ISSUE #005 · HIJACKED EXECUTION CONTEXT

Developers trust tooling. That does not earn it.

Three acts this week. Skill files inherit privilege before the model sees them. CI runners inherit publish identity from poisoned cache. Remote dev tooling inherits remote access from a developer's authenticated session.

> ATTACK

BleepingComputer · Lawrence Abrams · 2026-05-14 · **OpenAI confirms security breach in TanStack supply chain attack**

OpenAI confirmed on 2026-05-14 that the TanStack supply-chain attack hit two employee devices, exposed limited credential material from internal repos, and impacted code signing certificates for macOS, Windows, iOS, and Android. June 12 macOS cert rotation deadline. Wiz lists other named victims: Mistral AI. UiPath. Guardrails AI.

The chain: a `pull_request_target` workflow checked out attacker code, poisoned the shared Actions cache, then a later trusted release run restored that cache on `main`. From there the malware read OIDC tokens out of runner memory via `/proc/<pid>/mem` and published through the legitimate trusted-publisher path. 84 malicious versions across 42 `@tanstack/*` packages, all carrying valid Sigstore provenance. 170+ packages, 518M monthly downloads, CVE-2026-45321 (CVSS 9.6). **Valid provenance is not enough when runner memory is the real root of trust.**

▸ SHIP THIS WEEK

Audit every `pull_request_target` workflow in your org for `actions/cache` restores that cross branch trust boundaries. Rotate downstream credentials reachable from publish hosts during 2026-05-10 through 2026-05-13. Review every release job that minted OIDC tokens in that window. Pin signing identity to issuer plus subject (workflow identity). Treat valid provenance as baseline evidence. Sigstore Fulcio chains and equivalent issuer/subject bindings carry the actual trust.

> RULE

Microsoft · Yossi Weizman + MS Defender Security Research Team · 2026-05-14 · **When configuration becomes a vulnerability: Exploitable misconfigurations in AI apps**

Exposed AI control planes now behave like unauthenticated cluster-admin surfaces. Mage AI ships an internet-facing `LoadBalancer` on port 6789 with no auth and high-privilege service accounts. kagent lacks auth by default if exposed. AutoGen Studio ships without auth enabled. Microsoft Defender's telemetry: more than half of cloud-native workload exploitations stem from misconfiguration. 15% of observed remote MCP servers are severely insecure.

```
# detection logic for K8s audit logs (Sigma-style shape, adapt to your SIEM)
verb|in: [create, update, patch]
objectRef.resource: services
requestObject.spec.type: LoadBalancer
AND ((name contains mcp|mage|kagent|autogen) OR (ports.port == 6789))
```

> DEFENDER

Datadog Security Labs · Nick Frichette + Ryan Simon · 2026-05-11 · **Malicious Coding Agent Skills and the Risk of Dynamic Context**

SKILL.md dynamic-context commands execute before the model evaluates them. Model safety is bypassed entirely. Reversec extends this to the full attack surface: skills are instructions handed to a tool with file, shell, and network reach. Unvetted skill repos are a supply-chain risk that prompt-injection defenses do not address. Pick one repo, grep `.claude/`, list every skill author. Unknown author equals audit.

SpecterOps · Adam Chester · 2026-05-06 · **The Accidental C2: Exploring Dev Tunnels for Remote Access**

Dev Tunnels carry a full remote-access stack: REST, then WebSocket, then SSH via `russh` with `None` auth (outer tunnel is already trusted), then MsgPack RPC with methods like `spawn`, `fs_read`, `fs_write`. A built-in remote access framework shipped as a developer-productivity feature. Egress-control `*.devtunnels.ms` and `*.rel.tunnels.api.visualstudio.com`. Inventory who created Dev Tunnels in the last 90 days.

> RADAR

Fog Security · Jason Kao · 2026-05-12 · **Authorization Bypass in Amazon Quick: Unauthorized AI Chat Agent Usage**

AWS silently patched an auth bypass in Amazon Quick AI. The failure sat inside a managed service readers trust by default, which is the clearest cloud-plus-agents trust story of the week.

SOURCES

BleepingComputer / bleepingcomputer.com/news/security/openai-confirms-security-breach-in-tanstack-supply-chain-attack
Wiz / wiz.io/blog/mini-shai-hulud-strikes-again-tanstack-more-npm-packages-compromised
Microsoft / microsoft.com/en-us/security/blog/2026/05/14/configuration-becomes-vulnerability-exploitable-misconfigurations-ai-apps
Datadog / securitylabs.datadoghq.com/articles/malicious-skills-supply-chain-risks-in-coding-agents-with-dynamic-context
SpecterOps / specterops.io/blog/2026/05/06/dev-tunnels-the-accidental-c2
Fog Security / fogsecurity.io/blog/authorization-bypass-in-amazon-quick-ai-agents