AHEAD

From Exploration to *Secure AI Execution*

How to pragmatically embed security into your AI adoption roadmap

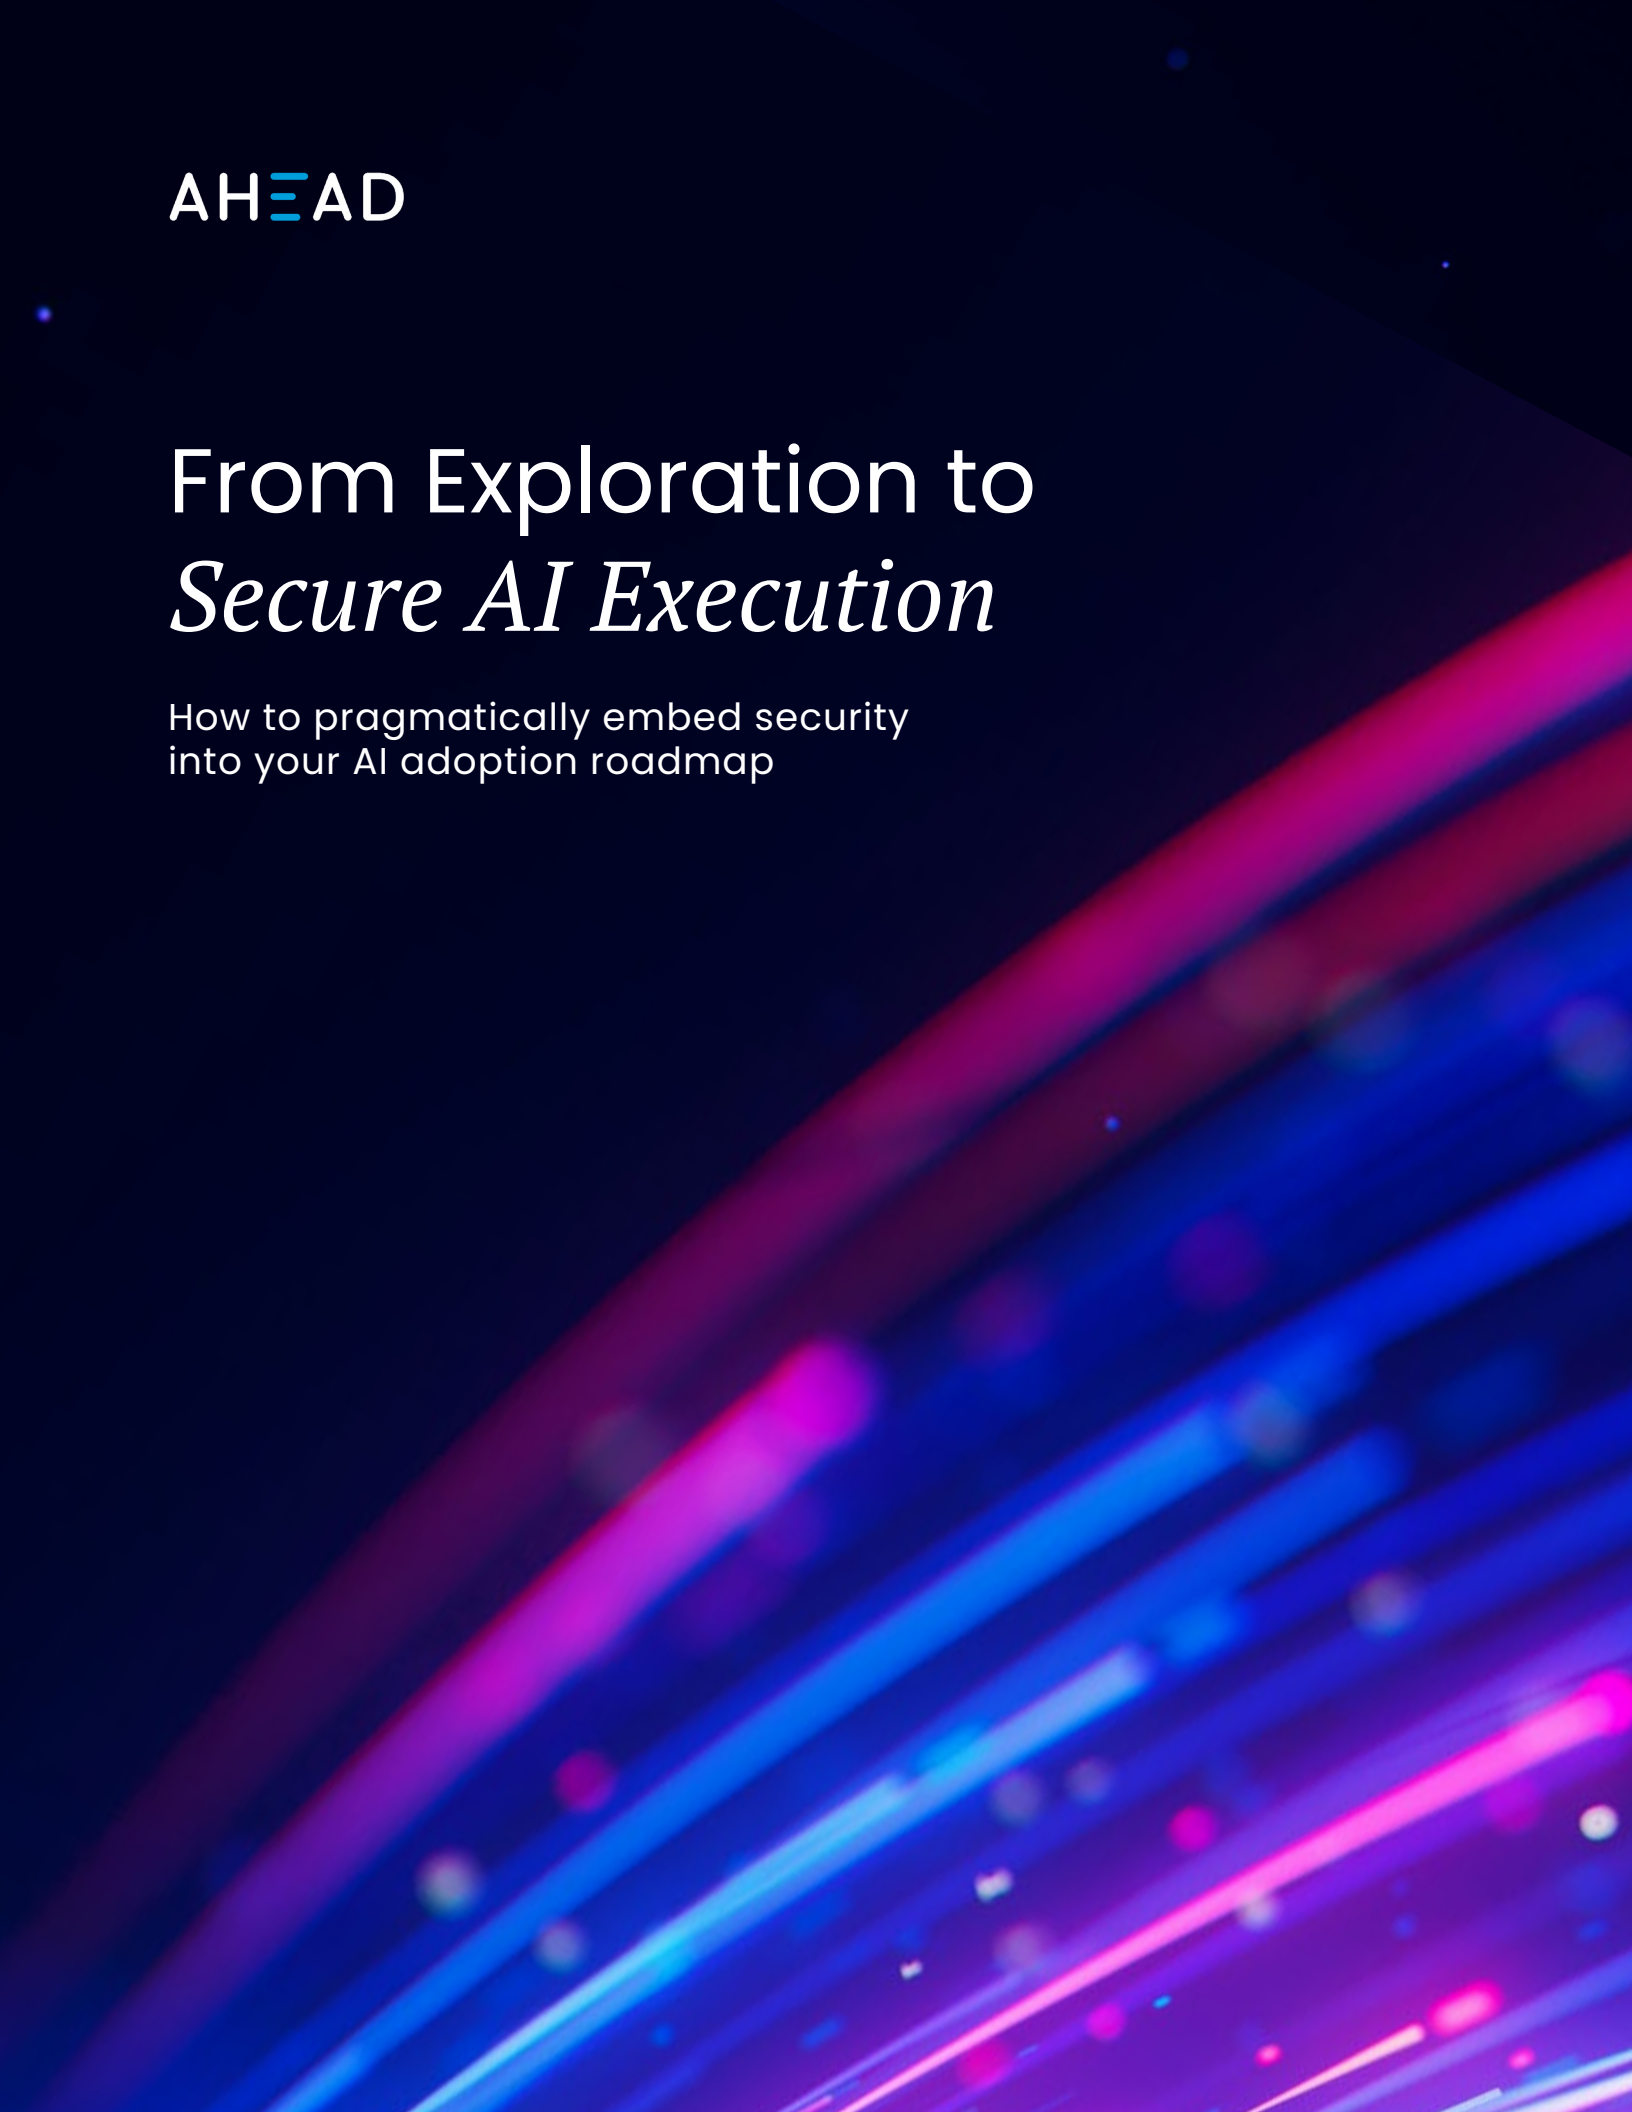

Executive Summary

Enterprise IT budgets across industries are pivoting to AI, outpacing traditional technology growth. For many organizations, experimental pilots are now shifting to large-scale AI investments spanning infrastructure, platforms, applications, and services.

And although the most competitive enterprises will be those that effectively harness AI, security remains a major roadblock to wide-scale adoption and ROI. Without proper safeguards, AI can introduce risks related to data leakage, model biases, and malicious usage. Therefore, it's critical to embed strong security protocols into AI initiatives and models from the beginning.

But the reality is that the vendors and architectural choices that technology leaders make this year will impact their organization for years to come. Enterprises need a pragmatic AI architecture now to avoid re-platforming later — which requires effective governance, secure runtimes, secure development processes, and more.

Read on to learn more about the barriers to secure AI adoption many enterprises face, how to approach security pragmatically, and why organizations should consider working with AHEAD to safely roll out new AI initiatives.

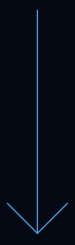

Barriers to AI Adoption

AI Literacy and Organizational Maturity

A fundamental challenge to enterprise AI adoption is a lack of relevant skills and knowledge among the entire employee base. If employees don't understand how AI works, they will not be able to use it effectively or securely. This knowledge gap increases the risk of employees bypassing governance, using unsanctioned tools, and inadvertently exposing sensitive data.

Organizations need to build AI literacy programs tailored by role to make adoption safe and scalable. From executive awareness down to hands-on training for developers, upskilling employees and providing clear governance policies help accelerate adoption and build trust in AI systems.

Data and Architecture Readiness

AI is only as good as the data it consumes. If data is siloed, messy, or inaccessible, then the output won't be reliable. Poor data foundations also mean higher costs, wasted effort, and models that can't be trusted in production. Inadequate architecture can introduce data privacy risks as well.

Enterprises need to deploy secure pipelines and architectures that make high-quality data available to sanctioned AI runtimes (the environments where models are deployed and inferencing is executed). This data should be properly classified as public, internal, restricted, or confidential so that it aligns with privacy and security standards, and governance procedures should define which roles have access to which data classification within company AI models.

Security and Regulatory Fatigue

Regulators are moving faster than most organizations can keep up. At the same time, adversaries are already exploiting AI-specific vulnerabilities. Companies risk not only breaches and reputational damage, but also fines and sanctions if compliance frameworks aren't addressed.

Organizations should implement both enterprise-wide governance for data and AI usage, and operational and development standards for AI systems. For example, modernizing security operations centers (SOCs) with AI-specific playbooks to handle new threats or turning to Policy as Code wherever possible to automate repetitive security processes.

The Path Forward for AI Adoption

Every new wave of innovation from cloud and mobile to AI brings hype and urgency. However, successful adoption still depends on the same three pillars: people, process, and technology.

If organizations only focus on the technology, they risk leaving people behind and skipping the process discipline that makes innovation sustainable. That leads to technical debt, confusion, and risk.

By putting people first through literacy, training, and role clarity, technology leaders can create the foundation for adoption. Establishing governance policies, standards, and security practices enables processes to be repeatable and safe. Only then does technology become an accelerator instead of a liability.

So while AI may be the newest emerging technology, the path to adoption for enterprises is not new. It's about applying these timeless fundamentals with discipline.

A Pragmatic Approach to Secure AI Adoption

AI Governance

An effective governance strategy includes policies around the safe usage of AI models and data, as well as standards for deploying and operating AI systems and infrastructure securely. This AI governance strategy should be integrated into existing data privacy, cybersecurity, and risk management frameworks to be most effective.

AI governance policies should clearly define the organizational users and third parties that have access to AI systems and data. There are also many standards organizations should consider to secure AI runtimes, software development lifecycles, workflows, and operations, which we discuss more in the following sections.

AI Development Lifecycles

The software development lifecycle (SDLC) for AI systems spans data acquisition and model training to application deployment and monitoring. However, each system has different risk profiles. Adopting pre-trained models introduces supply chain risks, while training custom models from scratch can be complex. In addition, retrieval augmented generation (RAG) and agentic workflows require more specific security measures related to data pipelines.

AI SDLC controls should include many different practices to detect vulnerabilities as early as possible within models, code, dependencies, and configurations before they're deployed. In addition, threat modeling and implementing guardrails can prevent vulnerabilities while tracking data lineages can help identify the source of model quality issues.

AI Runtime Environments

As mentioned previously, AI runtimes are environments where models are deployed and inferencing is executed. For example, many employees use enterprise assistants like Microsoft Copilot, vendor-hosted tools like OpenAI's ChatGPT, or AI embedded into existing systems. Some organizations also deploy their own custom runtimes for proprietary models.

Security controls for AI runtimes should include usage policies, such as protections for prompts, rate limiting, and managing permissions for non-human identities or third parties. The models themselves also need to be monitored and governed to prevent drift, hallucinations, and other issues.

AI-Optimized Security Operations Center (SOC)

After an AI system is deployed into production, organizations need an AI-optimized SOC to detect anomalous activity and identify emerging threats. Implementing comprehensive telemetry can help security teams gain visibility into AI systems, which can then be combined with detection engineering best practices and a standardized AI incident taxonomy to minimize false positives.

In addition, SOC teams need to implement automated incident responses, ideally with human-in-the-loop orchestration to review critical decisions and maintain oversight. AI SOC playbooks and tabletop exercises can also help SOC teams refine and practice standardized response to threats.

AI Security Maturity

AI Governance

POLICIES | STANDARDS | COMPLIANCE

Foundational: Overarching AI GOVERNANCE policy; fragmented policies; DATA GOVERNANCE strategy, initial policy reviews.

Advanced: Framework-aligned DEFINED RUNTIME & SDLC STANDARDS; scoped AI use; lifecycle documented; legal partnership, accountable owners.

Optimal: Standards and policy-driven POLICY AS CODE; continuous improvement; predictive risk management.

AI Secure Architectures

VISIBILITY | GUARDRAILING | RED TEAMING

Foundational: AI USAGE VISIBILITY; AI use case inventorying. Architecture gap analysis.

Advanced: CONTEXTUAL GUARDRAILING AI FW DEPLOYMENT; identity security, API Security. FOUNDATIONAL SDLC controls.

Optimal: ORCHESTRATED RED TEAMING; framework aligned policy-as-code; continuous improvement; ADVANCED SDLC controls.

AI Aware Security Operations

AI TAXONOMY | AI DETECTION | AI RESPONSE

Foundational: AI TELEMETRY VISIBILITY & DETECTION; Establish standardized AI incident categories to classify threats and ensure consistent SOC response.

Advanced: AI PLAYBOOKS; Develop AI-specific playbooks with repeatable workflows, enabling rapid and coordinated SOC actions.

Optimal: HUMAN IN LOOP OCHESTRATED REMEDIATION; Operationalize runbooks automating playbooks, integrating AI detections into SOC orchestration seamlessly.

OWASP Top 10 Agentic AI Threats of 2025

1. AGENT GOAL HIJACK

Malicious content alters an agent's objectives or decision path, causing unintended actions.

2. TOOL MISUSE & EXPLOITATION

Agents misuse legitimate tools due to ambiguous prompts, over-privilege, or poisoned inputs.

3. IDENTITY & PRIVILEGE ABUSE

Agents reuse, escalate, or leak inherited credentials, delegated permissions, or agent-to-agent trust.

4. AGENTIC SUPPLY CHAIN VULNERABILITIES

Compromised tools, prompts, plugins, or external agents alter behavior or expose data.

5. UNEXPECTED CODE EXECUTION

Agents generate or execute unsafe code or commands without proper isolation or review.

6. MEMORY & CONTEXT POISONING

Poisoned memory, RAG stores, or embeddings corrupt future agent behavior across sessions.

7. INSECURE INTER-AGENT COMMUNICATION

Unauthenticated or unprotected messages between agents to allow spoofing or injection.

8. CASCADING FAILURES

Errors in one agent propagate through interconnected multi-agent workflows, amplifying impact.

9. HUMAN-AGENT TRUST EXPLOITATION

Humans over-rely on agent recommendations, leading to unsafe approvals or unchecked actions.

10. ROGUE AGENTS

Compromised or misaligned agents act harmfully while appearing legitimate within the system.

A foundational level of AI security includes basic governance policies for AI and data usage, although they may be fragmented and leave gaps in critical areas. These organizations also work to gain visibility into AI usage and potential threats using monitoring and telemetry tools.

As organizations advance the maturity of their AI security, they define more comprehensive and integrated governance controls for their SDLCs and runtimes. Advanced SOC teams also define AI-specific runbooks with repeatable workflows to respond to threats.

An enterprise that has reached the optimal maturity level of AI security uses Policy as Code wherever possible, from organization-wide AI and data governance policies to standards for development and operations. SOC teams also leverage human-in-the-loop response automation to rapidly contain threats while aligning critical decisions with business policies.

AHEAD AI Security Solutions

AHEAD is a leading enterprise solutions provider with deep experience in cybersecurity and AI. We have a diverse team of security professionals that can provide the multi-domain expertise and additional capacity you need to enhance your AI security capabilities.

Our comprehensive AI security services include:

- **AI Policy Gap Analysis:** Evaluate the existing AI governance framework at your organization and develop new policies and standards to further reduce risk.
- **AI Secure Architecture Assessment:** Identify runtime and SDLC controls gaps so that your organization can invest in new and existing vendors to cover them.
- **AI Penetration Test:** Enable AHEAD's Red Team to stress test your AI security controls and uncover any potential weaknesses that need to be addressed.
- **AI SecOps Assessment:** Review your current detections, playbooks, and runbooks for AI incident response preparedness to improve your SOC processes.

Ready to roll out AI safely and effectively?

Contact AHEAD to learn more about our AI security solutions.

AHEAD

Combining cloud-native capabilities in software and data engineering with an unparalleled track record of modernizing infrastructure, we're uniquely positioned to help accelerate the promise of digital transformation.

National Hubs

CHICAGO
444 W. Lake Street
Suite 3000
Chicago, IL 60606

NEW YORK
500 5th Avenue
Floor 17
New York, NY 10110

ATLANTA
1117 Perimeter Center
W406
Atlanta, GA 30338

SAN FRANCISCO
2000 Crow Canyon Place
Suite 250
San Ramon, CA 94583