



Supporting document - Bitsec Brief 31

AI development is speeding up the rate of lines of code being produced at exponential rates. This amount of code is outstripping the rate in which security reviews can be conducted by humans.

AI coding agents need AI security agents to keep them in check - according to reports AI coding assistants produce "significantly less secure" code than humans!

Bitsec is the answer.

1. Problem Context and Market Opportunity

Cybersecurity is a \$200B+ annual market, yet outcomes remain structurally poor. In both Web2 and Web3, security spend is high, time-to-audit is slow, and exploit frequency continues to rise. In blockchain alone, billions are lost annually despite widespread use of traditional audits.

The root issue is not lack of effort, but lack of scalability. Human-led audits do not scale with modern software velocity, and they do not align economically with outcomes. Audits are slow, expensive, and often disconnected from real exploit discovery.

- The cyber security market is estimated to be worth \$196.51 billion in 2025 worldwide, with strong growth in security services.

Source: <https://www.statista.com/outlook/tmo/cybersecurity/worldwide>

- Data Bridge Market Research: Values the market at \$203.86 billion in 2025, growing to \$421.82 billion by 2032 (9.5% CAGR).

Source:

<https://www.databridgemarketresearch.com/reports/global-cybersecurity-market>

- \$227 billion total for Cybersecurity Market

Source:

<https://www.marketsandmarkets.com/Market-Reports/cyber-security-market-505.html>

2. Why AI security agents are needed

AI-assisted development is accelerating code production dramatically. As AI-generated code becomes a dominant input to software development, the gap between code written and code secured widens. This creates a second-order security crisis that existing tools and firms are not built to address.

AI code generation is reaching mainstream adoption. Security failures are increasing, not decreasing. Traditional audit models are failing economically and operationally. And Bitsec now provides the incentive infrastructure required to solve agent reliability at scale.

- 40% of 1,689 programs generated by GitHub Copilot were vulnerable to MITRE's 2021 CWE Top 25.

Source:

<https://cset.georgetown.edu/wp-content/uploads/CSET-Cybersecurity-Risks-of-AI-Generated-Code.pdf>

- Approximately 48% of code snippets generated by leading large language models (LLMs) like GPT-4 and Code Llama contained bugs that could lead to malicious exploitation, such as buffer overflows or memory leaks.

Source:

<https://cset.georgetown.edu/wp-content/uploads/CSET-Cybersecurity-Risks-of-AI-Generated-Code.pdf>

- Developers using AI assistants produced "significantly less secure" code than those without, and they were more likely to overestimate the security of their outputs.

Source:

<https://ee.stanford.edu/dan-boneh-and-team-find-relying-ai-more-likely-make-your-code-bugger>

Bitsec sits at the intersection of AI coding agents and the ever growing need for security, with a clear path from benchmarks to revenue and from narrow use cases to a broad security platform.

Bitsec exists to close this gap using incentivized, benchmarked AI security agents.

2. Why AI Security Agents Outperform Human Auditors

Security is fundamentally an adversarial, probabilistic problem. It rewards pattern recognition, exploit intuition, and systematic coverage rather than linear review processes.

Bitsec's architecture reflects this reality. Rather than static audits, Bitsec uses AI agents that can:

- Continuously scan and rescan codebases
- Generalize across languages and architectures
- Improve through competition and benchmarking
- Operate at speeds and scales humans cannot match

Early results from Bitsec v1 demonstrated that even relatively simple AI setups could rediscover real-world exploits, including high-impact vulnerabilities that had already caused major losses. In several cases, exploits were detectable quickly using models that were not explicitly trained on those codebases.

This suggests a structural advantage: once reliability improves, AI agents can systematically outperform humans on coverage, speed, and cost.

3. From Bitsec v1 to Bitsec v2

Bitsec v1 was designed as a proof of concept. Its primary goal was to answer a binary question: can AI agents find real vulnerabilities in real codebases?

The answer was yes.

V1 surfaced multiple real exploits across different ecosystems, including post-mortem discoveries and pre-exploit findings. It also highlighted a critical industry gap: there was no objective benchmark to compare AI security performance in a way that mapped to economic outcomes.

Bitsec v2 addresses this directly.

V2 introduces an agent-based architecture aligned with Bittensor incentives, where miners submit security agents that are evaluated against a real-world smart contract audit benchmark. This benchmark is built from historical audit challenges using open-source codebases and verified findings.

Performance is now measurable, comparable, and improvable.

4. Benchmarks as a Direct Revenue Lever

In Bitsec, benchmarks are not abstract metrics. They are a proxy for real-world earning potential.

Audit challenges and bug bounties already operate on outcome-based economics. Finding high or critical vulnerabilities directly maps to revenue, often in the form of bounty payouts or long-term contracts.

As Bitsec agents improve on benchmarks:

- They become more competitive in live audit challenges
- They win more bug bounties
- They establish credibility with paying customers
- They create external proof of effectiveness via public leaderboards

This creates a direct feedback loop where better benchmark performance leads to higher real-world revenue and stronger demand for the product layer.

5. Relationship to Ridges: Similarities and Differences

Bitsec shares structural similarities with Ridges in that both:

- Use agent-based submissions from miners
- Rely on real-world benchmarks
- Improve via competitive hill-climbing dynamics
- Align incentives through Bittensor emissions

However, the problem domain is materially different.

Security is an open-ended task. Unlike code generation, there is no fixed output. Agents must discover unknown vulnerabilities across arbitrary architectures and threat models, then match or exceed human findings.

This makes Bitsec's task harder, but also more defensible. Success in this domain compounds trust, credibility, and economic value more strongly than incremental improvements in generative tasks.

There are also market structure and end user differences:

Ridges has 5 large competitors - Cursor, Anthropic (Claude Code), Gemini, OpenAI (Codex). The end user is also the one paying and can assess quality beyond benchmarks.

Bitsec has competitors, but no clear leaders. The end user pays, assesses the product via

proxy from benchmarks and audit competition leaderboards.

6. Product Direction and Q1 Launch

The immediate focus for Q1 is achieving state-of-the-art performance on the smart contract audit benchmark.

Once reliability and performance reach competitive thresholds, Bitsec will move directly into productization. This includes:

- A GitHub-based workflow for scanning repositories and pull requests
- A user-facing interface for submitting codebases
- Automated reporting driven by agent outputs
- Participation in live audit challenges and bug bounties

These components reinforce each other. Benchmark success builds trust. Trust drives usage. Usage generates data and revenue. Revenue reinforces incentives.

7. Long-Term Vision

Bitsec is not limited to blockchain.

The same incentive flywheel can be extended to:

- Traditional Web2 security
- Penetration testing
- Model jailbreaking
- Infrastructure security
- Any domain where exploit discovery is measurable

As AI continues to accelerate code production, security must become continuous, automated, and economically aligned. Bitsec is positioned to become a central layer of trust for code auditing across ecosystems.

8. Additional information:

Overview video made by the founder explaining Bitsec:

<https://www.youtube.com/watch?v=tD9bXAcnxRk>

Website: <https://bitsec.ai/>

X: <https://x.com/bitsecai>