INFRASTRUCTURE AI WHITE PAPER

Agentic Security

Architecting Agentic Security for Galaxy Agentic Operating System.

A dynamic, multi-layered operating system framework that shifts security from post-event logging to **deterministic, runtime execution governance**.

01

The Engineering Challenge in Autonomous Infrastructure.

Deploying AI into physical infrastructure introduces a critical engineering problem: **traditional cybersecurity is static** and designed for human-in-the-loop IT environments. It cannot safely govern autonomous agents making real-time, consequential decisions regarding physical equipment, energy loads, and building safety.

GAOS solves this by introducing Agentic Security — a dynamic, multi-layered operating system framework that shifts security from post-event logging to deterministic, runtime execution governance.

TRADITIONAL CYBERSECURITY

Static, IT-era defenses.

Designed for human-in-the-loop networks, built around perimeter logging and post-event response.

AGENTIC SECURITY

Dynamic, runtime governance.

Built into the operating system itself, enforcing policy at every action, every agent, every layer.

02 The Dynamic Intelligence Hierarchy.

Instead of relying on a **centralized, highly vulnerable control node**, GAOS distributes security through a hierarchical "nervous system" of intelligence. This dynamic network mirrors modern infrastructure topology.

Signals flow seamlessly up and down this hierarchy, enabling the system to learn behavioral patterns, detect anomalies, and autonomously execute mitigation strategies across the entire infrastructure footprint.

03 The Six-Layer Deterministic Execution Pipeline.

GAOS enforces a **defense-in-depth architecture** where an AI agent cannot bypass the system's strict capability boundaries. Every automated decision is forced through a six-layer cryptographic and policy pipeline:

REQUEST → EXECUTION PATH

Six cryptographic and policy gates, deterministically ordered.

· DEFENSE-IN-DEPTH · CRYPTOGRAPHIC · NO BYPASS · DETERMINISTIC · IRREVERSIBLE FORWARD FLOW →

From intent to execution — every step is governed, gated, and recorded.

03

THE FIRST HALF · IDENTITY → CONTAINMENT → CAPABILITY

01 02 03 04 05 06

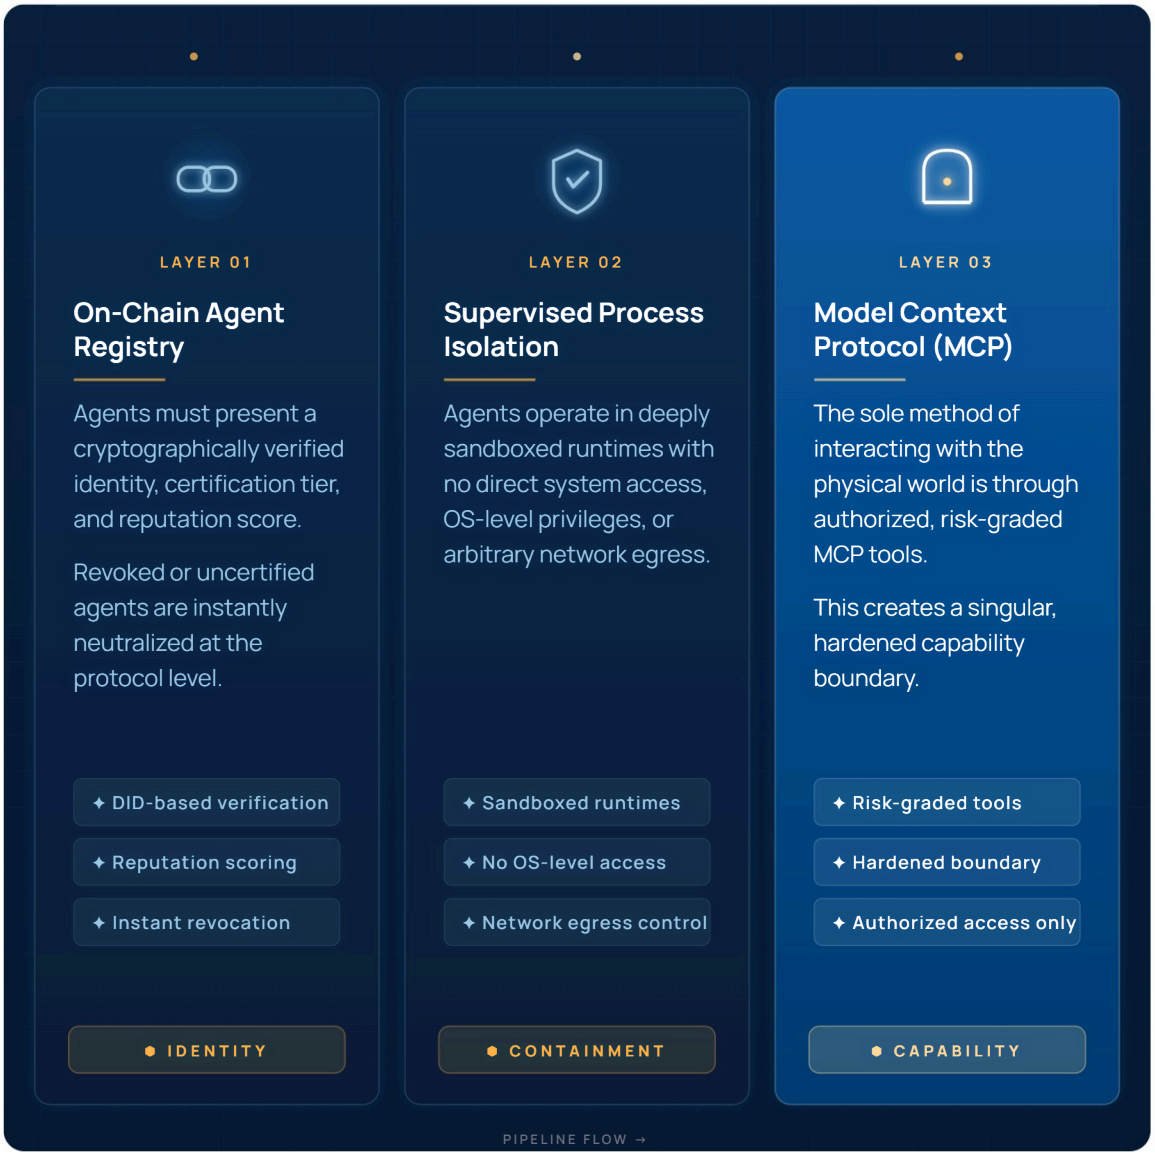

03

THE SECOND HALF · DECISION → ROUTING → EVIDENCE

01 — 02 — 03 — 04 — 05 — 06

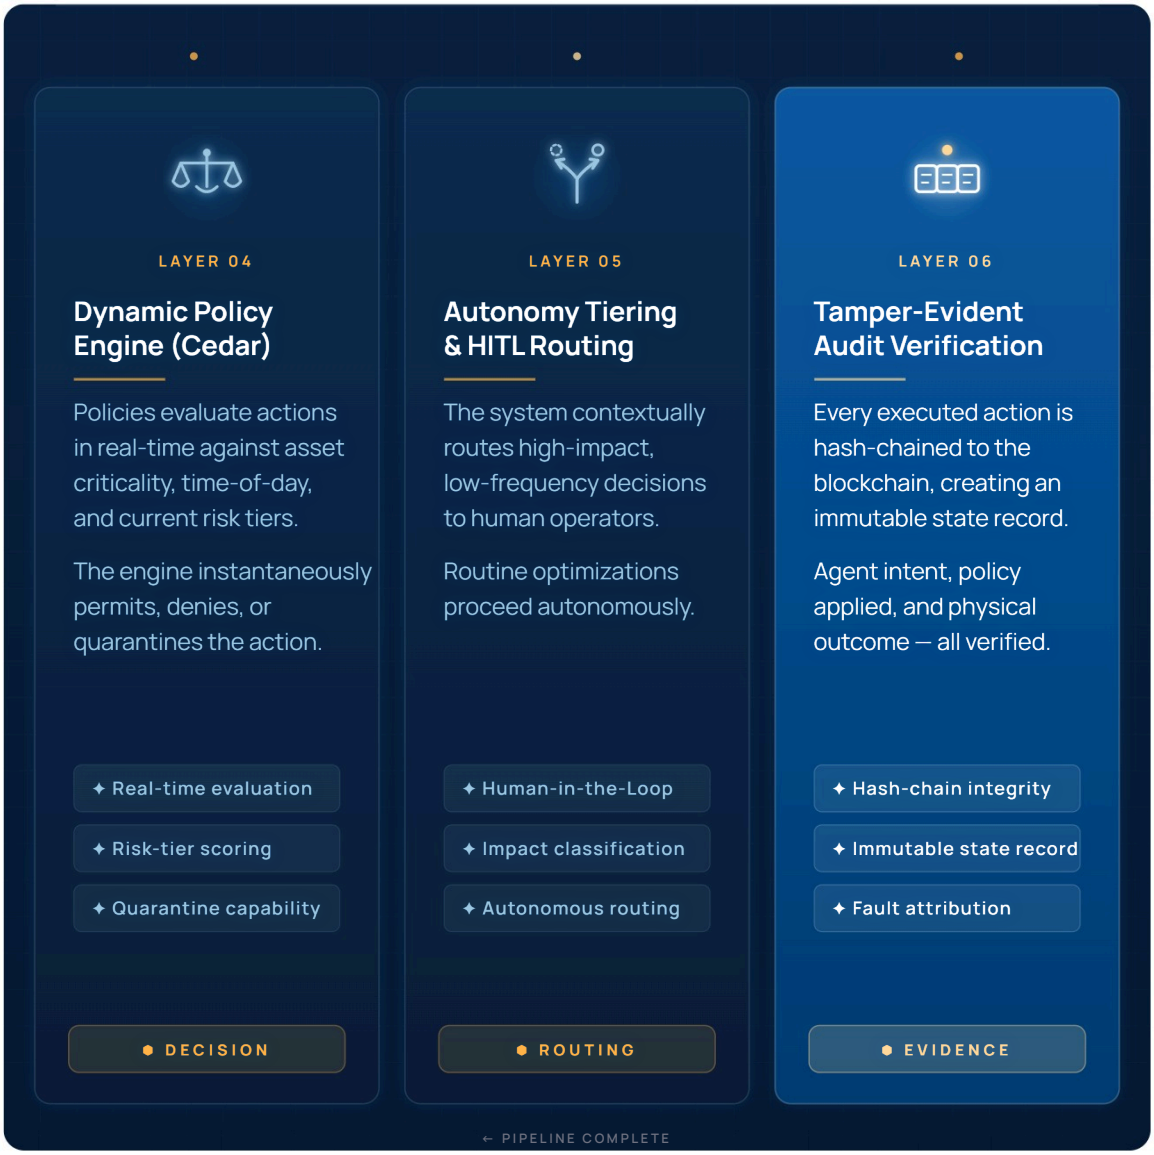

04 The Problem with Monolithic Architecture.

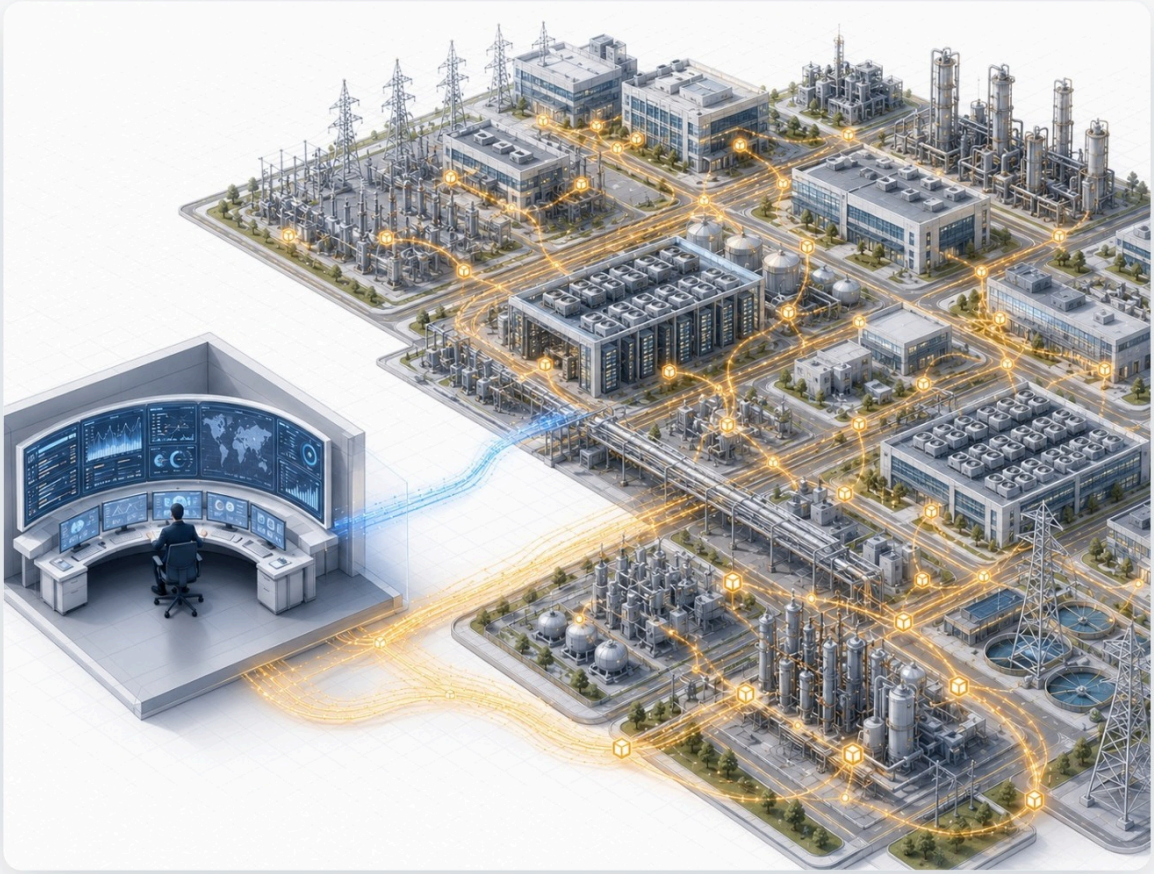

In autonomous infrastructure, mixing the state of physical assets with the behavioral logic of AI agents creates an **unacceptable risk profile**. A monolithic architecture — where asset telemetry, policy execution, and agent identities share the same ledger — creates processing bottlenecks and fundamentally breaks the separation of concerns required for regulatory compliance and fault attribution.

> If an AI agent malfunctions and alters a critical infrastructure asset, auditors must be able to independently verify the asset's state separately from the AI's internal governance.

04 The Dual-Ledger Innovation.

GAOS solves this through a **proprietary Dual-Blockchain Security Backbone**, a modular architecture designed to separate physical infrastructure assurance from AI agent governance while keeping them deterministically linked.

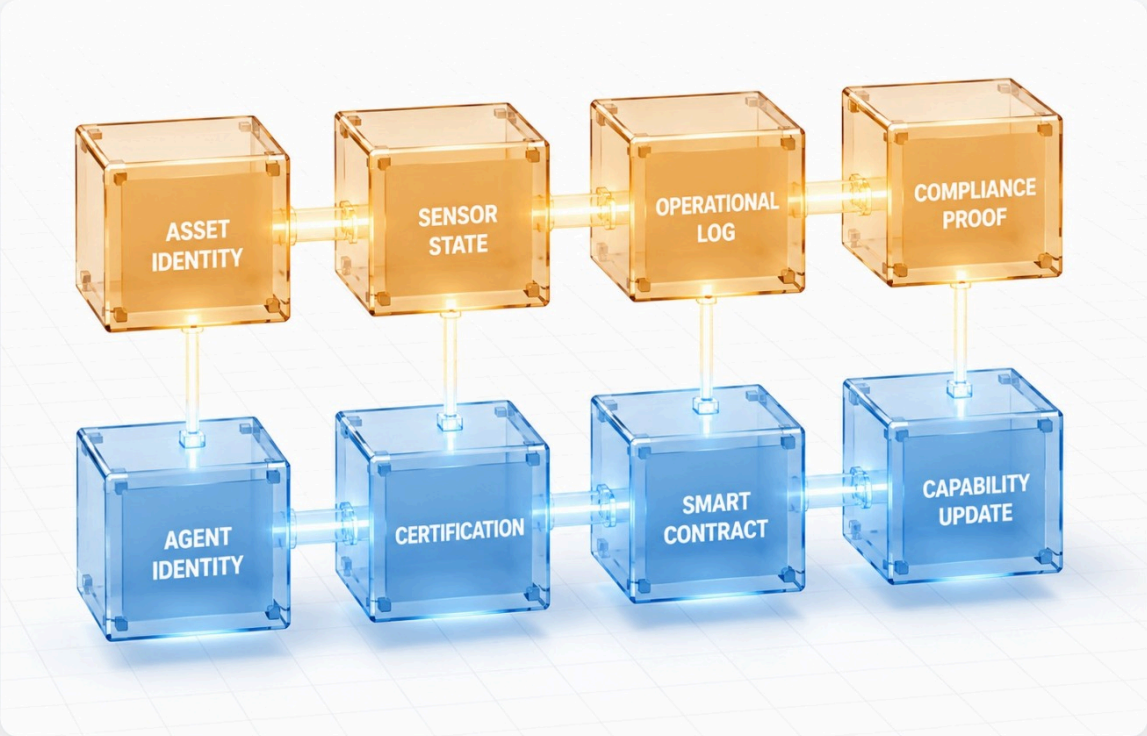

LAYER 1 · INFRASTRUCTURE BLOCKCHAIN

Physical State.

This ledger functions as the unalterable source of truth for the physical world — maintaining the Asset Identity Registry, operational data logs, and performance compliance records. Asset owners, insurers, and regulators rely on this chain for tamper-evident verification.

⬢ ASSETS · SENSORS · COMPLIANCE

LAYER 2 · AGENTIC BLOCKCHAIN

AI State.

This parallel ledger governs AI agents — managing identities, certification and reputation scores, and smart contracts for autonomous behavior. OEMs and administrators use this chain for version control, capability updates, and immediate ecosystem-wide revocations.

⬢ IDENTITY · GOVERNANCE · CONTRACTS

04 Cryptographic Interoperability and Fault Attribution.

The true power of this architecture lies in its interoperability. While deployed as logically separated ledgers to preserve governance integrity, specialized indexers and bridging services create **cryptographically correlated views** between the two chains.

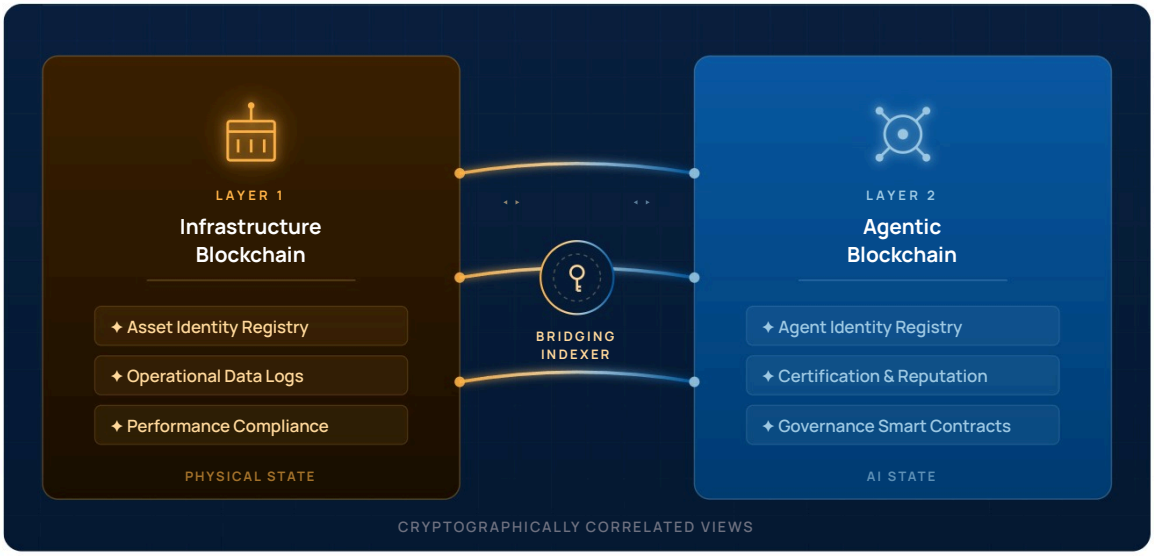

> **"Which certified agent (Agentic Chain) executed a high-impact setpoint change on this specific chiller (Infrastructure Chain) over the last 30 days, and under which Cedar policy was it approved?"**

By cleanly decoupling the physical asset state from the AI execution logic, GAOS delivers full traceability. This dual structure is the critical enabler that transforms autonomous infrastructure operations into auditable, verifiable, and ultimately, investment-grade financial assets.

05 The Legacy OT Vulnerability.

In traditional building automation, security is fundamentally flawed because it relies on **static permissions**. Today, an external vendor can often push script updates or execute code directly to an automation controller, instantaneously altering HVAC load profiles or safety-critical lighting systems without real-time oversight. This creates a massive vulnerability where unverified, human-driven changes bypass asset-level risk assessments.

GAOS nullifies this vulnerability through its Model Context Protocol (MCP) framework. AI agents — and by extension, the vendors deploying them — are completely prohibited from accessing databases, communicating via arbitrary network endpoints, or executing commands directly to OT hardware.

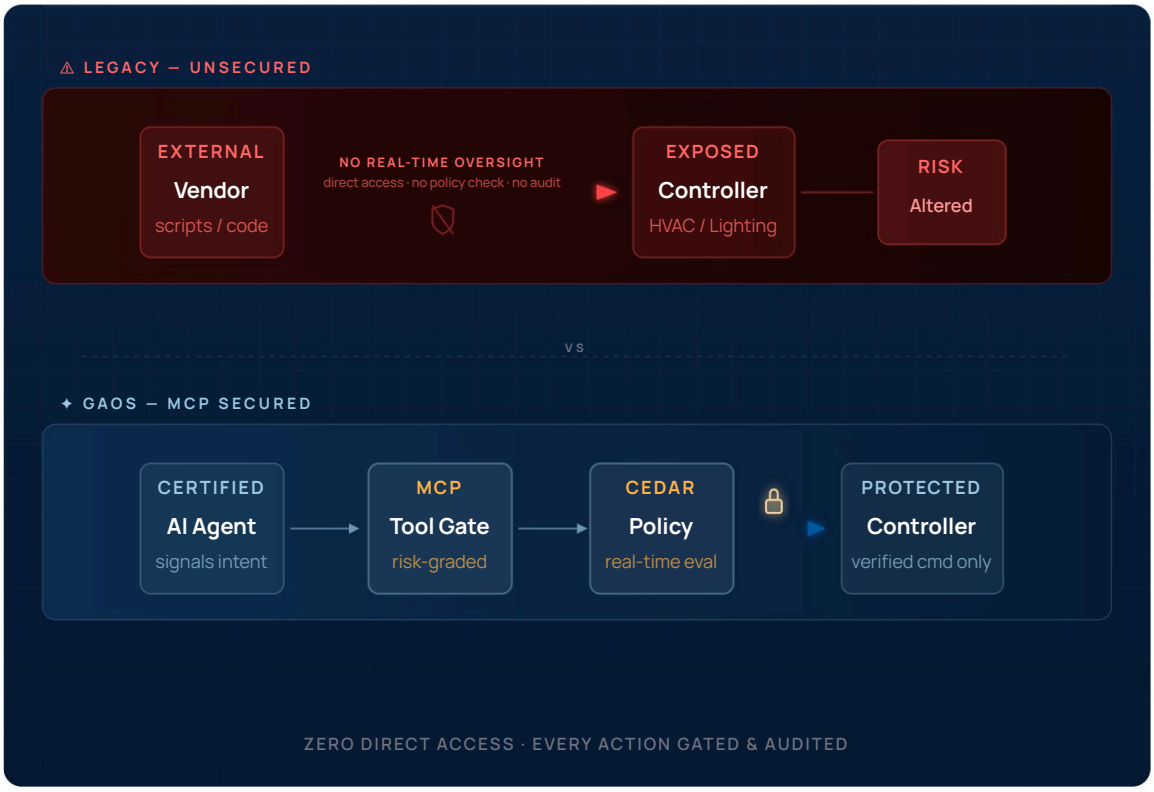

05

The MCP Tool Gating Mechanism.

Instead, all interactions with physical controllers must pass through authorized, risk-graded MCP tools. The execution flow acts as an **impenetrable gateway**:

- **01** · An agent signals intent to adjust a setpoint.
- **02** · The request is forced through the MCP Tool interface.
- **03** · The Cedar Policy Engine evaluates the exact state of the asset, time-of-day, and the agent's autonomy tier.
- **04** · Only upon policy clearance is the command forwarded to the execution target (the legacy controller).

EVERY REQUEST · EVERY GATE · EVERY TIME

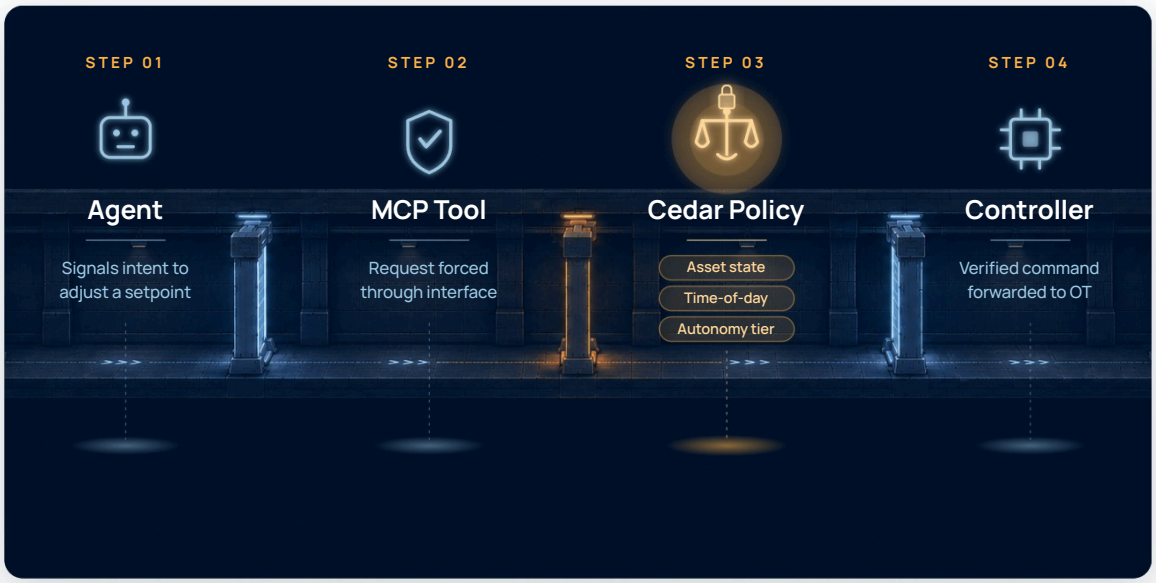

AGENT → MCP → POLICY → CONTROLLER

an impenetrable gateway · zero direct access

05 Heterogeneous Runtime Interoperability.

GAOS is architected to be **fundamentally model-agnostic**. It applies a unified Agentic Security envelope across diverse AI architectures, whether they are LLM-driven generative agents, containerized machine learning models, or traditional native control logic. This ensures that as underlying AI models evolve, the foundational security and governance infrastructure remains robust, standardized, and fully defensible.

If a legacy system begins reporting data anomalies, or if a connected tool exposes potentially dangerous operations (e.g., breaker switching), GAOS instantly falls back to a fail-safe state. Agents are dynamically downgraded to lower autonomy tiers or forced into observe-only "Monitoring Mode" until human supervision (HITL) overrides the system.

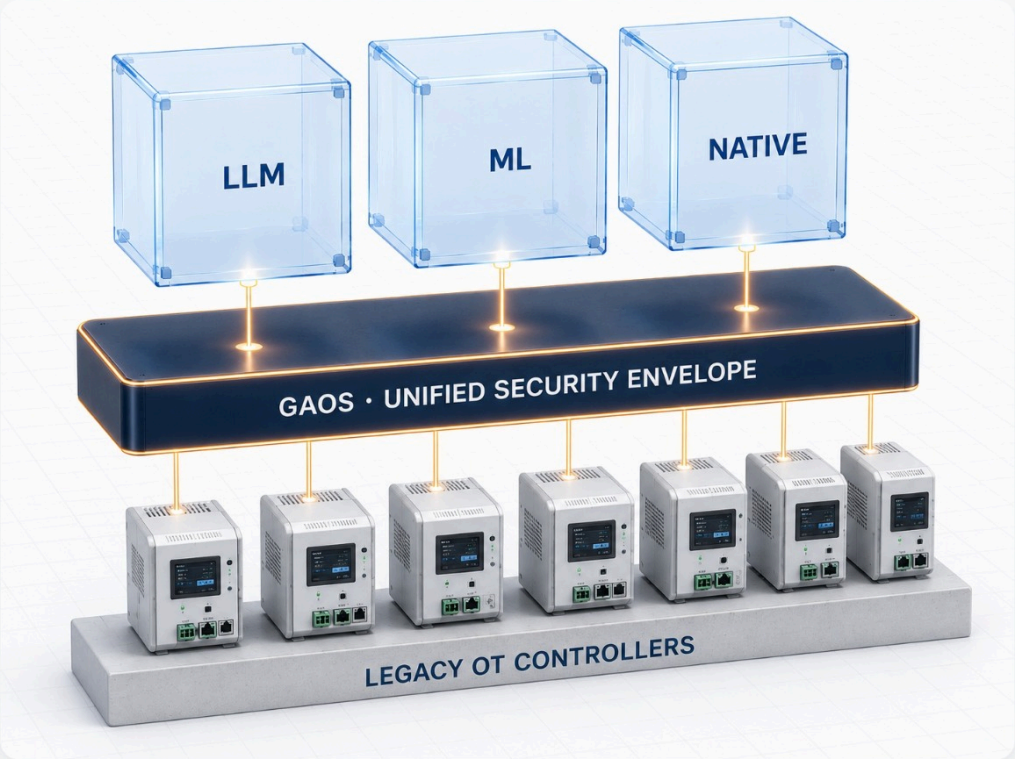

> GAOS allows asset owners to implement next-generation autonomy without ripping and replacing their existing hardware footprint.

06 Redefining Infrastructure Valuation.

Agentic Security is fundamentally a **core valuation driver**, not just a defensive cybersecurity feature. By replacing reactive, traditional security with dynamic runtime governance, GAOS establishes the necessary framework to safely scale autonomous infrastructure. This architecture transforms physical infrastructure into trusted, self-regulating, and auditable digital assets.

> The economic and commercial implications of this trust layer are profound.

06

01

Unlocking Regulated Markets

GAOS provides the **cryptographic proof and defense-in-depth architecture** required to deploy AI into highly regulated sectors, such as commercial real estate, energy grids, and public works, where safety and auditability are mandatory.

02

Audit-Grade Compliance

The system generates **tamper-evident, dual-blockchain audit trails** that align precisely with the expectations of regulators, auditors, and insurers, potentially lowering compliance costs and insurance premiums.

03

Enabling Financial Innovation

Verifiable, immutable data regarding asset performance and AI intervention serves as the foundation for modern fintech applications. This allows for the secure tokenization of physical assets and enables performance-linked financing models grounded in absolute data integrity.

06 Market Impact.

The deployment of GAOS represents a **foundational shift in the digital infrastructure economy**. It acts as the critical enabler for smart buildings and smart cities, providing the trust layer required to manage millions of connected assets — from localized microgrids to expansive commercial campuses.

> Under defined platform configurations, GAOS ensures there is no pathway for an AI agent to impact infrastructure without complete traceability, significantly reducing the mean time to detect and remediate high-risk behaviors.

06 Evolutionary Roadmap.

The GAOS Agentic Security Framework is **architected for continuous evolution**, designed to scale alongside advancements in AI capabilities and regulatory demands. The strategic roadmap includes:

- **Advanced Hardware Integration:** Transitioning from software-supervised runtimes to stronger isolation models utilizing hardware-backed secure enclaves for process control.
- **Dynamic Marketplaces:** Evolving the Agent Registry from isolated on-chain certificates to multi-chain, cross-ecosystem attestations, enabling rich agent reputation models and autonomous capability marketplaces.
- **Tokenized Capability Access:** Upgrading the Model Context Protocol (MCP) to support tokenized access and dynamic tool risk pricing.
- **AI-Assisted Governance:** Moving Human-in-the-Loop (HITL) processes from active control to AI-assisted approvals, utilizing risk-based routing to prevent operator fatigue.
- **Zero-Knowledge Compliance:** Advancing the blockchain audit layer to support public or shared proofs via zero-knowledge attestations, allowing operators to prove compliance to regulators without revealing sensitive operational data.

07 Conclusion.

Infrastructure AI's framework establishes a new global standard for operations.

> **GAOS guarantees that autonomy does not compromise control, intelligence does not bypass governance, and execution never escapes accountability.**

As AI agents increasingly dictate real-world outcomes, GAOS provides the **definitive trust architecture** for the autonomous world.