GRAYLOG ENTERPRISE EDITION

OVERVIEW

Graylog is a centralized log management (CLM) platform that seamlessly collects, enhances, stores, and analyzes log data. Logs are fundamental to any IT operations or security program, and placing them all in a single location greatly simplifies their use.

SYSTEM REQUIREMENTS

We recommend the following minimum system requirements:

- 4 CPU Cores
- 8 GB RAM
- SSD Hard Disk Space with High IOPS for Elasticsearch Log Storage

ARCHITECTURE

Graylog is composed of three components: Graylog, MongoDB, and Elasticsearch. All components can be installed on one server for evaluation or POC deployments. For production installations, we recommend that you separate the Elasticsearch component onto a separate server.

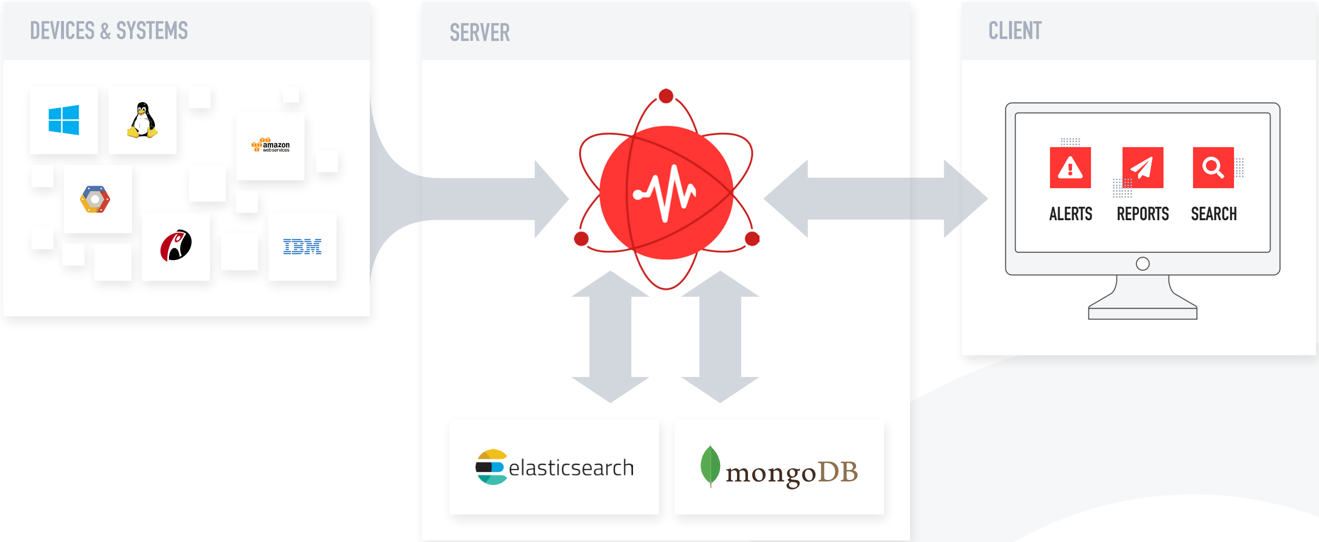

> Passing the logs from Microsoft and Linux devices was incredibly easy which made deployment a breeze. Since implementation we have found it to be absolutely invaluable."
>
> —Infrastructure Analyst

> Moving to Graylog was one of the best choices for managing log data, as generating reports is pretty fast, it can analyze data and send you alerts if a particular threshold has been reached. The interface is really good and it is supported by a really quick repository which is ElasticSearch."
>
> —Systems Administrator

> Moved to Graylog for ease of managing log data. Great for generating reports to deliver on business security and audit requirements. Ease of moving logs from Microsoft and Linux devices."
>
> —Data Storage Engineer

ABOUT GRAYLOG

We love logs. Purpose-built for modern log analytics, Graylog removes complexity from data exploration, compliance audits, and threat hunting so you can find meaning in data easily and take action faster. Founded in 2013, and primarily operating out of Houston, Texas and Germany, Graylog is used to ingest and analyze terabytes of log data across the globe by tens of thousands of users every day.

graylog

ARCHITECTURE	Role-based Access Control	Control who can access what data and Graylog capabilities, includes LDAP/Active Directory integration.
	REST API	Extract data and Graylog alerts to other systems to automate reporting, workflow, and research across your entire Operations Center tech stack.
	Content Packs	Import collections of parsers, alerts, dashboards, and reports tied to a specific data source from Graylog or the Graylog Marketplace. Export all or parts of the configuration of a Graylog instance to move easily from Test to Production. You can also share your custom solutions in the Marketplace.
	User Audit Logs	Log the log management system! Compliance and security best practices require tracking of who accessed what log data and what actions they took against that data.
COLLECTION & PROCESSING	Data Enrichment & Lookup Tables	Add data such as threat intelligence, WHOIS, IP geolocation, or other structured data to assist in analytics. Use to whitelist/blacklist data to remove noise or known “ignore this” situations, and perform faster research.
	Parsers	Convert raw machine data into structured data for storage, search, and analysis.
	Pipelines	Control data processing with staged processing rules to ensure the right parser, data enrichment, and lookup table(s) are applied.
	Sidecar	Centralize deployment and management of 3rd party, and custom log collectors.
	Streams	Route log messages into categories in real time while they are processed to make it easy to target queries, dashboards, and reports for faster results.
SEARCH & VISUALIZATION	Correlation Engine & Event Management	Be alerted via email, text, Slack, or other mechanisms based on a single event, many combined events, or even a lack of events.
	Interactive Dashboards	Customize data visualization using widgets with sophisticated data aggregation for result counts, histograms, statistical values, field value charts, stacked charts, pie charts, and pivot tables.
	Reports	Automate the delivery of key dashboard widgets to your inbox.
	Scalable Search	Build complex, even chained queries in minutes, with many different data visualization output options using Graylog’s web console. No proprietary query language needed. Use Elasticsearch with Graylog’s architecture of multiple processors, and multiple buffers on a single machine, that then multiplies that threaded search across the number of participating nodes in the cluster for immense scalability and speed.
	Views / Research Workflow	Rethink the intersection of dashboards, data queries, and workflows to greatly speed up an analyst’s job investigating errors, performance issues, and security threats. Using a single input parameter, initiate common multi-step analyses and present the results on one dashboard-like screen. Link views to use the results from one as the input parameter to another, then share Views with the rest of the team to save time, ensure consistency, and empower more junior team members.
	Parameterization	Searches are more efficient with parametrization. Enter one or more search criteria for a more comprehensive search. Streamline IT Operations, shorten customer service response time, and take back time for other tasks by saving and sharing searches that the team runs on a regular basis.

WHAT MAKES GRAYLOG UNIQUE

INCREDIBLE FLEXIBILITY

Graylog is built to open standards for connectivity and interoperability for seamless collection, transfer, storage, and analysis of log data. We now centrally manage any machine data collector--ours, custom, or 3rd party vendor--from the admin console, including stopping or starting any whitelisted system processes. Not only that, we can collect other structured data as well, such as DNS lookups from the wire.

RIDICULOUS SPEED

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. Graylog lets you search and investigate multiple issues at once with multi-threaded data retrieval, saving considerable time and delivering results much faster.

EASY EXPLORATION

Graylog lets you analyze data without having a knowing exactly what you are looking for before querying. Graylog expands and reveals more information as you go, delving deeper into the search results to explore further to find the right answers.

MASSIVE SCALABILITY

Horizontally scale to meet any size workload from a gigabyte to petabytes per day. Fault tolerance is built in, enabling distributed and load-balanced operations to prevent data loss.

TREMENDOUS VALUE

There are many facets to price--licensing, processing, storage, and system maintenance--and Graylog is more cost-effective than others across all of them. Graylog Enterprise is free up to 5 GB/day, and beyond that ingest rate, typically ⅓ to ½ the price of major competitors. And that lower price includes collection of all data across your environment. Throw in our top-notch customer experience from initial conversations to purchase to ongoing technical support and product enhancements, and your value skyrockets.

VIBRANT COMMUNITY

Graylog was created by developers for developers, and its start as an open source project has fostered an active and vocal community. Our supportive users offer frequent contributions to the product and roadmap, enthusiastically participate in discussions on the community forum, and have highly rated our customer service in multiple surveys.