

REPORT · Q3 SECURITY RE-BASE · STATE OF MCP

Measured consistently, the agentic web did not get safer this quarter

The first fresh deep scan since Q2. Where the method held steady, MCP security slipped or stalled: risky-code prevalence ticked up, and four in five sensitive servers still ship no authentication. Telling maintainers barely helped.

ONE NUMBER IS A RE-BASE, NOT A DELTA: READ BEFORE THE COMPOSITE

The supply-chain scanner changed method this quarter, so its exposure rate is not comparable to Q2 and its apparent improvement is a measurement artifact, not a real-world gain. Because that pillar feeds the composite Trust Index, we are **withholding the composite this quarter** rather than let a re-based pillar inflate it. The two pillars measured the same way as Q2 (code safety and identity) are reported as clean deltas below.

BY THE NUMBERS

37.5% ship a risky pattern up from 35.4%	78.6% sensitive, no auth flat from 79.8%	67 high-severity servers up from 55	2 of 47 warned servers fixed 6 got worse
--	--	---	--

Security & supply chain as of 2026-09-05, identity as of 2026-08-29. 2,408 servers deep-scanned, firsthand and aggregate-only. Live at majorlabs.co/security.

FINDING 1 · CODE SAFETY

More servers scanned, a higher share at risk

We deep-scanned **2,408** maintained servers, 407 more than Q2. The share carrying at least one risky code pattern rose to **37.5 percent** from 35.4, and the count at high severity climbed to **67** from 55. Growth is not diluting the risk; it is adding to it. The dominant pattern, by a wide margin, is server-side request forgery surface: 818 of the flagged servers.

Deep-scan severity, 2,408 servers

Each server counted once at its highest severity.



● Low **1,506** ● Elevated **835** ● High **67**


FINDING 2 · IDENTITY

Four in five sensitive servers still have no authentication

Of 2,330 servers scanned for identity posture, 733 take a sensitive action in their code: running shell commands, writing or deleting files, writing to a database, or sending mail or moving money. Of those, **576 show no authentication in source**:

78.6 percent, essentially unchanged from 79.8 in Q2. Narrow the 733 to the 350 that are reachable over the network (HTTP transport) and 267 show no authentication: 76.3 percent. This is the authorization gap agentic commerce keeps colliding with, and a full quarter of attention has not moved it. A server that takes sensitive actions and checks nobody's identity is the failure mode a signed mandate is supposed to prevent.

Correction, September 18, 2026: this finding originally said the 733 servers "touch sensitive data over the network." The scanner does not measure data sensitivity. It flags sensitive actions in source code, on any transport, and 350 of the 733 are network-facing. The counts and percentages are unchanged. What static analysis can and cannot show is set out at /identity.




Of every server that takes a sensitive action, only 157 authenticate and 576 do not. The ratio has not improved. Identity is still the pillar the ecosystem is not building.

FINDING 3 · REMEDIATION

We tracked 47 high-risk servers. Two fixed it. Six got worse.

This is the finding no one else can make, because it needs a cohort tracked over time. In Q2 we froze a cohort of 47 high-risk servers, contacted the first nine maintainers through coordinated disclosure, and re-scanned the same 47 this quarter. **Two fully fixed the issue. Four reduced it. Thirty-five were unchanged. Six got worse.** Net, the cohort did not improve. Disclosure is necessary and it is not sufficient; a warning that a stranger has to act on, unpaid and unenforced, mostly does not get acted on, and most of the cohort has not yet been warned at all. Correction, September 11,

2026: this finding originally said all 47 were contacted; only the first disclosure batch of nine had been sent. The scan numbers are unchanged. Update, September 18, 2026: the figures here are the August 29 re-scan this report was written from. The next re-scan, on September 5, moved one server from unchanged to worse: 34 unchanged and 7 worse. The current count is at /security.



SCORECARD

The two clean pillars, Q2 to Q3

MEASURED THE SAME WAY BOTH QUARTERS			
PILLAR	Q2	Q3	DIRECTION
Risky-pattern share	35.4%	37.5%	worse
High-severity servers	55	67	worse
Sensitive, no auth	79.8%	78.6%	flat
Servers deep-scanned	2,001	2,408	wider
Supply-chain exposure	75.2%	36.2%*	re-based
Composite Trust Index	38	withheld	see note

*Supply-chain method changed this quarter; the two figures are not comparable and the pillar is excluded from any delta claim.

Method, and why the composite is withheld

Every figure is a firsthand, read-only deep scan of maintained public MCP servers, aggregate-only, no per-server findings published. Code safety and identity were measured with the same instrument as Q2, so their movement is a real delta. The supply-chain scanner changed method this quarter, so its exposure rate is a re-baseline: 75.2% → 36.2% reflects how we measure, not what maintainers did.

The composite Trust Index equal-weights the three pillars. With one pillar re-based, the composite would move for a reason that is not real, so we withhold it for one quarter rather than publish a number we would have to caveat. From next quarter, all three pillars share a stable method and the composite resumes as a clean series.

Cite: Major Labs (2026). *The State of MCP: Q3 Security Re-base*. majorlabs.co/reports/state-of-mcp-security-q3.