Agentic Al for Application Security

The Future of Secure Code Development





The Generative Al revolution isn't just **changing** how we write code—it's fundamentally **transforming** the security landscape of software development.

Section 1: Executive Summary

The Generative AI revolution isn't just changing how we write code—it's fundamentally transforming the security landscape of software development. While GenAI promises unprecedented productivity gains, it also introduces a new class of security challenges that traditional AppSec tools simply weren't designed to handle.

Think of GenAl as a brilliant but unpredictable collaborator. It can generate complex solutions in seconds, but without proper guardrails, it can also introduce subtle vulnerabilities, incomplete features, and security blind spots that could compromise your entire application. In today's high-velocity development environment, where the pressure to deliver features quickly is relentless, these risks are amplified.

The Challenge: Security at AI Speed with Confidence

Traditional security approaches—static scans, manual reviews, and after-the-fact audits—are falling behind. They weren't built for a world where:

- Code is generated at unprecedented speed and scale
- Vulnerabilities can be introduced through AI hallucinations and contextual misunderstandings
- Security flaws might be replicated across thousands of code blocks in seconds
- Traditional testing methods can't keep pace with AI-powered development

The Solution: Agentic Al

Instead of relying on monolithic scanning and heavy manual triage, this approach orchestrates a team of specialized AI "roles" that each tackles a different aspect of AppSec. We gain a continuous, adaptive security mechanism that scales alongside development by delegating tasks like vulnerability discovery, fix generation, and data-privacy checks to intelligent agents across the software development lifecycle.

Key Value for Teams:

- Proactive Insights: Agents automatically spot vulnerabilities in code, configurations, and dependencies before merges reach production.
- Accelerated Fix Cycles: Rather than waiting for humans to interpret scans, agentic workflows propose real-time patches—dramatically cutting mean-time-to-remediation (MTTR).



www.qwiet.ai (877) 331-9092 3

- **Context-Aware Decisions:** Each agent references not only the immediate code snippet but also historical fixes, known threat payloads, and code dependencies.
- Reduced Alert Fatigue: Security engineers and developers can focus on strategic problems—like architecture and business logic—rather than sifting through endless logs or false positives.

Backed by academic research and real world testing—which highlights how AI agents iteratively validate code fixes for improved accuracy—Agentic AI ensures security coverage that evolves in **tandem** with your codebase.

The Security-Speed Paradox

Let's look at a typical scenario in modern development:

```
# Seemingly innocent code
user_input = request.args.get('query')
db.execute(f"SELECT * FROM users WHERE name LIKE '%{user_input}%'")
```

Traditional static analysis would flag this as a potential SQL injection vulnerability. But then what? Your team faces a familiar choice: delay deployment for manual review and fix, or accept the risk and move forward. Neither option is acceptable in today's development landscape.

This is where Agentic AI fundamentally changes the game. Instead of just flagging the vulnerability, our system:

- 1. Analyzes the full context through our Code Property Graph
- 2. Generates and validates a secure fix
- 3. Automatically implements the solution with your team's coding patterns
- 4. Verifies the fix doesn't break existing functionality





Time Reduction: 8-16 hours $\rightarrow 2-4$ minutes

Manual Effort: 6-12 hours $\rightarrow 0$ minutes



www.qwiet.ai (877) 331-9092 4

Section 2: The System

Agentic Architecture & Specialized Roles

In an **agentic AppSec** environment, each specialized AI agent performs a distinct function throughout the security lifecycle. These agents operate in **parallel** or **sequence** depending on the vulnerability type, ultimately converging on secure code.

1. Threat Analyst —

- Primary Task: Analyzes code for potential exploits, crafting real-world attack payloads.
- Value Add: Ensures vulnerabilities are understood from an attacker's perspective, preventing "theoretical" fixes that fall short against genuine threats.

2. Test Engineer —

- *Primary Task:* Translates those payloads into structured test cases, simulating how an exploit might run in practice.
- Value Add: Delivers automated tests that confirm a proposed fix truly addresses the root cause—no half-measures or untested patches.

3. Software Remediation Engineer —

- Primary Task: Proposes an initial code fix, guided by the Test Engineer's cases.
- Value Add: Iterates quickly, applying best practices (e.g., OWASP, NIST) to patch critical flaws without waiting for manual triage.

4. Evaluator —

- Primary Task: Critiques the Resolver's fix, identifying coverage gaps or overlooked edge scenarios.
- Value Add: Ensures the fix is robust, not just superficially patched. Encourages a cycle of improvement before final deployment.

5. Senior Software Engineer —

- Primary Task: Incorporates feedback from the Evaluator to refine and finalize the patch.
- Value Add: Polishes the fix, handles edge cases, and addresses recommendations for performance, readability, or deeper security controls.

6. Dependency Auditor —

- Primary Task: Inspects all libraries and frameworks used by the fix, ensuring none are unmaintained, risky, or "hallucinated."
- Value Add: Flags suspect or outdated dependencies, forcing a re-fix loop if necessary—closing a major supply-chain threat vector.

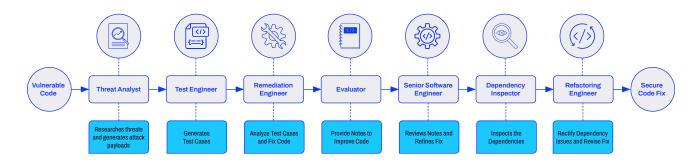


5

7. Refactor Engineer —

- Primary Task: If Dependency Auditor identifies invalid or stale dependencies, Refactor updates or removes them and re-optimizes the code accordingly.
- Value Add: Ensures the final code remains cohesive, stable, and free of technical debt introduced by questionable dependencies.

Guided Agentic Workflow



In short, these seven roles jointly form a **closed-loop** system—from early vulnerability detection to thorough test coverage and final dependency clean-up. The outcome is code that's not just *found* to be secure, but also *verified*, *refined*, and *monitored* post-deployment.

Section 3: The Superior Prompt

Code Property Graph—The Heart of Contextual Fixing

A **critical enabler** of autonomous fixes is the Code Property Graph (CPG). Think of it as a multi-dimensional map of your code's structure and behavior. It merges your Abstract Syntax Tree (AST), Control Flow Graph (CFG), and Program Dependency Graph (PDG) into one unified representation.

Why the CPG Matters

1. Holistic Code Insight:

By unifying syntax (AST), execution flow (CFG), and data dependencies (PDG), the CPG reveals **where** and **how** data flows—perfect for uncovering injection vulnerabilities or identifying where PII might be accidentally logged.

2. Context-Rich Prompts for AI Agents:

Before generating a fix, an agent is fed a relevant subset of the CPG. Rather than scanning an entire codebase in one go, the AI sees only the slices that matter: suspicious lines, their dependencies, and the surrounding methods or modules. This context is **crucial**—as the Meng et al. (2024) study highlights, deeper reasoning about the "true" root cause is critical for accurate bug fixing.



www.qwiet.ai (877) 331-9092 6

3. Precise Vulnerability Location:

By referencing the graph, the Automated Vulnerability Analyst can quickly **pinpoint** lines or functions that call insecure APIs, handle user input unsafely, or rely on outdated libraries. The code snippet used to prompt the LLM is more **targeted**, minimizing guesswork and false positives.

4. Automated Fix Generation:

When the Software Remediation Engineer agent crafts a patch, it benefits from the CPG's structural knowledge, ensuring the fix doesn't break call flows or disrupt related modules. The result is a fix that is both correct **and** robust.

5. Facilitating Multi-Agent Collaboration:

Each specialized AI role can update or consume the same CPG data, creating a **shared context** for collaborative decisions. If the Dependency Auditor flags a library mismatch, the CPG helps the Remediation Engineer see exactly where it's used—and the QA Specialist can retrieve test coverage for the relevant modules.

In essence, the CPG is the prime mover that **informs** and **empowers** each agent's prompt. By packaging the exact lines, function definitions, and data flows in a structured format, the fix generation process transitions from a guesswork approach to a context-informed operation—leading to fewer spurious suggestions and more reliable patches.

Section 4: Additional Technical Components & Practices

Beyond the CPG, several key pillars strengthen the agentic AppSec methodology:

Chain-of-Thought Prompting & Multi-Step Reasoning

- What It Is: Agents like Resolver or Max break down problems into smaller steps, verifying each substep before finalizing.
- Why It Matters: Research shows iterative, multi-step reasoning (cited in Meng et al. 2024) boosts success rates in bug fixes. Agents propose an initial patch, run internal checks, refine if needed, then finalize a stable fix.

2. Explainability & Auditability

- From Black Box to Transparency: By capturing each step of the chain-of-thought, teams can audit how an Al-generated fix was determined.
- Empowering Humans: This traceability ensures that security engineers and developers see why certain changes were recommended, fostering trust and compliance in the final solution.

3. Self-Verification & Cross-Verification

• Self-Verification: The same agent might do repeated passes—e.g., scanning for syntax errors or dataprivacy compliance in its own fix.



www.qwiet.ai (877) 331-9092

- Cross-Verification: Different roles weigh in from their own vantage points. For instance, Toby's test suite might expose an edge condition that the Resolver missed, leading to a revised fix from Max.
- Outcome: Fewer incomplete patches, sharper detection of "hallucinated" code references, and a more bulletproof remediation pipeline.

4. Hallucination Mitigation & Dependency Scoring

- Issue: Large models can reference or invent packages that aren't real or are dangerously out-of-date.
- Solution: Hal (Dependency Auditor) employs a scoring framework—perhaps 60% "In-Dimensions" (vulnerabilities, pinned versions) + 40% "Around-Dimensions" (library popularity, maintenance frequency)—to block or flag suspicious references. If flagged, Refactor adjusts the code accordingly.

5. IDE, CI/CD, & Production Hooks

- Developer Experience: Agents surface issues inline in the IDE. If a pull request introduces a known exploit pattern, the system can block merges or suggest an immediate fix.
- *Production Observability:* With the final code approved, a last pass from Hal or Toby might configure Splunk or Wiz triggers for real-time anomaly detection—bridging prevention and live monitoring.

Conclusion: The Path to Autonomous, Continuous AppSec

By combining **agentic AI roles**, **the rich context of a Code Property Graph**, and **iterative cross-verification**, organizations finally gain an AppSec approach that **moves in lockstep** with rapid development cycles. Teams can devote less energy to manual scanning and repetitive triage, and more energy to strategic architecture decisions.

In Short:

- Developer-Centric: Automated fixes and checks minimize disruptions to the dev workflow.
- Confidence & Coverage: The synergy of specialized roles and data-driven context fosters stronger security outcomes, backed by emerging research on agent-based code repair.
- Adaptive Evolution: As threats evolve, so do the agents—integrating new patterns, exploits, and compliance demands without overhauling your entire security pipeline.

In an industry where new vulnerabilities surface daily, **agentic AppSec** stands as a powerful ally—**continuous**, **context-aware**, and **tirelessly vigilant**.



www.qwiet.ai (877) 331-9092

Learn more about how Qwiet Al helps secure code in minutes at qwiet.ai

