

Endor Labs for Financial Services

Regulated AppSec that hits SLAs without breaking the platform

Financial services operates under near-zero tolerance for vulnerabilities and aggressive remediation SLAs from PCI DSS, FFIEC, and DORA — while the core banking, payments, and trading platforms the business runs on cannot tolerate broken upgrades. Endor Labs makes regulated AppSec operational: safe remediation, logic-flaw detection, and AI governance, so security keeps pace with regulators without breaking the platforms the business depends on.

6x

faster CVE remediation with safe, evidence-based fixes


95%

maximum false positive reduction for AI SAST

Every

AI coding agent action logged, governed, and auditable


How it works



PATCH

Fix what's critical, keep the platform up

- Upgrade Impact Analysis shows what an upgrade breaks before you commit
- Endor Patches backport fixes when an upgrade won't fit the change window
- Close transitive CVEs and meet PCI DSS, FFIEC, and DORA SLAs



DETECT

Catch the logic flaws that matter most

- AI SAST finds IDORs, authentication bypasses, and business logic errors
- Every pull request reviewed continuously — including AI-generated code
- Fewer false positives, so reviewers focus on findings that move risk

Agent activity		AUDIT LOG
✓		ALLOWED
✓		ALLOWED
✗		BLOCKED
✓		ALLOWED
✓		ALLOWED

GOVERN

Govern AI agents with audit-ready evidence

- Searchable audit log of every action AI coding agents take
- Real-time guardrails on dangerous commands, file access, and MCP servers
- Centralized visibility across Cursor, Claude Code, and other coding agents

WHY NOW

In the Mythos era, vulnerability discovery is trivially cheap

Frontier AI models like Anthropic's Mythos can discover vulnerabilities and generate exploits in open source components at unprecedented scale — attackers now build working exploits in hours, for under \$5. Finding volume is exploding while regulatory SLAs stay fixed. You can't patch everything: you need reachability-based prioritization and fast, safe fixes.

- Reachability evidence lets you defensibly deprioritize what isn't exploitable
- Principal Financial Group, Robinhood, Citi, and Nationwide trust Endor Labs

Built for regulated environments

PCI DSS v4 evidence QSAs accept — generated with every scan

What delivers each pillar

 Upgrade Impact Analysis Shows what an upgrade will break before you commit to it.	 Endor Patches Backported security fixes for the version you already run — no upgrade required.	 AI SAST Finds IDORs, authentication bypasses, and business logic errors legacy SAST misses.
 Secrets Detection Detects exposed secrets in code before they reach production.	 Agent Governance Real-time guardrails and audit logs for every AI coding agent.	 SBOM & VEX Audit-ready inventory and vulnerability records generated with every scan.

COMPLIANCE

PCI DSS v4: manage everything without fixing everything

Requirement 11.3.1.1 expanded vulnerability management to every finding in scope — not just highs and criticals. “Manage” means a defensible, documented assessment. Function-level reachability — accepted by leading PCI QSAs including Schellman — documents unreachable findings as managed, and Endor Patches close the rest inside the assessment window.

- Cut remediation workload by an average of 70–80% by filtering unreachable findings
- Up to 90% less container vulnerability noise across image OS layers
- SBOM, VEX, and vulnerability management records generate with every scan