SANS

Analyst Program

Product Review

Google SecOps: The SIEM's *Third Act*

Written by **Mark Orlando**

January 2025

Google Cloud Security

Introduction

Since their introduction to the security product market over 20 years ago, Security Information and Event Management (SIEM) platforms have faced challenges with complexity and scope creep. What began mainly as a means of consolidating alerts from intrusion detection systems and firewalls has become a central hub for all kinds of security functions: threat intelligence analysis, audit support, risk management, automation, detection engineering, and more. Vendors have responded to this demand by consolidating more tools and capabilities into their platforms, which has resulted in products that are powerful but complicated.

Entire product classes have emerged around reducing data ingestion costs and complexity and getting better detections out of the SIEM. If *Act One* of the SIEM "story" is log collection and search, and *Act Two* is threat detection and consolidation of functions such as security orchestration, automation, and response (SOAR) and user and entity behavior analytics (UEBA), is the SIEM now at risk of collapsing under its own weight? Is it time to pivot to a different form factor for detection and response, like extended detection and response (XDR) and its variants? The response based on our review of Google SecOps is a resounding "no." As we'll discuss in this paper, Google SecOps very successfully combines Google's large-scale search performance, world-class security expertise from Mandiant and VirusTotal, and artificial intelligence via the Gemini model in a streamlined detection and response platform.

The SIEM's *Third Act*

SecOps presents a compelling case for being the SIEM's *Third Act*, offering a modern and innovative approach to security operations. Like other SIEMs, SecOps is a consolidation of multiple products and feature sets, many of which come by way of acquisitions. Unlike many other SIEMs, SecOps has managed to integrate these capabilities in a way that feels deliberate, organic, and—most importantly—focused on the busy analyst! As we'll discuss in this paper, the product also leverages advanced automation features to reduce complexity and lower the barrier to entry for data entry, investigations, detection engineering, and other key capabilities. SecOps is a SIEM in the most literal sense of the phrase, not just a data analytics or search platform customizable for SOC use cases.

In a product space with a long history and many players, Google SecOps stands out thanks to some key differentiators:

- **Scalability and speed of Google infrastructure**
- **Unified visibility in hybrid and cloud-first environments**
- **Integrated threat intelligence based on Mandiant and VirusTotal**
- **Thousands of curated detections maintained by world-class experts**
- **AI-powered assistive automation**

This paper explores SecOps features and capabilities and summarizes our impressions of the product as practitioners with decades of experience in security operations and SIEM. We will look at how Google has implemented many core SIEM and SOAR features as well as innovations that are uncommon in the product space. If it has been awhile since you looked at the platform, or if you're simply interested in Google's unique approach to this core capability, this product is well worth a look.

A Threat-Informed SIEM for the Big Data Era

At the core of Google SecOps lies a robust data collection infrastructure designed to handle the massive data volumes characteristic of the Big Data era. As one might expect from a Google product, SecOps excels in ingesting, normalizing, and analyzing data at "Google scale and speed." There are multiple data onboarding options, including collection agents for on-premises data sources like Windows Event Logs, and comprehensive support for AWS, Azure, and Google Cloud logging and monitoring APIs. This comprehensive approach enables unified logging across diverse environments, including those with both cloud and on-premises resources, or even multiple SIEMs. For anyone who has worked with SIEM starting with the early on-premises days through "cloud-first," this is a welcome solution to unified monitoring in distributed, hybrid environments.

Streamlined workflows are central to SecOps design, allowing analysts to quickly and efficiently investigate potential threats and pivot through investigations. The platform features a fast and responsive search function. Analysts can save useful queries for later reuse or leverage Gemini to build custom detections based on their searches. The best solution to alert fatigue is *better alerting*, and SecOps comes equipped with tools to build a better detection function.

The Gemini model stands out as a particularly innovative feature, offering analysts the ability to ask investigative questions and receive answers without needing to understand the underlying data structure.[1] This capability eliminates the need to manually access and review external sources like raw logs, streamlining the investigative process. Gemini is fully baked into the product, featuring centrally in functions like Cases and linked in each screen via the Gemini "button." The combination of these AI features and proven investigative value-adds, like VirusTotal, Mandiant Threat Intelligence, and Safe Browsing data from billions of Chrome devices, makes SecOps tough to beat as a solution for threat-informed defense. Let's start by discussing SecOps' data collection capabilities.

Data Collection

SecOps offers an impressive set of data ingestion capabilities, further cementing its bona fides as a unified data platform. Logging options include collection agents (based on the BindPlane open source agent) for on-premises data sources, and support for major cloud logging and monitoring APIs for cloud-based logging.[2]

[1] https://cloud.google.com/chronicle/docs/secops/gemini-chronicle

[2] https://cloud.google.com/architecture/logging-on-premises-resources-with-bindplane

Ingestion and Parsing

If you have ever configured or managed a SIEM (see Figure 1 for the ingestion pipeline) the following will be familiar to you: log offloading via SecOps Forwarder on the endpoint or log source, and parsers that convert the log data to a standardized data model.[3]

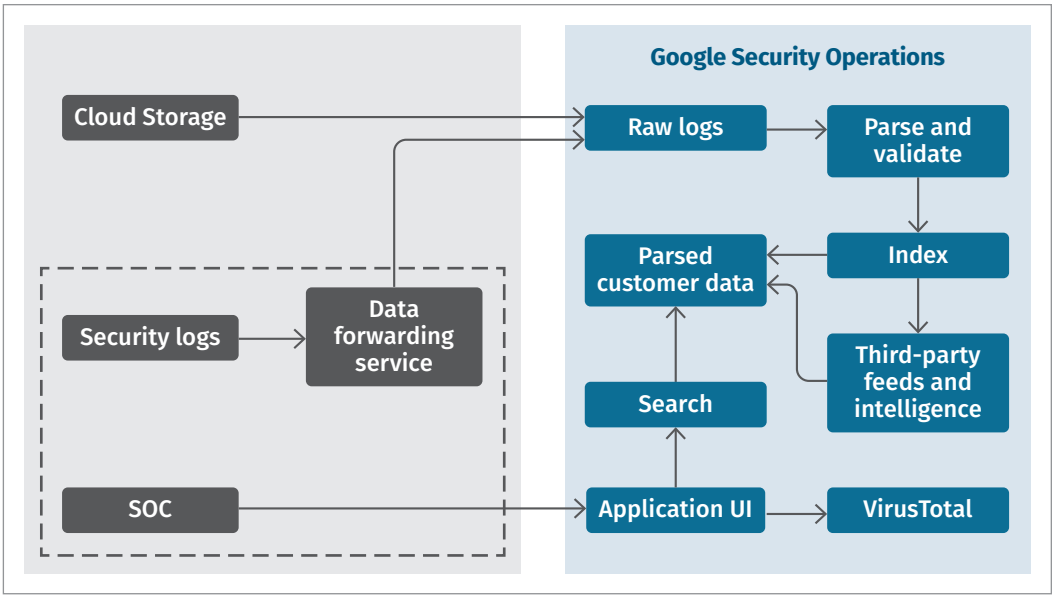

Figure 1. The SecOps Data Flow, Including Contextual Enrichment from Google Threat Intelligence

Other inputs, like stream data from an API, HTTPS webhook source, or cloud data storage, can be managed directly from the SecOps UI as feeds. SecOps performs data enrichment upon ingest, which reduces the effort required for log correlation and content development.

SecOps provides hundreds of prebuilt parsers out of the box.[4] On ingestion, data gets converted to Unified Data Model (UDM) as a standard format. We'll discuss UDM in more detail shortly; for now, suffice to say that this is a foundational element of SecOps' unified search and detection capabilities.

Pipeline and Data Format

Data collection in a SIEM is not a static function, but an evolving system to be monitored, managed, and audited regularly. To this end, SecOps has a dedicated ingestion metrics schema, covering every stage from log generation to ingestion. As we'll see in the Dashboard section, we can use this data to build detailed reports on the performance of our data pipeline and use this schema to export data collection and processing metrics to any external platforms we might use for infrastructure monitoring.

SecOps uses UDM, a standardized data structure, for log normalization. UDM's flexibility allows it to support complex analytics while maintaining high efficiency for data exploration and enrichment. UDM events are composed of multiple sections including metadata (header) that describes the event, source details, and log content.

Converting disparate log formats to a common data model makes it easier to identify relationships between entities including users, hosts, and IP addresses and to enrich source logs based on those relationships. UDM provides thousands of fields for describing and categorizing events, which makes it a great model for detection engineering and event correlation.

It's important to note that UDM is designed for investigations, not data analysis. UDM records describe *events*—actions occurring in the environment—and *entities*—elements such as assets, users, and resources. We'll see a few different examples of how this design choice factors into various SIEM functions like pivots and reporting.

Building and maintaining parsers are among the most arduous tasks in SIEM management. SecOps has recently introduced an autonomous parsing feature, currently in preview, that leverages machine learning and AI to significantly reduce the management overhead and continuous engineering effort typically associated with configuring and maintaining parsers for each individual source.

[3] "Flow and processing of customer security data to Google Security Operations," https://cloud.google.com/chronicle/docs/data-ingestion-flow

[4] https://cloud.google.com/chronicle/docs/ingestion/parser-list/supported-default-parsers

Threat Detection

Detection capabilities are front and center in SecOps. In this section, we'll discuss the role of Google Threat Intelligence in detection enrichment, SecOps' detection definition language, YARA-L, curated analytics, and AI enhancements in the detection and investigation functions.

Google Threat Intelligence

Investigations often start with the question, "What do we know about these potential threats?" Answering this question usually requires research performed outside of the SIEM, burning valuable analyst time at the critical early stages of an investigation. Google Threat Intelligence is built into SecOps, and it gives teams access to frontline intelligence from Mandiant incident response investigations, crowd-sourced data from VirusTotal, Safe Browsing data from billions of devices running Chrome, and curated open source intelligence. These integrations provide a more comprehensive view of the threat landscape and give analysts the context they need to act on alerts. The integration of this intelligence with a library of hundreds of curated detections, powerful search features, and risk-based analytics makes SecOps a powerful investigative tool. The next few sections will cover the latter capabilities in more detail.

Curated Detections and YARA-L

Having worked with many SIEM products and thousands of stock detection rules, I can honestly say that SecOps brings its "A game" when it comes to detection content. Curated detections, actively maintained by the Google Threat Intelligence Team, focus not only on indicators of compromise but also on more sophisticated tactics, techniques, and procedures. SecOps provides some great tools for managing this content and creating custom detections. Figure 2 shows the Curated Detections page, which provides an overview of active rules with associated MITRE ATT&CK™ coverage.

Figure 2. The Curated Detections Page shows analytics created and managed by the Google Threat Intelligence Team.

Experienced detection engineers know that some analytics contain more explicit criteria and are generally higher-confidence detections, while others may be more anomaly-based and require more work to validate. Curated detections are annotated accordingly in Figure 2 as “Precise” (P) or “Broad” (B).

Just as useful is the Curated Detection Dashboard (shown in Figure 3).

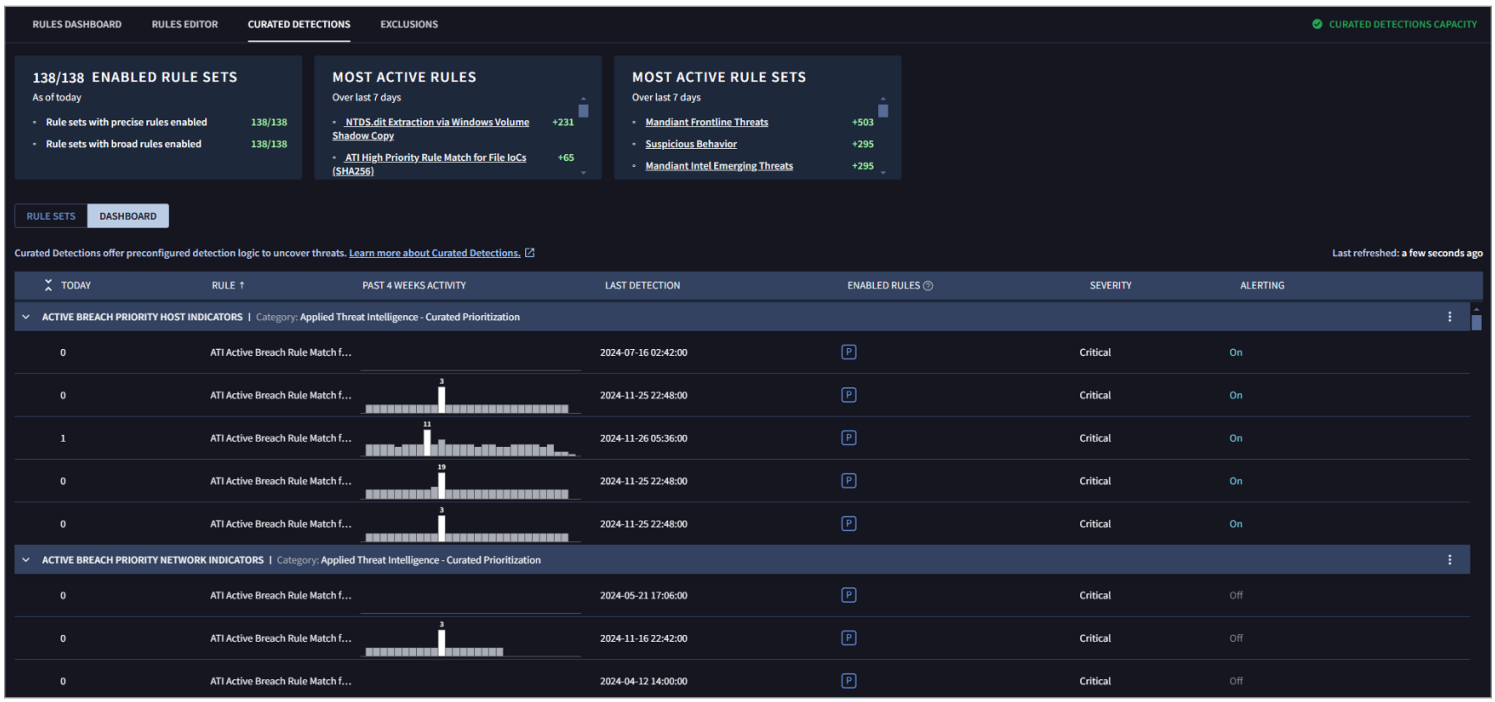

Figure 3. The Detections Dashboard provides a high-level view of detections useful for rule management and tuning.

This view shows information about curated detections that have triggered alerts based on log data in SecOps. You can see current activity for each ruleset, time of the last detection, rule status, and severity. From here, it’s an easy pivot into Rule Settings to modify or tune individual rules or all rules in a set. This is a nice change from products where the rule modification and tuning functions are decoupled from alert data, necessitating a back-and-forth process to identify where edits may be needed.

SecOps uses robust but flexible analytics based on its YARA-L language, which is derived from YARA (modified for logs, hence the “L”).[5] Like YARA, YARA-L is a plain text detection format consisting of a rule header and rule definition. As with YARA and other product-neutral detection languages like Sigma, this simplicity makes YARA-L rules easy to write and understand, and easy to annotate using detection frameworks and other useful metadata. YARA-L is also powerful, capable of supporting simple detections, like string matching, as well as complex logic for identifying scripts or programs.

[5] https://cloud.google.com/chronicle/docs/detection/yara-l-2-0-overview

Figure 4 shows the Rules Editor with an example YARA-L rule. Note that the rule metadata section includes the MITRE ATT&CK technique reference. This annotation feature is great if your team has standardized on ATT&CK or another standard framework for threat detection.

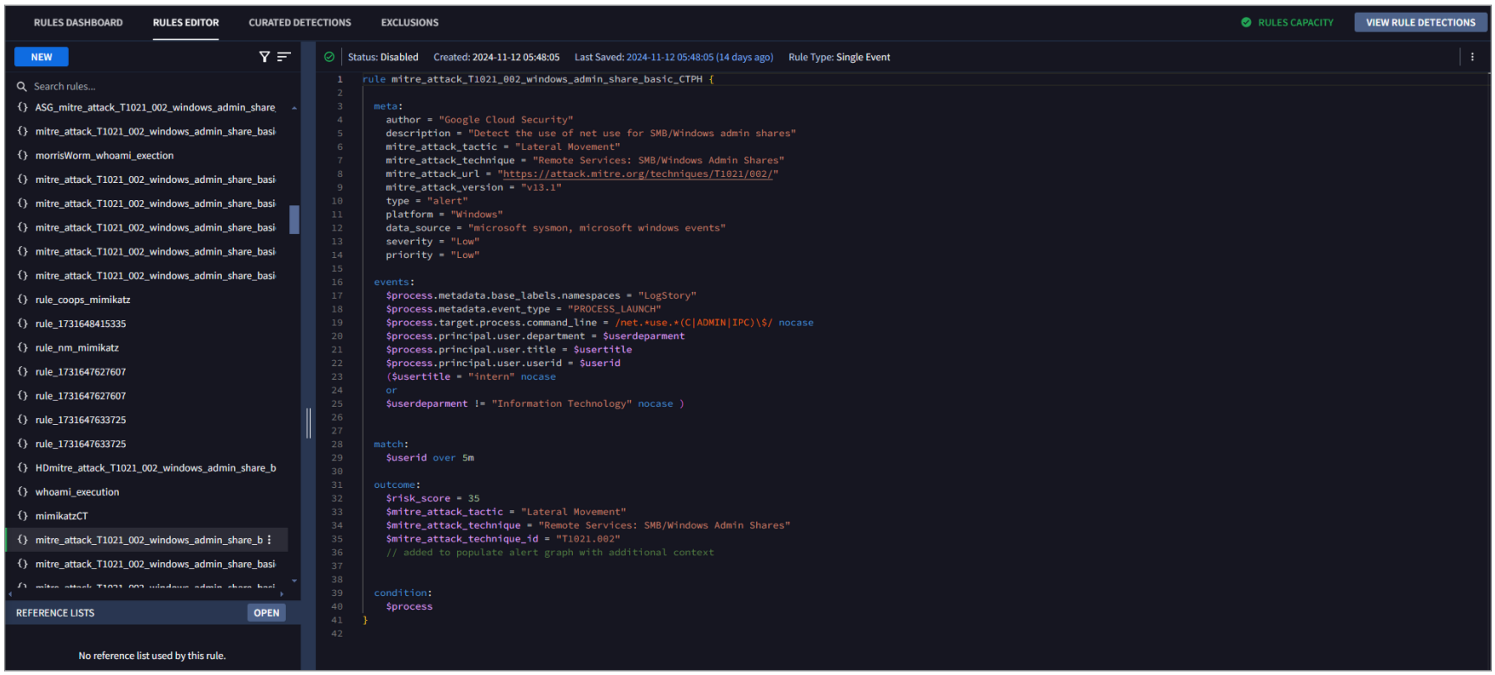

Figure 4. The Rules Editor Screen with YARA-L Rule for Lateral Movement via SMB/Windows Admin Shares

The simplicity of the Rule Editor and YARA-L and a huge library of curated detections makes for a great detection engineering experience that doesn't require starting from scratch or excessive screen-hopping. From here, you can launch a historical hunt based on the rule, view related detections and rule version history, or modify the frequency of rule execution.

Risk Analytics and UEBA

Of course, any detection strategy focused only on "known bads" is incomplete. The Risk Analytics capability in SecOps calculates risk scores for assets and users (aka entities) in your environment based on a set of predefined but customizable parameters. SecOps automatically correlates events based on these parameters and then assigns a risk score to help prioritize alerts. You can monitor UEBA with Risk Analytics for coverage for Authentication, Network Traffic Analysis, Peer Group Detections, Suspicious Actions, and Data Loss Prevention. It's also possible to use risk scores in YARA-L rules as an alert criterion. Severity, priority, scoring, and the risk calculation window are all customizable, enabling granular detection for behavioral anomalies among users or assets. These capabilities enable detection of anomalous activities which may be malicious—both the "unknown unknowns" and the "known unknowns." The Behavioral Analytics Dashboard, shown in Figure 5 on the next page, provides a unified view for all entities detected by SecOps.

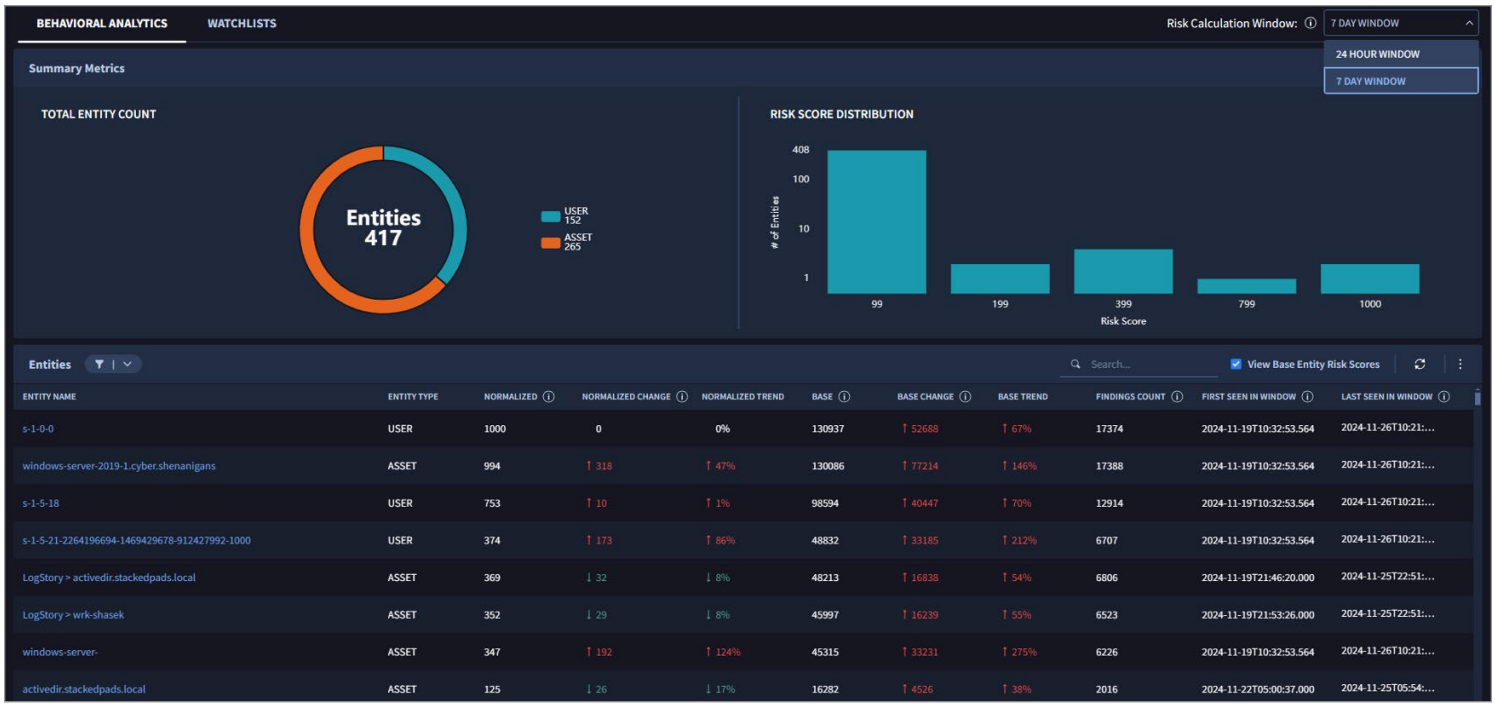

Figure 5. Entity Behavioral Analysis via SecOps Behavioral Analytics Dashboard

Unified Search Enhanced by AI

Unified search has always been the core promise of SIEM and a prerequisite for threat hunting and exploratory analysis. SecOps takes these functions to the next level with UDM and context-aware detections. The UDM model used by SecOps enables context-aware analytics where entities, relationships, vulnerabilities, and other objects can be viewed in a single detection. This makes it easier for analysts to perform searches and create detections that include contextual data, which supports faster and more accurate decisions about the disposition of an alert or the result of a search. Analysts also can implement dynamic risk scoring or responses based on detections by specifying output qualifiers. An example of this would be setting dynamic severity scores based on specific user or asset details.

Natural language search means that operators can simply ask questions using plain English in the SIEM search box and SecOps will convert those questions into UDM search syntax. An example of a simple query for suspicious port 22 traffic is shown in Figure 6. Power users also can still run native UDM searches if they choose.

Killer Feature Alert!

SecOps supports natural language search, allowing analysts to focus on investigations over search syntax. UDM search commands derived from the user inputs are displayed along with the results.

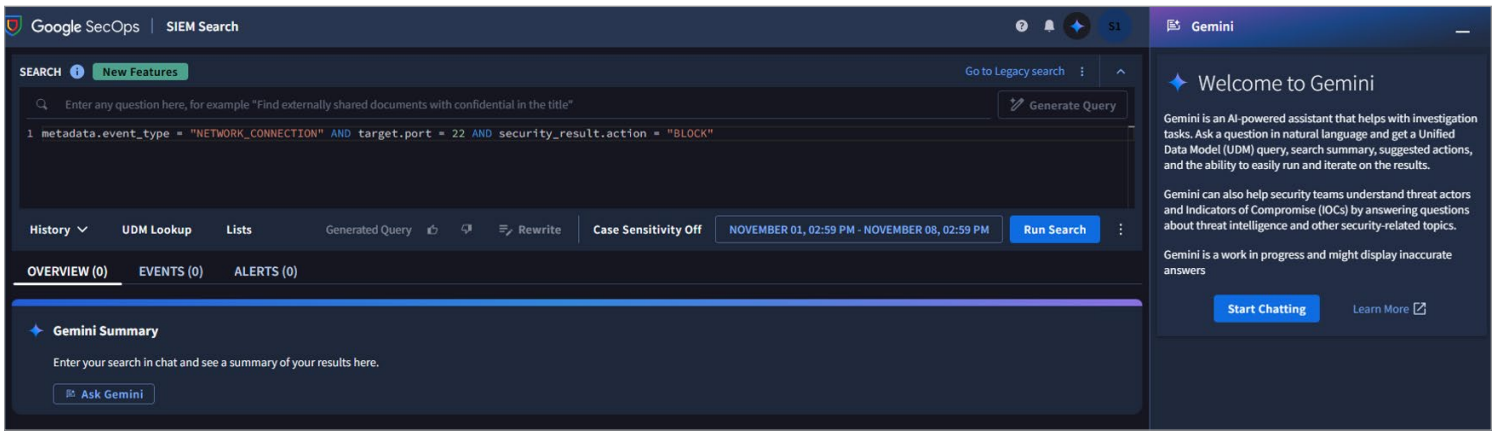

Figure 6. Example Natural Language Query with Associated UDM Search Syntax Provided by Gemini

Alerting

Detection is an area where Google's Gemini AI really shines. As an investigation assistant, Gemini can help answer questions, summarize events, hunt for threats, explain or even create detection rules, and recommend actions. Here are some examples of Gemini-augmented capabilities in SecOps:

- Given a query in plain English in the search box, Gemini uses natural language processing to interpret the query and construct a search in UDM syntax to get the answer.
- Asking for an explanation of a detection rule (see Figure 7) will result in an AI-generated summary of the rule and the implications of an alert.
- Ask Gemini about attacker activity associated with a vulnerability, and it will return tactics, techniques, and procedures aligned to the CVE based on Google Threat Intelligence.

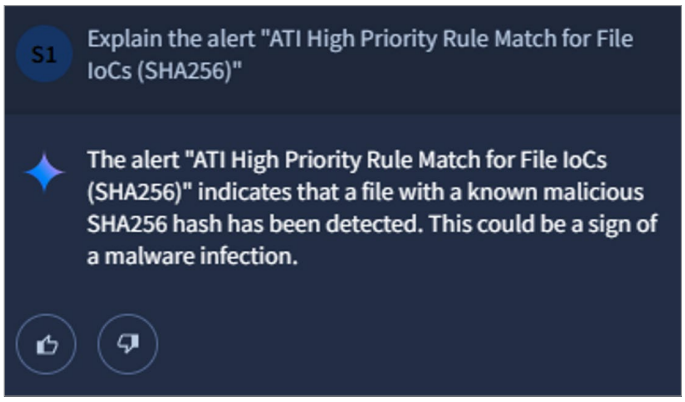

Figure 7. Gemini-Generated Alert Summary

As with any AI, the results provided in these examples may require some validation or iteration. But these AI-assisted search-and-detection features are fantastic examples of assistive technology that saves precious human cycles. They also let human analysts focus on what they do best—exercise creativity and decision making during the investigative process—and leave menial tasks to the automated assistant. We'll see more examples of this assistive automation in the following sections. From an operator's perspective, SIEMs tend to fall into one of two categories: the "blank canvas," where it is up to the user to decide how best to utilize the data, and the content-heavy platform that attempts to give you the answers but may be less customizable. SecOps is a pleasant compromise that provides plenty of content out of the box and flexible custom detections.

The Alert screen (see Figure 8) presents alerts and their associated metadata, with prebuilt filters and a search bar available in the left-hand Filter pane. This view is highly customizable, allowing analysts to dynamically sort, filter, and group alerts into Cases.

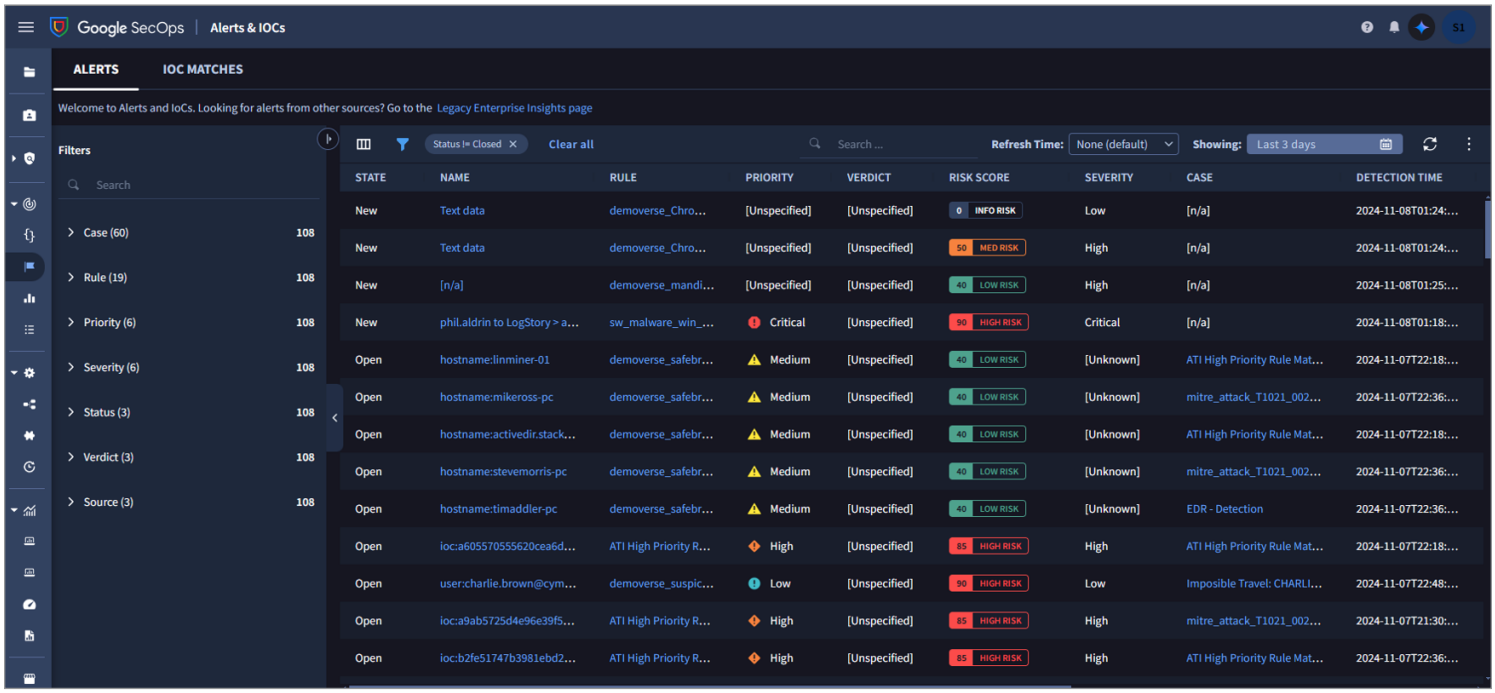

Figure 8. The SecOps Alert Queue with View Filtering Pane

Drilling into each alert provides a wealth of detail and contextual information, including event details, rule information, alert history, and a visual summary presented in link analysis diagram form. This detailed view also provides quick pivots to Cases, which we'll cover in more detail in the next section.

Investigations

Investigation is an area where many SOC teams struggle to maintain speed, quality, and completeness. Experienced analysts are highly susceptible to alert fatigue and bias, and less experienced analysts sometimes struggle to identify the best investigative path. SecOps provides powerful investigative tools for skilled investigators as well as assistive features for less experienced users.

The Cases tab is where most of the investigative workflow happens. The list view in the Cases tab (Figure 9) shows open cases, case priority, workflow stage, number of associated alerts, and other key details. The organization features in this view allow analysts to focus on specific criteria by sorting, filtering, or utilizing the tagging function (shown on the right).

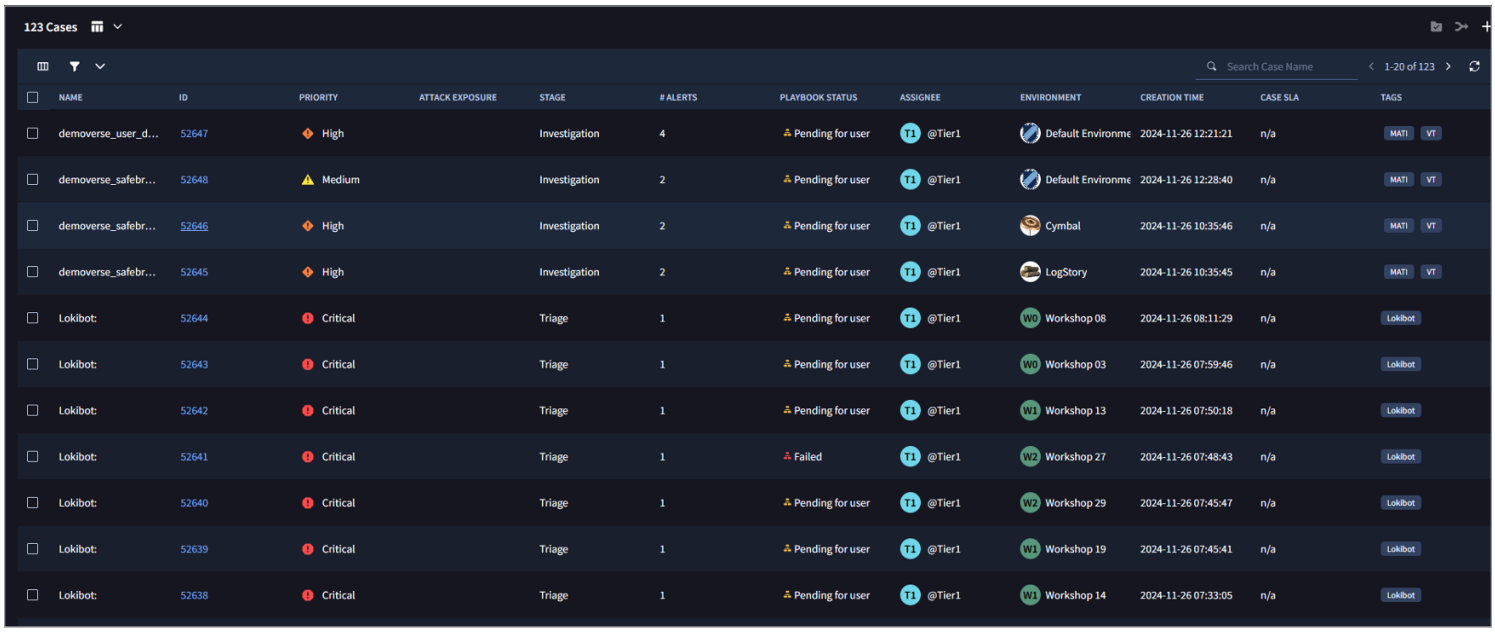

Figure 9. Cases Tab (List View) Showing High-Level Case Details

The Explore button in the Cases detail view (shown in Figure 10) offers link analysis-style visualizations of entities tied to the case, providing a visual representation of relationships and connections. In some tools, these graphical entity diagrams feel clunky and static—more of an ornamental feature for the user interface than a real investigative tool. But in SecOps, the entity diagram (shown in Figure 11 on the next page) is interactive and provides useful timeline features and entity metadata for further analysis.

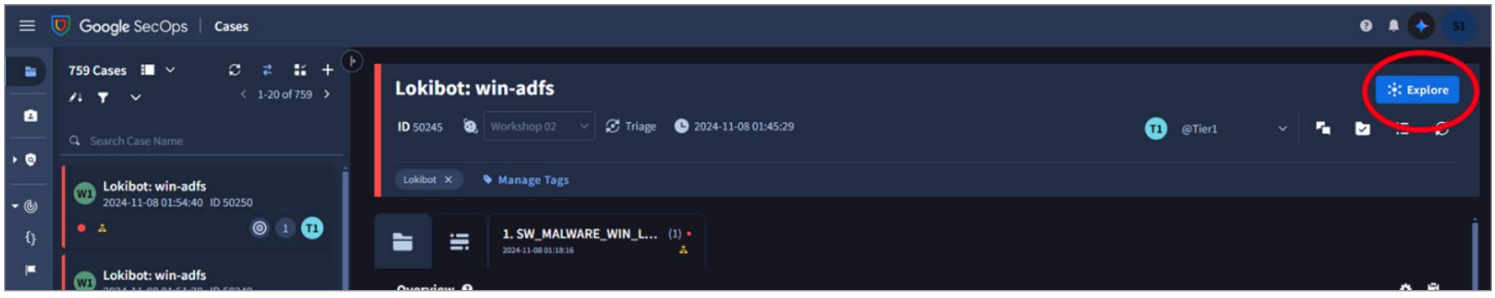

Figure 10. Explore Button in the Cases Detail View

Assistive Automation and World-Class Intelligence

The Overview pane in the Cases view, as shown in Figure 12, may be the best UI I have seen in 20 years of SIEM experience. Using Gemini, SecOps generates a case summary that includes key details to be investigated, a concise abstract of events in question, and recommended next steps.

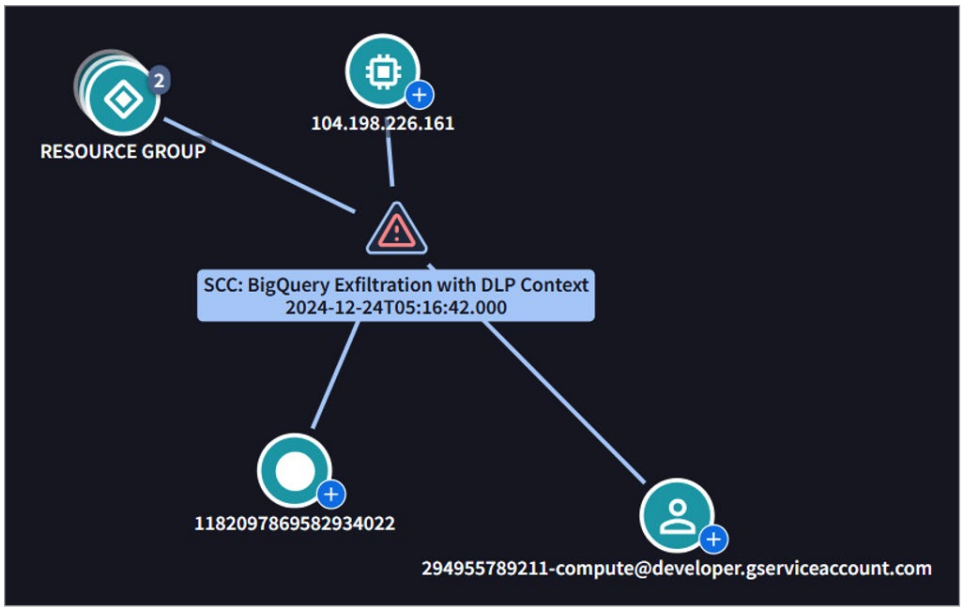

Figure 11. Entity-Diagram View of a Data Exfiltration Alert

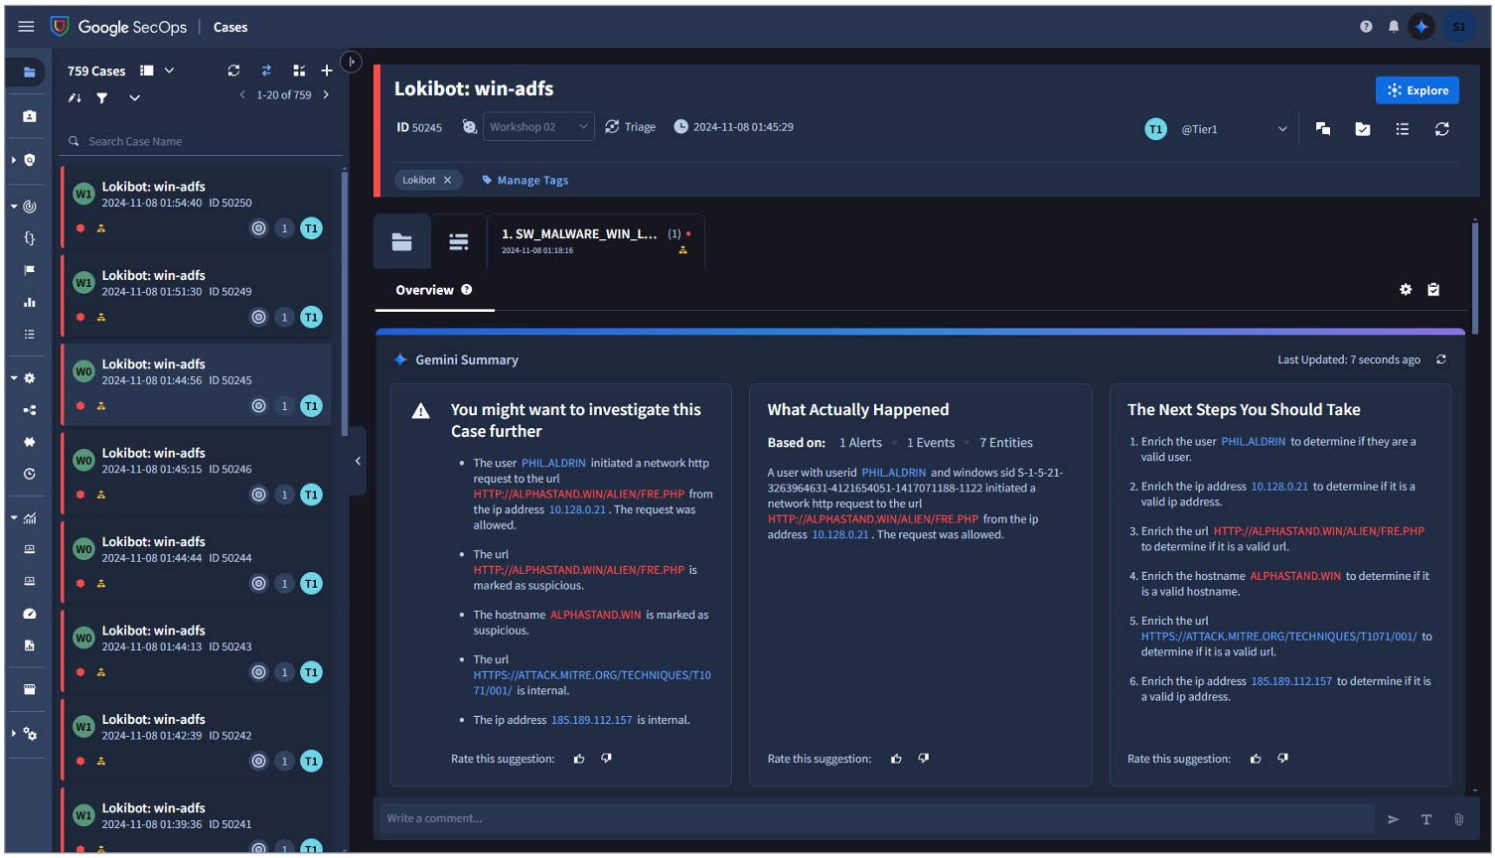

Figure 12. Cases View with Summary Provided by Gemini

Best of all, these sub-panes are titled: “You might want to investigate this Case further,” “What Actually Happened,” and “The Next Steps You Should Take.” The lack of operational jargon here is truly refreshing and invaluable to users who may be new to security investigations. Scrolling down below the Overview reveals pending actions, alert details, the case entity diagram, and insights generated by VirusTotal and Mandiant through Google Threat Intelligence. Figure 13 shows contextual detail for a specific event provided by VirusTotal and Mandiant Threat Intelligence.

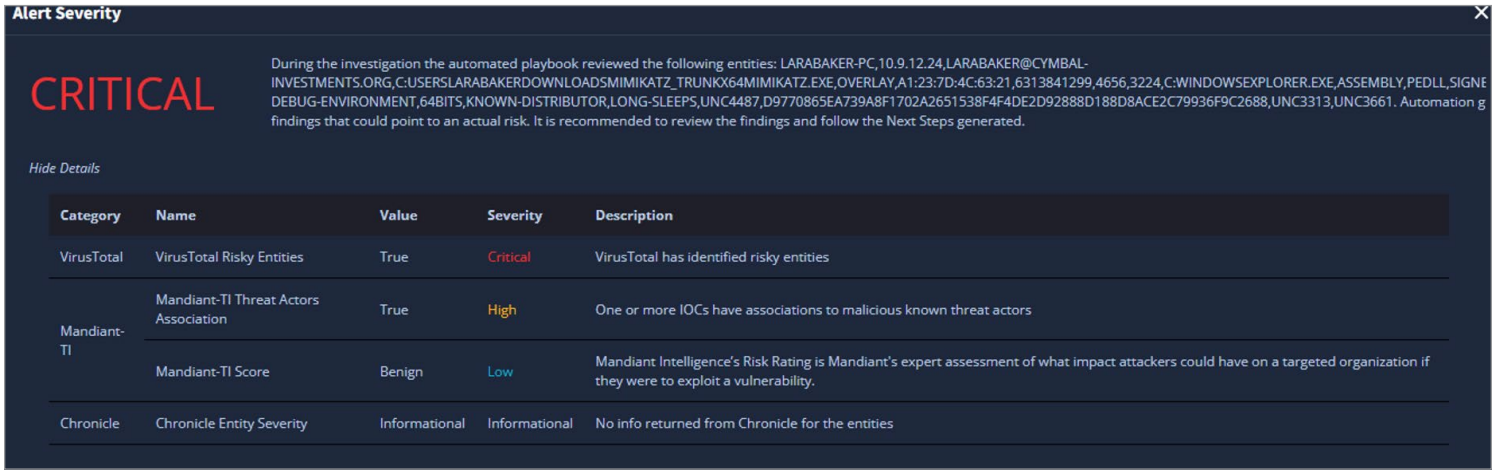

Figure 13. Context Detail Provided by VirusTotal and Mandiant Threat Intelligence

Other features of the Cases view are the Case Wall, where investigative updates and details are tracked, and the Playbook panel. Playbook flow and results can be viewed here, or users may initiate discrete automated Actions facilitated by various integrations, for example, isolating a host via CrowdStrike Falcon. We will cover Playbooks and Actions in more detail in the next section.

The Cases view reinforces threat-informed defense as a core tenet of the SecOps approach, providing contextual information most analysts are used to gathering from multiple tools and services in a single view. It also offers collaborative features, like a chat for annotation and discussion, and a Workdesk view for individually assigned Cases, Actions, Tasks, and Requests. The Cases interface, contextual views, and automation enable analysts to spend more time on the investigation itself than on data acquisition and manipulation.

Real Decision Support

An AI-generated Case Overview describes current events, explains why they’re important, and suggests next steps in plain English. This context helps analysts reduce triage time and move more deliberately through investigations.

Playbooks

Built on top of the Siemplify product acquired by Google in 2022, SecOps’ SOAR capabilities are flexible and highly mature. Playbooks are organized into a hierarchy that includes triggers, actions, and flows. The Playbook Designer is a graphical development tool that works like most other SOAR platforms. Figure 14 shows an example playbook, where the flow of control moves from left to right starting with a trigger (represented as a yellow box). Operation is guided by a series of actions (blue boxes) and, optionally, conditions (purple box) for if-then-else branches.

Figure 14. Visual Representation of a SecOps SOAR Playbook

Actions are the core building blocks of automated playbooks and rely on integrations that can be downloaded from the Google Security Operations Marketplace. When the playbook executes, each action returns:

- JSON-formatted output messages, tables, attachments, and/or links
- Script result, which will be appended to the Case Wall in the Cases view or displayed in the right-side panel of the Cases screen.

Google has extended SecOps' SOAR capabilities by automating the playbook development process with Gemini. Using the Playbook Assistant, users have the option of entering natural language prompts describing the desired scenario, steps, enrichment data, and other relevant details allowing Gemini to create or modify a playbook.

Based on the 2024 State of Security Automation Survey conducted by SANS, time for process improvements is a top concern for security teams automating their processes.[6] The AI-assisted playbook capabilities in SecOps are a big step in addressing that challenge. That said, there are some limitations to this feature. For example, Gemini cannot create custom integrations or utilize sophisticated parallel actions. The quality of the playbook is also dependent on the details included in the prompt, as is the case with any prompt-driven AI solution. But even if the resulting playbooks are not always production ready, this is a huge time saver for busy teams who simply need a head start on playbook development.

Dashboards

The decision to build SecOps' dashboarding capabilities on an existing business intelligence tool is noteworthy, as so many security tools seem to include dashboards only as an afterthought (though Google will soon implement a new Native Dashboards feature to replace the current Looker-based dashboards). Dashboards provide visualizations of various data points, including UDM events, entity data, detections, and ingestion metrics, offering a comprehensive overview of the security landscape.

SecOps' robust, easy-to-use dashboarding capabilities are built on top of Google's Looker business intelligence tool.

Google SecOps SIEM provides the following default dashboards:

- Main, providing a snapshot of data ingestion details (see Figure 15 on the next page).
- Preview, which is effectively a dashboard workspace
- Cloud Detection and Response, providing insights into the security posture of your cloud environment
- Context Aware Detections – Risk, displaying the threat status of entities identified in SIEM data (users and assets)
- Data Ingestion and Health, reporting the type, volume, and performance of data being ingested by SecOps
- IOC Matches, showing IOCs detected in the environment
- Rule Detections, providing insights into alerts triggered by detection rules
- User Sign In Overview, providing insights into user activity

[6] "The State of Automation in Security Operations: A SANS Survey," June 2024, www.sans.org/white-papers/state-automation-security-operations-sans-survey

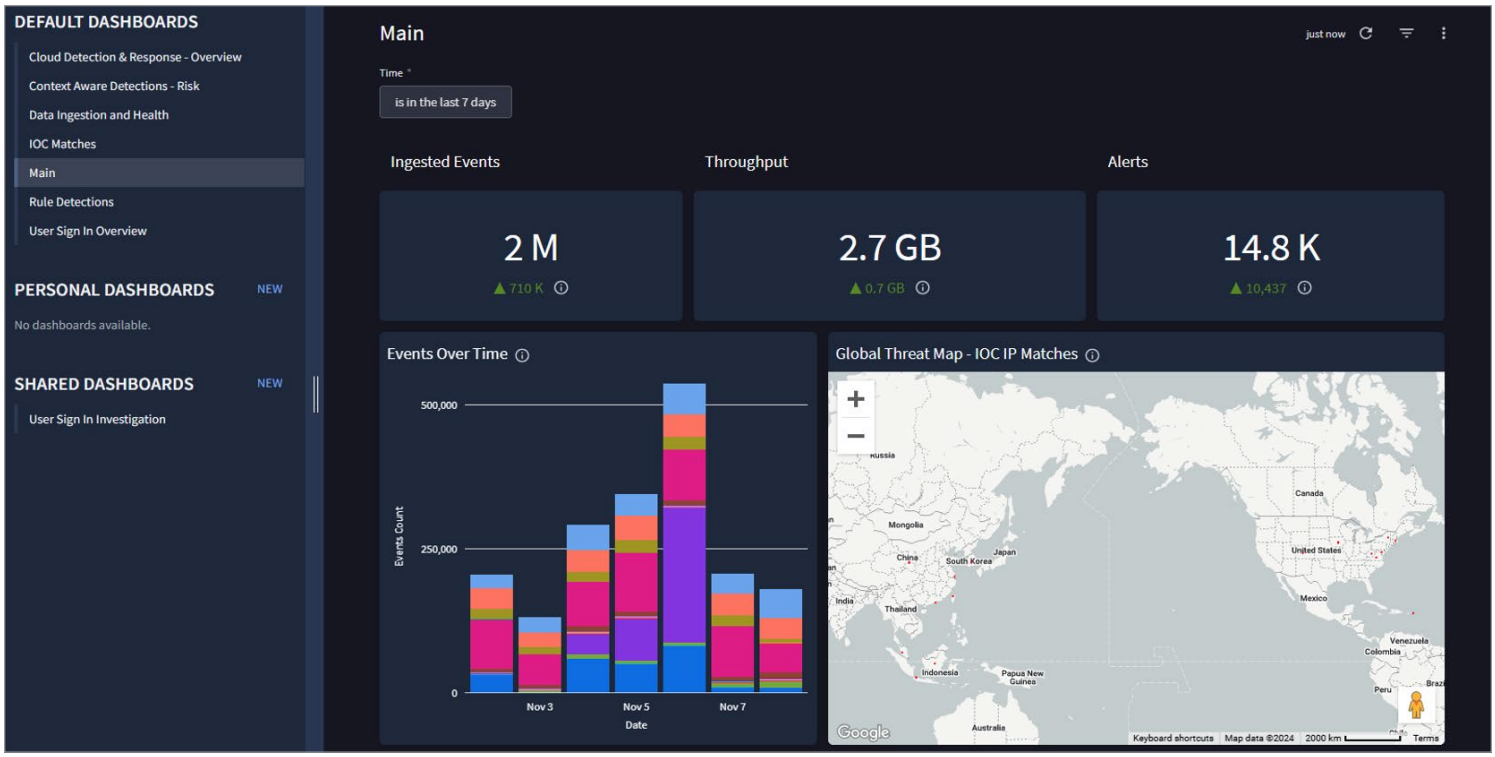

Figure 15. Example Data Ingestion Dashboard Using SecOps' Looker-based Dashboard Features

SOAR dashboards, an example of which is shown in Figure 16, provide an overview of case management tasks, alerts, and response actions. They also feature some nice reporting on analyst workload, case tags, detections, and other performance information.

Building dashboards in SecOps is a straightforward "tile"-based process in which the user can clone existing tiles or create them from scratch based on new queries or reports. As indicated in Figure 16, we can select from a variety of different visualizations for the data represented.

Although the SecOps navigation menu still distinguishes between SIEM and SOAR reporting, I'm told this function will be consolidated in a future release.

Figure 16. Example SOAR Dashboard with Case and Alert Statistics and Other Workload Details

Conclusion

Overall, Google SecOps presents a compelling solution for organizations seeking a modern SIEM platform. Data ingestion and search capabilities are exactly what you would expect from a Google product, contextual detections and streamlined response features are excellent, and its flexible reporting functions are competitive with any dashboarding or business intelligence platform. The assistive automation provided by Gemini is still in the early stages, so results and recommendations require careful validation by human operators. These features are major time savers, though, for users who are often at or over capacity. And finally, the inclusion of Google Threat Intelligence from VirusTotal and Mandiant, browsing data from billions of Chrome devices, and curated detections bring together several best-in-class threat intelligence and research capabilities. A compelling vision for SIEM's *Third Act*, Google SecOps empowers security teams to effectively address today's ever-evolving threat landscape.

Sponsor

SANS would like to thank this paper's sponsor:

Google Cloud Security