LDR550: **Leading Agentic Security Teams™**

2 Day Course	12 CPEs	Laptop Required

You Will Be Able To

- Redesign the security operating model around governed AI so teams can deliver more without losing accountability
- Set an appropriate level of autonomy for each security workflow
- Place human approval where an incorrect action would be costly or difficult to reverse
- Select AI capabilities the team can oversee without becoming dependent on one vendor
- Build a phased roadmap toward AI-native cybersecurity operations
- Prepare for automation bias, skill atrophy, and situations when AI is wrong or unavailable
- Explain the approach to executives and colleagues to gain organizational support

Business Takeaways

- Reduce analyst toil and cost through governed automation
- Accelerate security decisions without surrendering oversight
- Strengthen accountability for AI-driven actions, with a named owner for each one
- Lower concentration risk in the AI tools the team depends on
- Expand the depth and breadth of the security team's capabilities to improve outcomes

Who Should Attend

This is a leadership course for the people who set a security program's direction and run its teams. Students learn to design and govern the AI-native security function, and the course expects no hands-on work with the tools.

- CISOs and Senior Security Leaders
  - Leaders who own strategy, budget, and the security operating model
- Security Directors and Team Leads
  - Leaders who put the transition into practice within their teams

Security leaders are under pressure to do more with AI, but most add it piecemeal without changing how the team operates. This course teaches you to design the security function, so your team delivers more, with routine work handed to AI under your direction. You'll decide which work to delegate, at what level of autonomy, and who stays accountable, practice them in the SANS Cyber42 simulation, and start the plan for your own organization.

Most security teams add AI at the edges of the work, an assistant here and a review tool there, and the gains stay marginal because the operating model has not changed. AI can do more when it handles the routine work, with people setting policy, making the hard calls, and supervising the system.

You'll learn how to decide, for every security workflow, which actions AI may take on its own, who approves or overrides it, and who answers for the outcome. You'll learn to weigh each grant of autonomy against the damage a wrong action could do and how hard it would be to undo.

You'll practice on a fictional organization in the SANS Cyber42 simulation, where your decisions fill in the organization's Security Autonomy Matrix, the single table you can use to capture decisions. In the capstone you start the same plan for your own security function, from the budget, team, and tools you have today.

Author Statement

I've watched security teams add AI to their security program in incremental ways without meaningful results. The problem was the lack of an approach for putting AI at the core of the program without taking on undue risk.

I designed this course around two questions a leader faces when progressing to an AI-enabled security program. What may the AI do on its own, and who answers for the outcome? You'll learn how to make such calls for each aspect of the security team's work in a deliberate, measured way. This way, you'll place people where their judgment matters most, deploy AI where it offers value, and adjust AI agents' authority as they gain or lose trust.

—Lenny Zeltser

Section Descriptions

SECTION 1: What Good Looks Like, and the Framework for Getting There

Section 1 covers what a strong AI-native security program looks like and the decision framework behind it, the autonomy spectrum and the Security Autonomy Matrix. You'll baseline a security team's current autonomy, map its workflows, decide where human approval is required and apply the framework to security operations.

TOPICS:

- The AI-native security organization
- The autonomy spectrum and the decision framework
- Governance, accountability, and human oversight
- Applying the framework to security operations

SECTION 2: Applying the Framework and Leading the Change

Section 2 covers the use of the Security Autonomy Matrix beyond the SOC, the operating-model and team changes that follow, how to select AI capabilities, and how to manage the risks of running security on AI. You'll leave with a started matrix and a sketched 90-day roadmap for your own organization.

TOPICS:

- Applying the framework across the security function
- The operating model and the team
- Selecting AI capabilities for the security team
- Risks, failure, and the roadmap