

SOLUTION BRIEF

Contrast CVE Shield

Stop exploits before patches exist

Block CVE attacks in real time — with runtime-confirmed exposure, behavior-based interception, and zero disruption to legitimate application functionality.

51%

Of published
2024–2025 CVEs can
be reproduced by AI¹

28.3%

Of CVEs are exploited
within 24 hours
of disclosure^{2,3}

70%

Of security team
hours spent on
false positives⁴

The problem: The window between disclosure and patch

Security teams face a structural exposure gap that no amount of patching can fully close. The median time from public CVE disclosure to active exploitation has shrunk to hours³ — far shorter than any enterprise patch cycle. The average remediation cycle stretches to 241 days.⁵ The math does not work, and attackers know it.

AI has made this dramatically worse. Tools capable of generating working exploits for known CVEs can do so within minutes of disclosure, at an average cost under \$2,000.¹ In 2025, 28.3% of CVEs were being actively exploited within 24 hours of disclosure.³ No patch cycle covers that window.

Why runtime context changes everything

The fundamental problem with perimeter and static approaches is that they reason about what is present, not what is happening. A dependency scan tells you which libraries are in the build. It cannot tell you which of those libraries execute in a live request, which execution paths an attacker can trigger, or whether a specific exploit chain would succeed.

Runtime context answers the questions that matter: Is this vulnerable code path actually exercised in production? Has it been triggered in an exploitation attempt? Did the attempt succeed or was it blocked? Without this layer, security teams cannot distinguish a theoretically dangerous library from one that is actively under attack.

On average, only 38% of libraries in a running application are active in production. Only 30% of all CVEs present in an application are in code paths that actually execute.⁶ Without runtime visibility, security teams spend the majority of their remediation effort on vulnerabilities that could never be reached by an attacker — while the ones that can be reached may go unprotected during the disclosure-to-patch window.

The Contrast CVE Shield solution

Contrast CVE Shield is a runtime protection capability that closes the gap between vulnerability disclosure and patch deployment. It operates inside the running application — confirming which vulnerable libraries actually execute in production, detecting when a CVE is being actively exploited and blocking the dangerous system-level capabilities exploit chains require. Security teams move from reactive investigation to proactive enforcement.

Contrast CVE Shield protects against vulnerabilities including Log4Shell, Spring4Shell and Apache Commons Collections deserialization and hundreds of additional high- and critical-severity CVEs, with new protections delivered continuously as vulnerabilities emerge.

Key benefits

- **Immediate zero-day protection:** New shield definitions reach deployed agents within minutes of release. No emergency change window.
- **Runtime reachability confirmation:** On average, only 30% of CVEs in an application are active in production. CVE Shield shows exactly which ones are exploitable and which are noise.
- **Behavior-based blocking:** Intercepts what exploits do — OS commands, JNDI lookups, remote class loading — not what payloads look like. Novel AI-generated encoding variants are blocked automatically.
- **Near-zero performance impact:** +12 nanosecond overhead on active execution paths, while non-targeted code paths experience zero performance impact.
- **Automatic MITRE-mapped incidents:** Every blocked exploit generates a pre-enriched ADR incident with CVE ID, CWE, and MITRE ATT&CK tactic and technique. No manual triage required.

How it works

Deploy

One-time agent install. No code changes. The Contrast Kubernetes Operator can roll CVE Shield across an entire cluster in a single operation. Once deployed, all current and future shield definitions are applied automatically — no per-CVE configuration required.

Monitor

Run Monitor mode for 1-2 weeks. CVE Shield logs every capability violation without blocking, giving the team a baseline view of what it would have blocked over real production traffic. This confirms there are no false-positive concerns before enforcement is enabled. Operators advance individual shields from Monitor to Block independently, per CVE — no global toggle required.

Block

Enable Block mode per CVE, starting with the highest-severity vulnerabilities. CVE Shield enforces a per-CVE allow list of system-level capabilities — OS command execution, JNDI lookups, remote class loading, file writes, deserialization and more. The library continues to function normally for legitimate operations; only the specific capabilities required by an exploit chain are restricted. Because protection is behavior-based rather than signature-based, novel payload variants and AI-generated exploits are blocked automatically. Each shield is independent — any individual CVE can be rolled back without affecting others.

Respond

Blocked exploits auto-close as incidents in Contrast ADR. Every blocked attempt generates a pre-enriched incident with CVE ID, CWE classification and the specific capability that was stopped. SOC teams get a high-signal queue with complete forensic context — no manual log correlation. Successfully blocked exploits are auto-closed, reducing analyst workload and eliminating noise from low-priority findings.

How CVE Shield compares

Capability	SCA	WAF	eBPF tools	CVE Shield
Confirms library is actively used in production	No	No	No	Yes
Confirms CVE is in active code	No	No	No	Yes
Sees attacker behavior inside app	No	No	Limited — kernel only	Yes — full call stack
Blocks unknown / AI-generated exploits	No	Limited — signature-based	No — out-of-band only	Yes — behavior-based, inline
Distinguishes normal use from abuse	No	No	No	Yes — per-CVE allow list
Blocks exploits on undiscovered CVEs	No	No	No	Yes — blocks the exploit technique, not just the input
Performance impact	None	1-2ms per request	High — all syscalls	+12 ns on vulnerable paths only
Deployment	No agent	Network appliance	Kernel module required	Single agent, no code changes

SCA tells you what you have. CVE Shield tells you what is being attacked right now and stops it.

"Capabilities that connect CVE identification with runtime exploitability and active protection represent an important step toward more operationally relevant application security."

— Katie Norton, IDC Analyst

Conclusion

With CVE Shield, security teams gain a definitive answer to the question boards and auditors ask: "Are we protected from this CVE?" Not "we think so" — "we know so." CVE Shield is part of the Contrast runtime security platform, which equips AppSec, SecOps and development teams to proactively protect applications against evolving threats without slowing delivery.

¹[Contrast Security, Get Mythos ready: Secure your apps against AI](#)

²[Cloud Security Alliance, The Collapsing Exploit Window: AI-Speed Vulnerability Weaponization](#)

³[The Hacker News, 2026: The Year of AI-Assisted Attacks](#)

⁴[Security Boulevard, Why Your Security Team is Wasting 70% of Their Time on Phantom Threats And How to Fix It](#)

⁵[IBM, Cost of a Data Breach Report 2025](#)

⁶[Cyentia, A Visual Exploration of Exploitation in the Wild](#)