



Connect

Agentic AI en RHEL

**Seguridad empresarial, de la detección de CVEs a
remediaciones automáticas**

Josep Andreu Font

Senior Technical Account Manager

José Ángel de Bustos Pérez

EMEA Senior Specialist Solution Architect



Josep Andreu Font

Senior Technical Account Manager
Red Hat




José Ángel de Bustos Pérez

EMEA Senior Specialist Solution Architect
Red Hat



More vulnerabilities, less time to react



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.



The New AI Security Moment Has Arrived: i.e.: Mythos

Pours gasoline on the fire with increased volume of CVEs



May 25, 2026

[Security Week](#): Anthropic: Mythos Detected 23,000 Potential Vulnerabilities Across 1,000 OSS Projects

May 25, 2026

[Business Standard](#): Anthropic's Mythos finds 10,000+ vulnerabilities (high and critical), flags security bottleneck in apply fixes more than discovery

April 26, 2026

[Lyrie Threat Intelligence](#): Mythos Unleashed: How Anthropic's Restricted AI Discovered 2,000 Unknown Vulnerabilities in 7 Weeks




"Vulnerabilities are now
being exploited *before*
they are even reported."



The false positive (FP) problem

Primary objective: obtain reliable and accurate scanning results



Teams spend more time triaging false positives than fixing real vulnerabilities

Audit fatigue leads to real issues being missed

Scanner output without context erodes trust

Compliance teams reject patched containers



Why scanners need more context



Backported patches

RPM version !=
upstream version.

Scanners compare
strings, not patch
status



AppStream modules

Module streams create
version ambiguity
scanners can't resolve



Missing VEX context

Can't distinguish
"affected" from "not
affected" without
advisory data



Will not fix

Low-risk CVEs accepted by
vendor conflated with
unpatched vulnerabilities




The real-world impact



Wasted effort

Security teams spend 40%+ of triage time on false positives from scanners that don't understand backporting



Delayed deployments

Container images blocked by CI/CD gates due to phantom vulnerabilities already patched via backports




Audit failures

Compliance reports inflated with hundreds of false findings, obscuring the issues that actually need remediation



CVSS: Red Hat's adjusted scores



NVD provides a vendor-neutral CVSS assessment

Red Hat adjusts based on actual product exposure

Example:

NVD rates CVE as 9.8 Critical

Red Hat rates same CVE as 5.3 Medium

Reason:

Affected component not reachable in default config

Check access.redhat.com/security for RH-specific scores



Results comparison

Cutting through the noise



Raw Scanner output

847 total vulnerabilities
23 Critical, 156 High
No fix status context
Team overwhelmed by volume



After VEX (rh-cve-checker)

23 actionable vulnerabilities
54 with available fixes
770 confirmed false positives
Clear remediation priorities



Impact

91% noise reduction
Triage time cut from days to hours
Confidence in deployment decisions
Audit-ready reports




AI amplifies the scale problem




Increased Threat

Every AI agent, MCP server, and skill adds packages, dependencies, and attack surface




Exponential CVE counts

AI adoption adds a zero to your CVE count – 348 findings becomes 3,480



Scaling Limits

No human team can manually assess at this velocity



End-to-end Security

Automated end-to-end security vetting is no longer optional – it's survival



Key takeaways




Scanners are a start

Raw scanner output contains too much noise to be actionable. Validation against authoritative data is essential.



VEX/CSAF source of truth

Red Hat publishes machine-readable security data with per-product fix status, adjusted CVSS, and exploitability context.




Validation, not scanning

Tools like rh-cve-checker close the loop between scanning and action. Build validation into your CI/CD pipeline.



From CVE detection to remediation

Finding the real vulnerabilities is only "half the battle"



One agentic workflow – with human approval before production changes



Agentic Skill Pack for Site Reliability Engineers

- ▶ Tools designed for SRE tasks that works integrated via MCP to several tools: Lightspeed and AAP
- ▶ Discover your fleet, check what CVEs are impacting your infrastructure, generate playbooks and execute them, lastly validate your results

13 Specialized skills

- Fleet inventory
- Remediation
- CVE Impact & Validation
- Playbook generator
- Playbook executor
- Execution summary
- [...]

3 MCP Integrations

- Red Hat Lightspeed
- AAP Inventory Management
- AAP Job Management

AI Optimized documentation

- Semantic indexing
- Runs on Claude Code and Cursor



Navigating the AI-Driven Threat Landscape

Securing systems against zero-day exploits and rapid vulnerability discovery

01 / PERSISTENCE

Long-Term Support & Maintenance

With AI tools like Mythos discovering thousands of vulnerabilities in weeks, systems remain exposed longer. Long-term infrastructure stability acts as your baseline defense:

- Provides predictable maintenance and critical security patches over decades.
- Ensures change-averse deployments in telco, healthcare, and aerospace remain inherently secure.

02 / VISIBILITY

Rapid Threat Identification

When vulnerabilities are exploited before public reporting, immediate visibility is paramount to neutralizing risk:

- Instantly scan enterprise environments to pinpoint precisely which systems and versions are affected.
- Eliminate manual tracking bottlenecks through automated inventory analysis and immediate threat mapping.

03 / ACTION

Agile & Effective Remediation

Identifying threats is only half the battle. Swift, reliable patching execution minimizes the window of vulnerability:

- Deploy verified patches rapidly across distributed environments without operational disruption.
- Establish automated workflows to scale remediation from single systems to global infrastructure.



The baseline



SEE IT

Lightspeed

identifies vulnerabilities in your actual environment, contextualized to your workloads.

Available hosted online or on-premise with Satellite.



STAGE IT

Satellite

stages the patch in content views and promotes it through dev-test-prod lifecycle gates.



SHIP IT

Ansible Automation Platform

applies remediation fleet-wide with a full, policy-driven audit trail.

Available hosted online or on-premise with Satellite



Model Context Protocol (MCP) Ecosystem & Goose

RHEL MCP

Developer Preview

- Bridges LLM and RHEL infrastructure for real-time access
- Read-only tools cover CPU, memory, logs, and disk metrics
- Supports guarded execution with sandboxing and human-in-the-loop approvals

Satellite MCP

Technology Preview

- Integrates directly with Satellite API for large-scale environment visibility
- Simplifies inventory queries, patch reporting, and subscription compliance
- Inherits native Satellite permissions and roles

Lightspeed MCP

Developer Preview

- Connects AI clients securely to deep Red Hat analytics
- Provides centralized visibility across multiple Satellite instances
- Empowers use cases for Image Builder, Advisor, and Vulnerabilities

Goose CLI

Advanced Agent

- Optional, open-source AI agent available via RHEL Extensions
- Consolidates multiple MCP servers into a single interface
- Delivers faster streaming responses built for power users

Connecting AI capabilities securely with enterprise-grade Linux data.




What is MCP?

Model Context Protocol

- ▶ Model Context Protocol is an **open source protocol**, which enables 2-way-communication between AI applications and external services. It provides a **simplified and reliable way** for AI systems to access different data sources and tools.
- ▶ MCP **standardizes communication** on top of existing APIs. This makes it much easier to build complex AI workflows with large language models (LLMs) and connect those models to **real-world data**.
- ▶ MCP is a bridge between LLMs and your infrastructure. Enables LLM to have real time access to your systems actual state (logs, CPU usage, disk space, etc.).



Using a MCP Server



Which MCP components do we need?



MCP client

The AI application or system that requests access to external data or resources.



MCP host

The infrastructure that manages communication between the client and server.



MCP server

The component that provides specific tools, resources, and capabilities to the client.



Introducing Goose CLI

An advanced, open-source AI agent designed to orchestrate reasoning, consolidate multiple MCP servers, and deliver fast streaming responses directly in your terminal.

Available through the RHEL extensions repository.

INTEGRATION & CONTROL

Multi-MCP Orchestrator

Consolidates multiple Model Context Protocol (MCP) servers into a single command-line interface. Connects your local environment, Red Hat Satellite, and Red Hat Lightspeed simultaneously.

PERFORMANCE

Streaming Responses

Built specifically for power users, delivering ultra-fast streaming AI insights directly in your terminal. Accelerates command execution, diagnostic workflows, and real-time system analysis.

REASONING LOOP


Autonomous Action

Independently manages multi-step workflows. Safely evaluates intents, selects the appropriate RHEL tools, executes them securely, and monitors the outcomes under human-in-the-loop guardrails.



Goose CLI: Multi-MCP Architecture

Example deployment orchestrating RHEL, Red Hat Satellite, and Red Hat Lightspeed MCP servers simultaneously from a single local terminal environment.



Security Advantages of MCP & Goose CLI

How Agentic AI with Model Context Protocol transforms CVE discovery and mitigation

Proactive Visibility

Continuous CVE Discovery

Unlike traditional periodic scans, agentic AI leverages the Model Context Protocol to analyze live systems in real time:

- Identifies zero-day threats and newly reported vulnerabilities instantly.
- Correlates active system states against global CVE databases without human lag.

Guarded Execution

Sandboxed Security & MCP

Integrating LLMs with sensitive RHEL infrastructure requires rigid protection policies built natively into MCP:

- Enforces strict read-only execution for system state discovery and metric gathering.
- Leverages sandboxed environments to safely probe and analyze suspicious activities.

Controlled Action


Human-in-the-Loop & Goose

Goose CLI acts as an orchestrator, bridging AI reasoning with safe, enterprise-grade remediation workflows:

- Ensures no action is performed without explicit human-in-the-loop validation.
- Consolidates secure, multi-server operations into a single audit-tracked terminal.

Providing rapid threat containment while preserving strict access control and absolute operator visibility.





Os esperamos en el booth si quereis ver demos o hacer un lab:

- MCP Server for RHEL
- MCP Server for Lightspeed
- Step-by-step Security Updates
- Enhanced Malware Detection Coverage with CrowdStrike
- Examining System Performance
- Troubleshooting with AI
- Updating Software with Web Console
- ...



Tech Journey demos: Red Hat Security



Webinar

[Desplegando una fortaleza digital: no confíes, siempre verifica y encripta tus datos](#)

23 de octubre de 2026, 10:00 a.m.



Webinar

[Automatización y seguridad: estrategia de remediación integral con Red Hat](#)

13 de noviembre de 2026, 10:00 a.m.





Hands-on Workshop Week • November 30, to December 4, 2026

Red Hat Tech Journey workshop week

One week. Four areas shaping enterprise IT. Hands-on experience.



Discover how to apply open source technologies to build secure platforms, automate operations, and run modern applications across hybrid environments.

Resources

- ▶ [Model Context Protocol](#)
- ▶ [MCP Server for RHEL \(upstream\)](#)
- ▶ [Agentic skill pack for Red Hat customers](#)
- ▶ [Automate at scale with AI skills, agents, and MCP servers provided by Red Hat](#)
- ▶ [Red Hat Interactive Demos](#)
- ▶ [RHEL Interactive labs](#)
- ▶ [RH-CVE-CHECKER](#)





Connect

Thank you



linkedin.com/company/red-hat



facebook.com/redhatinc



youtube.com/user/RedHatVideos



x.com/RedHat

