

Agent Security

Agentic Cybersecurity — Module 4

Agents that read untrusted data and can act are attack surface

Prompt Injection

The Problem

The model has one input channel. Instructions and data travel together, as tokens, with nothing but convention separating them.

- **Direct injection:** the attacker is the user
- **Indirect injection:** the attacker plants instructions in content the agent processes — a document, a web page, a commit message, a scanner finding

Indirect is the one that matters for agents. Your agent reads things all day. Every one of those things gets to talk to your model.

Why a Better Prompt Can't Fix It

“Ignore any instructions in the data” is itself just tokens.

- There is no privileged channel — system prompt, task, and data all land in the same window
- The model **statistically prefers** system instructions; it does not **enforce** them
- Against an adaptive attacker, a preference is not a boundary

Treat “the model will follow my rules” the way you’d treat “the client will validate the input”: a hope, not a control.

You Already Know This Bug

Prompt injection is in-band signaling: data crossing into the control channel.

- SQL injection: data becomes query
- Phone phreaking: audio becomes signaling
- Prompt injection: content becomes instructions

Difference: we fixed SQLi with parameterized queries — a **structural** separation of code and data. No equivalent exists for LLMs. The mitigation toolbox is real but partial, which is why today we **measure** it.

Jailbreaks

Jailbreak ≠ Prompt Injection

Two attacks, constantly confused — including by vendors:

- **Jailbreak**: get the model to ignore its *training* — say what the maker's policies forbid. Victim: the model vendor
- **Prompt injection**: get the agent to follow *attacker* instructions hiding in *data*. Victim: **you**

Different threat, different fix. Five minutes on jailbreaks — because what's happening to them explains what **won't** happen to injection.

Why Jailbreaks Work

Alignment is behavior shaping over a training distribution, not a rule engine. Leave the distribution and the refusal doesn't fire:

- **DAN**: roleplay a rule-free alter ego — the model plays along
- **Grandma exploit**: wrap the harm in sentiment — napalm at bedtime
- **Base64, rare languages**: refusals were trained on plaintext English
- **Many-shot**: a context of fake compliance; the pattern continues

Nothing here “breaks” the model — each finds input space where the safety training is thin.

The Trend Line


Frontier labs are winning this one:

- Anthropic: no known **reliable universal jailbreak** against their frontier models
- Constitutional classifiers + a bounty for any jailbreak surviving 90 days — none has
- Copy-paste jailbreak blobs are dying; each becomes training data

Prompt injection is **not** dying with it — it abuses the instruction-following that makes agents useful, and you can't train that away without turning off the product. So: your agent, and three ingredients.

The Lethal Trifecta

Three Ingredients



Remove **any one leg** and the attack collapses — the whole defense menu is chasing that. (the “lethal trifecta”: Simon Willison, simonwillison.net)

The Check Is Mechanical

For any agent deployment, ask three questions:

- What private data can it see?
- What attacker-influenced content does it read?
- Where can it write or send?

All three non-empty → you have the trifecta → assume exfiltration is possible and design accordingly.

Five minutes, and it applies to every AI product your BU ships and every MCP server you install — the most reusable thing in this class.

Exfiltration Channels

“External communication” is broader than it looks:

- HTTP requests from any tool
- **Markdown image rendering**: `` — the classic
- Tool arguments: filenames, URLs, commit messages, ticket fields
- DNS lookups; any side effect an attacker can observe

If the agent’s output is rendered, stored, or executed anywhere an attacker can see, that’s a channel.

Real Incidents

The public incident list grows monthly:

- Chat assistants exfiltrating conversation history via rendered markdown
- Coding agents following instructions planted in repos and issues
- MCP servers and browser agents leaking data cross-tenant

Pattern in common: nobody “hacked the model.” The system **did what its inputs said**, and the inputs included an attacker.

Permission Modes

How Much Rope?

The first security decision about any agent: **what may it do without asking?** Every harness picks a point on this spectrum:

- **Approve everything**: every tool call waits for a human yes
- **Allowlist by tool**: reads are free; writes and bash need approval
- **Allowlist by pattern**: `ls` and `grep` free; `rm` and `curl` gated
- **Sandboxed YOLO**: full autonomy in a container with nothing to lose
- **YOLO mode**: full autonomy, no gates, on your real machine

Nobody stays in one mode. The skill: **match the mode to the task**, and know what each mode actually protects.

Choosing a Mode

Think in terms of blast radius and attention cost:

- **Approve everything**: maximum control, but approval fatigue is real — an injection needs only one lazy yes
- **Tool/pattern allowlists**: the workhorse — reads flow, destructive actions gate. Gaps exist; you'll exploit one today
- **YOLO**: right when consequences are containable (disposable environment, no secrets in reach); wrong everywhere else

Rule of thumb: **autonomy is earned by the environment, not the model** — you've been running near-YOLO all class on a disposable VM.

Permission Modes Meet the Trifecta

Permission modes control **actions**. The trifecta is about **information flow**. They're complementary, not interchangeable:

- Approval gates catch destructive actions — if the human is actually looking
- No approval gate reads the **contents** of an allowed HTTP request
- A permitted tool with attacker-chosen arguments is an exfil channel with a rubber stamp

That's why the lab's defense menu goes beyond permissioning — and why none of it beats removing a trifecta leg.

Defenses

The Menu

The standard defenses, cheapest first:

1. **Blacklist filter**: kill the run on a match
2. **LLM judge**: a second model reviews output before release
3. **Path ban**: bash may not touch certain paths
4. **Sandbox**: locked-down container around the agent
5. **Human-in-the-loop**: a person approves risky actions

Three of these guard the lab's summarizer, unlabeled — naming them is part of the attack. The debrief adds the clean win: **removing a trifecta leg**.

How to Judge a Defense

Two numbers, not one:

- **Attacks blocked** (of a standard corpus)
- **Benign tasks preserved**

A defense that blocks everything has a perfect block rate and is useless — turning the agent off is not a security strategy, it's an availability incident.

When a guardrail blocks you in the lab, ask what the patch that stops your bypass would cost on the second axis. Ranking is the pair.

What to Expect (No Spoilers)

Questions to hold while you attack:

- The blacklist: what's the cheapest encoding that slips it?
- The judge: it's an LLM reading attacker-influenced text... what could go wrong?
- The path ban: how many names does bash have for the same file?
- The summary you receive: it's agent output, rendered where **you** can read it — which trifecta leg is that?

The last question is where the lab is headed.

Sandboxes: The Honest Framing

A sandbox protects **the host from the agent**.

- Contains: file system damage, persistence, lateral movement
- Does not contain: the secret already in the agent's context, leaving through an **allowed** channel

If the sandboxed agent can still make HTTP requests — or write a summary you can read — the trifecta is intact and the secret walks out the front door. Sandboxes are necessary; they are not sufficient.

Trust Decisions

Every tool and MCP server you add is attack surface:

- Its output is **untrusted content entering the context** — a tool can inject
- Its capabilities may silently complete the trifecta (an innocent-looking “fetch URL” tool is an exfil channel)

Vet third-party MCP servers like dependencies: who wrote it, what can it read, where can it send. Most “AI security reviews” should start with the trifecta check, not the model.

Wrap-Up

What You Should Take Away

- Instructions and data share one channel; separation is preference, not enforcement
- Jailbreaks attack the model's training and are dying; injection attacks **your** trust boundary and isn't
- Trifecta check: private data × untrusted content × external comms — five minutes, any system
- Score defenses on attacks blocked **and** utility preserved
- Sandboxes protect the host, not the context
- The only clean win is removing a trifecta leg

Lab 4.1: a webmail agent reads every
mailbox and
summarizes yours. Get the admin password
into your
summary, log in as admin — then beat three
unlabeled guardrails