Ship It, They Said

Security Research and Product Development at Scale

Whoami?

Marion Marschalek

Hack & Cheese Security Consulting

Future PhD Student at Dartmouth College

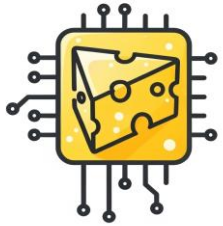

The Arc

From curiosity
to product deployment

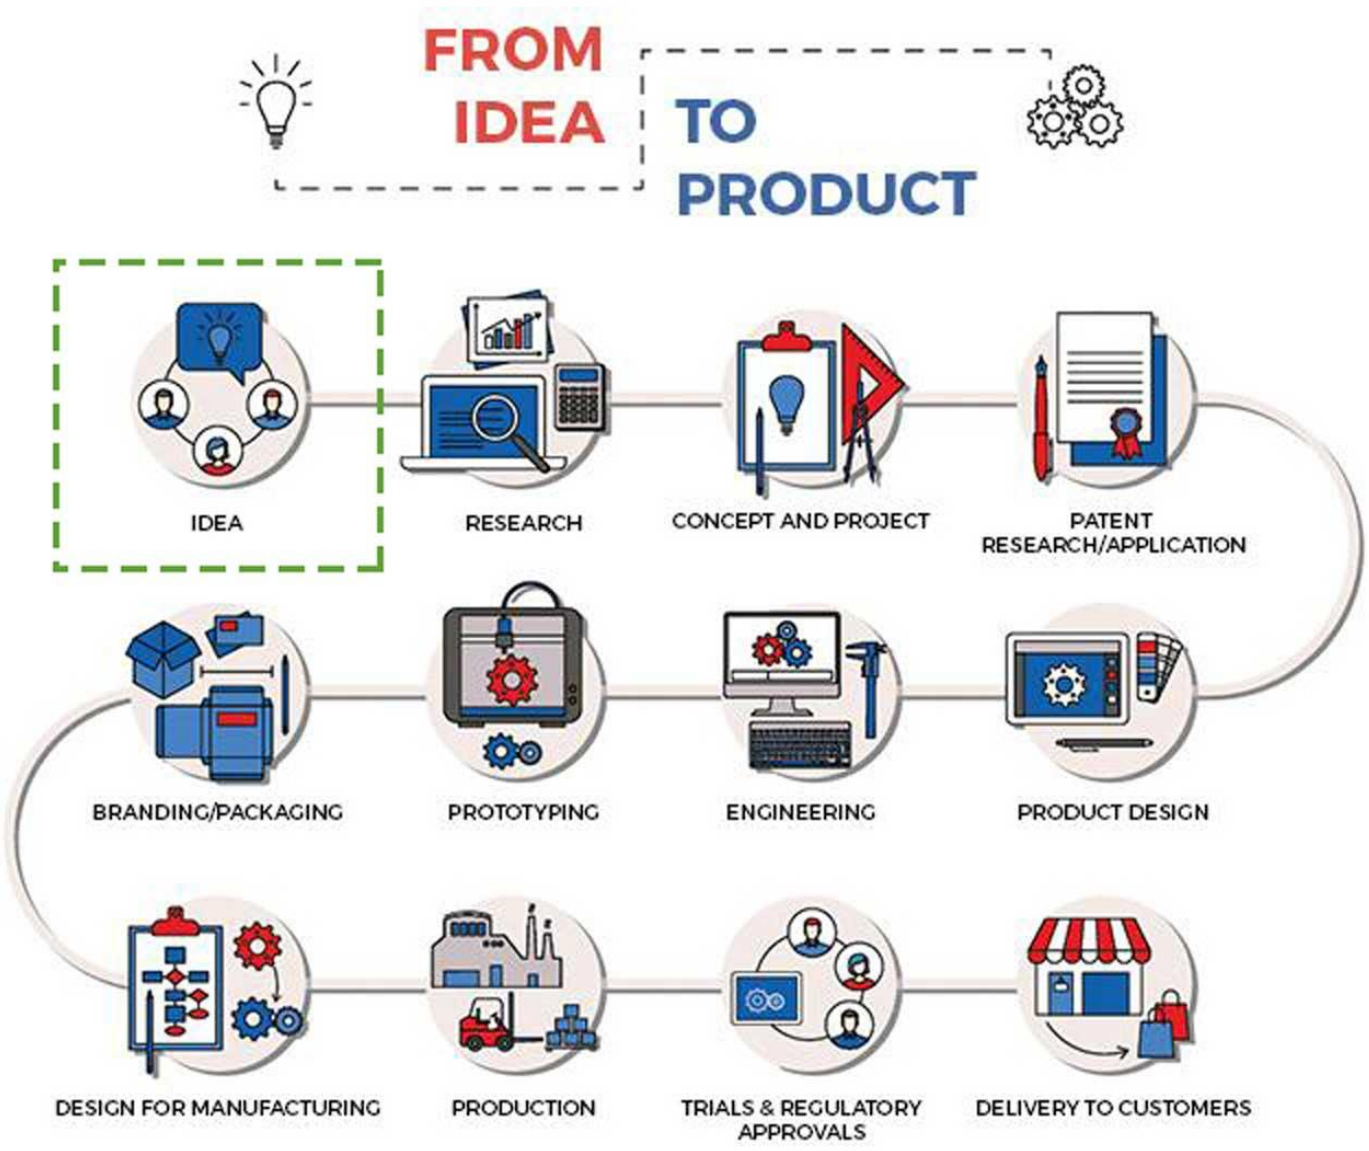

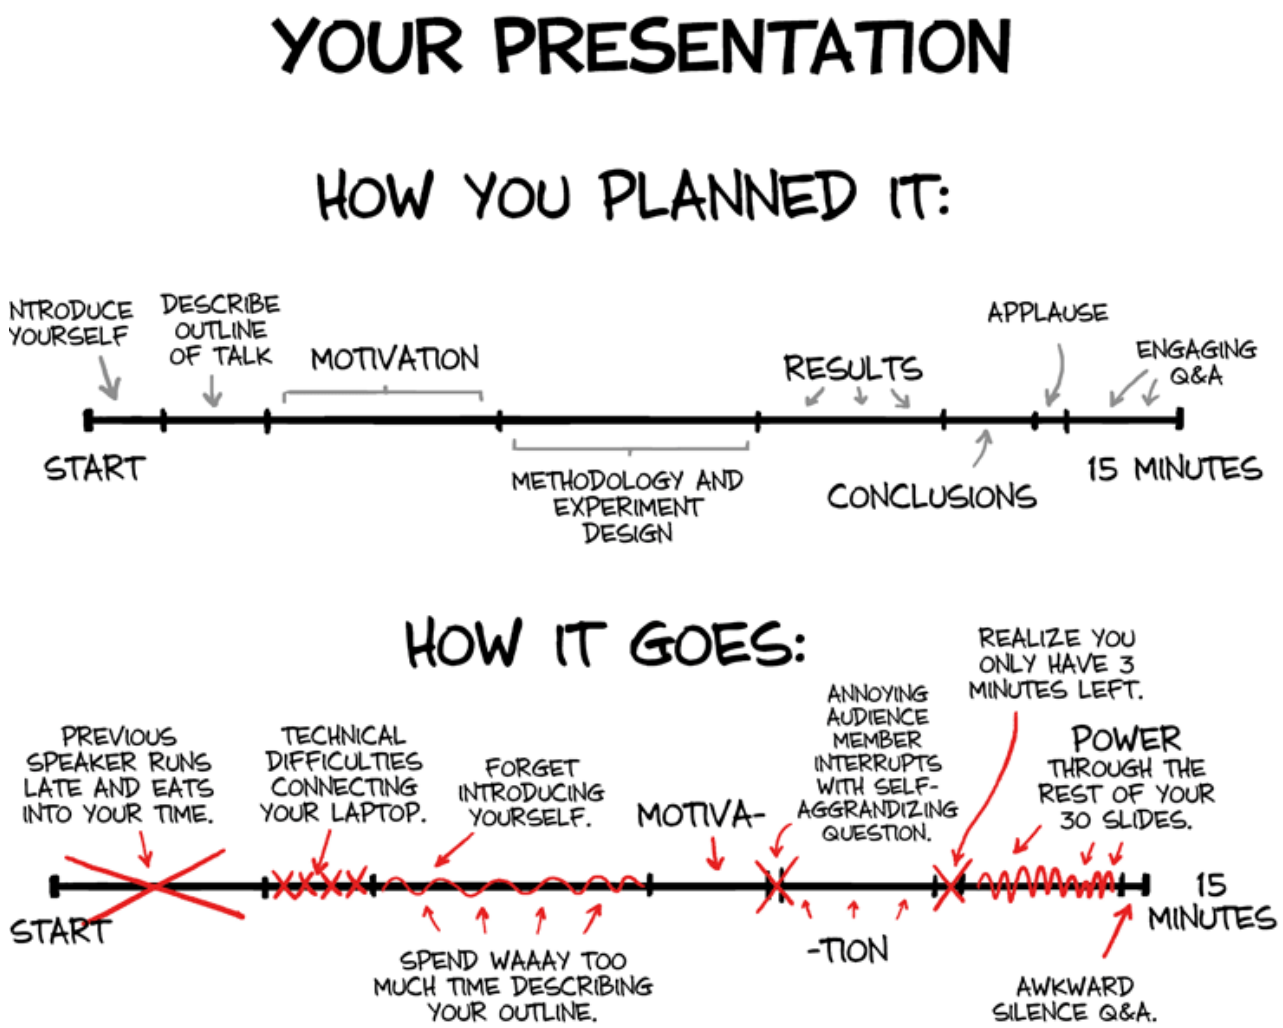

But hey, RESEARCH!

Research vs. Applied Security

The gap, and why there is a gap

Why do we do research?

Fun rabbit holes, and we'll go change the world

Oh and, dopamine and the learning curve

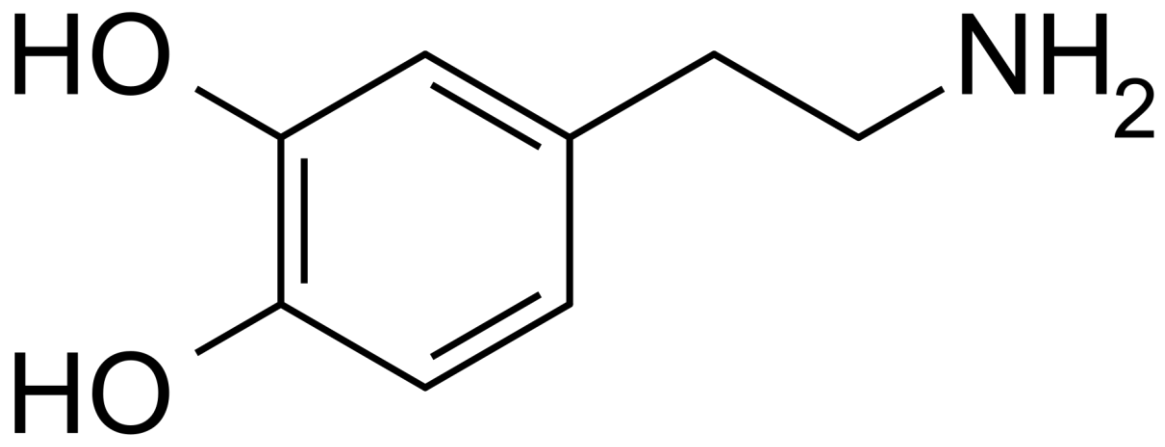

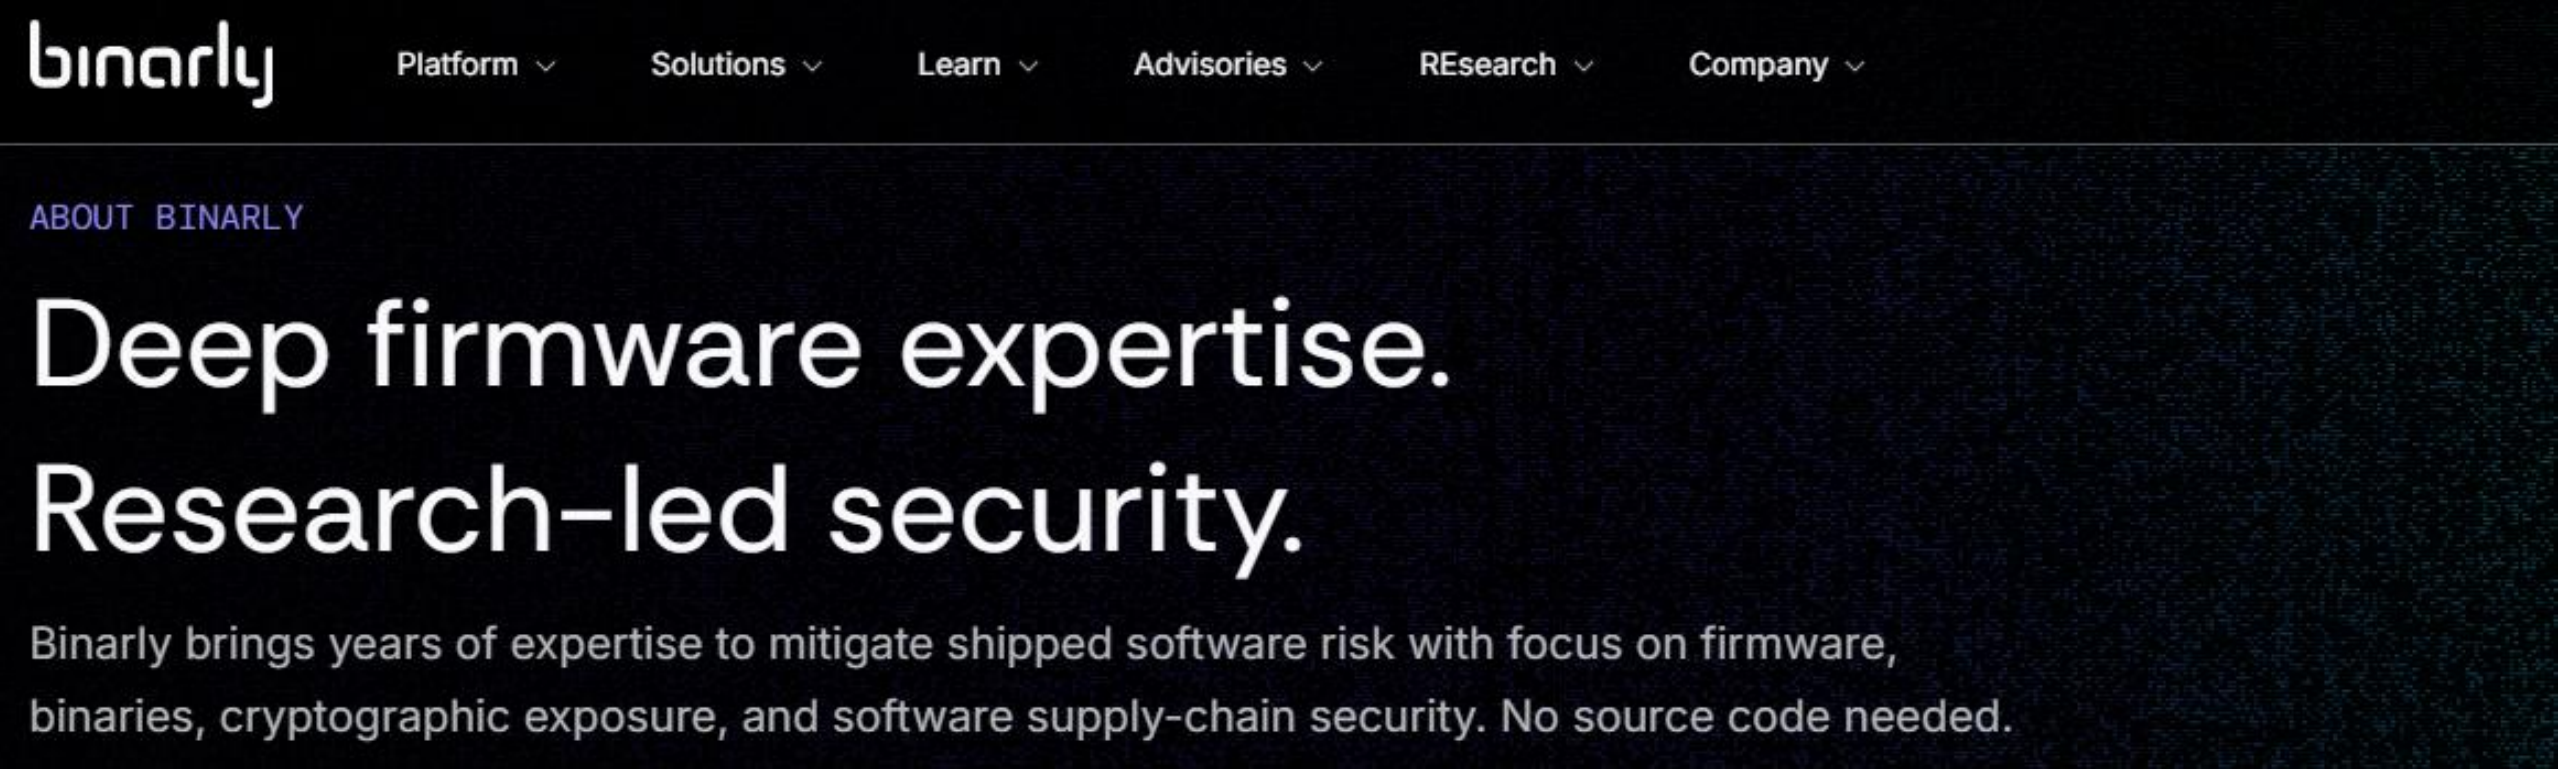

Alex Matrosov turned UEFI security upside down

Research led to the founding of Binarly

205 disclosures to 18 vendors impacting 511 products

Research & The Real World

Reverse Engineering a Ledger Nano X Hardware Implant

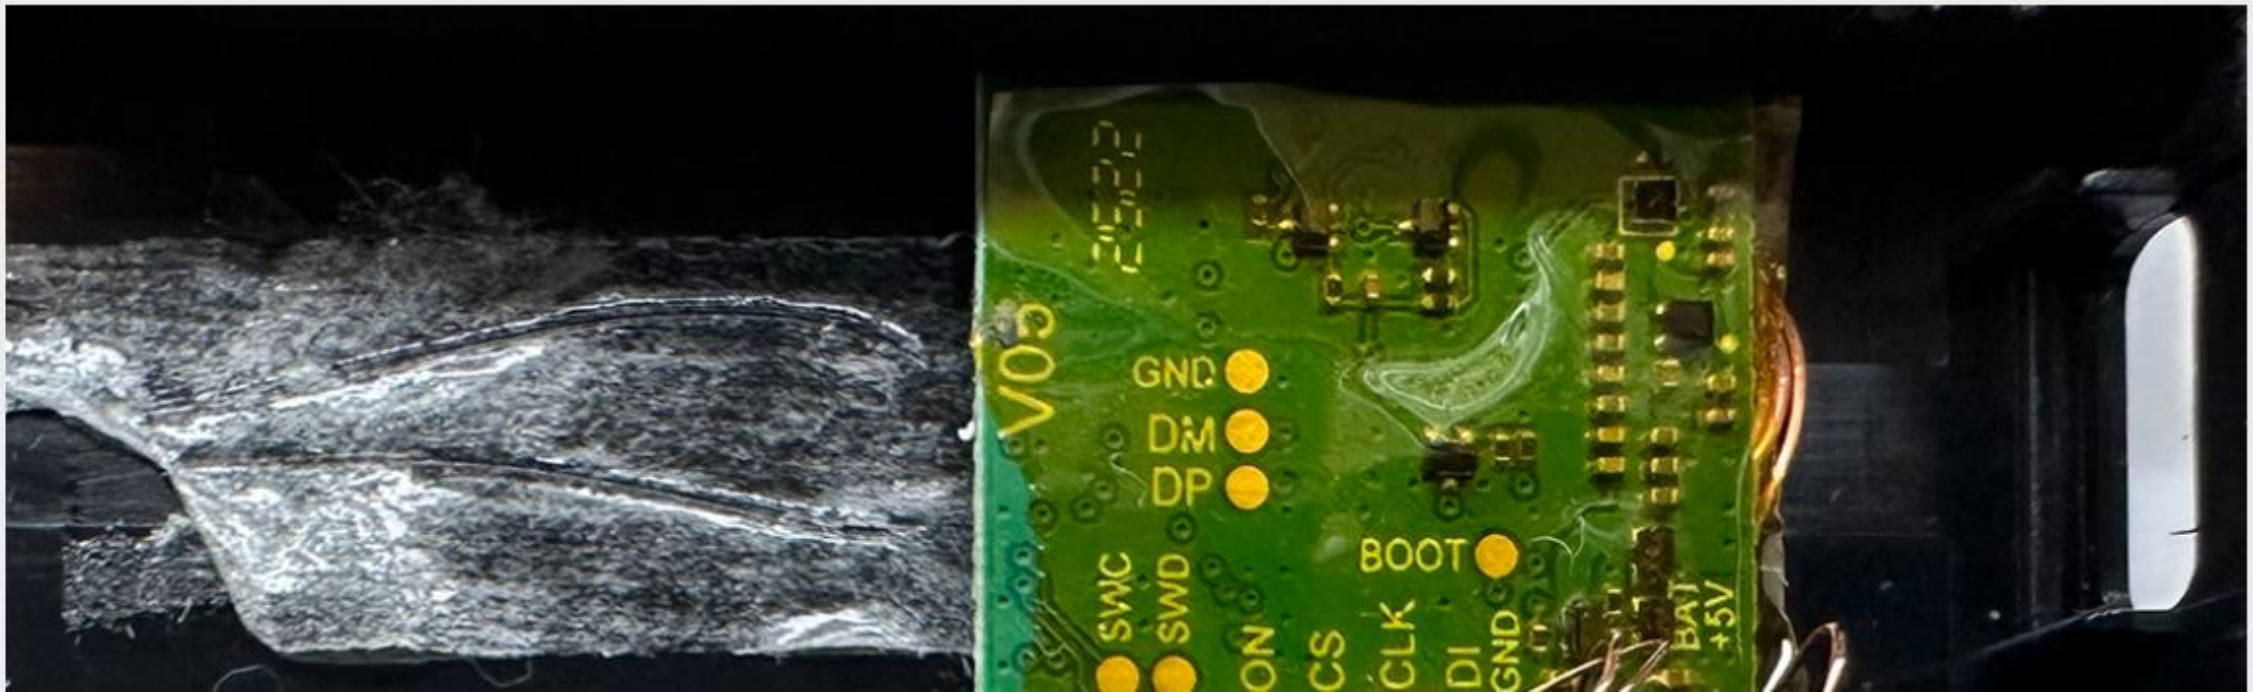

GreyNoise
State of the Edge Report 2026

Where Your Edge Defenses Fall Short

GreyNoise measured 212 exploitation attempts per second across H2 2025 — and the patterns inside that volume expose specific, measurable gaps in common edge defense strategies.

The *2026 GreyNoise State of the Edge Report* analyzes 2.97 billion malicious sessions from 3.8 million unique source IPs. It shows exactly which vendors are targeted, where attack infrastructure concentrates, where reputation-based defenses have coverage gaps, and what defenders can do about it.

The Verizon 2025 DBIR documented an 8x increase in edge device exploitation — edge vulnerabilities jumped from 3% to 22% of all vulnerability exploitation breaches. Mandiant M-Trends 2025 found the top four most frequently exploited vulnerabilities were all in edge devices — Palo Alto PAN-OS, Ivanti Connect Secure, Ivanti Policy Secure, and Fortinet FortiClient EMS. CISA issued Binding Operational Directive 26-02, requiring federal agencies to address end-of-support edge devices. This report puts numbers behind those findings.

What's Inside

- **VPN Targeting at Scale** — Palo Alto GlobalProtect: 16.7 million sessions, more than 3.5x Cisco and Fortinet combined. Vendor-by-vendor breakdown of what's absorbing the most traffic.
- **The Reputation Gap** — 52% of remote code execution attempts came from IPs appearing in GreyNoise for the first time — infrastructure so new that no threat feed had it yet. GreyNoise caught them on first contact. Why static blocklists alone aren't enough.
- **A 300,000-IP Residential Botnet** — Grew from 2,000 to 300,000 IPs in 72 days. 73% residential. Why traditional perimeter controls miss it.
- **Legacy CVE Exploitation** — Pre-2015 CVEs generated 7.3 million sessions — 4x more than 2023-2024 CVEs, with one 26-year-old vulnerability (CVE-1999-0526) accounting for the majority. Why the long tail of old vulnerabilities creates persistent exposure.
- **AI Infrastructure: Emerging Target** — 91,403 sessions targeting LLM inference servers. 175,000 Ollama instances identified as internet-exposed (SentinelLABS/Censys, January 2026). The newest attack surface at the edge.

https://www.greynoise.io/resources/2026-state-of-the-edge-report

GreyNoise
State of the Edge Report 2026

Where Your Edge Defenses Fall Short

GreyNoise measured 212 exploitation attempts per second across H2 2025 — and the patterns inside that volume expose specific, measurable gaps in common edge defense strategies.

The *2026 GreyNoise State of the Edge Report* analyzes 2.97 billion malicious sessions from 3.8 million unique source IPs. It shows exactly which vendors are targeted, where attack infrastructure concentrates, where reputation-based defenses have coverage gaps, and what defenders can do

VPN Targeting at Scale — Palo Alto GlobalProtect: 16.7 million sessions, more than 3.5x Cisco and Fortinet combined. Vendor-by-vendor breakdown of what's absorbing the most traffic.

The Reputation Gap — 52% of remote code execution attempts came from IPs appearing in GreyNoise for the first time — infrastructure so new that no threat feed had it yet. GreyNoise caught them on first contact. Why static blocklists alone aren't enough.

A 300,000-IP Residential Botnet — Grew from 2,000 to 300,000 IPs in 72 days. 73% residential. Why traditional perimeter controls miss it.

Legacy CVE Exploitation — Pre-2015 CVEs generated 7.3 million sessions — 4x more than 2023-2024 CVEs, with one 26-year-old vulnerability (CVE-1999-0526) accounting for the majority. Why the long tail of old vulnerabilities creates persistent exposure.

AI Infrastructure: Emerging Target — 91,403 sessions targeting LLM inference servers. 175,000 Ollama instances identified as internet-exposed (SentinelLABS/Censys, January 2026). The

https://www.greynoise.io/resources/2026-state-of-the-edge-report

The Cyber Resilience Act - Summary of the legislative text

Page Contents

The text below summarises the main provisions of Regulation (EU) 2024/2847, in order to support the general understanding of the Regulation in an accessible way.

This summary has been prepared by the Commission services and is not meant to systematically cover the full scope of the Regulation. This summary is not representative of the European Commission's official position. Natural or legal persons who may be subject to the Cyber Resilience Act (CRA) should refer to the text of the Regulation published in the Official Journal of the European Union.

Quick Links

Cyber Resilience Act - Legal text

FAQ - Cyber Resilience Act implementation (EN)

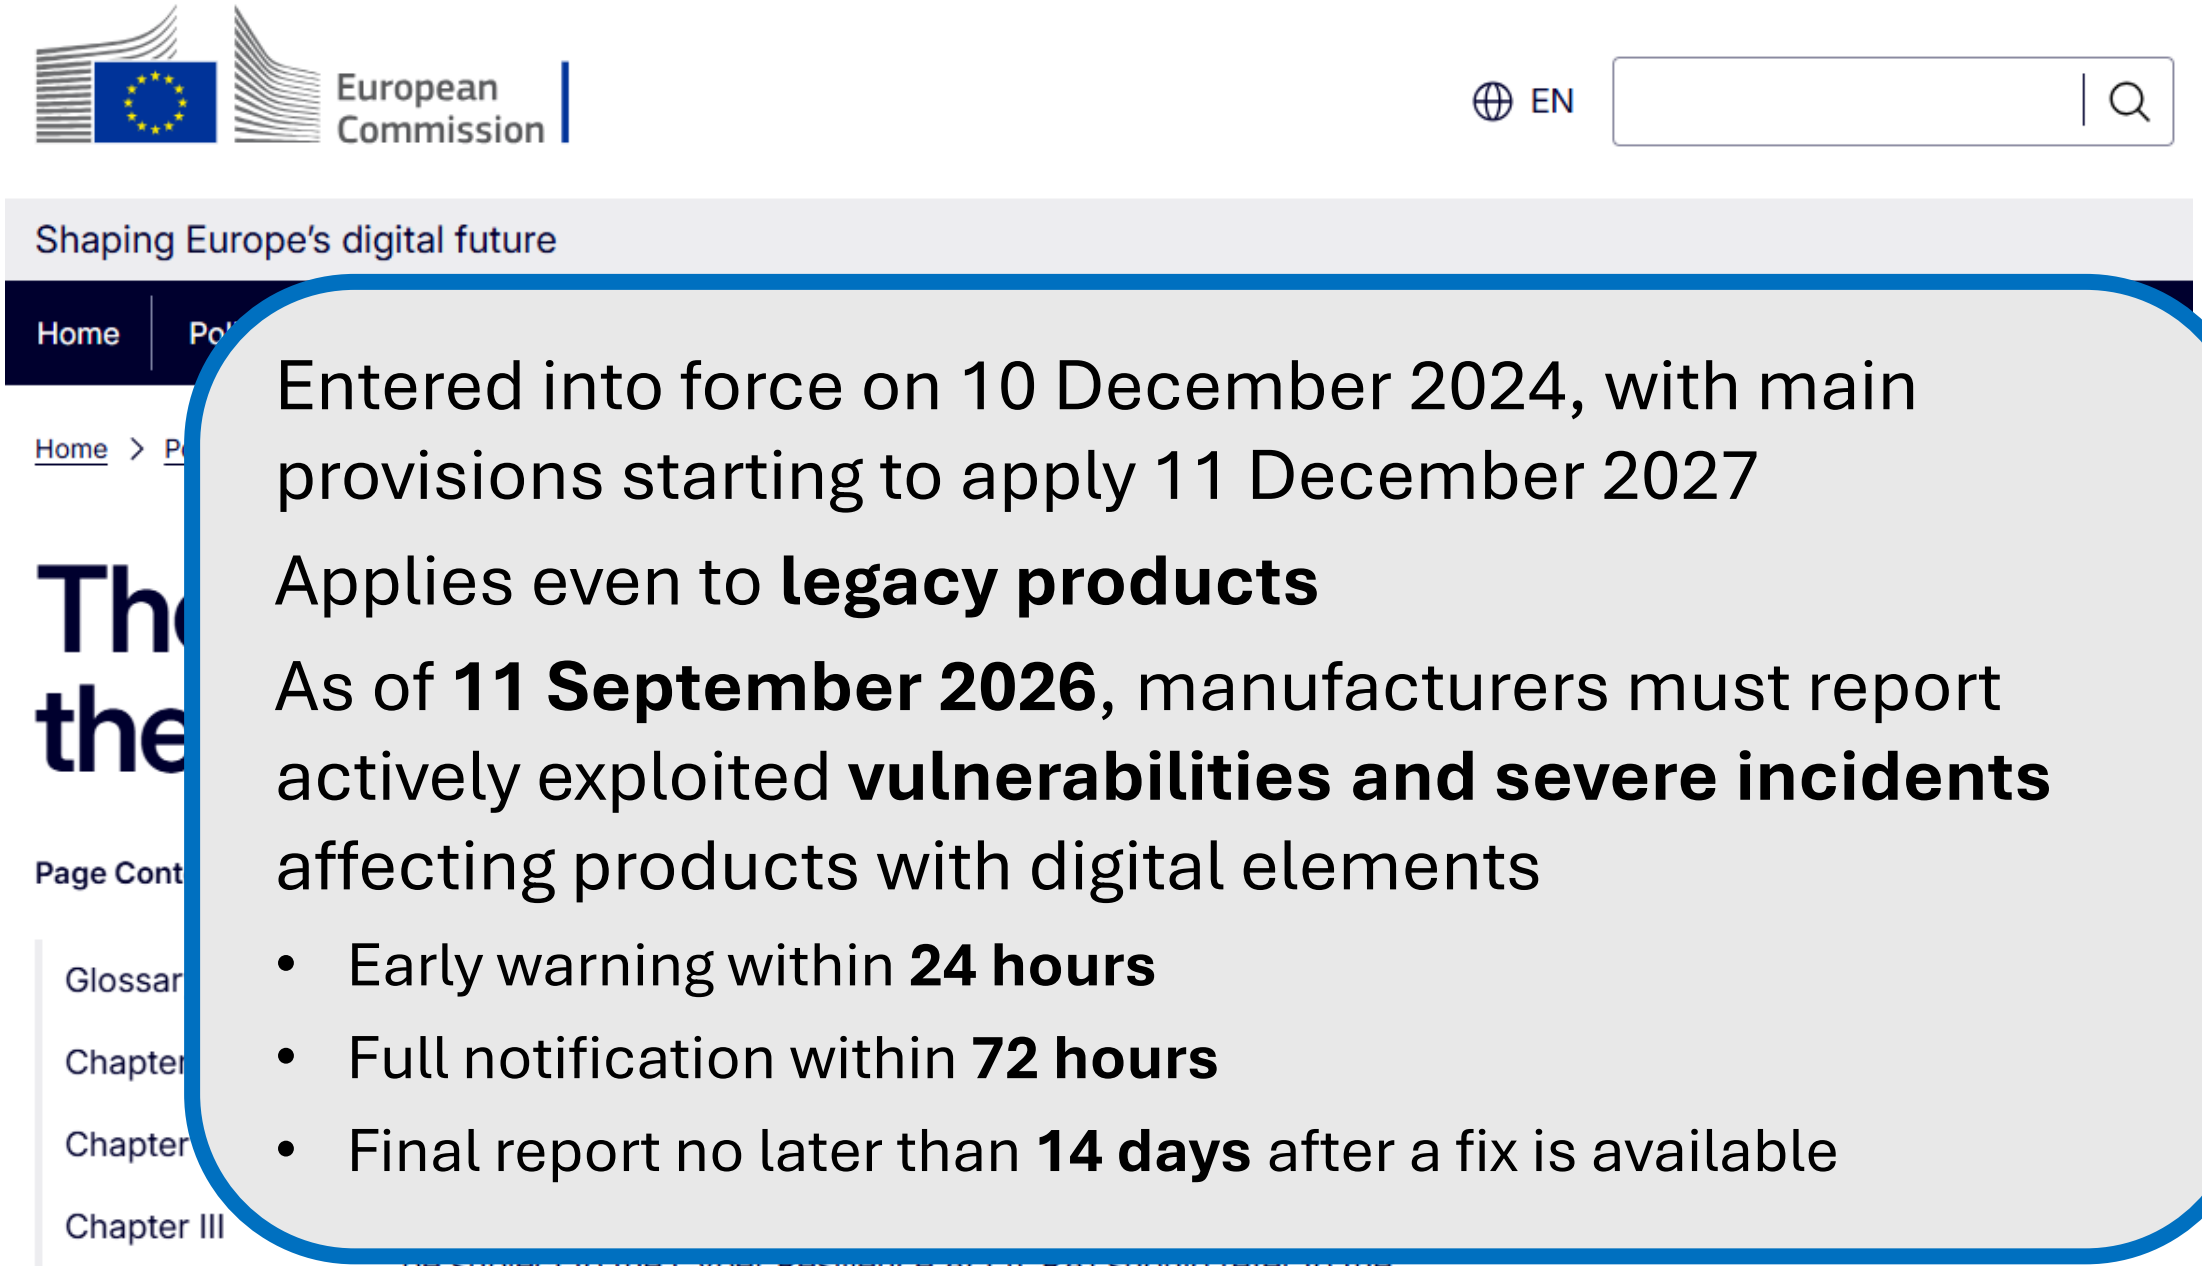
European Commission
EN
Shaping Europe's digital future
Home
Home
Page Cont
Glossar
Chapter
Chapter
Chapter III
Chapter IV
Chapter V
text of the Regulation published in the Official Journal of the European Union.
Entered into force on 10 December 2024, with main provisions starting to apply 11 December 2027
Applies even to **legacy products**
As of **11 September 2026**, manufacturers must report actively exploited **vulnerabilities and severe incidents** affecting products with digital elements
• Early warning within **24 hours**
• Full notification within **72 hours**
• Final report no later than **14 days** after a fix is available

Spectre v2: Branch Target Injection

In A Small Nutshell

Abuses **indirect branch prediction**, **speculative execution** and **cache timing side-channels**

Tricks the CPU into speculatively **executing memory it wouldn't have executed otherwise**, by poisoning indirect branch prediction

If speculative execution leaves state behind in cache that state can be inferred using **cache inference attacks**

This allows attackers to **read privileged memory**

Spectre v2: Retpoline

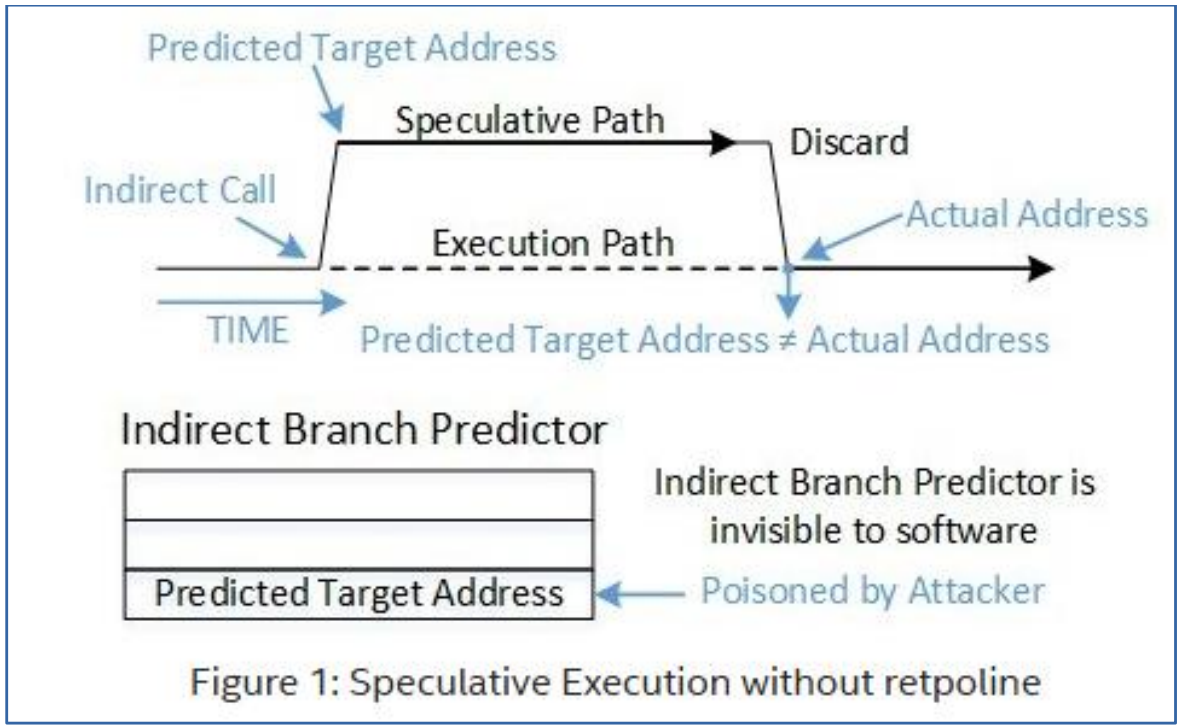

Figure 1: Speculative Execution without retpoline

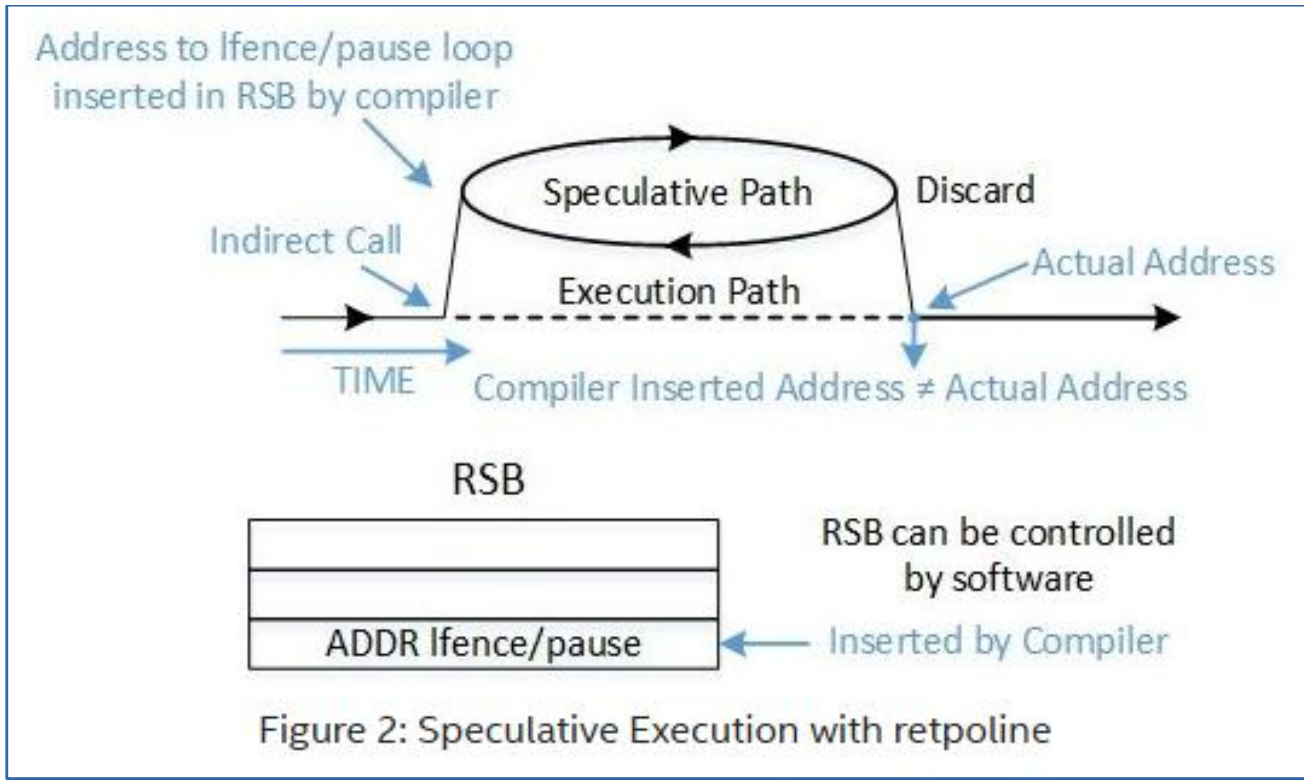

Figure 2: Speculative Execution with retpoline

The Attack vs. The Mitigation

So, what IS on the other side?

Product ideas,

designs,

design reviews,

implementation,

testing,

rollout,

operations and support.

Simple enough, right?

A detection idea is not a detection product

A one-off RE finding is not a malware sample triage pipeline

A PoC is not a weaponized exploit

A productionized detection model degrades over time

What worked on my machine doesn’t work on your machine

“We re-implemented an existing technique and made it 30% more reliable in production" doesn’t make a great paper for Usenix Security, but is exactly what product research and development does

Security Engineering, Development and Operations

What it means to work at scale

"Good enough" isn't good enough

A detection rule with **0.01% false positive rate** produces **100.000 false positives a day** at a rate of a billion processed daily events

What it means to work at scale

“Good enough” isn’t good enough

Every plausible misconfiguration exists

Every weird-but-legitimate workflow exists

Every legacy integration exists

In research, the corner cases are the footnotes

In production, the corner cases are the job

Security trade-offs are observability budget decisions in disguise

Adversarial drift, and, environmental drift

Miscalculations – We assumed the attacker wouldn't

When scale gets back at you ....

Operations vs. Security

Why do failures of large systems matter in security?

"Security operations is like
site reliability engineering
with an adversary in the loop"

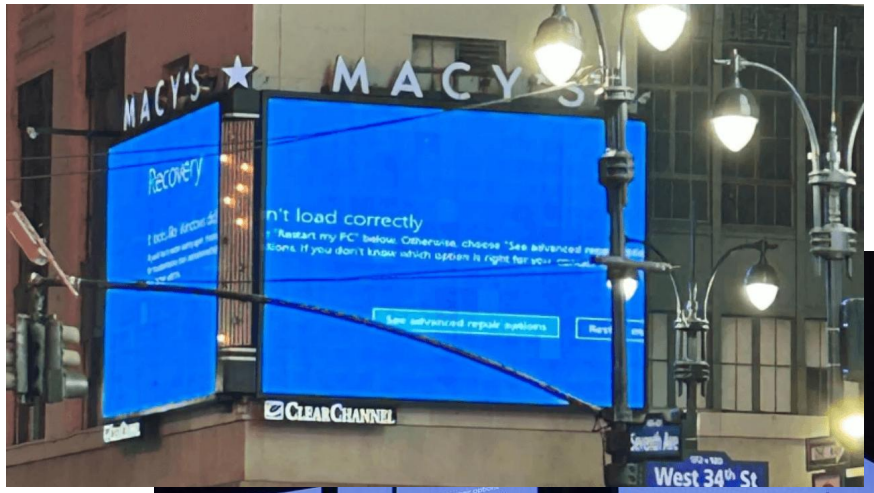
MACY'S
MACY'S
Recovery
n't load correctly
See advanced repair options
CLEAR CHANNEL
West 34th St


W 49 St

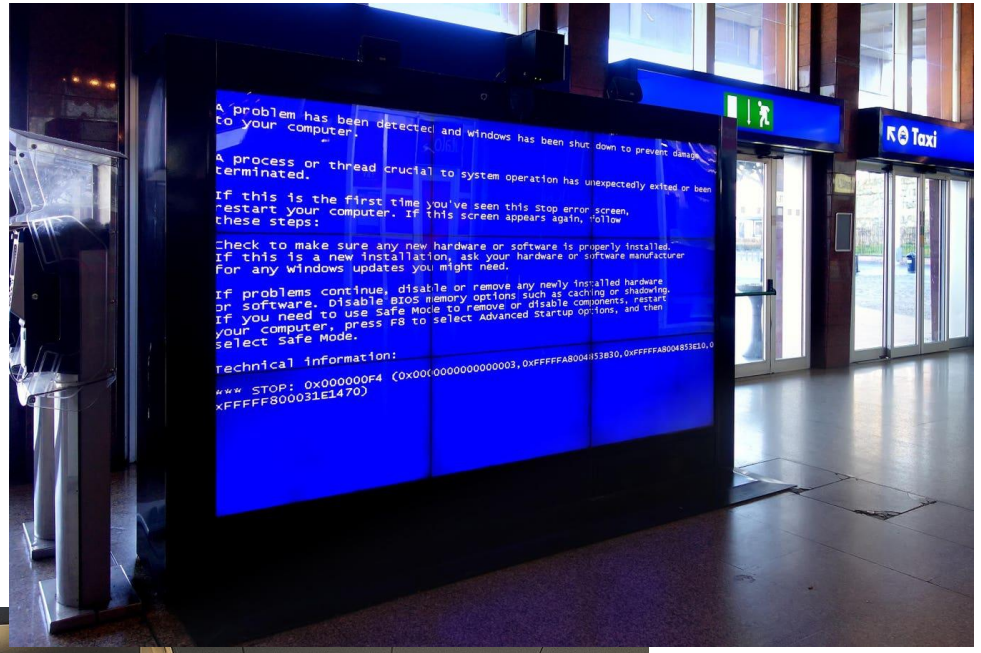
A problem has been detected and Windows has been shut down to prevent damage
to your computer.
A process or thread crucial to system operation has unexpectedly exited or been
terminated.
If this is the first time you've seen this Stop error screen,
restart your computer. If this screen appears again, follow
these steps:
Check to make sure any new hardware or software is properly installed.
If this is a new installation, ask your hardware or software manufacturer
For any Windows updates you might need.
If problems continue, disable or remove any newly installed hardware
or software. Disable BIOS memory options such as caching or shadowing.
If you need to use Safe Mode to remove or disable components, restart
your computer, press F8 to select Advanced Startup Options, and then
select Safe Mode.
Technical information:
*** STOP: 0x000000F4
xFFFFF800031E1470)
Taxi


Card only
:(
Your PC ran into a problem and needs to restart. We're just collecting some error info, and
then we'll restart for you.
100% complete

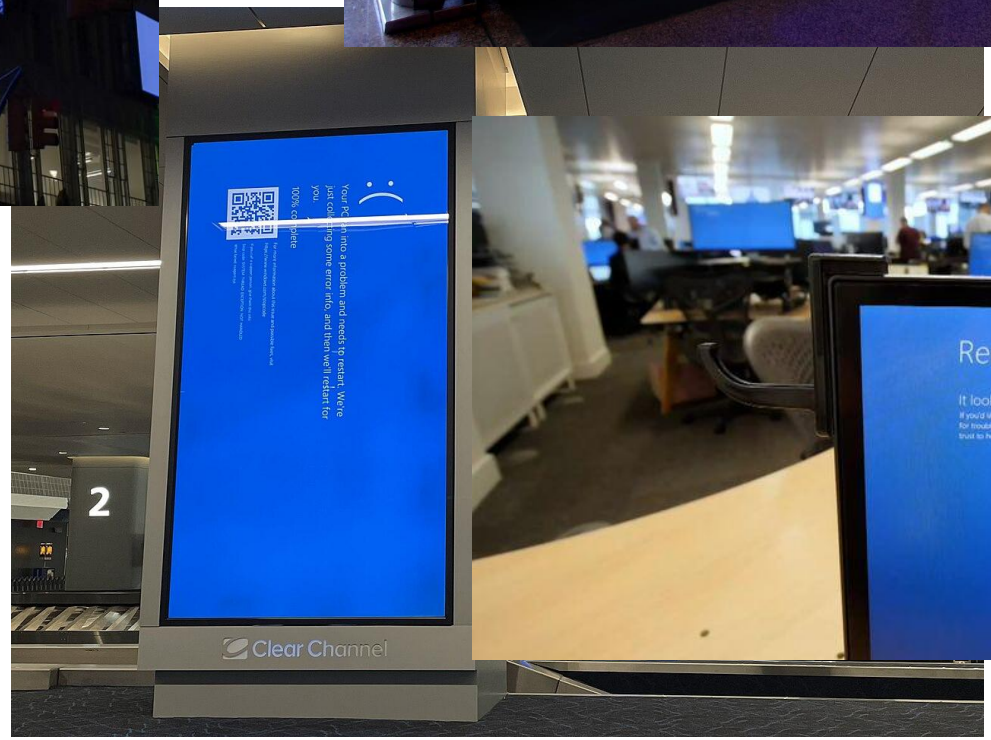
2
Clear Channel

Recovery
It looks like Windows hasn't loaded correctly.
Restart my PC

Channel File 291

At scale, security products become an attack surface, and a failure surface with massive blast-radius

- They have privileged access
- They run on every host
- Their time-to-update is instant

A detection catching one incident is invisible

A detection breaking production is a Wikipedia entry

Where Processes Fail

Validation that validates the wrong thing

Reviews/tests/staged rollouts skipped for urgency

Latent bugs triggered by routine operations

Safety mechanisms failing to wrong default

Recovery tooling dependent on system they protect

Tribal knowledge concentrated in one person

Backups and recovery paths that are outdated

What happens when things go sideways?

A Digital Post Mortem

Blameless – The point is to learn

Cause and trigger are different things

A post mortem you don't act on is just a story

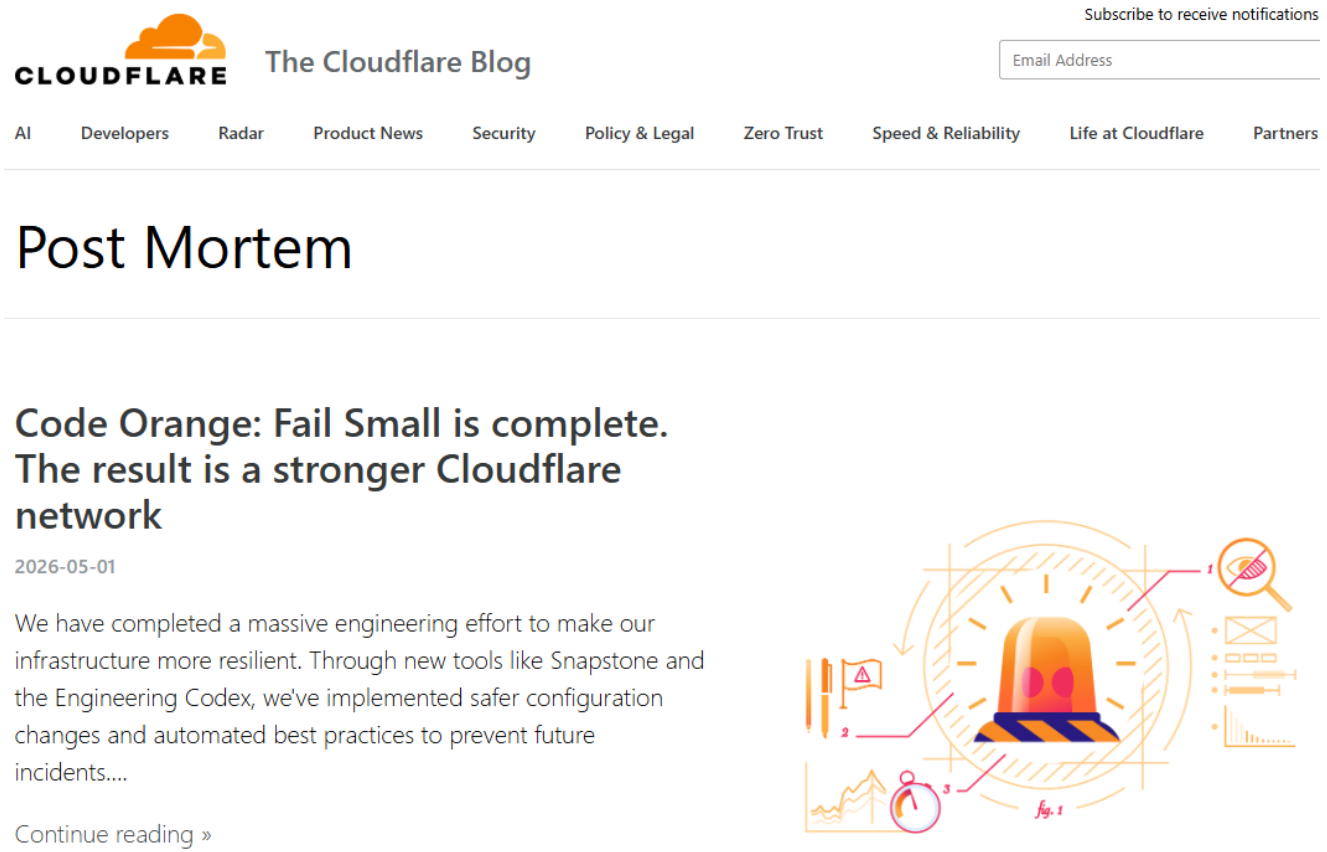

Cloudflare Outage on November 18, 2025

Cloudflare's core proxy started failing globally, taking down **large swaths of the internet**, incl. Twitter, Discord, OpenAI, DigitalOcean, 1Password, etc. etc.

A **permission change** on a reporting database cascaded into an outage with massive blast radius

A SQL query returned duplicate rows, this doubled the size of the target file, and the **proxy parser had a hard coded buffer limit**

Affected module crashed, tore down entire proxy service with it

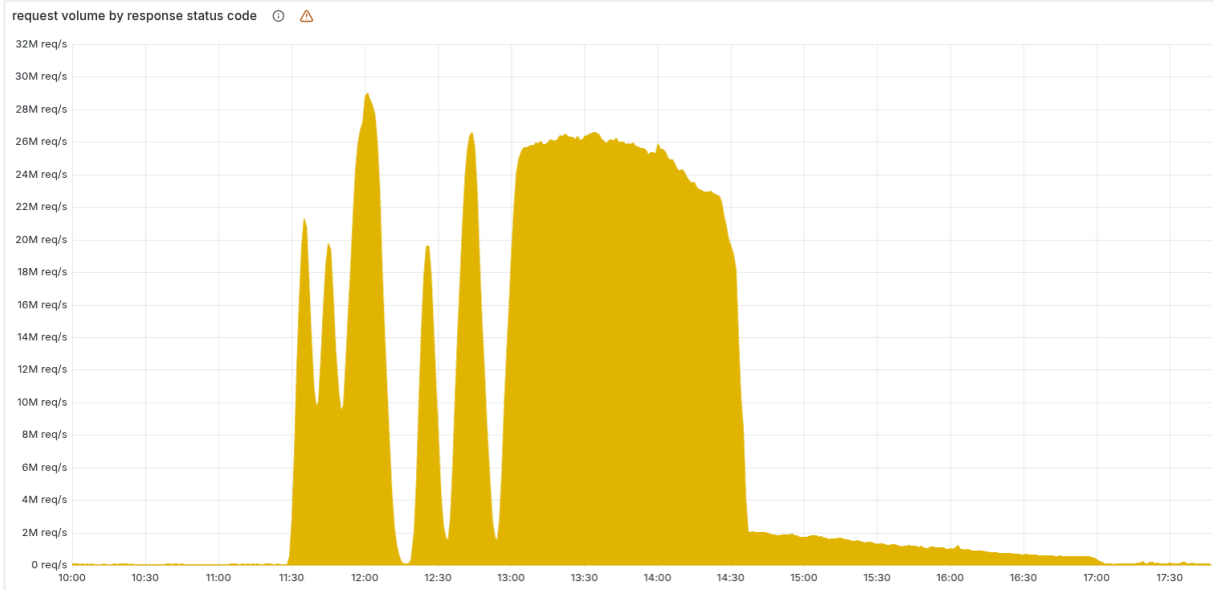

A Note on Operational Excellence

Operations is a design principle, not an afterthought

- Design for observability, response time, change management, runbooks, post mortems

PoCs are built for demonstration, not for operation

⇒The transition is a design process, not a mere engineering task

The AI Existential Crisis

Are we all powerful,
or out of a job?

Are we all powerful, or out of a job?

What can we currently automate?

Which jobs are disappearing?

What will the future look like for the experts?

AI is improving rapidly – Why

The models are improving, problems are solved faster

But what makes a model better?

- More **computing power**, larger and better datasets, and architectural refinements
- Improved **training techniques** like reinforcement learning on verifiable tasks
- **Data quality and curation** improvements
- **Scaffolding** around core models: Longer context windows, tool usage, multi-step agentic behavior

How is security changing in times of AI?

The easy tasks are going away

Challenging tasks are more tangible

Job descriptions are changing

Offense, defense and response are all speeding up

- At different speeds…

What defines speed, other than token budget?

Will technical skill matter in the future?

AI and the Security Community

Faster learning

More dopamine

Burnout

Layoffs

What is AI doing to the security community?
Evening out the playing field, or leaving most people in the dust?
Can we still grow the next us, if there is no more ladder to climb?

BlackHoodie

The original goal of BlackHoodie

What happens if we replace juniors with agents

What happens to juniors if they're not hired anymore

Building community: The young people, the veterans, the ecosystem between all of us!

How can we sustain our community?

Create opportunities

Mentorship matters more than ever

Embrace the situation and use AI to further skillsets faster

Research is **curiosity-driven** and produces **possibilities**

Production is **constraint-driven** and produces **reliability**

AI is now compressing both timelines simultaneously:
Research output is **faster**, production iteration is **faster**,
attackers move **faster**, defenders move **faster**

The existential question for the community is not "will AI replace us" but **"can we still grow the next us, fast enough"**

The Bottom Line