




The Mythos Moment: Project Glasswing and Project Lightwell

Marc Skinner

Setting the stage



Key Takeaways:

- 2020, teams had a **44-day head start** in preventing an exploit. By 2025 it was a **7-day deficit**.
- **Mean Time To Remediate (MTTR)** high/critical CVEs was **54.8 days** in 2025.

Source for CVEs: <https://www.cve.org/about/Metrics>

Source for MTTR: <https://info.edgescan.com/>

Source for MTTE: <https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026>



45%

of vulnerabilities in large organizations remain
unpatched after 12 months.

Source: Centre for Emerging Technology and Security (Alan Turing Institute), 2025 – cited in CETAS analysis, April 2026.

Then Mythos arrived.

- Anthropic, April 7, 2026.

What a single frontier AI model can now do.

Thousands

of zero-day vulnerabilities
autonomously discovered.

Every

major operating system and every
major web browser.

Hours

from public CVE to working exploit
– not weeks.

Source: Anthropic Frontier Red Team, "Claude Mythos Preview," red.anthropic.com, April 2026.



Project Glasswing

Securing critical software for the AI era

Securing the world's most critical software

- Amazon Web Services
- Anthropic
- Apple
- Broadcom
- Cisco
- CrowdStrike
- Google
- JPMorganChase
- Linux Foundation
- Microsoft
- NVIDIA
- Palo Alto Networks

Mythos Preview

- Jumpstart defensive security work
- Anthropic committing to \$100M usage credits to partners
- No OS or software is safe

Planned Outcomes

- Vulnerability disclosure processes
- Software update processes
- Open-source and supply-chain security
- Software development lifecycle and secure-by-design practices
- Standards for regulated industries
- Triage scaling and automation
- Patching automation

Project Lightwell

Red Hat and IBM securing the open
source supply chain

IBM Investment

\$5B

AI discovers vulnerabilities at **machine speed**

- Frontier AI models find critical zero-day vulnerabilities autonomously – including bugs hidden for 27 years – at under **\$50** each
- Less than **1%** of AI-discovered vulnerabilities have been patched
- The number of published CVEs has grown more than **520%** since 2016
- The National Vulnerability Database is shifting to **selective enrichment**
- **19.7%** of AI-generated code references hallucinated package names

You need a remediation **partner**

Remediating alone:

Your team triages the Mythos backlog internally. Three bad options for each finding:


- Upgrade and break production.
- Patch internally and carry a permanent fork.
- Accept the risk of a confirmed **exploit**.

Remediating with **Red Hat**:

We backport fixes to the exact versions you run in production:

- No upgrades required.
- SLA-backed delivery.
- Signed artifacts.
- Every patch contributed upstream so you're never stuck on a fork.


Lightwell – a clearing house for vulnerabilities



Shared work compounds member value.

- Trusted intermediary for patches, upstream coordination, and security embargos
- Intake, triage, and aggregation is deduplicated across members

Closing the vulnerability gap in production



Securing the ecosystems you depend on

- Curated package registries for Python, JavaScript, and Java that's scanned, patched, signed, certified
- Maintainer partnership program: funding upstream security work to fix bugs at the source
- AI-powered vulnerability triage with automated severity assessment and remediation prioritization
- Red Hat — Hardened Images secures the container. This offering secures the dependencies inside it. Together → full-stack coverage from OS to application layer

20,000 engineers to close the gap

- The largest engineering commitment to open source security since IBM's \$1B Linux investment in 1999
- **AI handles initial triage and acceleration. Engineers handle the judgment calls** – understanding community context, navigating responsible disclosure, and landing patches upstream
- **Every fix will be upstreamed if at all possible.** Commercial security and community health reinforce each other

Securing the **Open Source** Supply Chain

IBM and Red Hat are putting **20,000 engineers** to work making open source software safer

The Problem

96% of codebases contain open source code

9.8 trillion package downloads in 2025

581 avg. vulnerabilities per codebase, doubled in one year

AI finds critical bugs for under **\$50**. Less than **1%** get fixed

Attackers exploit flaws **7 days** before patches exist

What Red Hat Does

AI-powered scanning at scale across your full dependency tree

Patched packages at the **exact versions** you already run

SLSA Level 3 builds, HSM-backed signing, SBOMs for every package

24-hour critical CVE response SLA

Every fix contributed upstream.
No private forks.

AVAILABLE TODAY:

Hardened Images

Zero-CVE containers

OpenShift Security

Runtime protection

RHEL Subscriptions

Enterprise Linux foundation

Coming next:

Curated registries for Java, Python, JavaScript

No other vendor scans, fixes, signs, and certifies open source packages end-to-end.
Red Hat does.

What can you do now?





You only have 24 hours, how would you?

- Scan your entire application stack or infrastructure fleet for a specific file, text string, application version or other important artifact?
- Run a script, arbitrary command on any system or groups of systems to initiate a remediation or mitigation configuration changes?
- Push a patch for an application, OS or rebuild a container and redeploy from a secure, trusted source?
- Can you run a report, or validate all of the above actions have successfully completed?

Customers already own the answer.

OpenShift

The application platform

RHEL

The hardened floor

Satellite

The control plane

Ansible AAP

The remediation muscle



01

RHEL

The hardened floor.

- SELinux + ASLR out of the box
- Backported security fixes
- EUS & Extended Life-cycle
- Red Hat Lightspeed (Formerly Insights)

The very defenses that blocked Mythos exploits against the Linux kernel.

Red Hat triages upstream, fixes in place – no risky wholesale upgrades.

No end-of-life surprise on core policy or claims systems.

Proactive vulnerability, advisor, and compliance – already in your subscription.

02

Satellite

The control plane for the patch lifecycle.

- Content views & lifecycle environments

Staged dev→test→prod errata promotion. The audit trail examiners ask for.

- Errata management at scale

Apply a Red Hat security advisory across thousands of hosts in one action.

- Disconnected / air-gapped support

For core admin and claims systems regulators require isolated.

- OpenSCAP compliance reporting

Automated CIS, PCI, and DISA STIG benchmarking. Audit-ready, not audit-panicked.

03

Ansible AAP

The remediation muscle.

- Event-Driven Ansible

Lightspeed flags → EDA triggers → playbook patches. No ticket. No lag.

- Certified Content Collections

Vetted, supported automation – not random roles off Galaxy.

- Hybrid reach

RHEL, Windows, AIX, network, AWS, Azure, VMware, OpenShift – one control plane.

- Policy-as-code

Every change in Git. Every run logged. Audit evidence by default.



04

OpenShift

The application platform

- Build / Deploy

Develop, build and deploy container images or virtual machines

- Scale

Grow or shrink your workloads based on demand

- Secure

Leverage RBAC and hardened images to protect access and running applications

- Lightspeed AI

AI recommended optimization for system administrators

Before

Weeks

Manual cycles. Tribal knowledge. Tickets in flight.

After

Hours

Closed-loop remediation, audit-logged, across the hybrid estate.

— Illustrative: measured by customer-reported mean time to remediate a critical CVE.

Level Up

Security TAM

Security Expert

Add an Open Source Security
Expert to your team

Red Hat Security TAM

Goals

Benefits and Outcomes



Reduce your security risk



Keep your team informed with a single point of contact that provides personalized notifications, insights into current and future Red Hat products, and resolutions for critical security issues.



Effectively detect, respond to, and recover from incidents with minimal impact



Security TAMs provide increased responsiveness and can help speed up root cause analysis for security incidents.



Strengthen your security program in response to new threats and technologies



Meet your long-term security goals across your Red Hat portfolio with a Security TAM, who provides strategic and operational planning sessions.



Continuously build your existing team's expertise



Leverage our invite-only events to meet product experts and network with peers. Security TAMs provide ongoing education to ensure your team has the skills needed to mitigate risk.

Remember – Customer Focus

Key Needs: Speed + Innovation

Hardened Images

Trusted, micro-sized container images optimized to minimize security vulnerabilities upon release.

Trusted Libraries

SLSA Level 3 signed Python packages rebuilt from source for secure app dependencies. (Java, go, npm on roadmap)

Key Needs: Stability + Support

RHEL & Portfolio Security

Advanced security capabilities built-in help organizations to lower risk, save time, and mitigate downtime while supporting compliance. Plus options to contextualize risks, verify fixes, and automate remediation.

Project Lightwell

A \$5 billion commitment backed by new frontier AI capabilities and a global force of more than 20,000 engineers to help enterprises secure open source software they rely on.






Red Hat Hardened Images

Modern software, delivered as hardened, verifiable container images

Foundational set of open source languages, runtimes, databases, web, & utilities aiming to reduce toil around CVEs.

- Built on a **minimal, hardened Linux** userspace optimized for containerized environments.
- Available at **no cost**. Support is included with Red Hat Enterprise Linux & OpenShift Container Platform
- Delivered through Red Hat's SLISA 3 trusted supply chain.
- Visit images.redhat.com




Red Hat Trusted Libraries (RHTL)

Trusted Python Libraries. Built from Source. Verified End-to-End

What is it?

SLSA Level 3–verified, secure open-source libraries built, signed, and distributed through Red Hat's trusted supply chain. Starting with ~1,000 top Python packages (numpy, pandas, flask, requests) scaling to 10,000+ at GA (May 2026), with Go, npm, and Java planned.



The Red Hat Difference

Feature	Public Repos	Commercial Vendors	Trusted Libraries
Build Standard	Unknown / Opaque	SLSA Level 2 (Often)	SLSA Level 3 (Highest)
Provenance	None	Limited Metadata	Full Build-from-Source
Security Response	Reactive / Manual	Varies	Proactive (<24 hr builds)
Ecosystem	Isolated	Standalone Product	Native to RHADS, RHEL & OpenShift



Available as part of RHADS Subscription - packages.redhat.com

Thank You

