AI CONTROL

CONTROLCORE · AI FUTURE TRACK · BLOCKCHAIN LIFE 2026 · DUBAI · 1–2 DECEMBER

Human authority over autonomous intelligence.

AI may be autonomous in operation, but autonomy does not create authority. ControlCore puts a human-authorized decision in front of every consequential action.

Autonomy arrived. Authority never did.

Agents, copilots and robot planners now act on production systems, funds and the physical world. Identity systems answer who is asking. Nothing answers whether it may.

FAILURE MODES THE PLATFORM IS BUILT AROUND

- credential theft
- compromised developer
- prompt injection
- privilege escalation
- unknown AI channel
- unknown machine language
- compromised robot
- insider misuse
- policy misconfiguration
- telemetry manipulation
- cloud outage
- trust service outage

Authority is granted. Never assumed.

No authority by default

Unknown agents are blocked. An agent nobody registered is an agent nobody authorized.

Identity is not authorization

Knowing who an agent is says nothing about what it may do. Capability is not permission either.

Revoked stays revoked

A revoked or quarantined agent is blocked everywhere, at once, including at the edge.

High impact needs a human

Production writes, robot motion, physical actions and funds always wait for a person. Policy only ever makes a decision stricter.

Every request runs the same gauntlet

Each stage can only narrow what the previous one allowed. Every verdict is signed and kept in an append-only audit record.

REQUEST
An agent, copilot or robot planner asks to act.

IDENTITY
Who is asking, proven, not a name in a header.

CAPABILITY
Can this agent perform this class of action at all?

POLICY
Versioned, signed, auditable rules.

CONTEXT & RISK
Target, blast radius and live telemetry become a score.

HUMAN APPROVAL
High impact pauses for an authorized person.

VERDICT
Signed, then executed exactly as approved.

ALLOW
Proceed through the execution gateway.

LIMIT
Proceed with a narrower scope or rate.

QUARANTINE
Isolate the agent, keep the evidence.

BLOCK
Deny, and raise it to the SOC.

THE PLATFORM

One control plane, many enforcement points

ControlCore
The pipeline, as one authorize endpoint.

AI Firewall
Inspect and enforce every AI request in flight.

Policy Engine
Versioned, signed, auditable policy.

Risk Engine
Contextual risk as an input to policy.

AI Passport
Every agent's identity and authority, committed.

Trust Chain
Verifiable authority, anchored on Solieum.

Communication Intelligence
AI-to-AI traffic, decoded for humans.

Robot Gateway
Authorized commands only reach physical systems.

Edge
Local enforcement that survives outages.

SOC
Security operations for autonomous systems.

When humans cannot read the machine language

Agents talk to agents. The traffic is visible to authorized humans, and the AI Language Decoder turns it into something a person can judge.

Decode

Protocol discovery and machine-language parsing turn signal into structure.

Explain

Semantic intelligence turns structure into meaning, and behavior correlation shows what happened next.

Verify

A communication graph with replay, and forensic evidence records with integrity guarantees.

Prove every decision.

Identities, policies and control decisions are anchored on Solieum, the Layer 2 for Solana. Anyone holding a proof ID can check who authorized what, and when, without trusting the operator's word.

Registries

Agents, policies and approvers, committed.

Anchoring

Decision records rolled up and settled on Solana through Solieum.

Verification

Paste a proof ID into the Proof Explorer and check it yourself.

Why a chain at all: an audit log the operator controls can be rewritten by the operator. An anchored one cannot.

From human-only to controlled autonomy

Autonomy is earned per agent, per system, through policy. The high-impact classes keep a human in every mode.



Human-only

The agent proposes. A person performs every action.

Human-approved

The agent acts after an authorized person approves each request.

Supervised

Low-impact actions proceed under policy. Everything else waits.

Controlled autonomy

Policy, risk and limits decide. Humans still own production, funds and motion.

Wherever AI touches something that matters

- Enterprise AI & copilots
- Autonomous agents
- Robotics
- Humanoid systems
- Industrial systems & OT
- Cloud infrastructure
- Financial & digital asset infrastructure
- Data centers
- Healthcare
- Government & critical infrastructure
- AI research labs
- Multi-organization AI governance

Exchanges, trading platforms and market infrastructure are the digital-asset case: AI in payments, treasury and trading operations, with funds movements behind a human.

Built to be integrated, not bolted on

REST API

One authorize endpoint under /api/v1, with an OpenAPI 3.1 specification and signed webhooks for event delivery.

SDKs

TypeScript and Python today. Rust and Go planned.

Runs anywhere

Node backend, PostgreSQL and Redis, Docker Compose for the full stack, and an edge mode that keeps enforcing through an outage.

WHERE IT STANDS

Labelled honestly, stage by stage

AI Control Core · stage: Development

Concept

Described in the white paper.

Research

Being validated.

Development

Being built and tested.

Testnet-Pilot

Running with pilot partners.

Planned

On the roadmap.

Every feature on aicontrol.si carries one of these labels. The white paper is the canonical source. Nothing is offered as production infrastructure yet.

Intelligence should not mean unlimited authority.

WEB aicontrol.si

PIPELINE aicontrol.si/controlcore

DOCS aicontrol.si/docs

DEVELOPERS aicontrol.si/developers

STATUS aicontrol.si/status

Scan for the AI CONTROL brochure
solieum.com/brochures/ai-control

© 2026 AI CONTROL · NOT PRODUCTION INFRASTRUCTURE