layer, core elements like cloud services and graphics processing unit (GPU) chips come together in data centers to enable AI. Moving up to the intelligence layer, model developers build generative and non-generative models that can be delivered directly to enterprises or consumer end-users. In the enterprise context, AI integrators use these models to power AI systems like assistants, agents, and other tools. Finally, deployers put AI models and systems into use, and they are increasingly empowered with technical tools to control their AI deployments.

These layers aren't always clearly distinguished, which can be confusing. For example, individual companies may operate across multiple layers because of the breadth of their offerings. One way this can happen is when models are developed solely for first-party use in a company's own products. But leading companies that develop their own models also rely on third-party models in their AI products. This makes them both developers and integrators.

The AI value chain has evolved since the original conceptions of the developer-deployer model. There are several actors in the value chain that are not distinguished in early efforts to create AI policy frameworks. In part, that failure is a result of the predictable challenge of legislating around the technology that is so early in its evolution. Policymakers are operating from a small sample of potential technologies or applications, and consulting with a limited number of large technology companies. The result is an inaccurate view of the AI playing field that will ultimately lead to less innovation for both business and consumers.

When regulations fail to recognize integrators as a distinct category, practical risks and vulnerabilities emerge for both service providers and clients. The following scenarios show how these gaps could impact real-world projects and business outcomes.

1) Misplaced Compliance Burden: A consulting firm integrates a client-licensed AI chatbot into a customer service workflow. Because regulations only recognize developers and deployers, the firm must prepare extensive provider-level documentation and testing, even though it did not design or train the model.
2) Impossible Obligations: An integrator fine-tunes a third-party resume-screening model using client HR data, then hands the model back. Regulations require the integrator to monitor bias and data provenance after deployment, but the integrator no longer has access to the client's data or operational controls.
3) Over-Compliance and Project Delays: A systems integrator is asked to host a fraud detection model for a bank. Due to unclear regulatory expectations, the integrator over-documents integration steps and tests as if it were the model provider, resulting in higher costs and delayed project launch.
4) Accountability Gaps: An integrator configures an AI system for a healthcare provider, documenting every step. The client, however, does not monitor real-world use or provide necessary disclosures. If an error occurs, responsibility is unclear and accountability is misaligned.
5) Chilling Effect on Advisory and Managed Services: An integrator advises a financial institution on selecting an AI vendor but does not touch or configure the system. Broad regulatory definitions could still capture this advisory work, making the integrator hesitant to offer expert guidance.
6) Reduced Innovation and Market Access: Integrators avoid offering tailored fine-tuning or orchestration services for high-risk AI systems due to regulatory ambiguity. Clients lose access to advanced integration options, and the market for integrator-driven solutions shrinks.

Principles for Drafting AI Legislation, Regulation, and Guidance

As governments around the world contemplate artificial intelligence regulations, it is important to establish core principles to help guide that process. Among those principles, are the concepts around who are the parties the government seeks to include under its regimes, whether that be legislative rules or other types of AI guidance.

With that in mind, we would propose the following principles as part of that effort to help governments focus first on why it's regulating and who it's regulating,

1) Recognize the breadth of the AI value chain and define roles narrowly.
2) Distinguish between AI product categories and their impact on potential risks stemming from AI.
3) Target regulatory measures on the highest risk scenarios and assign accountability specifically.

As integrators, we believe in promoting the responsible and trustworthy use of AI, and that assigning accountability continues to ensure trustworthiness and transparency throughout the AI value chain.

Proposed Definitions

Transparency obligations are key for all actors in the value chain. Current AI policies often misallocate responsibility because they oversimplify the value chain into just “developers” and “deployers,” overlooking the many actors that adapt and integrate models into systems and applications. We argue that instead of adding a single new role definition, policymakers should conduct a full reappraisal of the AI value chain to create clearer, more precise roles that enable effective risk allocation and regulation.

AI Integrator

Companies considered AI integrators are integral to the AI value chain. We recommend that AI legislation or regulation include an explicit definition of “AI Integrators” to ensure clarity. By overlooking AI integration as a distinct and essential function, policymakers risk assigning responsibilities that do not align with how AI is brought to market or creating gaps that allow real risks to go unaddressed. AI integrators incorporate AI into their products and services to enhance product offerings, and customers can customize the system for their intended use. The role of an entity is typically product or service specific depending on the offering. It’s possible that an entity could have a different role depending on the product or service being offered.

Additionally, if legislation provides opportunities to participate in a sandbox environment, industry development programs, grant or tax incentive programs, or similar; like other categories, we believe AI integrators should also have the opportunity to participate.

An AI integrator should be defined as an entity that utilizes foundation models created by third parties to create AI systems and product features without substantially modifying the underlying AI model(s). The role of an AI integrator is distinct from the roles of both developer [of foundational models] and deployer [of AI systems].

Developer

A developer is any entity that designs, builds, trains, or otherwise creates an artificial intelligence model, and subsequently makes it available for distribution, commercialization, or further integration.

A developer is typically responsible for the technical architecture, training data, algorithmic design choices, and core functionality of the AI model, and is the actor that first places the technology “into the stream of commerce.”

Unless an entity intentionally makes substantial modifications to the underlying model that materially alters its performance, intended purpose, or risk profile in a way not contemplated by the developer, an entity should not be deemed a developer solely because they integrate, configure, or deploy a pre-existing AI system into an application or product.

Deployer

A deployer is any entity, acting in a commercial, organizational, or governmental capacity, that takes responsibility for putting an AI system, including a high-risk AI System into actual use.

Deployment involves not only the technical act of running or activating the system, but also the decisional authority to determine how, when, where, and for what purposes the AI system is applied. This role is distinct from that of a developer or integrator: a developer builds or trains the model; an integrator may utilize a foundational model in an AI system; but the deployer bears ultimate accountability for real-world operations, including the conditions of use, safeguards, oversight mechanisms, and the impact on end-users and affected populations.

Substantial Modification

Substantial modification means a change to an AI model or system, after it has been placed on the market or put into service that results in a new, reasonably foreseeable, and material risk including actions that:

1) Change the model's underlying mathematical framework in ways that could foreseeably impact the model's accuracy, robustness, or bias.
2) Remove specific safeguards or risk management measures, such as human oversight or fail-safe mechanisms.
3) Introduce self-learning capabilities to a previously static system.

Substantial modification is limited to changes impacting the underlying AI model and does not include changes to software architecture, controls, or interfaces in a software platform or system in which an AI model is integrated, or modifications to infrastructure, such as hardware, supporting the AI model's performance.

Downstream integration of AI models into AI systems or software platforms, and modification of AI systems, should not qualify as substantial modification because the underlying AI model remains intact. Existing best practices generally address modifications to software systems and platforms in which AI models are integrated, as well as hardware and other infrastructure supporting AI model performance; AI policies should narrowly focus discussion of substantial modification on those actions that impact underlying AI models as the activities potentially introducing novel risk.

As innovative developers seek to optimize the performance and computing resource demands of AI models and systems, a number of different optimization techniques have arisen. The chart below helps distinguish which techniques generally could be considered as substantial modifications based on their impact to underlying AI models.

Responsibilities Across the AI Value Chain

The AI ecosystem consists of a diverse array of actors playing different roles – including developers, integrators, deployers, and infrastructure providers – to enable end users to harness AI's power. These actors cannot effectively confront risk in isolation; each actor within the ecosystem must collaborate with others to produce and deploy effective, responsible AI systems. Accordingly, responsibility for managing risk across the AI ecosystem must be a shared undertaking.

When considering responsibility across the AI value chain, policymakers should operate according to the following principles:

- **Risk Identification.** Whether developer, integrator, or deployer, an organization beginning a new AI project should identify where new risks might be introduced and what mitigations may be appropriate.
- **Risk Ownership**. Actors introducing new risks must be accountable for addressing them, guided by organizational risk assessments. Risks created at the intersection of technologies and services should be appropriately shared.
- **Transparency-Enabled Responsibility**. Effective risk management depends on sharing technical documentation and incident reports across the value chain, and each actor throughout the value chain should be held responsible for sharing.
- **Systemic Risk and High-Risk Use Focus**. Effective AI policies will incentivize actors across the AI ecosystem to proactively work to reduce risk in all forms; however, given limited resources, regulatory policies should be targeted to address systemic risks and high-risk use cases.

Based on these principles, the below provides a set of key responsibilities for different roles across four categories: **data, security/privacy, risk/safety, and transparency and technical documentation**.

Developers

- Data: provenance, accuracy, objectivity, and governance of training data sets
- Security/Privacy: issue patches as necessary
- Risk/Safety: risk assessments, testing, and appropriate evaluations of models
- Transparency and Technical Documentation: transparency about model design, limitations, intended use cases, and potential risks; information sharing about known threats or vulnerabilities

Integrators

- Data: customer data handling, data lineage, data governance
- Security/Privacy: customer data privacy and security
- Risk/Safety: risk assessments and appropriate testing of integrated systems
- Transparency and Technical Documentation: on integrations, intended uses, and potential risks; information sharing about known threats or vulnerabilities

Deployers

- Data: access controls, data governance
- Security/Privacy: customer data privacy and security
- Risk/Safety: risk assessments and appropriate testing of integrated systems; responsibility for high-risk use cases; monitoring and auditing; human oversight
- Transparency and Technical Documentation: on use contexts, intended uses, and potential risks; notification to users of AI interactions; information sharing about known threats or vulnerabilities

Infrastructure Providers

- Data: encryption, access controls
- Security/Privacy: secure storage; protect compute resources
- Risk/Safety: ensure redundancy and resilience of underlying infrastructure
- Transparency and Technical Documentation: support for logging and traceability

Cross-Jurisdictional Compliance

As AI integrators increasingly operate across multiple regions, they face a complex and often conflicting landscape of AI regulations, such as the EU AI Act, US federal and state laws, and emerging frameworks in Asia-Pacific and other jurisdictions. The absence of harmonized requirements create significant challenges, including duplicative compliance efforts, increased operational costs, and heightened legal risk.

To address these challenges, we recommend that policy:

- Explicitly recognizes the cross-jurisdictional realities faced by integrators in legislative and regulatory guidance.
- Prioritizes harmonization of technical requirements and mutual recognition of compliance certifications to reduce friction for multinational integrators.
- Develops practical tools—such as model contractual clauses, compliance mapping frameworks, and centralized guidance—to help integrators navigate and reconcile overlapping requirements.
- Considers establishing a regulatory “one-stop shop” or similar mechanism to streamline compliance for organizations operating in multiple jurisdictions.

By addressing cross-jurisdictional compliance, policymakers can enable integrators to innovate and deliver AI solutions globally without undue legal uncertainty or operational burden.

Insights and Conclusions

Effective AI governance requires regulators and lawmakers to address the unique role of integrators in the AI ecosystem. The following insights set out specific actions for both current regulatory frameworks and future rulemaking.

- Regulators should publish interpretive guidance that recognizes integrators and assigns obligations based on actual control within the current developer/provider and deployer frameworks. This guidance should clarify an objective standard for when material modification converts a company into a developer or provider for that scope, and how deployers retain end-use responsibilities.
- Lawmakers should add “AI integrator” as a distinct role with proportionate duties that match actual control. Responsibilities for integrators should be defined separately from those of developers and deployers. New rules should include safe harbors tied to documented practices and apply proportional enforcement so liability follows control.

Proportional Enforcement. Enforcement should follow the principle of control. The entity with the ability to prevent or mitigate harm should take responsibility, based on documented role allocations and safe harbor criteria.

Ongoing Consultation. Policymakers should establish structured, recurring dialogue with system integrators, distinct from developers and deployers. This approach keeps rulemaking, guidance, and technical standards practical and enforceable as the market evolves.

Integrators responsibly bridge advanced AI technologies and real-world use. Treating them as full developers or deployers by default distorts accountability and raises costs without improving safety. Clearer rules—explicit recognition of integrators, control-based allocation of duties, contractual safe harbors, and proportional enforcement—will preserve incentives to innovate while ensuring responsibility rests with the parties best positioned to manage risk. This approach keeps innovation moving while making sure accountability lands where it belongs.