

CSA Agentic Trust Framework Compliance Attestation

Issued by: Akka Technologies, Inc.

Prepared by: Michael Nash, Chief Information Security Officer

Date of Attestation: 20 April 2026

Document Classification: Public

Review Cycle: Annual

Introduction

This attestation is issued by Akka Technologies, Inc. ("Akka") to confirm our compliance posture with respect to the Cloud Security Alliance Agentic Trust Framework ("CSA ATF" or "the Framework"). The CSA ATF, published by the Cloud Security Alliance in 2026, establishes requirements for trust, safety, and governance for agentic AI systems — AI systems capable of autonomous action, multi-step reasoning, and tool use on behalf of human principals.

Akka is a technology company headquartered in the United States, providing the Akka platform — a reactive microservices and distributed systems toolkit specifically designed to enable enterprises to build and operate agentic AI applications at scale. Agentic AI is central to Akka's product and mission, making the CSA ATF directly and comprehensively applicable to Akka's platform, architecture, and operations.

Akka has implemented all 18 CSA ATF controls. Of these, 15 controls are assessed at Low residual risk and 3 at Medium residual risk. No controls are assessed at High residual risk. This attestation documents Akka's implementation posture, key controls, and supporting evidence as of 20 April 2026.

Scope and Applicability

The CSA Agentic Trust Framework addresses the unique trust, safety, and governance challenges posed by AI agents — systems that can plan, reason, use tools, call external services, and take actions autonomously over extended workflows. The Framework applies to organisations that develop, deploy, or provide infrastructure for agentic AI systems.

Akka's scope under the CSA ATF is as a platform provider that:

- Develops and operates agentic AI capabilities embedded in the Akka platform
- Provides infrastructure, orchestration, and runtime environment for customer-built AI agents
- Deploys AI agents in its own internal operations
- Governs a portfolio of AI agents recorded in Akka's centralised Agent Registry

All 18 CSA ATF controls apply to Akka's platform and operations. Akka's compliance covers both its own agentic AI deployment and the platform-level controls it provides to customers to enable their own compliant agentic AI applications.

Compliance Posture

Akka is compliant with the CSA Agentic Trust Framework as of the date of this attestation. Akka's agentic AI platform was designed with trust, containment, and oversight as foundational architectural principles — not retrofitted compliance measures — making CSA ATF alignment a natural expression of Akka's product philosophy.

The 3 Medium-risk controls relate to areas where ongoing evolution of agentic AI capabilities and emerging industry standards continue to shape implementation. All Medium-risk controls have documented treatment plans and are subject to enhanced monitoring.

Agent Identity and Authentication

Akka implements robust identity and authentication controls for AI agents, a cornerstone of the CSA ATF:

Agent Identity Management: Every AI agent deployed on the Akka platform is assigned a unique, verifiable identity, recorded in the centralised Agent Registry. Agent identities are cryptographically bound to their scope and permission grants, preventing identity spoofing or impersonation.

Authentication for Agent Actions: When AI agents call external services, APIs, or internal platform components, they authenticate using service-specific credentials with scoped permissions. Credential management follows least-privilege principles, with agent credentials isolated from human user credentials.

Agent Registry: Akka's Agent Registry provides a comprehensive, auditable record of all AI agents in operation, their identities, assigned owners, scope boundaries, tool access grants, and operational status. The Registry is a foundational governance control enabling systematic management of the agent fleet.

Inter-Agent Authentication: In multi-agent workflows, agents authenticate to each other through platform-managed identity assertions, preventing unauthorised agent-to-agent communication or privilege escalation.

Agent Scope Limitation and Least Privilege

Akka's platform enforces strict scope limitation for AI agents, directly implementing the CSA ATF's containment requirements:

Explicit Permission Grants: AI agents on the Akka platform operate under explicit, narrowly scoped permission grants. Actions, tool calls, and data access are restricted to what is required for the agent's defined purpose. Default permissions are deny-all, with explicit grants required for each capability.

Scope Boundary Enforcement: The platform enforces scope boundaries at runtime, preventing agents from exceeding their authorised action space. Attempts to call unauthorised tools, access data outside the agent's scope, or escalate permissions are blocked and logged.

Tool Access Control: Each agent has a declared set of permitted tools and APIs. Tool access is enforced by the platform runtime, not the agent's own code, ensuring that even a compromised or misbehaving agent cannot access tools it was not explicitly granted.

Minimal Data Access: Agents are provisioned with access to the minimum data necessary to complete their task. Data access is contextual — agents receive data relevant to their current task rather than broad access to data stores, implementing data minimisation at the agentic layer.

Human Oversight of Autonomous Decisions

Akka's platform is built around the principle that meaningful human oversight must be available and enforceable at all levels of agentic AI operation:

Configurable Human-in-the-Loop: The platform provides configurable human-in-the-loop controls at the workflow, task, and action levels. Customers can require human approval before an agent takes actions above defined significance thresholds, ensuring that consequential autonomous decisions are subject to human review.

Override and Halt Mechanisms: Human operators can override, pause, or halt AI agent execution at any point in a workflow. Override actions are logged with timestamp and operator identity, providing an auditable record of human interventions.

Escalation Paths: When an agent encounters a situation outside its defined scope or confidence threshold, configurable escalation paths ensure that the situation is routed to a human operator rather than the agent taking unilateral action.

AI Review Board Oversight: Akka's AI Review Board exercises governance over the entire agent fleet, reviewing new agent deployments, significant scope changes, and escalated incidents. The Review Board ensures that human oversight of autonomous AI decisions is maintained at the organisational level.

Audit Trails for Agent Actions

Comprehensive, tamper-resistant audit trails are a core trust requirement of the CSA ATF, and a primary capability of the Akka platform:

Action Logging: All agent actions — tool calls, API invocations, data accesses, decisions, and outputs — are logged with sufficient detail to reconstruct agent behaviour for audit, investigation, or compliance review purposes.

Audit Trail Integrity: Audit logs are written to isolated, append-only storage with integrity protections, ensuring that agent action records cannot be modified or deleted by agents or application code.

Retention and Accessibility: Audit logs are retained for defined periods in accordance with Akka's data retention policy and applicable regulatory requirements. Logs are accessible to authorised operators for investigation and review.

Audit Log Monitoring: Automated monitoring of agent audit trails identifies anomalous patterns — unexpected action sequences, scope boundary violations, high-frequency tool calls, or unusual data access patterns — and alerts security operations teams for investigation.

Customer Audit Access: Customers can access audit trail data for AI agents running on the Akka platform through defined interfaces, enabling customers to conduct their own audits and satisfy regulatory obligations requiring AI audit trail evidence.

Data Minimisation and Secure Agent Communication

Akka implements the CSA ATF's data governance and communication security requirements:

Data Minimisation by Design: Agent workflows are architected to pass the minimum data required at each step. Agents receive task-specific data contexts rather than broad data access, and intermediate data is cleared from agent memory when no longer needed for the current task.

Secure Communication Channels: All communication between agents, and between agents and external services, is encrypted in transit using current TLS standards. Communication endpoints are authenticated, preventing man-in-the-middle interception or injection of unauthorised instructions.

Prompt Injection Defences: Akka's platform includes defences against prompt injection attacks — attempts by malicious content in agent inputs to redirect agent behaviour contrary to the operator's instructions. Input sanitisation, instruction boundary enforcement, and anomaly detection are applied to agent input processing.

Data Residency and Sovereignty: Customer data processed by AI agents respects configured data residency requirements, ensuring that agent-processed data does not transit regions in violation of customer compliance obligations.

Supporting Evidence

The following evidence supports this attestation:

- Akka ISMS — CSA ATF control framework, 18 controls, all Implemented (15 Low, 3 Medium)
- Agent Registry (internal governance document)
- AI Review Board charter and records
- Platform architecture documentation covering agent identity, scope enforcement, and audit logging
- SOC 2 Type II report (available under NDA), covering security and availability controls for the platform infrastructure
- AWS infrastructure SOC 2 and ISO 27001 documentation
- EU AI Act compliance framework (82 controls, all Implemented)
- NIST AI RMF 1.0 compliance records
- Incident response procedures including AI agent incident classification

Conclusion

Akka's compliance with the CSA Agentic Trust Framework reflects both a regulatory and a product commitment. As the agentic AI platform provider, Akka's commercial success depends on customers trusting that AI agents built and deployed on the platform are safe, contained, auditable, and subject to meaningful human oversight. The CSA ATF's requirements are therefore aligned with Akka's own engineering and governance philosophy.

With all 18 CSA ATF controls implemented and a strong risk posture (15 Low, 3 Medium, 0 High), Akka is well-positioned to serve customers in regulated industries — including financial services, healthcare, and critical infrastructure — who require evidence of rigorous agentic AI governance before adopting AI platform services.

This attestation is available to customers and prospects on request and is subject to annual review. Questions regarding Akka's CSA ATF compliance programme should be directed to:

Michael Nash

Chief Information Security Officer

Akka Technologies, Inc.

michael.nash@akka.io

Signed: Michael Nash, CISO, Akka Technologies, Inc.

Date: 20 April 2026