FIXING SIEM FATIGUE:

A Practical Guide to Smarter Security Ops

graylog

Like Don Quixote in Man of La Mancha, many security operations centers (SOCs) feel that they are dreaming the impossible dream as they attempt to fight what can seem like an unbeatable foe. Most analysts struggle daily to maintain defenses and respond to alerts. As malicious actors continue to bombard organizations across various attack vectors, SOC teams work feverishly to improve their detections so they can reduce the number of false positives they have to chase down.

For most security analysts, the work is a continuous barrage of notifications and investigations. According to the VikingCloud 2024 Cyber Threat Landscape Report: Cyber Risks, Opportunities, & Resilience, these challenges impact overall security posture. The report noted:

33%

of companies were late to respond to cyberattacks because they were dealing with false positives.

63%

of cyber teams spend 4 or more hours per week, approximately 208 hours per year, dealing with false positives.

15%

of teams spend more than 7 hours per week, approximately 364 hours per year, managing false positives.

In today's perimeterless digital world, SOCs need solutions that aggregate, enrich, correlate, and analyze vast amounts of log data. Cloud technologies can generate terabytes of data every day, yet not all data is equally important. For example, while organizations collect DNS and DHCP logs, they often fail to integrate these sources into their security monitoring because it increases the costs associated with their security information and event management (SIEM).

The SANS 2024 SOC Survey: Facing Top Challenges in Security Operations found that only 38% of respondents send everything into the SIEM. While this number marks an increase over 2023's 29%, it also highlights that a majority of organizations may have limited visibility into their security because they lack the necessary data. Simultaneously, a SIEM has become a "must have" solution rather than a "nice to have" tool, ranking 4th for "highest satisfaction" across 47 different listed technologies.

Additionally, the SANS SOC Survey found that the two most-used models for determining necessary SOC capabilities were the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) and the MITRE ATT&CK framework, enabling organizations to report these top five metrics around SOC service levels:

- Number of incidents handled
- Time from detection to containment to eradication
- Thoroughness of eradication (no recurrence of original or similar compromise)
- Time to discover all impacted assets and users
- Incident occurrence due to known vs. unknown vulnerability

For many SOCs, the dual challenge of false alerts and reporting metrics stems from the original setup decisions. While most SOC teams use a SIEM, the pricing models often force them to make trade-offs around logging what they need and reducing costs. When coupled with the complex, slog of deploying a new SIEM, most SOCs feel that they have no choice but to work with the tools they have, no matter how bulky, time-consuming, and costly they are to manage. Teams struggle to balance the need to manage costs while ensuring they configure the log sources necessary to achieve maximum coverage.

To add to the operational challenges, many SIEMs fail to meet security analysts where they are, ultimately causing additional security risks and diminishing leadership's confidence in the team. Security teams struggle as they integrate more tools whileand still suffering from the effects of the cybersecurity skills gap. VikingCloud's report supports this, noting that 55% of companies believe modern cybercriminals are more advanced than their internal team. Without the right skills, the SOC team is unable to optimize its SIEM's capabilities.

In an attempt to overcome these dual challenges, they seek to leverage generative artificial intelligence (GenAI) to bridge the skills gap. Again, these SOCs face a challenge where they sought a solution. GenAI models that lack training on security data may fail to provide the appropriate value. Further, feeding sensitive log data into a third-party solution can create another potential security risk. Unfortunately, SOCs often have no insight into the GenAI models and the data used to train them. The organization invests in a tool that fails to meet its needs, spending its limited budget without receiving the intended benefits.

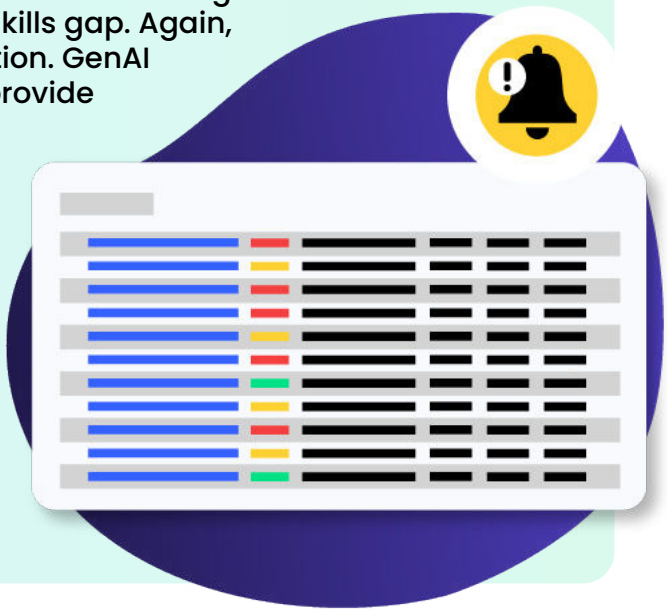

For SOC teams, Graylog Security provides the built-in capabilities that improve their threat detection, investigation, and response processes while helping them navigate the many challenges they face daily.

What are the primary SOC challenges when using a SIEM?

Any SOC team will admit, "the struggle is real." SIEMs have earned a reputation for being behemoths that require specialized skills, leaving many SOC teams frustrated. Despite the belief that security analysts and technologies are less sophisticated than attackers, the reality is that the tools SOC teams use are often difficult to manage and maintain.

SIEMs are not "set it and forget it" tools. Once a SOC onboards the tool, it needs to maintain the data and optimize the configurations. These activities remove the SOC team from its primary tasks, protecting the organization and responding to real security incidents.

While no list of challenges will ever be comprehensive, the main reasons that SOCs struggle with their SIEMs include:

- **Setup:** time consuming, especially tuning it correctly
- **Log capture:** choosing sources based on security use cases and missing important security information
- **Costs:** balancing storage and use costs
- **Platform and data management:** managing the SIEM and the data separately
- **Index life cycle management:** retaining different data to achieve diverse security and compliance objectives
- **False positives:** difficulty tuning rules and investigating too many alerts
- **Cloud log coverage:** leaving logs in vendor solution that leads to missing resources and configurations
- **Required skills:** needing to know proprietary query languages
- **Parsing and normalization:** troubleshooting data when parsed incorrectly

How should SOCs evaluate a SIEM?

Despite these challenges, organizations and their security teams know that they need some tool that aggregates, enriches, correlates, and analyzes log data. High-fidelity alerts and rapid investigations are mere table stakes. The GigaOm Radar Report admits that the well-adopted and well-implemented capabilities for all SIEMs are:

- Multiple ingest streams
- Flexible storage
- Configurable alarms
- Root cause analysis
- Dashboards and visualizations
- Certifications, compliance, and audits

As the security tool market expands, finding the right solution becomes more difficult. A SIEM is a long-term financial investment that can impact the organization's security and staffing strategies. When many solutions appear similar, organizations should consider the following capabilities beyond basic requirements:

- **Alarm fidelity and self-tuning:** defining detection rules and automatically tuning them to reduce false positives
- **Data enrichment:** providing additional context to logs from various sources, including threat intelligence feeds and user directories
- **Collaboration and case management:** enabling security analysts to share information, assign tasks, and communicate within the solution
- **Automation:** offering set up actions like prepackaged connectors, threat enrichment, or contextual information extraction
- **Threat hunting and retrospective analysis:** supporting analyst-driven searches of historic data for suspicious activity that evaded real-time detection
- **Monitoring ephemeral resources:** integrating with tools that monitor containerized and ephemeral environments, like serverless functions
- **Data analysis and risk scoring:** leveraging different machine learning models to analyze data, look for anomalies, and identify threats with a risk score that considers the threat's impact on the real-world environment

Even after eliminating some tools based on these initial capabilities, an organization may still have several contenders. As it seeks to narrow the field further, some additional considerations will include:

- **Documentation and support:** offering comprehensive technical documentation and support services
- **Scalability:** serving large deployments and responding to changes when ingesting data
- **Professional services**: offering additional support for activities like deployment configuration, calibration, incident response, assessing security posture, or digital forensics
- **Licensing:** managing costs in a transparent and predictable way
- **Cost optimization:** reducing costs across the complete tech stack and operations not just a pricing level
- **Security content:** going beyond pre-packaged rules sets and out-of-the-box integrations to deliver security content as soon as new threats, vulnerabilities, and technologies are discovered and deployed

Even researching SIEMs has become a full-time job. With the plethora of options available on the market, almost every technology can appear, at a glance, to look the same. Unfortunately, many organizations make decisions then face buyer's regret later on, especially since switching to a new solution creates additional costs beyond the basic sticker price.

Why Graylog Security?

Graylog Security was built by security professionals who understand the challenges that SOC teams face, especially the struggles they face working with legacy SIEMs. At every step, Graylog Security's development started by identifying a problem that teams faced with the current solutions on the market then worked to build out a true solution that solved them.

Rapid Setup with Automation, Documentation, Support, and Professional Services

Most organizations have some centralized security monitoring. Even if the current setup fails to provide the desired outcomes, the potential costs of moving to Graylog may seem like an administrative burden.

With Graylog's transitioning strategy, documentation, support, and professional services, SOCs can take a fresh look at their current alerts and dashboards to ensure that they transition what matters most rather than packing up everything.

Transition Process

Graylog's three-phased transition process enables SOC teams to achieve a faster return on investment. Just like moving from one residence to another, SOC teams can review their current setup and only bring those alerts and reports that enhance security and compliance.

Phase 1 is taking stock of current needs and configurations by:

- Assessing current security goals and sponsorship
- Inventorying data sources and determining what needs to be migrated.
- Identifying how data is consumed and the requirements for replicating or improving these processes in Graylog.

During this phase, the organization gets to do a comprehensive review of its program to understand what works and remove data or alerts that impede efficient security monitoring and response.

Phase 2 transitions the data and alerts into Graylog by:

- Transitioning ingested data while maintaining security continuity.
- Translating existing saved queries, dashboards, and integrations to the Graylog platform.
- Following Graylog recommendations or replicating established business processes.

This phase enables the organization to build upon what was already working for it so that the SOC team can improve its detections and response capabilities faster.

Phase 3 reviews the value of the transition by:

- Ensuring that the transition aligns with the organization's security goals.
- Optimizing the value derived from Graylog's capabilities, focusing on ease of use and efficiency.

By evaluating the outcomes earlier, the SOC team can prove a return on investment and continue to enhance its processes.

Documentation, Support, and Professional Services

Graylog Security is more than a tool that SOCs can use. Graylog is a partner that ensures everyone who uses the platform has the support they need.

With robust documentation and an active online community, new users have the help they need right at their fingertips. When customers need additional, personalized help, they can turn to Graylog Support services, built with a mission to help users gain competence, capability, and confidence in the solution.

As the SOC team gains experience, analysts can enroll in Graylog Academy with innovative and practical courses built by practitioners. Users can choose from free on-demand courses or purchase live instructor-led trainings to uplevel their skills and gain more value from their Graylog deployment.

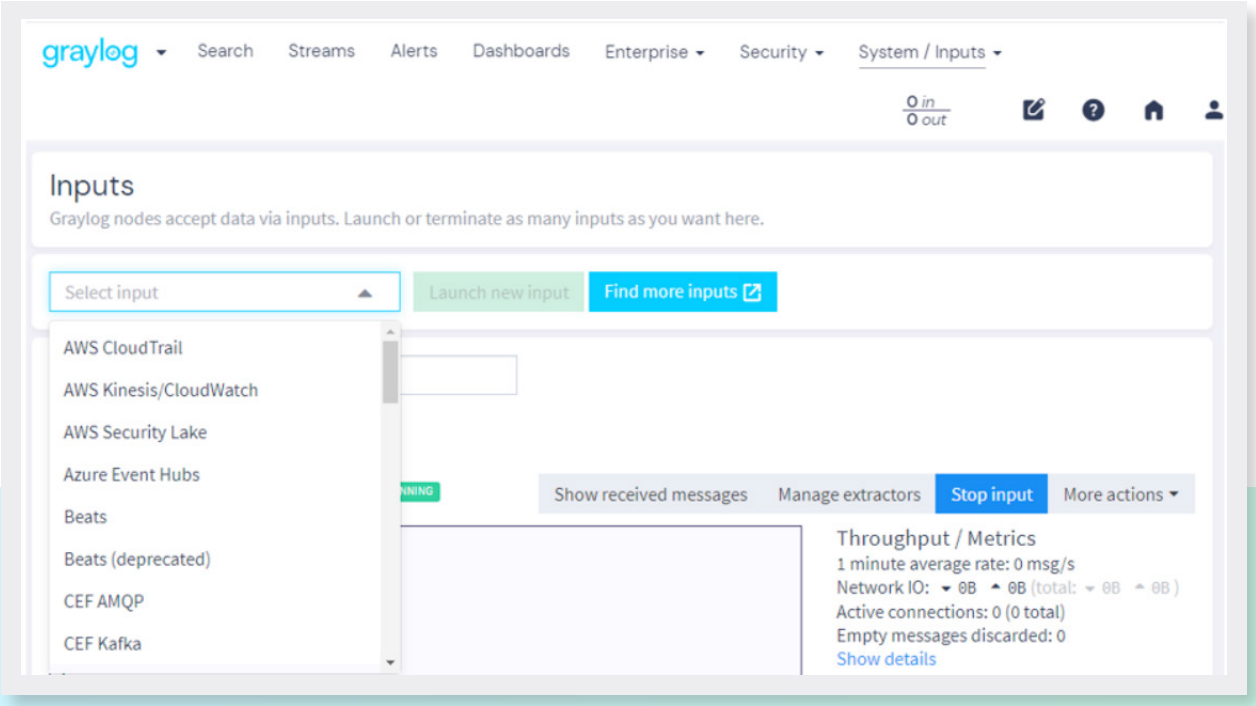

Graylog Security eases SOC team burdens further by providing Illuminate Content that processes logs using a standard methodology and leveraging the Graylog Information Model (GIM) schema to improve log data quality by ensuring events have all the required categories and subcategories.

Illuminate's processing hierarchy is broken into three key areas:

- Processing packs that identify logs, perform parsing and normalization, identify specific messages types, and enrich event messages.
- Illuminate core processor that provides common processing logic, identifies common or reserved IP addresses, enriches event messages with category, subcategory, and event type data, enables geolocation and ASN enrichment, and allows for GIM enforcement.
- Spotlight packs that operate on parsed and processed logs to provide dashboards as well as Sigma Rules and Event Definitions for detecting unusual activity

Streamline Log Capture by Parsing and Normalizing Data Using Native Connectors, Data Pipelines, and Data Enrichment

At its core, Graylog's mission is to ensure that organizations can centralize their security telemetry to gain insights. Using Graylog's listeners and pull inputs, SOCs can configure Graylog to ingest data from all technologies across even the most complex environments. The Graylog interface streamlines these processes by offering drop-down menus that make selecting the right inputs easier.

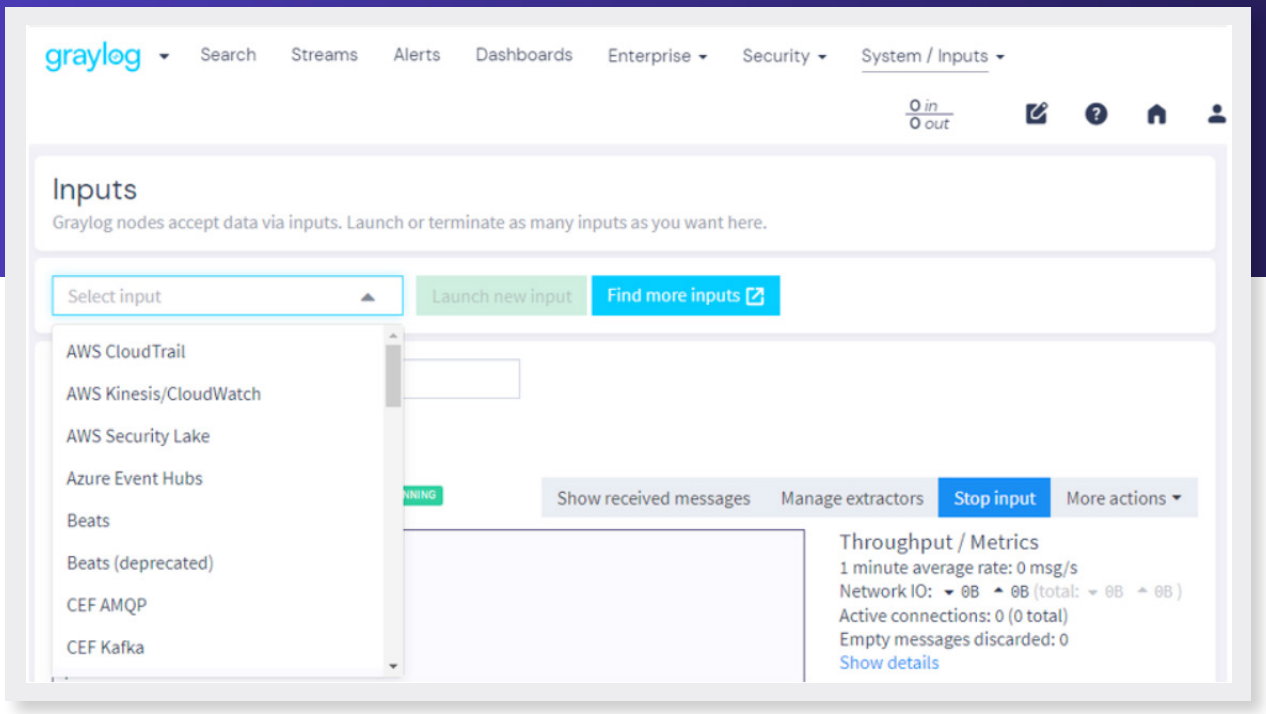

Graylog Security eases SOC team burdens further by providing Illuminate Content that processes logs using a standard methodology and leveraging the Graylog Information Model (GIM) schema to improve log data quality by ensuring events have all the required categories and subcategories.

Illuminate's processing hierarchy is broken into three key areas:

- Processing packs that identify logs, perform parsing and normalization, identify specific messages types, and enrich event messages.
- Illuminate core processor that provides common processing logic, identifies common or reserved IP addresses, enriches event messages with category, subcategory, and event type data, enables geolocation and ASN enrichment, and allows for GIM enforcement.
- Spotlight packs that operate on parsed and processed logs to provide dashboards as well as Sigma Rules and Event Definitions for detecting unusual activity

Scalable Cost, Data, and Index Life cycle Management with Data Routing and Data Tiering to Optimize Costs, Improve Threat Hunting, and Enable Transparent Licensing

As an organization's IT and security technology stacks grow, the environment generates more data. With Graylog, SOCs have the versatility necessary to build data pipelines that respond to both their security and budgeting needs. Graylog's data pipelines automate security telemetry collection, transformation, and delivery by:

- Ingesting data from various sources
- Processing data by cleaning and organizing it
- Transforming data by applying a standardized format
- Enriching data with additional context

Data Routing

Graylog's Data Routing begins by deciding what data goes directly to the data lake or to Graylog, allowing teams to ingest unlimited data by categorization data by immediacy. After identifying the data that they want to keep but not use right away, SOCs can route the compressed data to a data lake, reducing storage costs without impacting the SIEM license costs.

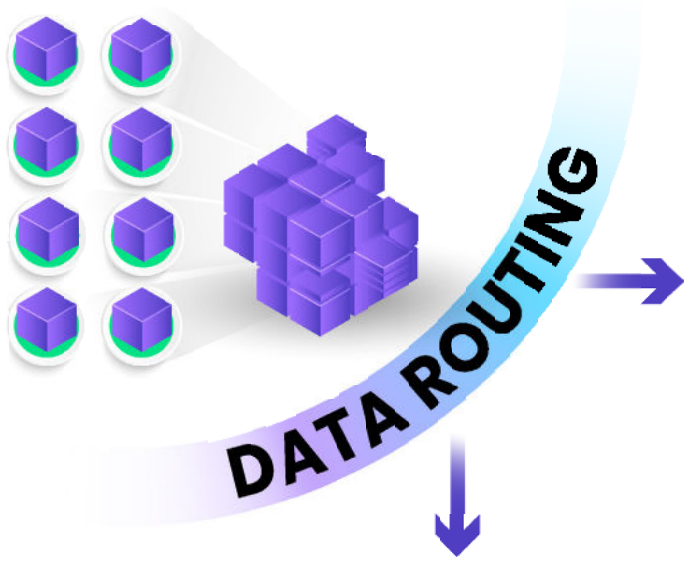

Data tier classification may be based on:

- Performance requirements
- Frequency of use
- Cost efficiency

Graylog pre-processes all data and enriches all fields during the routing process so that SOC teams have immediate access to ready-to-analyze data when they need it. Additionally, Illuminate helps SOCs identify the logs necessary for improved detection based on various use cases, including:

- Data exfiltration
- Privilege escalation
- Insider threats

Data Tiering

Data sent directly to Graylog begins its life cycle in the Hot tier then automatically rolls down to the Warm and Archived tiers using the policies set for its type and log source.

Graylog defines the three data tiers as follows:

- **Hot tier:** high performance tier for frequently searched data that requires easy access and search but increases operating costs
- **Warm or "stand by" tier:** data stored in searchable snapshots, reducing costs for data that requires infrequent access like logs from recent weeks
- **Archived tier:** compressed log data stored in local file or S3-compatible storage object to save money when storing historical data to meet compliance objectives

By selectively pulling security telemetry from a security data lake on an as-needed basis, SOC teams gain the scalability necessary for investigating events, like:
Pulling historical data during an incident investigation
Threat hunting to identify IoCs

SOCs gain more control over and insight from their log data while optimizing spend. Graylog's pricing and licensing offer transparency into how teams can reduce costs while ensuring high-fidelity alerts based on critical data.

Overcome the Cybersecurity Skills Gap and Reduce False Positives with Security Analytics, Risk Scoring, and Continuously Updated Security Content

Graylog is purpose-built to meet security analysts across all experience and skill levels.

Anomaly Detection

Based on an understanding of the organization's unique log data, Graylog's Anomaly Detection identifies outliers by running Artificial Intelligence/ Machine Learning (AI/ML) behavioral analysis to send alerts whenever activity deviates from usual behavior or operates outside normal levels. Graylog feeds enriched data into the Anomaly Detection tool and uses an anomaly index to generate alerts based on the SOC team's configurations.

To help security teams build high-fidelity alerts and detections, Graylog Security offers a Sigma Rule processor for connecting and importing rules then matching messages against them.

Risk Scoring

As part of Graylog Security, teams benefit from two types of risk scores:

- **Event:** risk posed to the environment which can be used to trigger additional research or a security investigation
- **Asset:** risk that an event poses to an asset calculated based on event risk and vulnerabilities associated with the asset

Automated Investigation Triggers

With detections for suspicious or malicious activity matched to log events, SOCs can build automated investigation triggers based on a new alert which improve incident response time for security analysts of all experience and skill levels.

Graylog Security makes it easy to data without requiring a proprietary query language. SOCs can rapidly investigate alerts, reducing alert fatigue, with a collection of dashboards, logs, searches, and events to:

- Create investigations based on timelines, data sets, events, and alerts.
- Associate events with investigations
- See associated data points grouped in a single location.
- Create and reuse investigations to save time and effort.
- Quickly narrow down results
- Add evidence to investigations

Visibility into Threat Coverage

To understand the current security posture, SOC teams need at-a-glance visibility that helps them identify potential compliance and security risks. Graylog maps Sigma Rules to MITRE ATT&CK tactics and techniques so teams have insight into current threat coverage to improve monitoring without requiring specialized security skills.

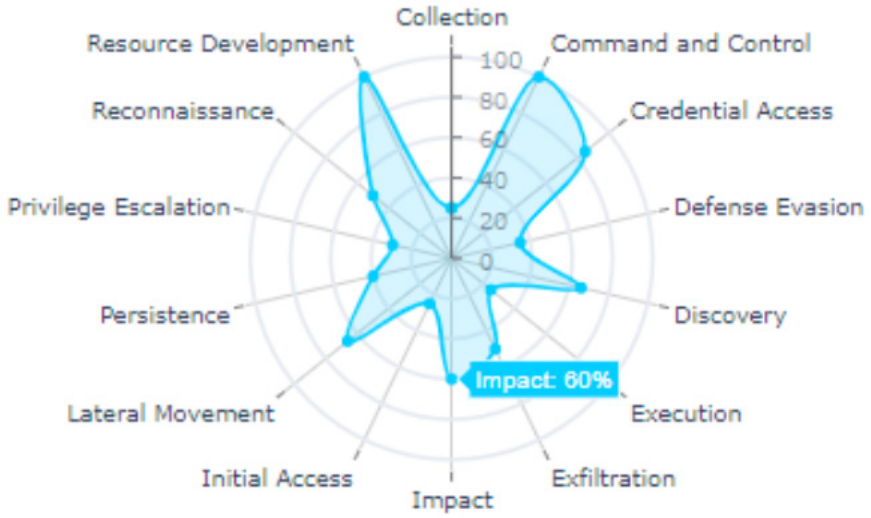

Cybersecurity Focused GenAI

GenAI in the cybersecurity realm is a double-edged sword. It can reduce the time spent on manual tasks, yet it can also create new risks as attackers seek to compromise the models. Graylog Security GenAI-powered reporting that SOC teams can use purposefully without creating additional risks.

Graylog AI-powered investigation reports were purpose-built to maximize value while minimizing risk. Graylog's GenAI reporting is trained on the platform's investigation module, enabling SOC teams to review the reports and compare them with the log data if they need to verify information. The reports analyze the submitted events and logs to generate a detailed report that includes key findings and recommended defensive actions.

Graylog's reporting capability enables junior analysts to understand an event's technical aspects better. Meanwhile, keeping all data within Graylog enables organizations to secure the security telemetry and mitigate exposure to the platform because they no longer need to send it to a third-party application for summarization.

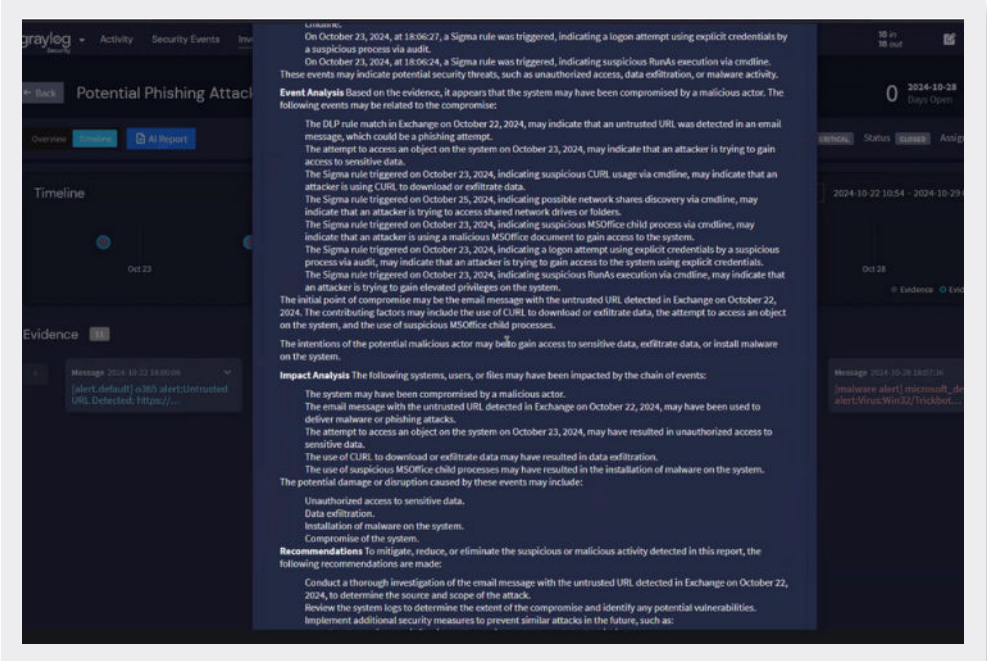

Centralize All Security Activities in a Single Interface to Ease Collaboration and Case Management

Graylog was built so that multiple users could collaborate on investigations by sharing, assigning, and notifying assignees. When beginning an investigation, SOCs can select the responsible users or teams directly in the platform. With all threat detection and incident response data and activities in a centralized location, everyone across the SOC and any other teams have the same information so that they can investigate an incident faster.

Graylog enables cross-functional collaboration by providing two types of permissions related to investigations:

- Investigations Manager: With this role, you have full control over investigations.
- Investigations Reader: With this role, you have read access to investigations only.

For example, by providing read-only access to IT or compliance functions, the SOC reduces the time spent responding to questions without compromising the evidence's integrity.

Graylog's Timeline feature enables SOC teams to understand the incident's history by bringing together all related events and messages in chronological order. Every dot in the Timeline represents an evidence card that includes the details related to log messages and events. With everyone working from the same information, teams reduce the time spent investigating the incident so that they can contain and eradicate a threat faster.

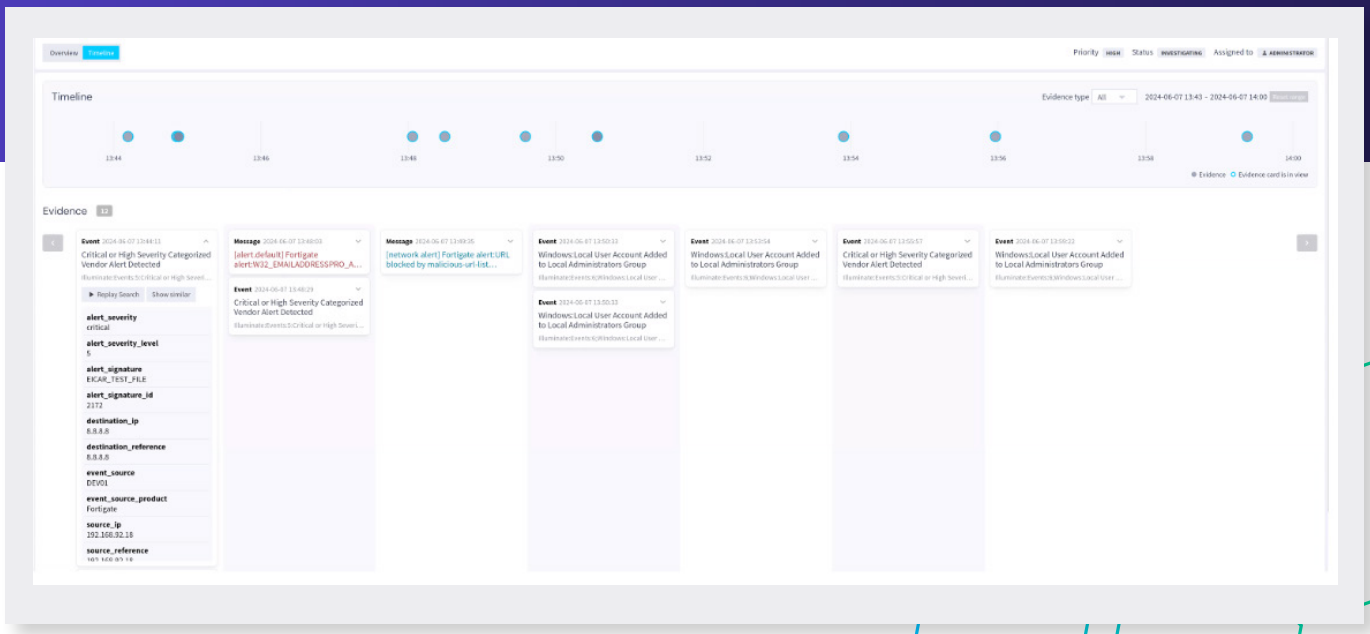

Graylog offers four default status settings:

- **Open:** to document the start of an investigation
- **Investigating:** to track the ongoing investigation activities
- **Closed:** to document the end of the investigation
- **False Positive:** to resolve an investigation that did not require containment, eradication, or system recovery

However, SOC teams can also create new status types or delete existing ones to customize their experience and align to their internal processes or naming conventions. To further customize the experience, SOCs can connect Graylog to ticketing systems or communication tools to ensure that everyone remains informed about the investigation's status.

ABOUT GRAYLOG

Graylog is the no-nonsense SIEM that cuts through noise and complexity. It delivers what security teams need most: full visibility, faster investigations, and smarter detection—without trade-offs or surprise costs. From automated workflows to correlation and anomaly detection, Graylog helps analysts move faster and stay focused. With a product suite spanning Graylog Enterprise, Security, API Security, and Open, we support everyone from large enterprises to lean teams. Trusted by over 60,000 organizations around the world. Learn more at graylog.com or connect with us on Bluesky and LinkedIn.

www.graylog.org
info@graylog.com | 1301 Fannin Street, Suite 2000, Houston, TX 77002