

Defend against AI-detected threats

66

Every week we have to do security scans ... list vulnerabilities in packages that are on the operating systems. Any time that occurs, we can immediately go to the Red Hat site, and get full explanations of the vulnerabilities, which help determine if we're affected. Nobody else does that for other Linux distributions. It would quickly become a nightmare trying to run a production government environment without it because we would have to do all that research every month ourselves.²

Ryan Kimbrell
Senior Cloud Operations
Engineer at Zoom

Defending against AI vulnerabilities

The threat landscape has fundamentally shifted. AI-powered discovery tools are now finding potential vulnerabilities at a rate that is exponentially higher than historical norms. Enterprise security teams were already being pushed to the brink in 2025 with an average of 132 CVEs discovered per day.¹ In Q1 of 2026, that number has already climbed 26%, and continuing advances in AI vulnerability discovery tools like Mythos may make traditional approaches to managing CVEs unsustainable.²

In 2020, teams typically had a 44-day head start remediating a vulnerability before it was exploited.³ By 2025, this head start became a 7-day deficit.⁴ Now, bad actors can take advantage of new vulnerabilities a week before a new CVE is even made public.

Meanwhile, enterprise defenders continue to operate at human speed—manual triage, change management meetings, and patching SLAs measured in weeks or months.

Traditional IT security approaches are unable to keep pace with the triage required to manage the rising volume of CVEs. On average, at large organizations, 37% of the vulnerabilities discovered in a year remain unpatched, and the software industry average for mean time to remediation (MTTR) is 55 days.⁵ This security gap is unsustainable, leaving organizations struggling.

4 critical challenges facing enterprise security teams today

Organizations are currently struggling with 4 compounding pressures:

1. **Faster machines.** AI models are discovering new vulnerabilities at speeds that human defenders and manual processes can not match.
2. **Volume inundation.** Massive surges in vendor patch releases can overwhelm security teams, triggering severe alert fatigue and straining traditional vulnerability management workflows.
3. **Stability versus security.** Teams delay patching to avoid breaking production applications.
4. **Configuration drift.** Traditional environments and over-permissioned accounts can map into AI-exploitable paths far faster than traditional manual assessment.

The answer is not more security software. For existing Red Hat customers, the integrated remediation capabilities required are likely already in their portfolio—they only need to be activated.

Build a security-focused foundation with Red Hat Enterprise Linux

Red Hat® Enterprise Linux® can help minimize the effects of CVEs while organizations assess, test, and deploy patches. Image mode for Red Hat Enterprise Linux delivers the platform as a signed, versioned container image with a software bill of materials (SBOM), while [Security-Enhanced Linux](#) (SELinux) and [fapolicyd](#) (File Access Policy Daemon) add runtime layers that help block unauthorized execution and confine compromised processes.

Identify real exposure with Red Hat Lightspeed

Before remediation can happen at scale, teams need to know where they are actually exposed—not in a generic vulnerability database, but in their specific environment, against their real configurations and active workloads.

Red Hat Lightspeed helps organizations prioritize CVEs on Red Hat Enterprise Linux systems by combining package and advisory data with Red Hat-curated security rules, workload grouping, and configuration-aware assessment—including controls such as SELinux and service exposure. Unlike version-based scanning, it distinguishes systems with exploitable vulnerabilities from those that are vulnerable, but not affected. This allows teams to focus first on risks where they can take action. In testing on 90 hosts, automated detection and remediation using Red Hat Lightspeed reduced hands-on administrator time by up to 79% for CVEs and 86% for known operational issues, compared with manual scripting.⁶

This contextual precision matters now more than ever. The sheer volume of newly discovered vulnerabilities has created a new form of operational paralysis: alert fatigue. When hundreds of CVEs are disclosed in a compressed window, no team can evaluate them all manually. Red Hat Lightspeed cuts through that noise, allowing teams to focus remediation effort on the findings that will significantly affect their specific environment.

Vet and control patches with Red Hat Satellite

Knowing what to patch is half the challenge. Getting the right patch tested, verified, approved, and into production without creating new instability requires a controlled, repeatable process.

Red Hat Satellite provides on-premise errata management and content staging at scale across local, hybrid, and fully air-gapped infrastructure. Satellite creates content views—verified software snapshots at a specific point in time—and promotes them through gated lifecycle environments (from development to testing to production). Patches must pass through defined approval gates before reaching production so they can be fully vetted before deployment.

Satellite also helps protect the software supply chain: Packages that enter the environment are cryptographically signed and sourced from Red Hat's verified repositories. For organizations operating in highly regulated or disconnected environments, Satellite delivers the same capabilities to environments without direct internet access—providing control over what software is permitted into the estate. Compliance reporting is automated using the Open Security Content Automation Protocol (OpenSCAP) scanning framework, generating audit-ready reports against continuous integrations (CI), the Defense Information Systems Agency's Security Technical Implementation Guides (DISA STIG), and payment card industry (PCI) developer sandbox baselines.

Remediate at scale with Red Hat Ansible Automation Platform

Once patches are vetted and staged, the final challenge is applying them consistently across a complex hybrid estate spanning thousands of servers, multiple cloud environments, on-premise

datacenters, and air-gapped environments—without creating a manual backlog that re-creates the velocity gap.

Red Hat Ansible® Automation Platform applies remediations at scale across the entire hybrid estate, using Policy as Code (PaC) to provide consistency, auditability, and repeatability. Event-Driven Ansible (EDA) closes the loop by automating the remediation workflow itself: When Satellite finishes a repository sync containing a new security patch, a webhook automatically signals Ansible Automation Platform to identify vulnerable systems and trigger the remediation playbook. High-risk systems can be automatically quarantined before a patch is even deployed, limiting exposure during the critical response window. Each action is executed using Red Hat Ansible Certified Content Collections, maintaining a clear PaC audit trail that helps satisfy compliance requirements without additional tooling.

The Red Hat closed-loop remediation solution is not a patchwork of separate products. It is a single integrated portfolio delivered under 1 subscription and supported by 1 team. For organizations already running Red Hat solutions, the tools required to defend against AI-driven threats are likely already in place. Connecting them shortens the mean time to remediation from weeks of manual effort to automated hours—with a comprehensive, provable audit trail across each system in the estate.

Read more about Red Hat Lightspeed

Ready to accelerate your CVE remediation? [Learn more about Red Hat Lightspeed.](#)

Learn more about Red Hat Satellite

To read more about how Red Hat Satellite streamlines infrastructure management at scale, [visit the product page.](#)

Explore a path to fully optimized automation

To discover how Red Hat Ansible Automation Platform helps reduce human error in repeat processes, [visit the product page.](#)

1. CVE Program. ["Metrics."](#) CVE: Common Vulnerabilities and Exposures, accessed June 2026.
2. Red Hat case study. ["Zoom for Government enhances security with a Red Hat platform."](#) 10 March 2023.
3. Google Cloud Blog. ["How Low Can You Go? An Analysis of 2023 Time-to-Exploit Trends."](#) 15 Oct. 2024.
4. Google Cloud Blog. ["M-Trends 2026: Data, Insights, and Strategies From the Frontlines."](#) 23 March 2026.
5. Edgescan. ["2026 Vulnerability Statistics Report"](#) Edgescan, 2026
6. Principled Technologies, sponsored by Red Hat, ["Save administrator time with the automated remediation capabilities of Red Hat Lightspeed,"](#) Nov. 2025 (revised).



About Red Hat

Red Hat is the open hybrid cloud technology leader, delivering a trusted, consistent and comprehensive foundation for transformative IT innovation and AI applications. Its portfolio of cloud, developer, AI, Linux, automation and application platform technologies enables any application, anywhere—from the datacenter to the edge. As the world's leading provider of enterprise open source software solutions, Red Hat invests in open ecosystems and communities to solve tomorrow's IT challenges. Collaborating with partners and customers, Red Hat helps them build, connect, automate, secure, and manage their IT environments, supported by consulting services and [award-winning](#) training and certification offerings.

North America	888-REDHAT1	www.redhat.com	in @red-hat
Asia Pacific	+6564904200	apace@redhat.com	℥ @redhat
Latin America	+5443297300	info-latam@redhat.com	▶ @redhat
Europe, Middle East, and Africa	+0080073342835	europa@redhat.com	📷 @red_hat

Copyright © 2026 Red Hat. Red Hat, the Red Hat logo, Ansible, and OpenShift are trademarks or registered trademarks of Red Hat, LLC or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. The OPENSTACK logo and word mark are trademarks or registered trademarks of OpenInfra Foundation, used under license. All other trademarks are the property of their respective owners.