CARDINALOPS

Detection Posture Management

THIRD ANNUAL REPORT

State of SIEM Detection Risk

Quantifying the MITRE ATT&CK gaps that lead to undetected attacks in production SIEMs

Table of Contents

Using the 196 Techniques in MITRE ATT&CK V13 as the baseline, we found that actual detection coverage remains far below what most organizations expect and what SOCs are expected to provide.

We found that, on average:

Enterprise SIEMs only have detections for

24% **of all 196 techniques** in the MITRE ATT&CK v13 framework

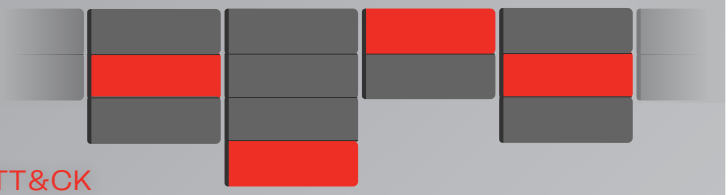

Enterprise SIEMs already ingest sufficient data to cover

94% **of all MITRE ATT&CK techniques**

In other words, we don't need to collect more data, but rather scale our detection engineering processes to develop more detections faster.

SIEMs are embracing "detection-in-depth" with detections for multiple Security Layers

But monitoring of Containers lags far behind at only:

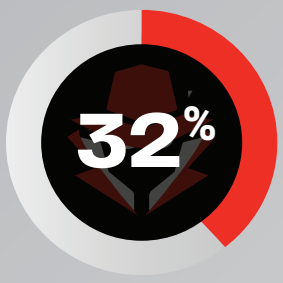

The most common layers are Windows, Network, and IAM (all at 96%).

12% **of all SIEM rules are broken and will never fire due to common issues such as misconfigured data sources and missing fields.**

1 Executive Summary

> "Use cases are the core of security monitoring activities. Organizations need a process to identify, prioritize, implement, and maintain security monitoring use cases. These processes cannot be too complex because security monitoring requires fast and constant changes to align with evolving threats."
>
> **Dr. Anton Chuvakin,** Office of the CISO, Google Cloud | Gartner Blog Post

Not much has changed since Anton wrote this blog post in 2016 – except that complexity and the rate of change have dramatically increased.

In this 3rd annual data-driven report, CardinalOps set out to gain visibility into the current state of detection coverage and use case management in enterprise SIEMs.

What did we find? Using the 196 adversary techniques in MITRE ATT&CK v13 as the baseline, we found that actual detection coverage remains far below what most organizations expect and what SOCs are expected to provide.

Worse, organizations are often unaware of the gap between the theoretical security they assume they have and the actual security they have in practice, creating a false impression of their detection posture.

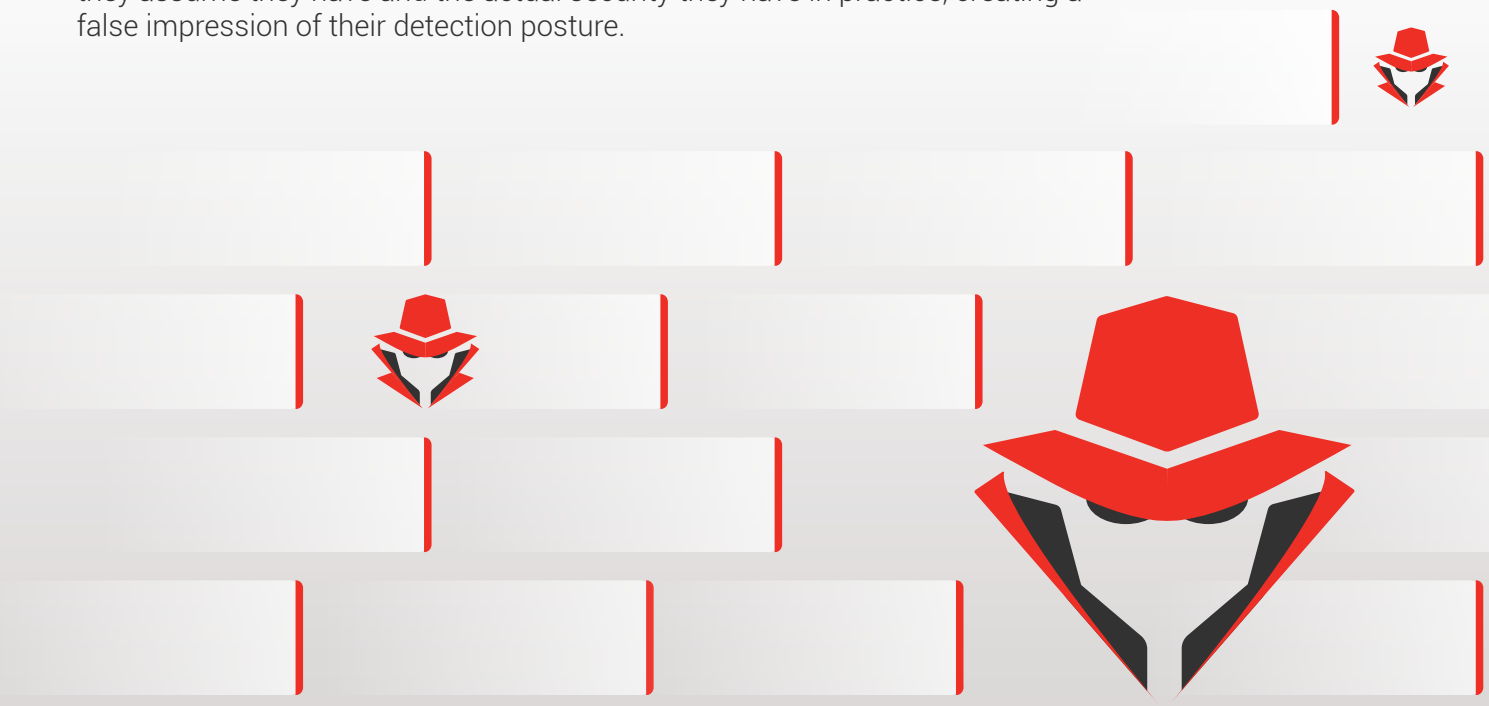

In particular, we found that, on average, enterprise SIEMs:

- **Only cover 24% of all MITRE ATT&CK techniques.**
  In other words, they're missing detections for 76% of MITRE ATT&CK techniques that adversaries use to breach their environments.

- **Are already ingesting sufficient data to potentially cover 94% of all MITRE ATT&CK techniques.**
  This suggests we don't need to collect more data, but rather we need to scale our detection engineering processes to develop more detections faster.

- **Are implementing "detection-in-depth" across multiple Security Layers**
  with the most common layers being Windows, Network, and IAM. In the middle: Linux/Mac, Cloud, Email, and Productivity Suites. At the bottom: Containers.

  The low result for containers is surprising because, according to Red Hat, 68% of organizations are running containers. Yet, our data shows that containers are generally not being monitored in the SIEM, perhaps because it's challenging to write high-fidelity detections to uncover anomalous behavior in these highly-dynamic assets.

- **Have 12% of their rules that are broken**
  and will never fire an alert due to common issues such as misconfigured data sources, missing fields, and parsing errors. This results in increased risk due to additional gaps that adversaries can exploit to breach organizations.

Most common Security Layers

96% Windows

96% Network

96% IAM

87% Linux/Mac

83% Cloud

78% Email

63% Productivity Suites

32% Containers

While both coverage and rule health metrics improved by a few percentage points since 2022, we attribute this primarily to differences in sampling, and also because we had a higher number of more mature organizations in this year's dataset.

What are the reasons for this disparity between actual and expected coverage?

- **Complexity**
  The average enterprise has more than 130 distinct security tools (endpoint, network, cloud, email, IAM, etc.). Each of these tools has its own log format, event types, and/or alert types, with each requiring unique detections to be developed based on a detailed understanding of how they function.

  As a result, according to Ponemon, more than 80% of security professionals rate the complexity of their SOC as very high, and less than 40% assess their SOC as highly effective.

- **Constant change**
  in infrastructures, security tools, attack surfaces, adversary techniques, and business priorities (e.g., cloud).

- **No "one-size-fits-all"**
  every enterprise is unique, making it impractical to copy-and-paste generic content from SIEM vendors, MSSPs, open source communities, and marketplaces.

- **Manual and error-prone processes**
  that are highly dependent on individual "ninjas" with specialized expertise, making it difficult to effectively scale and maintain high-quality detections.

- **Challenges in hiring and retaining skilled personnel**
  who can develop detections across diverse scenarios and log source types.

In section 2 of the report, we provide a series of best practice recommendations to help CISOs and detection engineering teams address these challenges and be more intentional about how detection coverage is measured and continuously improved over time. These recommendations are based on the experience of our in-house security team and SIEM experts like Dr. Anton Chuvakin, Office of the CISO at Google Cloud and former Gartner Research Vice President and Distinguished Analyst.

It is our goal with this report to help the security community move forward in recognizing the importance of bringing automated, repeatable, and consistent processes to detection engineering, and to provide independent benchmarks enabling CISOs and SOC leaders to answer the question "How prepared are we to detect the highest priority threats?"

2 Methodology

Rather than rely on subjective survey-based data, CardinalOps analyzed configuration metadata from real-world production SIEM instances to gain visibility into the current state of detection coverage in modern SOCs.

We examined aggregated and anonymized data across:

- **Diverse SIEM solutions** including Splunk, Microsoft Sentinel, IBM QRadar, and Sumo Logic.
- **More than 4,000 detection rules** with the largest SIEM we analyzed having more than 600 rules!
- **Diverse verticals** including banking and financial services, insurance, manufacturing, energy, media & telecommunications, professional & legal services, and MSSP/MDRs.
- **Nearly 1M log sources**
- **Hundreds of unique log source types**

Many of these organizations represent multibillion dollar, multinational corporations – making this the largest recorded sample of real-world SIEM data analyzed to date.

3 Continuing Importance of the SIEM

According to Forrester Research: "Ultimately, the SIEM remains the operating system of the security operations center, and it isn't going away …"

In fact, according to the SANS 2023 SOC Survey, SIEMs and EDR are the top two technologies considered critical to having an effective SOC.

This view is also supported in a Twitter poll that Anton conducted less than a year ago:

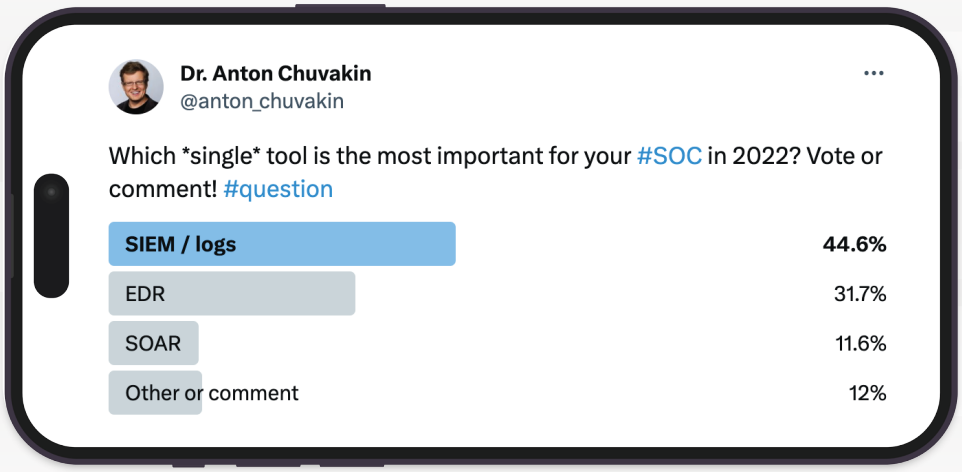

As one security leader recently explained to us, even if you're using EDR to detect and block malicious activity at the endpoint layer, you still need a SIEM with custom detections that act as a critical "backstop" to catch attacks that EDR solutions miss.

This can occur for several reasons including that sophisticated adversaries have figured out a way to disable or bypass EDR controls; relevant EDR alerts have been disabled due to excessive noise; or adversaries have devised a way to "hide in the noise" of untuned alerts.

So how do we measure and continuously improve our coverage for the threats most important to our organizations? The MITRE ATT&CK framework can help.

SEE NEXT SECTION

4 Why MITRE ATT&CK Matters

> "Since its creation in 2013, the MITRE ATT&CK framework has been of interest to security operations professionals. Based on ESG research, MITRE ATT&CK usage has now reached an inflection point. After nine years, MITRE ATT&CK and its use cases have evolved well beyond a reference architecture. In many ways, MITRE ATT&CK has become a "lingua franca" of security operations."
>
> **Jon Oltsik,** ESG Distinguished Analyst | CSO Online

As the standard framework for understanding adversary playbooks and behavior, MITRE ATT&CK now describes more than 500 techniques and sub-techniques used by threat groups such as APT28, the Lazarus Group, FIN7, and LAPSUS$.

According to ESG research, 89% of organizations currently use MITRE ATT&CK to reduce risk for security operations use cases such as determining priorities for detection engineering, applying threat intelligence to alert triage, and gaining a better understanding of adversary tactics, techniques, and procedures (TTPs).

The biggest innovation introduced by MITRE ATT&CK is that it extends the traditional intrusion kill chain model to go beyond static IOCs (like IP addresses, which attackers can change constantly) to catalog all known adversary playbooks and behaviors (TTPs).

These TTPs are grouped into both:

Tactics

Why
an adversary is performing an activity, such as Initial Access or Privilege Escalation.

Techniques

How
they are executing that activity, such as by exploiting a public-facing application or modifying a domain policy.

MITRE ATT&CK has also standardized our taxonomy vocabulary for both offensive and defensive teams. As Rick Howard, former CISO for Palo Alto Networks and now Chief Analyst at The CyberWire says in a recent CyberWire podcast:

> "Where the Lockheed Martin kill chain model is conceptual, the MITRE ATT&CK framework is operational. [Before the framework], we were all looking at the same activity and couldn't talk about it collectively in any way that made sense because each vendor and government organization had their own language and any intelligence coming out of those organizations couldn't be shared with anybody else without a lot of manual conversion grunt work. Talk about the Tower of Babel!"
>
> **Rick Howard,** Former CISO for Palo Alto Networks and now Chief Analyst at The CyberWire

It's also become the standard way to communicate to executive leadership about defensive posture and how it relates to recent attacks and vulnerabilities they heard about in the news (like Microsoft Outlook vulnerabilities, Follina or Okta PassBleed) — as well as answer the classic question "How prepared are we to detect the highest-priority threats?"

In this report

and in the CardinalOps Detection Posture Management platform

We use the MITRE ATT&CK framework to measure an organization's detection coverage across all these TTPs. The platform also helps organizations prioritize new detections to address gaps for the techniques that matter most to them, and delivers deployment-ready detections for their existing SIEMs, among other use cases.

5 Coverage for MITRE ATT&CK Techniques

Our data shows that enterprise SIEMs, on average:

Only have detections for

24% **of all 196 techniques**
in the MITRE ATT&CK v13 framework

This implies that adversaries can execute around 150 different techniques that will be undetected by the SIEM. Or stated another way, SIEMs are only covering around 50 techniques out of all the techniques that can potentially be used by adversaries.

Are already ingesting sufficient data to potentially cover

94% **of all MITRE ATT&CK techniques**

This suggests that we don't need to collect more data, we need to scale our detection engineering processes to develop more detections faster.

Of course, collecting more data from various security layers – such as network, cloud, IAM, and email – is a good thing and will provide deeper coverage via "detection-in-depth" because it results in a given technique being covered in multiple ways rather than just via one detection at a single layer such as the endpoint.

6 Health Metrics

Our data shows that enterprise SIEMs, on average:

- **Have 12% of their rules that are broken and will never fire an alert**

  due to common issues such as misconfigured data sources, missing fields, and parsing errors.

This commonly occurs due to ongoing changes in the IT infrastructure, vendor log format changes, and logical or accidental errors in writing a rule. Adversaries can exploit gaps created by broken detections to successfully breach organizations.

Here are some specific examples (for Splunk SPL) of some of the ways rules can break:

- **Sourcetype does not exist**
- **Index does not exist**
- **Lookup does not exist**
- **Scheduling has time gaps leading to missed alerts**
- **Sourcetype has not reported logs in the past X days**
- **Index has not reported logs in the past X days**
- **Sourcetype <–> Index are mismatched**
- **Logical operators are not in uppercase**
- **Parsing is incorrect**

Data quality issues:

- **Process Command Line is not being logged in Windows**
- **Key Vault changes are not being logged in Azure**

7 Most Common Security Layers in Enterprise SIEMs

This year we are reporting on the most common Security Layers being monitored in the SIEM.

Developed by CardinalOps, MITRE ATT&CK Security Layers extends the concept of ATT&CK coverage by measuring the "depth" of detection coverage for the first time. It does this by mapping each detection to a specific security layer – such as endpoint, network, email, cloud, containers, and IAM – and then enumerating the number of distinct layers covered for a given technique.

This enables SecOps teams to ensure they have "detection-in-depth" at multiple layers for the techniques that matter most to them.

Additionally, Security Layers enable organizations to link their coverage to desired business outcomes by immediately identifying blind spots related to crown-jewel assets such as their most sensitive applications and data. It also reveals missing telemetry and data sources that can be incorporated into their detection strategy to increase depth of coverage.

Our data shows that the most common security layers in enterprise SIEMs are:

1. **Windows — 96%**
2. **Network — 96%**
3. **Identity & Access Management (IAM)[1] — 96%**
4. **Linux/Mac — 87%**
5. **Cloud[2] — 83%**
6. **Email – 78%**
7. **Productivity Suites (primarily Office 365) – 63%**
8. **Containers — 32%**

[1] The IAM layer includes logs from Active Directory and Okta, for example.
[2] The Cloud layer includes public cloud logs such as AWS CloudTrail, AWS GuardDuty, GCP Audit Logs, and Azure audit logs, for example.

The low result for containers is interesting because, according to Red Hat research, 68% of organizations are running containers. Yet, our data shows that they are generally not being monitored for suspicious or anomalous behavior.

One explanation for this might be that, due to the dynamic nature of microservices-based application environments, monitoring them can be a hefty challenge and they are likely to bring a significant volume of data to SIEM platforms.

Another explanation might be that detection engineers are challenged by the prospect of writing high-fidelity detections to alert on anomalous activity for these highly-dynamic assets.

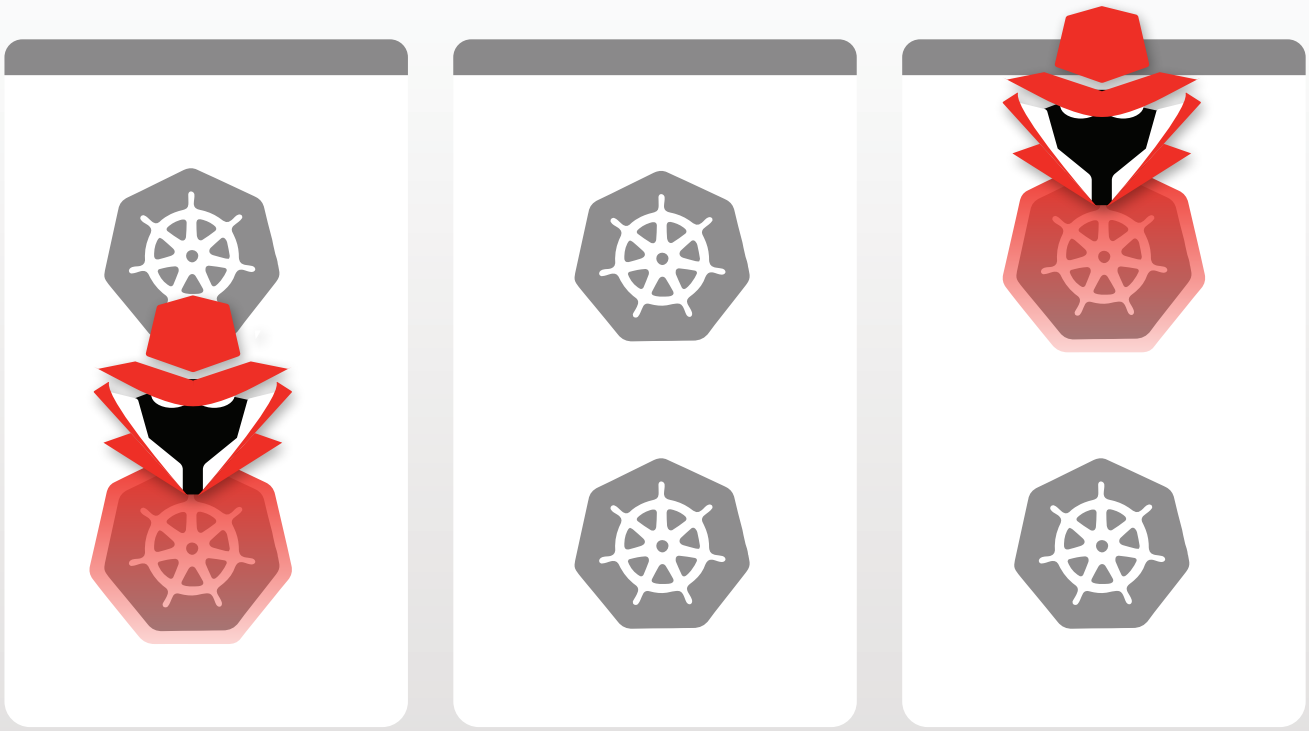

8 Best Practices for Detection Posture Management

> "Organizations need to become more intentional about detection in their SOCs. What should we detect? Do we have use cases for those scenarios? Do they actually work? Do they help my SOC analysts effectively triage and respond?"
>
> **Dr. Anton Chuvakin,** Office of the CISO, Google Cloud | SANS webinar on "The Future of SIEM"

Here are a series of best practice recommendations for enhancing detection coverage and detection quality in your SOC.

8.1 Review current SIEM processes

- **What is the approach for finding false negatives – and what adversary techniques, behaviors, and threats are currently being missed?**

- **How are use cases managed and prioritized? Typically, we find they're added to the backlog via an ad-hoc process driven by a combination of:**
  - Threat analysts & threat intelligence
  - Breach and attack simulation (BAS) tools
  - News about the latest high-profile attacks or vulnerabilities
  - Manual pen testing
  - Red teaming

- **How are detections developed today and what is the process for turning threat knowledge into detections?**

- **How long does it typically take to develop new detections?**

- **Is there a systematic process to periodically identify detections that are no longer functional due to infrastructure changes, changes in vendor log source formats, etc.?**

8.2 Become more intentional about how you develop and manage detection content

Focus on effectiveness, coverage, and improvements. Ask your team questions such as:

Do I?	What?	Are We?
Do I really detect it?	What do I need to detect based on our business priorities, crown jewel assets, industry sector, etc.?	Are we missing data sources that would improve our coverage in high-priority areas?
Do I detect it well?	What do I detect today?	
Do I triage and respond correctly?		

8.3 Build or refresh your use case management processes

Choose 3-5 enhancements to address the questions from the last section, with an agreed-upon timeline.

8.4 Measure and continuously improve

Detection engineering processes are no different than other security and IT management processes. As IT modernizes and uses DevOps and SRE approaches, so should the SOC. You can't improve what you can't measure. Many SOC metrics - focused on people, process, and technology - are needed for consistent improvement.

Set organizational goals around how to increase detection coverage and reduce the time to detect non-functioning rules.

9 CardinalOps Platform Overview & Top Use Cases

> "Security operations teams are at a crossroad. Organizations need unprecedented security operations scale and efficiency but continue to be dragged down by manual processes, skills shortages, and suboptimal technology usage. The MITRE ATT&CK framework can help, as it introduces an adversary view and structure for security operations. Organizations seeking to operationalize MITRE ATT&CK as a framework for identifying and remediating control gaps may want to consider detection posture management with CardinalOps."
>
> **Jon Oltsik,** Senior Principal Analyst and Fellow | Enterprise Strategy Group, "Operationalizing MITRE ATT&CK with Detection Posture Management"

Backed by security experts with nation-state expertise, the CardinalOps platform uses automation and MITRE ATT&CK to continuously assess your detection posture and eliminate coverage gaps in your existing SIEM – so you can easily implement a threat-informed defense.

What's more, it improves detection engineering productivity by 10x, reduces the need to hire additional personnel which are in short supply, and reduces mundane tasks for detection engineers.

Native API-driven integrations include Splunk, Microsoft Sentinel, IBM QRadar, Google Chronicle SIEM, CrowdStrike Falcon LogScale, and Sumo Logic.

Here are the top 4 use cases for the platform, with a description of all use cases shown here.

9.1 Map all your detections to MITRE ATT&CK

Most organizations are still using spreadsheets or other manual tools to understand their ATT&CK coverage. This is a mundane and time-consuming activity that takes your engineers and analysts away from more strategic activities – plus your visibility into your actual MITRE ATT&CK coverage is always out of date.

In order to map your ATT&CK coverage, our platform starts by connecting via the native API of your existing SIEM. It then ingests all your rules as well as metadata about your log sources (your sensitive log data never leaves the SIEM).

The platform then uses specialized, ML-based analytics and feature extraction to map your detections to the most appropriate ATT&CK technique and sub-technique, producing a heatmap and coverage score that's continuously updated whenever you add detections or the ATT&CK framework gets updated.

The heatmap and metrics can easily be filtered based on selected variables including APT groups, ATT&CK matrices, security layers (endpoint, network, IAM, cloud, etc.), and whether you want to examine covered or uncovered techniques.

MITRE ATT&CK coverage showing coverage and health metrics at top, and selected filters at bottom left.

9.2 Gain new detections to address critical gaps faster

Once you've identified your top priorities for eliminating coverage gaps – such as specific APT groups, Tactics and Techniques, or log source types – the platform delivers curated, high-fidelity detections to close the gaps.

Rules are delivered deployment-ready, meaning they're in the native query language of your SIEM and have been pre-validated and auto-customized for your environment, including your data sources, naming conventions, and indexes.

The platform makes it easy to quickly review, test, and push new rules into your SIEM with the click of a button (via its native API).

Plus, you gain access to a searchable rule catalog containing thousands of rules – covering hundreds of diverse data sources – including for the latest high-profile threats and vulnerabilities.

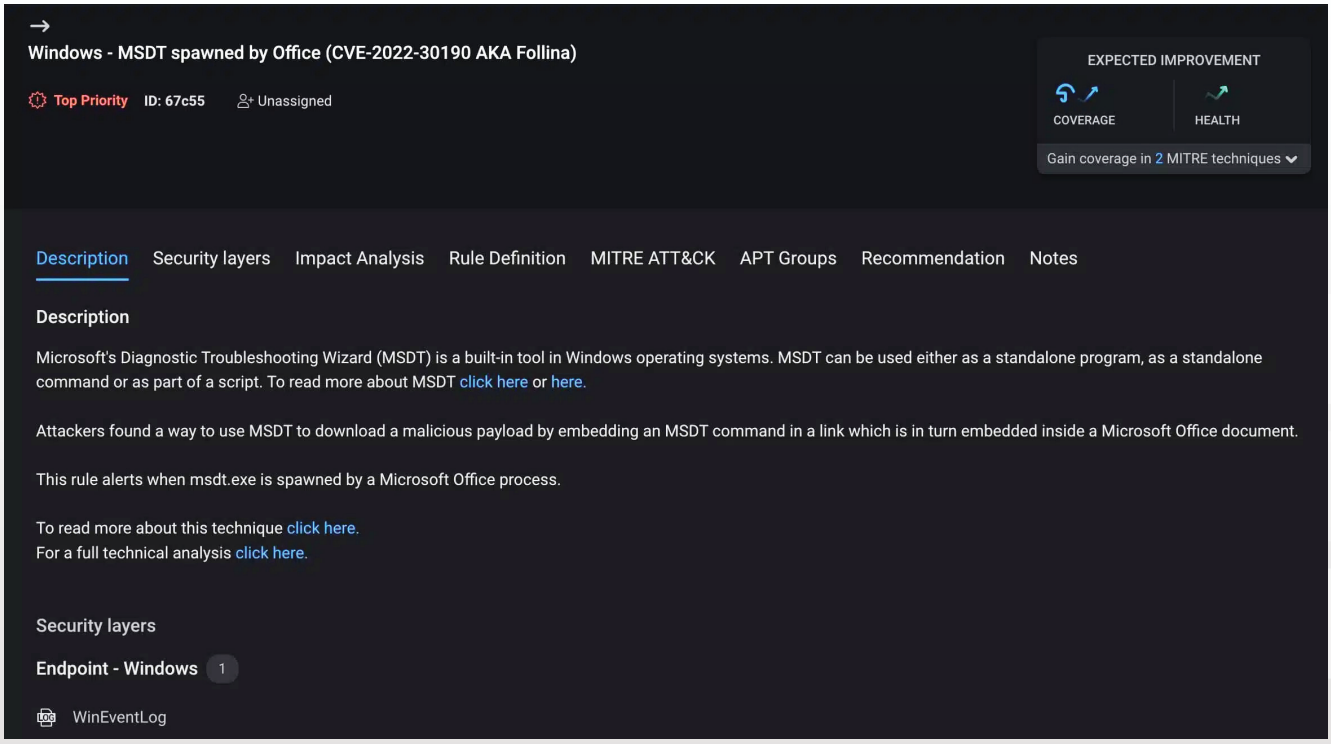

Rule Definition

```
index=wineventlog sourcetype=WinEventLog EventCode=4688 New_Process_Name=*msdt.exe Creator_Process_Name IN
(*windword.exe, *excel.exe, *powerpnt.exe ) "msdt" | stats values(Creator_Process_Name) as Creator_Process_Name,
values(Account_Name) as Account_Name,values(Process_Command_Line) as Process_Command_Line, values(dest_ip) as dest_ip
count by ComputerName
```

MITRE ATT&CK

TACTICS	TECHNIQUES	SUB TECHNIQUES
Execution (TA0002)	Exploitation for Client Execution (T1203)	
	User Execution (T1204)	Malicious File (T1204.002)

APT Groups

TECHNIQUES	SUB TECHNIQUES	APT GROUPS
Exploitation for Client Execution (T1203)		admin@338 (G0018), Andariel (G0138), Aoqin Dragon (G1007), APT12 (G0005), APT28 (G0007), APT29 (G0016), APT3 (G0022), APT32 (G0050), APT33 (G0064), APT37 (G0067), APT41 (G0096), Axiom (G0001), BITTER (G1002), BlackTech (G0098), BRONZE BUTLER (G0060), Cobalt Group (G0080), Confucius (G0142), Darkhotel (G0012), Dragonfly (G0035), Elderwood (G0066), Ember Bear (G1003), EXOTIC LILY (G1011), Higaisa (G0126), Inception (G0100), Lazarus Group (G0032), Leviathan (G0065), MuddyWater (G0069), Mustang Panda (G0129), Patchwork (G0040), Sandworm Team (G0034), Sidewinder (G0121), TA459 (G0062), The White Company (G0089), Threat Group-3390 (G0027), Tonto Team (G0131), Transparent Tribe (G0134), Tropic Trooper (G0081), FIN11, APT9
		admin@338 (G0018), Ajax Security Team (G0130),

Example of a new detection showing a description of the attack that is being detected by this rule; the full rule in the native syntax of your SIEM; and which Techniques and APT groups are covered by this detection. Once the rule has been manually reviewed and automatically tested using the past 30 days of historical log data (this is configurable), it can be pushed directly into the SIEM via the SIEM's API.

9.3 Continuously identify and fix broken rules

If you're like most detection engineering teams, you're continuously adding new detection rules to your SIEM. But over time, your environment has changed in different ways.

Your network has changed, your security tools have been upgraded to newer versions and log formats, older log sources have been retired, and your monitoring targets have changed.

And you may even have added generic rules that were copied and pasted from open sources or by an MSSP (and might contain RegEx errors that prevent proper parsing).

The result? Broken rules that will never fire due to misconfigured data sources, missing fields, parsing errors, and other data quality issues – creating additional gaps in your coverage.

This leads to a false sense of security because your CISO and SecOps team thinks they're protected — but then are surprised when your Red Team (or worse, an adversary) finds a hidden gap in your defenses and exploits it.

The CardinalOps platform uses specialized analytics to continuously analyze all your rules to ensure they have all required prerequisites to fire (log data, field values, etc.). But it doesn't just identify issues with broken rules, it delivers remediated rules that you can review, test and instantly deploy into your SIEM.

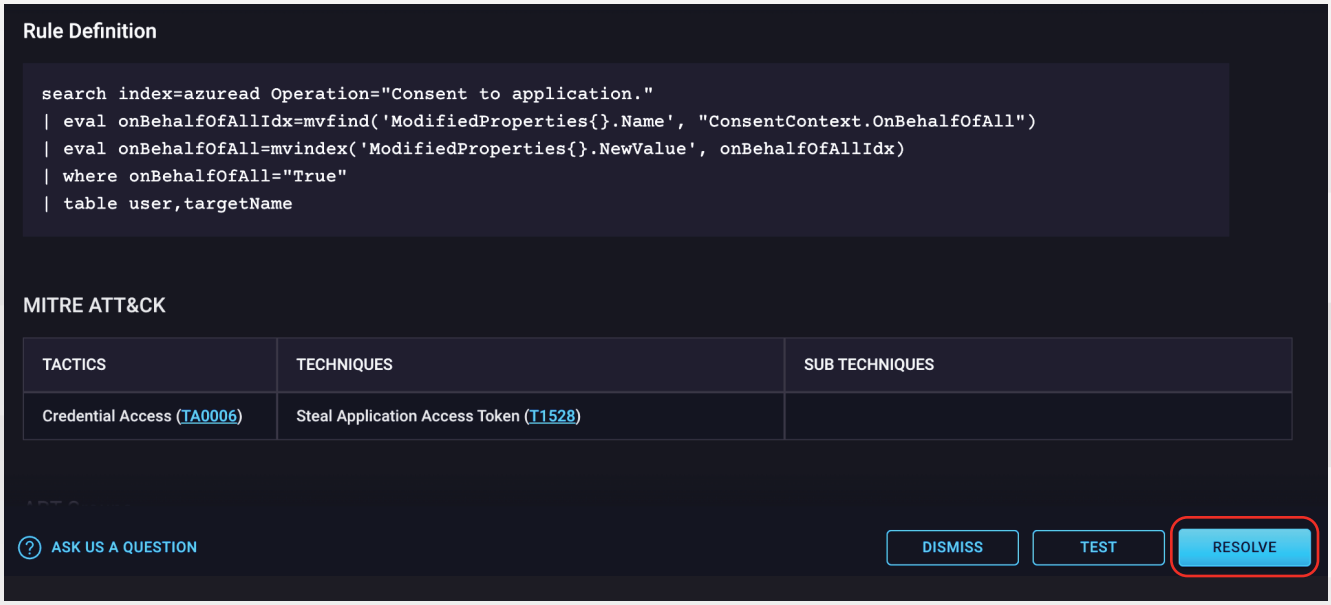

New and remediated rules can be pushed directly into the SIEM after manual review and automated testing.

And if remediation requires actions outside your SIEM – such as enabling event logging options on endpoints that were previously turned off – it delivers detailed recommendations on how to fix it, with links to technical documentation you can share with your IT team.

9.4 Automate to reduce need for additional personnel and eliminate mundane tasks

While automation has delivered significant benefits to multiple areas of the SOC – such as anomaly detection and incident response – the detection engineering function remains stubbornly manual and typically dependent on "ninjas" with specialized expertise.

With CardinalOps, you can now apply automation and analytics to supercharge the operational efficiency of your team and streamline the end-to-end process of researching, testing, and delivering new detections. Address the latest vulnerabilities. Onboard new log sources. And respond to ongoing requests from your Red Teams and threat intelligence teams.

You can also leverage automation to address more mundane tasks such as mapping your rules to ATT&CK, identifying and fixing broken detection and data sources, and tuning noisy detections.

The benefits? Higher productivity, greater agility, and cost avoidance from a reduced need to hire additional personnel. Plus happier staff members that are less likely to leave because they can now spend their time on more interesting activities such as threat hunting and researching new and novel attack techniques.

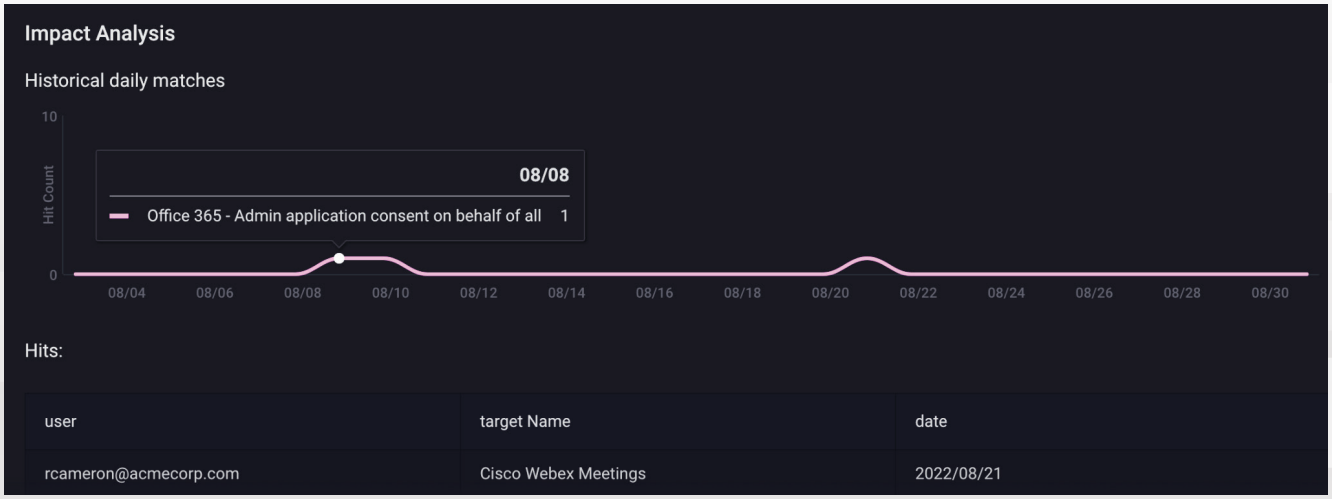

Impact Analysis shows an automated test workflow typically executed before deploying any new rules. The test shows if and when the new rule would have fired, had it been in place for the past 90 days, as a way to ensure the rule is not too noisy and/or to determine appropriate exclusions.

10 What Customers Are Saying About CardinalOps

"CardinalOps delivers the strategic expertise and automation we need to ensure our SOC is operating at maximum effectiveness and efficiency."

CISO, National Stock Exchange

"CardinalOps has been transformational for my team. Plus, time-to-value was extremely short. In our complex environment, it's not easy for vendors to get their solutions into production – at scale – but they promised us quick API-level integration with Splunk, and they delivered."

VP of Global Security Engineering & Architecture, Fortune 50 Manufacturer

"Splunk is the backstop we rely on to catch attacks our other security tools (like EDR) miss. CardinalOps ensures all our custom detections are working as intended and we aren't missing detections for the MITRE techniques and APTs most relevant to our organization. Plus the platform saves us a ton of time on MITRE mapping, and their team has been incredibly responsive."

Director of Information Security, $3B Global Corporate Law Firm

About CardinalOps

Backed by security experts with nation-state expertise, the CardinalOps platform uses automation and MITRE ATT&CK to continuously ensure you have the right detections in place to prevent breaches, based on a threat-informed strategy.

What's more, it improves detection engineering productivity by 10x, reduces the need to hire additional personnel which are in short supply, and reduces mundane tasks for detection engineers.

Native API-driven integrations include Splunk, Microsoft Sentinel, IBM QRadar, Google Chronicle SIEM, CrowdStrike Falcon LogScale, and Sumo Logic.

Learn more at CardinalOps.com

Learn more

CARDINALOPS
Detection Posture Management