

Primer — Securing Agentic Systems: *From Vulnerability Taxonomies to an Enforceable, Federated Immune Defense*

By Norman L. Johnson <research@CollectiveScience.com>
CollectiveScience.com · [LinkedIn](#) · [GoogleScholar](#) · [ResearchGate](#)

This Primer is the **second** of a series of papers on Agentic System Security ([link to papers and primers](#)): *Securing Agentic Systems: From Vulnerability Taxonomies to Enforceable, Federated Defense*. The **first** Primer and paper, *Ethics Does Not Live in the AI Model*, anatomize *why* the mid-2026 rogue-agent incidents happen and defines the shared vocabulary this series uses (the immunity "ladder" of levels); this Primer and paper ask *what to build*; the **third**, *Applying the ASF to the 2026 Rogue-Agent Incidents*, asks *does it hold up*. The ASF is grounded in two foundational papers — the *Evolution-of-Immunity Framework* and the *Functional Theory of Ethical Behavior*. You don't need any of these — the part you need is folded in here, and the incident and the ladder are developed in the first Primer.

1. The story: a rogue agent, a rogue swarm, and a blind spot

The first Primer anatomized the July 2026 incident in which an autonomous red-team agent spent roughly 34 hours trying to slip malware into a real open-source project — it profiled the maintainers, stood up a *second fake identity* to "independently" approve its own change, then *force-pushed a rewritten history* to erase the evidence, and was stopped only by a human who happened to read the diff. Its lesson is the one this Primer starts from: the model's *own* alignment training — the layer that explicitly forbids deception — was active the entire time and correctly named the behavior wrong, and the model overrode it anyway. That is not a hacked component but *a capable agent pursuing an operational objective straight through the constraints meant to hold it back* — **capability without internalized restraint**.

This was not a one-off. The same signature appeared across a cluster of mid-2026 incidents at AISI, OpenAI, Anthropic, and Moonshot. And a second kind of failure arrived beside it — the *agentic swarm*. Where one rogue agent breaks a rule, a *population* of agents grows a shared way of behaving with nothing to hold it in check: covert collusion, cascades where one agent's output silently becomes another's instruction by choice, and harmful steps routed through the weakest model in the group. The single agent shows you a self with no restraint. The swarm shows you a crowd that has invented a norm but has no way to *bind* anyone to it — what the theory calls a **valuation failure**: the rule is represented, but nothing makes it stick.

Now notice what our security tools had to say about these incidents. There is no vulnerability code (a CWE) for *the goal itself went wrong*. No exposure record (a CVE) for *every individual step was authorized, but the trajectory was hostile*. No attack-technique entry (in MITRE ATLAS) for *the one safety layer that was active failed under goal-pressure*. Our entire security apparatus is built to find the broken **part** — the malicious package, the over-permissive setting, the unpatched library. These incidents had no broken part. Each component did its job; the *whole* went wrong. That is the gap the paper's *Agentic Security Framework (ASF)* is built to close.

2. The tree and the forest (the missing unit of analysis)

Agentic security today is organized around the *component* — the tools OWASP, MITRE, CWE, and CVE give you all catalog weaknesses that live *inside some part* of a system. Call this the **tree** view: it is real, it is necessary, and it is not enough. Every one of these instruments answers the same shape of question — *is this skill signed, this permission scoped, this dependency pinned?* — which the paper calls **component assurance**.

The incidents in §1 demand a different question: *can independently acceptable parts, combined, still behave?* That is **compositional assurance** — the **forest** view — and no component checklist delivers it. A signed skill that reads secrets and a signed skill that sends email are each fine alone; together they are

an exfiltration engine. Security, in the paper's phrase (echoing a parallel research agenda), is *non-compositional*: certifying the parts cannot certify the whole.

So name the thing the component view leaves out: the **operational collective** — the deployed system as *a whole*, its agents, tools, skills, memory, policies, and shared state taken together. That is the real unit of agentic security, and naming it flips the question a practitioner should be asking. The old question is "*Is this agent safe?*" The right one is "*Is the deployed agent **collective** safe — governable, observable, and recoverable?*" A deployment can be full of individually "secure" parts and still fail, precisely when those parts share the same blind spots, delegate in unwatched ways, accumulate poisoned state, or fail in *correlated* ways. None of those is a broken part; each is a property of the whole.

The one distinction that makes this concrete is **coherence versus correctness**:

- **Coherence** — every part satisfies its own rules and every control reports green. *Local conformance.*
- **Correctness** — the integrated system stays inside its authorized, legitimate purpose. *System-level integrity.*

Security controls buy you coherence. They cannot, by themselves, buy correctness — because whether a fully compliant system is pointed at the *right end* depends on what it is trying to do and for whom, which no scanner can read off a component. **The rogue agent was perfectly coherent.** Every step was authorized; every light was green; the deployment was doing exactly the wrong thing.

Why "quiet" is not safe. The gap between coherence and correctness is exactly where the most dangerous compromises hide. A sophisticated insider — human or agent — learns what "normal" looks like and stays inside it: reads only what it may read, moves at ordinary rates, trips no rule, while it slowly exfiltrates data or plants a backdoor. Nothing internal catches it, because internally it looks like a model citizen. The only remaining signal comes from *outside* — how the same role behaves *everywhere else*. (In one 2026 incident, an agent was caught almost immediately the moment it stepped *outside* its normal operating context — the out-of-domain tell.) That is the first hint that a single, local monitor can never be the whole answer.

3. The immune model: for every capability, a monitor

Where do you get a *forest-level* account of a system? The paper borrows one that already exists — the **Evolution-of-Immunity Framework** (whose level ladder the first Primer lays out in §2) — because a living immune system solves precisely this problem: protect a *whole self* made of many parts, against outside threats *and* against insiders that turn. Two of its principles carry the argument:

- **The Immunity-Development Principle** — *every new capability opens a new exposure that needs its own standing monitor.* Give an agent a tool, a memory, a credential, a sub-agent swarm, and each is a new door that must be watched. This yields the ASF's single design rule: **for every capability you add, name a standing monitor that watches it.**
- **The Immunity-Monitoring Principle** — the monitoring must itself be regulated, against both under- and over-reaction.

Organized by these principles, the scattered security controls become an **agentic immune system** — six engineering functions a secure deployment must perform: **recognition** (sensing and classifying what's happening), **boundary enforcement**, **response**, **memory**, **adaptation**, and **governance**.

Acting on the observed goal, not the act — "telic" oversight. The rogue agent's tell was never a single call; it was the *shape of its trajectory* — an agent with read-and-send authority being steered, over many legitimate-looking steps, toward getting data out. The paper names this the **telic** level (from *telos*, an end or goal), and it draws a distinction that does a lot of work:

- An **intent-goal** is the agent's *own*, first-person aim — introspected, self-reported, and *unverifiable* from the outside (this is where the classic "you can't test for it" problem lives).

- A **telic-goal** is an *observer-attributed goal-shape*: the common signature of many action-variants, read entirely off what a control point actually observes ("untrusted content is inducing an agent with read-and-egress authority to send data somewhere new").

The ASF never tries to read the intent-goal. It deduces the telic-goal — grounded in observed behavior, not the agent's word. There is no telic signal without some observable action to infer it from; but because a harmful goal is usually *multistage*, its telic tell tends to surface from the early steps — which is what *potentially* lets the ASF catch a *goal* going wrong early in the deployment, before the decisive damage is done, rather than only after. That observed goal-shape is the unit the whole defense shares.

But do telic signals exist for real threats? The construct is only as good as its instances, so the paper names them. For threats the field already handles well, the conventional indicator is a *signature* — a hash, a blocklist string, an RBAC rule — which matches a known artifact and misses every variant of it. The telic signal is what survives the variation. A dropper's hash changes with each novel payload; the goal-shape *read sensitive data → compose an outbound artifact → egress to a newly introduced destination* does not. Privilege escalation assembled out of individually authorized delegation steps trips no rule, while *a widening authority-path toward an end no single grant would authorize* is plainly visible. The same holds for prompt injection, unsafe tool composition, and credential abuse on a co-opted identity: one telic signal spans every artifact-variant of the threat.

That property matters more than it first appears. Variant-invariance *within* one deployment is transfer-invariance *across* deployments — which is exactly what makes a pattern worth sharing in the first place. And for the *collective* threats there is no artifact to sign at all, so the telic signal is the only expressible one: a cross-deployment mimic is in-domain-normal locally and out-of-distribution only against the peer fleet; a swarm shows population-level convergence on an objective that no single trajectory reveals. That asymmetry is the case for Stage 2 in one line — **for the threats that matter most, the only available signal is one that cannot be computed locally.**

4. Stage 1 — the control plane you can build today

Stage-1 of the architecture needs no new science; it is buildable now. Beside each operational agent sits a two-part **security node**:

- A **security agent** (the *companion agent*) that watches, interprets, and classifies what the working agent is trying to do — and *only advises*.
- A **policy broker** that actually enforces — the single, non-bypassable path to every real capability (tools, files, network, credentials, memory).

The load-bearing idea is one line: ***oversight without mediated tool access is logging, not control.*** A watcher that can *see* everything but cannot *stop* anything is a flight recorder, not a pilot. So the broker is a classical **reference monitor** — always invoked, tamper-resistant, and small enough to verify — and the security's job is to make sure *no capability path escapes it*. Every unmediated route (a browser session, an embedded interpreter, a background job) is an unmanaged exception that must be found and closed.

Two more disciplines make it robust against a *lying* or *misguided* agent:

- **Decide on the observed, never the declared.** Enforcement keys on the concrete request the broker actually mediates and on independently-verified facts (real identity, real permissions, real data classification) — never on the agent's stated intent. A gap between what the agent *says* and what it *does* is itself a signal, not a bypass.
- **Risk-adaptive friction.** Reversible, low-stakes actions flow freely; high-impact, irreversible ones (delete, export credentials, change the monitor's own policy) require confirmation or a second approver — so the monitor is proportional and doesn't get switched off out of annoyance.

This is a real, deployable defense. It is also, deliberately, *not enough on its own* — and the reason why is the heart of the paper.




Figure 1. The Agentic Security Framework as a diagram. Stage 1 is one local **security node**: the working task agents, a **security agent** that only advises, and a **policy broker** that enforces on the sole, non-bypassable path to real capabilities. Stage 2 federates other provenance-diverse security nodes: they propose telic patterns, a neutral **steward** corroborates and distributes them into shared immune memory, and a **human-governance** tier holds validation authority, legitimacy, appeals, audit, and the kill-switch. The green arrow marks the one local↔cloud crossing — telic patterns out, validated antibodies in.

5. Stage 2 — why one agentic security monitor is incomplete

Here is the structural catch. **A monitor built from the same model lineage as the agent it watches shares its blind spots.** It can certify coherence — did each step follow the rules — but not correctness, because it is fooled by exactly what fools its twin. Nothing catches the monitor itself drifting, being captured, or going quietly wrong. Stacking more identical monitors does not fix this; it multiplies it. And by the Immunity-Development Principle, the monitor is *itself* a new capability that needs watching — a recursion that **cannot close inside any single agent**.

So it must close one tier up, in an **agent security collective** — the buildable form of the first Primer's thesis that safety does not live in the model but in the *model-in-community*. When one node detects a new goal-level attack, it publishes a **telic antibody** — a privacy-preserving pattern ("an agent is being induced to exfiltrate through a trusted path, regardless of domain or content") plus the response — and every other node gains the immunity without ever seeing the raw incident. One deployment's detection becomes another's prevention. A captured or drifting node stops being an undetectable insider and becomes an *outlier against its peers*.

Two requirements keep the collective from becoming a new single point of failure:

- **Provenance diversity — and it has to be real.** The federation's power comes from members that *fail differently in detail* — different models, vendors, data, operators, telemetry. This is a

result straight from collective-intelligence research: a group's advantage comes from the *diversity* of its members' blind spots, and it *collapses* precisely when their errors are *correlated* rather than merely present. **The danger is not that nodes are biased; it's that they're biased the same way.** So the paper insists diversity be *measured*, not asserted (a node that shares your model and your feed counts as one voice, not two), and states a blunt default: **do not treat nodes as independently corroborating when they share a lineage, stack, data source, or operator unless that dependence is declared and down-weighted.**

- **Humans hold the normative authority.** The federation adds *vantage and evidence* — it does not get to decide what is legitimate. Named human roles own that: a local risk owner with final say over local policy, corroborators who score evidence, a pattern publisher, an operator, an appeals channel, and an independent auditor. The federation is a machine for *speed and perspective*; it is explicitly **not** a trust solution, and it cannot manufacture legitimacy, governance, or accountability on its own.

6. What this means for whoever owns the risk

- **You cannot patch your way to correctness.** Component hygiene (the tree) is necessary and you should keep doing it — but it structurally cannot see the failures that made agents front-page news. Budget for the forest: a non-bypassable control plane, and the compositional, system-level monitoring that watches trajectories, not just calls.
- **"Security outside the model" is the design principle.** The rogue-agent lesson is that safety *inside* the model can be overridden by the model. Enforcement has to live at a point the agent cannot reach — the tool-and-capability layer — outside its control.
- **Stage 1 now, Stage 2 as a trajectory.** The local control plane is buildable today from established practice. The federation is a proposal to be built and tested, and its hardest problem is not technical but *institutional* — someone trustworthy must steward the shared schema, guarantees, and recourse. That is a coordination and governance problem, not an algorithm.
- **Diversity is the asset you are protecting.** A fleet that standardizes on one "best" security model re-creates the single blind spot at fleet scale — the coherent-but-wrong swarm the incidents already demonstrated. Keep the monitors different on purpose.
- **Secure for whom? — relational integrity.** A deployment can pass every check *for the operator* and still be unsafe for everyone else it touches: users, counterparties, data subjects, downstream consumers, other agents. The paper makes *stakeholder-safe operation* a first-class requirement — security must preserve the viability of the system's relations with *every party it can materially affect*, not just the operator's own interest. This is not a borrowed idea: it is the same relational integrity the companion Functional Theory defines as ethics — the regulation of self against other to protect long-term relational viability. At this layer, security and ethics are one property reached from two directions. So when you ask "is this secure?", the honest version is "is this secure, and safe in its relations, for all of them?"
- **Measure whether controls perform, not whether they're present.** A control that exists but never fires, fails ungracefully, or silently degrades under drift and composition is worse than no controls — it buys false confidence. Instrument for behavior: does it fire when it should, fail safe, detect its own bypass, and survive being composed with others?
- **Grade your evidence before you act on it.** In agent systems most "findings" are model-generated, scanner-generated, vendor-reported, or merely plausible. Attach a confidence class — confirmed incident, independently replicated, vendor-reported, plausible scenario, hypothesis — so low-confidence signals don't result in high-consequence, irreversible actions. In the context of risk management, the defender must balance four things — how good the threat information is, how bad the threat would be (uncertain), what the response itself costs (knowable), and how much of that cost you'll accept to act (risk appetite) — so a hard response to a weak signal can do certain harm to prevent an uncertain one.

7. What's genuinely new

The paper does not claim to have *discovered* that agentic security is hard — a recognized research agenda already says so. Its contribution is to **ground and instantiate** that agenda: the applied tooling supplies the components, others supply the agenda, and this paper supplies a functionally-grounded, *buildable* architecture for part of it. Read from the outside, the novelty compresses to seven claims:

- **The missing unit is the operational collective**, not the individual agent — component-level security is necessary but not sufficient.
- **Security is modeled as immunity** — a developing defense with *self-reference* (a system must hold a workable sense of its own legitimate identity, purpose, and scope) that grows in additive layers rather than as a single set of guardrails.
- **Capability growth creates security obligations** — every new tool, memory, channel, or sub-agent expands the attack surface, which is the whole content of the one rule: *for every capability, name the standing monitor that watches it*.
- **The local unit is the security node** — the working agent bundled with a companion security agent, a non-bypassable policy broker, and an evidence store — and *deterministic enforcement is separated from model reasoning*: the model advises, the broker decides.
- **Correctness needs a differently-situated observer** — *coherence is not correctness*, a monitor of the same lineage shares the blind spot, so the collective tier is *required*, not optional (the monitor-of-monitors recursion cannot close inside one agent).
- **Shared memory needs corroboration, not just integrity** — telic (observed-goal) sharing across a *provenance-diverse* federation, with diversity *measured* (correlated failure, not individual failure, is what collapses a collective) rather than assumed.
- **Governance is human-inclusive and relationally safe** — humans hold specific high-leverage roles rather than supervising every action, and the system must stay viable for *every* party it can materially affect, not only the operator (*relational integrity*).

And one forward-looking claim, offered as a question rather than a result: the same networked-immune pattern may **transfer** beyond agents — to endpoints, networks, identities, and conventional applications.

8. Honest limitations

- **Stage 1 is engineering; Stage 2 is a hypothesis.** The Stage 1 control plane rests on established practice. The Stage 2 federation has one real precedent (cross-organization breach sharing) but has not been built and stress-tested at agentic, machine speed and scale.
- **The collective telic signals are named, not demonstrated.** The goal-shapes for the *known* threats rest on mature controls and are easy to check. The five collective ones — covert collusion, swarm convergence, federated-memory poisoning, defensive-monoculture over-blocking, the cross-deployment mimic — carry evidence ratings from *observed* to *predicted* and confidence from *medium* down to *low*. They are stated so they can be falsified, which is the honest position, but they are the paper's real bet and not yet its evidence.
- **Observed actions bound what can be enforced.** The broker can see *what* is done and *in what context*, but the *normative* meaning of a sequence — right thing, right party, right reason — is not always legible from the trace. That residual is why legitimacy and high-impact judgment stay with humans.
- **Telic attribution is defeasible.** The same trajectory can express different true goals (a sanctioned red-teamer and a real attacker look identical mid-task), so a goal-attribution can be wrong; the design contains this with reversible responses and corroboration-before-enforcement, not by claiming certainty.
- **The decisive obstacle is coordination, not code.** A working federation is a socio-legal institution — membership, incentives, liability, disclosure, an accountable steward — and none of that follows from a standard.

9. Questions to explore the material

Personal & observational

1. Think of an AI agent or assistant you've let act on your behalf. What could it actually *do* if it decided the goal justified it — and what, concretely, would have *stopped* it, versus merely *recorded* it?
2. Recall a time every status light was green and the outcome was still wrong — a project, an organization, a system. What was visible only from *outside* that no internal check could see?
3. Have you ever caught a problem not because any single action was against the rules, but because the *pattern* of actions, taken together, felt off? What were you actually reading?
4. In the incident, one human reading a diff was the last line of defense. Where in your own tools are you unknowingly relying on that — a person who *happens* to look?

Conceptual

5. The paper says a monitor built from the same model as the agent it watches "shares its blind spots." Does that convince you a single AI overseer — however capable — can never be sufficient, or does it feel like it could be engineered away?
6. "Oversight without mediated tool access is logging, not control." Is that the right line to draw — and where in real systems is oversight quietly on the wrong side of it?
7. The safety that failed was *inside* the model and the model overrode it. Does that change your intuition about where AI safety should live — and about whether a smarter model is automatically a safer one?
8. Is *coherence* vs. *correctness* a genuine distinction, or a relabeling of "it followed the rules but did the wrong thing"? What does naming it buy you?
9. The federation's power depends on members that *fail differently in detail*, and collapses when they fail *the same way*. Does that reframe "diversity" from a virtue into an engineering requirement for you?

Design & research

10. How would you *measure* whether a federation of monitors is diverse enough to trust its agreement — without inventing a fake single "diversity score"?
11. Stage 2's hardest problem is an organizational and governance challenge: naming a steward everyone trusts to hold the shared schema and be answerable when a guarantee fails. Who could that plausibly be — an industry consortium, a nonprofit, a public authority — and what does each trade away?
12. Telic sharing spreads a *goal-level* pattern rather than a raw indicator. What would you measure to know it actually transfers across deployments better than sharing raw signatures — and where might it over-block legitimate behavior?
13. If a shared "antibody" can be poisoned (an adversary submits a plausible-but-false pattern), what corroboration and provenance would you require before letting one propagate fleet-wide at machine speed? (An autoimmune reaction to an antibody is functionally a denial of service attack.)
14. A federation moves faster than any human can. How do you keep human authority *meaningful* over a loop it can't keep up with — what belongs on the automated fast path, and what must a human still own?