

AI Agent Production Readiness

Operational Design, Governance, Risk, Compliance & Security

Team Training: Onsite or Live Virtual
2 Days | 14 Hours



Over 300,000
business and IT
professionals trained

Many AI agent initiatives stall between specification and production - not because the technology fails, but because organizations cannot answer the fundamental question- Is this agent safe, governed, and operationally ready? Without structured governance and operational design, agents reach pilot but never reach production.

Inteq's AI Agent Production Readiness training course provides a structured, business-oriented approach for designing the complete governance architecture and operational framework that make AI agents production-ready.

Participants learn to design comprehensive guardrail architectures, specify transparency and explainability requirements for different stakeholder audiences, map regulatory and compliance obligations to testable agent behaviors, embed ethical and fairness requirements with measurable testing protocols, and conduct AI-specific threat modeling.

The course then bridges directly into operational design - resilience planning, multi-agent orchestration, lifecycle management with governance gates, structured feedback loops with learning boundaries, and capacity planning that accounts for governance overhead.

Governance decisions made on Day 1 of the course become the constraints that operational design must accommodate on Day 2 of the course - mirroring reality, where you cannot operate what you cannot govern

Without a disciplined governance and operations methodology, organizations face costly deployment failures, regulatory exposure, and stakeholder resistance. That's truly unfortunate because there is a clear path to getting production readiness right.

Based on decades of business analysis and process improvement experience, Inteq has developed a comprehensive set of governance, risk, compliance, and operational design frameworks.

Participants utilize these frameworks to systematically govern, secure, and operationalize AI agents across their organization's business processes.

You will learn:

- To design comprehensive agent guardrails with defense-in-depth enforcement that ensures agents cannot violate policy, regulations, or ethical norms
- To specify transparency and explainability requirements tailored to different stakeholder audiences - from technical teams to regulators
- Techniques for mapping regulatory and compliance obligations to concrete, testable agent behaviors and audit trail specifications
- To conduct AI-specific threat modeling covering prompt injection, data poisoning, and privilege escalation
- To design resilience architectures, multi-agent orchestration, and lifecycle governance that take agents from pilot to production
- ...and much more

Course Outline

Agent Guardrails, Constraints, and Safety Architecture

- Hard & soft constraints, authority limits, and scope boundaries
- Defense-in-depth enforcement architecture
- Guardrail specification templates and testing requirements
- Override protocols and escalation triggers

Trust, Transparency, and Explainability Requirements

- Six-level transparency spectrum
- Audience-specific explainability designs
- Explanation generation requirements by decision type
- Transparency-governance alignment

Audit, Compliance, and Regulatory Requirements

- Regulatory obligation mapping to agent behaviors
- Audit trail specification and evidence preservation
- Compliance testing protocols and monitoring
- Cross-jurisdictional regulatory considerations

Ethical, Responsible, and Trustworthy AI Requirements

- Bias detection and fairness testing frameworks
- Harm prevention protocols and accountability structures
- Ethical requirements specification with measurable criteria
- Ethics-security intersection analysis

Agent Security and AI-Specific Threat Modeling

- AI-specific attack vectors and threat taxonomy
- Prompt injection, data poisoning, and model manipulation
- Security control specification and testing requirements
- Threat model integration with governance architecture

Stakeholder Analysis and Governance Model Design

- Governance bodies, decision rights, and RACI matrices
- Oversight scaling and accountability structures
- Governance model completeness assessment
- Authority mapping across guardrails, compliance, ethics, and security

Agent Resilience and Multi-Agent Orchestration

- Failure mode and effects analysis (FMEA)
- Graceful degradation tiers within governance constraints
- Multi-agent orchestration topologies and conflict resolution
- Business continuity planning for agent systems

Agent Lifecycle Management and Learning Governance

- Lifecycle stages with governance gates
- Version management and change control
- Feedback loop design with learning boundaries
- Drift detection and autonomous evolution governance

Scalability, Capacity Planning, and Production Readiness

- Capacity planning with governance overhead accounting
- Performance modeling and cost-effectiveness at scale
- Scalability risk assessment
- Guardrail enforcement at production volume

Case Study

Participants carry a single AI agent opportunity through the complete operational and production readiness lifecycle - producing a comprehensive Production Readiness Package. This case study provides an invaluable template for production readiness work in your organization.

Who should attend?

- Business Analysts and Business Systems Analysts
- Compliance Officers, Regulatory Analysts, Risk Managers and Security Architects
- Business and Operations Managers
- Ethics and Responsible AI Leads
- Legal Counsel, Audit Managers and Program Managers
- Developers, Solution Architects and DevOps Leads
- Subject matter experts and business professionals involved in AI agent governance, compliance, and operational deployment

Prerequisites:

Prerequisites: Inteq's Discovering Agentic AI Opportunities and Analysis and Inteq's Specifying AI Agent Business Requirements courses – or equivalent experience with AI agent concepts presented in these courses.

This course integrates seamlessly with Inteq's Business Systems Analysis courses, Business Process Management courses and Inteq's other Agentic AI courses

What's included:

- Digital badge and personalized certificate of completion
- Continuing Education Units (CEUs)
- IIBA Professional Development Units (PDUs)
- Electronic comprehensive course manual
- Supplemental course materials including templates, frameworks, and scoring worksheets

TEAM TRAINING: (Onsite or Live Virtual): Inteq AI Agent Production Readiness training course can be tailored to your organization's specific training needs and objectives and can be combined with other Inteq training courses to create 3, 4 and 5-day hybrid training programs.



What is the Next Step?

Let's start a conversation to discuss your objectives in more detail.
Contact Chandra Galloway: 800.719.4627 | cgalloway@inteqgroup.com
Copyright © | The Inteq Group, Inc.

www.inteqgroup.com