Lumos

A perspective on AI maturity in security

The Cyber AI Transformation Matrix

LUMOS · FOR SECURITY LEADERS · 2026

WHY THIS EXISTS

A maturity map for AI in security

We built maturity models to help companies move to the cloud. We do not have the same map for AI in cybersecurity yet. That's what this is about: a Cyber AI Transformation Model. It's a simple way to see where your security org is today, and where it needs to be next quarter, not five years from now.

The model splits each security domain across four levels of AI capability, from Thought Partner (a chat interface that helps you reason) to The System (autonomous operation at machine speed). The same four levels apply whether you are talking about identity, endpoint, network, or app security.

HOW TO USE THIS

Treat this as a domain-by-domain assessment, then decide where machine-speed defense matters most. Maturity will not move evenly across security, and that is to be expected. Identity might be at Level 3 while network security is still at Level 1.

THE FOUR LEVELS

L1
Thought Partner
A chat interface that helps you research, draft, and reason. The human still owns the work.

→

L2
Assistant
AI is grounded in your environment. It reads your data and surfaces context the human acts on.

→

L3
Teammate
Agents take action. Humans review and approve, but the work is being done by the system.

→

L4
The System
Autonomous defense at machine speed. The human sets policy and reviews outcomes.

THE MATRIX

Where each domain stands across the levels.

Eight domains, four levels of AI maturity. Use this as a self-assessment grid: highlight the cell that best describes where your organization is today, then circle where it needs to be by the end of the next quarter.

Security domain	LEVEL 1 Thought Partner	LEVEL 2 Assistant	LEVEL 3 Teammate	LEVEL 4 The System
Identity Security	"Who has access to what?" via chat.	Role mining and entitlement classification using your org's data.	Agents run access reviews and route access requests.	Continuous least-privilege enforcement; auto-remediate toxic permission combos.
Endpoint Security	Research threat indicators, analyze suspicious processes.	Enrich alerts with device, user, and asset context.	Agents isolate compromised devices and execute remediation.	Self-healing fleet; coordinated response across devices and identities.
Network Security	Analyze logs, draft firewall rules.	Detect traffic anomalies using your actual topology.	Agents tune firewall rules and adjust segmentation.	Perimeter adapts autonomously during active incidents.
Cloud Security	Research misconfigs, ask compliance questions.	Score cloud posture across your accounts.	Agents fix config drift and enforce posture.	Compliance orchestrated across clouds; blast radius auto-contained.
Data Security	Ask about classification rules, draft retention policies.	Discover and label sensitive data using business context.	Agents enforce DLP and flag anomalous access patterns.	Data flows governed in real time across systems.
SecOps	Draft incident reports, map to ATT&CK.	Triage alerts using your SIEM history and environment.	Agents investigate alerts and handle tier-1 triage.	Full loop: detect, investigate, respond, close - without waiting for a human.
Email Security	Analyze suspicious headers and URLs.	Score senders against your internal mail patterns.	Agents detect BEC by writing style and auto-quarantine.	Mail, auth, and click signals coordinated for pre-human response.
App Security	Research CVEs, ask about secure patterns.	Scan code with awareness of your ownership and dependencies.	Agents patch vulnerabilities in CI/CD pipelines.	Continuous patching with autonomous rollback on failure.

Read across a row for the journey of one domain. Read down a column for the texture of one maturity level across security as a whole.

A POINT OF VIEW

Choose where Level 4 matters most.

Maturity will move at different speeds across security. Network might sit at L1 while identity is already at L3. The discipline is deciding explicitly, which domain earns autonomous defense first.

Our take: identity should sit at the top, because most major breaches start with a compromised account, and access volume has outgrown human review. Start there.

RECOMMENDATION

30 Day Game Plan

1. **Map every domain to a current level.**
   Walk the matrix with the team that owns each domain. Be honest about where you actually are, not where the vendor deck says you are.

2. **Pick one domain to push to Level 4.**
   Concentrating depth beats spreading effort. The right choice is usually the domain with the highest volume of decisions humans cannot keep up with.

3. **Set a quarterly target, not a five-year plan.**
   Identify the next level for every domain by the end of next quarter, and write it down. AI capability is moving faster than annual planning cycles.

CUSTOMER SPOTLIGHT

netskope

> Nearly 80% of breaches stem from excessive access. Being able to tighten our risk posture with Lumos while improving the user experience has been huge.

RK

Rehman Khan
Chief Information Security Architect, Netskope

See how to get to Level 4 in identity security.

Lumos secures every agent in your environment, whether human, NHI or AI agent with the world's most powerful agentic identity management platform.

Book a demo →