
A PRACTICAL GUIDE FOR TECHNOLOGY LEADERS

Secure AI-Assisted Development.

Protecting enterprise code, intellectual property, and privacy in the age of private AI coding agents.

● PRIVATE LLMS

● OWASP TOP 10

● IP PROTECTION

● ENTERPRISE AI

WRITTEN FOR

CTOs · CDOs · Technology Leaders

Decision-makers adopting AI coding tools

PRODUCT

AgentOne

by Iterate.ai

SUBJECT

Private AI Coding

Security, privacy & IP



ABOUT THIS PAPER

Capture the productivity of AI coding without giving up your code.

The rapid adoption of AI-powered coding assistants is transforming software development. For the enterprise, that transformation carries real risk: intellectual-property leakage, code-security vulnerabilities, regulatory-compliance gaps, and the slow erosion of competitive advantage.

This paper provides a practical framework for harnessing AI coding acceleration without compromising security, privacy, or IP — and shows how private, on-premises infrastructure delivers enterprise-grade productivity with zero data exposure.

WHO THIS PAPER IS FOR

Technology leaders navigating AI coding adoption: **CTOs** evaluating secure tool strategies, **CDOs** responsible for data governance and IP protection, **Engineering VPs** balancing developer productivity with code security, **business leaders** assessing ROI, and **security architects** designing safe AI integration frameworks.

HOW TO READ IT

Front to back, or by part. The threat landscape (Part One) motivates the frameworks and architecture (Part Two); the practice and product sections (Parts Three & Four) make it operational.

PUBLISHED BY

Iterate.ai — San Jose, CA & Denver, CO. Private AI infrastructure for the enterprise.

iterate.ai/agentone



CONTENTS

What's inside.

From the threat landscape to a working deployment roadmap. The framework comparison on p.8 and the tool evaluation on p.15 are the two most leaders return to.


PART ONE · THE THREAT LANDSCAPE		04 — 07
00	Executive Summary The stakes, the numbers, and the thesis in one page.	04
01	The AI Coding Assistant Threat Landscape Five critical threat vectors facing every enterprise using public AI tools.	05
02	The Rise of Vibe Coding Hallucinated dependencies, security bypasses, and hardcoded secrets.	06
03	Secure Vibe Coding Practices Treat AI code as untrusted — and the four habits that make it safe.	07
PART TWO · FRAMEWORKS & ARCHITECTURE		08 — 12
04	Enterprise Security Frameworks OWASP, SHIELD, NIST AI RMF, MAESTRO, and the 3Rs.	08
05	Public vs. Private AI Coding A strategic choice with implications for IP, compliance, and cost.	09
06	Defense in Depth Five layers that secure the AI coding pipeline end to end.	10
07	OWASP-Level Security by Default Eighteen tests, six categories, run at the moment code is created.	11
08	Secure Coding Architecture Three layers: developer environment, AI engine, security & governance.	12
PART THREE · PUTTING IT INTO PRACTICE		13 — 15
09	The Secure AI Development Framework Six pillars and a six-stage lifecycle any organization can adopt.	13
10	Compliance & Governance in Practice SOC 2, GDPR, HIPAA — and a governance operating model that works.	14
11	Evaluating AI Coding Tools The enterprise-critical capabilities that should drive procurement.	15
PART FOUR · AGENTONE · COLOPHON		16 — 18



EXECUTIVE SUMMARY

The acceleration is real. So is the exposure.

AI coding assistants are transforming software development — but most popular tools work by sending your proprietary source code to external servers. That architecture creates an attack surface most organizations are only beginning to understand.



Research shows that **48% of AI-generated code contains security vulnerabilities**, while **35% exhibits licensing irregularities** that expose organizations to legal risk. For regulated enterprises, sending code to external services can itself be a compliance violation.

This paper examines the threat landscape, establishes best practices for secure AI-assisted development, and explores how private, on-premises AI coding infrastructure delivers enterprise productivity gains with **zero data exposure**.

THE INDUSTRY SIGNAL

"67% of enterprises pursuing data sovereignty have already shifted to private AI infrastructure to strengthen regulatory compliance and maintain control over their most sensitive assets."

Enterprise AI Security Survey, 2025

The fundamental question is no longer *whether* to adopt AI coding assistance, but **how to adopt it without exposing the organization's most valuable digital assets**. That is a strategic business decision — not merely a technical one.



01 MAP THE RISK BEFORE YOU ADOPT

Most AI coding tools send your code outside your security boundary.

Any code pasted into a prompt, any file opened for context, any repository indexed for autocomplete becomes data that exists outside your perimeter. Five threat vectors follow directly from that architecture.

<p>01</p> <p>Code Exfiltration</p> <p>Proprietary code sent to third-party servers for processing.</p>	<p>02</p> <p>IP Contamination</p> <p>Licensed open-source code injected into your codebase.</p>	<p>03</p> <p>Vuln Injection</p> <p>48% of AI-generated code carries security flaws.</p>	<p>04</p> <p>Model Training</p> <p>Your code used to train models your competitors then use.</p>	<p>05</p> <p>Compliance Breach</p> <p>HIPAA / GDPR / SOC 2 regulatory violations.</p>
--	---	---	--	---

Five critical threat vectors facing enterprises using public AI coding tools.

Code Exfiltration

Many standard terms of service grant providers the right to use submitted data to improve their models. In a worst case, fragments of your proprietary code could be suggested to developers at competing firms. This is not theoretical — it is an architectural reality of how most tools work.

Vulnerability Injection

AI code carries a higher rate of flaws than human-written code — SQL injection, XSS, insecure deserialization, path traversal. Developers also scrutinize AI suggestions *less* than their own work, compounding the risk.

IP Contamination

Assistants trained on open-source repos can emit code that mirrors GPL or AGPL material — roughly 35% of samples show licensing irregularities. Contamination has already forced complete codebase rewrites at Fortune 500 companies.

Regulatory Exposure

For HIPAA, GDPR, SOC 2, or PCI DSS organizations, sending code or data schemas externally may itself be a violation. The 2026 OWASP update adds Agent Goal Hijacking (ASI01) — manipulating autonomous agents through poisoned inputs.



02

A NEW CATEGORY OF RISK

The rise of vibe coding.

Building software primarily through natural-language prompts has moved from side projects into enterprise workflows. The speed is real — and so are risks traditional security frameworks were never designed to catch.

45%

Of AI-generated code contains classic OWASP Top 10 vulnerabilities. (*Veracode, GenAI Code Security Report 2025*)

7 / 10

Instances of LLM-generated Java code contain security flaws.

20%

Of vibe-coded applications have serious vulnerabilities or configuration errors. (*Wiz*)

VIBE CODING THREAT VECTORS

1

Hallucinated dependencies.

Socket.dev analyzed 576,000 code samples and found **20% of AI-recommended packages do not exist** — 205,000 unique phantom names. Attackers register these on NPM and PyPI; an unwitting `npm install` can pull a malicious payload. A new supply-chain attack, uniquely enabled by AI.

2

Hallucinated security bypasses.

A model accidentally removes a security check while generating or refactoring — deleting an auth keyword, stripping input validation. The code still compiles and appears to work, so the regression slips through review. The AI is not malicious; it simply does not understand what it removed.

3

Hardcoded secrets.

Models frequently emit code with hardcoded API keys, tokens, and credentials — patterns learned from millions of public repos where developers committed secrets. They look real enough to pass casual review but represent significant exposure in production.

4

Missing regulatory context.

AI assistants have no awareness of HIPAA, PCI DSS, or GDPR data-handling mandates. Required storage and processing methods will not appear in generated code unless explicitly specified — and even then, compliance is not guaranteed.



03

KEEP THE SPEED, LOSE THE RISK

Secure vibe coding practices.

Four practices mitigate the risks while preserving the productivity benefit. The first is the foundation everything else rests on.

FOUNDATIONAL PRINCIPLE

Treat all AI-generated code as untrusted.

Compiling and running does not mean secure. Every AI-generated function, module, and config should get the same scrutiny as third-party library code — because that is what it is: code written by an external entity with no knowledge of your requirements.

TWO-STAGE GENERATION

Use the AI self-reflection pattern.

First, ask the AI to build the feature. Then, in a separate prompt, instruct it to “act as a Security Engineer” and review its own code for injection, path traversal, auth bypasses, and hardcoded secrets. This catches many issues single-pass generation misses.

SPEC-DRIVEN

Define security before any code.

Specify the policies the AI must satisfy: no public database access, unit tests per feature, all user input sanitized, no hardcoded keys, established auth libraries only. Ground these rules in the OWASP Top 10 as a minimum baseline.

SUPPLY CHAIN

Verify every dependency.

Before installing any AI-recommended package, confirm it exists on the official registry, check its download count and maintenance status, and review its source. Automated dependency scanning can flag hallucinated or suspicious packages before they enter the codebase.



Security scanning is necessary but not sufficient. The safest workflow automates the practices teams defer under pressure — tests, diagrams, formatting, and docs.



04 ALIGN TO INDUSTRY STANDARDS

Enterprise security frameworks.

No single framework covers all risks. Together they give technology leaders a robust toolkit for governing AI-assisted development — and a vocabulary regulators and auditors recognize.

OWASP · 2026

Top 10 for Agentic Applications

Peer-reviewed by 100+ experts. Four major risks: **memory poisoning**, **tool misuse**, **privilege escalation**, and **cascading failures** in multi-agent systems. The baseline threat model for AI coding agents.

PALO ALTO · UNIT 42

The SHIELD Framework

Built for vibe coding: **Separation of duties**, **input/output validation**, **security-focused helper models**, and **continuous monitoring**. Designed for AI-generated code, not adapted from traditional AppSec.

NIST

AI Risk Management Framework

Four functions map to the coding lifecycle: **Govern** (policies), **Map** (where AI code lives), **Measure** (vuln density), and **Manage** (controls). Demonstrates due diligence to regulators.

MULTI-AGENT

MAESTRO Threat Modeling

For swarm architectures where multiple agents work in parallel. Structures threat identification at agent boundaries, communication channels, and shared resources — risks single-agent models miss.

THE 3RS OF RESILIENT SECURE CODING · 2026

Rotate

Secrets, keys, and certificates are ephemeral and rotated regularly — never hardcoded.

Repair

Vulnerable code and dependencies patched promptly — within hours, not weeks.

Repave

Infrastructure rebuilt from known-good templates rather than patched in place.



05

A STRATEGIC CHOICE

Public vs. private AI coding.

Public tools run on shared cloud infrastructure; your code leaves your perimeter, however briefly. For many organizations, that alone is disqualifying. The distinction is not merely technical — it is strategic.

PUBLIC AI CODING TOOLS	PRIVATE AI CODING TOOLS
✘ Code sent to external servers	✔ Code never leaves your network
✘ May train on your proprietary code	✔ Zero data used for model training
✘ No built-in security scanning	✔ OWASP security tests built in
✘ Shared infrastructure	✔ Dedicated on-prem / VPC deployment
✘ Limited audit trail	✔ Full observability and audit trails
✘ Compliance gaps for regulated industries	✔ SOC 2, GDPR, HIPAA, PCI aligned

Side-by-side comparison of risk profiles between public and private AI coding approaches.

WHY PRIVATE LLMS MATTER

Beyond security, private deployment eliminates per-token API costs, reduces latency by keeping inference local, and ensures availability independent of external providers. For large codebases, extended context windows across sessions change what AI can do — from completing lines to understanding entire architectures. **67% of enterprises pursuing data sovereignty have already shifted to private AI infrastructure.**

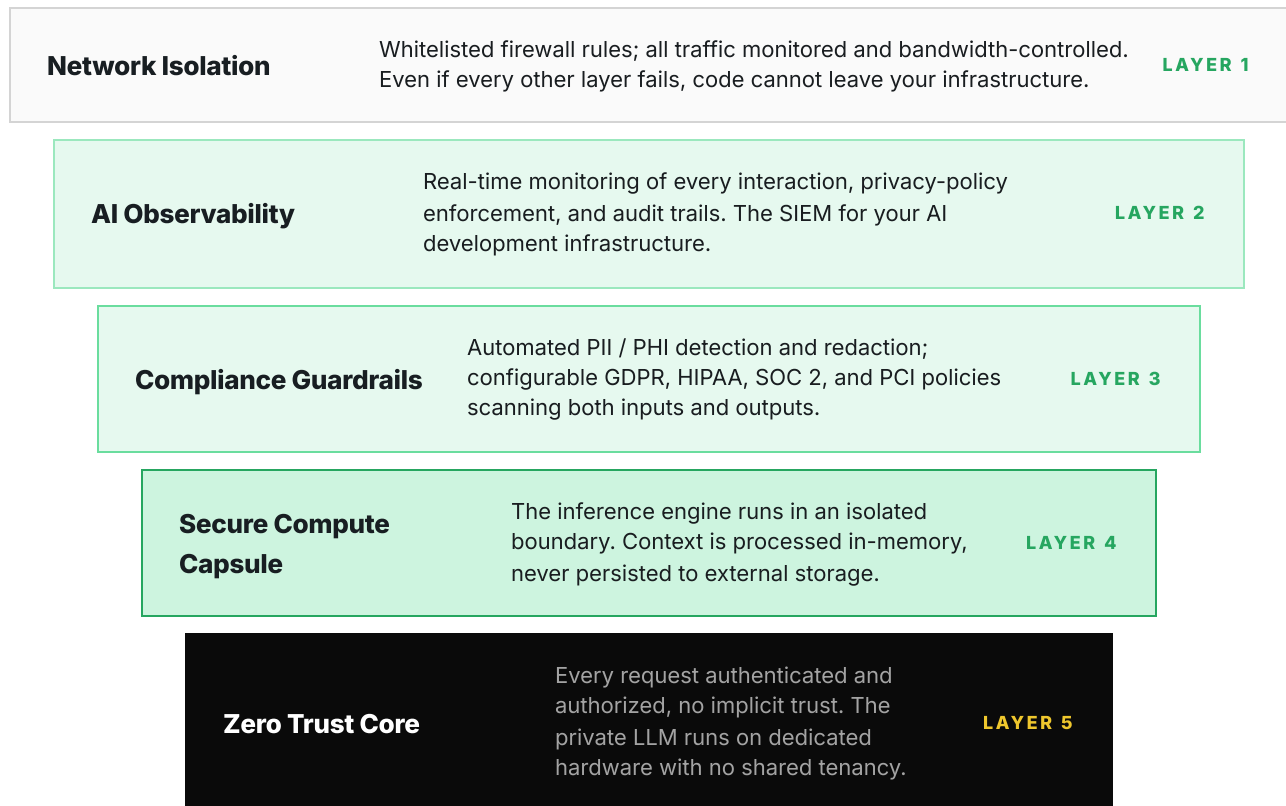


06

NO SINGLE CONTROL IS ENOUGH

Defense in depth for the AI coding pipeline.

Like traditional AppSec, AI coding infrastructure needs layered defenses — from the network boundary down to the inference engine itself. Each layer holds even if another fails.



Five-layer defense-in-depth model for private AI coding infrastructure — outermost network boundary to innermost zero-trust core.



07 SCAN AT THE MOMENT OF CREATION

OWASP-level security by default.

AI produces code too fast to wait for CI/CD scans — that window lets vulnerable code be committed and merged first. A mature environment runs **18+ tests across six categories** on every generation event.

<p>01 Injection Prevention</p> <p>SQL, NoSQL/GraphQL, command, template (SSTI), and XML external entity (XXE) attacks — the patterns AI most often produces by concatenating user input into queries.</p>	<p>02 Data Handling</p> <p>Unsafe deserialization, path traversal, and input validation. Subtle flaws easily missed in review — path traversal in particular appears often in AI file-handling code.</p>
<p>03 Authentication & Sessions</p> <p>JWT flaws, session weaknesses, and auth-bypass patterns. AI often generates auth that works but uses insecure defaults — weak hashing, missing token expiry.</p>	<p>04 Cryptographic Validation</p> <p>Weak ciphers, insecure random-number generation, poor key management. Critical because crypto mistakes are invisible to functional testing — the code “works.”</p>
<p>05 Configuration Security</p> <p>Hardcoded secrets, API keys, insecure dependencies, and missing security headers — AI is notorious for emitting hardcoded credentials pulled from training data.</p>	<p>06 Static Code Analysis</p> <p>Dead-code detection, complexity analysis, and secure-pattern enforcement — ensuring AI code is not only secure but maintainable.</p>

Beyond scanning: a truly secure workflow also automates the engineering practices teams defer under pressure — running full unit-test suites after every generation, producing architecture diagrams, applying style guides, and generating API documentation. Automated, these ensure AI code meets the same quality bar as hand-crafted code, consistently and at scale.



08 THREE LAYERS, CLEAR BOUNDARIES

Secure coding architecture.

Whether you build or buy, the core principles hold: separate concerns into three layers, each with clear responsibilities and security boundaries.

Developer Environment	AI Engine Layer	Security & Governance
<ul style="list-style-type: none"> IDE / Editor Plugin Code Chat Interface Command Interface MCP Connectors Terminal Access Browser Tools 	<ul style="list-style-type: none"> Controller / Router Orchestration Loop Plan & Execute Parallel Agents Private LLM Code generation 	<ul style="list-style-type: none"> OWASP Scanning Observability Monitor PII / PHI Detection Audit Trail Logging Compliance Engine Zero Trust Architecture

Three-pillar architecture: Developer Environment, AI Engine, and Security & Governance.

Developer layer

Integrates into existing workflows — native IDE, chat, and CLI. All interactions route through a secure local proxy, never directly to external APIs. The proxy enforces policy, logs, and strips sensitive context.

Engine layer

Routes requests by sensitivity, orchestrates multi-step tasks, manages context, and coordinates parallel agents. Supports multiple private and commercial models with plan-then-execute approval gates.

Governance layer

What separates enterprise-grade from consumer tools. Operates independently of the engine — it cannot be bypassed by prompt injection. Security is enforced at the infrastructure level, not the application.



09

THE BASELINE ANY TEAM CAN ADOPT

The Secure AI Development Framework.

1 Assess Evaluate risks & compliance	2 Isolate Deploy private LLM on-prem	3 Govern Set policies & guardrails	4 Build Code with secure AI tools	5 Scan OWASP security analysis	6 Monitor Observability & audit
--	--	--	---	--	---

The six-stage Secure AI Development Lifecycle.

1 · Risk Assessment & Classification

Classify your codebase by sensitivity. Public docs can safely use cloud AI; crown-jewel IP must never be processed externally. Map repos to risk tiers and enforce them through tooling, not policy documents developers ignore.

3 · Governance & Policy Enforcement

Design governance in from day one. Automated guardrails enforce policy in real time: PII/PHI detection, GDPR and HIPAA checks, configurable content filters — continuous, not quarterly reviews that are always out of date.

5 · Observability & Audit Trails

Track who uses AI, which models, token consumption, cost allocation, and every prompt-response pair. Enables compliance audits, incident response, and optimization of developer productivity.

2 · Environment Isolation

Run private LLMs on dedicated hardware — Intel Gaudi, NVIDIA A100/H100, or equivalent — inside your network. A single server can power 15–20 developers at sub-second latency. Security by architecture beats security by policy.

4 · Continuous Security Scanning

Integrate OWASP-level scanning at the moment code is generated, not in CI/CD. Results visible to the developer immediately, creating a tight feedback loop that improves code quality.

6 · Continuous Improvement

Review code-quality metrics, update scanning rules for emerging threats, and invest in developer training on secure AI collaboration. Treat adoption as a process, not a one-time deployment.



10

A PREREQUISITE, NOT A FEATURE

Compliance & governance in practice.

AI coding tools create new control points that existing governance frameworks were not designed to address. Three regimes matter most for regulated industries.

SOC 2 TYPE II

New control points

Where is code processed? Who has access? Are prompts and responses logged? Can audit trails be produced on demand? Ensure your AI infrastructure is in scope — and that the vendor can provide their own SOC 2 report.

GDPR

Data processing & transfer

If prompts or context contain personal data, that data is being “processed” — possibly in another jurisdiction. Private deployment sidesteps transfer concerns, but you must still document AI data flows in your ROPA.

HIPAA

PHI in code

Healthcare code references PHI in schemas, endpoints, and fixtures. Sending it to a cloud AI without a BAA is a violation. Even with one, inadvertent PHI exposure makes private deployment the recommended approach.

BUILDING A GOVERNANCE OPERATING MODEL

Three capabilities — or governance becomes theater.

1 · Policy framework

Defines what data can touch AI, which models are approved, and what review applies to AI code.

2 · Technical enforcement

Implements those policies through automated guardrails — not manual checklists.

3 · Monitoring & reporting

Gives leadership ongoing visibility into AI usage, risk metrics, and compliance status.

Without all three, policies exist on paper but are not enforced in practice.



11 PRIORITIZE RISK FIT OVER FEATURE COUNTS

Evaluating AI coding tools for the enterprise.

The comparison below focuses on enterprise-relevant capabilities that directly impact security, scale, and operational control.

CAPABILITY	AGENTONE	COPILOT	CURSOR	CLAUDE CODE	WINDSURF
Deployment options	On-Prem / VPC / Cloud	Cloud only	Cloud only	Cloud only	Cloud only
Large codebases (100K+ files)	✓ Yes	Degrades	✗ 20–50K	✓ Yes	✗ No
Automatic security reviews	✓ Yes (18+)	✗ No	~ Limited	✗ No	✗ No
Private LLM support	✓ Yes	✗ No	~ Limited	✗ No	~ Limited
MCP server integration	✓ Full + Externals	✗ No	✓ Yes	✓ Yes	✓ Yes
Multi-model support	✓ Yes (20+)	✗ No	✓ Yes	✓ Yes	✗ No

Feature comparison across enterprise-critical capabilities. ✓ full · ~ limited · ✗ unsupported.

Deployment flexibility is the most consequential decision — on-prem/VPC keeps all code within your boundary. Table-stakes for regulated industries.

Private LLM support lets you run open-source or custom models on your own infrastructure — eliminating data-exposure risk and per-token costs.

Security scanning determines whether you catch vulnerabilities at generation time or in production. Built-in OWASP scanning is a different posture entirely.

Large-codebase handling matters: many tools degrade past 100K+ files. Test any tool against your actual codebase, not a demo project.



How AgentOne addresses these challenges.

Rather than bolting security onto a consumer assistant, AgentOne was built as a private, enterprise-first AI coding platform. Here is how it implements the principles in this paper.

Private by architecture

Runs entirely within your infrastructure — on-prem or VPC. A single server powers up to 20 developers with no cloud dependency. There is no setting to accidentally expose your code; the network path simply does not exist.

Enterprise-scale context

Up to 2 million tokens of context across sessions — enough to understand 100K+ file repos with dependency mapping. Six patented technologies, including persistent workspace memory with LFU eviction.

Observability with AgentWatch

Org-wide visibility: who uses AI, which models, token-level cost by team, centralized key management, budget enforcement, and complete audit trails — without slowing teams down.

Built-in security scanning

18+ OWASP-aligned tests across six categories on every generation event, with findings shown immediately. The full OWASP WSTG v4.2 is covered across 12 security categories.

Swarm intelligence

For complex tasks, 26 parallel micro-agents explore strategies simultaneously, then assemble the optimal solution — cutting completion from 45–60s to 10–15s with higher quality.

Compliance verified

Continuous security screening through Vanta covering SOC 2 Type II controls — secure development, access control, cryptography, data management, incident response. Continuously monitored, not point-in-time.



A PHASED PATH, NOT ALL-OR-NOTHING

Your roadmap to secure AI development.

Each phase builds on the last — demonstrating value incrementally while establishing governance foundations.

Phase 1 Assessment

Weeks 1–2

Audit current AI tool usage across teams. Classify codebases by sensitivity. Identify compliance requirements and risk tolerance. Map current security gaps.

Phase 2 Pilot deployment

Weeks 3–6

Deploy a private AI coding environment for a pilot team of 5–10 developers. Connect your preferred private LLM. Configure security policies and guardrails. Establish baseline metrics.

Phase 3 Scale & govern

Weeks 7–12

Roll out to additional teams. Deploy observability for org-wide visibility. Implement budget controls, team-level policies, and multi-tenant isolation. Conduct the first compliance audit against the new infrastructure.

Phase 4 Optimize & expand

Ongoing

Fine-tune models on domain-specific patterns. Expand integrations to enterprise systems. Leverage automated documentation, testing, and diagramming. Measure ROI against Phase 1 baselines.

IN CONCLUSION

Organizations that fail to adopt AI coding tools will fall behind. But those that adopt carelessly will expose their most valuable IP and accumulate security debt that takes years to remediate. The path forward is clear: **private AI infrastructure that delivers the full power of modern LLMs without any of the risk.**

Assess. Isolate. Govern. Scan. Observe. Improve. That is how you capture the productivity of AI-assisted development while protecting the code that defines your competitive advantage.

ABOUT THIS PAPER

Ready to secure your AI development pipeline?

AgentOne by Iterate.ai is the only AI coding platform that combines private deployment, OWASP security scanning, 2M-token context, swarm intelligence, and full observability in a single, enterprise-ready package. To learn more or schedule a demo, visit iterate.ai/agentone or contact our enterprise team.

This paper reflects publicly available research as of 2026 and is intended for educational purposes. It is not legal advice and should not be relied upon as such.

REFERENCES

- 01 Veracode, *GenAI Code Security Report 2025* — 45% OWASP Top 10 vulnerabilities.
- 02 Wiz — 20% of vibe-coded apps with serious vulnerabilities.
- 03 Socket.dev — 576K samples; 205K hallucinated package names.
- 04 OWASP Top 10 for Agentic Applications (2026).
- 05 Palo Alto Networks Unit 42 — SHIELD Framework.
- 06 NIST — AI Risk Management Framework (AI RMF).
- 07 MAESTRO — multi-agent threat modeling framework.
- 08 IBM — *Cost of a Data Breach* & continuous-learning commentary.
- 09 *Enterprise AI Security Survey, 2025* — data-sovereignty shift.
- 10 OWASP Web Security Testing Guide (WSTG) v4.2.

PREPARED BY

Iterate.ai — AgentOne Team

Prepared by Iterate.ai's product and engineering teams, with Claude (Anthropic) and Iterate's *Generate* platform. Trusted by Fortune 500 companies for enterprise AI.

TECHNOLOGY PARTNERS

Intel · NVIDIA · Qualcomm · Dell · IBM · NetApp

PRODUCT

AgentOne
Private AI Coding

HEADQUARTERS

San Jose, CA
& Denver, CO

ON THE WEB

iterate.ai/agentone
info@iterate.ai

EDITION

2026
© 2026 Iterate.ai