SALT

CISO RESOURCE • COMPLIANCE

EU AI Act
CISO Checklist

A practical compliance checklist for security leaders preparing for full EU AI Act enforcement. Mapped to every applicable article. Aligned to Salt Security capabilities.

ENFORCEMENT DATE
August 2, 2026

APPLIES TO
High-risk AI system providers & deployers

MAXIMUM PENALTY
€35M or 7% global turnover

CONFIDENTIAL

HOW TO USE THIS CHECKLIST

Your EU AI Act readiness checklist

This checklist maps the mandatory compliance obligations for high-risk AI systems under Regulation (EU) 2024/1689 to concrete security actions. Items marked **urgent** require immediate attention ahead of the August 2, 2026 enforcement date. Items marked **active** represent ongoing operational requirements. Use this as both a readiness assessment and a standing compliance operating procedure.

☐ Urgent — act immediately ☐ Active — ongoing operational requirement ☐ Foundation — complete if not already done

STEP 01 Scope & Classification

- ☐ **Classify all AI systems against Annex III of the EU AI Act.** Identify which deployments qualify as high-risk based on use case, sector, and function. — Annex III
- ☐ **Map every AI agent and its API/MCP server connections.** Any AI system that invokes APIs to take real-world actions is in scope under cybersecurity and logging mandates — regardless of model type. — Art. 15
- ☐ **Determine your role: provider, deployer, or both.** Providers and deployers have overlapping but distinct obligations under Articles 9–17, 26, 72, and 73. — Art. 26
- ☐ **Document third-country applicability.** If your AI system outputs are used within the EU, the Act applies regardless of where your organization is headquartered. — Art. 2

STEP 02 Risk Management System

- ☐ **Establish a continuous, iterative risk management system** covering identification, estimation, evaluation, and mitigation of risks across the AI lifecycle — not a point-in-time assessment. — Art. 9
- ☐ **Maintain a live risk inventory** of all AI agents, MCP server connections, and APIs with severity scoring. Flag new connections and configuration drift as emerging risks in real time. — Art. 9
- ☐ **Document quality management system (QMS) procedures** covering data governance, logging, post-market monitoring, and cybersecurity — with evidence that Salt's monitoring outputs feed the QMS directly. — Art. 17

STEP 03 Data Governance & Record Keeping

- [ ] **Implement deep observability into data traversing the action layer at inference time.** Detect anomalous data access patterns and malformed API responses that could introduce compromised inputs or data poisoning. — Art. 10
- [ ] **Deploy tamper-evident, immutable logging of every AI-to-API interaction** — request payloads, response data, timing, authentication context, and anomaly flags. Logging must have been running for at least 6 months prior to enforcement. — Art. 12
- [ ] **Verify log retention policies meet Article 12 minimums:** 6 months standard; 24 months for biometric or law enforcement AI systems. — Art. 12
- [ ] **Validate that logs are tamper-evident** and can demonstrate provenance to a national supervisory authority. Test integrity controls before enforcement. — Art. 12

STEP 04 Technical Documentation & Transparency

- [ ] **Generate and maintain a complete interface inventory** — every AI agent, MCP server, and API endpoint in scope — as continuous input to Article 11 technical documentation. This must exist before market placement or deployment. — Art. 11
- [ ] **Provide deployers with a complete interface inventory and behavioral visibility map** showing the full capabilities, limitations, and interfaces of the AI system. This is a legal obligation on providers under Article 13. — Art. 13
- [ ] **Keep technical documentation current.** Any new agent deployment, MCP server connection, or API integration that affects the risk profile requires documentation update before the change goes live. — Art. 11

STEP 05 Human Oversight

- [ ] **Implement real-time alerting on anomalous agent behavior** that gives security operators full context — exactly what APIs an agent is calling — and enables termination or quarantine of agent sessions exhibiting unauthorized behavior. — Art. 14
- [ ] **Conduct tabletop exercises on human oversight procedures.** Operators responsible for Art. 14 oversight must be able to demonstrate that they can identify a problem and intervene in a real-world scenario. — Art. 14
- [ ] **Establish and test corrective action procedures.** When behavioral anomaly detection surfaces an incident, the organization must have documented procedures for immediate risk response under Article 20. — Art. 20

STEP 06 Cybersecurity Resilience

- [ ] **Deploy AI-powered behavioral baselines** for every agent and API interaction. Detect deviations consistent with data poisoning, adversarial input injection, model evasion, and unauthorized data exfiltration. — Art. 15(3)
- [ ] **Monitor east-west traffic between agents, MCP servers, and internal services.** Identify prompt injection patterns in API responses, anomalous agent-to-agent interactions, and unauthorized data access across the multi-agent surface. — Art. 15(5)
- [ ] **Map multi-agent architectures for compliance traceability.** Recitals 99 and 100 extend the compliance boundary to every agent in a chain that performs a high-risk function. Salt's east-west monitoring attributes each action to the originating agent. — Recitals 99–100

STEP 07 Deployer Operational Obligations

- [ ] **Implement continuous operational monitoring of deployed AI systems.** Deployers are independently responsible for monitoring — regardless of what the provider does. The immutable audit trail must satisfy the deployer's logging obligation separately. — Art. 26
- [ ] **Maintain post-market monitoring infrastructure from day one of deployment.** Every behavioral baseline, anomaly detection, and posture change is a post-market monitoring data point under Article 72. — Art. 72
- [ ] **Establish incident detection and reporting capability before enforcement.** Organizations cannot report what they cannot detect. Reporting windows are: 24 hours (life/safety risks), 72 hours (other serious incidents), 15 days (malfunctions). — Art. 73
- [ ] **Test Article 73 reporting procedures end-to-end.** Validate that incident detection, evidence collection, internal escalation, and external notification can all occur within the required timeframes. — Art. 73

STEP 08 Compliance Documentation Package

- [ ] **Compile your compliance evidence package.** At minimum: Annex III classification decision, risk management system documentation, interface inventory export, log archive (6+ months), Art. 14 oversight procedures, and Art. 73 incident reporting procedures. — Arts. 9–17
- [ ] **Engage external legal counsel to review compliance documentation** against the specific requirements of Regulation (EU) 2024/1689. This checklist is an operational guide, not legal advice. — General
- [ ] **Determine whether conformity assessment requires external auditor involvement.** For certain high-risk AI system categories, third-party conformity assessment is mandatory before deployment. — Art. 43
- [ ] **Register high-risk AI systems in the EU database** as required for applicable system categories before deployment or market placement. — Art. 71

Salt deploys in days. The clock is running.

Salt Security's Agentic Security Graph closes the compliance gaps this checklist identifies — across LLM, MCP, and API layers. Every day of monitoring is a day of compliance evidence.

Request a demo

SALT