# Secure By Design AI agent architecture for government markets

Government agencies at both federal and state levels face increasing pressure to deliver higher-quality services faster while reducing administrative overhead. Artificial intelligence (AI) automation, including agents and agentic workflows, is now essential to meeting these efficiency targets.

To implement AI agents and agentic solutions, government agencies rely on an expansive supply chain of commercial vendors. However, these contractors face frustrating and time-consuming barriers to meeting regulatory requirements:

1. No consensus on what a mature, well-defined AI agent or agentic setup should look like.
2. Lack of a clear risk-mitigation plan.
3. While FedRAMP, FISMA, DoD CC-SRG, and CMMC lay out controls, they don't explain how to put those controls in place for AI agent systems and agentic workflows.

These obstacles delay government market entry and stymie any commercial or enterprise AI agent product launch requiring secure-by-design architecture—where security controls are built into AI agent systems from the start, rather than added after implementation.

# Coalfire's agentic security solution is the answer

Coalfire, an Amazon Web Services (AWS) Partner with 20+ years of compliance engineering expertise, provides three connected solutions to help teams deploying AI agents overcome the obstacles defined earlier.

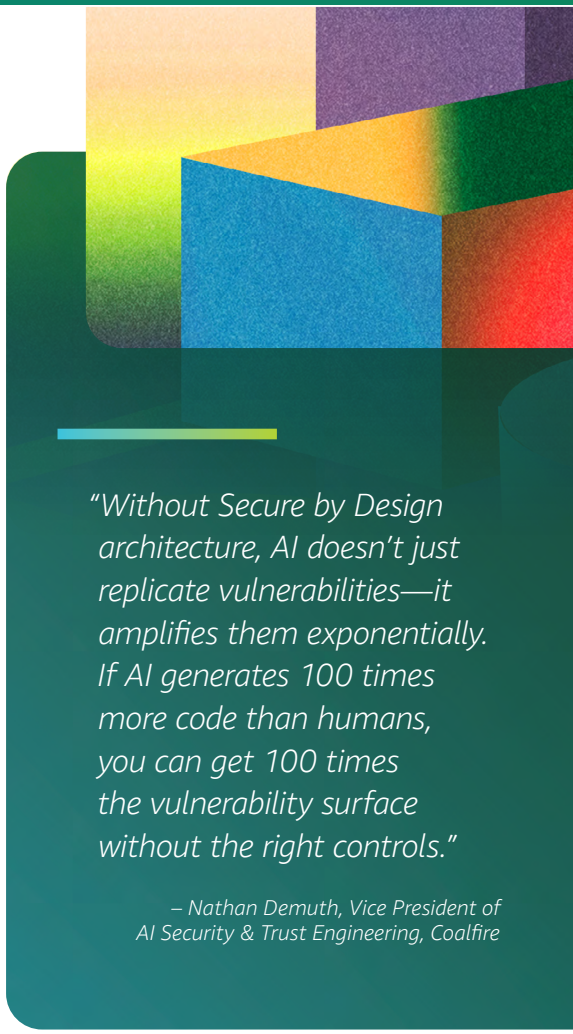

*"Without Secure by Design architecture, AI doesn't just replicate vulnerabilities—it amplifies them exponentially. If AI generates 100 times more code than humans, you can get 100 times the vulnerability surface without the right controls."*

*– Nathan Demuth, Vice President of AI Security & Trust Engineering, Coalfire*

## GuardianAI™ | A proprietary framework for production-ready AI architecture

No common definition exists for a mature, production-ready AI agent architecture. When contractors build AI agents, assessors ask, "Where is training data stored? How do you prevent information leakage?" Most teams can't provide clear answers.

Coalfire's proprietary GuardianAI framework provides that missing standard. It maps to FedRAMP, CMMC, FISMA, ISO/IEC 42001, NIST AI RMF, and OWASP AI standards, translating abstract requirements into concrete architectural definitions—specifying where data encryption occurs, how to implement audit logging, and how to structure access controls.

## LegionAI™ | Licensed, customizable AI agents that automate compliance operations

CISOs are told to "do more with less," yet 30-45% of security budgets go to compliance documentation that doesn't reduce risk. As AI agents expand the threat surface, teams need to redirect resources toward higher-value activities.

LegionAI provides licensed AI agents that automate governance, risk, compliance, and security operations. These agents handle high-volume tasks like collecting compliance evidence from AWS deployments, generating compliance updates, and tracking control implementation. Organizations see 15-20% operational cost savings—one contractor reduced their hiring budget by 45%. Beyond compliance, these agents support threat hunting and security operations, enabling CISOs to scale capabilities typically too expensive to staff.

## ForgeAI™ | Project-based consulting to implement AI trust patterns

Product teams often receive compliance frameworks but struggle to translate them into system architecture. When procurement offices reject proposals due to unclear security approaches, the real issue isn't understanding requirements—it's knowing how to build them.

ForgeAI delivers engineering consulting using GuardianAI to implement AI trust patterns: architectural blueprints that translate security controls into buildable designs. Organizations using ForgeAI cut remediation timelines in half.

# Getting started: From compliance burden to competitive advantage

Here's how to build trust when creating AI agent solutions:

1

### Establish your AI agent maturity baseline

Document your current business policies and procedures before deploying any AI automation. AI agents perform only as well as the instructions, memory, and guardrails you provide. If humans handle processes inconsistently, AI implementations will fail. First, define what "production-ready" means for your organization.

2

### Map architecture to regulatory controls

Identify which frameworks apply (FedRAMP, CMMC, FISMA, DoD CC-SRG) and document how your AI agent and workflow architecture address each control with specific engineering decisions. Create a remediation roadmap with concrete solutions for gaps between regulatory requirements and your current system. Organizations typically begin with a workshop or pilot project on an initial system.

3

### Automate to reduce operational costs

Analyze your security and compliance workflows to identify high-volume tasks that consume budget without adding security value. Deploy automation for generating things like compliance management, tracking control implementation status, or preparing quarterly audit materials. Redirect your security team to threat analysis and incident response.

4

### Demonstrate security to procurement teams

Procurement teams evaluate AI security through specific questions: How is data isolated? What happens if your model hallucinates? How do you prevent unauthorized access? Prepare clear, jargon-free answers with diagrams showing your testing procedures, incident response plans, and data handling protocols.

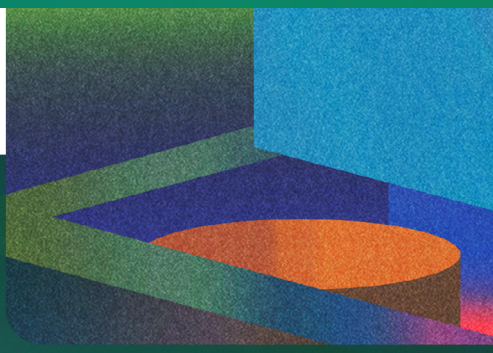

## *Is Coalfire right for your organization?*

- *Are you implementing AI agents in highly regulated industries?*
- *Has your procurement process stalled because you can't clearly explain your AI security approach?*
- *Is your security team spending more time on documentation than actual security work?*
- *Do you need to demonstrate AI security to government customers but lack a clear framework?*

*If you answered yes to any of the above, Coalfire may be a good fit.*

*"If your processes and procedures are ill-defined or poorly thought-out, then you're going to see a very poor success rate in your AI implementations."*

*– Nathan Demuth, Vice President of AI Security & Trust Engineering, Coalfire*

## *Why AWS and Coalfire?*

### *Core services*

***Amazon Bedrock and Amazon Bedrock AgentCore** provide managed AI runtime with built-in guardrails that enforce compliance policies programmatically.*

***Amazon Nova** delivers AWS-native intelligence that's optimized for infrastructure automation, understanding Amazon Simple Storage Service (Amazon S3), Amazon Elastic Compute Cloud (Amazon EC2), and CLI operations.*

***AWS's security foundation** provides compliance via AWS Key Management Service (AWS KMS), Amazon Virtual Private Cloud (Amazon VPC), Amazon CloudWatch, and AWS X-Ray.*

### *Partnership benefits*

- *33 geographic regions supporting government data residency*
- *AWS Well-Architected Reviews and Security Reference Architecture (SRA) alignment*
- *Integration with Salesforce, ServiceNow, Jira, Splunk, and other enterprise tools*

aws | COALFIRE

Get started by contacting the Coalfire team.
Visit Coalfire's AI Security and Trust Engineering page and see how to go from innovation to authorization. →