WP-020 | DORA and MiCA AI Compliance | June 2026

DORA-Compliant AI Governance: Zero-Knowledge Proof Records for Financial Services ICT Audit

AffixIO Research June 2026 WP-020 Download PDF

Financial firms now run AI inside DORA's ICT perimeter, MiCA's crypto rules, and the EU AI Act at once. Mutable logs satisfy none of them under scrutiny. We show how zero-knowledge proof records give BaFin-ready audit trails without storing customer prompts or model outputs.

CONTENTS

1. The Five-Regime Financial Services Compliance Problem

No sector faces a more complex regulatory intersection than financial services in 2026. A bank operating in the EU that uses AI systems for credit decisions, fraud detection, or customer service simultaneously faces compliance obligations under DORA (ICT risk management), the EU AI Act (high-risk AI system requirements), NIS2 (cybersecurity for critical entities), Schrems II (data transfer restrictions), and, if it offers crypto-asset services, MiCA. Each regime has distinct audit trail, documentation, and governance requirements. None of them is fully aligned with the others.

Financial services AI compliance in 2026 is therefore not a single compliance programme. It is five overlapping compliance programmes with different authorities, different documentation requirements, different timelines for enforcement, and different penalties for non-compliance. The EU AI Act alone imposes penalties of up to 7% of global annual turnover for providers of non-compliant high-risk AI systems; for financial institutions processing billions in transactions, this represents hundreds of millions in potential fines.

The practical challenge for financial services compliance teams is not understanding each individual framework but finding architectural solutions that satisfy multiple frameworks simultaneously. AffixIO's ZK proof-based audit architecture is designed for precisely this challenge: a single cryptographic infrastructure that produces audit records satisfying all five regulatory regimes through a shared technical layer.

---

2. DORA: ICT Governance and AI Systems

The Digital Operational Resilience Act (DORA, Regulation (EU) 2022/2554) has applied to financial entities since 17 January 2026. DORA's scope covers banks, investment firms, insurance undertakings, crypto-asset service providers, and critical ICT third-party providers that serve financial entities. Its requirements are organised around five pillars: ICT risk management, ICT incident reporting and classification, digital operational resilience testing, ICT third-party risk management, and information sharing.

DORA's ICT risk management requirements, primarily in Articles 5 through 16, establish obligations for governance, risk identification, protection, detection, response, and recovery of ICT systems. These obligations apply to all ICT systems used by covered financial entities, including AI systems used for decision-making, customer interaction, or operational functions.

DORA Article 9: Protection and Prevention

DORA Article 9 requires financial entities to implement ICT security policies, procedures, protocols and tools to protect against ICT risks. Specifically, it requires policies for "ICT-related incident management, including detection capabilities and logging arrangements." For AI systems, this includes requirements for audit trails that enable reconstruction of AI-related incidents and demonstrate that the AI system operated within its authorised parameters.

---

3. BaFin's AI-in-DORA Guidance (January 2026)

In January 2026, Germany's Federal Financial Supervisory Authority (BaFin) issued supervisory guidance clarifying how AI systems, particularly generative AI and large language models, must be integrated into DORA-compliant ICT governance frameworks. The guidance applies to all BaFin-supervised institutions but is expected to influence the approach of other national competent authorities across the EU.

BaFin's key conclusions are significant. First, AI systems used in financial services operations constitute ICT systems for the purpose of DORA and are therefore subject to all DORA ICT governance requirements. Second, the use of generative AI or LLMs does not exempt an institution from DORA's logging and audit trail requirements; rather, it creates additional logging obligations because AI system behaviour is more variable and context-dependent than traditional software. Third, model updates and configuration changes to AI systems must be documented and tested under DORA's change management and testing frameworks.

The practical implication of BaFin's guidance is that financial institutions using AI systems cannot treat those systems as exempt from DORA's audit trail requirements. Every AI decision that has a material impact on a customer, a counterparty, or the institution's risk position must be recorded in a way that enables retrospective audit, incident investigation, and regulatory review.

---

4. MiCA and AI in Crypto-Asset Services

The Markets in Crypto-Assets Regulation (MiCA, Regulation (EU) 2023/1114) creates a harmonised EU framework for crypto-asset service providers. Crypto-asset service providers (CASPs) authorised under MiCA are also subject to DORA from January 2026, creating a dual compliance obligation that reflects the intersection of crypto-asset operations with financial services ICT risk management.

AI systems used by CASPs in trading, valuation, risk management, or customer service are subject to both MiCA's record-keeping and transparency requirements and DORA's ICT governance requirements. MiCA Article 72 requires CASPs to maintain records of all services, activities, and transactions for a minimum of five years. Where AI systems participate in these activities, the AI system's decision records must satisfy MiCA's retention requirements.

The combination of MiCA and DORA creates a compliance infrastructure requirement that is difficult to satisfy with conventional logging: records must be retained for five years (MiCA), be tamper-evident (DORA incident reconstruction), support real-time monitoring (DORA detection capabilities), and enable forensic investigation of incidents (DORA incident response). ZK proof-based audit records satisfy all four requirements through a single infrastructure.

---

5. EU AI Act High-Risk Financial Services AI

The EU AI Act's high-risk classification includes AI systems used in credit scoring, access to financial services, insurance underwriting, and recruitment and employment decisions. These are precisely the AI use cases most commonly deployed by financial institutions. From August 2026, providers of such systems must comply with Articles 11 through 15, covering technical documentation, record-keeping, transparency, human oversight, and accuracy.

EU AI Act Article 12 requires high-risk AI systems to have logging capabilities enabling post-market monitoring and retrospective audit. The logging requirements must capture "at a minimum, the period of each use, the reference database against which the input data has been checked, the input data, and the natural or legal persons involved in the verification." For AI systems making

credit decisions, this means recording the model version, the input features (in a privacy-preserving form), the output decision, and the authorisation context for the decision.

AffixIO's ZK proof records satisfy Article 12's logging requirements through a cryptographic mechanism: the proof commits to the model version, the input feature commitment, the output decision, and the authorisation record, without exposing the customer's personal data in the audit trail. This satisfies both Article 12 (logging) and GDPR Article 5 (data minimisation) simultaneously.

---

6. Why Mutable Audit Logs Fail DORA Requirements

DORA's requirement for audit trails is not merely a requirement for records to exist; it is a requirement for records that can support incident investigation, regulatory inspection, and enforcement action. Mutable application logs, the typical mechanism for recording AI system activity, fail this requirement in several respects.

Tamper-Evidence and Non-Repudiation

DORA requires that audit records be sufficient to reconstruct ICT incidents. If an AI system's audit logs can be modified, deleted, or selectively filtered before regulatory inspection, they cannot serve as reliable evidence of what the AI system did. Non-repudiation, the cryptographic property that a party cannot deny having produced a record, is not achievable through mutable logs alone. A ZK proof signed by the AI system's session key at the time of the decision provides non-repudiation: the signed proof cannot be retroactively modified without invalidating the signature.

Incident Reconstruction Gaps

DORA's incident investigation requirements assume that audit records are complete and accurate. In practice, mutable log systems have gaps: logs may not be retained for the full five-year period required by MiCA, log entries may be dropped under high load conditions, and log formats may not be compatible with the forensic tools used during investigation. ZK proof chains, with their hash-

linked structure and append-only anchoring, provide a complete, gap-free record that is structurally incompatible with selective deletion.

---

7. ZK Proof Records as DORA-Compliant Audit Trails

AffixIO's ZK proof records are designed to satisfy DORA's ICT governance and audit trail requirements for AI systems. Each AI decision generates a proof recording five properties without disclosing the underlying data:

- The AI model version used (as a model commitment)
- The input feature commitment (a hash of the input features, not the features themselves)
- The output decision (in a form that allows reconstruction without revealing customer data)
- The policy scope at decision time (the authorised parameters within which the AI was operating)
- A session nonce preventing replay attacks on the audit record

These five properties, recorded in a ZK proof signed by the AI system's session key, satisfy DORA's logging requirements for AI decisions without exposing customer personal data to the audit record. The audit record can be inspected by regulators, reconstructed by incident investigators, and retained for five years or more without creating a data protection liability.

---

8. Mapping ZK Proofs to DORA Article Requirements

DORA Article	Requirement	ZK Proof Mechanism
Article 9(4)(d)	Logging arrangements for ICT systems	Per-decision ZK proofs with session anchoring

DORA Article	Requirement	ZK Proof Mechanism
Article 10(2)	Detection of anomalous activities	Proof validation failures flag out-of-policy AI behaviour in real time
Article 11(1)	ICT incident response and recovery	Proof chain enables precise incident localisation to specific decisions
Article 13(6)	Digital operational resilience testing evidence	Test run proofs provide cryptographic evidence of testing execution
Article 17(1)	ICT-related incident classification and reporting	Proof chain provides the forensic record needed for incident reporting
Article 28(3)	Third-party ICT service monitoring	Third-party AI services generate proofs using AffixIO API; financial entity holds verification keys

9. Critical ICT Third-Party Provider Governance

In November 2025, the European Supervisory Authorities (ESAs) published their first list of 19 designated Critical ICT Third-Party Providers (CTPPs), including Amazon Web Services, Google Cloud, Microsoft, Oracle, SAP, and Deutsche Telekom. These providers are now subject to direct EU oversight, including annual risk assessments, on-site inspections, and mandatory reporting.

For financial institutions that rely on cloud-based AI services from CTPPs, DORA creates a third-party governance obligation: the institution must monitor the AI services it uses from CTPPs and maintain audit records of their use. This is practically challenging because the financial institution does not have direct access to the CTPP's infrastructure and cannot generate audit records of AI decisions made by the CTPP's systems.

AffixIO's architecture addresses this through a verifiable audit API: the CTPP runs the AffixIO proof generation layer alongside its AI inference service, generating per-decision proofs that are returned to the financial institution alongside the AI decision. The financial institution holds the verification key and can independently verify each proof without requiring access to the CTPP's

infrastructure. This satisfies DORA's third-party monitoring requirement without requiring the CTPP to disclose its model architecture or the financial institution to exercise intrusive access rights.

---

10. Incident Reporting and Proof-Based Evidence

DORA requires financial entities to report significant ICT-related incidents to their national competent authority. The incident report must include a description of the incident, its impact, the measures taken to address it, and supporting evidence. For AI-related incidents, the supporting evidence must demonstrate what the AI system did, when, under what configuration, and whether its behaviour was within authorised parameters.

Conventional application logs provide the "what" and "when" of AI incidents but often cannot provide the "under what configuration" or "whether within authorised parameters" with sufficient reliability for regulatory use. AffixIO's proof chain provides all four elements for each AI decision in the incident timeline: the model commitment establishes the configuration, the policy scope proof establishes the authorised parameters, and the proof chain's integrity properties ensure that the record has not been modified since the incident occurred.

Incident reports supported by ZK proof evidence are substantially more defensible in regulatory proceedings than reports supported by application logs alone. The cryptographic non-repudiation of signed proofs means that a regulated entity can demonstrate with certainty what its AI system did during an incident, rather than relying on potentially contestable log records.

---

11. Implementation for Financial Services AI

Implementing AffixIO's ZK proof audit architecture in a financial services AI context requires integration at three layers: the AI inference layer, the decision record layer, and the compliance reporting layer.

AI Inference Layer Integration

The AffixIO attestation companion integrates with existing AI inference infrastructure as a middleware component. For each AI decision, the sidecar receives the model commitment, the decision output, and the authorisation context from the inference service, generates a ZK proof, and appends it to the session proof chain. The integration is compatible with all major inference platforms including Azure OpenAI Service, AWS Bedrock, and Google Vertex AI.

Decision Record Layer

The proof chain is stored in a tamper-evident record store, using hash-linked proof entries anchored to an append-only log. Records are retained for the full MiCA five-year period. The record format is compatible with CycloneDX audit record schemas, enabling integration with existing ICT governance tooling.

Compliance Reporting Layer

AffixIO provides a compliance reporting interface that maps proof chain data to DORA's required logging fields, EU AI Act Article 12 record-keeping fields, and MiCA Article 72 transaction record fields. Regulatory inspection requests can be fulfilled through the reporting interface without requiring inspectors to access live AI system infrastructure.

---

12. Conclusions

Financial services AI governance in 2026 is defined by a convergence of five regulatory regimes, each with distinct but overlapping requirements. DORA mandates ICT governance and audit trails for all ICT systems including AI. BaFin has confirmed that LLMs and generative AI are fully within DORA's scope. The EU AI Act imposes additional logging, transparency, and human oversight requirements for high-risk financial services AI. MiCA extends these obligations to crypto-asset service providers. NIS2 adds cybersecurity controls requirements.

The common thread across all five regimes is the requirement for reliable, tamper-evident, retrospectively verifiable records of AI system behaviour. AffixIO's ZK proof-based audit records satisfy this requirement through a single cryptographic infrastructure that produces non-repudiable,

privacy-preserving, independently verifiable evidence of AI decisions, configurations, and policy compliance. By grounding audit records in cryptography rather than application logging, AffixIO provides financial services AI governance that is structurally resistant to the most common causes of regulatory audit failure: mutable logs, incomplete retention, and insufficient forensic detail.

Related reading

- WP-007: EU AI Act and NIS2 Compliance in One Architecture
- WP-003: The Proof-Not-Log Paradigm for AI Audit Trails
- WP-006: PII-Free KYC by Design with Zero-Knowledge Identity Circuits

---

Frequently asked questions

Does DORA apply to LLMs?

BaFin's January 2026 guidance treats generative AI and LLM deployments as in-scope ICT systems requiring governance and audit evidence.

Can ZK proofs replace traditional ICT logs?

They complement logs by providing tamper-evident, independently verifiable records of AI policy evaluation without retaining sensitive transaction data.

How does this relate to MiCA?

Crypto-asset service providers using AI for risk scoring or customer communication need the same verifiable governance artefacts as traditional banks.

All White Papers

- About
- Solutions
- Legal
- Trust & Security

Contact

truth layer | yes | no | proof