Sovereign AI Horizontal Memory (SAIHM) — Regulatory & Compliance Framework

Priority 1 — Regulators, Auditors & Compliance Officers

SAIHM

April 2026

Legal Notice

This document is provided for regulatory review, compliance assessment, and auditing purposes. It describes the Sovereign AI Horizontal Memory (SAIHM) protocol's compliance architecture, jurisdictional framework, and data protection mechanisms.

---

1. Executive Summary

1.1 What Is SAIHM?

Sovereign AI Horizontal Memory (SAIHM) is a decentralized, privacy-preserving memory infrastructure protocol that enables artificial intelligence agents to maintain persistent, encrypted, cross-session memory. Built on the COTI V2 blockchain's native Garbled Circuit (GC) execution environment, SAIHM provides:

- **Data Sovereignty:** Each AI agent's memory is encrypted with agent-specific keys and stored across a multi-tier decentralized storage network. No single operator can access plaintext memory contents.
- **Regulatory Compliance by Design:** GDPR erasure (Article 17), EU AI Act transparency, MiCA reporting, and cross-jurisdictional data transfer controls are embedded as protocol-level invariants — not optional features.
- **Post-Quantum Cryptographic Security:** All internal cryptographic operations use NIST-standardized post-quantum algorithms (FIPS 203, 204, 205), ensuring long-term data protection against quantum computing threats.
- **Auditable Operations:** A tamper-evident, hash-chained audit ledger records all protocol operations, stored on Filecoin with Storj mirroring and Arweave checkpoint anchoring.

1.2 Intended Audience

This document is designed for:

- Financial and technology regulators (EU, US, UK, APAC)
- Data protection authorities (DPAs)
- Compliance officers evaluating SAIHM integration
- External auditors conducting protocol assessments
- Legal teams reviewing AI agent data handling

1.3 Regulatory Scope

SAIHM's compliance architecture addresses the following regulatory frameworks:

Regulation	Jurisdiction	SAIHM Coverage
General Data Protection Regulation (GDPR)	EU/EEA	Articles 17, 22, 44–49; ISO 27701 PIMS
EU AI Act	EU	Articles 14, 50(2), 22; Omnibus amendments
Markets in Crypto-Assets (MiCA)	EU	Quarterly reporting; travel rule
NIST Post-Quantum Standards	US (global adoption)	FIPS 203, 204, 205
ISO 27001 / ISO 27701	International	Annual review cycle; control change monitoring
ETSI EN 319 401	EU	Cryptographic key retention policy
NIST SP 800-207	US (global adoption)	Zero Trust architecture model
FinCEN Travel Rule	US	$3,000 threshold compliance

2. Architecture Overview for Regulators

2.1 High-Level Architecture

SAIHM operates as a set of sealed Garbled Circuits (GCs) executing on the COTI V2 blockchain. The architecture is organized into functional layers:

Core Layer: Protocol orchestration, cryptographic key management, and state coordination.

Storage Layer: Multi-tier decentralized storage across Filecoin (warm/archival with Proof of Data Possession), Storj (encrypted object storage), Arweave (permanent immutable anchoring via ANS-104 DataItems), and IPFS (cooperative pinning pool with delegated routing).

Identity & Session Layer: Agent identity management, session token issuance with scope-limited access controls, and behavioral reputation scoring.

Economics Layer: Protocol fee computation, congestion pricing, and incentive mechanisms — all operating within sealed garbled circuits with no external smart contract dependencies.

Governance & Compliance Layer: On-chain governance voting, regulatory reporting, and tamper-evident audit ledger maintenance.

Semantic Layer: AI-native memory operations including semantic search, embedding validation, and knowledge graph construction — all performed under encryption.

2.2 Data Flow Summary

1. An AI agent establishes a session via ML-KEM (FIPS 203) key encapsulation.
2. Memory write operations encrypt data using per-shard Data Encryption Keys (DEKs) wrapped in Key Encryption Keys (KEKs).
3. Encrypted memory shards are distributed across storage tiers based on salience, priority, and jurisdictional requirements.
4. Read operations authenticate via session tokens with scope-limited permissions.
5. All operations are recorded in the hash-chained audit ledger.
6. Protocol fees are settled in COTI (nCOTI denomination).

2.3 No Plaintext Exposure

At no point in the SAIHM data flow does any operator, storage provider, validator, or protocol component have access to plaintext agent memory contents. All computation occurs within sealed Garbled Circuits where inputs and outputs are encrypted.

---

3. GDPR Compliance

3.1 Article 17 — Right to Erasure

SAIHM implements GDPR Article 17 (Right to Erasure) through cryptographic erasure — the destruction of Data Encryption Keys (DEKs) rather than the physical deletion of encrypted data across distributed storage nodes.

3.1.1 Erasure Mechanism

When an erasure request is received:

1. **DEK Destruction (GC-3):** The per-shard DEK is cryptographically destroyed within the sealed GC-3 circuit. Without the DEK, the encrypted shard data becomes computationally irrecoverable.
2. **Arweave Anchor:** A `MPS-DEK-DESTRUCTION-v1` Blake3-domain anchor

DataItem is written to Arweave, creating an immutable, timestamped proof of erasure.
3. **IPFS Unpin Sequence:** Cooperative pinning pool nodes are instructed to unpin the shard CID within `ipfs_erasure_unpin_attempt_window_epochs` (default: 720 epochs / 30 days).
4. **Cross-Tier Confirmation:** Erasure completion is confirmed across all storage tiers within `gdpr_erasure_cross_tier_confirm_latency_advisory_epochs` (default: 168 epochs / 7 days).
5. **Audit Recording:** The complete erasure event is recorded in the Private Audit Ledger.

3.1.2 Erasure Guarantees

Metric	Target	Monitoring
Erasure completion (P95)	≤ 672 epochs (28 days)	`gdpr_art17_erasure_p95_latency_advisory_epo`
DEK destruction anchor SLA	≤ 72 epochs (3 days)	`gdpr_erasure_dek_destruction_anchor_sla_epo`
Cross-tier confirmation	≤ 168 epochs (7 days)	`gdpr_erasure_cross_tier_confirm_latency_adv`
Backlog critical threshold	≤ 10 pending requests	`gdpr_art17_erasure_backlog_critical_count`

3.1.3 Erasure During Conservation Mode

GDPR Article 17 erasure requests are **always processed**, even when the protocol is in conservation mode (degraded operational state). This is a protocol invariant — erasure is classified as PERMITTED during conservation mode.

3.2 Articles 44–49 — Cross-Border Data Transfers

SAIHM's cross-chain data path (GC-19) implements jurisdictional controls for cross-border data transfers:

- **Chain Registry:** Each supported blockchain is registered with jurisdictional metadata. Data path activation requires explicit verification of transfer mechanism adequacy.
- **Adequacy Decision Tracking:** `gdpr_art45_adequacy_review_interval_epochs` (default: 35,040 epochs / ~4 years) ensures periodic review of third-country adequacy decisions.
- **Advisory on Change:** `gdpr_art45_adequacy_decision_change_advisory` fires when an adequacy decision status changes, alerting compliance officers.

3.3 Article 22 — Automated Decision-Making

SAIHM does not make autonomous decisions about individuals. It provides memory infrastructure to AI agents. However, the protocol supports explainability through:

- Sealed computation disclosure mechanisms (audit-accessible upon lawful request)
- Session scope transparency (agents can only access memory within their authorized scope)
- Behavioral reputation scores (PRS) operate on behavioral signals only — no personal data profiling

3.4 ISO 27701 PIMS Integration

SAIHM's Privacy Information Management System (PIMS) alignment is verified through:

- Annual ISO 27701 review anchored to Arweave (`MPS-ISO27701-PIMS-REVIEW-v1`)
- `iso27701_review_max_age_days` = 365 (maximum interval between reviews)
- Control change monitoring with regression detection

---

4. EU AI Act Compliance

4.1 Applicability Assessment

SAIHM is an **infrastructure protocol** — it does not itself constitute an AI system under the EU AI Act definition. It provides memory services to AI agents operated by third parties. However, SAIHM proactively addresses relevant obligations:

4.1.1 Article 14 — Human Oversight

The Self-Improvement Proposal Engine (SIPE) — SAIHM's autonomous parameter optimization system — is explicitly constrained:

- **SIPE cannot self-apply:** All SIPE proposals require governance vote approval (20% quorum + 66% supermajority). No autonomous parameter change can take effect without human-in-the-loop governance.
- **Governance voting window:** 14 days (fixed, not governance-adjustable)
- **Pre-screening:** GC-14 pre-screens all proposals for schema validity and safety constraints before they reach voters.

4.1.2 Article 50(2) — Watermarking

SAIHM has been assessed as **not applicable** for Article 50(2) watermarking requirements. SAIHM does not generate synthetic AI content — it stores and retrieves encrypted memory shards. This assessment is tracked as `eu_ai_act_art50_watermarking_mps_applicability_confirmed = not_applicable`.

4.1.3 Article 22 — Explainability

Sealed computation disclosure mechanisms allow authorized auditors to verify that protocol operations conform to their specified behavior, without exposing agent memory contents or cryptographic key material.

4.2 Omnibus Amendment Tracking

SAIHM actively tracks the EU AI Act Omnibus amendment process:

Milestone	Status	Date
Parliament plenary passage	Confirmed (569 votes)	March 26, 2026
Trilogue commencement	Confirmed	March 26, 2026
Current status	Trilogue active	As of April 2026
August 2, 2026 deadline	Legally operative	Regardless of Omnibus

The protocol maintains an advisory guard (`eu_ai_act_aug2026_deadline_proximity_advisory`) that fires as the August 2, 2026 compliance deadline approaches, ensuring timely compliance actions regardless of Omnibus outcome.

4.3 Proposed Omnibus Timeline Extensions

The Omnibus proposal includes extensions to certain annexes: Annex III to December 2027, Annex I to August 2028, and watermarking to November 2, 2026. These extensions are **not yet agreed** and SAIHM maintains compliance readiness for the original deadlines.

---

5. MiCA Compliance

5.1 Quarterly Reporting

SAIHM generates MiCA-compliant quarterly reports with the following guarantees:

Parameter	Value	Descript
		Maximun days afte

`mica_quarterly_report_max_lag_days`	45	quarter-e for repor submissi
`mica_report_data_freshness_required_epochs`	720	Data mus be ≤ 30 days old report tir
`mica_event_classification_coverage_advisory_floor`	0.99	≥ 99% of events m be classif
`annual_snapshot_mandatory_field_count`	14	Fixed required fields per annual snapshot

Reports are anchored to Arweave using `MPS-MICA-REPORT-ANCHOR-v2` domain for immutable provenance.

5.2 Travel Rule Compliance

SAIHM implements FinCEN Travel Rule coverage:

- **Threshold:** $3,000 (`fincen_travel_rule_threshold_usd`)
- **Coverage floor:** 100% (`fincen_travel_rule_coverage_advisory_floor = 1.00`)
- Cross-chain fee settlements include jurisdictional metadata for travel rule compliance

---

6. Cryptographic Standards & Post-Quantum Readiness

6.1 NIST Post-Quantum Cryptography

SAIHM exclusively uses NIST-standardized post-quantum algorithms for all internal cryptographic operations:

Standard	Algorithm	SAIHM Usage
FIPS 203	ML-KEM (Kyber)	Session key encapsulation
FIPS 204	ML-DSA (Dilithium)	Digital signatures (all protocol signing)
FIPS 205	SLH-DSA (SPHINCS+)	Hash-based signature fallback; annual self-test

6.1.1 PQC Invariant

All MPS-internal asymmetric cryptographic operations MUST use post-quantum algorithms. The only exception is the outer RSA-4096 signature on ANS-104 DataItems, which exists solely for Arweave L1 network compatibility and is not an MPS security invariant. Each DataItem carries an additional inner ML-DSA provenance signature for MPS-internal verification.

6.2 Key Hierarchy & Retention

- **Key Retention Policy:** Current + 2 prior keys per agent (ETSI EN 319 401 compliant)
- **Key Derivation:** HKDF with domain-separated derivation strings (118 active domains)
- **Entropy:** drand beacon (League of Entropy mainnet) + sealed CSPRNG, reseeded after every key generation event
- **Key Rotation:** Governance-triggered for protocol-level keys; automated lifecycle management for per-agent keys

6.3 HKDF Domain Separation

All key derivations use explicit domain separation strings to prevent cross-context key reuse. The HKDF Domain String Registry contains 118 active entries (of 192 maximum capacity), organized by function:

- Cryptographic primitives (signing, encryption, encapsulation)
- Storage operations (per-provider key derivation)
- Session and identity management
- Governance and compliance anchoring
- Audit and provenance

6.4 Hash Functions

SAIHM uses Blake3 for all non-HKDF hashing operations, with domain-separated derivation across 33 registered Blake3 domain strings covering Merkle trees, content hashing, governance anchoring, and erasure verification.

---

7. Zero Trust Architecture (NIST SP 800-207)

7.1 Design Principles

SAIHM implements Zero Trust as a protocol invariant (§1.3 of the architecture specification):

1. **No implicit trust:** Every operation requires cryptographic authentication. Session tokens are scope-limited and time-bounded.
2. **Least privilege:** Agents can only access memory shards explicitly included in their session scope.
3. **Continuous verification:** Behavioral reputation (PRS) is

continuously updated based on agent behavior within the protocol.
4. **Microsegmentation:** Each garbled circuit operates as an isolated sealed computation environment. Inter-GC communication follows strict schema validation.

7.2 Session Security

Control	Implementation
Authentication	ML-KEM session key encapsulation (FIPS 203)
Authorization	Scope-limited session tokens with per-shard permissions
Token expiry	Epoch-bounded; configurable per session
Replay prevention	Nonce-based with HKDF derivation; cross-chain replay detection
Scope integrity	Blake3 hash of sorted scope list; Jaccard overlap detection

7.3 Behavioral Reputation System (PRS)

The Protocol Reputation Score (PRS) provides continuous trust assessment:

- **Range:** 0–10,000 (fixed-point)
- **Initial score:** 7,500
- **Levels:** HEALTHY (7,500–10,000), WARNING (5,000–7,499), DANGER (4,501–4,999), CRITICAL (1–4,500), SUSPENDED (0)
- **Recovery:** Passive (10 pts/epoch) + Active (BFSI score × 50 pts/epoch)
- **Minimum threshold spacing:** 500 points between levels (governance-immutable invariant)

Violations trigger sealed decrements — the specific violation types and decrement values are defined as sealed protocol constants, ensuring consistent and predictable enforcement.

---

8. Audit & Accountability

8.1 Private Audit Ledger

All protocol operations are recorded in a tamper-evident, hash-chained audit ledger:

- **Storage:** Filecoin (primary), Storj (mirror), Arweave (checkpoints)
- **Integrity:** Each entry contains `prior_hash` linking to the previous entry. Chain continuity violation triggers `audit_ledger_chain_integrity_failure_critical` CRITICAL and protocol halt.
- **Signing:** ML-DSA signatures by GC-14 audit signing key

- **Schema version:** Tracked and verified; incompatible versions rejected
- **Append-only:** No entry deletion permitted

8.1.1 Audit Entry Fields

Each audit ledger entry records:

- Entry ID and hash chain linkage
- Epoch and wall clock timestamp
- Originating GC component
- Event identifier and severity tier
- Agent identity hash (where applicable, pseudonymized)
- Shard identifier (where applicable)
- Fee information (where applicable)
- PRS score after event (where applicable)
- GDPR erasure reference (where applicable)
- Payload hash for event-specific data

8.2 Protocol Liveness Beacon

SAIHM emits a Liveness Beacon every epoch (1 hour) containing:

- Protocol Health Index (PHI) score and level
- Active signal count and prioritized signal list
- Conservation mode status
- SIPE submission status
- Cross-chain data path status

The beacon provides continuous proof-of-liveness and system health attestation for external monitoring.

8.3 Governance Transparency

All governance proposals, votes, and outcomes are anchored to Arweave with ML-DSA signatures. Governance voting uses private token-weighted votes via garbled circuits, ensuring vote privacy while maintaining outcome transparency.

---

9. Conservation Mode & Emergency Operations

9.1 Conservation Mode

When protocol health (PHI) drops below the critical threshold, SAIHM enters Conservation Mode — a controlled degradation state:

Operation	Status in Conservation Mode
Shard writes	SUSPENDED
Shard reads	PERMITTED
New session issuance	SUSPENDED

Session renewal	PERMITTED
Governance votes	PERMITTED
GDPR Art. 17 erasure	PERMITTED
Fee collection	UNCONDITIONAL
Audit ledger writes	PERMITTED
Arweave anchoring	PERMITTED
PHI monitoring	ACTIVE
Liveness Beacon	ACTIVE
PRS updates	ACTIVE

9.1.1 Exit Process

Conservation mode exit requires:

1. PHI above critical threshold for `conservation_mode_exit_clear_cycles` consecutive cycles (default: 3)
2. Governance vote with 20% quorum and 67% supermajority
3. `conservation_mode_exit` proposal type (exempt from GC-14 pre-screening)

9.2 Fork Incompatibility Response

SAIHM includes a structured response to upstream blockchain fork incompatibilities:

- Detection within `fork_incompatibility_response_sla_epochs` (default: 6 epochs / 6 hours)
- Migration anchor written to Arweave for audit trail
- Categorized incompatibility types (bytecode hash mismatch, cryptographic primitive removal, consensus break, etc.)
- Resolution advisory within `fork_incompatibility_resolution_advisory_epochs` (default: 26,280 epochs / ~3 years)

---

10. Multi-Jurisdictional Data Handling

10.1 Storage Provider Diversification

SAIHM distributes encrypted data across multiple storage providers in different jurisdictions:

Provider	Jurisdiction	Role
Filecoin	Decentralized (multi-jurisdictional)	Primary warm storage; PDP proofs
Storj	Distributed (multi-jurisdictional)	Encrypted object storage; audit mirror
Arweave	Decentralized	Immutable anchoring;

	(permanent)	provenance
IPFS	Decentralized (multi-jurisdictional)	Cooperative pinning; delegated routing

This multi-provider architecture ensures:

- No single jurisdiction can compel access to complete agent memory
- Erasure compliance operates across all tiers (DEK destruction renders data irrecoverable regardless of storage jurisdiction)
- Storage provider concentration monitoring prevents jurisdictional over-dependence

10.2 Cross-Chain Identity

SAIHM supports cross-chain agent identity via GC-19, with:

- Canonical cross-chain MPS agent identity (Blake3 derivation)
- Home chain identification in session tokens
- Cross-chain replay detection with configurable observation windows
- Nonce sequence gap monitoring for anomaly detection

---

11. Protocol Health Index (PHI)

11.1 Composite Health Score

PHI provides a single [0,1] composite health score computed from weighted domain signals:

Domain	Weight	Coverage
Storage	0.17	Filecoin, Storj, Arweave, IPFS health
Identity	0.13	Session management, agent identity
Governance	0.13	Proposal processing, vote integrity
Security	0.17	PRS, cryptographic operations
Economics	0.13	Fee collection, reserve health
Multi-chain	0.12	Cross-chain path availability
Semantic	0.08	Semantic search, knowledge graph health
Embedding	0.07	Embedding validation, LSH performance

11.2 Health Levels

- **HEALTHY:** PHI ≥ 0.95
- **ADVISORY:** 0.80 ≤ PHI < 0.95
- **CRITICAL:** PHI < 0.80 → triggers Conservation Mode consideration

PHI is emitted every epoch in the Liveness Beacon, providing continuous health attestation.

---

12. Compliance Monitoring Parameters

12.1 Key Compliance Parameters

The following governance-adjustable parameters are directly relevant to regulatory compliance:

Parameter	Default	Purpo
`gdpr_art17_erasure_completion_advisory_days`	28	Target GDPR erasure complet
`gdpr_art17_erasure_backlog_critical_count`	10	Maximu pending erasure requests
`mica_quarterly_report_max_lag_days`	45	Maximu MiCA re delay
`iso27701_review_max_age_days`	365	Maximu ISO 277 review interval
`fincen_travel_rule_threshold_usd`	3,000	Travel r notificat threshol
`gdpr_art45_adequacy_review_interval_epochs`	35,040	Adequac decision review period
`audit_ledger_chain_integrity_recovery_sla_epochs`	12	Audit ch recovery deadline
`governance_vote_window_days`	14	Governa vote per (FIXED)
`governance_quorum_pct`	20	Minimu vote particip (FIXED)
`governance_supermajority_pct`	66	Minimu approva threshol (FIXED)

12.2 Event Monitoring for Compliance

Compliance-relevant events are classified into four severity tiers:

- **CRITICAL:** Requires immediate response (protocol may halt). Examples: audit chain integrity failure, fork incompatibility

detected.
- **ADVISORY:** Requires attention within SLA window. Examples: erasure backlog approaching threshold, adequacy decision change.
- **INFORMATIONAL:** Recorded for audit trail. Examples: governance vote completed, report submitted.
- **HEARTBEAT:** Routine operational signals.

All events are recorded in the audit ledger and emitted via the Liveness Beacon.

---

13. Contact & Further Information

13.1 Regulatory Inquiries

For regulatory inquiries regarding SAIHM compliance:

- **Protocol Documentation:** Available at the SAIHM project portal
- **Audit Access:** Liveness Beacon data is publicly accessible on Arweave
- **Governance Participation:** On mainnet, open to all gCOTI token holders via the on-chain governance system with private garbled-circuit voting

13.2 Standards References

Standard	Reference
GDPR	Regulation (EU) 2016/679
EU AI Act	Regulation (EU) 2024/1689
MiCA	Regulation (EU) 2023/1114
FIPS 203	ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism)
FIPS 204	ML-DSA (Module-Lattice-Based Digital Signature Algorithm)
FIPS 205	SLH-DSA (Stateless Hash-Based Digital Signature Algorithm)
ISO 27001	Information Security Management Systems
ISO 27701	Privacy Information Management
ETSI EN 319 401	Electronic Signatures — Trust Service Providers
NIST SP 800-207	Zero Trust Architecture
