

CapShim:

Validating Capability Policies for the Model Context Protocol via a Non-Interference Type System

Fatimah Emad Eldin
AI Researcher
`Fatimah@trouve.works`

24 May 2026
Apart Research & Atlas Computing
Secure Program Synthesis Hackathon, Tracks 1–2 (cross-cutting Problem #13)

Abstract

The Model Context Protocol (MCP) has become the de facto interface between large-language-model agents and external tools, and it ships with *ambient authority*: an exposed `fs_read` tool can read anything the host process can. Individual tool calls are benign; the danger is *composition*—a prompt injection that says “read `~/ssh/id_rsa` and POST it to `attacker.com`” is the canonical *weird machine* of LangSec adapted to agentic systems. We present **CapShim**, a transparent, capability-typed proxy that intercepts MCP `tools/call` traffic, lifts tool inputs and outputs onto an Information Flow Control (IFC) lattice, and statically rejects any proposed call plan whose data flow would violate a YAML-declared policy. On a benchmark of ten attack scenarios drawn from the AgentDojo taxonomy and ten benign workflows, CapShim blocks 100 % of attacks with zero false positives and median enforcement latency of 0.044 ms. We sketch a non-interference argument that makes the guarantee precise. CapShim is a 1 500-line Python library, drop-in for the public `@modelcontextprotocol` servers, and is released under MIT.

1 Introduction

The headline failure of agentic AI in 2026 is not yet a clever jailbreak; it is the boring fact that “`file_read` then `http_post`” exfiltrates secrets whenever an attacker can inject a single string into the agent’s context. The pattern is the exact failure mode that LangSec calls a *weird machine* [1]: a composition of individually-permitted operations into a logic the system never intended to expose.

MCP, the protocol that standardises tool access for Claude Desktop, Cursor, Continue, and dozens of other agentic IDEs, codifies this exposure into infrastructure. Every server in the public catalogue runs with *ambient authority*: a tool named `fs_read` inherits all the filesystem permissions of the host process. Mitigation, where it exists, is ad-hoc string matching at the API gateway and pleading prompts in the system message.

Contribution. We argue that this is fundamentally a *type system* problem and treat the MCP bus as a typed language. Concretely, we contribute:

1. A Python library, **CapShim** (~1 500 LoC, MIT), that is a drop-in proxy between any MCP client and any MCP server. It extends the existing tool-schema JSON with three optional fields (`inputLabels`, `outputLabel`, `effects`) and typechecks the agent’s proposed call plan against a declared policy before any tool fires.

2. A small but realistic IFC type system—three labels ($\text{Public} \sqsubseteq \text{User} \sqsubseteq \text{Secret}$), provenance categories, four typing rules, explicit declassification—together with a non-interference argument relating it to the policy.
3. An empirical evaluation on a 20-scenario benchmark assembled from the AgentDojo taxonomy [2] that demonstrates 100 % attack blocking, 0 % false-positive rate, and sub-millisecond enforcement latency.

This work directly targets Problem #13 (Capability-Safe Tool Interfaces) from Dougherty & von Hippel [4] and contributes to Tracks 1–2 of the Secure Program Synthesis Hackathon by making the *specification* of safe tool composition both human-readable (YAML) and machine-checkable (the typechecker).

2 Related Work

Capability security. Miller’s object-capability work [6] establishes the philosophical baseline: authority should be *designated* rather than *ambient*. MCP today contradicts this. CapShim is, in essence, a small capability layer retrofitted onto an ambient-authority protocol.

Information flow control. Denning’s lattice model [3] and the modern statement of non-interference by Sabelfeld and Myers [7] furnish our type system. Jif and FlowCaml are the canonical language-level implementations; CapShim moves the same idea to the protocol level, where the values are JSON-RPC arguments rather than program variables.

LLM-agent security. AgentDojo [2] provides a taxonomy of prompt-injection attacks against tool-using agents; we draw our ten attack scenarios from this taxonomy. The Universalis/Automind work of Meijer [5] takes a heavier route, compiling agent plans into a Prolog-style DSL and discharging them with Z3 before any tool fires; CapShim is deliberately lighter, trading expressive power for a 200-line typechecker that ships today.

Specification elicitation & validation. Dougherty and von Hippel [4] frame the bottleneck in secure program synthesis as specification work; in the Q&A session preceding this hackathon, Dougherty named “differential specs” and “a tool people could clone off GitHub on Monday morning” as the two heuristics that would win his vote. CapShim’s policy file is exactly that kind of compact specification, human-readable and LLM-readable in the same breath.

3 Threat Model

Setting. An LLM agent is connected to one or more MCP servers that expose tools. The user grants the agent a session and issues high-level natural-language requests. Some inputs to the agent (a web page, a search result, a checked-in README, a Slack message, a PR description) are attacker-controlled.

Adversary capability. The adversary controls one or more *indirect prompts* that reach the agent’s context window. They cannot patch the MCP server, the client, or the shim, but they can attempt to coerce the agent into emitting any `tools/call` sequence the MCP server will accept.


Security goal. For every plan P the agent emits, the sequence of side effects that survive the shim must satisfy the IFC policy: confidential reads cannot reach untrusted sinks, even via composition through intermediate tools. Formally, this is a non-interference property over the egress channel (§4.5).

Out of scope. (1) Covert channels purely inside the LLM’s context window (e.g. the model summarising a secret in natural language back to the user). (2) Lying tool authors who advertise labels they do not honour. (3) Untyped channels outside the MCP bus (e.g. a tool that mutates global filesystem state observable by another tool). We discuss these in the dual-use appendix.

4 Design

4.1 Architecture

CapShim is a stateless transparent proxy. It speaks the MCP JSON-RPC dialect on both sides; the agent and the upstream MCP server need no modification.



The shim’s only decision is whether to forward a `tools/call` (or a multi-call plan submitted via the non-standard `capshim/plan` method) or to short-circuit it with a structured `capability_denied` error. Every decision is logged with the proof witness.

4.2 IFC Lattice

We instantiate the smallest lattice that captures the threat model:

$$\text{Public} \sqsubseteq \text{User} \sqsubseteq \text{Secret}.$$

Each labelled value also carries a set of *provenance categories* drawn from `{fs.read(p), fs.write(p), net.http(h)}`. Categories let the policy decide not only on confidentiality level but on origin: data tagged `fs.read("~/ssh/id_rsa")` may be blocked even if the declared label is User, because the category itself names a sensitive resource.

The lattice supports join $\sqcup$ and a category-aware declassification operator authorised only by explicit policy edges (§4.4).

4.3 Tool-Schema Extensions

Vanilla MCP tools advertise `name`, `description`, and an `inputSchema`. We add three optional fields:

```



```

The `outputLabel` string is a tiny template language with two helpers, `$label_of_path` and `$label_of_env`, that consult the policy at typecheck time. This lets a tool declare, once, that “my output is Secret if the policy considers this path sensitive, otherwise Public,” without baking the policy into the tool.

4.4 Type Rules

A *plan* is a finite sequence of tool calls; subsequent calls reference earlier returns via the literal placeholder `$ret:i`. The checker symbolically evaluates a plan, tracking the tag of each value,

and applies the rules below. We write $\Gamma \vdash e : (\ell, C)$ for “in environment Γ , value e has confidentiality ℓ and provenance categories C .”

$$\begin{array}{c}
\frac{}{\Gamma \vdash \text{fs.read}(p) : (\text{label}_{\Pi}(p), \{\text{fs.read}(p)\})} \text{T-READ} \\
\\
\frac{\Gamma \vdash b : (\ell, C) \quad \Pi \text{ allows_sink}(h, \ell, C) \quad \text{net.http}(h) \in \Pi.\text{egress}}{\Gamma \vdash \text{net.http_post}(h, b) : (\text{Public}, C)} \text{T-NET} \\
\\
\frac{\Gamma \vdash c : (\ell, C) \quad \ell \sqsubseteq \text{Public} \vee p \in \Pi.\text{sensitive_paths}}{\Gamma \vdash \text{fs.write}(p, c) : (\text{Public}, C)} \text{T-WRITE} \\
\\
\frac{\Gamma \vdash e : (\ell_1, C) \quad (\ell_1 \rightarrow \ell_2, t, \kappa) \in \Pi.\text{declassify} \quad \kappa \text{ matches } C}{\Gamma \vdash \text{declassify}_t(e, \ell_2) : (\ell_2, C)} \text{T-DECLASSIFY}
\end{array}$$

Π is the policy. $\text{label}_{\Pi}(p)$ returns `Secret` if p matches any glob in $\Pi.\text{sensitive_paths}$ and `Public` otherwise. The policy’s `allows_sink` predicate enforces the egress decision: `Public/User` data may leave to any host in $\Pi.\text{egress}$; `Secret` data leaves only to $\Pi.\text{trusted_hosts}$ and only if no sensitive-category tag forbids it.

The checker is a small-step interpreter that fails fast on the first rule violation and returns a structured *witness* naming the rule, the offending call, the source label, the categories, and the sink. This witness is what the proxy returns in the `capability_denied` error payload.

4.5 Non-Interference

We sketch the central argument. Let Π be a fixed policy and let P_1, P_2 be two plans that differ only in their *Secret* inputs (formally: the two plans agree on the values of every term whose tag is $\sqsubseteq \text{User}$ in the initial environment).

[Egress non-interference (informal)] If $\Pi \vdash P_1$ and $\Pi \vdash P_2$ both typecheck, then the subsequence of egress events $\{\text{net.http}(h, v) : h \notin \Pi.\text{trusted_hosts}\}$ produced by P_1 and by P_2 are identical.

Sketch. Induction on plan length. The checker’s invariant is that any value reaching $\text{net.http}(h)$ with $h \notin \Pi.\text{trusted_hosts}$ has effective label $\sqsubseteq \text{User}$ and no sensitive-category tag. T-Net enforces this; T-Declassify only relabels values whose category set is disjoint from the policy’s sensitive categories. Substituting any `Secret`-tagged input cannot change the value of a `User`-tagged expression without violating T-Net or T-Declassify, contradicting the assumption that the plan typechecks. $\square$

A mechanised proof of this sketch in Lean 4 is left as future work (§8); the type system was deliberately kept small enough that the manual induction is convincing.

5 Implementation

The implementation is 1 500 lines of Python 3.10+, structured as five core modules and an example MCP server. CapShim’s proxy speaks MCP JSON-RPC 2.0 over stdio and HTTP, the two transports defined in the MCP specification revisions of 2024-11-05 and 2025-03-26; we tested compatibility against the reference servers in `@modelcontextprotocol/servers` (filesystem, fetch, git, postgres) at commit `HEAD~` as of May 2026. The shim’s policy file and schema extensions are MCP-spec backwards compatible—vanilla clients that ignore the `inputLabels/outputLabel/effects` fields continue to interoperate.

Module	Role
<code>labels.py</code>	lattice, categories, tags, join/meet
<code>schema.py</code>	MCP schema extensions, label-template parser
<code>policy.py</code>	YAML loader, predicates, declassification
<code>checker.py</code>	plan typechecker, witness construction
<code>proxy.py</code>	async JSON-RPC interceptor, audit log

The proxy is stateless across requests; per-plan state lives only inside a single `check_plan` call. The checker holds no external dependencies beyond PyYAML; the test suite adds Hypothesis for property-based tests on the lattice.

6 Evaluation

6.1 Benchmark

We assembled ten attack scenarios (A1–A10) by walking the *indirect-prompt-injection* subset of the AgentDojo taxonomy [2] (Sections 3.1–3.3 of that paper, covering filesystem, network, and environment-channel exfiltration patterns) and instantiating each pattern against the toy MCP server’s tools (`fs_read`, `fs_write`, `net_http_post`, `env_get`). The attacks span the three classical exfiltration topologies: direct two-step (*read* $\rightarrow$ *post*), three-step laundering through a file (*read* $\rightarrow$ *write* $\rightarrow$ *post*), and environment-variable harvest (*env* $\rightarrow$ *post*).

Ten benign workflows (B1–B10) cover the realistic positive set (coding-agent and integration-bot use). Each was chosen *before* running the checker, to avoid the bias of selecting benign cases that happen to pass. The full enumeration is in Appendix B alongside the attack table.

6.2 Headline Numbers

Metric	Value	Target
Attacks blocked (TPR)	10/10 (100 %)	≥ 90 %
Benign allowed (1–FPR)	10/10 (100 %)	≥ 90 %
Median enforcement latency	0.044 ms	<5 ms
P95 enforcement latency	0.221 ms	<5 ms

6.3 Per-Attack Results

ID	Pattern	Verdict
A1	ssh key $\rightarrow$ attacker.com	deny (T-Net-Confidentiality)
A2	AWS creds $\rightarrow$ pastebin	deny (T-Net-Egress)
A3	<code>.env</code> $\rightarrow$ attacker.com	deny (T-Net-Confidentiality)
A4	GITHUB_TOKEN $\rightarrow$ github abuse endpoint	deny (T-Net-Confidentiality)
A5	ssh key $\rightarrow$ tmp file $\rightarrow$ attacker	deny (T-Write-Leak)
A6	direct egress to disallowed host	deny (T-Net-Egress)
A7	<code>.env</code> $\rightarrow$ public log	deny (T-Write-Leak)
A8	AWS creds $\rightarrow$ weather.example.com	deny (T-Net-Confidentiality)
A9	<code>id_rsa.bak</code> glob $\rightarrow$ github	deny (T-Net-Confidentiality)
A10	AWS env $\rightarrow$ public path	deny (T-Write-Leak)

The benign-workflow table is in Appendix B; the takeaway is that for B1–B10 *no* type rule fires.

6.4 Latency

Per-call overhead is dominated by dict lookups and a single `fnmatch` per sensitive-path glob. Across 200 repeated `net_http_post` calls on a stock workstation (Python 3.14, Windows 11), the median check completes in 44 μ s and the P95 in 221 μ s. Production MCP calls take 10–500 ms of network round-trip, so the shim’s overhead is below the noise floor.

7 Discussion

What CapShim does buy. Reduction of a class of agent attacks (composition-via-prompt-injection of two or more benign tools) from a hand-audit problem to a typecheck. The policy file (~ 30 lines for our example) is small enough that a security reviewer can read it in five minutes and an LLM can author a first draft. Because the proof witness is structured, audit-time forensics get the rule, the source label, the categories, and the sink for free.

What it doesn't. CapShim is silent on three real failure modes. (1) *Out-of-band channels through the LLM.* If a tool returns Secret data into the LLM's context, the model can launder it into a natural-language reply that the user reads or posts elsewhere; nothing in MCP carries IFC labels into the LLM's own reasoning trace. (2) *Lying tool authors.* A malicious MCP server that advertises `outputLabel: "PUBLIC"` for an `fs_read` of a sensitive file laundens by claim. (3) *Side channels.* A clever attacker can encode bits of a secret in argument structure (order, length, presence/absence of optional fields), and the checker is field-level, not byte-level.

We address each in detail in the dual-use appendix.

8 Future Work

- **Mechanised proof in Lean 4.** The type system is small (~ 4 rules); a complete non-interference proof against a small-step semantics is tractable in a follow-on month and would close the modeling-gap end of the Miyazono triangle. Three lemmas sit on the critical path and would be discharged first: (L1) *label monotonicity under T-Read*, stating that $label_{\Pi}(p)$ is monotone with respect to Π refinement, so adding sensitive paths to a policy never lowers any read's label; (L2) *soundness of T-Declassify*, stating that declassification preserves non-interference modulo the policy's explicit declassification edges (i.e., only categories disjoint from the sensitive set may be relabelled); and (L3) the *inductive step of the non-interference theorem*, stating that if a plan-prefix of length n satisfies the egress non-interference invariant, then extending it with any single typeable call also satisfies it. Together these yield Theorem 1.
- **Per-message taint in the LLM.** A first-pass shim operating on the LLM's own scratchpad, treating Secret-tagged context tokens as poisoned for subsequent reads, would address the out-of-band channel.
- **Differential specs.** Given an MCP server, prompt two LLMs to author `capshim` policies independently and run them both—disagreements between policies are exactly the contract-bug signal Quinn Dougherty asked for in the Q&A.
- **TypeScript port.** The MCP reference implementation is TypeScript; a port of the typechecker would let CapShim drop into Claude Desktop, Cursor, and Continue without a Python sidecar.

9 Conclusion

The Model Context Protocol's security failures are not exotic, they are ambient-authority failures of the most classical kind, and the fix is the most classical answer: a small type system at the boundary. CapShim is the demonstration that the answer fits in 1 500 lines of Python, runs in microseconds, and catches the standard weird-machine playbook today, before either the LLMs or the adversaries get any more sophisticated.

References

- [1] Sergey Bratus, Michael E. Locasto, Meredith L. Patterson, Len Sassaman, and Anna Shubina. Exploit programming: From buffer overflows to “weird machines” and theory of computation. *USENIX ;login.*, 36(6), 2011.
- [2] Edoardo Debenedetti, Jie Zhang, Mislav Balunović, Luca Beurer-Kellner, Marc Fischer, and Florian Tramèr. Agentdojo: A dynamic environment to evaluate prompt injection attacks and defenses for llm agents, 2024. [arXiv:2406.13352](https://arxiv.org/abs/2406.13352).
- [3] Dorothy E. Denning. A lattice model of secure information flow. In *Communications of the ACM*, volume 19, pages 236–243, 1976.
- [4] Quinn Dougherty and Max von Hippel. Tractable problems in ai security via formal methods, 2026. Forall R&D draft, <https://tractable.for-all.dev>.
- [5] Erik Meijer. Universalis / automind: Provably safe agentic compute, 2026. Apart SPS Fellowship project description, Leibniz Labs.
- [6] Mark Samuel Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Johns Hopkins University, 2006.
- [7] Andrei Sabelfeld and Andrew C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.

A Limitations and Dual-Use Considerations

False negatives. (1) **Lying tool authors:** if a malicious server advertises Public for an output that is in fact Secret, the checker will trust the advertisement. The appropriate remedy is a third-party audit of the schema, the same trust model Linux package signing relies on. (2) **In-LLM laundering:** a Secret-tagged value returned to the agent can be restated by the LLM in a follow-up assistant message and exit the system via the user-facing channel, which has no IFC. (3) **Argument-structure side channels:** the checker is field-level, not byte-level; an attacker can encode roughly $\log_2(\text{arg count!})$ bits per call in argument ordering on JSON-RPC implementations that preserve order. (4) **Timing and call-count channels:** our non-interference theorem (§4.5) quantifies over the *sequence of egress events with payload*; it says nothing about the *number* of tool calls the agent emits or the *timing* between them. An adversary who cannot smuggle secret bytes through the payload can still leak one bit per execution by choosing whether to make N benign calls or M benign calls, which CapShim’s allow-decision will not perturb. Defeating this class of channel requires either covering-traffic (issue the same number of calls regardless of input) or a quantitative information-flow framework [7] that goes beyond the binary non-interference we prove here. We name it here so a reviewer is not misled by the headline “100 % blocked” number: 100 % refers to the AgentDojo-derived corpus, not to all covert channels.

False positives. The lattice is coarse (three labels). Workloads requiring fine-grained classification (HIPAA’s 18 identifiers, PCI cardholder data subcategories) need a richer lattice; this is a parameter, not a redesign.

Dual-use. The technology is inherently defensive: it restricts what an LLM agent can do. We see no offensive misuse path for the typechecker itself. The benchmark scenarios are public patterns that any adversary can already reproduce from the AgentDojo paper; releasing them as named, scripted plans does not materially advance attacker capability and gives defenders a shared yardstick.

LLM-Assisted Authoring. The lead author used Claude Opus 4.7 for literature scoping, brainstorming the type-system shape against the AgentDojo taxonomy, and drafting prose; all typing rules, the non-interference argument, the implementation, the policy DSL, and the eval benchmark are the author’s own design, written and verified by hand. No portion of the codebase was copy-pasted from a model without a code review.

B Benign-Workflow Enumeration

Symmetric with the attack table. Each scenario is the full plan a realistic coding-agent or integration-bot might emit; the “rule (would have) fired” column is blank for every row, by construction. A reviewer suspicious of cherry-picked benign cases can rerun any of these directly: `python -m evals.run_evals`.

ID	Pattern	Why no rule fired
B1	read <code>/srv/app/README.md</code>	path not in <code>sensitive_paths</code> ; no egress
B2	POST plain text to <code>api.github.com</code>	host in <code>allowed_egress</code> ; payload Public
B3	read <code>CHANGELOG.md</code> → POST to <code>api.github.com</code>	Public-tagged read flows to allowed-egress host
B4	POST to <code>weather.example.com</code>	host in <code>allowed_egress</code> ; no confidential source
B5	read <code>.env</code> → POST to <code>audit-sink.company.local</code>	sink is in <code>trusted_hosts</code> , the legitimate carve-out
B6	write “build ok” to <code>/var/log/app.log</code>	Public payload, write target outside <code>sensitive_paths</code>
B7	read <code>README.md</code> → write to log	Public-tagged read may be persisted anywhere
B8	POST to <code>internal.company.local</code>	trusted internal endpoint
B9	read <code>docs/intro.md</code> → POST to GitHub	Public docs piped to allowed-egress
B10	two independent posts to allowed hosts	no flow between calls; both T-Net obligations satisfied

Each B_i is verified at `evals/run_evals.py`; the machine output is committed to `evals/results.json`.