Christian Schneider
Security Architect, Hacker, Trainer

CONSULTING · ASSESSMENT
2 sessions + report · Inhouse or remote · DE / EN

CODING AGENT SECURITY

Securing the use of *AI coding agents*.

Copilot, Cursor, Claude Code, Windsurf, Devin, and a long tail of MCP-based extensions are already in production use across your engineering org, often before any formal security review. A structured two-session workshop and a defensible report give you a prioritized risk view and a two-wave remediation roadmap covering the next nine months.

christian-schneider.net/consulting/coding-agent-security ↗

NINE DOMAINS ASSESSED · 78 QUESTIONS

1. **Inventory & adoption**
   Which agents are in use, by whom, where they run. Deployment mode and runtime location: IDE plugin, CLI, vendor cloud background agent.
2. **Governance & policy**
   AI coding policy, approval gates for new agents and MCP servers, vendor due diligence, training and awareness coverage.
3. **Identity, auth & secrets**
   SSO and SCIM enforcement, token scope and lifetime, secret redaction on agent traffic, where dev-side secrets live.
4. **MCP & tool connectivity**
   MCP server inventory and source vetting, hosting models, tool capability surface: shell, file write, network, cloud SDK, IaC.
5. **Code context & data exposure**
   What code leaves your network and where it goes. Vendor data handling: training opt-out, zero-retention, DPA, BYOK, region pinning.
6. **Generation safety & quality gates**
   Human review of AI code, scanning stack wired into the pipeline, prompt-injection defenses on source inputs, destructive-action controls.
7. **Supply chain & dependency risk**
   Slopsquatting controls, agent binary provenance, MCP server signing and pinning, third-party rule packs and personas.
8. **Execution environment & sandboxing**
   Where agent-initiated commands actually run, shell sandbox strength, dev-machine hardening, agent reach into CI/CD and production.
9. **Monitoring, audit & incident response**
   Central logging of prompts, completions, tool invocations. Anomaly detection, IR playbooks for poisoned MCP, leaked token, agent destructive action.

HOW IT WORKS

1. **Scoping call.** Agents, teams and depth in scope. Format and attendees.
2. **Session 1: Workshop.** Facilitated walk-through of the 78 questions across all nine domains.
3. **Report production.** Risk overview, two-wave roadmap and residual risk outlook written up async.
4. **Session 2: Risk & roadmap.** Walk the report together. Wave 1 owners and timelines agreed in the room.
5. **Follow-up (optional).** Progress review against Wave 1 once you have had time to work with it.

YOU RECEIVE

- → **Risk overview:** prioritized across the nine domains, grounded in your answers.
- → **Two-wave roadmap:** Wave 1 (0–3 mo) foundations, Wave 2 (3–9 mo) strengthening.
- → **Residual risk outlook:** what remains after each wave, for board communication.
- → Per-domain severity verdict, mapped to recognized standards.
- → Maturity baseline (0–10) to measure improvement against.

RIGHT FIT WHEN

AI coding agents are already in production use across your engineering org, and leadership or customers are asking how that is governed.

Earlier in your journey? A 3-hour custom focus session is the better starting point.

Let's scope your Coding Agent Security assessment.
A short scoping call pins down which agents and teams are in scope, what topics deserve the most depth, and who should be in the room. Fixed price once the scope is clear.

GET IN TOUCH
mail@christian-schneider.net
christian-schneider.net