



# 2022 State of Logs

An in-depth look at the use and management of logs

# Overview

Managing logs is critical for every business in every industry. With proper management and practices in place, logs have the power to help software engineers optimize the performance of systems and operations, identify and resolve technical issues, better manage resources, and strengthen security.

The New Relic State of Logs report provides context and insights into the growth and maturity of logs management. The report is based on petabytes of data gathered from millions of applications within the New Relic observability platform. To create the report, New Relic anonymized and deliberately coarse-grained the appropriate data to give a general overview of how logs are used and managed. Any detailed information that could help attackers and other malicious parties was deliberately not included in the report.

## The following categories were examined:

- › [The growth and centralization of logs](#)
- › [The impact of open source tools for logging](#)
- › [The power shift happening inside AWS](#)
- › [The most used languages and frameworks](#)



## Software engineers don't want logs to exist in a vacuum

We see a 35% year-over-year increase in logging data. As the volume of log files grows, a trend is emerging with software engineers wanting to have log data available in one place to speed up the time to detecting and responding to transactions, errors, and security incidents. The practice of centralized log management was created out of the frustration and time commitment felt from software engineers in examining thousands of log files across a number of sources to pinpoint and resolve incidents. Even for relatively small companies, managing multiple logging sources and tools becomes increasingly complex, creating information silos and data that is not always adequately parsed or accessible.

Not only does a decentralized approach to log management make detecting and responding to transactions, errors, and security incidents slow and costly, but also using separate tools to accomplish these tasks is confusing and outdated.

Our data indicates that software engineers need access to log data alongside application performance monitoring (APM) and infrastructure monitoring to pinpoint the cause of errors when error rates increase, or when memory usage spikes. The following chart shows that 56% of customers use infrastructure monitoring with logs together, and approximately 14% use logs alongside APM monitoring. Collecting application logs and making them useful has been a pain point for software engineers. With the introduction of features like New Relic [automatic logs in-context](#) and forwarding from APM agents, it is now easier to find the problematic needle in the haystack. As a result, we are seeing a 68% year-over-year increase in customers using APM alongside logs and this percentage is expected to rise.

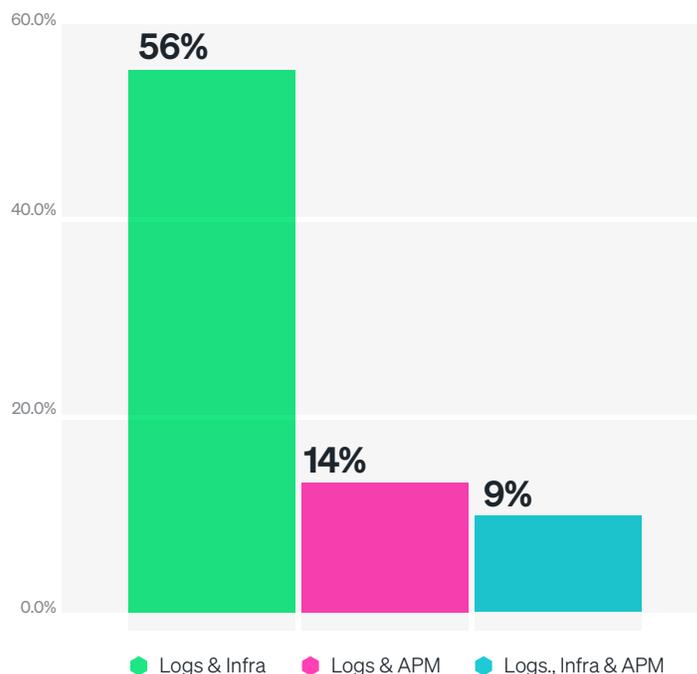


Figure 1: Use of logs monitoring with infrastructure monitoring and APM monitoring



## Fluent Bit is the most used open-source tool for logs

Recent years have seen changes in the tools software engineers use to collect and forward logs. Here we explore the top 10 most notable open-source logging frameworks used to collect log data.

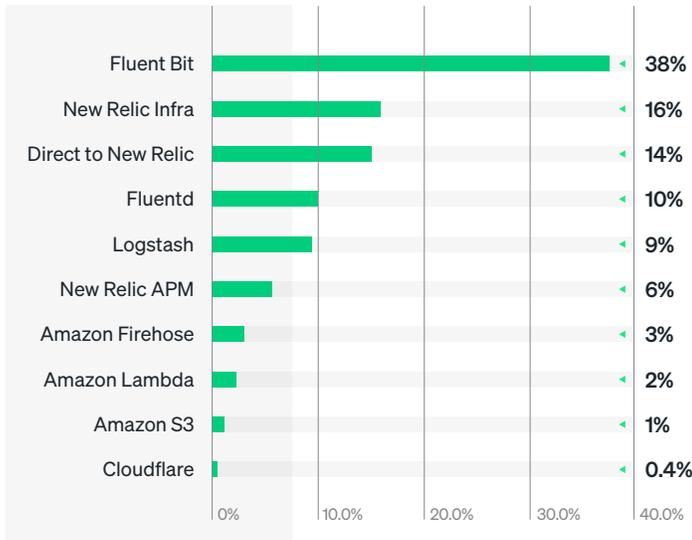


Figure 2: Top 10 plugin type for logging frameworks by percentage of data ingested

The chart above shows that Fluent Bit is the most used open-source processor and forwarder tool amongst New Relic users (38% use it). Fluent Bit started in 2017 to help solve logging challenges in containerized environments like Docker and Kubernetes. It is now considered an industry standard technology and has a massive following today because of its ability to manage observability data at scale across cloud-native, Internet of Things (IoT), and bare metal environments. Its popularity will continue due to its ease of use in Kubernetes environments. Following in usage is the New Relic infrastructure agent (16%), followed by sending log data directly to New Relic (14%) via an HTTP endpoint. New Relic infrastructure agents, APM agents, and the sending of log data directly via custom code illustrate the importance of making data more accessible to developers. With our newest APM language agents, developers can

start collecting logs along with metrics, traces, and events without the need to set up or configure anything. As more application development teams seek integrated log forwarding, we expect APM agents to continue to rise in popularity.

Cloudflare just broke into the top 10 list, which could be a result of our partnership to create stronger integrations.

While OpenTelemetry log sources do not make the top 10, we expect it to rise in popularity as OpenTelemetry adoption continues to grow and the logging specification reaches stable status. We believe this is consistent with a desire to consolidate tools, since OpenTelemetry can be considered a valued standard for adding instrumentation to cloud-native applications and provides a common format across all services.

## NGINX is the most common type of log

When looking at the most popular log type, NGINX is the most used, capturing 38%, followed by Syslog (25%), ALB (20%), and Microsoft IIS W3C (9%). Web servers, load balancers, and content delivery networks (CDNs) are often used by software engineers to perform wellness checks on applications. They are not easily monitored with agents, which is why these three are the most common type of log sent to New Relic.

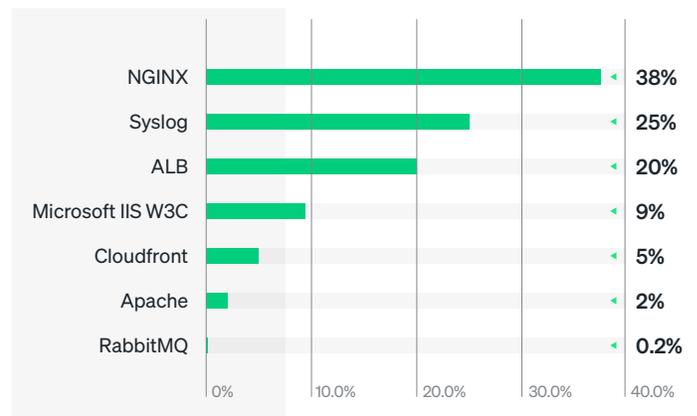


Figure 3: Most common framework by log type

## Firehose will soon be the de facto log forwarder for AWS serverless users

When looking at the most used Amazon Web Services (AWS) in cloud environments, the following table shows Amazon Lambda in the top spot, with 46% of all AWS users in our customer pool adopting the technology. While Lambda has been the main serverless tool for software engineers to use to transfer logs and metrics into CloudWatch automatically, the growth rate of users adopting Amazon Firehose is staggering and we anticipate Firehose taking the lead spot later this year or in early 2023.

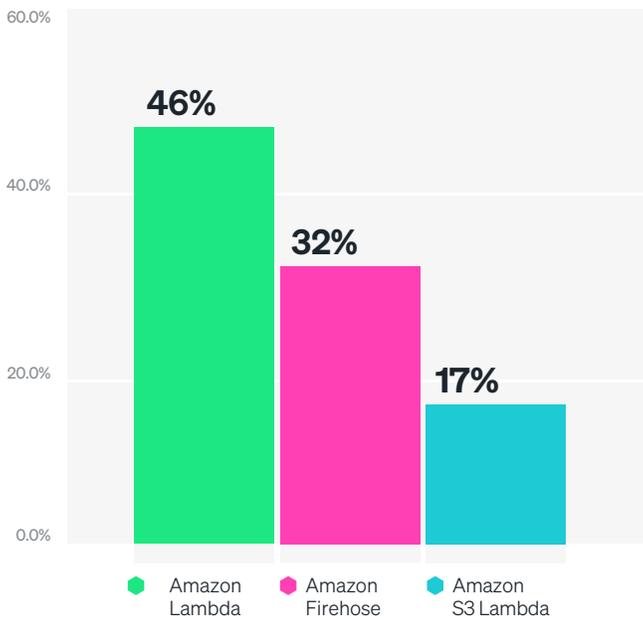


Figure 4: Most used AWS services in cloud environments

Currently, only 32% of New Relic accounts that use AWS have adopted the Firehose technology, but its usage has increased 62% year-over-year compared to 23% year-over-year growth for Lambda. Our data indicates that software engineers are continuing to forward logs from Lambda, but that new deployments increasingly favor forwarding from Firehose. The rise in popularity of AWS Firehose can be attributed to a significant investment in integrations, with broad availability across its partner ecosystem.

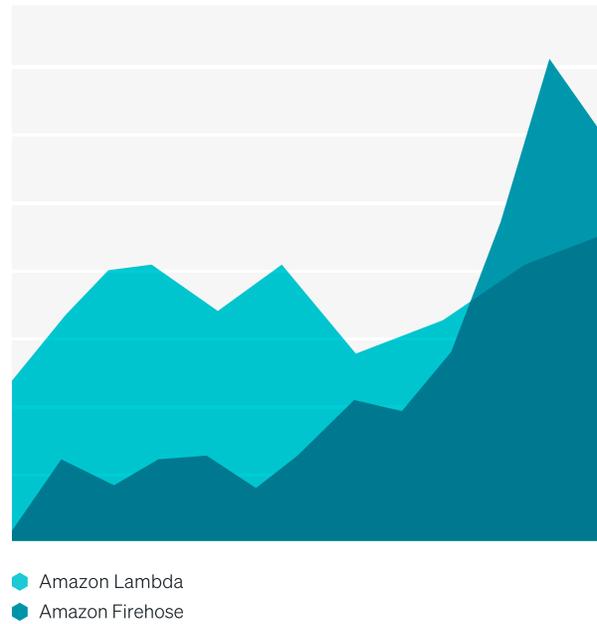


Figure 5: Year-over-year growth of AWS Firehose vs AWS Lambda



## Java has a commanding lead with the most application log data

When examining popularity around languages, the data shows that 50% of all logs ingested by language agents comes from Java. Java has a commanding lead over .Net (26%), Ruby (22%), and Node.js (2%). The Java lead aligns with the overall popularity of Java with software developers and the adoption of the Apache Log4j.

For more insight into trends happening with Java, read the [New Relic State of the Java Ecosystem report](#).

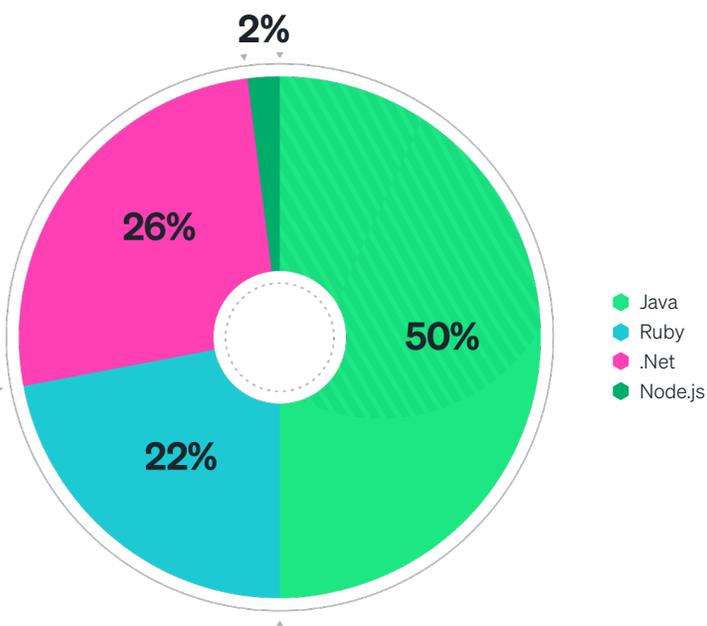
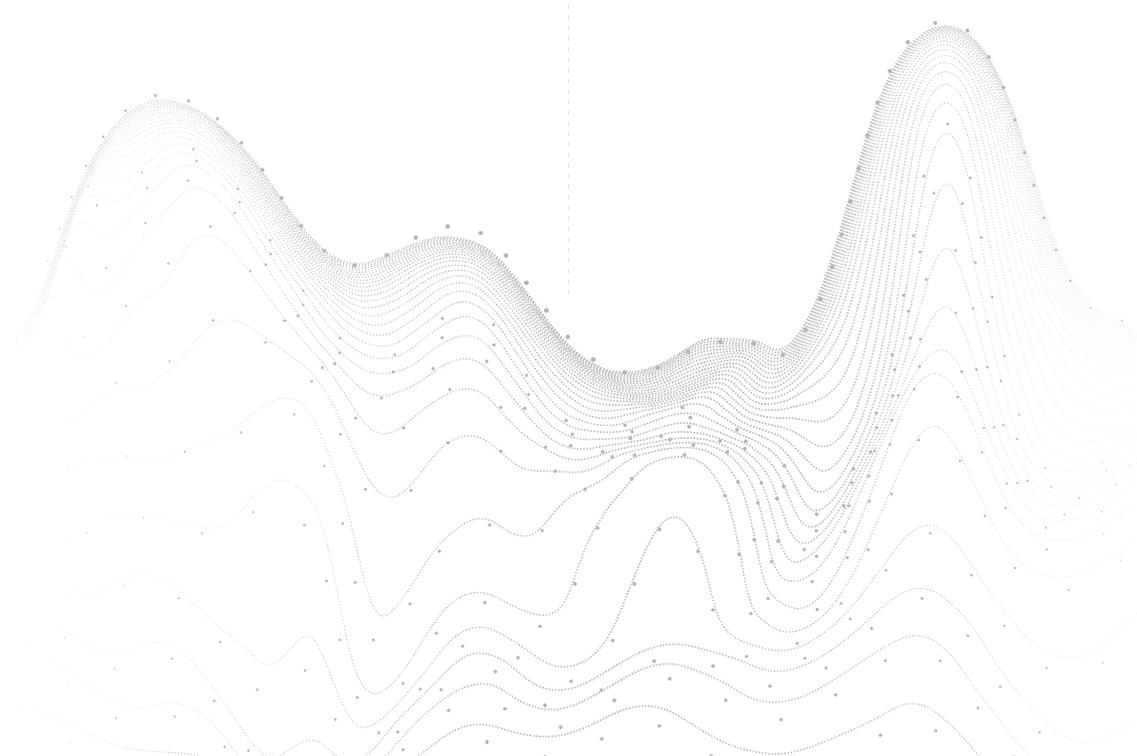


Figure 6: Breakdown of logging ingest by programming languages



# Summary

Log management is increasingly becoming a critical factor to improve operational efficiency by helping find the right information fast, but also to reduce time to market, reduce complexity, and improve customer experiences. New Relic has seen a strong 35% increase in the volume of this essential telemetry data in the last year alone.

In this report we present a snapshot of the global log management and the technologies shaping the future of the practice. This can help you understand the technology trends and help organizations see the big picture as they implement new technology in their development and operations practices.

There are many factors to consider when deciding what technology to use for a log monitoring solution. Expansive selections of cloud technologies, containers, and microservices architectures bring a variety of logs. The pursuit of full-stack observability requires broad access to these many sources of log data. And logs do not operate in a vacuum, as seen by the number of New Relic users who use log data alongside their application and infrastructure performance data to help find the right information fast. We hope this data helps you gain further insights into modern development environments, the tools necessary to build quality digital services at high velocity, and monitoring best practices.

## Research Methodology

The data for this report was drawn entirely from applications reporting to New Relic in July 2022 and August 2022. New Relic anonymized and deliberately coarse-grained the appropriate data to give general overviews of the deployment and management of logs. Any detailed information that could help attackers and other malicious parties was deliberately not included in the report.



## About New Relic

As a leader in observability, New Relic empowers engineers with a data-driven approach to planning, building, deploying, and running great software. New Relic delivers the only unified data platform that empowers engineers to get all telemetry—metrics, events, logs, and traces—paired with powerful full-stack analysis tools to help engineers do their best work with data, not opinions. Delivered through the industry's first usage-based consumption pricing that's intuitive and predictable, New Relic gives engineers more value for the money by helping improve planning cycle times, change failure rates, release frequency, and mean time to resolution (MTTR). This helps the world's leading brands including adidas Runtastic, American Red Cross, Australia Post, Banco Inter, Chegg, GoTo Group, Ryanair, Sainsbury's, Signify Health, TopGolf, and World Fuel Services (WFS) improve uptime, reliability, and operational efficiency to deliver exceptional customer experiences that fuel innovation and growth.

Start monitoring logs with New Relic today.

[Sign Up Now](#)

