AKKA

Akka.io | info@akka.io | 1-877-989-7372
580 California St, #1231 San Francisco, CA 94104

NIST AI 600-1 Attestation

Attestation of Alignment with the NIST AI Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1)

Issuer: Akka Date: 2026-04-20 Subject: Alignment with NIST AI 600-1 — Trustworthy and Responsible AI: Generative AI Risk Management

Introduction

Akka develops and operates an agentic AI platform that leverages large language models and generative AI capabilities to enable enterprise customers to build, deploy, and operate AI-powered applications. As the provider of a platform on which generative AI systems are built and run, Akka is directly subject to the risk categories that NIST AI 600-1 addresses: hallucination, prompt injection, data poisoning, harmful content generation, transparency, and accountability throughout the AI lifecycle.

NIST AI 600-1 ("Trustworthy and Responsible AI: Artificial Intelligence Risk Management: Framework: Generative Artificial Intelligence Profile") is a voluntary guidance document published by the National Institute of Standards and Technology. It establishes a risk management profile for organizations developing, deploying, and managing generative AI systems, with particular emphasis on the unique risks introduced by large language models and foundation models.

Akka has implemented all 53 controls defined in its NIST AI 600-1 control set, achieving full implementation coverage across the framework. This attestation summarizes the measures Akka has in place to align with the framework's objectives and demonstrates our commitment to responsible AI governance.

Scope of Compliance

AKKA

Akka.io | info@akka.io | 1-877-989-7372
580 California St, #1231 San Francisco, CA 94104

This attestation covers Akka's alignment with all ten risk domains of NIST AI 600-1, organized below by topic.

1. AI Governance and Organizational Accountability

Akka has established a formal AI governance structure with designated roles and responsibilities for AI risk oversight. An AI Governance Committee provides organizational accountability for decisions affecting the design, deployment, and monitoring of AI systems on the Akka platform. Akka's AI governance policies define the scope of AI activities subject to risk management, document organizational responsibilities, and establish escalation paths for AI-related incidents and ethics concerns.

Akka's AI risk management approach is aligned with the NIST AI Risk Management Framework (AI 100-1) and the EU AI Act, both of which provide foundational governance structures that directly satisfy NIST AI 600-1 governance requirements.

2. Generative AI Risk Assessment and Mitigation

Akka conducts structured AI risk assessments covering the unique risk categories of generative AI, including hallucination, data memorization, model collapse, monoculture risk, and downstream impacts from large foundation model providers. Risk assessments follow the NIST AI RMF MAP, MEASURE, and MANAGE functions and are documented in Akka's AI risk register.

A hierarchy of identified generative AI risks is maintained and reviewed regularly. Mitigation strategies are documented for each identified risk category, with particular attention to risks arising from Akka's role as an AI platform provider whose outputs reach end users through customer applications.

Akka.io | info@akka.io | 1-877-989-7372
580 California St, #1231 San Francisco, CA 94104

Akka maintains alignment with the NIST AI RMF 1.0, ISO/IEC 42001, and EU AI Act risk management requirements, each of which maps to and covers the corresponding NIST AI 600-1 risk assessment and mitigation controls.

3. Data Governance and Training Data Quality

Akka has implemented policies governing training data collection, retention, curation, and quality assurance. These policies address the NIST AI 600-1 requirements for ensuring training data is appropriate, representative, and documented. Data minimization practices are in place to limit the collection and retention of personal information in AI training contexts.

Bias detection and mitigation processes are applied to identify and address harmful biases in AI-generated content, consistent with EU AI Act Article 10 (training data governance) requirements and ISO/IEC 42001 data quality controls. Periodic monitoring of AI-generated content for privacy risks — including detection of personally identifiable information in model outputs and prevention of data memorization — is conducted using approved tooling.

4. Generative AI-Specific Content and Safety Controls

Akka has implemented the following controls specific to the risks of generative AI:

Harmful content prevention: Akka's acceptable use policies prohibit the use of the platform to generate violent, illegal, or harmful content. Controls aligned with EU AI Act prohibited use provisions (Article 5) are in place to prevent misuse of generative AI capabilities for prohibited purposes, including chemical, biological, radiological, and nuclear (CBRN) applications.

Akka.io | info@akka.io | 1-877-989-7372
580 California St, #1231 San Francisco, CA 94104

Content provenance and labelling: Akka supports content provenance mechanisms including watermarking and metadata tagging of AI-generated content, consistent with the EU AI Act's transparency and labelling requirements (Article 29). Controls for identifying AI-generated content, including deep fakes, are implemented in alignment with EU AI Act Article 30 disclosure obligations.

Acceptable use policies: Documented acceptable use policies define the categories of queries and use cases that the platform will and will not support. These policies are communicated to users prior to interactive AI activities, satisfying the transparency disclosure requirements of EU AI Act Article 28.

Bias and fairness: Bias detection and mitigation processes apply across model outputs to identify and address demographic, representational, and allocational biases. Results are reviewed by interdisciplinary teams with diverse backgrounds and capabilities, consistent with ISO/IEC 42001 governance requirements.

5. Transparency, Explainability, and User Awareness

Akka provides transparency into the origin, limitations, and appropriate use of AI-generated content. Documentation of AI system capabilities and limitations is maintained and made available to users. Explainability features and model transparency mechanisms are documented and available on the Akka platform.

User-facing acceptable use policies and instructions for use are published in accordance with EU AI Act Article 13 (transparency) and Article 14 (human oversight) requirements. Users are informed of their ability to provide feedback and seek recourse in relation to AI-generated outputs.

Training and awareness programs covering generative AI risks, responsible use, and organizational policies are provided to relevant Akka personnel.

Akka.io | info@akka.io | 1-877-989-7372
580 California St, #1231 San Francisco, CA 94104

6. Testing, Validation, and Adversarial Robustness

Akka conducts regular testing of its AI systems to evaluate risk-relevant capabilities and verify the robustness of safety measures. Testing activities include:

- **Adversarial testing:** Regular red-teaming and adversarial testing exercises to identify vulnerabilities in generative AI outputs, including susceptibility to prompt injection, jailbreaking, and harmful content generation. This directly addresses the CBRN and offensive cyber misuse testing requirements of NIST AI 600-1.

- **Stress testing:** AI system resilience is tested under adverse conditions, including unusual inputs and edge cases, consistent with EU AI Act Article 9 risk management requirements.

- **Failure mode analysis:** Potential failure modes of generative AI components are systematically identified and documented as part of the AI risk management process.

- **Capability evaluation:** The risk-relevant capabilities of generative AI models integrated with the Akka platform are assessed prior to deployment, with results documented and reviewed.

7. Monitoring, Audit Trails, and Continuous Improvement

Akka maintains comprehensive audit trails of AI system interactions and operational events. Continuous monitoring of AI system behaviour is in place using approved security and observability tooling. Monitoring results feed into Akka's feedback loops and risk management processes, enabling ongoing identification of new impacts and emerging risks from generative AI deployments.

Regular internal audits review AI governance controls for effectiveness and compliance. An ethics review process is conducted by Akka's AI Governance Committee, supported by external perspectives from advisors reflecting diverse viewpoints. User feedback mechanisms

Akka.io | info@akka.io | 1-877-989-7372
580 California St, #1231 San Francisco, CA 94104

are in place to collect and act on reports of problematic AI-generated content, with documented instructions for users to submit feedback and seek recourse.

Document retention policies ensure that test, evaluation, validation, and digital transparency records are maintained for appropriate periods, consistent with ISO/IEC 42001 documentation requirements.

8. Privacy and Data Protection in AI Contexts

Akka conducts privacy impact assessments for AI processing activities involving personal data, aligned with EU GDPR requirements. Controls are in place to detect and prevent the inclusion of personal information in AI-generated outputs and to prevent training data memorization that could expose personal data.

Monitoring of AI-generated content for privacy risks is conducted periodically, with results reviewed and acted upon. Akka's privacy and data protection framework is certified under ISO/IEC 27001 and independently assessed under Akka's SOC 2 Type II audit programme.

9. Security Controls Supporting AI Risk Management

Akka's information security controls — independently assessed under SOC 2 Type II and aligned with ISO/IEC 27001/27002 — provide foundational security assurance that underpins AI-specific risk management:

- Encryption of data in transit and at rest protects AI model inputs, outputs, and training data.

- Access controls restrict access to AI systems and model infrastructure to authorized personnel.

Akka.io | info@akka.io | 1-877-989-7372
580 California St, #1231 San Francisco, CA 94104

- Change management processes govern modifications to AI models and platform components.

- Incident response procedures address AI-specific incidents, including harmful output events and model integrity issues.

- Audit trails provide a complete record of AI system interactions for review and investigation.

10. Third-Party and Supply Chain Risk

Akka manages the AI-related risks introduced by third-party foundation model providers and AI service dependencies. Third-party AI providers are assessed as part of Akka's vendor risk management programme, with particular attention to their own AI governance, bias management, and security practices. Contractual provisions require third-party AI providers to notify Akka of material changes to models and to cooperate in incident investigations.

Supporting Third-Party Evidence

Akka's compliance with NIST AI 600-1 is supported by the following independent assurance:

- **SOC 2 Type II**: Akka holds an annual SOC 2 Type II report covering the Security, Availability, and Confidentiality Trust Service Criteria. The report provides independent evidence of the operational effectiveness of Akka's security controls.

- **ISO/IEC 27001**: Akka is ISO/IEC 27001 certified. The certificate and statement of applicability are available on request.

- **EU AI Act alignment**: Akka maintains a documented alignment with the EU AI Act, which covers a significant portion of the NIST AI 600-1 control set. EU AI Act compliance documentation is available to customers on request.

- **NIST AI RMF alignment**: Akka maintains alignment with the NIST AI Risk Management Framework (AI 100-1), which provides the governance and risk management foundation for NIST AI 600-1 compliance.

Conclusion

Akka has implemented all 53 controls in its NIST AI 600-1 control set, achieving full implementation coverage across the framework's generative AI risk domains. The majority of these controls are satisfied through Akka's existing compliance with higher-priority frameworks including the EU AI Act, NIST AI RMF 1.0, ISO/IEC 42001, ISO/IEC 27001/27002, NIST CSF, and EU GDPR, demonstrating the depth and coherence of Akka's AI governance programme.

This attestation is subject to annual review and is available to customers, prospects, and auditors on request. For further information regarding Akka's NIST AI 600-1 alignment, please contact the Akka security team.

Authorized Signatory: Michael Nash

CISO/Chief of Staff

Akka

580 California Street, #1231

San Francisco, CA 94104

michael.nash@akka.io