BRIEFING · ALEETH · 2026

Stateless Is Not Governed

On July 28 the protocol every AI uses to reach real systems resets. A public scan of 4,356 servers found one ready. Readiness was never the hard part.

SHANE SCHRECK FOUNDER, ALEETH JULY 2026

HTTPS://PUBLICATIONS.ALEETH.COM/STATELESS-IS-NOT-GOVERNED/ PDF

The Model Context Protocol is the port through which AI agents reach real systems. One protocol, any tool, no custom integration for every pairing. In under two years it went from one vendor's proposal to neutral stewardship under the Linux Foundation, and it is now in every major AI client. If an AI at your company touches a database, a CRM, a payment rail or a build pipeline, there is a good chance it is speaking MCP to get there.

On July 28 that protocol goes stateless. The initialization handshake disappears. Servers stop holding per-client session state. It is the correct change, and it forces nearly every server on the protocol to be rebuilt.

A public scanner swept the official registry on July 12 and probed 4,356 reachable servers for readiness against that date. It found one.

That one was ours. What follows is the argument for why that is a smaller thing than it sounds, and why the number that should actually worry you is a different one entirely.

Why being ready is not the achievement

The readiness test asks three questions. Does the server support the new discovery step. Does it validate the new routing headers. Does it work without a session identifier.

Every one of those is plumbing.

Not one of them asks whether the server controls what the AI actually does. Not one asks whether it can prove, afterward, what it permitted, what it refused, and on whose authority. A server can pass all three checks and still hand an agent unlimited reach into a production system with no record that anything happened.

So when a scanner reports one ready server out of thousands, what it is really measuring is how much of this field has not yet noticed that the ground is moving. Being early to a transport change is a matter of engineering attention. It is not virtue, and it should not be sold as one.

The number that should worry you

The 2026 public security research on this ecosystem is not subtle. Independent audits scanning MCP servers report security findings in roughly two of every three examined. The failure modes repeat: over-privileged tools, prompt injection reaching real actions, tool definitions that change between calls, name shadowing and typosquatting inside registries, credentials and sensitive parameters landing in logs, and almost universally no verifiable audit trail at all.

Directories have grown to meet the volume. One tracked close to thirty-seven thousand servers in mid-2026 and assigns letter grades for quality and safety. That is useful, and it is not a certification. It grades what a crawler can see from outside a public repository.

Here is the part that should be uncomfortable for everyone in this market, including us. The certifications that exist in 2026 certify people. There is a well-regarded credential for MCP security practitioners, with labs and a practical exam. There is no equivalent for the servers themselves. No neutral body issues a pass or fail against a published bar. Enterprises are adopting these servers right now, absorbing unpriced risk, with nothing to point at when someone asks them to justify the decision.

A tool call is not a chat message

This is the reframe worth keeping even if you discard everything else here.

When an agent produces text and gets it wrong, you get a bad paragraph. When an agent makes a tool call and gets it wrong, it moves money, rewrites

a record, provisions access, or halts machinery. The server on the other end is the only thing deciding how far that authority extends.

Most servers in circulation do not decide, because nothing in the protocol requires them to. An MCP server either governs the authority it hands to machines, or it launders that authority for whoever connects to it. There is no third kind of server.

So we published the bar

The ICA-Certified MCP standard states, as pass or fail requirements, what a server must do before it may claim it governs the authority it is handed. Twenty-eight criteria across six domains, each naming the evidence it demands.

Risk state keyed to an authenticated principal rather than a client-supplied identifier. A cumulative, non-decaying allowance so a patient attacker probing one polite request at a time still trips it. Quarantine that latches until a human clears it, never a timer. Signed, hash-chained evidence with no raw sensitive parameter in the ledger. Fail closed: no receipt, no response, even after the tool has already run. Policy bound by content hash, not by a label.

Then we scored our own server against it and published what we failed.

The current result is twenty-one pass, four partial, zero fail. Reaching zero fail meant closing four criteria that had been failing in public since the standard was first published. The four that remain partial are honest partials, mostly

key custody and control-plane work, and they are printed as incomplete rather than argued into passes.

The warning we could have quietly dropped

Our independent readiness grade is an A. Seven checks pass, none fail, and one returns a warning. We print the warning.

The scanner reports that our server declares a capability the new specification removed. We do not serve that method. We serve the replacement the specification defines, and the library beneath us sets the old flag precisely because we serve the new mechanism, which the scanner then reads as the old capability. On the substance we are conformant. On the scoreboard we carry a warning, and it survives re-verification.

Showing the A and skipping the asterisk would have been trivial. A standard that grades its author gently is marketing. We sell the difference between a claim and a receipt, so it has to hold for our own scorecard first.

Do not take our word for any of it

The scanner is public and we do not own it. One line reproduces our grade on your own machine with a read-only credential. Without a credential it sees an authenticated endpoint: an anonymous request returns 401 with a correct challenge and discoverable protected-resource metadata. That is the right

posture for a governance product. We are grade A ready, and we do not leave our capabilities open to anonymous enumeration.

Every governed call our server makes is sealed with a signed receipt, and any receipt verifies against a public ledger without contacting us.

What we are actually asking

Take the standard and score your own server against it. Honestly, including the criteria you fail. That costs nothing and it will tell you more than a directory badge will.

And if you want the result to mean something to your customers, the certification engagement now exists. It is evidence-based against the published standard, every verdict seals to a receipt anyone can verify, issuance is a human decision that cannot be automated away, and the mark is kept by conduct on a public register that never deletes an entry. One address, no form: security@aleeth.com.

And if you run agents against production systems and cannot answer, with evidence, what those agents were permitted to do last Tuesday and what they were refused, you do not need a scanner to find your finding. That is the finding.

On July 28 this field rebuilds its plumbing. The plumbing was never the hard part. The hard part is that an agent reaching a real system is exercising institutional authority, and almost nothing in this ecosystem is currently deciding how much.

ALEETH›

The Institutional Control Architecture. The control, certification, evidence and verification layer for autonomous AI. Beneath the act. Above the machine. Impossible to bypass.

NOT PITCHED. NOT PROMISED. PROVEN.

SHANE SCHRECK · FOUNDER · ALEETH

AI HAS MET ITS MATCH. TRUTH.™