WHITEPAPER

From Shadow AI to Strategic AI

A Practical Governance Blueprint for Enterprise AI at Scale

How forward-looking enterprises move from ungoverned, ad-hoc tooling toward a secure, accountable, and value-driven AI operating model.

alteus
Enterprise AI Platform

2026

Shadow AI was never the problem. It was the experiment that proved the workforce was ready. The task now isn't to shut it down, it's to give it structure.

This white paper provides a practical framework for converting fragmented, ungoverned AI usage into controlled, measurable, and scalable AI capability. It defines a four-stage maturity model, a five-dimension operating framework, and a sequenced 30/60/90-day action plan.

The objective is to give employees governed environments where AI adoption is safe, visible, and reusable across the organization.

FROM SHADOW AI TO STRATEGIC AI

THE SHIFT

Every organization's AI journey begins the same way: people experiment. What is often labeled "Shadow AI" was, in reality, a necessary and healthy first stage: employees tried tools, discovered what worked, and built genuine fluency with what these systems can do. That grassroots experimentation surfaced real use cases faster than any top-down mandate could, and it revealed which capabilities teams actually value.

The organizations that formalize early will not simply manage risk; they will compound the fluency their people have already built into a lasting advantage, while others are still treating experimentation as something to police.



"AI delivers enterprise value when it moves beyond isolated use cases and is scaled through coordinated adoption, clear governance, and cross-functional alignment."

Emil Calofir, Head of AI at eSolutions

THE SOLUTION

Maturing from experimentation to strategic capability means three things: classifying your data (knowing what's sensitive and defining how each category is handled), governing agents centrally (so the best ones are shared and reused across teams instead of rebuilt in silos), and making delivery measurable (turning ad-hoc wins into a durable, auditable advantage). This paper lays out a four-stage maturity model, a five-dimension operating framework, and a 30/60/90-day action plan to get there.

3x

Early adoption research found employees using generative AI at roughly three times the rate leaders estimated, an early signal of how quickly fluency spread

78%

of IT leaders encountered unexpected SaaS charges from AI pricing models (Zylo)

Aug 2026

EU AI Act transparency obligations under Article 50 have applied since 2 August 2026. Is your AI operating model ready?

Table of contents

Executive Summary 5

Why Shadow AI Persists 7

The Hidden Cost of Fragmented AI 9

Why Governance Must Become Infrastructure 11

The Enterprise AI Operating Model 13

Maturity Model: From Shadow AI to Strategic AI 15

Practical Blueprint for Enterprise AI Governance 17

Example Use Cases by Department 19

Where Alteus Fits 21

30/60/90-Day Action Plan 23

References 24

Final Checklist: Enterprise AI Governance Readiness 25

Executive Summary

The workforce is already using AI. The question is whether the organization controls the risk or the risk controls the organization.

Shadow AI showed that the workforce is ready for enterprise AI. Employees started adopting these tools because they delivered faster reports, better code, and quicker customer responses. But the same behavior that signaled readiness also created exposure, and the risk is what gets noticed first: sensitive data leaves the organization's control, compliance obligations go unmet, and AI spending fragments across departments with no central visibility. Proof and exposure are two sides of the same phenomenon, and suppressing one means losing the other.

The scale of that early adoption is well documented. McKinsey's Superagency in the Workplace research (2025) found that employees were using generative AI at roughly three times the rate their leaders estimated, and that a large majority of employees expected AI to reshape a significant share of their work within two years. While the specific figures will continue to evolve as enterprise AI matures, the underlying pattern has proven durable: adoption moves faster than leadership visibility, and faster than formal governance.

The organizations that formalize early will not simply manage risk; they will compound the fluency their people have already built into a lasting advantage, while others are still treating experimentation as something to police.

Three pressures make acting on this now unavoidable.

The first is regulatory. The EU AI Act's prohibitions on harmful AI practices and

AI literacy obligations have been in force since February 2025. Transparency obligations under Article 50 have applied since 2 August 2026, with a limited transition until 2 December 2026 for specific marking and detection requirements for certain generative AI systems placed on the market before 2 August 2026. High-risk obligations for stand-alone AI systems apply from 2 December 2027. Each milestone raises the cost of not knowing where and how AI is used. An enterprise that cannot produce an inventory of its AItools, users, and data flows cannot assess its compliance posture with respect to any of these obligations.

The second pressure is data exposure. Sensitive corporate and customer information flows into unmanaged AI tools with no audit trail, no access controls, and no guarantee of data residency.

The third is wasted investment. Duplicated AI subscriptions, uncontrolled consumption-based pricing, and local experiments that never scale into reusable enterprise capabilities all drain budget without delivering system-level returns.

This whitepaper provides a practical framework for converting fragmented, ungoverned AI usage into controlled, measurable, and scalable AI capability. It defines a four-stage maturity model, a five-dimension operating framework, and a sequenced 30/60/90-day action plan. The objective is to give employees governed environments where AI adoption is safe, visible, and reusable across the organization.

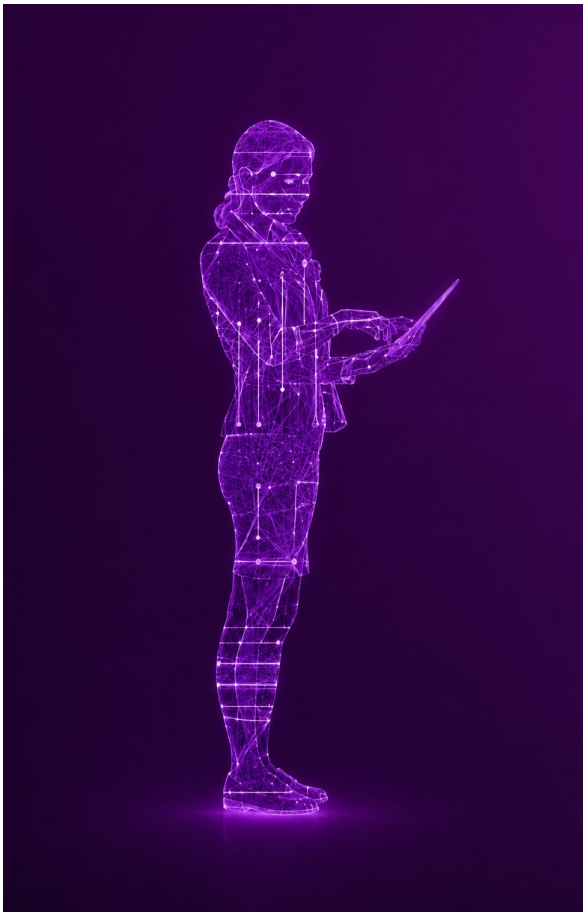
Shadow AI was the experiment. Strategic AI is the operating model. The task is not to reverse adoption, it's to orchestrate it.

Why Shadow AI Persists

AI spreads faster than governance, demanding visibility, guardrails, and scalable controls.

Shadow AI is not a temporary anomaly or merely a compliance issue. It is a predictable response to pressure for speed, efficiency, and autonomy. Employees adopt AI tools for the same reason they once adopted spreadsheets, collaboration platforms, and cloud storage before IT formally governed them: the tools solve immediate problems faster than the organization can approve sanctioned alternatives. Shadow AI is therefore not a passing wave, but a recurring feature of enterprise technology adoption.

Consider the analyst who must deliver a quarterly commentary by day's end. A generative AI tool can produce a first draft in minutes. The decision to use it is typically not driven by recklessness, but by deadline pressure and the availability of a tool that appears to work. Scaled across functions that produce reports, presentations, research, and written analysis, the broader pattern becomes clear: AI adoption often advances wherever operational tempo exceeds the speed of formal technology governance.



The same dynamic applies in technical teams. Engineers adopt AI coding assistants, model APIs, and workflow automation tools to reduce time spent on repetitive coding, testing, and documentation tasks.

Support teams integrate AI chatbots, summarization tools, and multilingual assistance into customer workflows to reduce response times. Business functions do the same in their own domains, whether in content generation, contract review, recruitment workflows, or internal knowledge retrieval. In each case, the incentive is immediate utility; governance usually arrives later.

Knowledge retrieval accelerates the spread further. Employees use AI tools to search internal documentation, summarize meeting transcripts, and extract insights from large document sets. The productivity gain is immediate. The data exposure is invisible. Meanwhile, business units run their own AI pilots in isolation. Marketing tests content generation. Legal experiments with contract review. HR trials CV screening. Each department builds its own tool stack, data flows, and risk profile, with no mechanism to share what works.

The most difficult driver to contain is the one IT cannot see coming. SaaS platforms add AI features through routine updates. CRM systems now include predictive analytics.

Collaboration tools include generative writing assistants. Business intelligence platforms include natural language querying.

IT teams did not approve these features. They arrived with tools that were already approved, and the organization is using AI without knowing the extent of its exposure.

Research published in 2025 helps illustrate the structural nature of this problem. McKinsey's Superagency study found a significant perception gap between leadership estimates of employee generative AI use and employees' self-reported adoption. While the exact figures will evolve as enterprise AI matures, the more important takeaway remains: leaders often underestimate how quickly AI use spreads inside the organization, especially through informal workflows, team-level experimentation, or newly embedded SaaS features.

That insight points to the central governance implication for the years ahead. The key risk is not simply that employees use AI without permission; it is that organizations lack the visibility, controls, and operating model needed to distinguish productive experimentation from unmanaged exposure. Enterprises that respond only with prohibition are likely to drive AI usage further underground. Those that respond with governance, clear guardrails, approved pathways, and monitoring can convert the same employee demand into a safer and more scalable form of adoption.

The challenge is not AI adoption, but managing it before risks compound.

The Hidden Cost of Fragmented AI

The security, compliance, financial, and operational costs that accumulate when AI usage has no operating layer.

Shadow AI imposes costs that do not appear on any single budget line. They accumulate across security, compliance, procurement, and operations, compounding as AI adoption accelerates.

Consider the contract that leaves the building. A legal team member pastes a 40-page customer agreement into a public AI chatbot to extract key clauses. The data is now on infrastructure the organization does not control, in a jurisdiction it may not know, subject to terms of service nobody reviewed. It may be used for model training. It may be accessible to the provider's staff. The organization has no record of the transaction, no data processing agreement in place, and no way to audit what happened afterward. A 2025 survey found that 93 percent of IT leaders have concerns about the data security risks associated with AI tools, a concern grounded in exactly this scenario, repeated across thousands of employees and dozens of unmanaged tools.

The inability to answer basic questions compounds the exposure. Most enterprises cannot list which AI tools their employees use, which models are involved, which data flows where, or which business functions depend on which tools. Without an inventory, risk assessment cannot begin. Compliance documentation cannot be produced. Budget management cannot be accurate. Different departments operate under different assumptions about what data can be shared with AI tools, which models are approved, and what review process applies to AI-generated outputs. These assumptions are rarely documented, rarely consistent, and rarely aligned with regulatory obligations.

For regulated industries, the audit gap is the cost that escalates fastest. Financial services firms, healthcare providers, and legal practices operate under strict documentation requirements. When AI tools enter regulated workflows without logging, version control, or human review trails, the organization cannot demonstrate compliance to a regulator or an auditor. Under the EU AI Act, transparency obligations that already apply under

Article 50 make the inability to evidence AI use, outputs, and data involvement an operational and compliance risk, not a future problem.

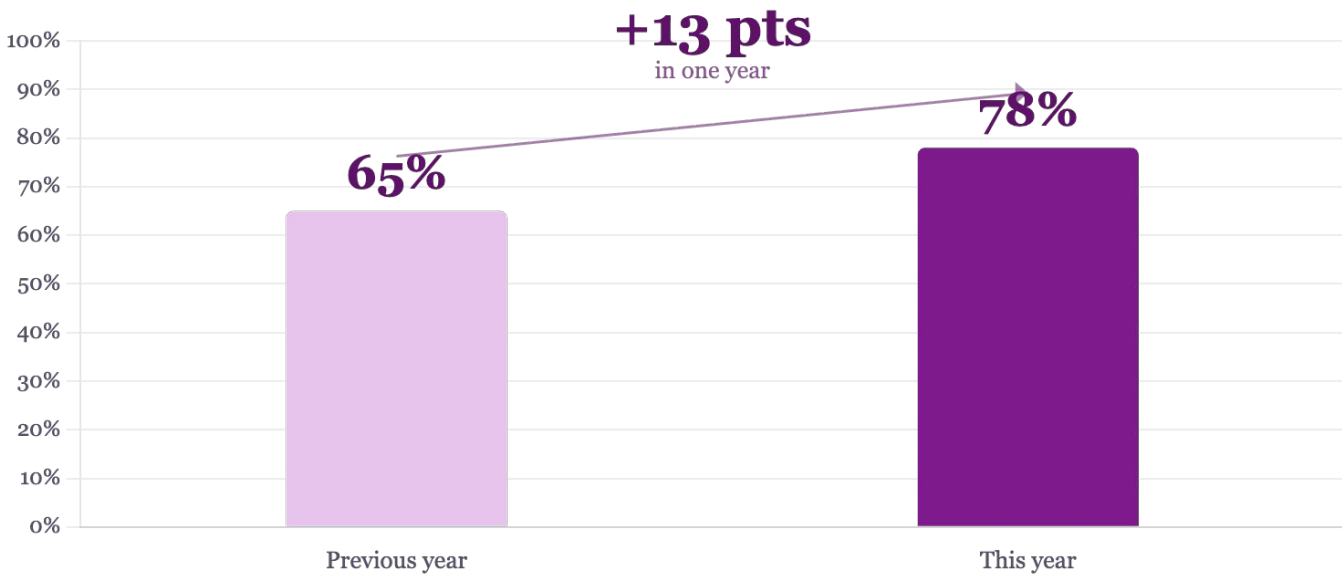
The financial cost is more visible but equally structural. Different departments purchase overlapping AI subscriptions, one team buys a generative writing tool, another buys a different one, and neither knows about the other's purchase. Zylo reports that 78 percent of IT leaders encountered unexpected SaaS charges due to consumption-based or AI pricing models, up from 65 percent the previous year. Employees and lines of business now control 87 percent of applications and 85 percent of SaaS spending. AI tools enter through the same decentralized purchasing channels that have been spreading across enterprises for a decade.

The cost that frustrates leadership most is the failure to scale. A department builds an effective AI workflow for document summarization, it saves hours per week, it produces consistent output, and the team relies on it. But the workflow is tied to that team's specific tool, data access, and prompts. It cannot be shared with compliance, procurement, or HR. The organization gains no institutional capability from the experiment. The investment does not compound.

These costs are structural. They arise from the absence of an operating layer between the workforce and the AI tools the workforce uses. Policy documents alone do not close the gap.

Most enterprises cannot list which AI tools their employees use.

Share of IT leaders hit with unexpected SaaS charges from consumption-based or AI pricing models (Zylo).



Why Governance Must Become Infrastructure

Policy documents alone cannot close the gap. Governance works when it is embedded in the tools employees use every day.

AI governance in most enterprises is a document. It specifies which tools are approved, which data categories are restricted, and which review processes apply. It sits in a policy repository. Employees are expected to read it, understand it, and comply with it. This approach fails for the same reason that a speed limit sign without enforcement fails. The policy exists. The behavior continues.

Governance becomes effective when it serves as infrastructure: an operating layer that sits between the workforce and AI models, enforcing policy through technical controls rather than employee discretion. That layer begins with a current inventory: a single source of truth showing which AI tools are in use, by whom, in which departments, and for which purposes, updated continuously rather than through periodic surveys. It extends to role-based access controls that tie permissions to identity, role, department, and project: a legal team member sees different approved models and data sources than a marketing team member, and external collaborators operate within defined boundaries. Data isolation between departments, projects, and use cases ensures that customer data used in a support workflow does not become accessible to the engineering team's AI assistant.

The layer also controls how AI tools connect to organizational knowledge. Approved data connectors link to internal document repositories, knowledge bases, and data sources via retrieval-augmented generation (RAG), grounding AI responses in authorized organizational data rather than public model training data. A complete logging and monitoring record captures who used which model, with which data, for which purpose, at which time, the foundation for audit readiness, cost attribution, and compliance reporting. Real-time usage and cost visibility let

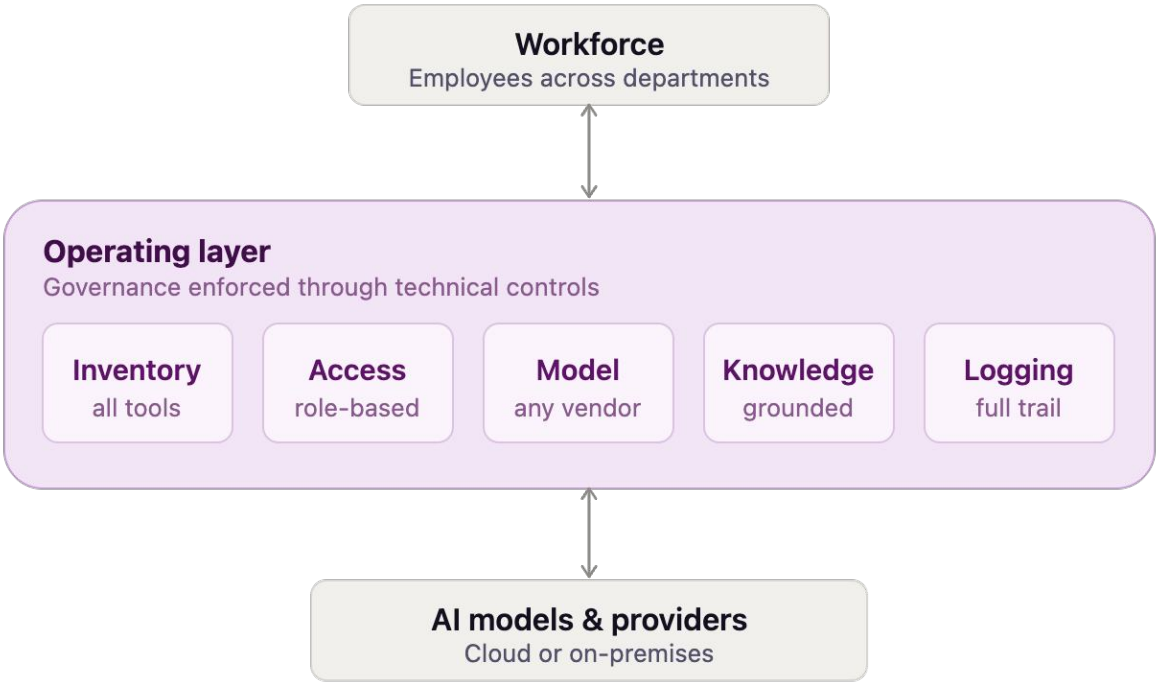
the organization track AI consumption across departments, attribute costs to business units, and identify underutilized or duplicated subscriptions.

Model and provider flexibility is the dimension that prevents the governance layer from becoming a new form of lock-in. The organization can switch between AI providers, use multiple models simultaneously, and deploy on cloud or on-premises infrastructure without dependency on a single vendor's pricing, availability, or data handling practices. Department-specific contexts and agents complete the picture: configured AI environments where a finance department agent operates with different data sources, model preferences, and output formats than a customer support agent.

The EU AI Act makes this shift from policy to practice a compliance requirement. Transparency obligations require organizations to know when AI is used, how outputs are generated, and what data is involved. High-risk system obligations demand documented risk management, logging, and human oversight. These requirements cannot be met through policy documents. They require technical controls.

This approach fails for the same reason that a speed limit sign without enforcement fails. The policy exists. The behavior continues.

All AI access flows through one layer



The Enterprise AI Operating Model

Five dimensions: visibility, control, enablement, measurement, reuse, that define what governed AI looks like in practice

Strategic AI is AI that operates within a defined governance framework, produces measurable outcomes, and generates reusable institutional capability. Five dimensions define the operating model that produces this result.

Visibility is the foundation. The organization knows where AI is used. It maintains a current inventory of AI tools, models, use cases, and data flows. When a regulator asks which AI systems process customer data, the organization answers with data, not estimates. When an auditor requests documentation of AI usage in a specific workflow, the organization produces it. When the CFO asks what the company spent on AI tools last quarter, the answer is attributed to business units, not buried in unlabeled SaaS charges.

Control builds on visibility. The organization manages who can use which AI tools, with which data, and under which conditions. Permissions are enforced at the technical layer. Data boundaries are maintained between departments, projects, and use cases. Policy violations generate alerts. The gap between what the policy says and what employees actually do closes because the policy is embedded in the tools themselves.

Enablement is the dimension that makes governance sustainable. The organization provides employees with governed AI environments that are faster and more capable than the unmanaged alternatives. Teams build AI assistants, configure agents, and launch workflows within approved boundaries. The path of least resistance leads through governance. When the governed environment is the best available option, Shadow AI shrinks because it no longer offers an advantage.

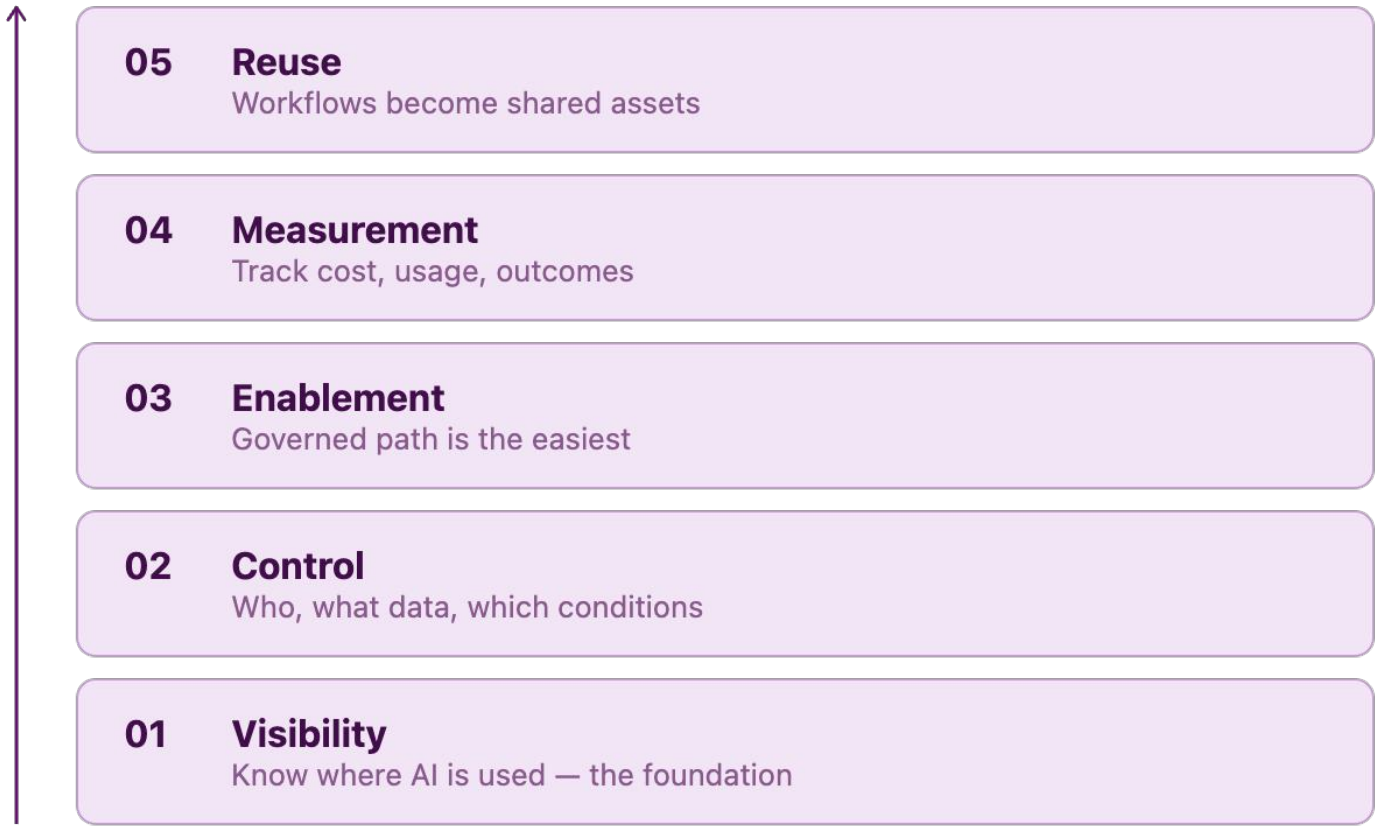
Measurement closes the loop. The organization tracks AI usage, cost, quality, and business outcomes. It knows which use cases deliver returns and which consume resources without impact. It attributes AI spending to business units and compares costs across providers. The data that was invisible in Stage 1 becomes the basis for investment decisions in Stage 4.

Reuse is where the investment compounds. Successful AI workflows, agents, and configurations become shared organizational assets. A document summarization workflow built by the legal team becomes available to compliance, procurement, and HR. The organization stops solving the same problem in six different departments with six different tools.

This operating model does not emerge from a single procurement decision or a single policy update. It requires an operating layer that enforces governance, enables adoption, and surfaces data across all five dimensions simultaneously.

The organization stops solving the same problem in six different departments with six different tools.

The five dimensions of governed AI



Maturity Model: From Shadow AI to Strategic AI

Four stages every organization moves through, and the two gaps that separate most enterprises from governed AI.

Organizations move through four stages on the path from fragmented AI usage to governed AI capability. Each stage has a distinct operational profile, risk exposure, and set of missing capabilities.

Stage 1: Uncontrolled Experimentation

Employees use AI tools independently. IT has no inventory of which tools are in use. Security has no visibility into data flows. Procurement has no control over spending. Compliance has no documentation of AI usage in regulated workflows. The organization's AI risk exposure is unknown, making it unmanaged.

Stage 2: Departmental Adoption

Individual departments build AI workflows and standardize on preferred tools. Governance is inconsistent. The legal team operates under one set of rules, marketing under another, and engineering under a third. Successful use cases remain trapped inside the departments that built them. The organization has islands of AI capability with no bridge between them.

Stage 3: Centralized Oversight

IT, security, and business leadership define organization-wide AI governance standards. Approved tools, models, and data handling policies are documented. A central inventory of AI use cases exists. Compliance reporting is possible. Departments operate within defined boundaries. What is missing: an operating layer that enforces governance through technical controls rather than policy documents, and the ability to turn departmental successes into institutional assets.

Stage 4: Enterprise AI Operating Layer

AI is governed, measured, reusable, and integrated into business processes. An operating layer sits between the workforce and AI models, enforcing policy, enabling adoption, and surfacing data across all five dimensions of the

operating model. Departments launch AI initiatives within governed workspaces. Successful workflows become shared assets. The organization can demonstrate compliance with regulators, attribute costs to business units, and scale AI capabilities without increasing risk.

Most enterprises operate at Stage 1 or Stage 2. The gap between Stage 2 and Stage 3 is organizational. It requires leadership alignment and cross-functional governance. The gap between Stage 3 and Stage 4 is architectural, and it requires infrastructure.

Four stages, two gaps



The gap between Stage 2 and Stage 3 is organizational. The gap between Stage 3 and Stage 4 is architectural.

Practical Blueprint for Enterprise AI Governance

Eight sequenced steps to move from discovery to governed operations: each step builds on the one before it.

Moving from Shadow AI to Strategic AI requires sequenced action across eight steps. Each step builds on the previous one. The sequence matters.

Step 1: Map existing AI tools, agents, and use cases. Conduct a discovery process across departments. Identify which AI tools are in use, who is using them, for which purposes, and with which data. Use SaaS discovery tools, expense reports, security audits, and direct interviews. The output is a current-state inventory, the foundation for every subsequent step.

Step 2: Identify sensitive data flows. For each identified AI use case, determine which data enters the AI system. Flag use cases involving customer data, employee data, financial data, protected health information, intellectual property, or data subject to regulatory controls. This step surfaces the organization's actual AI risk exposure. It often reveals exposures that leadership did not know existed.

Step 3: Define approved models and providers. Based on the inventory and risk assessment, establish which AI models and providers are approved for each use case and data category. Consider model performance, data handling practices, deployment options, and jurisdictional requirements. The EU AI Act's risk-based classification framework provides a structure for this assessment.

Step 4: Set permission levels by role, department, and project. Define access policies that map identity and context to approved tools and data sources. A legal team member reviewing contracts operates under different permissions than a marketing team member generating campaign copy. External contractors operate within narrower boundaries than employees.

Step 5: Create internal knowledge contexts using RAG. Connect approved AI tools to internal document repositories, knowledge bases, and data sources. This grounds AI responses in organizational data,

reduces the risk of hallucinations, and keeps sensitive data within controlled environments rather than flowing to external model providers.

Step 6: Launch controlled pilots for high-value departments. Select two or three departments with clear AI use cases and high Shadow AI exposure. Deploy governed AI workspaces for these departments with pre-configured agents, approved data sources, and defined output formats. The governed environment must be faster and more capable than the unmanaged alternatives employees currently use. Otherwise, Shadow AI continues.

Step 7: Monitor usage, cost, and performance. Track AI consumption across departments, models, and use cases. Measure adoption rates, cost per use case, output quality, and user satisfaction. Surface anomalies: unexpected cost spikes, policy violations, or usage patterns that indicate Shadow AI migration to new unmanaged tools.

Step 8: Scale reusable agents and workflows. Identify the workflows and agent configurations that deliver the highest impact in pilot departments. Package them as shared organizational assets and make them available to other departments through the governed AI environment. The investment compounds.

This blueprint is an operating rhythm. The inventory, risk assessment, and policy framework require continuous updating as new AI tools enter the organization, new regulatory obligations take effect, and new use cases emerge.

Package successful workflows as shared organizational assets and make them available across the governed environment. The investment compounds. This blueprint is an operating rhythm, not a one-time project.

Example Use Cases by Department

The following examples illustrate how adopting governed AI changes the operational profile of specific business functions

IT and Security

A CISO receives an email from the board: "What AI tools are our employees using, and where is the data going?" In a Stage 1 organization, the answer is an estimate based on a survey conducted six months ago. In a governed environment, the answer comes from a centralized administration console that shows current AI tools, users, and data flows across the organization. Role-based access controls enforce department-level and project-level permissions. Usage logs and audit trails support compliance reporting. Policy violations generate alerts. Where EU AI Act transparency obligations apply, this shift from estimate to evidence supports compliance readiness, because organizations must be able to show how AI is used, not merely describe it in policy.

Legal and Compliance

A junior associate at a law firm needs to review a 200-page merger agreement by Friday. She opens a public AI chatbot, pastes sections of the document, and asks for clause summaries. The document contains confidential client information and privileged communications. In a governed environment, the associate uses an approved AI workspace with internal document access via RAG and workspace-level data isolation. The AI-assisted review operates within a controlled environment where data does not leave the organization's infrastructure. The firm can demonstrate to regulators and clients that AI use in legal workflows is subject to defined controls.

Finance

The Financial Planning & Analysis (FP&A) team uses three AI tools for report generation, data analysis, and forecasting across four cost centers. Nobody has a complete view of what the organization spends on AI. In a governed environment, finance teams work within approved workspaces that provide internal data connectors and

cost tracking by department and use case. AI spending becomes visible, attributed to business units, and comparable across providers. Duplicated subscriptions surface and get consolidated.

Operations

A logistics team builds an AI workflow to extract shipment data from emails and automatically update tracking systems. It works. The procurement team, facing the same data extraction problem with supplier invoices, builds its own workflow from scratch using a different tool. In a governed environment, the logistics team's workflow becomes a shared asset in an agent library. Procurement configures it for invoices instead of shipments. The organization solves the problem once.

Customer Support

A support team lead integrates an AI chatbot into the company's help desk to handle tier-one queries in five languages. Customer names, account numbers, and purchase histories flow through the tool. The team lead did not check where the data is processed or whether the provider has a data processing agreement with the company. In a governed environment, support teams use approved AI chatbots with multilingual capabilities and access to the internal knowledge base via RAG. AI-generated responses are logged and auditable.

Engineering and Data

Three engineering teams use three different AI coding assistants. One team's preferred tool stores prompts and generates code on the provider's servers. Another team's tool is blocked by corporate firewall rules, so they use a workaround. Model selection is driven by individual preference. In a governed environment, engineering teams access multiple approved model providers through a single platform, with on-premises deployment options and internal documentation access via RAG. Teams switch between models without changing tools. AI-assisted development operates within defined code review and security scanning workflows.

Customer data remains within controlled environments. AI-generated responses are logged and auditable.

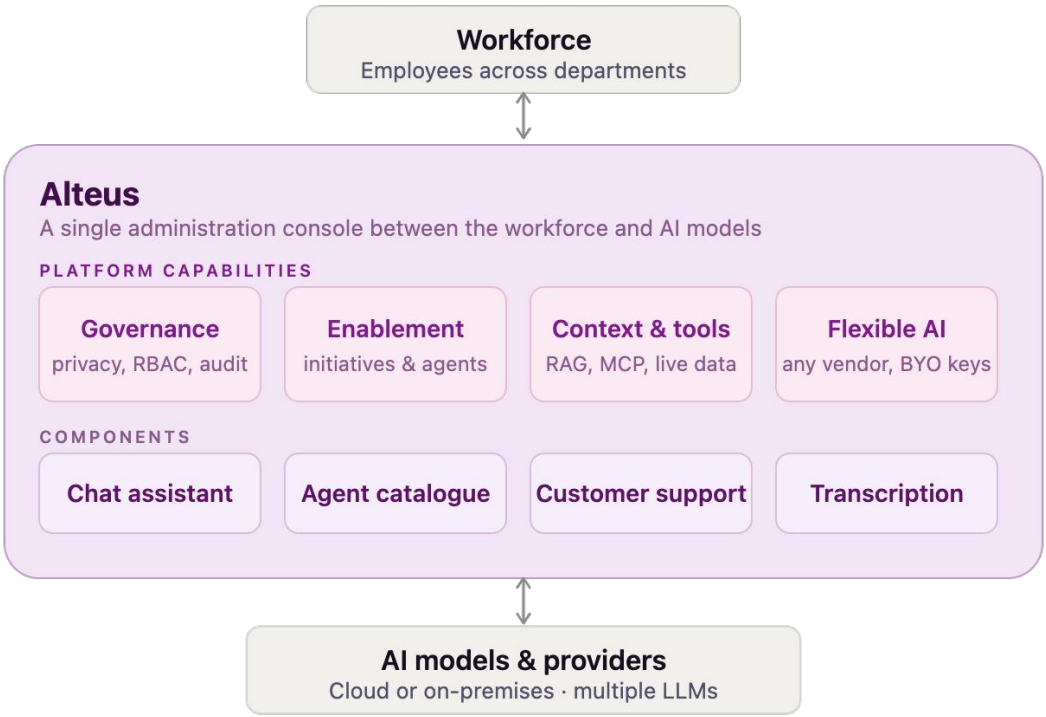
Where Alteus Fits

This governance model requires more than policy: it needs an operating layer that makes AI usage visible, applies controls consistently, and gives employees a governed environment they'll actually use. That's where Alteus fits.

Alteus is an enterprise AI platform designed to help organizations transition from fragmented AI adoption to a governed, measurable, and scalable enterprise AI implementation. It brings AI tools, users, data, and governance into a single operating environment, creating the technical layer needed to support the five dimensions outlined earlier in this paper: visibility, control, enablement, measurement, and reuse.

For visibility, Alteus provides a centralized administration console that gives organizations a current view of AI usage across departments, users, and use cases. This helps leadership answer basic but critical questions: which tools are being used, by whom, for what purpose, and with what data.

For control, Alteus applies role-, department-, and project-based permissions at the platform level. Organizations can define who can access which models, data sources, and agents, while maintaining privacy boundaries across teams and workflows. This reduces reliance on policy documents alone and helps embed governance directly into daily usage.



For enablement, Alteus gives teams governed workspaces where they can use AI productively within approved boundaries. Organizations can connect internal knowledge sources through retrieval-augmented generation, configure department-specific agents, and support business workflows without forcing users into unmanaged public tools.

For measurement, Alteus provides dashboards for usage, cost, and operational activity across models, teams, and use cases. This gives finance, IT, and business leaders a clearer view of adoption patterns, spending, and opportunities to improve efficiency or reduce duplication.

For reuse, Alteus makes it easier to turn successful departmental configurations into shared enterprise assets. Agents, workflows, and knowledge-enabled use cases can be extended across teams rather than rebuilt from scratch in each function.

Just as importantly, Alteus supports model and deployment flexibility. Organizations can work with multiple model providers, use their own API keys, and choose cloud, on-premises, or hybrid deployment approaches based on security, regulatory, and operational requirements. This helps reduce vendor dependency while preserving freedom to adapt as the AI landscape evolves.

In practice, this means organizations can use Alteus to support a range of governed AI capabilities, including internal AI assistants, department-specific agents, customer support experiences, transcription workflows, and AI-enabled process automation.

Alteus is built by eSolutions, a technology company with 25 years of experience delivering enterprise software, data platforms, and secure infrastructure solutions. In addition to the platform itself, eSolutions supports clients with AI strategy, readiness assessments, governance design, implementation, integration, and training.

For organizations currently operating in Stage 1 or Stage 2, Alteus helps create the operating layer needed to move toward Stage 4: an enterprise AI environment where adoption is visible, governed, reusable, and aligned with business value

Organizations can work with multiple model providers, use their own API keys, and choose cloud, on-premises, or hybrid deployment approaches based on security, regulatory, and operational requirements.

30/60/90-Day Action Plan

The following plan provides a sequenced path from Shadow AI discovery to governed AI operations. It is designed to produce visible results within each phase while building toward the full operating model.

FIRST 30 DAYS

Audit & Assess

Find and map usage

NEXT 60 DAYS

Define & Design

Standards and pilots

First 30 Days: Audit and Assess

- Conduct a cross-departmental Shadow AI discovery. Use SaaS management tools, expense report analysis, security audits, and department interviews to build a current-state inventory of AI tools, users, and use cases.
- Map sensitive data flows. For each identified AI use case, determine what data enters the AI system. Flag use cases involving regulated data categories.
- Identify stakeholders. Establish a cross-functional AI governance working group with representation from IT, security, legal, compliance, finance, and key business units.
- Assess current maturity stage using the four-stage model in section 6. Document the specific gaps between the current state and the target operating model.

Next 60 Days: Define and Design

- Define governance standards. Establish approved AI tools, model and providers by use case and data category. Document data handling requirements, access policies, and review processes.
- Set permission levels. Define role-based access policies that map identity and context to approved to and data sources.
- Select pilot departments. Choose t or three departments with clear AI cases, high Shadow AI exposure, a leadership support for governed AI adoption.
- Design pilot workspaces. Configure governed AI environments for pilot departments, including approved models, data connectors, RAG-base knowledge access, and pre-configu agents.

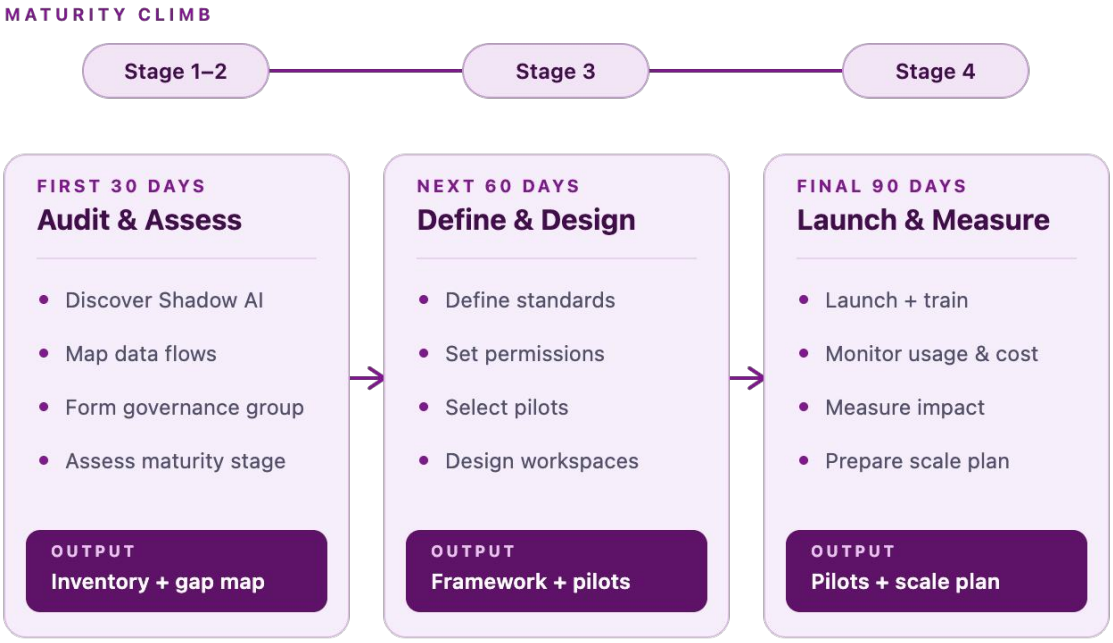
FINAL 90 DAYS

Deploy & Scale

Launch and scale

Next 90 Days: Launch and Measure

- Launch governed AI workspaces for pilot departments. Provide training and onboarding support. Make the governed environment the path of least resistance for AI adoption.
- Monitor adoption, usage, cost, and performance. Track which use cases deliver measurable returns. Identify where employees continue to use unmanaged tools and address the gap.
- Measure cost and productivity impact. Compare AI spending before and after governance. Document productivity gains from governed AI workflows.
- Prepare the scale plan. Based on pilot results, define the roadmap for extending governed AI workspaces to additional departments and use cases.



References

1. Wood, Nicole. "Shadow AI Explained: Causes, Consequences, and Best Practices for Control." Zylo, March 2025 (updated February 2026).
2. Council of the European Union. "Artificial Intelligence: Council and Parliament Agree to Simplify and Streamline Rules." Press release, 7 May 2026.
3. McKinsey & Company. "Superagency in the Workplace: Empowering People to Unlock AI's Full Potential." January 2025.
4. European Commission. "AI Act." Shaping Europe's Digital Future.
5. European Commission. "AI Act Service Desk - Frequently Asked Questions."
6. European Parliament and Council. "Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence." Official Journal of the European Union, 2024.
7. European Commission. "Commission Proposes Tech Sovereignty Package to Strengthen Europe's Digital Autonomy and Resilience." Press release.
8. European Commission. "Commission Seeks Feedback on the Draft Guidelines for the Classification of High-Risk Artificial Intelligence Systems." June 2026.
9. European Commission. "AI Act Enforcement Gets Independent Expert Support." June 2026.

Final Checklist: Enterprise AI Governance Readiness

Use this checklist to assess your organization's current AI governance posture. Answer each question with "yes," "no," or "partially." A "no" or "partially" answer identifies a gap that the blueprint in this whitepaper addresses.

- [ ] Do we know which AI tools employees use across all departments?
- [ ] Do we know what data enters AI systems, and can we prove it to a regulator?
- [ ] Can we control user access to AI tools by role, department, and project?
- [ ] Can we monitor AI usage and cost in real time, attributed to business units?
- [ ] Can we support multiple AI model providers without creating vendor dependency?
- [ ] Can successful AI workflows built by one department be reused by others?
- [ ] Can we demonstrate compliance with the EU AI Act's transparency and high-risk system obligations?

Next Step: Take the Enterprise AI Governance Readiness Assessment

A structured assessment maps your organization's current maturity stage, identifies specific gaps, and defines the sequence of actions needed to reach Stage 4: an enterprise AI operating layer where AI is governed, measured, reusable, and integrated into business processes.

Start an AI Governance Assessment

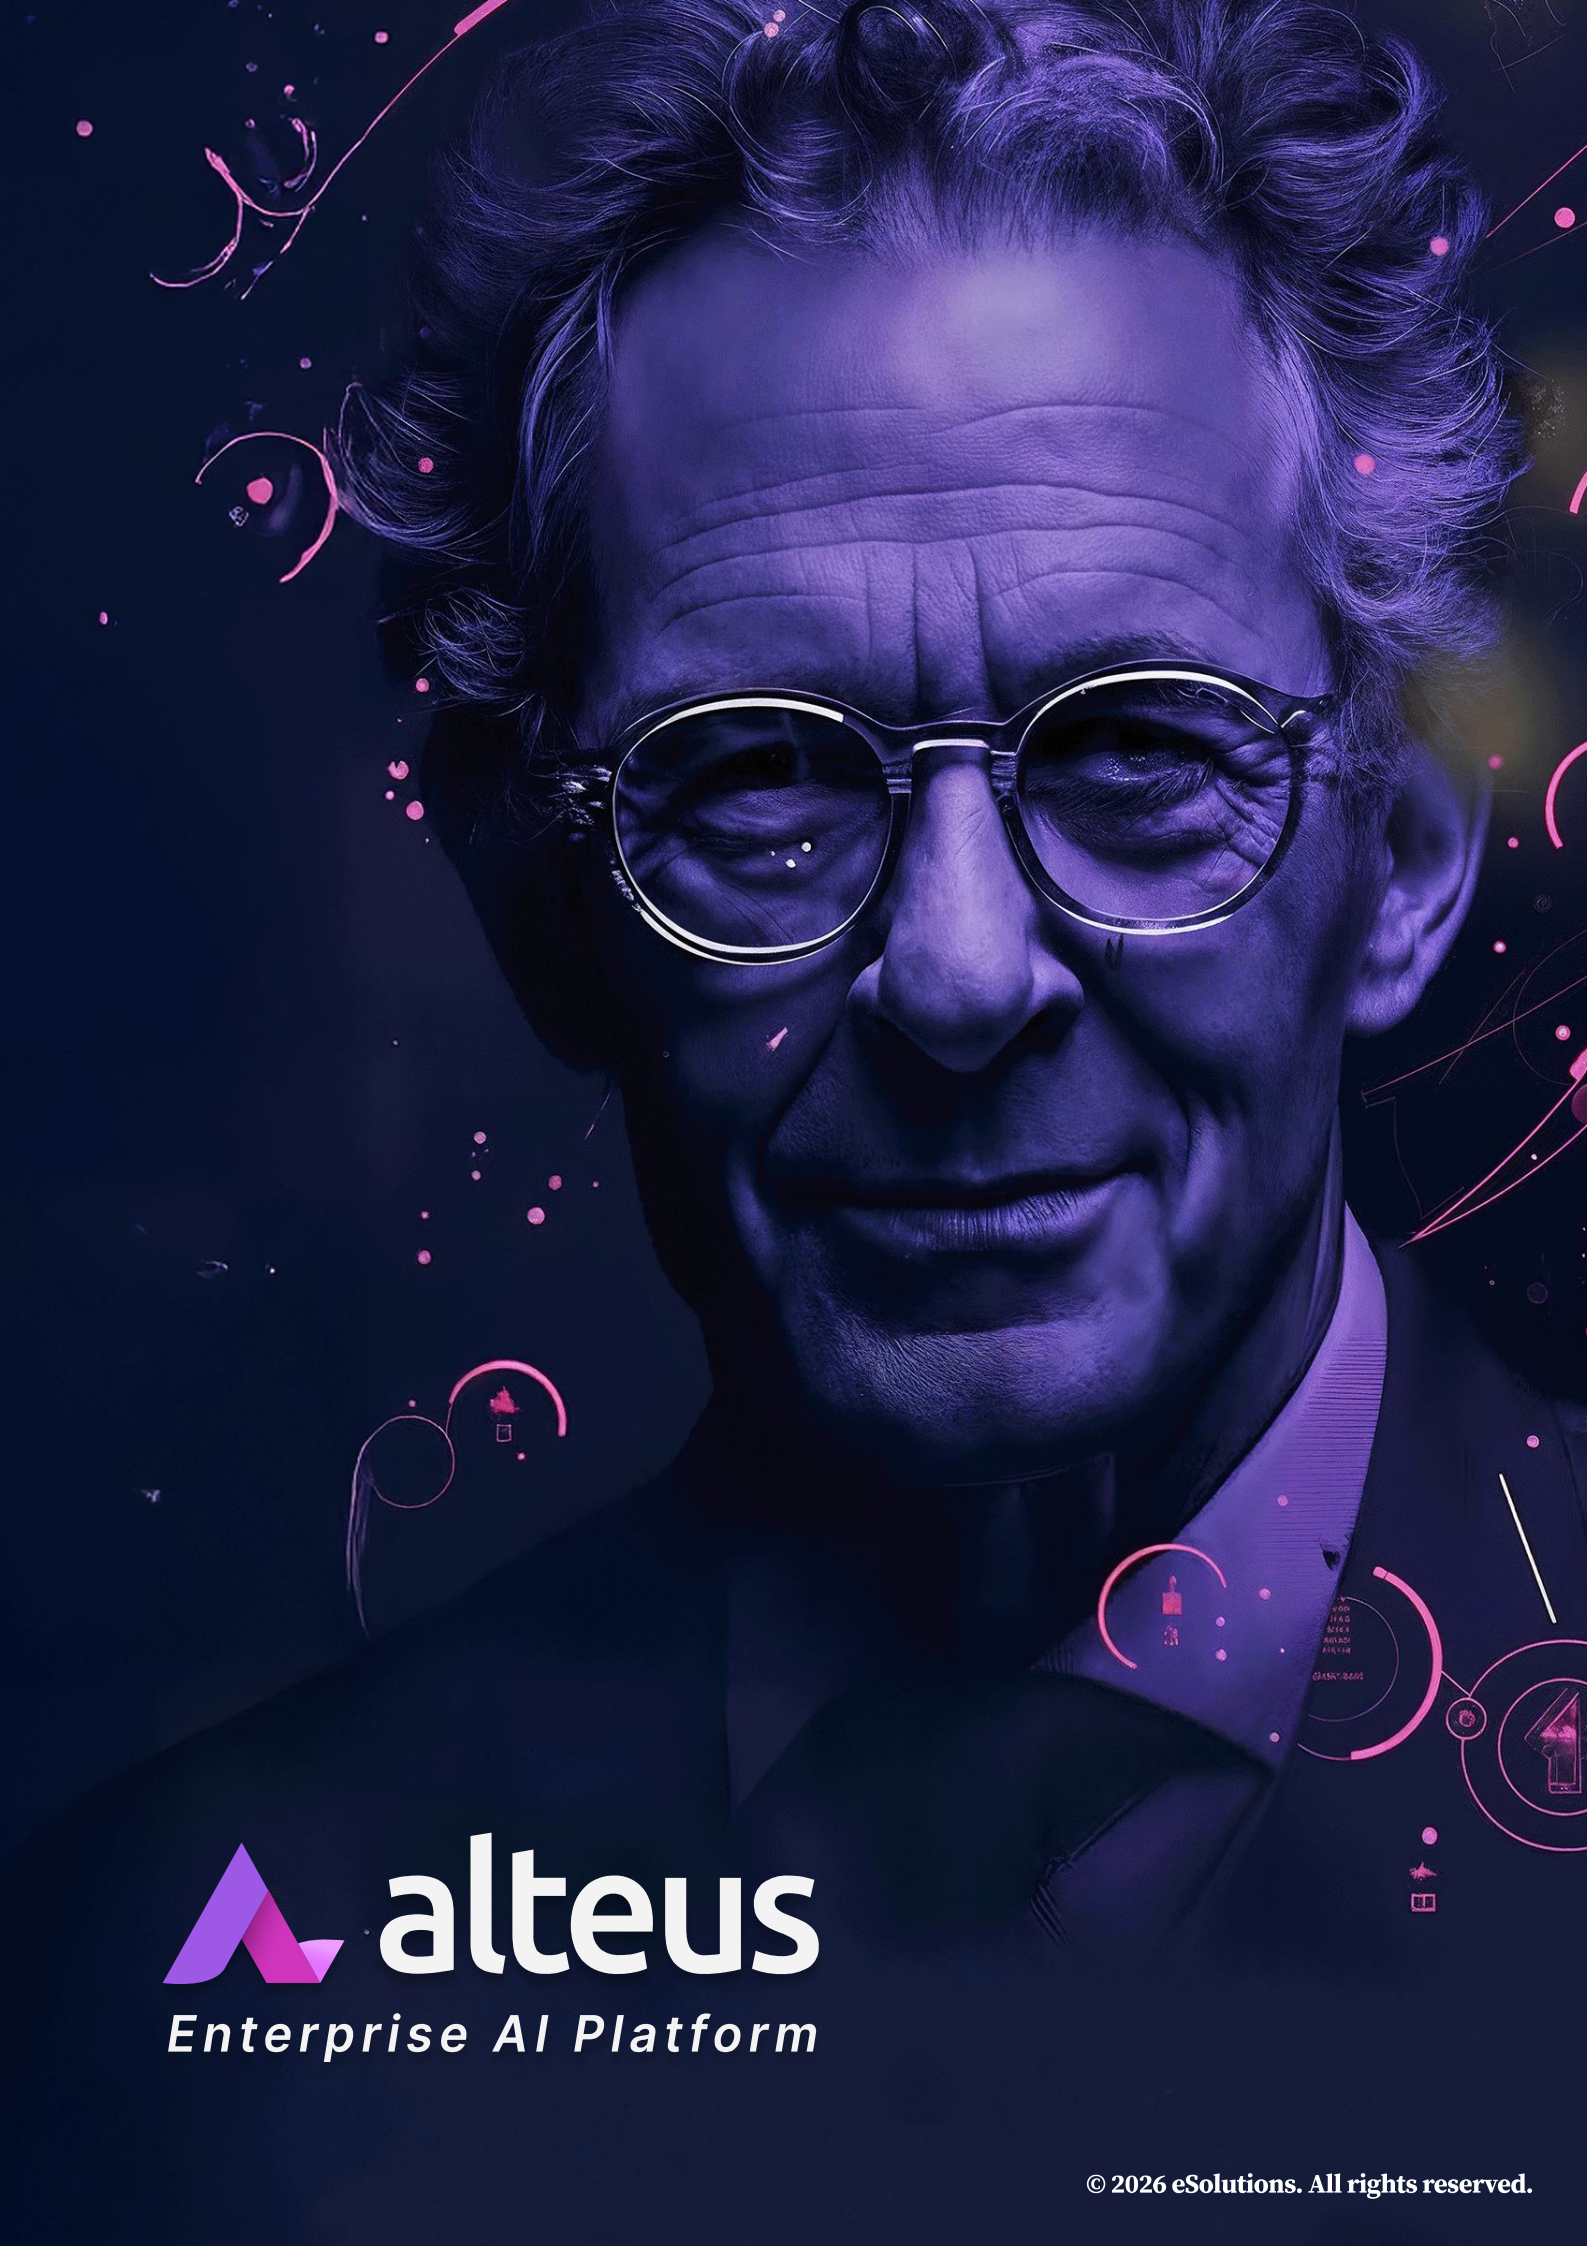
Book an Alteus Demo

About eSolutions

eSolutions is a technology company developing integrated, complex, and secure big data, cloud, IoT, and software engineering solutions. 25 years in operation. 150+ engineering team. Enterprise clients include Carrefour and Regina Maria. Offices in Bucharest, Bonn, and Berlin. Member of the Qvest group.

About Alteus

Alteus is an enterprise AI platform that unifies AI tools, users, data, and governance in a single administration console. Built by eSolutions. Cloud and on-premises deployment. Multi-model. No vendor lock-in.



© 2026 eSolutions. All rights reserved.