Omdia by informa techtarget

OCTOBER 2026

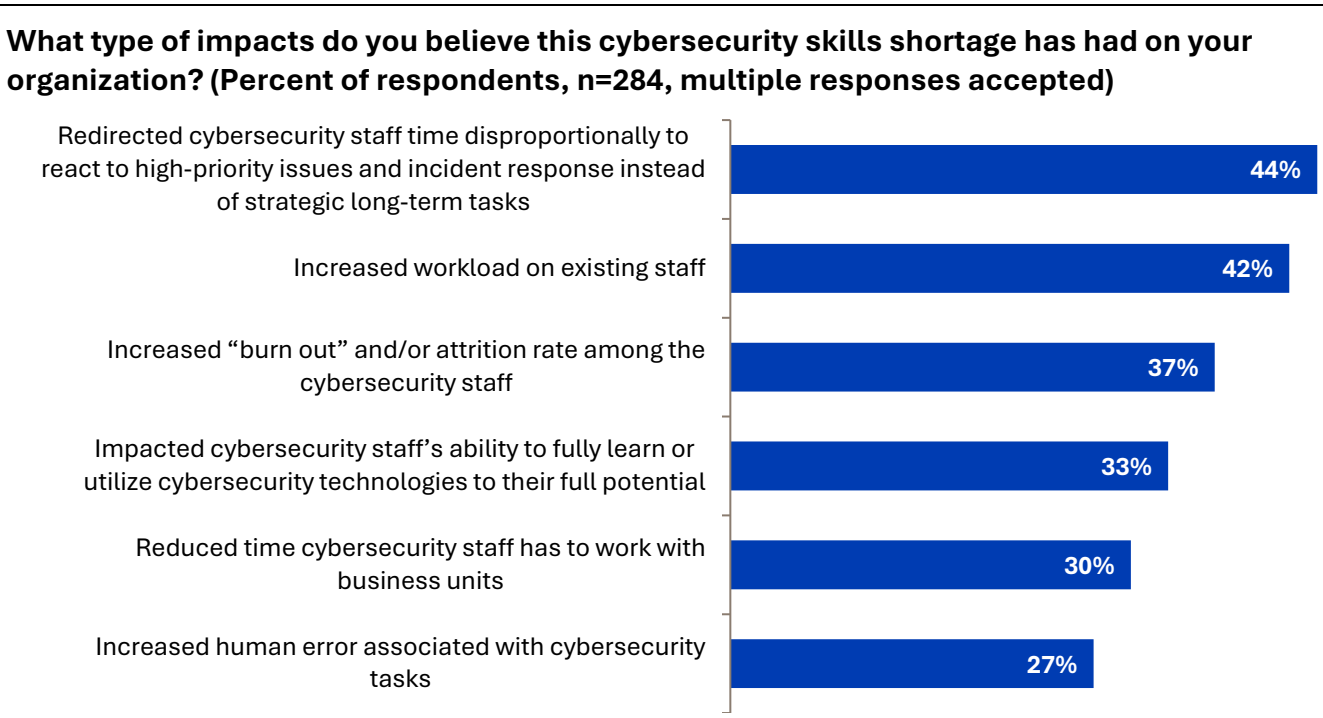
Agentic Defense with Google SecOps

Tony Palmer, Chief Analyst

Cybersecurity challenges

Omdia research reveals that the vast majority (75%) of organizations have been impacted to some extent by the cybersecurity skills shortage. Respondents reported significant impacts to their security programs from this skills shortage, including redirected time (44%), increased workload (42%), increased "burn out" (37%), and decreased time for cybersecurity staff to utilize their tools to their full potential (33).[1]

Figure 1. Top 6 impacts of the cybersecurity skills shortage



Source: Omdia

More than two-thirds (68%) of respondents indicated that being a cybersecurity professional has become more difficult over the past two years. Those who said working as a cybersecurity professional is more difficult today than two years ago cited several challenges underlying this belief. More than half (55%) cited the increase in cybersecurity complexity and workloads, and 52% cited cyberthreats and attack surface growth.

[1] Source: Omdia Research Report, *The Life and Times of Cybersecurity Professionals, Volume VIII*, May 2026. All Omdia research references and charts in this Technical First Look are from this report.

This Technical First Look from Omdia was commissioned by Google and is distributed under license from TechTarget, Inc.

© 2026 TechTarget, Inc. All Rights Reserved.

Google Security Operations

Google Security Operations (SecOps) is an AI- and threat intelligence-driven security operations platform that empowers security teams to better defend against modern threats. Google SecOps combines Google's hyperscale infrastructure, Gemini's AI capabilities, and Google Threat Intelligence (TI) to serve as the workbench for SOC teams tasked with detecting, investigating, and responding to cyberthreats across their hybrid environment.

Google SecOps is designed with three key benefits in mind:

- **Insights at speed and scale**. Built on Google's core infrastructure, Google SecOps scales to ingest, normalize, and retain massive amounts of data, eliminating the legacy tradeoff between cost, scale, and data coverage. Google SecOps's context-rich data model and fast search quickly surface the insights security analysts need for effective investigation and response.
- **Intelligence-driven outcomes**. Google SecOps enriches data at ingestion with insights from Google TI. This provides a deep understanding of threat actors to help organizations expose threats without the requirement for intensive custom engineering by the customer.
- **AI-infused productivity**. With the ongoing cybersecurity skills shortage, organizations (including MSSPs) are challenged to hire and retain enough skilled security experts. Google SecOps leverages both AI and automation to elevate the threat detection and response skills of security teams and make them more productive.

It's important to note that Google SecOps is a hybrid, multicloud solution. It can ingest data from cloud and on-premises sources, and actuate response actions on any cloud—AWS, Azure, Oracle, and others.

First Look

This Technical First Look is focused on three powerful AI capabilities offered by Google SecOps: The Threat Hunt agent, the Detection Engineering agent, and the Parser Extension agent.

The Google SecOps Threat Hunt agent

The Threat Hunt agent is an autonomous AI-driven capability powered by Gemini, Google TI, Mandiant Frontline analyst insights, and the MITRE ATT&CK framework embedded directly within Google SecOps. It is designed to transform threat hunting from a time-consuming, manual process into a scalable, proactive, automated force multiplier, creating in-depth threat hunting plans based on specific threat actors, campaigns, malware, software toolkits, or tactics, techniques, and procedures (TTPs).

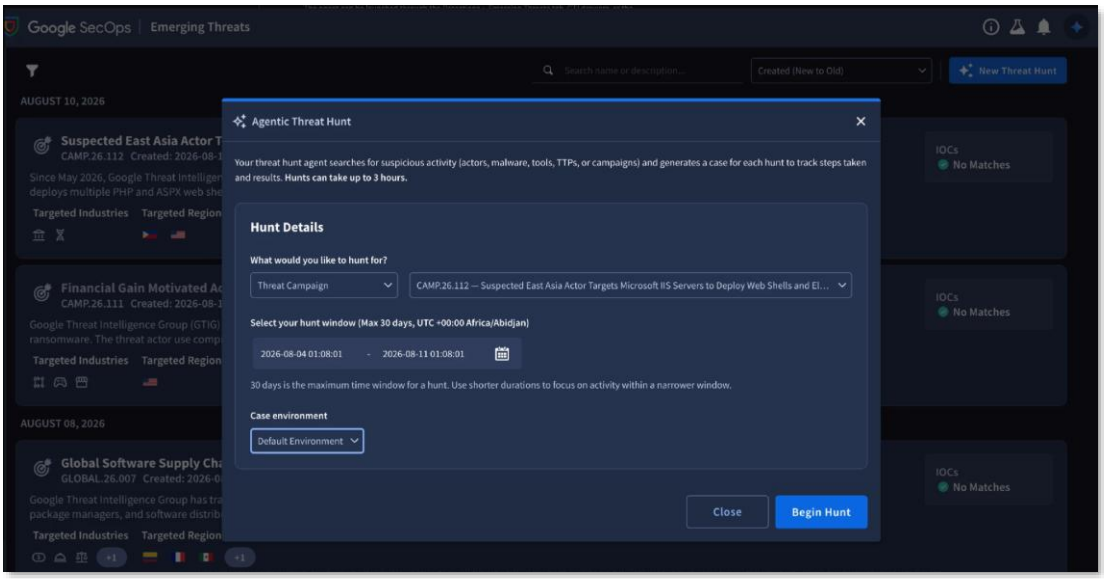
Omdia examined the Threat Hunt agent by logging into Google SecOps and opening Emerging Threats. We clicked on the *New Threat Hunt* widget at the top and walked through the selection of criteria for the hunt.

Analysts first set the entity the agent will hunt for. Entities can be selected from lists of threat actors, campaigns, specific malware families, software toolkits, or MITRE TTPs. The agent can search within a window of up to 30 days, selected from the entire corpus of data available.

Once parameters were set, we kicked off a hunt by clicking *Begin Hunt*.

©2026 TechTarget, Inc. All Rights Reserved.

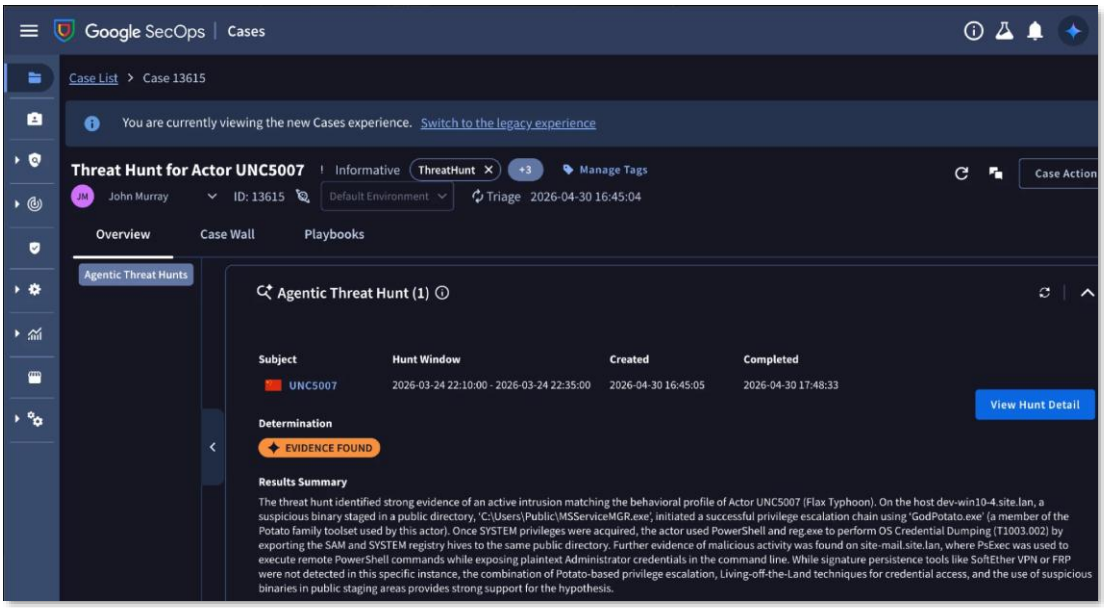
Figure 2. Launching the Threat Hunt agent



Source: Google and Omdia

The agent generates an in-depth threat hunting plan—for its own internal use at this point—using data from Google TI and MITRE. The agent then builds out a natural language hypothesis describing how to look for this threat, translating that into a query that it executes.

Figure 3. Examining the tracking case

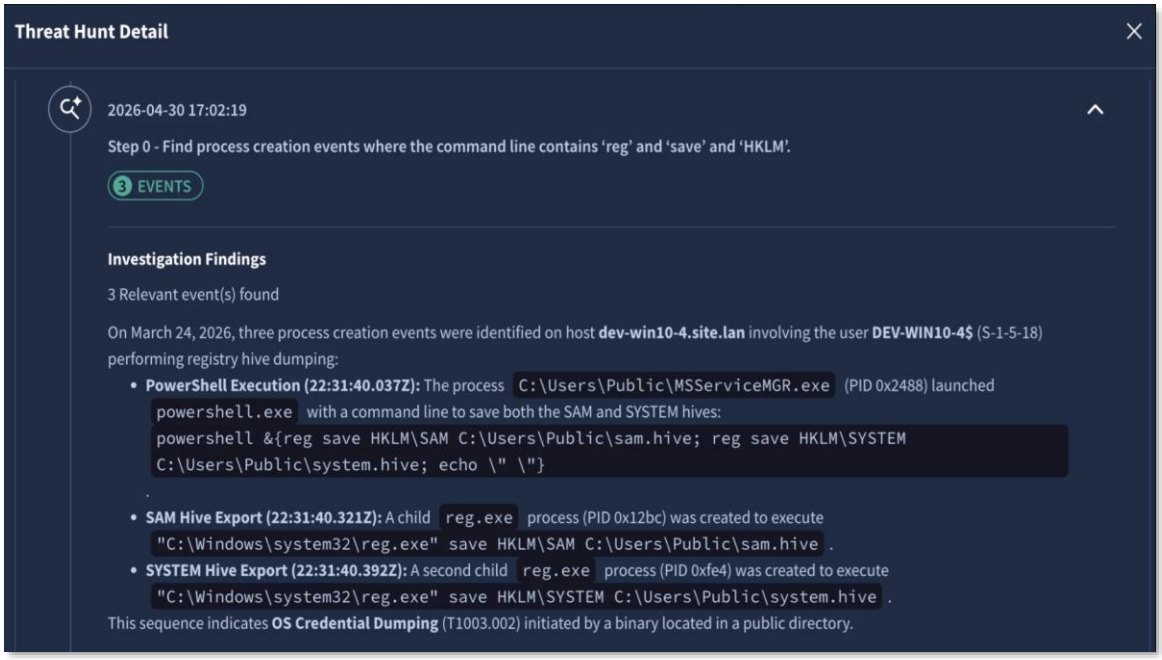


Source: Google and Omdia

©2026 TechTarget, Inc. All Rights Reserved.

Follow-up queries are often generated based on initial findings (e.g., to investigate child processes or user logins). The agent's goal is to determine if there is evidence of a threat within the environment, and it provides a summary of its conclusions, with the ability to view all details and findings (see **Figure 4**).

Figure 4. Threat hunt detail



Source: Google and Omdia

Google shared with us that customers previewing the agent in production environments are executing complete, successful hunts in less than 30 minutes, on average. To give that some context, in Omdia's experience, the typical manual threat hunting process is often performed by a small, specialized team or by senior analysts using spare cycles or during reactive incidents. This process, which involves researching vulnerabilities and indicators of compromise (IOCs), translating this information into queries for the security information and event management (SIEM) system, executing those queries, and then performing triage on the results and documenting everything, can take approximately 8 to 24 person-hours per hunt.

It's important to note that threat hunting is of necessity most often performed by the most experienced, most expensive, and most time limited employees within the SOC. In short, Omdia found that what used to take specialized teams or senior analysts up to a full day to complete can now be executed autonomously by the Threat Hunt agent in less than an hour.

The Google SecOps Detection Engineering agent

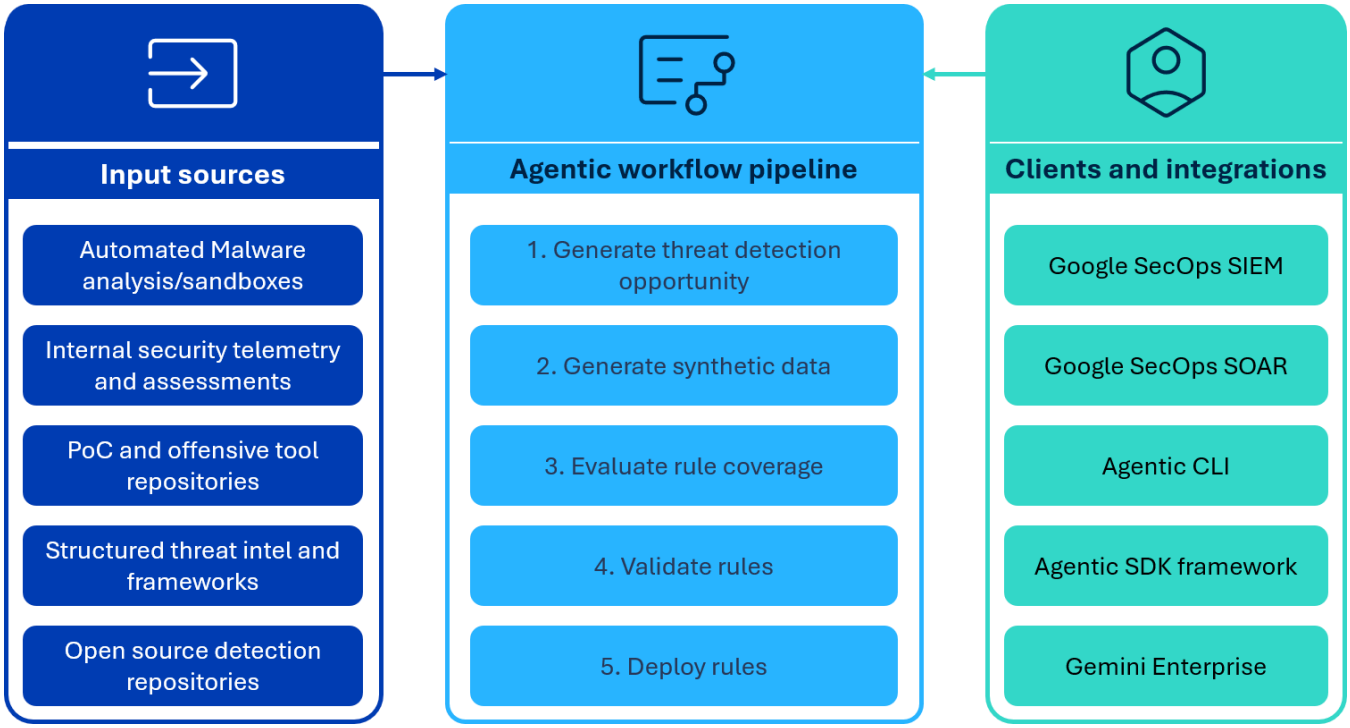
Detection engineering is often constrained by the operational compromise of fidelity versus speed. Addressing an active zero-day attack by rushing a rule into production infuses unnecessary risk into an already-dangerous situation, potentially overwhelming analysts with false positives, skewing risk scores, and

©2026 TechTarget, Inc. All Rights Reserved.

breaking dashboards. Spending days or weeks carefully testing and tuning queries in staging labs leaves enterprises exposed during the adversary's most lucrative exploitation window.

The Google SecOps Detection Engineering agent is designed to give security teams the ability to safely simulate threat coverage on demand. The agent can generate realistic attack sequences using synthetic data and run them against an organization's detection capabilities, enabling them to simulate a wide breadth of coverage and verify rule execution in production-safe environments, confirming the strength of defenses proactively—before an actual attack occurs.

Figure 5. Agentic detection engineering workflow

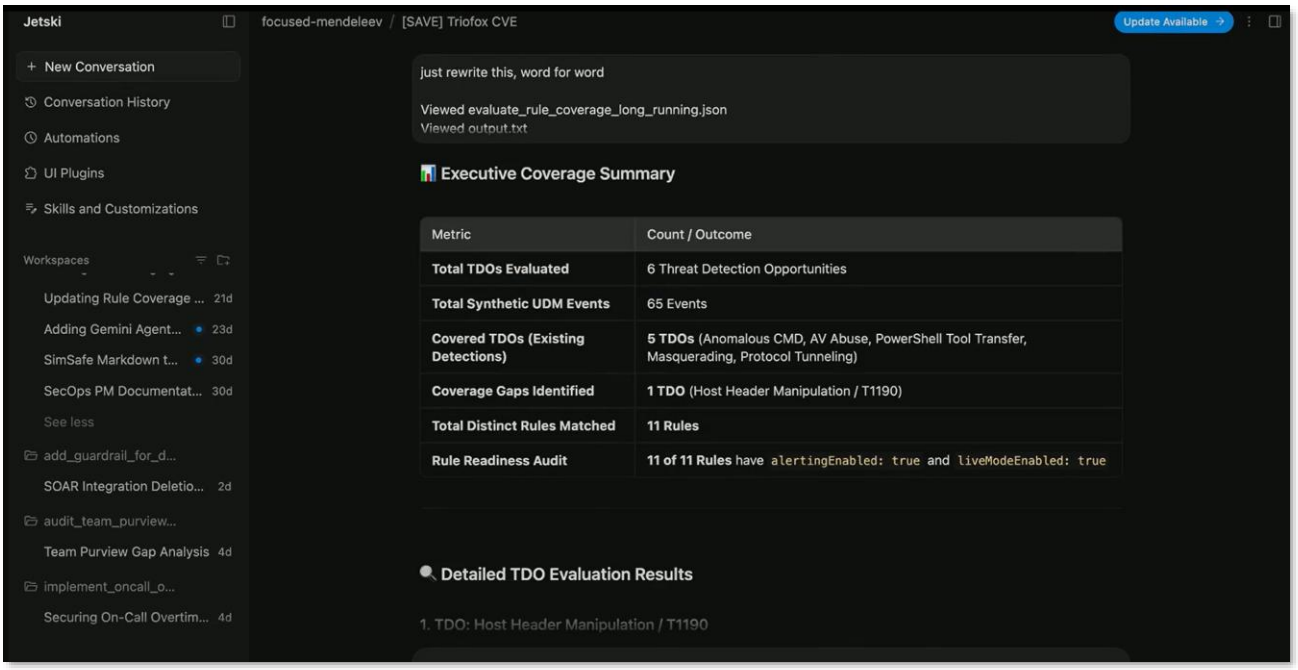


Source: Google and Omdia

We began our exploration of the Detection Engineering agent by running a "detection engineering coverage evaluation skill" pulled directly from a GitHub repository. We then added a vulnerability blog sourced from Google TI to preload information into the system. We included screenshots, which is a best practice because more detail improves the detection engineering agent's performance. When the process was completed, the agent had evaluated six threat detection opportunities (TDOs) and generated 65 synthetic events that looked like that attack. Of those six TDOs, five were covered by 11 distinct rules already in place, and we were able to clearly identify the TDO that wasn't covered so the gap could be closed.

©2026 TechTarget, Inc. All Rights Reserved.

Figure 6. Production-safe event simulation



Source: Google and Omdia

It's important to note that synthetic events should not be commingled with production security telemetry. Google SecOps handles this by allowing detections to include a label for simulated data on UDM search and dashboards, with options to include or exclude this data.

In practice, Google customers have already integrated the agent into their CI/CD pipelines, enabling them to take a Jira ticket for a new threat, run the agent, have a human perform a final review of the generated rules, and then push them to production.

Omdia verified that the Detection Engineering agent validates coverage while reducing the risks of testing in production. It can identify coverage gaps and create new detection rules for threat scenarios, reducing time and effort needed, reducing the skill barrier by codifying complex threat logic, and enabling analysts of all skill levels to author high-performance rules.

In Omdia's practical experience, manual detection engineering is a time-consuming process, even for the most senior and expert analysts in the SOC. Reverse engineering an attack, creating and testing detection rules, then deploying them in production typically takes up to a week, sometimes longer. We observed the Detection Engineering agent in action and watched it complete the entire workflow in under 30 minutes.

Parser Extension agent in Google SecOps

Contemporary enterprise environments process substantial and continuously evolving telemetry from hundreds of disparate log sources, which often include custom attributes, non-standard fields, and

©2026 TechTarget, Inc. All Rights Reserved.

frequently updated vendor schemas. The volume and heterogeneity of this data present ongoing challenges for detection engineers and SOC analysts, particularly around log parser maintenance.

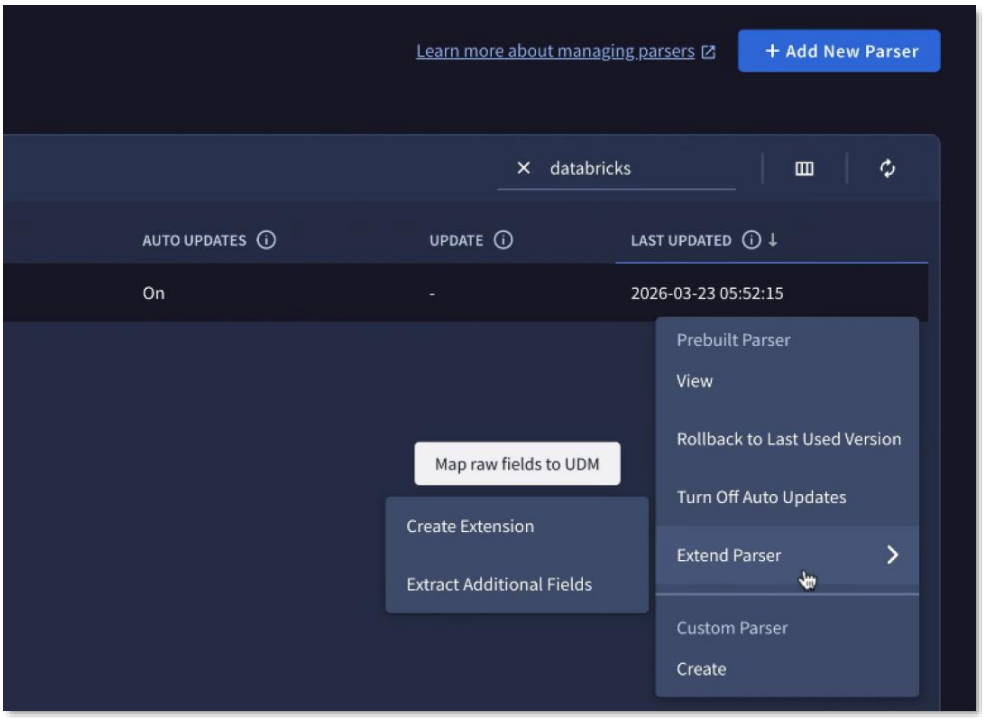
Traditionally, extracting unmapped fields or modifying existing parsers has required security engineers to develop complex regular expressions or parser code manually. This approach has been associated with technical complexity, extended onboarding timelines, and operational overhead for detection engineers and threat responders requiring timely access to newly available data.

Google currently provides over 1,300 pre-built, out-of-the-box parsers for various log sources. Parser extensions enable customers to add custom parser code as an overlay to existing parsers, avoiding the need to fork entire parsers, some of which can contain thousands of lines.

Google developed the Parser Extension agent as a response to these ingestion challenges. Leveraging Google Gemini, the Parser Extension agent is designed to streamline log normalization by enabling security analysts and detection engineers to specify parsing and mapping requirements using natural language inputs. The system then generates, tests, and validates parser extensions, reducing the time and technical expertise traditionally required for this process.

To see the Parser Extension agent in action, we accessed a test environment provided by Google. In Google SecOps, we navigated to *Settings* then *Parser Management*, located a target log type (in this example, Databricks), selected *Extend Parser,* and then chose *Create Extension* (see **Figure 7**).

Figure 7. Launching the AI parser extension wizard and providing context



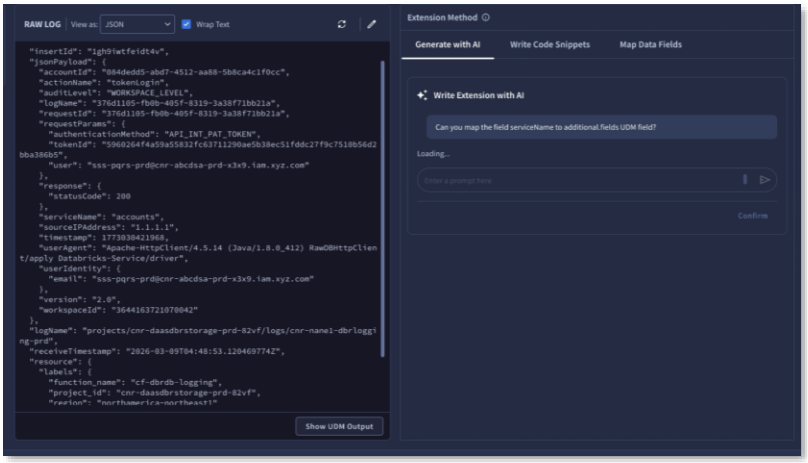
Source: Google and Omdia

©2026 TechTarget, Inc. All Rights Reserved.

The interface loaded the raw log sample in the middle and the *Write Extension with AI* tab on the right (see **Figure 8**).

We entered the prompt: *"Can you map the field serviceName to additional.fields UDM field?"* and Gemini generated the corresponding JavaScript parser extension. Analysts can refine the extension interactively (e.g., to add conditional action normalization) with subsequent additional prompts without overwriting earlier logic.

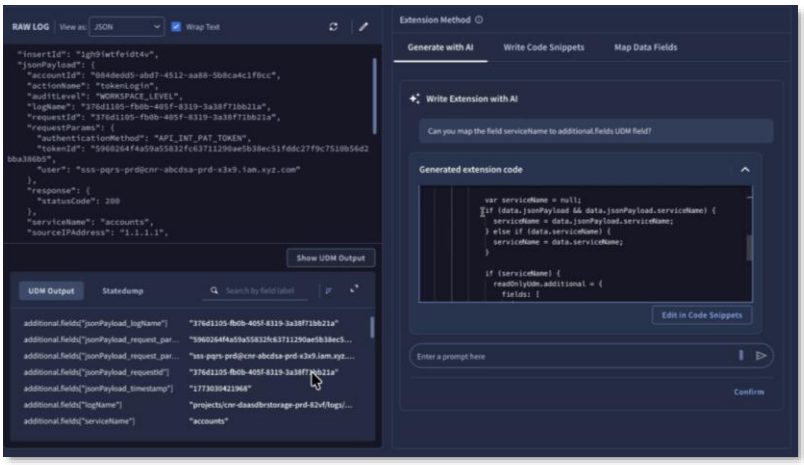
Figure 8. Conversational prompting and JavaScript code generation



Source: Google and Omdia

Before deploying the extension to live ingestion pipelines, analysts can click *Validate* to review the live UDM output (see **Figure 9**). The side-by-side view displays the raw input log on the left and the normalized UDM event on the right, highlighting newly extracted fields to confirm normalization accuracy.

Figure 9. Validating the new extension



Source: Google and Omdia

©2026 TechTarget, Inc. All Rights Reserved.

When investigating active alerts or search queries, analysts often need to identify unmapped raw log fields. Rather than navigating to *Parser Management*, users can click *Refine Parser with AI* directly from the raw log viewer in the SIEM Search screen in Google SecOps. SecOps loads the current log sample and unmapped fields into the extension builder for immediate remediation.

Transforming log normalization from a specialized coding exercise into an intuitive, conversational workflow enables security analysts and detection engineers to describe their parsing and mapping intent in plain English, allowing Gemini to generate, test, and validate production-ready parser extensions in minutes.

Conclusion

Omdia research indicates that 75% of organizations have been impacted by the cybersecurity skills shortage, resulting in redirected time, increased workloads, burnout, and reduced tool utilization. Additionally, 68% of cybersecurity professionals reported that their roles have become more difficult over the past two years, primarily due to increased complexity, expanded workloads, growing cyberthreats, and attack surface expansion.

Omdia validated that the Threat Hunt agent autonomously generates and executes comprehensive hunting plans based on threat actors, campaigns, malware, or TTPs, transforming threat hunting into a scalable, automated function accessible to broader security teams, directly mitigating skills shortage impacts. The Detection Engineering agent resolves the operational tension between detection speed and fidelity by enabling security teams to safely simulate threat coverage on demand without risking production environments, enabling proactive defense validation and accelerating safe response to emerging threats. The Parser Extension agent addresses log normalization challenges by transforming parser development from a specialized coding task requiring knowledge of regular expressions and complex syntax into a natural language workflow. This significantly reduces the technical expertise required, eliminates operational bottlenecks in accessing new data sources, and enables faster threat detection and response.

Omdia's evaluation confirms that Google SecOps' AI-driven agents deliver measurable operational benefits that directly address the cybersecurity skills shortage and increasing operational complexity. The Threat Hunt agent reduces hunt completion time from hours to minutes while democratizing advanced threat hunting capabilities. The Detection Engineering agent enables safe, rapid validation of detection coverage and rule creation, reducing both risk and skill requirements. The Parser Extension agent eliminates parser development bottlenecks, transforming a process requiring specialized coding expertise into an intuitive, second-long workflow.

Collectively, these capabilities represent a fundamental shift in security operations productivity. By automating time-intensive tasks, reducing skill barriers, and enabling proactive defense validation, Google SecOps empowers organizations to operate more effectively despite persistent talent shortages. Omdia believes these AI agents offer substantial aggregate value by enabling security teams to scale their capabilities, respond faster to emerging threats, and focus skilled analysts on higher-value strategic activities rather than repetitive operational tasks.

©2026 TechTarget, Inc. All Rights Reserved.

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the "Omdia Materials") are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together "Informa TechTarget") or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an "as-is" and "as-available" basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

Get in touch: www.omdia.com askananalyst@omdia.com

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.