spectrum · DELTA

THE DETECTION DEBT REPORT

2026

RULES DEPLOYED

RULES THAT WORK

29.2%

of rules are broken on average

Spectrum's analysis of detection performance, based on customer data.

SPECTRUM.SECURITY

CONTENTS

01 · EXECUTIVE SUMMARY

Many security teams report excellent detection coverage, but do those detections actually work? We measured the detection profile of our customers before they started working with us, and what we found might surprise you: **nearly 3 in 10 detection rules are broken**.

556
RULES PER ENVIRONMENT
The average deployed, enabled rule count.

29.2%
BROKEN
Deployed and enabled, but unable to fire.

~26k
MISSED ALERTS / WEEK
Per environment, behind a broken rule.

WHAT WE FOUND

We define broken rules as rules that are deployed, enabled, and unable to produce an alert. Some examples of the root cause for a broken detection could be a missing index, a field that has never been mapped, or a log source that drifted or stopped showing up.

Two thirds of these failures are data-selection problems, and not driven by bad logic. This points to a systemic broken feedback loop between the SIEM and detection engineering.

Interestingly, these broken rules clustered into some high-value rules: MFA disabled, mailbox forwarding rules, audit logging stopped were three examples broken in the majority of organizations that deployed them.

WHERE THIS DATA COMES FROM

SOURCE
Real customer data from production SIEM deployments.

SAMPLE
Organizations across financial services, healthcare, pharmaceuticals, software, retail, energy and managed security, spanning six different SIEM platforms.

COVERAGE
More than **250 distinct log sources** across endpoint, cloud, SaaS, network and identity.

WINDOW
7 days.

METHOD
Every rule was verified against the customer's **live SIEM**, and then replayed against real data to count the alerts it would have raised.

WHY THIS DIFFERS FROM WHAT YOU HAVE READ BEFORE

Previous benchmarks analyzed **exported rule files**, which could catch a malformed query. We use the context of customer's live systems to surface rules broken due to the data they are collecting.

Detection estates are large and fairly inert.

556
RULES PER ENVIRONMENT

The average estate is big enough to feel covered.

About 162 of them, on average, cannot produce an alert under any circumstance.

16.2%
BEST ENVIRONMENT

Even the healthiest estate carried real debt.

Minimum = 16.2%, Median = 27.8%, Max = 45.7%.

r = 0.08
RULES VS. BROKEN RATE

Bigger rule sets are not better rule sets.

Rule count and broken rate are effectively uncorrelated. Fewer rules doesn't mean you'll have a lower broken rate.

67%
DATA SELECTION

Rules break because of data, not syntax.

Two thirds of broken rules fail on the data they read rather than the logic they contain. Genuine syntax errors account for under 1% — those get caught.

COMMON REASONS RULES CAN'T FIRE

1 **The field moved** — 30%
The value still exists, but it now lives somewhere else. Usually a vendor changing its logs, or a parser change.

2 **The field value changed** — 26%
Logs flowing and fields populated, but the value changed underneath a field the SIEM still reports as healthy.

3 **The logs were never there, or went dark** — 22%
Never ingested, or stopped arriving more than 30 days ago.

ANONYMIZED EXAMPLE: MFA NOT FIRING

Here's a real example: a customer had a rule written to detect multi-factor authentication being disabled. Originally the field `target_type` should match `"account_name"`. One day the vendor started writing these same logs with a value of `"user_name"` instead.

```
VendorAuditLogs_CL
| where action == "mfa_disabled"          ← still matches
| where target_type == "account_name"     ← now written as "user_name"
```

Detection deployed
Enabled, healthy, reported as covering MFA-disable.

The vendor changed a value
Same field, still populated — `target_type` now reads `"user_name"`.

30 days later
8 MFA-disable events occurred, but **0 alerted** because of this change.

03 · DETECTION DEBT BY SIEM

We found some interesting variation across different SIEMs, and between the same SIEMs at different customers.

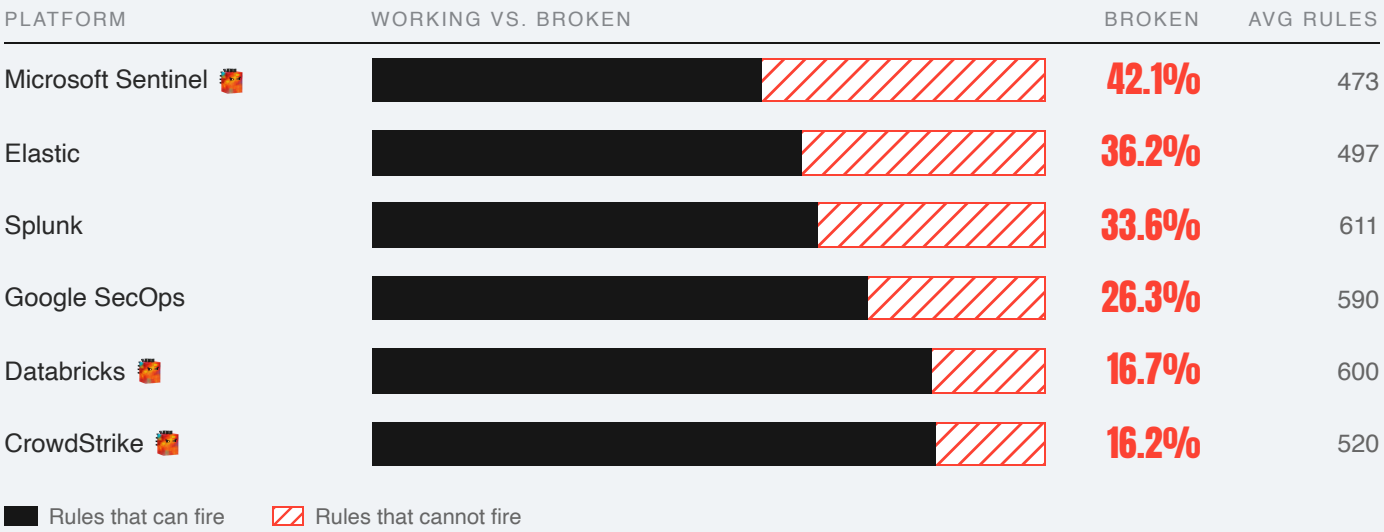

PRACTICE BEATS PLATFORM

The same SIEM at different customers had huge spreads in broken detection rates. Within Elastic, environments ranged from **16.9% to 45.7%** broken. Google SecOps ranged from 18.8% to 39.9%, and Splunk from 18.0% to 39.2%.

Whatever separates a healthy estate from a broken one, it is not the logo on the console.

CURATED BEATS CLONED

The two healthiest environments had something in common: almost no vendor-shipped content. Nearly every rule had been written for that specific environment, against telemetry the team knew it was collecting.

The worst were the opposite: large prebuilt rule packages switched on against data streams that weren't properly configured.

Reading this chart. Bars show the share of each platform's pooled rule inventory that can and cannot fire; "avg rules" is the mean deployed rule count per environment on that platform. Platform-level rates are pooled across environments and are indicative, not conclusive — particularly for the platforms marked .

04 · MOST WIDELY BROKEN DETECTIONS

Detection debt is not randomly distributed: our team noticed the same critical detections fail across a variety of organizations.

	DETECTION THEME	ENVIRONMENTS AFFECTED	RULES BROKEN	PLATFORMS
1	**MFA disabled / login without MFA**	88%	52%	Elastic · SecOps · Sentinel · CrowdStrike
2	**Inbox forwarding / mailbox rule created**	75%	70%	Elastic · SecOps · Sentinel
3	**Audit logging stopped / logs cleared**	83%	57%	Elastic · SecOps · Sentinel
4	**Guest / external user invited or enabled**	80%	89%	Elastic · SecOps · Sentinel
5	**Cloud root-account activity**	50%	25%	Elastic · SecOps

THE BUSINESS-EMAIL-COMPROMISE CHAIN HAS A DEAD LINK EVERYWHERE

Phishing-delivered detections were broken in **every** environment that deployed them (61% of rules). Impossible-travel: two thirds. Put the chain together and **every organization that deployed it had at least one stage that could not fire**.

PREBUILT ENDPOINT CONTENT FAILS ALMOST EVERYWHERE IT LANDS

Windows persistence rules (autorun, scheduled task, WMI) were broken at an **85%** rate, while lateral-movement tooling was broken 73% of the time. Ransomware behavior and credential dumping were also broken at over 50%. These are Sigma-derived packs that just didn't have the right data they needed.

DETECTIONS FOR AI

Rules covering **LLM and AI usage** (Bedrock, Claude, OpenAI) were broken 50% of the time, and source-control platform security rules were broken in every single environment.

05 · ALERTS & MISSED ALERTS

We replayed broken rules against seven days of real data to show exactly how many alerts would be missed.

~165k
ALERTS / WEEK
Average per environment.

~26k
MISSED / WEEK
Average per environment, behind a broken rule.

15.9%
OF ALL ALERT VOLUME
Sits behind a rule that cannot fire.

95.6%
WORST ENVIRONMENT
Almost nothing its SIEM should raise, reaches anyone.

TWO ENVIRONMENTS ARE EFFECTIVELY BLIND

At one environment, **95.6%** of the alerts its SIEM should have raised sat behind a broken rule. At another, **92.7%**. Both consoles reported a healthy, active estate throughout.

SOME RULES DRIVE MOST OF A COMPANY'S ALERTS

Separately from what is missing, what does arrive is startlingly narrow. At one financial services firm, **a single rule generated 98.5% of the week's entire alert volume**. At three others, one rule accounted for over 84%.

MOST MISSES CANNOT EVEN BE COUNTED

Only about **8% of broken rules even return a replay count**, as the rest fail before they can even execute. This means that these missed alerts values are just a **floor**, and actual values are much higher.

THE MOST EXPENSIVE BROKEN RULES IN THE SAMPLE

	DETECTION	ENVIRONMENT	CAUSE	MISSED (7 DAYS)
1	Suspicious shortcut execution from an unusual directory	Managed security provider	Data selection	**70,000**
2	Multiple sign-in locations — identity anomaly	Energy-infrastructure manufacturer	Data selection	**56,232**
3	High-risk cloud workload activity	Direct-to-consumer apparel brand	Data selection	**53,521**
4	Data-access baseline deviation	Software company	Data selection	**43,545**
5	Account tampering — suspicious failed-logon reasons	Healthcare services company	Data selection, syntax	**10,079**

How this was measured. Each broken rule was replayed against the customer's live data over the same 7-day window, while deployed, enabled and unable to fire. Counts are alerts that should have reached an analyst and did not.

06 · METHOD & DEFINITIONS

WHAT "BROKEN" MEANS

A **broken** rule is one that is deployed, enabled, and determined by Spectrum's Health Check to be incapable of firing in the customer's live environment. Four failure classes qualify:

— the rule reads from an index, table, or dataset that does not exist;
— it references a field that is not present in the data it reads;
— its query cannot execute against the live system;
— its schedule is dead, so it never runs.

A rule that runs correctly and finds nothing is **not** broken.

HOW RULES WERE VERIFIED

Verification ran against each customer's **running SIEM**, not an exported rule file: field presence and data-source liveness were checked, and queries were executed. Broken rules were then replayed against seven days of the customer's real data to count the alerts they should have raised.

BENCHMARK

Comparisons are to CardinalOps, *2025 State of SIEM Detection Risk* (5th annual report), which reports 13% of rules broken based on rule-export analysis across a larger set of environments.

ABOUT DELTA

Delta is the research team behind Spectrum Security. They research and develop AI techniques for better detection engineering practices and outcomes.

spectrum · DELTΔ

SPECTRUM.SECURITY