CyberRisk ALLIANCE

Agentic AI Executive Reporting Toolkit

AI is now a board-level priority, and security teams are expected to clearly articulate the risks, controls, and governance considerations that come with adoption. This template provides a framework for communicating AI risk to executive leadership, along with a 90-day roadmap to help guide implementation efforts.

Objective

Give leadership clear visibility into adoption, risk, and readiness.

Board & Executive Reporting Topics

- Current agentic AI adoption
- High-risk use cases
- Security and privacy controls
- Third-party AI exposure
- Incident trends
- Policy exceptions
- Regulatory and compliance considerations
- Business value delivered
- Program maturity roadmap

Executive-Level Questions to Answer

1. Where are we using agentic AI today?
2. Which use cases carry the most risk?
3. Who owns those risks?
4. What controls are in place?
5. How do we know agents are behaving as intended?
6. What could go wrong?
7. Are we prepared to respond?
8. What investment is needed to scale safely?

Core Deliverables

- Board briefing deck
- Executive risk dashboard
- AI maturity scorecard
- Investment roadmap

90-Day Implementation Roadmap

Days 1–30 Establish Foundation

1. Create program charter
2. Form governance team
3. Inventory current AI use
4. Define risk tiers
5. Draft acceptable use policy
6. Identify pilot candidates

Days 31–60 Build Controls

1. Finalize review process
2. Create security control framework
3. Implement logging requirements
4. Define identity and access standards
5. Launch role-based training
6. Begin controlled pilots

Days 61–90 Operationalize and Scale

1. Review pilot outcomes
2. Approve initial production use cases
3. Integrate monitoring into SOC workflows
4. Launch executive dashboard
5. Conduct tabletop exercise
6. Publish quarterly roadmap

CISO Success Factors

To make the program effective, CISOs should focus on five principles:

1. **Treat agents as identities, not tools.**
   Every agent needs an owner, permissions, logs, and accountability.
2. **Control autonomy before scale.**
   The higher the autonomy, the stronger the approval, monitoring, and rollback requirements.
3. **Secure the data supply chain.**
   Agentic AI risk often begins with excessive access to sensitive data.
4. **Build governance into adoption.**
   Business teams need a clear path to safe experimentation, not vague restrictions.
5. **Measure both value and risk.**
   Security should help the organization prove where agentic AI is useful, safe, and worth scaling.

Accelerator Resources

CyberRisk Collaborative Accelerator Program Subscribers get exclusive access to pre-built templates:

- Agentic AI Implementation Roadmap
- Agentic AI Inventory Tracker
- Vendor AI Assessment Scorecard
- Agentic AI Kill Switch Framework
- Human in the Loop Control Matrix

ABOUT CyberRisk Collaborative

The CyberRisk Collaborative is a members-only community that offers a full suite of leadership, technical, strategic, and personal brand building services for CISOs and their teams. Cybersecurity executives leverage this peer-to-peer community model to validate their strategies and integrate proven best practices from other CISOs into their security programs.

Membership Inquiries