LOGICALIS
Architects of Change

CISCO
PARTNER

Executive Brief

Defending Against the AI-Powered Adversary

5 strategies to secure the organization in the face of AI-related vulnerabilities



Executive Summary

In just a few short years, AI has moved from an experimental technology to an everyday business tool, a trend that is fundamentally reshaping the enterprise security landscape. The same AI capabilities that help organizations work faster are now helping cyber attackers adapt and strike with unprecedented speed. This brief examines how AI-powered adversaries operate today and how teams can better defend themselves.

83% of organizations plan to deploy agentic AI

Source: Cisco State of AI Security 2026

29% feel ready to do so securely

Source: Cisco State of AI Security 2026

AI-powered adversaries are becoming increasingly autonomous. Tasks that once required skilled human operators, such as reconnaissance, phishing, and exploitation, can now be carried out by AI with little human involvement. Attackers, for instance, now deploy AI agents that can scan for vulnerabilities, harvest credentials, and move through networks with minimal human direction.

The enterprise's newest assets have become its newest vulnerabilities. Agentic AI, which organizations are deploying across critical workflows, creates a new attack surface, while shadow AI spreads ungoverned tools faster than IT can track them. Non-human identities, which now far outnumber human users, have become a primary target for attackers, and gaps in AI governance can now carry real regulatory consequences.

The advantage now belongs to defenders who can see and respond fastest. The strongest defenses combine automated, AI-accelerated detection and response with continuous AI-asset discovery, automated AI governance, and identity-first controls that extend to non-human identities. To keep pace with the rapidly evolving security landscape, many teams are turning to managed services offerings to free internal resources and fill critical security gaps.

The sections that follow cover five key AI threats that enterprises are facing today and practical strategies for defending against them.

1. AI-Enabled Attacks

Attacks are escalating in speed and scale

Attackers are using AI tools to sharpen phishing, social engineering, and attack campaigns. Modern intrusions often extend beyond a single network boundary, using compromised identities and trusted connections to reach across the IT estate to cloud services, endpoint management systems, and internet-facing infrastructure. Offensive AI tools are even being sold as subscription services, which puts advanced capabilities within reach of low-skilled attackers.

62% of organizations reported experiencing a deepfake attack within the previous 12 months

-Gartner

In the coming years, the pace will only accelerate. Gartner predicts that by 2027, AI agents will cut the time it takes to exploit exposed accounts in half by automating deepfake-driven social engineering, credential abuse, and more.

Defense strategies

- Match AI-assisted threats with automated detection, cross-environment correlation, and rapid containment, supported by behavioral analytics that identifies unusual identities, devices, and access patterns.
- Harden identity verification with phishing-resistant MFA, independent confirmation for high-risk requests, and awareness training that covers AI-enabled phishing, voice cloning, and deepfake impersonation.
- Prioritize protection of backups, identity services, endpoint-management systems, and virtualization control planes, because compromise of these systems can disrupt operations and undermine recovery.

2. Agentic AI

Agentic AI is creating a new attack surface

AI agents are moving from pilot to production faster than security controls can mature. Enterprises are giving agents authority to access databases and modify code, which introduces risks that include prompt injection, tool misuse, excessive privileges, memory poisoning, and more. Attackers may also target the agentic supply chain by planting malicious code or instructions in repositories, AI skills, and MCP servers that agents rely on.

Security operations teams are already adapting: Gartner predicts that by 2028, 50% of all enterprise incident response efforts will focus on incidents involving custom-built AI applications, many of which may be deployed before security testing and governance have fully matured. As regulation and procurement expectations evolve, agent-specific controls such as least-privilege access, approved tool integrations, action logging, and human approval for high-impact actions are likely to become more common.

Defense strategies

- Secure agents as privileged non-human identities: inventory their models, data sources, skills, tools, and MCP integrations. Assess connected components, sandbox high-risk actions, and enforce runtime governance.
- Enforce least-privilege tool authorization and scoped, short-lived delegated credentials, with human approval for high-impact actions.
- Deploy end-to-end telemetry and behavioral anomaly detection across agent interactions, including prompts, tool calls, identities, data access, and actions.

3. Non-Human Identities

Non-human identities have become a primary entry point for attackers

Service accounts, API keys, tokens, and AI agents can far outnumber human users in many enterprises, and attackers increasingly target credentials and permissions to gain access. With the proliferation of AI agents, traditional identity models leave organizations with limited visibility into agent actions across connected systems. Agents also raise the stakes further because they use dynamic, delegated, or short-lived credentials that many identity programs were not designed to govern at scale.

Identity sprawl, automation, and agentic adoption are continuing to accelerate, making governance a critical concern for security teams. Agent-identity controls such as signed actions, delegated scopes, attestation, and detailed audit logging are likely to become more common in security reviews and procurement requirements.

Defense strategies

- Build a dedicated program for non-human identities that inventories each identity, assigns an owner, scopes permissions, shortens credential lifecycles, and flags anomalous behavior after authentication.
- Extend Zero Trust to agents by discovering, authenticating, authorizing, and continuously evaluating agentic identities as rigorously as human users.
- Monitor exposed infrastructure, including unauthorized or insecure MCP servers reachable from the public internet, and restrict unmanaged agent-to-agent connections.

4. Shadow AI

Ungoverned AI use is now the norm inside enterprises

Shadow AI has become mainstream, as workers adopt AI tools faster than IT can govern them. Exposure arrives continuously and often quietly, as a browser extension, a niche vertical tool, a personal AI account, or a new AI feature in existing SaaS can appear between review cycles without anyone noticing. Each of these can then open a path for sensitive data to leave the organization through channels security teams cannot see or control.

57% of CIOs say employees' AI use jeopardizes data security

-Logicalis 2026 CIO Report

Left unchecked, the risk compounds. Gartner predicts that by 2030, more than 40% of enterprises will experience security or compliance incidents linked to shadow AI. The result: the longer organizations wait to address the governance imperative, the larger their attack surface grows and the more costly remediation becomes.

Defense strategies

- Replace blocking with discovery and visibility, since banning AI without a usable sanctioned alternative doesn't stop usage, it simply drives it further underground.
- Run continuous AI-asset discovery across cloud, endpoint, and network.
- Apply data loss prevention (DLP) and prompt and output inspection to both sanctioned and newly discovered AI tools on a continuous governance cycle.

5. AI Security Governance

Governance shifts from best practice to enforceable law

In 2026, AI governance moved from recommended practice to legal obligation. Since August 2026, the EU AI Act has required that organizations clearly label AI-generated content, such as deepfakes, so people know when they are seeing or hearing something created by AI. Stricter rules for high-risk AI, such as systems used in hiring, credit decisions, or critical infrastructure, take effect in December 2027 and require organizations to prove those systems are properly documented, monitored, and overseen by people. To comply, organizations need a clear view of where AI is in use, how it is being used, and whether it meets these requirements.

The cost of falling behind is steep. Gartner predicts that through 2027, relying on manual AI compliance processes will put 75% of regulated organizations at risk of fines exceeding 5% of their global revenue. As a result, governance is evolving from a one-time project into an ongoing necessity that covers AI inventory, risk management, and monitoring.

Defense strategies

- Maintain a live AI inventory as the foundation for both compliance and security, since organizations cannot govern AI systems they cannot see.
- Automate compliance processes and establish a cross-functional governance board spanning security, legal, data, and HR to provide continuous oversight.
- Extend third-party risk management to the AI supply chain, covering models, datasets, MCP servers, and agent skills.

Keep Pace with AI-Powered Adversaries

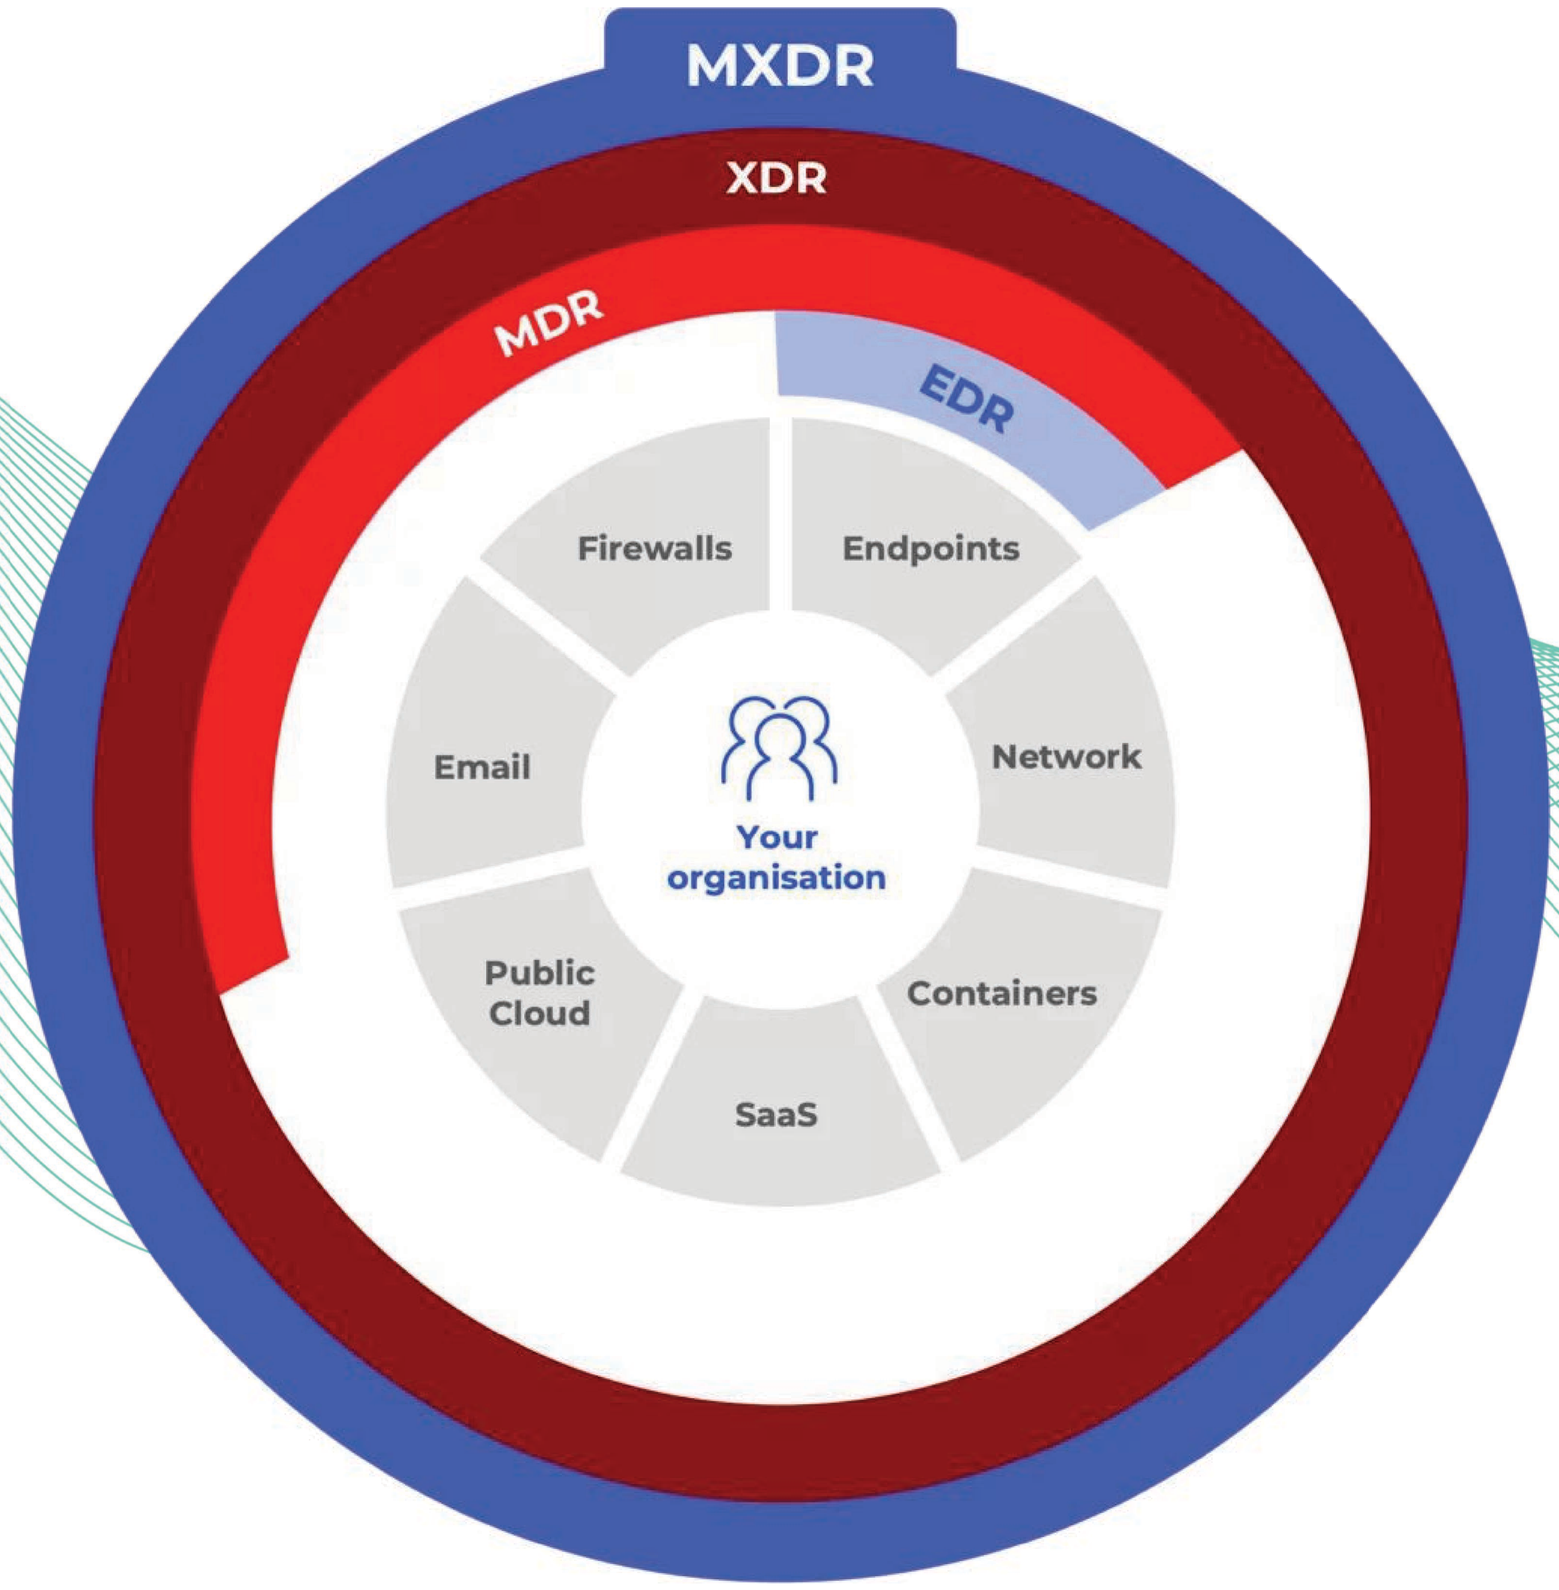
Cisco MXDR from Logicalis

Defending against AI-powered adversaries requires round-the-clock detection, rapid response, and specialist expertise that few security teams can maintain on their own. With advanced threats and skills shortages, more enterprises are turning to managed security services to stay ahead of AI threats and keep internal resources focused.

Cisco MXDR from Logicalis combines Cisco's AI-driven, network-led security platform with Logicalis' deep security expertise and global network of SOCs. Organizations gain 24/7 protection, direct access to experienced security analysts and threat hunters, and a fully managed extension of the in-house security team.

- **Defend against AI-enabled attacks** with AI-driven detection, cross-environment correlation, and automated response playbooks that contain threats in minutes, not days.
- **Secure agentic AI** with end-to-end telemetry and behavioral anomaly detection across agent identities, tool calls, and data access, extended by Logicalis threat hunters who investigate suspicious agent behavior.
- **Protect non-human identities** through identity-focused analytics that surface anomalous credential and access behavior, backed by 24/7 monitoring and Cisco Talos threat intelligence.
- **Expose shadow AI** with network-level visibility that reveals traffic to unsanctioned AI tools and data movement that DLP and governance programs can then act on.
- **Strengthen AI governance** with detailed incident records and compliance-ready reporting that support AI inventory, audit, and regulatory requirements.





Cisco MXDR from Logicalis

The service builds upon Cisco XDR, a cloud-native foundation that unifies threat detection, investigation, and response across network, endpoint, email, cloud, and identity, with integrations spanning Cisco and third-party security tools. Cisco XDR correlates telemetry to prioritize threats by risk, uses agentic AI to verify attacks and build investigation plans, and triggers automated response playbooks to contain threats quickly. Logicalis' SOC analysts add proactive threat hunting and global threat intelligence to investigate anomalous identity behavior, flag unusual data movement such as traffic to unsanctioned AI tools, and produce detailed incident records that support AI governance and compliance reporting.

Contact Logicalis to learn how Cisco MXDR can help your organization stay ahead of AI threats.

About This Report

Logicalis and Cisco created this report to help security and business leaders build a practical defense strategy against AI-powered adversaries. It draws on independent industry research to outline the most pressing AI threats facing enterprises today and the strategies organizations can use to stay ahead of them.

About Logicalis

- First global Cisco partner to launch Cisco XDR as a managed service
- First Cisco XDR Cisco Powered Services Specialization Partner to offer the service globally
- Cisco partnership for 25+ years, spanning networking, security, and managed services
- Global Managed Services Provider operating in 30 territories and serving more than 10,000 clients
- More than 7,000 employees, including specialized SOC teams of analysts, incident responders, threat hunters, and compliance auditors

CISCO PARTNER

Cisco Preferred Cloud and AI Infrastructure Partner
Cisco Preferred Collaboration Partner
Cisco Preferred Networking Partner
Cisco Preferred Security Partner
Cisco Preferred Services Partner
Splunk Partner