THE GRAYLOG ADVANTAGE

A centralized log strategy has become an essential component in the administration of today's complex IT environments. Since log management data is central to the analysis of securing IT enterprise services, understanding operational performance of IT infrastructure and applications, and key to meeting data privacy regulations, it is critical that organizations understand the positives and negatives associated with the market-leading vendors.

Although Graylog is not as well-known as some in the market, we have a number of advantages over our better-known competitors. Graylog has considerably faster analysis speeds, lower management overhead, simplified query writing and dashboard creation, and petabyte-level scaling with an elegant design of only three components.

This paper examines the critical criteria to consider when evaluating tools for managing your data. There are compelling reasons why Graylog is the best choice for log management and analysis, regardless of your event and data collection stack, technologies, and configurations.

EASIER DATA EXPLORATION AND EXPANDED INSIGHT

Analysts are being inundated with the amount of log data from all of the connected devices, sensors, systems, and applications in their environments. Their work days are full as they address threats from multiple attack vectors. Analysts must be able to do what the current leading tool is not built to do: analyze data without having a complete plan prior to searching and reveal more information as they go to explore data further and find the right answers.

ANALYST EFFICIENCY

The Graylog interface is designed around custom dashboards that visualize a variety of metrics and trends on a single page, where Splunk drives how users discover and explore data. If the user wants to explore multiple topics and manipulate the presentation of the data being explored in Splunk, then multiple screens and apps are required.

From one screen, Graylog allows you to aggregate data from multiple sources, initiate a search across multiple parameters, and analyze, visualize, and report on the data while saving the search to reduce repeatable tasks. You can also trigger alerts when certain thresholds are exceeded or suspicious patterns are emerging, and use those views to respond to those alerts. All this can be done with zero training. Eliminating the need to jump from screen to screen is significantly more efficient, saving considerable time and ending frustration.

FASTER THREAT HUNTING, GREATER AFFORDABILITY

If users want to merge or use complex data manipulation, they must drop into Splunk's query language. The argument is that you can find anything you're looking for since the language is flexible and can perform a number of transformations on your data. However, this approach requires such precision in a search that the user must know exactly what is being looked for prior to creating the query. But this is the very opposite of what you need when looking for threats—if you already knew the answers, you wouldn't be searching for what might be out of place.

Graylog is built for a new wave of data explorers and threat hunters. It allows the user to explore the data without having a complete plan prior to engaging in the search. The power of Graylog's search lies in its ability to expand and reveal more information, delving deeper into the search results to explore the data further to find the right answers. The Graylog approach is considerably more powerful as users are generally not sure of the extent or breadth of an issue prior to the investigation.

Additionally, the knowledge needed to perform searches with Splunk requires both experience and training to understand the language. A typical training investment for Splunk administrators is $3,000 - 5,000 per user—more for advanced certifications—and results in Splunk experts rather than creating IT, Security, or DevOps experts.

Graylog is typically 30 - 50% the cost of Splunk, and any user can run queries and extract meaningful data. Since Graylog began as an open source product, a simple and intuitive user experience was a requirement because there was no commercial support available for a long time. As a result, the Graylog approach is designed to easily transform an analyst into a knowledge expert as opposed to a Splunk expert.

BETTER FLEXIBILITY & SCALABILITY

Threat actors are attacking over continually lengthening periods with more nuanced attacks. As a result, the volume of data required for analysis is growing by terabytes, if not petabytes. Having a scalable, reliable, and manageable infrastructure will be a necessity as the volume of data continues to increase.

BUILT-IN FAULT TOLERANCE

An enterprise environment should be fault-tolerant at each stage of the data pipeline from ingest to presentation. Splunk does not include a component to commit data to disk to facilitate this functionality. To account for this shortcoming, Splunk best practice recommends installing another piece of software, a syslog infrastructure, to receive and commit the data to disk. Additionally, at the data storage and replication stage, Splunk adds yet another piece of software, an accounting server that adds more complexity and requires more management.

Starting at data ingest, Graylog uses a message journal to commit data to disk, preventing data loss in case of a network outage or if indexes or searches fail. Graylog seamlessly provides data replication and data recovery without requiring additional components. This approach allows the user to be confident that there will be no data loss, even in periods of peak or unusual usage. This fault-tolerant capability is critically important to any business that cannot afford to have missing data.

MASSIVE SCALABILITY

Graylog scales to petabyte levels with an elegant design of only three components. Splunk can execute load balancing only through dramatic changes in management of the infrastructure as well as the addition of another server role. Graylog is fault-tolerant by design and requires no further components for it to operate in a distributed, load-balanced manner. Therefore, the Graylog environment easily scales up to meet any size workload from one or two gigabytes to several terabytes per day, then can scale back down to normal data flow as appropriate based on data load.

ANALYZE EVERYTHING AT HALF THE COST

For complete data fidelity and replication from ingest to presentation, the complexity of an environment using Splunk increases quickly. To include data replication, remote host data source management, and search high availability, a Splunk installation requires a minimum of six unique server roles, increasing the administrative burden on its administrators, and therefore the cost. A similar environment in Graylog requires only two roles.

Graylog is built to open standards for connectivity and interoperability for seamless collection, transfer, storage, and analysis of log data. Graylog is also SIEM-agnostic by design—Graylog log streams can pass unaltered or enriched data to any application in your monitoring, alerting, and analysis stack.

FASTER ANALYSIS

When working with enterprise-scale data, every second—or millisecond—matters. The longer it takes to analyze data coming in, the longer it takes to find and resolve issues. Therefore, the ability to simultaneously explore multiple indicators of compromise is crucial to speeding up analysis.

Splunk operates on a limited number of search threads per processor, and a limited number of threads per Splunk process, running a single threaded search per data segment. This design results in a search taking significantly more time and fails to optimize the utilization of the architecture. Splunk has realized the challenge presented by their approach and has implemented some improvements in this area, but these mitigations require further knowledge of the application architecture, another set of commands, and possible escalation to Splunk Support.

The design of Graylog's data storage and retrieval architecture inherently allows for multithreaded and distributed search across the environment. Each search uses multiple processors, multiple buffers, and multiple buffers on a single machine, then multiplies that threaded search across the number of participating nodes in the cluster. This approach gives much faster results, which allows the analyst to work through the dataset without having to schedule, save, or "walk away" from a search to continue at a later time.

A BETTER CUSTOMER EXPERIENCE

From the moment research begins to find a log management solution and throughout the product life cycle, customers rely on a vendor partner to deal fairly and honestly, respond to feature needs on a timely basis, and provide the needed support to meet their particular mission.

The Graylog founders were log management customers before they started the company, so understand the challenges faced by users. They built an organization that recognizes and adapts to meeting customer needs. We understand there are situations where a customer exceeds their contracted data limits and we have never shut down a customer for exceeding their data ingest limits. We are flexible in dealing with data volume growth over time or anomalous data spikes, and help you find the best solution to address your environment's changing needs.

The original Graylog use case dealt with terabytes of data, and we could not find a long-term affordable solution. Our open source core allows us to provide far more value to the customer since our product is priced considerably lower than our competition. That lower price includes collection of all data, enabling you to have an enterprise-wide standardized log platform for handling multiple use cases across security, IT Ops, DevOps, and compliance.

Graylog's interaction with the modern open source ecosystem has driven ease of installation to the point where there are currently 35,000 installations operating worldwide, with only a handful ever requesting on-site help with installation—even in the most complex environments. With Graylog Marketplace, you no longer have to wait for vendor updates to parse logs. The open-source community can help resolve issues.

Solving customer issues immediately is a core tenet of quality customer care. Every ticket opened with Graylog is handled by an engineer; we do not employ Level 1 support staff, nor contract to third parties. Graylog tackles and resolves those issues quickly, rarely ending the week with any open support tickets. As a result, Graylog has been rated #1 by our users, received a 100/100 on service, and a net promoter score of 85 in a recent Info-Tech Research Group report.

Ultimately, it is customer satisfaction that drives a successful software business. We are delighted when prospects research our product and service. The Graylog fan base is large and vocal, consistently trumpeting its support for the Graylog solution widely over social media and in the open-source community.

CONCLUSION

Graylog is purpose-built and designed to deliver the best log collection, storage, enrichment, and analysis. The simplicity in searching, exploring, and visualizing data means no expensive training or tool experts are required. The design, scalability, and flexibility of the Graylog solution allows for faster data analysis, a less expensive infrastructure, and an easier-to-manage environment. And doing business with Graylog is second to none. From product research to post-sale, we provide customer value and delight across the board.

In the final analysis, there are only a few capable solutions for log management at scale on the market today. Graylog is the best solution to address increasing data volumes, reduce complexity, enable flexible analysis, and let you do more with your security and performance data while providing a scalable solution that supports current and future technologies.

graylog

www.graylog.com | info@graylog.com | 708 Main Street, Houston, TX 77002