TAG

ENHANCING SECURITY USING AI-ENABLED SOFTWARE AUTO-FIX FROM QWIET

DR. EDWARD AMOROSO,
CEO, TAG
RESEARCH PROFESSOR, NYU

qwiet AI

TAG

ENHANCING SECURITY USING AI-ENABLED SOFTWARE AUTO-FIX FROM QWIET

DR. EDWARD AMOROSO, CEO, TAG,
RESEARCH PROFESSOR, NYU

Agentic AI-based auto-fix solutions, such as those provided by commercial cybersecurity vendor Qwiet AI, are not just emerging; they are transforming the fight against exploitable software vulnerabilities. These advanced capabilities are dramatically shifting how software is secured, tested, and remediated, ushering in a new era of cybersecurity. This report provides an overview of these methods with emphasis on the approach taken by Qwiet AI.

INTRODUCTION

A common thread across most cyber incidents is exploitable software somewhere in the attack chain. TAG analysts have long observed that such vulnerabilities are often introduced during the development process through human error. While tools have existed to analyze, test, and scan software for decades, few have adequately addressed the core problem: vulnerable code making it into production.

The rise of AI has inspired a fresh approach to this software problem, one where it is no longer just about finding the vulnerabilities, but rather about fixing them autonomously. This is the promise of auto-fix: AI systems will detect and generate validated and compliant code patches. Such systems not only significantly improve security posture but also liberate developers from the grind of routine remediation, providing a sense of relief and allowing them to focus on more strategic tasks.

This TAG report explores the emerging category of vendor-supported auto-fix platforms, focusing on the comprehensive and integrated solution provided by commercial vendor Qwiet AI. Qwiet AI's modern platform is designed to meet all your needs, providing a complete solution you can rely on. We also touch on capabilities available from other CI/CD ecosystem vendors, such as GitHub, while emphasizing the importance of CI/CD platform-agnostic support for developers.

[1] *TAG Infosphere provides advanced research and advisory support in cybersecurity and artificial intelligence for enterprise teams, government agencies, public policy lawmakers, academic researchers, and commercial vendors. See* ***https://www.tag-infosphere.com/****.*

[2] *NYU's Center for Cybersecurity (CCS) is an interdisciplinary academic center in which leading edge cybersecurity research, teaching, and scholarship are directed into meaningful real-world technology, platforms, and policies. See* ***https://www.cyber.nyu.edu/****.*

UNDERSTANDING AUTO-FIX

The technique that we refer to as AI-based auto-fix refers to the autonomous detection, diagnosis, and remediation of vulnerabilities in application source code. These solutions typically integrate into the software development lifecycle (SDLC), where they leverage code analysis engines, often built on representations like the Code Property Graph (CPG), to understand software context deeply.

Once a software vulnerability has been identified, these modern auto-fix systems generate patches using large language models (LLMs) and verify their correctness using test frameworks and compliance rules. The best platforms operate through agentic AI, collaborating personas that replicate the expertise of threat analysts, test engineers, and remediation developers.

Qwiet AI's commercially available AutoFix capability demonstrates this approach well. Using the CPG to contextualize software structure, control flow, and data flow, its agent personas can reason through vulnerabilities, craft safe code changes, and validate results using automated test cases, all within minutes. For experienced software engineers, supporting these functions will undoubtedly be a welcome advance.

CASE EXAMPLE: CROSS-SITE SCRIPTING (XSS)

To illustrate, let's consider a basic and typical example where a web application fails to sanitize user input in an HTTP response, creating an XSS vulnerability. Traditional SAST tools would flag this issue and provide remediation suggestions. But an auto-fix engine like Qwiet's would go further. Its agentic AI workflow would perform the following tasks in support of the software lifecycle:

- Confirm the vulnerability by generating a test case.
- Craft a patch using the appropriate sanitization library (e.g., escape-html).
- Validate that the test case no longer triggers the vulnerability.
- Ensure the fix follows OWASP guidance and doesn't break other logic.
- Deliver the patch to the developer as a merge-ready suggestion.

In this model, the vulnerability is discovered, confirmed, and fixed in under five minutes, —transforming what was once an hours-long manual triage into an autonomous, compliant workflow.

AGENTIC AI AND THE CODE PROPERTY GRAPH

Let's delve into the key technical innovations that drive the Qwiet solution. First, 'agentic AI' involves multiple, cooperating AI agents, each assuming a specific security persona such as threat hunter, test engineer, remediation developer, code reviewer, etc. This approach ensures context-rich detection, precise remediation, and multi-stage verification. Fixes must pass through a validation loop, reducing the risk of hallucinated or incorrect code changes.

In addition, the 'Code Property Graph' is the data structure that unifies syntax, data flow, and control flow into a single model. AI agents can traverse the code's logical structure to reason about vulnerabilities. For example, the CPG can reveal whether user-supplied input reaches a dangerous sink function without sanitization, thus detecting an exploitable flaw rather than a theoretical one.

Qwiet's implementation of both Agentic AI and CPG allows deep and scalable analysis, enabling detection of multi-function vulnerabilities, such as chained injections or access control bypasses. It also significantly reduces false positives by filtering for actual reachability. While seemingly novel today, we would expect that these capabilities will gradually emerge as requirements in every software development environment.

PLATFORM ADVANTAGE

While several commercial vendors now offer some version of auto-remediation, including GitHub via Dependabot, Qwiet AI's preZero platform provides a comprehensive security capability beyond basic software patching. It integrates SAST, SCA, secrets detection, container security, IaC scanning, and SBOM generation into a single development workflow.

Another functional capability is that the AutoFix is triggered by any detected issue, whether in code, configuration, or dependency. This allows the system to generate patches for Terraform scripts, Dockerfiles, third-party libraries, and application code. All fixes are tagged with the appropriate compliance context (e.g., OWASP A1, NIST 800-218), and logs provide a complete audit trail.

Also, 'AutoFix' is delivered asynchronously via cloud-based processing, meaning developers receive fixes as non-blocking pull requests. AutoFix minimizes disruption and allows DevSecOps teams to embed security into CI/CD pipelines across GitHub, GitLab, Bitbucket, or Azure DevOps. This is an important differentiating feature for Qwiet, and one that we expect to be relevant to most development environments.

REAL-WORLD IMPACT

TAG analysts reviewed a recent case study of Qwiet AI deployed into a large telecommunications company. The results reported from the use of the platform for development included the following:

- 90% false positive reduction from legacy scanners, thanks to reachability analysis.
- 95% faster scan times compared to the incumbent SAST solution
- An order-of-magnitude reduction in remediation time for high-priority vulnerabilities.
- Savings of over $300K annually in developer remediation hours.
- Enhanced compliance with OWASP, NIST, and ISO 27001 based on audit-ready logs.

These results highlight the potential for auto-fix to transition application security from a bottleneck to a value enabler.

CONCLUSION

The future of secure software development lies in automation, and auto-fix capabilities like those from Qwiet AI provide a critical step forward. Qwiet's AutoFix delivers safe, explainable, and compliant code remediation within minutes by combining agentic AI reasoning with deep contextual analysis via the Code Property Graph. As we have explained, this combination will gradually emerge as standard functional requirements.

Also, as CI/CD pipelines become the norm, integrating platform-agnostic, intelligent auto-fix will be essential for modern DevSecOps. And while other vendors offer partial solutions, Qwiet AI's focus on security-first auto-remediation positions it at the forefront of this shift. TAG recommends that enterprise security leaders look seriously at platforms like Qwiet as part of their roadmap to autonomous application security.

ABOUT TAG

Recognized by Fast Company, TAG is a trusted next generation research and advisory company that utilizes an AI-powered SaaS platform to deliver on-demand insights, guidance, and recommendations to enterprise teams, government agencies, and commercial vendors in cybersecurity and artificial intelligence,.