

Agentic AI Governance At Scale: Securing Autonomous Systems in a Regulated World

As AI capabilities rapidly expand across industries, organizations face a critical challenge: how to innovate responsibly while meeting evolving security, ethical, and regulatory demands. This session explores what it takes to govern AI at scale—without slowing momentum.

Presented by: Peter Holcomb

720.266.9589 


<https://optimoit.io> 

ABOUT US


Optimo AI is a Denver-based security-first Managed Intelligence Provider (MIP) and System Integrator. We provide custom AI engineering solutions, fractional chief information security officer (CISO) leadership oversight, AI security & governance, and audit readiness services for scaling, highly regulated AI-driven businesses.

Founder & CEO of Optimo AI: Peter Holcomb

B.S. in IT from CSU, AI security & tech enthusiast, avid concertgoer, enjoys the outdoors, cooking, and supporting Colorado's vibrant cybersecurity community.



AI Governance Principles



What is AI Governance?

AI governance is a dynamic framework of policies, principles, and practices that guide how we build, deploy, and live with artificial intelligence.

Frameworks & Standards

Commonly Used:

- ISO 42001
- NIST AI RMF
- EU AI Act
- AIUC-1
- CSA—AICM
- SB-205
- GDPR
- CCPA



Increased Efficiency

Data-Driven Decisions

Competitive Advantage

Automating Repetitive Tasks

PRIVACY CONSIDERATIONS


Assessing Data Quality



7 PRINCIPLES OF PRIVACY BY DESIGN

1. Proactive not reactive
2. Privacy by default
3. Privacy embedded in design
4. Full functionality
5. End-to-end security
6. Visibility and Transparency
7. Respect for User privacy

The Responsible AI Life Cycle



Stakeholder Engagement

AI governance is not owned by any single team. It is an interdisciplinary effort, drawing on business leaders with a range of knowledge and experience. By mapping AI needs and impacts across an organization, these leaders can make the right investments to support their broader business goals. This means aligning IT, financial, cybersecurity, and other business operations to help ensure they drive toward the same vision of success.

Audit and Assessments:

Ensure AI systems comply with internal and external standards through regular evaluations and assessments.

Cybersecurity:

Protect AI systems from cyber threats and vulnerabilities, preserving data integrity and security.

Enterprise Risk Management:

Convene teams and stakeholders to track, prepare for, and mitigate AI risk holistically across the organization.

Information Technology:

Support the technical infrastructure to help ensure AI systems are reliable, scalable, and aligned with standards and requirements.

Operational:

Integrate AI into business processes efficiently and safely, while equipping employees with the necessary resources to maintain responsible AI standards.

Third-Party:

Manage relationships with AI vendors and partners; validate responsible AI practices throughout the third-party ecosystem.



Communications:

Disseminate information regarding AI initiatives; promote transparency and foster trust through clear and consistent messaging.

Fairness and Inclusion:

Evaluate the impact and potential AI bias toward certain groups of individuals, often by performing AI risk assessments

Human Resources:

Facilitate user adoption by addressing workforce impacts, providing training, and promoting a culture of responsible AI use.

Legal:

Navigate the increasingly complex regulatory landscape, while maintaining organizational compliance with laws and ethical guidelines to mitigate legal risk.

Sustainability:

Promote environmentally safe development goals and energy-efficient practices.

Understanding AI Systems


 Definitions

Automation vs. AI workflow vs. AI agent

	Automation	AI workflow	AI agent
Definition	 A program that executes predefined, rule-based tasks automatically	 A program that calls an LLM via API for one or more steps	 A program designed to perform non-deterministic tasks autonomously
Core foundations	 Boolean logic	 Boolean logic  Fuzzy logic	 Fuzzy logic  Autonomy
Tasks	Deterministic, predefined tasks	Deterministic tasks requiring flexibility	Non-deterministic, adaptive tasks
Strengths	• Delivers reliable outcomes • Fast to execute	• Better handling of complex rules • Great for pattern recognition	• Highly adaptive to new variables • Simulates human-like behavior and reasoning
Weaknesses	• Limited to tasks explicitly programmed • Cannot adapt to new scenarios • Struggles with complexity	• Requires data to train models effectively • Harder to debug and interpret	• Less reliable, may produce unpredictable undesired outcomes • Slower to execute
Example	Send a Slack notification every time a new lead signs up on our website	Analyze, score and route every website inbound lead using ChatGPT	Perform a full internet search on every inbound lead and update infos

Agentic AI Architecture Design Patterns

Agentic Design Patterns



Agentic Design Patterns






Figure 1: 7 Layer Reference Architecture for Agentic AI

What are AI agent protocols?

AI Agent Protocols

- » Agent2Agent (A2A) protocol
- » Agent Communication Protocol (ACP)
- » Agent Network protocol (ANP)
- » Model Context Protocol (MCP)
- » Agent-User Interaction (AG-UI)



Here are some key aspects to keep in mind when assessing agent protocols:

Efficiency, **Reliability**, **Scalability**, **Security**

Agentic AI Threat Modeling

Agentic AI introduces new promise—and new peril. With agents taking autonomous actions and collaborating with other agents or users, our security mindset needs to evolve. But with structured threat modeling, playbook-driven mitigations, and security built into the AI lifecycle, we can design resilient, trustworthy AI systems.

- AI-specific threat modeling
- Model lifecycle risk assessments
- Secure data pipelines & lineage tracking
- Red team simulation for AI abuse cases

CSA MAESTRO

(Multi-Agent Environment, Security, Threat, Risk, and Outcome), a novel threat modeling framework designed specifically for the unique challenges of Agentic AI



OWASP MAS

The Multi Agent System guide builds on the OWASP Agentic AI – Threats and Mitigations publication, our master agentic threat taxonomy



AGENTIC AI THREAT MITIGATIONS

RISKS


-  Memory Poisoning
-  Privilege Compromise
-  Impersonation
-  Human Manipulation
-  Hallucinations



MITIGATION

-  Memory Validation
-  Access Controls
-  Hallucination Guardrails
-  Human Oversight
-  Multi-Agent Security

The Importance of AI Observability and Debugging



```
lindy-observability.s3.us-east-2.amazonaws.com/logs/test.json

{
  "timestamp": "2025-07-28T21:48:00-06:00",
  "task_id": "6883246cb2924b9d2e3cb3dd",
  "agent_name": "Lindy KB Agent",
  "agent_id": "6880283bc2e47a48988e4dd1",
  "status": "working",
  "user_query": "Tell me how to set up AI observability",
  "response_time_ms": 15000,
  "token_usage": {
    "input_tokens": 4102,
    "output_tokens": 490,
    "total_tokens": 4592
  },
  "cost_cents": 97,
  "knowledge_base_search": {
    "query": "AI observability setup implementation monitoring",
    "results_count": 2,
    "relevance_scores": [0.7232922315597534, 0.711804211139679],
    "sources": ["WEB_SCRAPED", "WEB_SCRAPED"]
  },
  "workflow_steps": [
    {
      "step": "message_received",
      "timestamp": "2025-07-25T06:30:05.370Z",
      "duration_ms": 13
    },
    {
      "step": "knowledge_base_search",
      "timestamp": "2025-07-25T06:30:05.426Z",
      "duration_ms": 1995
    },
    {
      "step": "llm_processing",
      "timestamp": "2025-07-25T06:30:10.629Z",
      "duration_ms": 14805
    },
    {
      "step": "response_sent",
      "timestamp": "2025-07-25T06:30:25.607Z",
      "duration_ms": 0
    }
  ],
  "performance_metrics": {
    "total_execution_time_ms": 20237,
    "knowledge_base_hit": true,
    "knowledge_base_relevance": "limited",
    "response_quality": "comprehensive_despite_kb_limitations",
    "user_satisfaction_indicators": {
      "acknowledged_limitations": true,
      "provided_alternative_guidance": true,
      "offered_follow_up": true
    }
  },
  "anomalies_detected": [
    {
      "type": "knowledge_base_gap",
      "description": "Agent KB lacks comprehensive AI observability content",
      "severity": "medium",
      "impact": "Agent provided general guidance outside KB scope"
    }
  ],
  "recommendations": [
    "Expand knowledge base with AI observability best practices",
    "Consider specialized AI observability agent",
    "Monitor for similar out-of-scope queries"
  ]
}
```

Observability and Debugging

- Individual Agent Monitoring
- Inter-Agent Communication Tracing
- Shared State Analysis
- Root Cause Analysis
- Explainability

observability
MessagesAdd canvas+
peter 9:41 PM joined #observability. Also, Lindy joined. AI Observability Agent APR 9:49 PM AI Observability Agent Alert - New Task Analysis
Monitor Details
Monitor Name: Agent Task Resolution Monitor\
Type: Workflow Performance Analysis\
Timestamp: July 28, 2025 9:48 PM MDT\
Alert Severity: Medium (Knowledge Base Gap Detected)
Affected Component/Service
Agent: Lindy KB Agent (ID: 6880283bc2e47a48988e4dd1)\
Task ID: 6883246cb2924b9d2e3cb3dd\
Service: AI Knowledge Base Query Processing
Performance Summary
- Response Time: 15.0 seconds (within normal range) - Token Usage: 4,592 total tokens (4,102 input, 490 output) - Cost: \$0.97 in credits - Status: Working
Key Metrics & Trends
Workflow Execution Breakdown:
- Message Processing: 13ms - KB Search: 1.995ms - LLM Processing: 14,805ms (largest component) - Response Delivery: <1ms
Knowledge Base Performance:
- Search Query: "AI observability setup implementation monitoring" - Results: 2 documents found - Relevance Scores: 0.72, 0.71 (moderate relevance) - Issue: Limited to Lindy-specific documentation
Anomaly Detected
Type: Knowledge Base Content Gap\
Description: Agent's KB lacks comprehensive AI observability content, forcing reliance on general LLM knowledge\
Impact: Agent provided helpful response despite KB limitations, but this indicates potential service gap
Tags/Labels
agent-performance knowledge-base ai-observability content-gap user-query-analysis
Related Resources
- Dashboard: Agent Performance Dashboard - Task URL: https://chat.lindy.ai/optimo/lindy/lindy-kb-agent-6880283bc2e47a48988e4dd1 - Logs: Available in task execution blocks
Recommendations
- Expand KB Content: Add comprehensive AI observability documentation - Monitor Pattern: Track similar out-of-scope queries for content planning - Consider Specialization: Evaluate need for dedicated AI observability agent
Overall Assessment: Agent handled query effectively despite KB limitations, maintaining transparency and providing valuable guidance.

Agentic AI Guardrails

RBAC

Limit what actions each agent is allowed to perform based on roles

Resource Limits

Constratin resources such as CPU or LLM tokens that each agent can consume

Rate Limiting

Prevents agents from making too many requests within short timeframe

Validation Rules

Define rules that all agent outputs must satisfy to ensure alignment with requirments

Authentication

Require agents to authenticate their identities preventing unauthorized actions

Data Isolation

Isolate certain types of data between agents to ensure authorized use of data

Human Override

Provide human operators to intervene and manually take control of key systems

Anomaly Detection

Use advanced alorithms to detect unusual agent behavior that may indicate issues or risks

Audit & Compliance

Enable auditing and reporting tools to document actions and ensure compliance



Common Agentic AI Use Cases

- Contract Evaluation
- Financial Statements
- Process Automation
- Lead Generation
- Compliance
- Email Assistant
- Meeting Notetaker

Created by me

Shared with me

New Folder

New Agent

...

Optimo IT PROD Agents

Name	Last Run	Integrations	Enabled
@ Agentic Security Hub	7 days ago		
AI Security Content Creator	About 23 hours ago		
Compliance KB Agent	13 days ago		
Enrich New Leads	13 days ago		
Lindy KB Agent	About 1 month ago		
PROD AI Observability Agent	28 days ago	HTTP	
PROD Email Assistant	10 minutes ago		
PROD Lead Outreach	13 days ago	+2	
PROD Meeting Notetaking Agent v1.0	About 15 hours ago	+6	
PROD Meeting Prep Agent	10 minutes ago	+4	
User Access Review Agent	28 days ago	HTTP	
Website Customer Support	2 months ago		

The Concept of Swarming



Transportation

Multi-agent systems can be used to manage transportation systems.



Healthcare & public health

Multi-agent systems can be used for various specific tasks in the healthcare field.



Supply chain management

Multi-agent systems can use their vast informational resources, versatility and scalability to connect the components of supply chain management.



Defense systems

Multi-agent systems can aid in strengthening defense systems.



Education, Training & Awareness

Education and awareness initiatives can prepare employees to integrate AI into their work responsibly and effectively. These efforts may involve both general and role-specific training on using AI, understanding policies and guidelines, and meeting compliance and reporting requirements. In addition to formal education and training, successful integration requires embedding AI into an organization's culture and processes through cross-functional buy-in and a collaborative approach.



Introduction to AI

These courses are designed to ensure a basic foundational understanding for employees to learn the basics of AI. They outline the basics of AI technology, potential applications, and provide knowledge to be conversant on the topic.



AI Governance

These courses help develop an understanding of the best practices for planning, guiding, and maintaining AI systems and instilling accountability for AI systems and technologies.



Privacy, Data Protection, and Security

These courses introduce privacy concerns regarding the responsible use of AI and data protection education for secure development and deployment of AI systems.



AI Ethics

These courses are designed to give the learner an understanding of ethical principles to consider when developing and/or using AI solutions.



AI Bias

These courses are designed to provide education on different types of bias, how they can be introduced into AI systems and how users of systems can ensure that fairness is incorporated when utilizing AI in business decisions or analysis.



AI Transparency and Explainability

These courses are designed to educate on the importance of transparency when designing and deploying AI systems. They also include content related to challenges of explaining why an AI system produced a particular output.

The Return on Investment of AI

Key insights

Chapter 01

AI agents are now being deployed at scale.

52%

of executives whose organizations use gen AI also have adopted AI agents in production²

Agentic AI early adopters enjoy a significant advantage.

88%

of executives from agentic AI early adopter orgs see ROI on at least one gen AI use case³

Gen AI continues to deliver returns.

74%

of executives see ROI on at least one gen AI use case⁴

ROI is thriving across use cases.

39%

of executives saw ROI on gen AI use cases for productivity, customer experience (37%), sales and marketing (33%)⁵

Executive backing drives AI success.

78%

of executives from organizations with C-level sponsorship report seeing ROI now on at least one gen AI use case⁶

Challenges still present roadblocks to implementation.

#1

Data privacy and security is the top consideration for companies when evaluating LLM providers⁷

² Total (global): n=3466; Question text: (1) How are AI agents leveraged across the enterprise? (2) How many AI agents does your company currently have deployed in production across your organization?

³ Agentic AI early adopters: n=460; Question text: In what timeframe do you expect gen AI to deliver return on investment (ROI) to the following areas of your business?


⁴ Total (global): n=3466; Question text: In what timeframe do you expect gen AI to deliver return on investment (ROI) to the following areas of your business?

⁵ Executives whose organization is currently using or planning to use gen AI on use cases: Individual productivity: 2024: n=1451, 2025: n=3271; Customer experience and field service: 2024: n=1445, 2025: n=3285; Sales and marketing: 2024: n=1454, 2025: n=3245 Question text: In what timeframe do you expect gen AI to deliver return on investment (ROI) to the following areas of your business?

⁶ Executives who report their organization has comprehensive C-level sponsorship and clear corporate vision for gen AI objectives: n=1269; Question text: In what timeframe do you expect gen AI to deliver return on investment (ROI) to the following areas of your business?

⁷ Total (global): n=3466; Question text: Which of the following factors are MOST important to your company when considering LLM providers? Please select the top three.

ARTIFICIAL INTELLIGENCE REVOLUTION: 10 U.S. STATES LEADING THE PACK



1. California
2. Florida
3. Massachusetts
4. Texas
5. New York
6. Illinois
7. North Carolina
8. **Colorado**
9. Washington
10. Pennsylvania

8. *Colorado* leads the nation in business AI adoption, the analysts say, because of its aerospace and renewable energy sectors underscoring its growth potential.

Final Takeaway On The Artificial Intelligence Revolution

“What's fascinating about our analysis is how it challenges the notion that AI innovation belongs exclusively to Silicon Valley,” says Danny Veiga, AI Technology Strategist and CEO of **Chadix**. “We're seeing states like Florida leap into second place overall, while Colorado leads the nation in AI business adoption.”

Overcoming Challenges in AI Implementation

Agentic AI orchestration, while powerful, presents several challenges:

Common Barriers:

- **High initial costs for AI tools**
- **Data privacy and security concerns**
- **Need for employee training and adaptation**
- **Debugging and Security**
- **Complexity and ethical considerations**

Solutions:

- **Start with affordable, scalable tools**
- **Implement data protection measures**
- **Offer training sessions on AI usage**
- **Implement robust security practices**
- **Ensure responsible use of AI systems**

Risk Assessment Resources & Certifications


Take advantage of readily available AI security assessment tools to determine your AI security governance posture. You'll gain valuable insight into your overall maturity score, room for improvement and understand gaps for continuous monitoring towards successful AI adoption.

Tools:

- Cloud Security Alliance - AI Controls Matrix
 - Mapping to:
 - ISO 42001
 - NIST AI 600
 - BSI AIC4
- Trusted AI Safety Expert Certificate Program
- NIST AI RMF Playbook



Future of AI in Productivity



Enhanced automation tools

Modular systems allow seamless integration of new components, quickly adapting to evolving needs. This agility is powered by an AI Orchestration platform.



AI in remote work management

Dynamic resource allocation optimizes resource utilization, minimizing waste and driving cost savings.



Hyper-personalization in customer experience

Distributed architecture enables seamless scaling. New agents handle growing workloads through dynamic delegation.

Outlook

AI adoption will continue to evolve, reshaping productivity across all sectors.

THANK YOU

FOR YOUR ATTENTION AND PARTICIPATION

Please reach out to us if you'd like to learn more about how you can safely adopt and implement AI within your business. We're ready to be your strategic managed intelligence partner on your AI transformation journey.

Contact Us:

 720.266.9589

 <https://optimoit.io>

 peter@optimoit.io

 4045 Wadsworth Blvd
Ste. 260
Wheat Ridge, CO 80033



Resources

1. [AI Governance](#)
2. [Agentic AI - Threats and Mitigations](#)
3. [Multi-Agentic System Threat Modeling](#)
4. [A practical guide to building agents](#)
5. [OWASP Top 10 for LLM Apps](#)
6. [CSA Agentic AI Threat Modeling Framework: MAESTRO](#)
7. [Beyond Rules: Agentic Orchestration](#)
8. [Top Agentic AI Design Patterns](#)
9. [NIST AI Risk Management Framework](#)
10. [EU AI Act](#)
11. [General Data Protection Rule \(GDPR\)](#)
12. [What is a multi-agent system](#)
13. [What are AI agent protocols](#)
14. [The ROI of AI 2025](#)
15. [AI Revolution: 10 U.S. States Leading the Pack](#)