PRASAGA FOUNDATION · PUBLIC BRIEFING · JUNE 2026

SagaChain and SagaAIoT™

Why the Structural Difference Matters

A plain-language guide for executives, partners, and decision-makers

THE PROBLEM

Every Enterprise Faces the Same Governance Gap

For decades the answer has been **log files**. Systems record what happened, and auditors reconstruct a story from those records.

This fails when AI systems make decisions at machine speed on physical infrastructure, when multiple parties need to agree on what happened, and when regulators require proof that a policy was
actually enforced — not just that a log claims it was.

Speed
Thousands of AI decisions per second

Multi-party
No shared verifiable source of truth

Regulatory
Policy enforcement must be structural

Physical
Device commands need pre-execution auth

The core issue is not that organizations lack governance policies. They have them. The issue is that those policies live in documents and databases — not in the decisions themselves. The action and the authorization are never structurally inseparable.

CURRENT LANDSCAPE

What Existing Technologies Actually Deliver

Legacy / Centralized Systems

- ⚠ Governance lives in policy documents
- ⚠ Audit is log files — inherently disputable
- ⚠ Insiders can alter or selectively disclose records
- ⚠ Not designed for AI agent authorization
- ⚠ Authorization separated from execution

Conventional Blockchains

- ⚠ Strong for financial asset settlement
- ⚠ Smart contracts add programmability
- ⚠ No native meta-class object model — types are compiler conventions, not chain state
- ⚠ Business context not natively recorded
- ⚠ Impractical throughput for IoT telemetry

SagaChain + SagaAIoT™

- ✓ Open, permissionless, trustless Layer-1
- ✓ Patented object model — native hierarchies
- ✓ Authorization IS the governance object
- ✓ AI agent constraints are structural
- ✓ Device payloads off-ledger; policy on-chain

The Structural Difference:

A Patented Object Model at the Core of a Layer-1 Blockchain

SagaChain is an open, permissionless, trustless, decentralized Layer-1 blockchain.

What makes it categorically different is what sits at its core: a patented data and object model.

That model is what enables purpose-built platforms — like SagaAIoT™ — to deliver capabilities not previously possible on any general-purpose blockchain.

EVM smart contracts

No native type system on-chain. Solidity structs are compiler conventions that flatten to storage slots. Cross-contract references are raw addresses. Type definitions live in off-chain ABIs and must be redeployed to evolve.

SagaChain meta-class model

Classes, inheritance hierarchies, field schemas, and method dispatch are registered as on-chain state via SagaRegisterClasses. Objects reference each other by typed LOID — not by address lookup. The type system itself is chain state: composable, extensible, queryable.

SagaAIoT™ is built ON SagaChain. SagaChain is not purpose-built for IoT — it is the open Layer-1 whose patented object model unlocks governed AI and device command use cases that no other chain can address cost-effectively.

PLATFORM ARCHITECTURE

SagaAIoT™: Three Integrated Planes

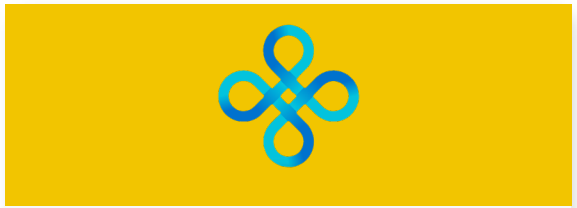

Data Plane

SagaIoT™

Transport-agnostic device messaging — XMPP and MQTT, both first-class protocols. High-throughput, enterprise-scale. Device payloads never traverse the ledger.

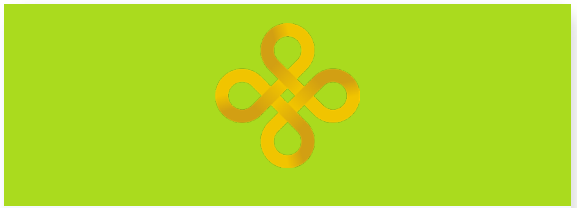

Trust Plane

SagaChain™

The governance ledger. Device registry, command authorization, audit evidence, and policy objects as immutable XBOM objects. Independent of any vendor cloud.

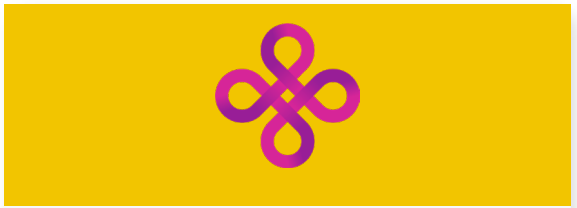

Agent Plane

SagaAI™

AI agent identity and authorization. Structural enforcement of what agents are permitted to do — not documented, not advisory, not bypassable without an immutable record.

Core rule: device payloads travel on the high-throughput messaging network. Only policy records and audit evidence go to the ledger.

HOW AUTHORIZATION WORKS

Following One Authorization Event Through Each System

Legacy System

- Agent receives API key or certificate
- Calls a function; system checks a DB table
- Result logged by the same system that acted
- Log is the organization's testimony about itself
- Insider with DB access can alter or suppress

Authorization separated from execution. Audit is disputable.

Conventional Blockchain

- Transaction submitted; contract checks state
- Result immutably recorded on-chain
- BUT: no governance policy context recorded
- No record of which risk assessment was in scope
- No human oversight link; no agent boundary

Transaction immutable. Business context absent.

SagaChain + SagaAIoT™

- Access token checked (scoped, time-limited)
- OAuth → RBAC → ABAC/PBAC evaluated in sequence
- Audit record created BEFORE the device acts
- Human co-sign required for safety-critical verbs
- Token + session + policy + command = one chain record

Enforcement IS the object. Independently verifiable.

CAPABILITY COMPARISON

Side-by-Side: What Each Approach Delivers

Requirement	Legacy / Centralized	Other Blockchains	SagaChain + SagaAIoT™
Non-bypassable authorization	Policy docs + perimeter controls	Custom engineering on flat state	**Authorization is the chain object**
Indisputable audit records	Centralized, alterable logs	Tx immutable; context absent	**Consensus-verified governance objects**
AI agent governance	API keys; post-hoc audit	No mature standard	**Structural constraints via SagaAI**
Physical device commands	Application logs; no pre-auth	Impractical throughput	**Pre-execution gate; payloads off-ledger**
Portable governance records	Vendor-locked	Chain-portable in principle	**Self-hosted; moves with organization**
Human oversight enforcement	Disconnected approval workflow	Not chain-native	**Named co-sign required; structural**
Quantum-resistant security	Retrofit required; unclear path	Emerging; few full stacks	**NIST-aligned; opt-in today**

SAGAAI GOVERNANCE MODULE

Four-Layer Access Control: OAuth, RBAC, ABAC, PBAC

Non-Bypassable

Every access decision is a SagaChain transaction. No path for an agent or privileged operator to act without an immutable record.

Least-Privilege

Granular role-based permissions with full inheritance hierarchy. Time-bound elevation for contractors or temporary access.

Delegated Auth

OAuth 2.0 patterns on-chain: short-lived scoped tokens, hashed references (raw secret never on-chain), immediate revocation.

Unified Audit Trail

Every token, role grant, session, and guard evaluation on the same ledger as your regulatory framework records.

Composable Stack

Integrates directly with SagaAI guard infrastructure, MCP tool invocations, and the SagaAIoT governed command path. ABAC/PBAC context attributes evaluated per-action: risk score, time window, device attestation state, playbook step — producing permit / deny / require-human-approval outcomes.

Regulatory Aligned

Directly addresses EU AI Act Art.12/14, ISO/IEC 42001 Annex A, and NIST AI RMF Govern + Measure functions.

Status: OAuth/RBAC + ABAC/PBAC implemented as SagaChain objects · 10/10 tests each · ARA Phase 4+4b shipped (fixture + registry client) · Live registry E2E pending chain domain load

REGULATORY ALIGNMENT

Why This Matters for Regulated Industries

EU AI Act

Article 12 (Logging) · Article 14 (Human Oversight)

High-risk AI systems require documented evidence of human oversight and traceability of decisions. SagaChain satisfies this structurally: every access decision, command dispatch, and human co-sign is a consensus-verified chain object. Evidence for auditors is created in real time by the enforcement mechanism itself — not reconstructed afterward.

NIST AI RMF

Govern + Measure Functions

The Govern and Measure functions require organizations to demonstrate that governance controls are operational, not just documented. SagaChain governance objects provide queryable, provenance-rich evidence aligned directly to these functions — with no dependency on the organization's own testimony.

ISO/IEC 42001

Annex A (Access Controls) · Clause 9 (Performance Evaluation)

The AI Management System standard requires access controls and performance evaluation records. The SagaAI four-layer stack (OAuth/RBAC/ABAC/PBAC) implements both as chain objects — not as database entries in a system the organization controls unilaterally. Records are independently verifiable by third-party auditors.

SECURITY POSTURE

Security Testing: What Was Done and What Was Not

Completed (Internal)

- **Release Gate**
  make platform-vuln-bundle → 18/18 pass (re-verified 27 June 2026)
- **Adversarial P1**
  Authorization matrix + traversal → 3/3 pass
- **Adversarial P2**
  Fuzz + command bridge → 3/3 pass
- **Adversarial P3**
  Governed workflow lab → 4/4 pass, 1 skip
- **Production Config Audit**
  Live sagaiot.net → 13/13 pass
- **Container CVE Scan**
  Trivy on sagaiotd → 0 unfixed HIGH/CRITICAL
- **Static Analysis**

Not Yet Completed

- **P4 Federation + Chain Adversarial**
  3/3 pass (complete)
- **External / third-party penetration test**
  Not conducted — required before strongest external security claims
- **Mosquitto/EMQX full source audit**
  Broker images not fully scanned
- **Production digest parity**
  Broker image/prod digest not fully verified
- **ARM MCU hardware audit**
  Pending for post-quantum attestation

Organizations conducting security due diligence should factor this scope boundary into their evaluation.

HONEST ROADMAP

Where Things Stand: Shipped vs. In Development

Capability	Status	
High-throughput device messaging (XMPP + MQTT, co-equal)	Shipped & certified at enterprise scale	Shipped
Security validation — internal test suite (18/18 + P1–P4)	Shipped — internal only; no external pen test	Note
Governed command path with pre-execution policy gate	Shipped (opt-in)	Shipped
Device attestation with tamper-evident audit anchoring	Shipped (opt-in)	Shipped
Human co-sign requirement for safety-critical commands	Shipped	Shipped
OAuth + RBAC + ABAC + PBAC as chain objects (four-layer stack)	Code complete · 10/10 tests each · ARA Phase 4+4b shipped	Shipped
Quantum-resistant cryptography (NIST FIPS 203/204/206)	Shipped (opt-in)	Shipped
Public TestNet	In development	In Dev
SagaAI live registry E2E on loaded chain node	Pending chain domain load	In Dev
Production MainNet	In development — no launch date published	In Dev

REAL-WORLD APPLICABILITY

Industry Applications

Manufacturing & Utilities

AI agents govern machine commands and grid operations. Pre-execution policy gates prevent unauthorized actuation. Audit records survive any dispute about what was authorized and when.

Healthcare

High-risk AI classification under EU AI Act demands human oversight records and traceable decisions. Every clinical AI action linked to the authorization policy that permitted it.

Finance & Supply Chain

Supply chain events (GS1 EPCIS, DSCSA), pricing transparency, and transaction monitoring governed by agents whose scope and role are structurally enforced and independently auditable.

Smart City & Critical Infra

Multi-stakeholder IoT estates with competing governance requirements. Portable governance objects travel with the organization — no hyperscaler lock-in, no vendor cloud dependency.

ENGAGEMENT

How to Engage with PraSaga Foundation

Pilot Program

Pilots run on your infrastructure. No dependency on public TestNet. Governance records, device data, and audit evidence remain in your environment.

- Messaging layer deployment alongside your existing device estate
- Governed command path integration for one device class or workflow
- Pre-execution policy gate and tamper-evident audit record demonstration
- Mapping your governance framework to SagaChain objects

PraSaga Foundation

prasaga.com

sagaiot.net

code.prasaga.com/sagachain

sagascan.prasaga.com

Disclosure: SagaChain is currently in public development testnet and has not launched to production MainNet. Capabilities marked 'In Development' are not shipped. All governance modules are assistive infrastructure; final authority and regulatory compliance responsibility remain with designated institutional actors.




Authorization should not be the weakest link in AI governance.

It should be one of the strongest proofs of trustworthiness.

SagaChain makes governance enforcement structurally inseparable from the decision itself.

Not a log. Not a policy document. The enforcement is the object.

prasaga.com · sagaiot.net · code.prasaga.com/sagachain · sagascan.prasaga.com