Guardare

USE CASE · ENDPOINT · EMAIL · HUMAN RISK

Three unrelated alerts. One user. One decision.

How Guardare threaded an endpoint exposure, an email campaign and a human-risk signal into a single remediation task — in seconds.

Illustrative example drawn from a live customer environment. Identifying details have been anonymized.

Nothing in the stack was wrong

The endpoint tool was right.

The mail gateway was right.

The awareness platform was right.

Every one of them reported accurately, on time, into its own console. The gap wasn't detection — it was that nobody read them in the same sentence.

Three signals arrived separately, each scored as a routine medium, each owned by a different queue. Individually, they were three tickets someone might get to this week. Together, they described a specific person, on a specific machine, about to be handed a specific exploit.

SIGNAL 01 · ENDPOINT

A critical CVE exposure in Google Chrome

The endpoint was running a Chrome build below the vendor's fixed version, placing it inside two published NVD vulnerable ranges. The flaw is a use-after-free in PDFium: opening a malicious PDF is sufficient for remote code execution.

DETECTED BUILD	SAFE THRESHOLD	CVES
151.0.7900.0	≥ 151.0.7922.72	CVE-2026-17993 · CVE-2026-18012
SEVERITY	**EXPLOIT PATH**	**SCOPE**
CVSS 8.8 — high	Opening a PDF	1 endpoint · risk score 65

IN ONE LINE

The browser build is two CVEs behind the fix, and the exploit path is simply opening a PDF.

SIGNAL 02 · EMAIL

A spike of malicious PDFs aimed at that same domain

Read alongside mail provider and gateway telemetry, Guardare identified an active campaign targeting the domain: mailbox flooding to bury the inbox, followed by voice calls impersonating IT support to pressure the user into granting remote access or opening an attachment.

SOURCE	PATTERN
Mail provider + gateway defense	Email bombing → vishing as IT
PAYLOAD	**TARGET**
Malicious PDF attachments	The affected endpoint's domain

IN ONE LINE

The exact file type that exploits that browser was arriving in the mail flow, at volume.

SIGNAL 03 · PERSON

A privileged user on that browser who fails phishing simulations

The browser belonged to an account with administrative access. The awareness record showed a 60% failure rate in simulated phishing and an incomplete *Social Engineering Red Flags* module — the training that covers this precise attack pattern.

USER	PHISH-PRONE SCORE
Administrator account (anonymized)	60
TRAINING STATUS	**COMBINED RISK**
Incomplete	95 · MITRE ATT&CK T1566

IN ONE LINE

The person most likely to open that PDF is an admin, on the unpatched browser, who hasn't finished the training that covers it.

THE JOIN

Same browser. Same inbox. Same person.

Separately: three medium-severity tickets, in three consoles, with three different owners.

Together: one line of exposure — with a name, an owner and a due date on it.

Hours to thread. Seconds to find.

The manual path isn't technically difficult. It's scattered — and it only happens if someone thinks to go looking.

THE WORK	BY HAND, ACROSS CONSOLES	WITH GUARDARE
Match the detected build to NVD vulnerable ranges	Endpoint inventory export → NVD lookup · 25–40 min	AUTOMATIC
Confirm severity and the real exploit path	Read two CVE records, translate to plain risk · 20–30 min	AUTOMATIC
Check whether that payload is actually arriving	Mail gateway queries, campaign correlation · 30–45 min	AUTOMATIC
Tie the endpoint to a named user and their risk	Directory lookup → awareness platform reports · 20–30 min	AUTOMATIC
Decide the remediation, the owner and the deadline	Judgment call, usually a meeting · 30–60 min	AUTOMATIC
Total	**2–4 hours, three consoles**	**Seconds**

Timings are illustrative and will vary by environment and team structure. The more useful question is usually not how long this takes, but how often it happens at all.

WHAT YOU GET BACK

Not another alert. A recommendation that arrives as work.

01

One task, not three

Both CVEs handled as a single patch action, with the email campaign and the user's risk profile attached as context.

02

One owner and a due date

Assignment and deadline are part of the recommendation, not a follow-up meeting.

03

A sourced explanation

Every conclusion names the record it read — the detected build, the NVD range, the CVE IDs and CVSS score, the simulation result, the incomplete module. The judgment is ours; the underlying facts are yours, and you can click into any of them.

04

Current by construction

Builds are matched against published NVD ranges and vendor fixed versions, and each recommendation carries its own modified and due dates.

Where this pattern shows up

Guardare doesn't replace the tools in that stack. Each of them was right in this story. It reads them together — because the gap isn't detection, it's correlation.

Three or more security tools are in place, with no layer reading across them.

Security is a part-time responsibility for someone whose title says something else.

Administrators use the same devices, browsers and habits as everyone else.

Nothing in your stack was wrong. It just took three tools and most of an afternoon to find out you had a problem. We'd have told you in seconds.

Run this against your own environment — guardare.com/demo