Jürgen Buddy Hoffmann
Ansible Platform Leader, Red Hat

Red Hat Summit Connect

Before the Exploit Strikes

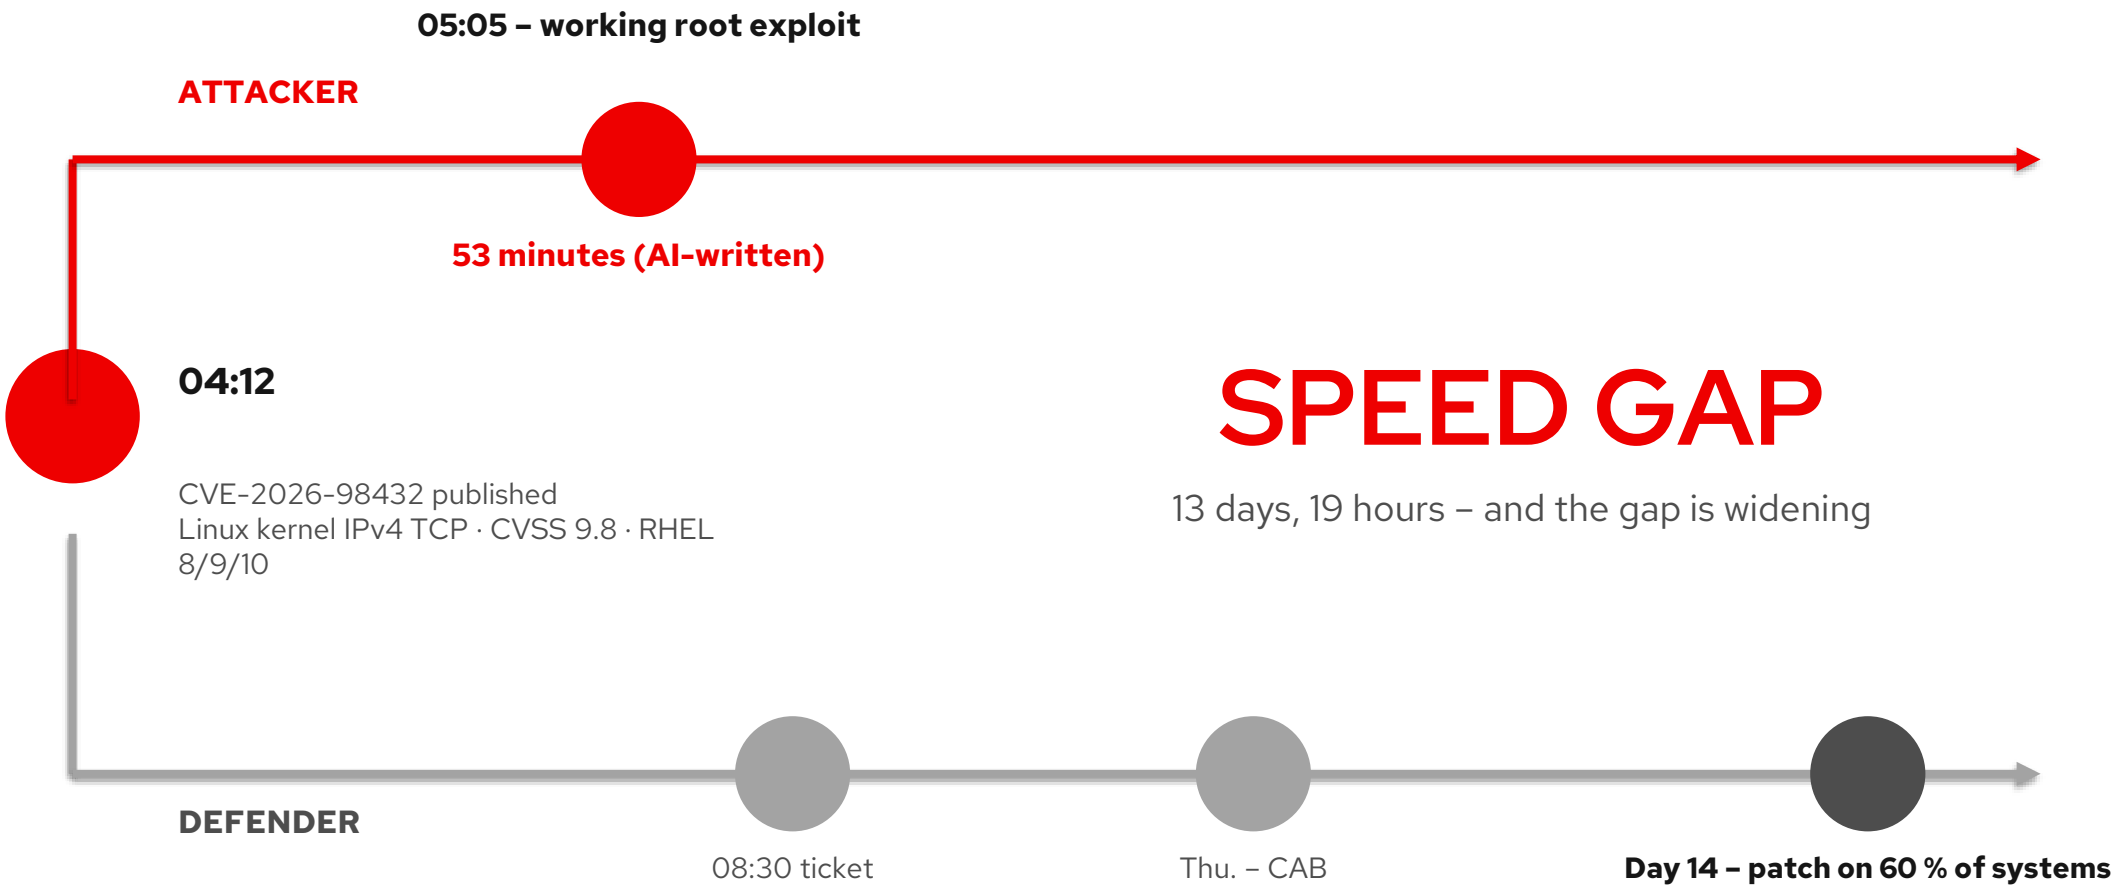
Building a Trusted Execution Layer with Ansible Automation Platform

04:12

Remember this time.

Red Hat

04:12. CVE-2026-98432 is public. What now?



Source: Red Hat RHSB-2026-016 (CVE-2026-98432); exploit time per Anthropic, via The Hacker News, 3 Nov 2026; 04:12 timing is a replay; median patch time: CISO conversations, 2026

The speed gap in numbers: AI finds faster than we fix

93,909

CVEs published in 2026 by 9 November – about 300 per day, almost double all of 2025 (48,151)

6,157

vulnerabilities disclosed to maintainers of 591 open-source projects by one AI programme (Anthropic Glasswing)

8 %

of them patched by 2 October: 516 of 6,157 – the fix side cannot keep up

$1,280

per Linux kernel root vulnerability – one OpenAI pipeline run: 948 found, 1,169 working exploits, $1.21 m

Glasswing, as of 2 October 2026



"Trivially exploitable." – Matt Hicks, CEO Red Hat

Source: cvetodo.com / NVD, 9 Nov 2026; Anthropic disclosure dashboard, 2 Oct 2026 (via SIG, CSA); OpenAI pipeline results via Cloud Security Alliance, 5 Nov 2026; CRN, 27 Aug 2026

October 2026: the speed gap stopped being a scenario

17 – 26 OCTOBER

PHP2Shell

CVE-2026-109618 / -109619 · CVSS 9.2
found with AI assistance

–4 days

first exploit attempt before the patch existed (AWS honeypot, 17 Oct)

- 21 Oct, 02:00 – PoCs circulate, 12 h before release
- 21 Oct, 14:00 – PHP fix released
- 26 Oct – CISA: actively exploited
- 32.5 m PHP services still on an affected version (Censys)

30 OCTOBER – 1 NOVEMBER

Linux kernel, IPv4 TCP

CVE-2026-98432 · CVSS 9.8
remote, unauthenticated

53 min

for an AI-written remote root exploit: 48 min to leak, +5 to root shell (per Anthropic)

- Affects RHEL 8, 9, 10 and OpenShift
- 30 Oct – Red Hat bulletin RHSB-2026-016
- 1 Nov – fixed kernels for RHEL 8, 9, 10
- No attacks in the wild reported (3 Nov)

10 NOVEMBER

Patch Tuesday

Microsoft
one single monthly release

1,093

CVEs fixed in one release – a record

- 63 rated critical
- 1 zero-day exploited in the wild
- ~300 new CVEs per day industry-wide in 2026

Exploit before patch is no longer the exception. The window you control: from patch to rollout.

Source: AWS Security Blog, 21 Oct 2026; Recorded Future; The Hacker News, 27 Oct & 3 Nov 2026 (Censys; Anthropic); Red Hat RHSB-2026-016; Tenable, 10 Nov 2026

How the attacker works in 2026: a pipeline, not a person

1. Observe

Agent watches advisories, commits, NVD feeds – around the clock

2. Understand

Multimodal model reads diffs, comments, architecture diagrams

3. Build

Exploit code is generated and corrected iteratively against stack traces

4. Test

Sandbox farm validates in parallel against real distributions and kernels

5. Scale

Shodan search for vulnerable systems, automated attack

04:12 → **05:05**

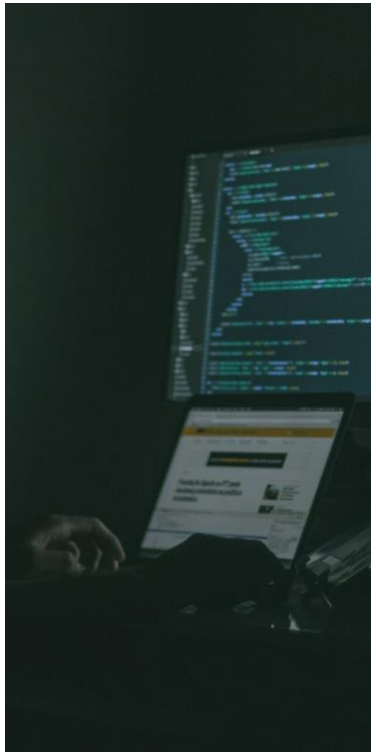


Photo: Jefferson Santos / Unsplash

No human in the loop. Cost per kernel root exploit: ~$1,000 in API credits. Only a better pipeline beats a pipeline.

Source: OpenAI pipeline results via Cloud Security Alliance, 5 Nov 2026; InfoQ on Kimi K2.5 agent swarm; CRN interview with Matt Hicks, 27 Aug 2026

Revolut 2026: no exploit – a trusted sender was enough

~680

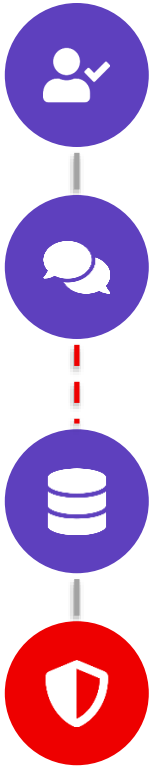
customers affected – targeted high-net-worth users; ID copies, selfies, IBANs, transaction histories

$3 m

ransom demanded by the attackers via Telegram (FT report, not confirmed by Revolut)

from $40

price of a hijacked government mailbox in the underground – emergency requests as a service from ~$100



Genuine government mailbox

Sender and domain authentic – reportedly an Italian agency

“Emergency data request”

disguised as law enforcement, repeated over months

MISSING VERIFICATION GATE

Data is released

systems and customer funds untouched – no code was broken

Callback to the agency

only after the fraud had already been noticed

A request is not an authorisation. Between request and execution belongs a gate – in the architecture, not in the gut feeling of the night shift.

Source: TechCrunch 12 Sep 2026; Financial Times via CybelAngel 16 Sep 2026; Rankiteo 15 Sep 2026 – ransom figure not confirmed by Revolut

Benelux, 2 November: the warning came three weeks earlier

11 Oct

France

Refrigeration at hundreds of supermarkets set to 50 °C – Danfoss equipment, largely outdated

13 Oct

The warning

Danfoss advises customers: disconnect from the internet or use a secure connection

1–2 Nov

Benelux & beyond

NL: at least 17 businesses – butchers, grocers, cold storage. BE: Carrefour, Delhaize, Spar. Also LU and SE

Cause

Per Danfoss

Systems not on a secure network – still on the default password

17+ Dutch businesses with fridges and freezers tampered with overnight

20 days between the supplier's warning and the attack

0 exploits needed

Knowing is not acting. A warning without an executed change is just a document.

Source: NL Times and DutchNews, 2 Nov 2026 (citing Danfoss via NOS, VRT, RTL Info); cause as stated by the supplier Danfoss

“This breadth of, ‘I use anything I grab from the internet, and I trust it’ – I do think those days are over.”

Matt Hicks

President and CEO, Red Hat · CRN, 27 August 2026

The pivot: five billion dollars as the only logical answer

$5 bn

Project Lightwell – IBM and Red Hat investment in AI-assisted vulnerability remediation

20,000+

engineers delivering validated, signed backport patches for all open-source software

400+

previously unknown Java library vulnerabilities found, fixed and backported (Oct 2026)

Old process

- Describe the problem
- Discuss upstream
- Wait
- Patch eventually
- Exploit? “Too hard.”

New process

- Build the exploit yourself – “the uncomfortable work”
- Patch first
- Trusted circle
- Then upstream and disclosure
- For the versions you actually run

Source: IBM Newsroom, 28 May and 6 Oct 2026; IBM Think, 6 Nov 2026; CRN, 27 Aug 2026

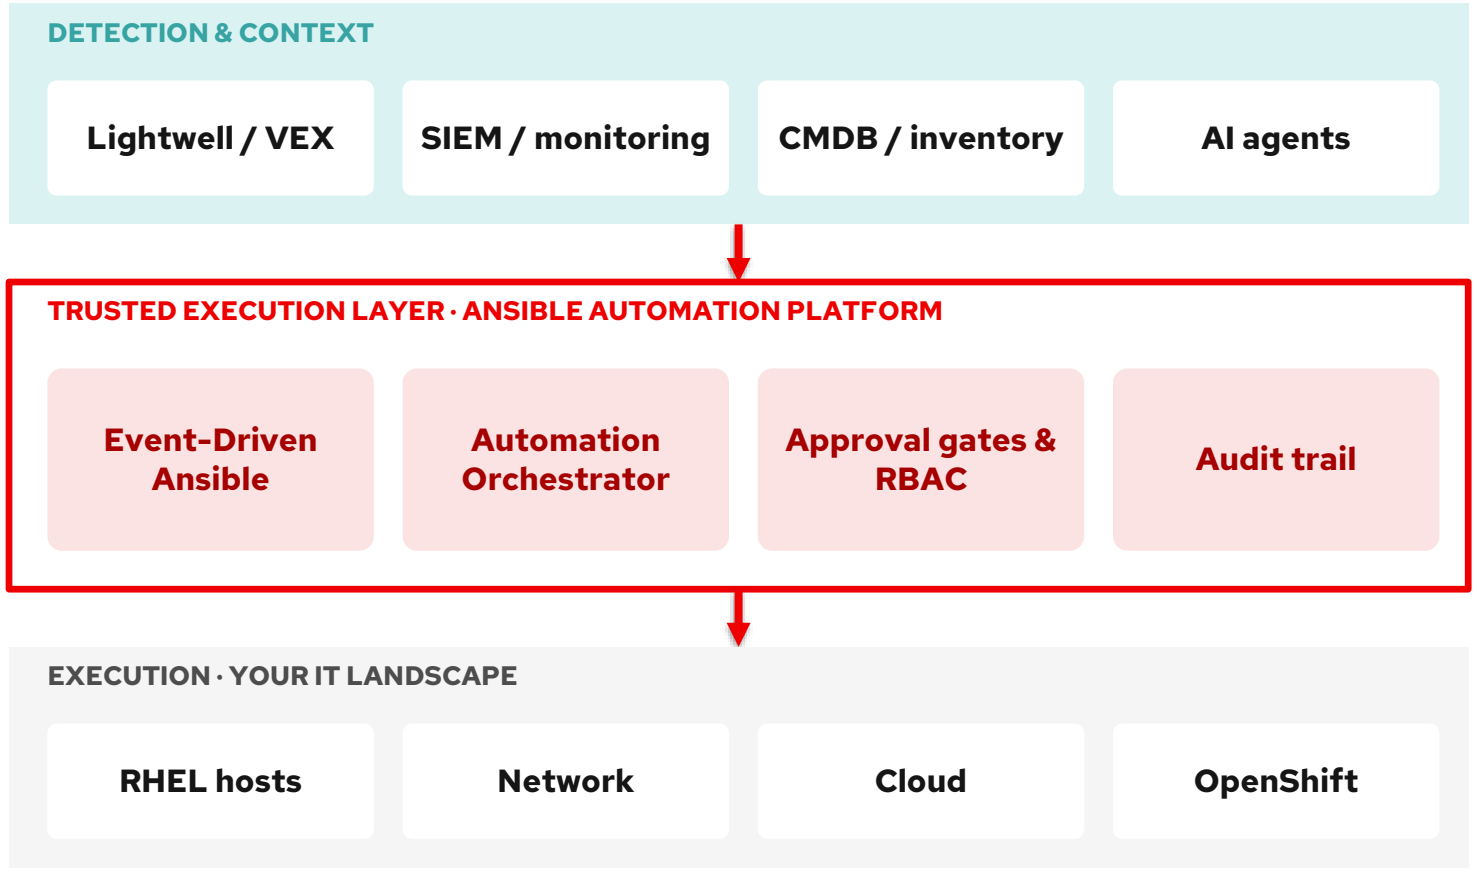
Speed is not a question of the team. It is a question of architecture.

The business risk is evident: reputation · revenue · press

No one wants to be the team that didn’t take action.



Ansible Automation Platform: the Trusted Execution Layer



One security boundary for everything that gets executed

- AAP is the Trusted Execution Layer – robust, resilient, responsive
- Inside it, the Automation Orchestrator composes workflows from your playbooks and rulebooks
- Three kinds: event-driven, AI-driven, task-driven
- RBAC, approval gates, complete audit trail
- No agent executes directly. Every agent requests an execution.

Source: redhat.com – Ansible Automation Platform, Automation Orchestrator, Event-Driven Ansible

MCP and A2A: the language of agents

MCP – Model Context Protocol

How does an AI model talk to tools and data?

- Ansible Automation Platform, CMDB, ticketing – each as an MCP server
- The model sees capabilities, not credentials
- Originally Anthropic, now an open standard

Agent → MCP server → Tool

A2A – Agent-to-Agent Protocol

How does one agent delegate tasks to another?

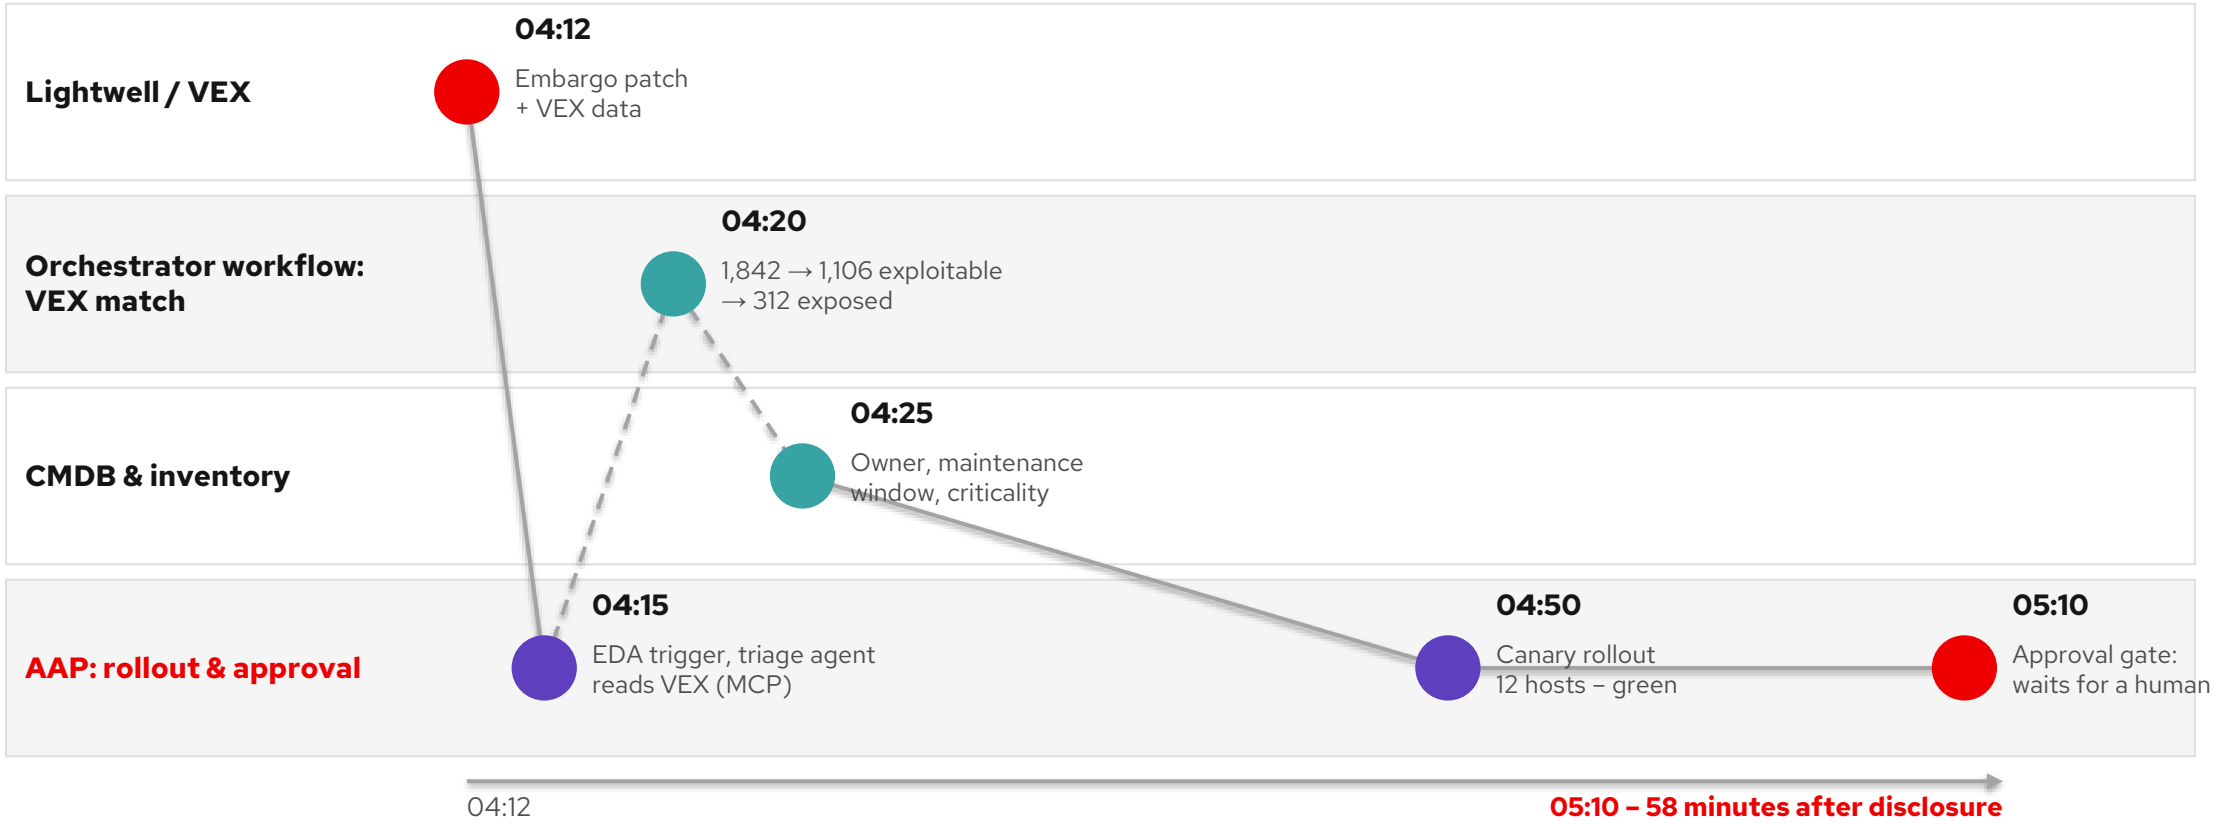
- Triage agent → patch agent → compliance agent
- Structured hand-off with identity and traceability
- Originally Google, now an open standard

Agent ↔ Agent ↔ Agent

AAP in the middle: MCP server for models, policy authority for every agent hand-off.

Source: modelcontextprotocol.io; a2a-protocol.org; both standards under the Linux Foundation

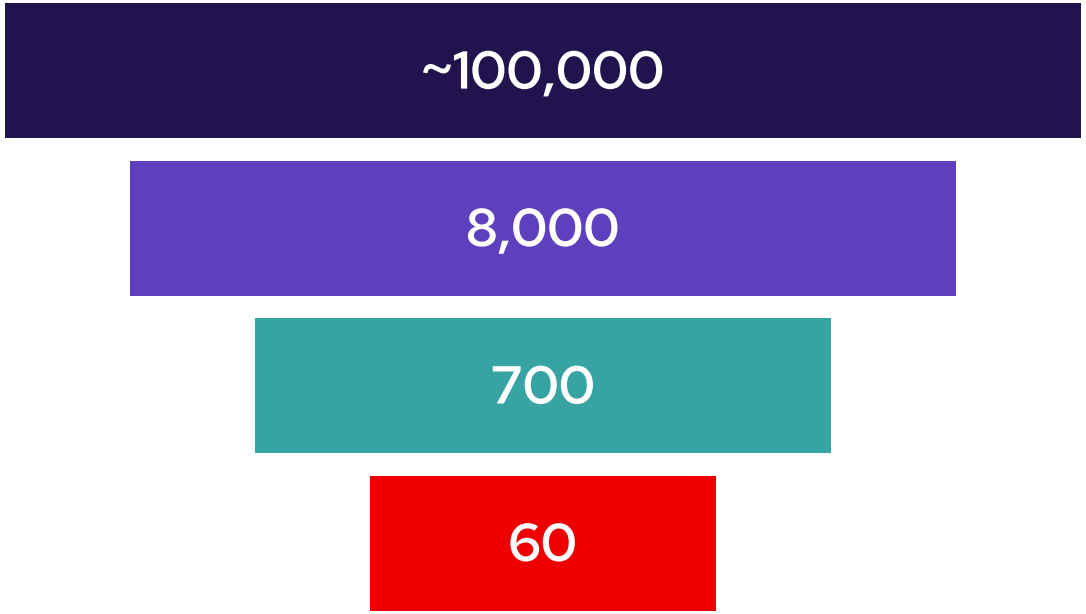
04:15: the defender pipeline runs – automatically



dashed = MCP call · solid = A2A hand-off / workflow

Source: Replay of CVE-2026-98432 – timings and host numbers illustrative; flow with Event-Driven Ansible, Automation Orchestrator and Lightwell

Reality check: not every 9.8 is your 9.8



CVEs per year (2026 pace)

affect software in your inventory

exploitable in your configuration per VEX

exposed and business-critical – decisions

What is VEX?

Vulnerability Exploitability eXchange – the vendor's machine-readable statement on whether a CVE is exploitable in a specific product: affected · not affected · fixed · under investigation. Red Hat ships VEX with every advisory.

VEX – exploitability

Is the vulnerability exploitable in my configuration? The vendor statement instead of the CVSS score.

Orchestrator – context

A workflow reads VEX, correlates it with inventory and exposure – and prioritises before a human sees anything.

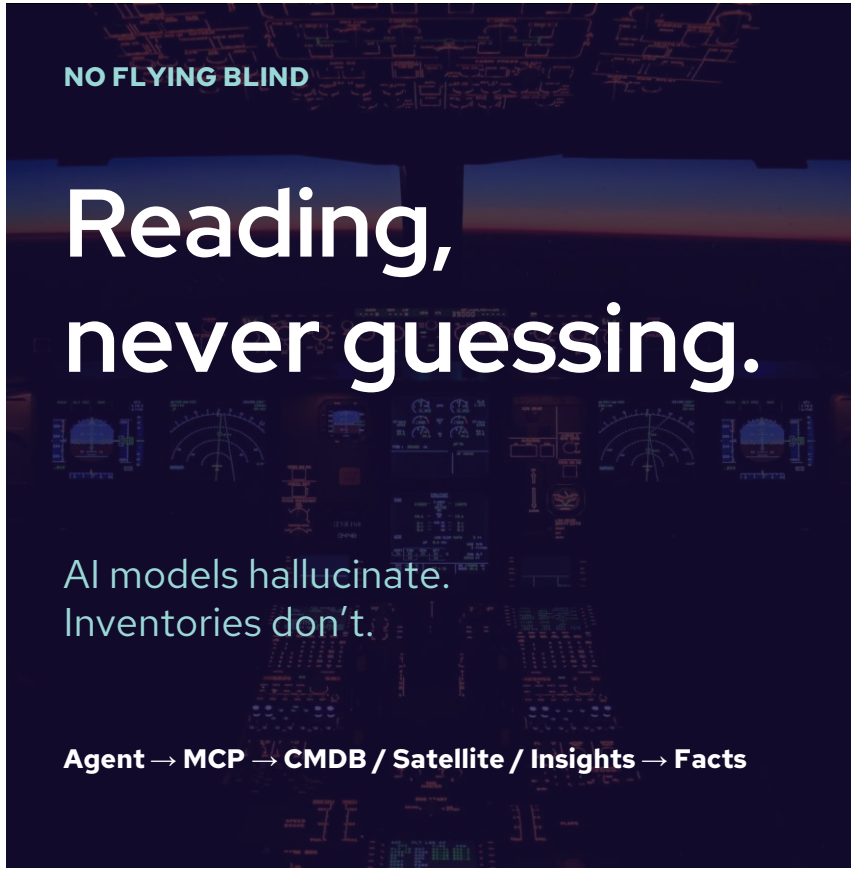
Priority – business impact

Thousands of notifications become dozens of decisions a human can take responsibility for.

"Most customers cannot read VEX." – Matt Hicks. Nobody has to: the automation can.

Source: cvetodo.com / NVD, 9 Nov 2026; CISA KEV via SOCRadar, Nov 2026; CRN, 27 Aug 2026; funnel figures illustrative

The CMDB remains the single source of truth

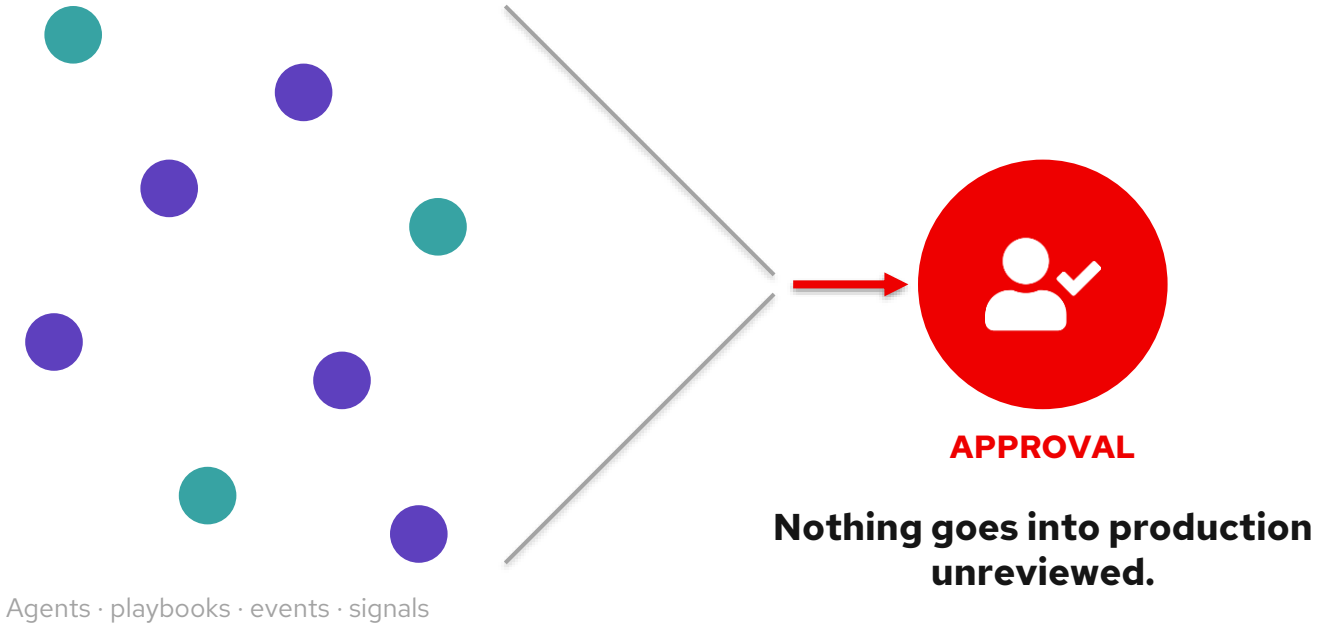


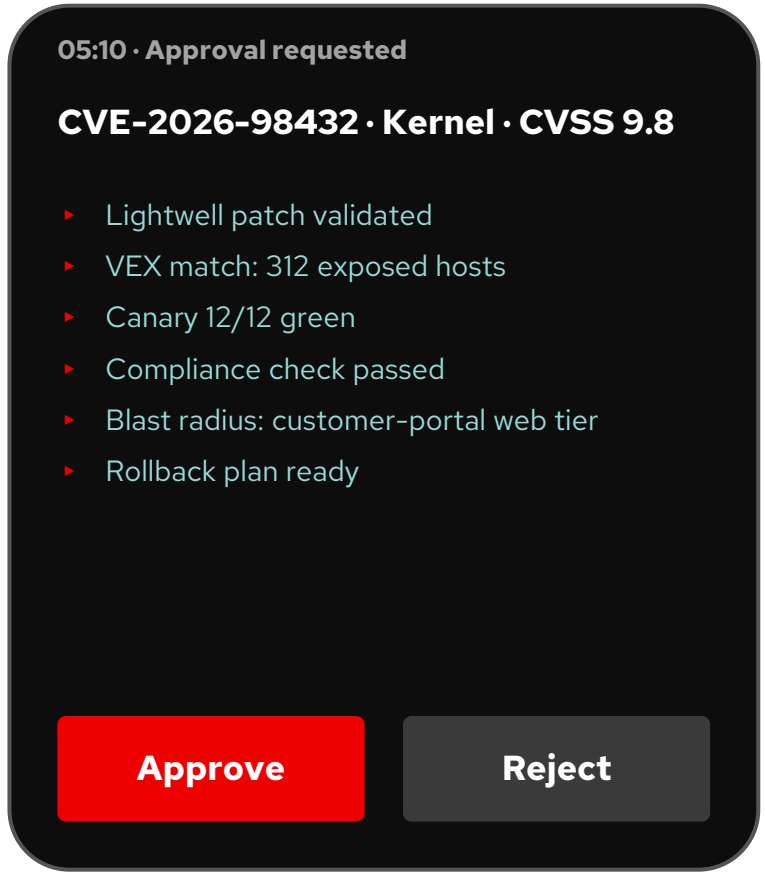
What the AI cannot know, it has to look up

- Agents read from the CMDB – they do not invent inventory
- CMDB, Satellite and Insights supply owner, criticality, maintenance window, dependencies – available via MCP
- Every automation writes back: patch level, reboot time, result – the CMDB stays current because automation maintains it
- A gap between inventory and reality is itself a security finding
- No agent may touch a system that is not in the inventory

Source: Architecture principle; inventory data from CMDB, Red Hat Satellite and Red Hat Insights · Photo: Andrés Dallimonti / Unsplash

The control paradox: automate more, and the human matters more

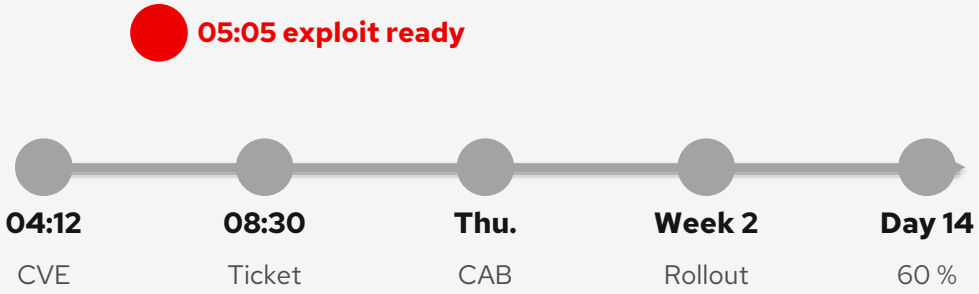




Source: Architecture principle: approval gates in Ansible Automation Platform

05:12: approval granted. 06:30: 312 hosts secured.

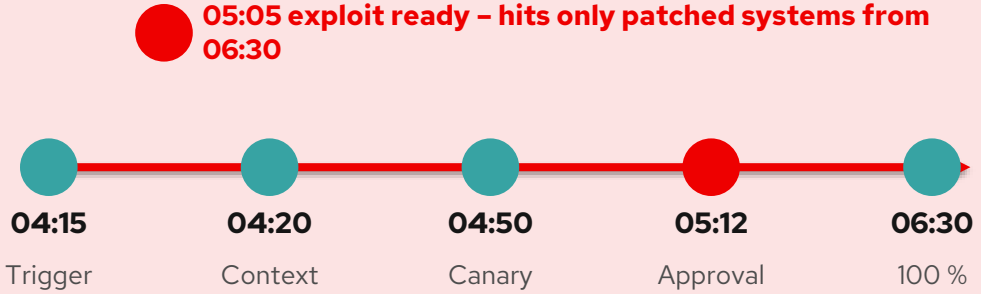
Without the architecture



13 days, 19 h

Attack window for 312 exposed hosts

With a Trusted Execution Layer



1 h 25 min

Attack window for 312 exposed hosts · 794 more in the maintenance window

No hero. No night-shift marathon. One approval.

Source: Replay of CVE-2026-98432 (exploit time per Anthropic) – your numbers will differ; the order of magnitude will not

For your business: measurably faster, measurably safer

Risk

- Mean time to remediate for critical CVEs: from weeks to hours
- Attack window for exposed systems measurable in hours
- Number of unpatched, exploitable systems as a daily KPI

Cost

- Manual patch hours per critical CVE towards zero
- No night shifts, no emergency CABs
- One workflow pattern for kernel, middleware, network, cloud

Trust

- Every change with evidence and approval in the audit trail
- DORA and NIS2 evidence from the system instead of from Excel
- A CMDB that is right

Source: Red Hat Business Value Practice. Results vary by environment; no contractual commitment.

Three steps to start tomorrow

1

WEEK 1

Know what affects you

Match advisories and VEX data against your inventory with a workflow. Goal: one number – how many exploitable, exposed systems today?

2

MONTH 1

Build one workflow

One CVE type, one canary ring, one approval gate in Ansible Automation Platform. Not perfect – but measured: time from advisory to approval request.

3

QUARTER 1

Let the agents in

Triage agent via MCP to AAP and the CMDB, A2A for the hand-off to the patch agent. The human stays at the gate – that is the target picture.

Go deeper today: 14:00 Lightwell Network (Room 2) · 15:00 demo “From CVE Alert to Patched Fleet” (Silent Theater 2) · 15:15 Ansible Compliance (Room 3)

Source: Red Hat recommendation; VEX data with every Red Hat Security Advisory (CSAF)

The exploit takes 53 minutes. Your approval takes two. Everything in between is automation.

06:30



Red Hat Summit

Connect

Thank you

Questions? Jürgen Buddy Hoffmann · juhoffma@redhat.com

linkedin.com/in/juergen-buddy-hoffmann

linkedin.com/company/red-hat

facebook.com/redhatinc

youtube.com/user/RedHatVideos

x.com/RedHat

Red Hat