PraSaga Foundation

Your AI Is Already Acting on Your Physical World.

The question is whether anyone can prove it was authorized to do so.

SagaAIoT™ is the governance and accountability platform for AI systems that operate physical infrastructure.

AI Systems Are Making Consequential Decisions Right Now — Automatically.

Without a governance layer, every automated action is an unmanaged liability.

Industrial

A predictive maintenance AI issues a shutdown command to a production line at 2 a.m. The line stops. Who authorized it? Under what rules? What if it was wrong?

Smart City / Utility

An AI agent reroutes traffic signals during an emergency — or adjusts power distribution. Can the city manager explain to the public why that decision was made, and prove it was legitimate?

Facilities / Healthcare

A building AI changes environmental conditions in a clinical setting. When a regulator asks for the authorization record — who approved this, under what policy, and was it logged — what do you show them?

The common thread: the AI acted. But the organization cannot prove it was authorized. That is the exposure.

Three Questions You Cannot Currently Answer.

If an AI agent causes an incident on your physical infrastructure, regulators and counsel will ask:

1

Who authorized that action?

Not "the AI system" — that is not an answer that satisfies a regulator, an insurer, or a court. Who approved the policy that governed it? Was that policy in effect at the time? Is there a record?

2

What were the rules, and were they followed?

What constraints governed that command? Was there a parameter limit? A human approval requirement? Was the AI operating within its authorized scope — and can you prove it?

3

What happened — and can you prove it?

Did the command execute? Was it acknowledged? What was the outcome? If the answer lives in eight different log files across three vendors' clouds, you don't have an audit trail. You have an archaeology project.

Regulators Are Not Waiting. The Rules Are Already Here.

Organizations deploying AI on physical infrastructure are subject to frameworks that are active today.

HIGH

EU AI Act

Requires documented controls, human oversight mechanisms, and verifiable audit trails for AI systems that interact with physical infrastructure. Non-compliance carries fines up to €30M or 6% of global turnover.

MED

NIST AI Risk Management Framework

Establishes provable governance controls — not self-attestation — as the standard for high-risk AI applications. U.S. federal contractors and critical infrastructure operators are increasingly expected to align.

HIGH

NERC-CIP / Utility Regulations

Grid operators face mandatory audit requirements for automated control actions. Automated decisions affecting power distribution require records that demonstrate authorized operation.

MED

ISO 27001 / GDPR / HIPAA

Audit trail requirements and accountability obligations extend to AI-driven operations. Absence of a verifiable record is itself a compliance failure — not just a documentation gap.

The cost of retrofitting governance after an incident or regulatory inquiry is substantially higher than building it in. The window for proactive action is closing.

SagaAIoT™ Does Four Things. Nothing More.

In plain language, for any executive audience:

01

Checks every AI command before it reaches a physical device.

Before any automated instruction touches a machine, valve, thermostat, signal, or piece of equipment — SagaAIoT checks: Is this agent authorized to do this? Is this within the approved limits? Does this require a human to approve it first? Commands that fail the check never reach the device. There is no workaround.

02

Creates a permanent, tamper-proof record of what was authorized and why.

Every authorized command generates a single, correlated record: which policy approved it, who or what issued it, whether it executed, and what the device reported back. This record cannot be altered after the fact. It is the answer to every question a regulator, insurer, or attorney will ask.

03

Governs multi-step operational playbooks — not just individual commands.

Complex sequences — HVAC procedures, safety shutdowns, utility coordination, incident responses — run step by step under the same governance as single commands. Every step is individually policy-checked. Every step generates its own record. Procedures requiring human approval pause and wait. The entire sequence is one coherent audit record.

04

Connects to any device estate — without disrupting operations.

SagaAIoT includes its own purpose-built messaging layer that connects to industrial, cloud-connected, and on-premises device estates. No equipment replacement, no changes to field teams, no dependence on any specific vendor's platform.

What Changes When SagaAIoT Is in Place.

WITHOUT SAGAAIOT™	WITH SAGAAIOT™
AI commands reach devices with no policy check. Any agent that can connect can issue any instruction.	Every command is checked against policy before it reaches a device. Unauthorized commands are blocked — not logged after the fact, blocked.
When something goes wrong, teams spend days correlating logs across multiple systems trying to reconstruct what happened.	A single, correlated record links the instruction, the authorization, the execution, and the outcome. Available on demand.
Human approval for high-stakes automated decisions is ad hoc or absent. There is no enforceable mechanism.	High-stakes commands require a named human approver before the device receives anything. Configurable by operation type.
Audit evidence for regulators is assembled manually after the fact, from incomplete records across disconnected systems.	Audit evidence is generated automatically, tamper-evident, and available before any inquiry is made — not reconstructed in response to one.

The Cost of Inaction Is Not Hypothetical.

What organizations are actually exposed to when AI operates physical systems without governance:

€30M

or 6% of global turnover maximum EU AI Act fine for high-risk system violations

3–5×

higher cost to retrofit governance after an incident vs. building it in proactively

15–25

year typical device lifecycle in utilities and infrastructure — every year of exposure compounds

0

workaround paths through SagaAIoT's policy layer — non-compliant commands never reach devices

The strategic framing:

Every day AI agents operate on physical systems without a governance layer, your organization is accumulating liability exposure that cannot be unwound retroactively. The question is not whether to govern AI actuation — it is whether to do it before or after an incident forces the issue.

What This Means for Each Member of Your Leadership Team.

SagaAIoT addresses a different but connected concern for every executive function:

CEO

BOARD ACCOUNTABILITY

When the board asks whether AI automation is being governed responsibly, the answer is documented, not promised. Governance is built into operations, not bolted on after a crisis.

CFO

LIABILITY & INSURANCE EXPOSURE

Ungoverned AI actuation is an unquantified liability on the balance sheet. SagaAIoT creates a verifiable audit record that defines the boundary of organizational responsibility and supports insurance underwriting for AI-driven operations.

COO

OPERATIONAL CONTINUITY

When an AI system causes an unplanned event, SagaAIoT provides the complete chain of what was authorized and when — eliminating the days-long log archaeology that currently precedes any remediation decision.

CLO / General Counsel

REGULATORY DEFENSE

Every governed command generates a tamper-proof record before the device acts. That record is the response to a regulatory inquiry, a litigation hold, or a third-party audit — produced on demand, not reconstructed under pressure.

CISO

ATTACK SURFACE & IDENTITY

AI agents operate as identities on your network. SagaAIoT enforces that every AI identity is verified before it can issue a command, every command is within defined bounds, and every action is recorded cryptographically. The ARA ships a four-layer hybrid access-control stack (OAuth/RBAC/ABAC/PBAC) — context-aware policy evaluation producing on-chain access decisions with frozen context snapshots. Phase 4+4b fixture mode shipped; live registry E2E pending domain load.

CTO / CIO

ARCHITECTURE & VENDOR LOCK-IN

SagaAIoT is a complete platform — PraSaga's own messaging layer plus governance — deployed entirely within the organization's infrastructure. Governance records are self-hosted and portable. No hyperscaler dependency. No vendor lock-in on the audit record.

You Own the Governance. No One Else Does.

Every other governance option ties your audit records and policy controls to a vendor's platform. SagaAIoT does not.

Your policies, your infrastructure.

Authorization rules, audit records, and compliance evidence are stored in your environment. You do not call a vendor's API to prove what your AI did.

No cloud dependency for compliance.

If your hyperscaler has an outage, your governance record is still intact. If you change cloud providers, your governance model moves with you.

Connects to any device estate.

SagaAIoT's purpose-built messaging layer connects to industrial, cloud-connected, and on-premises device environments. No equipment replacement, no firmware changes, no field team disruption.

Deployable in any regulated environment.

Sovereign cloud, air-gapped networks, on-premises data centers, jurisdictions with strict data residency requirements — SagaAIoT operates in all of them.

What you get with a hyperscaler IoT platform:

✗ Audit records stored in their cloud

✗ Policy controls defined by their APIs

✗ Governance model changes when they change their platform

✗ Lock-in risk if you need to migrate

✗ Limited options for sovereign or regulated deployment

What Is Built and Verifiable on the Development TestNet Today.

Built, CI-Verified, and Live — Reviewable at sagaiot.net and sagascan.prasaga.com:

✓ Policy-Gated AI Commands

Every AI command is checked against defined policy before the device receives it. Commands outside the approved scope are blocked — not logged after the fact, blocked. Live on the production reference at sagaiot.net.

✓ Governed Operational Playbooks

Multi-step procedures run under the same governance as individual commands — each step individually policy-checked, audited, and held for human approval where required. One coherent audit record for the entire sequence.

✓ Human Approval Enforcement

Designated command types or procedure steps pause and wait for a named human approver. A hard gate — the device receives nothing until approval is given. Configurable per command type and per procedure step.

✓ Single Correlated Audit Record

Every governed command and every governed procedure step produces one complete record: authorization, issuing identity, dispatch, device confirmation. Created before the device acts.

✓ Self-Hosted Deployment

The full platform runs in your environment — cloud, on-premises, or sovereign infrastructure. Your governance records do not leave your perimeter.

✓ Automated Response Rules

When a device enters a defined fault state, a response workflow initiates automatically — with a full audit record of every step. Operations teams focus on exceptions, not noise.

✓ SagaAI four-layer access control (OAuth/RBAC/ABAC/PBAC) — ARA Phase 4+4b shipped (fixture + registry client). ABAC/PBAC context-aware escalation: risk, time window, device attestation, playbook step evaluated per-action. Live registry E2E pending chain domain load.

We Tell You What Is Not Finished.

PraSaga does not represent planned capabilities as shipped. These are in progress:

What this means for discovery and pilots:

IN PROGRESS

SagaChain registry-mode governance rollout

In plain terms: Live at sagaiot.net (June 2026). Governed commands, playbooks, and attestation CI-verified (platform-vuln-bundle: 18 pass / 0 fail). Enterprise pilots available now on customer infrastructure. Chain-registered governance mode and SagaAI guard-certificate wiring are the remaining sequencing items.

Pilot impact: Pilot impact: Pilots can begin now. Organizations engaging during registry-mode rollout shape pilot design and gain priority access as full chain governance activates.

IN PROGRESS

Full autonomous AI agent identity verification

In plain terms: The platform already blocks unauthorized AI commands, generates the full audit record, and governs multi-step procedures. The OAuth + RBAC Governance module (scoped tokens, role hierarchy, on-chain access guard) is code complete with 10/10 tests passing — pending SagaChain domain load. The ARA Phase 4+4b ships the hybrid OAuth/RBAC/ABAC/PBAC guard + guard certificate in fixture mode, with a registry client for txn submission. The remaining step is live registry E2E on a loaded chain node — a domain-load dependency, not a missing capability.

Pilot impact: Human-operator governed commands and governed playbooks are fully verifiable today. Full autonomous agent governance requires one additional development milestone.

IN PROGRESS

Live ledger anchoring of audit records

In plain terms: Audit records are generated and cryptographically secured today. The step of committing them to the permanent ledger is built and tested, but runs in verification mode by default — not live submission.

Pilot impact: For pilot and discovery purposes, the audit record is fully available and verifiable. Live ledger anchoring is an activation step for production deployment, not a re-architecture.

From Discovery Today to 90-Day Pilot — Available Now.

"Can we demonstrate — to a regulator, an insurer, or a board — that our AI operates under enforceable governance?"

Now

Discovery and Code Review

Review the live production reference at sagaiot.net (XMPP, MQTT TLS, operator console — deployed June 2026) and the governance record structures at sagascan.prasaga.com. Technical and compliance teams examine governance objects, command records, and procedure audit trails directly. No sales process required.

Outcome: Informed decision on whether to pursue a pilot. Questions that come from the code review are the right ones to bring to PraSaga.

Month 1–2 (Customer Infrastructure)

Connect and Demonstrate Governance

Connect to one device class on customer infrastructure, enable the policy gate for one command type, define one multi-step procedure, and walk through all governance outcomes: rejection, approval, human cosign, full procedure audit record. TestNet governance records available in parallel for compliance review.

Outcome: Live demonstration of governed AI actuation and governed playbook execution. Complete audit records in hand for all outcomes.

Month 3 (Evaluate and Decide)

Evaluate and Decide

Review pilot findings against the organization's compliance and operational requirements. Expand scope if warranted. Deliver evaluation report to executive sponsor with production readiness recommendation.

Outcome: Evidence-based production decision. Production scope defined or specific gaps identified with timeline.

PraSaga Foundation

The governance question is not going away.

Every week you operate AI on physical infrastructure without a governance layer, you are extending a liability that cannot be corrected retroactively.

1. Review the live platform: sagaiot.net (production reference, June 2026) and sagascan.prasaga.com — governance objects, command records, and procedure audit trails publicly visible. Pilots available now on customer infrastructure.
2. Schedule a discovery session: 60-minute technical review of what is built, with the engineering team walking through the governance model and audit record structure.
3. Define a pilot use case: identify one device class, one command type, and one multi-step procedure — scope the pilot before public TestNet opens so you are first in line.

Strategic Partnerships and Enterprise Development | www.prasaga.com

sagaiot.net · sagascan.prasaga.com · sagascan.prasaga.com/wizard