SALT

SOLUTION BRIEF • COMPLIANCE

EU AI Act compliance: the agentic security imperative

How Salt Security's Agentic Security Graph delivers the technical controls required for high-risk AI system compliance under Regulation (EU) 2024/1689.

ENFORCEMENT DEADLINE
August 2, 2026

MAXIMUM FINE
€35M or 7% of global turnover

APPLIES TO
Providers and deployers of high-risk AI systems

CONFIDENTIAL

OVERVIEW

The compliance clock is running

Regulation (EU) 2024/1689, the EU AI Act, entered into force on August 1, 2024. The most consequential provisions for enterprise security teams apply as of **August 2, 2026** — the date on which mandatory requirements for high-risk AI systems become fully enforceable.

Penalties are severe. Non-compliant providers and deployers of high-risk AI systems face fines of up to **€30 million or 6% of global annual turnover**, whichever is higher. For violations of prohibited practice provisions (Article 5), that ceiling rises to **€35 million or 7% of global turnover**.

€35M	7%	Aug 2026
MAXIMUM FINE FOR PROHIBITED PRACTICE VIOLATIONS	OF GLOBAL ANNUAL TURNOVER, WHICHEVER IS HIGHER	FULL ENFORCEMENT DATE FOR HIGH-RISK AI PROVISIONS

Who is in scope

The EU AI Act applies to providers placing AI systems on the EU market, deployers operating high-risk AI systems within the EU, and third-country providers and deployers whose AI system outputs are used in the EU. If your organization uses AI systems in hiring, credit decisions, critical infrastructure management, biometric identification, law enforcement, or essential services, you are subject to high-risk compliance requirements.

> **CRITICAL SCOPE NOTE**
>
> If your AI agents invoke APIs — including internal services, third-party platforms, or MCP servers — that action layer is in scope under the Act's cybersecurity and logging mandates. Traditional API security tools and traditional AI security tools each cover only half the picture. Salt Security was built to close that gap.

THE GAP MOST ORGANIZATIONS ARE MISSING

Securing the wrong layer

Most enterprise AI security investment is concentrated at the model layer: responsible AI guardrails, output filtering, LLM red-teaming, and model governance. These efforts are necessary. They are not sufficient for EU AI Act compliance.

Article 15 of the EU AI Act is explicit: protection must extend to the actions an AI system takes, not only the outputs it generates. The threats enumerated in Article 15(5) — data poisoning, adversarial manipulation, confidentiality attacks — target the interfaces through which AI systems interact with the world. In modern AI architectures, those interfaces are APIs.

! **The structural compliance problem:** A traditional API security tool sees API traffic but has no understanding of the AI system behind it. A traditional AI security tool understands the model but has no visibility into the APIs the agent calls. Neither provides the control plane the EU AI Act requires. Salt Security was purpose-built for this gap.

REGULATORY REQUIREMENTS

What the EU AI Act actually requires

The Act establishes a tiered risk framework. Understanding where the hard obligations fall is essential for building a defensible compliance program. Article 5 prohibitions have been in effect since February 2025. The following high-risk AI system mandates become fully enforceable August 2, 2026.

Article-by-article compliance mapping

The table below maps all applicable EU AI Act obligations to Salt Security capabilities. Coverage levels are characterized precisely across three tiers.

ARTICLE	REQUIREMENT	SALT CAPABILITY	HOW SALT SUPPORTS COMPLIANCE	COVERAGE
Art. 9	Risk management system — continuous and iterative throughout AI lifecycle	AG-SPM Continuous Discovery	Automatically discovers all AI agents, MCP server connections, and APIs. Builds a live risk inventory with severity scoring. Flags new connections and configuration drift as emerging risks in real time.	STRONG

ARTICLE	REQUIREMENT	SALT CAPABILITY	HOW SALT SUPPORTS COMPLIANCE	COVERAGE
Art. 10	Data governance — secure handling; prevent unauthorized access and data poisoning at inference time	API Data Flow Visibility	Deep observability into data traversing the action layer at inference time. Detects anomalous data access patterns and malformed API responses that could introduce compromised inputs.	STRONG
Art. 11	Technical documentation — description of all components and interfaces before market placement	Interface Inventory Export	The Agentic Security Graph produces a continuous, exportable inventory of every AI agent, MCP server, and API endpoint in scope — direct input to Article 11 technical documentation.	EVIDENCE
Art. 12	Record keeping — automatic logging; tamper-evident logs retained 6 months (24 months for biometric/law enforcement)	Immutable Audit Trail	Captures a complete, tamper-evident log of every AI-to-API interaction: request payloads, response data, timing, authentication context, and anomaly flags. Retention policies configurable to Article 12 minimums.	STRONG
Art. 13	Transparency — providers must supply deployers with sufficient information on capabilities, limitations, and interfaces	Agentic Security Graph	Generates a continuous, exportable map of every interface the AI system uses. Gives deployers the complete interface inventory and behavioral visibility the Act requires.	STRONG
Art. 14	Human oversight — ability to stop the system and intervene in operations	AG-DR Real-Time Alerting	Surfaces anomalous agent behavior in real time with full context. Enables security operators to see exactly what APIs an agent is calling and to terminate or quarantine agent sessions exhibiting unauthorized behavior.	STRONG
Art. 15(3)	Cybersecurity resilience — technical robustness	Behavioral Threat Protection	Builds AI-powered behavioral baselines for every agent and API interaction. Detects deviations consistent with	STRONG

ARTICLE	REQUIREMENT	SALT CAPABILITY	HOW SALT SUPPORTS COMPLIANCE	COVERAGE
	against adversarial attacks by unauthorized third parties		data poisoning, adversarial input injection, model evasion, and unauthorized data exfiltration.	
Art. 15(5)	Specific threats — protection against data poisoning, adversarial examples, confidentiality attacks, and model evasion	East-West Traffic Analysis	Monitors lateral API traffic between agents, MCP servers, and internal services. Identifies prompt injection patterns in API responses, anomalous agent-to-agent interactions, and unauthorized data access.	STRONG
Art. 17	Quality management system covering data governance, logging, post-market monitoring, and cybersecurity	Compliance Evidence Platform	Provides the cybersecurity monitoring, logging, and posture evidence that a QMS requires. Salt is the technical control layer whose outputs feed the provider's QMS documentation and continuous improvement processes.	PARTIAL
Art. 20	Corrective actions — providers who know or have reason to believe their system presents a risk must act immediately	Threat Detection and Alerting	Salt is the detection mechanism that triggers the Article 20 obligation. Behavioral anomaly detection surfaces incidents that constitute reason to believe a risk exists. The immutable audit trail provides the evidentiary basis for required notification.	STRONG
Art. 26	Deployer obligations — monitor operation, maintain logs, ensure human oversight, and report serious incidents	AG-SPM + AG-DR for Deployers	AG-SPM provides continuous operational monitoring. AG-DR enables the human oversight the deployer is legally responsible for maintaining. The immutable audit trail satisfies the deployer's logging obligation independently of the provider's.	STRONG

ARTICLE	REQUIREMENT	SALT CAPABILITY	HOW SALT SUPPORTS COMPLIANCE	COVERAGE
Art. 72	Post-market monitoring — established, documented, and actively implemented monitoring system from day one of deployment	Continuous Behavioral Monitoring	Salt's continuous monitoring of the action layer is a direct implementation of Article 72 post-market monitoring. Every behavioral baseline, anomaly detection, and posture change is a post-market monitoring data point.	STRONG
Art. 73	Serious incident reporting — 24 hours (life/safety risks), 72 hours (other serious incidents), 15 days (malfunctions)	Incident Detection and Evidence Trail	Salt's behavioral detection identifies incidents as they occur — the precondition for meeting Article 73's timeframe requirements. Organizations cannot report what they cannot detect; Salt closes that gap.	STRONG

TECHNICAL ARCHITECTURE

The Agentic Security Graph: a system of record for Article 15 compliance

The EU AI Act's cybersecurity requirements do not describe a point-in-time assessment. Article 15 calls for AI systems designed and developed to be resilient against the attack categories it enumerates. Demonstrating that resilience to a national supervisory authority requires a system of record — continuous evidence that the action layer is monitored, governed, and protected.

The Agentic Security Graph operates across three layers that directly correspond to the Act's scope of concern:

	WHAT SALT SEES	EU AI ACT RELEVANCE

01 **LLM / model layer**	Model identity, version, and configuration; prompts and completions; token consumption patterns; agent orchestration instructions	Arts. 9, 11, 13 — Risk inventory, technical documentation of system components and interfaces, and transparency evidence for the AI system itself
02 **MCP server layer**	WHAT SALT SEES All MCP server connections; tool definitions exposed to agents; authorization scopes; east-west traffic between agents; prompt injection patterns in server responses	EU AI ACT RELEVANCE Arts. 15(3), 15(5) — Cybersecurity resilience controls and detection of prompt injection, adversarial examples, and unauthorized manipulation via the tool-use layer
03 **API layer**	WHAT SALT SEES All APIs invoked by AI agents; authentication and authorization context per call; data accessed and returned; behavioral baselines; anomaly detection; full request/response logging	EU AI ACT RELEVANCE Arts. 10, 11, 12, 14, 15, 20, 26, 72, 73 — Data governance, interface inventory for technical documentation, immutable logging, human oversight enablement, cybersecurity resilience evidence, corrective action detection triggers, deployer operational monitoring, post-market monitoring implementation, and incident detection enabling Article 73 reporting timelines

COMPLIANCE ROADMAP

The path to August 2, 2026

August 2, 2026 is weeks away. Organizations that began their compliance programs in Q3 2025 — following the preparation timeline below — are now finalizing documentation and entering the enforcement period with a continuous monitoring record behind them. Organizations still building their programs face a compressed timeline: steps that ideally unfolded over the past year must now be executed in parallel and at pace. The roadmap below shows both the recommended preparation timeline and the urgent actions required now.

PHASE	COMPLIANCE ACTIVITIES	SALT SECURITY CONTRIBUTION
Q2–Q3 2025	Classify AI systems against Annex III. Identify which deployments qualify as high-risk. Conduct a gap assessment against Articles 9, 10, 11, 12, 13, 14, 15, 17, 20, 26, 72, and 73.	AG-SPM discovery run generates initial AI agent and API inventory. Identifies undocumented interfaces and shadow AI connections that expand compliance scope.
Q3 2025	Implement technical controls for high-risk systems. Establish logging infrastructure. Begin building 6-month Article 12 log archive. Implement access controls on AI-facing APIs. Begin Article 9 risk management iteration and Article 26 operational monitoring for deployers.	Salt logging infrastructure captures all AI-to-API interactions from deployment. Posture governance enforces least-privilege access policies across the agentic environment.
Q4 2025	Complete technical documentation per Article 11. Conduct tabletop exercises on Article 14 oversight procedures and Article 20 corrective action procedures. Validate logging completeness and tamper-evidence controls. Draft Article 17 quality management system documentation.	Agentic Security Graph export supports Article 11 technical documentation. AG-DR interface validated as the operational Article 14 oversight and Article 20 corrective action mechanism. Salt monitoring data feeds directly into Article 17 QMS documentation.
NOW	Commission internal compliance audit. Remediate gaps urgently. Engage external assessors if conformity assessment requires it. Compile compliance documentation package. Organizations starting now must compress and parallel-track the Q3–Q4 2025 preparation activities.	Salt's continuous monitoring record — however many months accumulated — is the primary Article 15, 72, and 73 compliance evidence. Behavioral threat detection records demonstrate cybersecurity resilience. Organizations deploying Salt now begin building that record immediately.
AUG 2, 2026	High-risk AI Act provisions fully enforceable. Organizations with documented, continuous compliance programs are in the strongest position relative to supervisory authority scrutiny.	Organizations running Salt from Q3 2025 enter the enforcement date with a documented record of continuous monitoring, posture governance, and behavioral threat detection.

MULTI-AGENT ARCHITECTURE NOTE

Recitals 99 and 100 of the EU AI Act address AI systems that interact with other AI systems. In agentic architectures where multiple AI agents collaborate, the compliance boundary extends to every agent in the chain that performs a function classified as high-risk. Salt's east-west traffic monitoring identifies agent-to-agent communication patterns, maps the full chain of API interactions across a multi-agent workflow, and attributes each action to the originating agent — providing the traceability required for both Article 12 logging and Article 14 oversight obligations.

Every month you delay is a month of monitoring data you will not have

The August 2, 2026 enforcement date is fixed. Salt Security's Agentic Security Graph can be deployed in days and begins building your continuous compliance record from day one.

Request a demo

SALT

Regulation (EU) 2024/1689 (EU AI Act) entered into force August 1, 2024. High-risk AI system provisions apply from August 2, 2026. GPAI model provisions applied from August 2, 2025. This document is for informational purposes and does not constitute legal advice. Organizations should consult qualified EU legal counsel regarding their specific compliance obligations.