NOVACOV

Outreach Brief
Updated 2026-02-04 • For business/technical discussion only (not legal advice).

Verifier → Permit → Fail-Closed Gateways for Agentic AI execution (network egress, device I/O, accelerator dispatch).

One-line: Agents stay blocked from real actions (egress / device I/O / dispatch) until they present a verifiable permit + audit-ready receipts.

WHY NOW (POLICY PULL)

Agentic AI turns model output into actions. Regulators, auditors, and procurement increasingly ask for provable pre-action controls and audit-ready evidence. NOVACOV is built to generate receipts, not promises.

NOVACOV IN 3 CONTROL PRIMITIVES

Verify-to-Activation Append-only verifiable log publishes signed heads. Runtime validates freshness, inclusion, and append-only evolution.	**Permit-before-Action** On pass, mint a signed, expiring permit bound to audience (agent/tenant/mission) and record a receipt in an IAL.	**Fail-Closed dispositions** On fail, return HOLD / QUARANTINE / DENY / ESCALATE. Gateways require a valid permit identifier before subsequent actions.

PROOF ARTIFACTS (BUYER-FACING)

- Signed head (tree size, root hash, timestamp, signatures; optional witness co-signatures)
- Inclusion proof (ICC commitment is included in the current head)
- Consistency proof (append-only evolution when the head advances)
- Permit token (signed; audience-bound; nonce + monotonic counter; expiry)
- IAL record (hash-chained receipt per decision; exportable to audit systems)

REGULATORY / POLICY MAPPING (HEADLINE)

Framework	Timing / buyer ask	NOVACOV mapping
EU AI Act	High-risk rules staged 2026/2027; full roll-out by 02 Aug 2027	Pre-action gating + audit receipts (proofs + IAL).
Korea AI Basic Act	In force 22 Jan 2026	Traceable approvals (tier/quorum) + evidence pack.
US Federal AI governance (OMB M-24-10)	Procurement wants proof of controls	Controls enforced before action; exportable receipts.
NIST AI RMF + ISO/IEC 42001	Common audit language for AI risk management	Receipts bind policy + approvals + execution evidence.

DESIGN PARTNER OFFER (6–12 WEEKS)

- Week 0–2: align threat model + target gateway; agree evidence schema + reason codes.
- Week 3–6: integrate verifier MVP in sandbox; show permit-before-egress + receipts export.
- Week 7–12: move enforcement to driver/firmware/hypervisor; hardening plan + pilot-to-prod option.

Success condition: identify the platform owner and schedule a 30-minute deep dive with Security present.

Two questions: (1) highest-risk action surface? (egress/dispatch/actuation) (2) who consumes evidence? (Security/audit/GRC/regulators)

Acronyms: ICC = Immutable Constraint Chain; IAL = Immutable Arbitration Ledger; ELV = Encrypted Logic Vault; GLG = Licensing Gate.