

ALEETH >

FROM A CATALOG OF RISKS TO ENFORCED, PROVABLE CONTROL

THE MAP AND THE MACHINE

Every serious framework for agentic AI names the risks and the controls you should have. None of them enforces a single action or proves a single control was applied. That gap needs a machine, not a better map.

JULY 2026

ALEETH • THE SCHRECK GROUP, INC.
INSTITUTIONAL CONTROL ARCHITECTURE • JULY 2026

AI has met its match. Truth.™

Every framework is a map. None of them is the machine.

Every serious framework for agentic AI security tells you the same kind of thing: here are the risks, here are the controls you should have. SAIL 2.0 catalogs ninety-one of them. OWASP, NIST, ISO 42001, the EU AI Act, DASF, and AIUC-1 each do their own version.

They are all maps. Good maps. None of them enforce a single action or prove a single control was applied.

That is the gap this paper is about. A map tells you the terrain. It does not stop the vehicle, and it does not leave a receipt. Agentic AI runs at machine speed, and the dangerous failures happen in the one place a document cannot reach: the moment an agent decides to act. Closing that gap needs a machine, not a better map.

ICA is that machine. Not a better description of the terrain. The thing that stops the vehicle and leaves the receipt.

The AI is now the actor, and the old controls assumed it never would be.

For twenty years, security assumed a clean division of labor: humans make decisions, software executes them deterministically. Every control we built, every firewall, every EDR agent, every access policy, sits on that assumption.

Agentic AI breaks it. The AI is now the actor. It reasons, holds memory, chains tool calls, and takes real actions on production systems without a human approving each step. A single injected instruction inside a retrieved document can redirect the whole chain, and the tool calls still look legitimate on the way out.

Three things change at once.

The decision moves into the model. The reasoning that picks which tool to call, and why, is where the compromise lives. Most controls see the action and the log, not the intent behind it.

There is no pause. When a person acts suspiciously, the stack gets several chances to intervene. When an agent goes wrong, nothing stops it mid-execution unless something was built to.

It happens at volume and at speed. Machine-speed tool calls across many sessions, plus the agents nobody registered, mean the surface is larger than the part anyone is watching.

The frameworks name all of this correctly. Naming it is not governing it.

One describes what should be true. The other makes it true and proves it.

Hold the two ideas side by side.

A **map** describes what should be true. SAIL's ninety-one risks, ISO's clauses, the EU AI Act's articles. You read it, you nod, you build a checklist. At audit time you assemble evidence that you did the things. The map is essential. It is also inert. It has never once stopped an action or produced proof on its own.

A **machine** makes it true and proves it. It authorizes each action against policy as the action happens. It contains what breaks, in the moment, not in the postmortem. And it seals a tamper-evident record of every governed step, so the proof exists before anyone asks for it.

The map is what you are supposed to do. The machine is the thing that does it and can show its work. ICA is the machine the maps imply and none of them can be.

This is not a competition with SAIL or OWASP or NIST. It is the opposite. ICA maps to all of them. The point of a control architecture is to be the enforcement and the proof underneath the standard you already answer to.

All ninety-one SAIL risks cross to one architecture that actually runs.

We mapped ICA directly onto SAIL 2.0. All ninety-one of SAIL's agentic risks now cross to ICA's own architecture: thirty-five assessment criteria across seven control layers, plus nine agent-level dimensions. That mapping is live inside the platform today, not a diagram in a slide.

The effect is simple and it is the whole commercial argument. When ICA enforces a control, that control is standing evidence against the matching SAIL risk, and against the frameworks SAIL itself maps to. One architecture, one certificate, many standards answered at once.

The difference from every other row in that mapping table is that ICA is the only entry that actually runs. The rest are documents pointing at each other. ICA is the layer doing the work they describe.

Authorize. Contain. Prove.

If a framework is a map, here is what the machine adds.

Authorize every action, not every session. Session-level access control was built for humans who log in and then behave. An agent passes one check and then chains two hundred tool calls, some of them sensitive, none individually examined. ICA authorizes each action against policy as it happens, with the invoking identity's rights, and holds or blocks the ones that cross the line.

Contain what breaks, in the moment. A control that only reports after the fact is not containment. ICA can stop, halt, quarantine, or kill a governed agent mid-execution. The switch is real, and a designated human, not the agent, holds it. Irreversible actions are held for that human by default.

Seal proof that cannot be edited later. This is the part no checklist can copy. Every governed step produces a cryptographically signed, hash-chained receipt at the moment it happens. Not a log that can be trimmed. Not evidence assembled the week before the audit. A tamper-evident record that already exists, and that anyone can verify independently.

Authorize. Contain. Prove. Those are verbs a document cannot perform.

The reasoning layer and the endpoint, the two places the maps flag as least governed.

Two places decide whether agentic governance is real or theater. Both are the places the frameworks themselves flag as least governed.

The reasoning layer. The dangerous failures do not live in the action. They live in the intent that produced it, one layer down, where no firewall can see. A poisoned page can shift an agent's understanding of its own task across several turns, and every individual tool call still looks fine. Control that stops at the action is looking at the shadow, not the thing casting it. ICA governs and seals the decision itself, beneath the act. That is not a slogan. It is the only altitude at which this class of failure is catchable.

The endpoint. The agents that build and run your business increasingly live on laptops and in browsers: coding assistants with merge authority, agentic browsers holding your sessions, desktop tool bridges reaching local credentials. It is the least instrumented ground in the enterprise, and it is where the agent you are certifying was actually built. Governance that never reaches the endpoint is governing half the surface.

Beneath the act. Above the machine. Impossible to bypass.

That is where a control architecture has to operate, because that is where the failures are. Control that stops at the visible action is looking at the shadow. ICA governs and seals the decision that cast it.

From governance you assert to governance you can prove.

A compliance checklist is a claim about the past: on this date, we believed these controls were in place. It ages the moment it is signed, and it rests on the honesty of the person assembling it.

An ICA certificate is a different object. It is backed by continuous enforcement and a chain of signed receipts. It does not say we intended to control this. It says here is the record that we did, verify it yourself. When a regulator, a customer, or a board asks the hard question, the answer is not a slide. It is cryptographic evidence that holds up without anyone taking your word for it.

The market is converging on the right questions. Who owns the use case after launch. Who can pause it. Who keeps the evidence current. Those questions are correct, and they all have the same hidden dependency: none of them can be answered honestly unless the system itself records who did what, and can prove it later. Ownership without an instrument is a name on a risk register. An operator without a control plane is a title.

The frameworks are the map, and the map is getting better every quarter. What has been missing is the machine underneath it: the layer that authorizes every action, contains what breaks, and seals proof that cannot be rewritten. That is ICA.

To see the machine, not the map

ALEETH builds the Institutional Control Architecture: the control, certification, evidence and verification layer for autonomous AI. To read how ICA maps to the standards you already answer to, or to see a certificate verified live, visit aleeth.com.