Log Analytics for CloudOps

Making Cloud Operations Stable and Agile

By Kevin Petrie

July 2021

Research Sponsored by

About the Author

Kevin Petrie's passion is to decipher what technology means to business leaders and practitioners. He has invested 25 years in technology, as an industry analyst, writer, instructor, product marketer, and services leader. A frequent public speaker and accomplished writer, Kevin has a decade of experience in data management and analytics. He launched, built and led a profitable data services team for EMC Pivotal in the Americas and EMEA, implementing data warehouse and data lake platforms for Fortune 2000 enterprises. More recently he ran field training at the data integration software provider Attunity, now part of Qlik.

Kevin has co-authored two books, *Streaming Change Data Capture: A Foundation for Modern Data Architectures*, OReilly 2018, and *Apache Kafka Transaction Data Streaming for Dummies*, Wiley 2019. He also serves as data management instructor at eLearningCurve. Kevin has a B.A. from Bowdoin College and MBA from the University of California at Berkeley. A bookworm and outdoor fitness nut, Kevin enjoys kayaking, mountain biking and skiing with his wife and three boys.

About Eckerson Group

Eckerson Group is a global research and consulting firm that helps organizations get more value from data. Our experts think critically, write clearly, and present persuasively about data analytics. They specialize in data strategy, data architecture, self-service analytics, master data management, data governance, and data science. Organizations rely on us to demystify data and analytics and develop business-driven strategies that harness the power of data.
Learn what Eckerson Group can do for you!

RESEARCH CONSULTING ADVISORY

GET • MORE • VALUE • FROM • YOUR • DATA

About This Report

This report is sponsored by ChaosSearch, who has exclusive permission to syndicate its content.

Table of Contents

Executive Summary

The cross-functional discipline of Cloud Operations (CloudOps) applies ITOps and DevOps methodologies to the management of cloud applications and infrastructure. Executed well, CloudOps helps enterprises maintain stable and yet agile cloud environments. But this requires effective analytics of logs—those tiny files that capture events such as user actions, application tasks, and compute errors, as well as the messages that applications and cloud components send to one another. CloudOps teams that analyze their logs effectively can achieve stability by optimizing performance, controlling costs, and governing data usage. They can stay agile by responding to events that require speed, scale, or innovation.

But the challenge of log analytics for CloudOps boils down to a simple analogy: you need to see the forest through many, many trees. To make sense of all the trees—including trees that remain on premises—you need to ingest, transform, index, and store millions of logs at high scale and low latency. This generates processing overhead that can choke pipelines or create sky-high compute costs. To overcome this challenge, you must either throttle your logs or widen and extend your analytics pipeline to process more logs. New indexing tools help widen your analytics pipeline, making CloudOps more effective.

CloudOps teams should identify and remediate bottlenecks in their log analytics pipelines, starting with those that cause the most business pain. They should decide whether to reduce the logs or widen the pipeline, based on their staff expertise and available tools. Finally, they should establish common methodologies and tools among CloudOps engineers, DevOps engineers, developers, site reliability engineers, and IT managers.

Introduction

At the turn of the 20th century, enterprises shut down their clunky generators and started buying electricity from new utilities such as the Edison Illuminating Company. In doing so, they cut costs, simplified operations, and made profound leaps in productivity. The promise of modern cloud computing invites easy comparisons to those first electric utilities: outsource to them, save money, and simplify.

Alas, Jeff Bezos is not Thomas Edison, and cloud computing delivers incremental gains rather than profound leaps. For starters, you cannot outsource all your on-premises systems to a cloud provider, thanks to technical debt, data gravity, and sovereignty laws. What you do outsource can generate surprisingly high cloud compute bills. In addition, you still need to carefully manage your application performance and govern your data on the cloud—and connect back to persistent on-premises infrastructure. The results of your transition to the cloud, while promising, depend on how well you manage these challenges.

Cloud operations (CloudOps) helps enterprises manage these challenges and realize the intended results of cloud computing. CloudOps is a cross-functional discipline that encapsulates the processes and tools you need to run applications and infrastructure on platforms offered by AWS, Azure, and Google. You can view CloudOps as the following.

- **Configuration and monitoring of cloud infrastructure**. CloudOps also includes the configuration of virtualized cloud resources; monitoring and control of capacity utilization; and monitoring and response to performance or cost issues. This aspect of CloudOps is essentially IT operations (ITOps) for the cloud, minus physical hardware implementation and management, because cloud providers handle that themselves.

- **DevOps for cloud applications**. CloudOps includes the development, testing, release, monitoring, and management of applications on cloud infrastructure. This aspect of CloudOps is essentially development operations (DevOps) for the cloud.

CloudOps seeks to achieve two classic, sometimes opposing, goals of IT: stability and agility.

- **Stability**. CloudOps seeks to apply consistent methods to meet service level agreements (SLAs), support chargeback efforts, and assist compliance. These objectives derive from ITOps.

- **Agility**. CloudOps seeks to address urgent market demands, innovate, and release competitive enhancements. These objectives derive from DevOps.

The discipline of CloudOps depends on three types of tools. First, **DevOps tools** such as GitHub help build, test, and release software on cloud infrastructure. They provide collaborative capabilities such as version control, bug tracking, and feature requests. Second, **application monitoring tools** such as Dynatrace and AppDynamics monitor the performance of applications, spot issues, help identify root cause, and then assist remediation. These tools monitor both the application layer and infrastructure layer.

Third are **log analytics tools**, the focus of this report. Log analytics tools such as ChaosSearch and the ELK stack (comprising the open-source tools Elasticsearch, Logstash, and Kibana) integrate and analyze logs that capture application and infrastructure events. These tools study a wider array of logs—including message queues, content delivery networks, and virtual private clouds—than application monitoring tools. They also contribute to observability projects that study logs alongside other data points such as telemetry signals.

Log analytics contributes to CloudOps' twin objectives of stability and agility by helping you understand the intricate workings of IT on the cloud.

What is a log?

Each day, the average enterprise's cloud applications, containers, compute nodes, and other components throw off millions of tiny logs. Each log is a file whose data describes an event such as a user action, service request, application task, web page view, or compute error. It might simply summarize the event, the timestamp, and the component(s) involved, or warn of an error or failure that needs urgent attention. Logs also capture messages that applications and other cloud components send to one another.

You perform log analytics when you ingest, transform, search, and query all those files of events and messages. By studying logs, you can identify patterns, anomalies, and trends, and derive metrics to characterize those trends. You can use tools such as Kibana to visualize these outputs and generate threshold-based alerts that guide smart action. Enterprises use log analytics to assist disciplines such as ITOps, DevOps, and now CloudOps.

CloudOps engineers that analyze their logs effectively can make cloud computing more stable and agile. They maintain stability by optimizing performance, controlling costs, and governing data usage. They stay agile by responding to events that require speed, scale, or innovation. In addition, log analytics can reduce the typical tradeoffs between stability and agility. For example, IT teams can maintain SLAs and chargeback controls even while rapidly releasing new software versions. They also can quickly add users or resources without losing visibility into SLA performance or compliance risk.

Log analytics therefore offers a key to realizing the promise of cloud computing. Figure 1 illustrates the CloudOps objectives of stability and agility, and the role of log analytics.

Figure 1. Log Analytics Helps Enterprises Achieve CloudOps Objectives

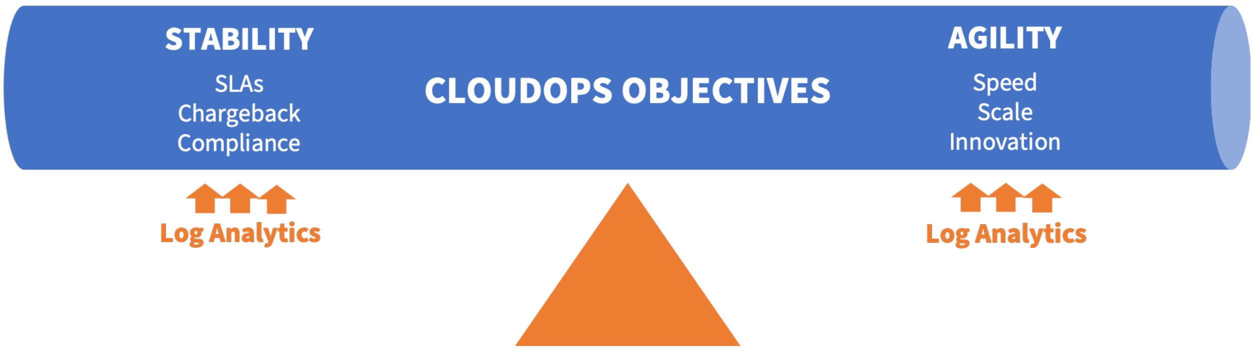

By using log analytics to achieve these CloudOps objectives, enterprises can realize significant benefits. They can improve the productivity of their IT organization, meeting the dynamic needs of the business with less effort and at a lower cost. They can determine and meet service level agreements (SLAs) with greater confidence. They can reduce the operational risk of their cloud initiatives, and increase the likelihood those cloud initiatives will deliver the intended benefits to the top and bottom lines. Finally, enterprise IT teams can commit to future cloud initiatives with greater confidence they will meet deadlines and budgets. This report explores why and how enterprises can use log analytics to achieve such benefits with CloudOps.

Effective log analytics enables enterprises to achieve both stability and agility with CloudOps.

Market Overview

All this matters because the cloud matters. Cloud computing shapes the ability of enterprises to transform themselves and compete in the 2020s. By renting elastic cloud resources, enterprises can support new customer platforms, distributed workforces, and back-office operations. CloudOps helps them apply discipline to cloud initiatives like these, and log analytics makes them smarter as they do so. Here are examples of how log analytics can help.

- **Stability**. The CloudOps engineer monitors log trends to identify an application issue, then parses and correlates logs from that application as well as its container, compute, and storage resources. This helps manage performance and meet SLAs. They also use log analytics to track user actions with sensitive data to improve compliance.
- **Agility**. The CloudOps engineer learns from their container logs that their latest application version consumes more compute cycles for certain workloads than the last version. They also might learn from their compute cluster logs that performance becomes erratic during bursts of user logins. These insights help them rapidly build and release a better application version, with the right allocation of compute resources.

CloudOps and log analytics also help manage two distinct aspects of cloud environments: highly automated processes, and virtualized hardware that you never touch.

Highly automated processes. When IT engineers or business managers subscribe to cloud services such as a software as a service (SaaS) application, development workspace, or compute cluster, they kick off automated processes. AWS, Azure, and Google Cloud provision those services by activating workflows that execute tasks across components of the cloud environment: compute, storage, containers, etc. Cloud providers also automate processes to remediate issues that arise, helping enterprises meet performance and availability SLAs. In these ways, automated cloud processes make things more stable and more agile.

However, automated cloud processes can hurt stability and agility in other ways. CloudOps engineers and teams must keep a close eye on what is happening to help compliance officers maintain governance standards. They must audit how containers and applications authenticate users, and how those users handle Personally Identifiable Information (PII). CloudOps engineers also need to understand how automated cloud processes and on-premises processes impact one another. For example, you might use a SaaS CRM application that integrates with a legacy on-premises customer support application. If the on-prem application stops updating records, are users of the SaaS application automatically notified? Sales reps need to know if their top customer just complained to the support desk. Log analytics helps CloudOps teams understand and control these factors.

Virtualized hardware resources. Who needs a wrench? Rather than taking weeks to install cabinets and plugging cables into ports, CloudOps teams spin up virtualized cloud resources through their cloud provider's portal in a matter of hours. This makes enterprise IT more stable because cloud providers assume the responsibility and liability of this cumbersome physical work. It makes enterprises more agile by helping them scale rapidly.

However, virtualized storage, compute, and network resources still need a lot of oversight. CloudOps teams must study how users, applications, and containers utilize those resources to support chargeback and prevent cost overruns. They also must study resource utilization to optimize the performance of applications they develop and release onto a cloud platform. Log analytics helps CloudOps teams understand and control these factors as well.

Figure 2 illustrates the role of log analytics in on-premises and cloud environments.

Figure 2. Log Analytics on Premises and in the Cloud

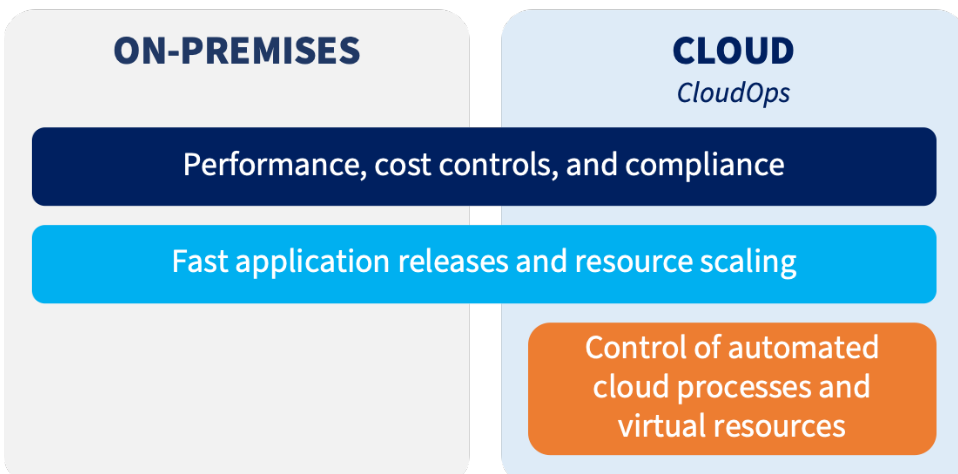

Effective log analytics enables CloudOps engineers to control automated processes and virtual resources on the cloud.

Challenges

The challenges of log analytics for CloudOps boil down to a simple analogy: you need to see the forest through many, many trees. To make sense of all the trees—including trees that remain on premises—you need to process millions of logs at high scale and low latency. This generates processing overhead that can choke pipelines, making log analytics less timely and accurate. It also can create sky-high compute costs that surprise planners and break budgets. These challenges force enterprises into a tough tradeoff and choice:

- **Option 1**: They can process fewer logs to stay within budget, which makes CloudOps less effective.
- **Option 2**: They can choose to process all the logs, which makes CloudOps more effective but might break the budget.

The challenge for CloudOps teams is to somehow process all the right logs while staying within their budgets.

Architecture

CloudOps architectures have three dimensions: the cloud environment, the logs they generate, and the pipelines that process those logs for analytics.

The cloud environment

The primary components of a cloud environment are users, applications, and containers, as well as compute, storage, and network resources. Here is an overview of each component and how they work together.

- **Users**. Professionals and consumers have daily routines that depend on the cloud. Knowledge workers use applications, host files, and communicate on cloud platforms. In their free time, people chat, share pictures, and get driving directions from cloud-reliant applications.
- **Applications**. Modern applications rely on cloud platforms in numerous ways. For example, SaaS applications such as Microsoft Office 365, Dropbox, and Slack all run on the cloud. Mobile apps such as Facebook, Instagram, and Google Maps enable users to interact with each other and distribute content over cloud platforms.
- **Containers**. A container bundles an application with whatever it needs to run, including system tools, libraries, and configuration files, creating a modular package that supports workloads with fewer cloud resources. Containers such as Docker, and container orchestration tools such as Kubernetes, make applications more efficient, more portable, and easier to deliver.
- **Compute**. Clusters of compute nodes, often virtualized, perform the computations necessary to process application workloads. For example, Amazon Elastic Compute Cloud (EC2) offers CPU and memory for applications to consume. They can scale up or down based on changing requirements.
- **Storage**. Enterprises that use cloud platforms typically store their data and metadata in objects, rather than hierarchical files, on Amazon Simple Storage Service (S3), Azure Storage, or Google Cloud Storage.
- **Network**. Cloud providers shuttle data across their platforms using hubs, routers, and switches. While cloud users typically don't see this traffic, it can impact their service levels and compliance standing.

Figure 3 illustrates the components of a cloud environment.

Figure 3. Components of a Cloud Environment

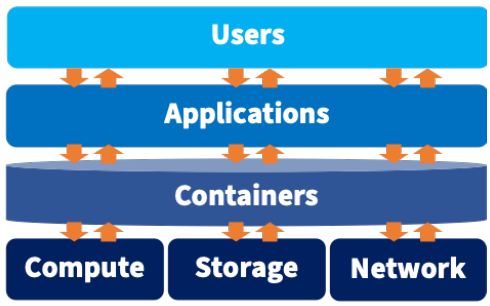

Logs and more logs

All these components generate logs as they operate, interoperate, and communicate with one another. Logs capture operational events related to user activity, content distribution, resource utilization, and IP network traffic. They also capture the messages that components send one another to communicate and interoperate. You can ingest all these logs from the SaaS or custom applications themselves, as well as other sources that include audit logs, message queues, content delivery networks, containers, container orchestration tools, load balancers, and virtual private clouds.

Table 1 provides examples of these sources and what their logs capture.

Table 1. Log Sources, Examples, and Descriptions

Source	Examples	What the logs capture
Applications	Software as a service (SaaS) Custom applications	User actions, application tasks, and security or compliance-related events
Audit logs	AWS CloudTrailGoogle Cloud Audit Logs	User actions that consume cloud resources
Message queues	Amazon Simple Queue Service SQS) Azure Queue Storage	Messages that components send to one another
Content delivery networks	FastlyCloudflare	Content distribution to online and mobile users
Containers	Docker	Resource consumption by containers and applications
Container orchestration tools	Kubernetes	Resource consumption by containers, container groups (pods), applications, and clusters
Load balancers	AWS Elastic Load Balancing (ELB) Google Load Balancer	Traffic and workloads on containers, compute resources, and other components
Virtual private clouds	Amazon Virtual Private Cloud (VPC) Azure Virtual Network	IP traffic on network interfaces

Log pipelines

Log analytics tools contain data pipelines that ingest, transform, index, store, analyze, and visualize logs. The composition and sequence of these pipelines varies by tool and environment, but here is a summary of the typical steps:

- **Ingest**. Data pipelines ingest logs of events and messages from sources such as message queues, content delivery networks, containers, and load balancers. They consolidate these logs for central access and processing.
- **Transform**. Pipelines often convert logs from various sources into a common structure or format, such as Open Database Connectivity (ODBC) or JavaScript Object Notation (JSON). They normalize data to improve integrity, eliminate duplicates, and streamline processing.
- **Index**. Log analytics depends on a central index to locate data without searching all the logs every time users have a query. Indexes aim to make searches more efficient, but can struggle to process high numbers of logs.
- **Store**. Pipelines store logs and indexes either before or after the transformation/indexing stages. Increasingly they rely on new cloud-native object stores such as S3, as described in the prior section. These scalable repositories help enterprises analyze long periods of historical data, and comply with regulatory requirements for data retention.
- **Analyze**. CloudOps teams search and query logs to understand patterns and breaks in those patterns. For example, they monitor operational events and derive metrics about resource utilization. These activities support use cases such as operational performance and customer analytics.
- **Visualize**. Charts, graphs, and dashboards help CloudOps engineers understand the workings of their cloud environments in a simple, visceral way. They boil thousands of logs down to clean metrics and threshold-based alerts, helping drive smarter actions.

Figure 4 illustrates a typical data pipeline for CloudOps log analytics.

Figure 4. Data Pipeline for CloudOps Log Analytics

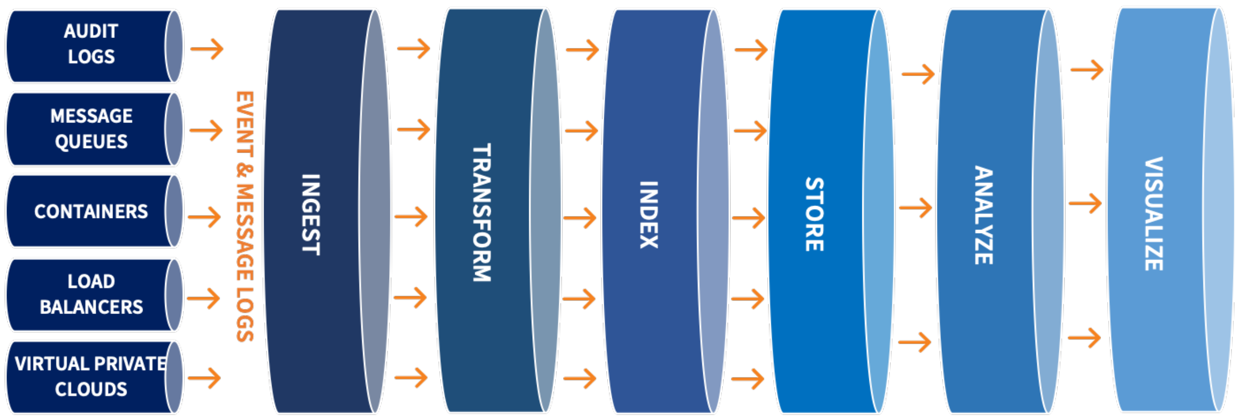

The ELK stack pipeline

Many CloudOps teams use the Elasticsearch engine as part of the so-called ELK stack, which also comprises Logstash, Lucene, and Kibana. You can think of these components as interlocking LEGO pieces. Logstash ingests and transforms the logs, then works with the Lucene database to index them. Elasticsearch queries and searches the logs for analysis, and Kibana helps visualize the results. These pipelines typically leverage cloud-native object stores and integrate with a variety of user interfaces via open APIs. They also integrate with Logstash alternatives such as Fluentd and Vector.

Given the popularity of Elasticsearch and Kibana, various commercial vendors integrate their software with ELK stack components. For example, AWS offers Elasticsearch on Amazon EC2, with the option to ingest logs via Amazon Kinesis Firehose or Logstash. It also offers Amazon Cloudwatch, a log analytics tool that can integrate with Elasticsearch. The vendor Elastic, meanwhile, offers a full ELK stack package. The ELK stack and its components help cost-effectively address basic log analytics use cases for CloudOps.

However, as log volumes continue to rise, the ELK stack can struggle to transform data at scale. Perhaps the greatest chokepoint is indexing. As log volumes rise, they can inflate the Lucene index, which lacks the compression needed to properly handle larger workloads. Each schema change also requires logs to be re-indexed. These factors drive up processing overhead and can choke search and query workloads. Lucene users also must spend time setting up—and scaling—their clusters, schemas, and shards.

Figure 5 illustrates this overload on the ELK stack pipeline for CloudOps log analytics.

Figure 5. ELK Stack Pipeline for CloudOps Log Analytics

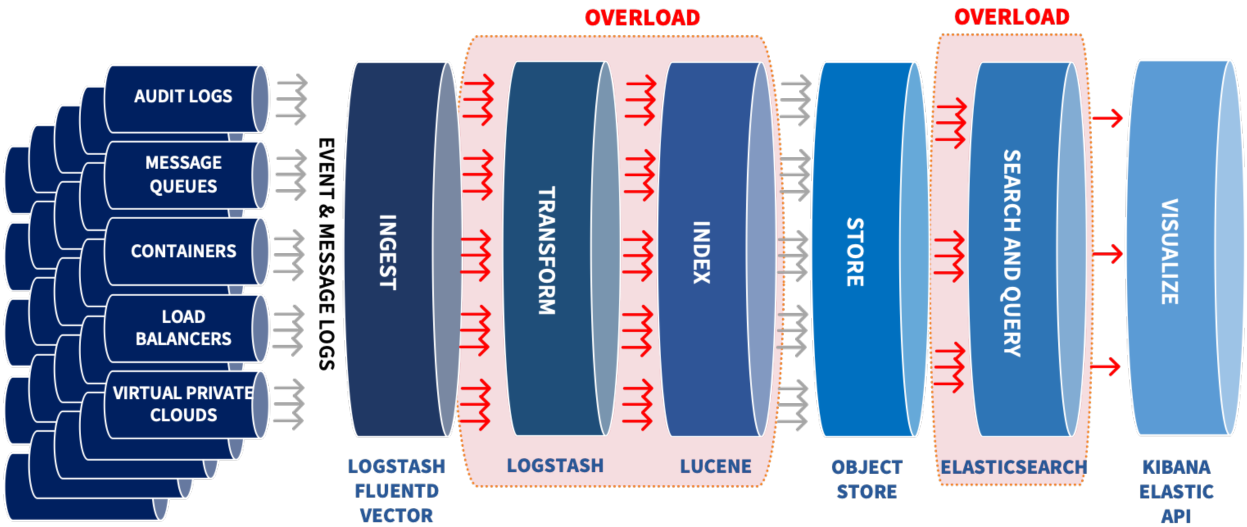

Enterprises have two options to avoid this overload: reduce the volumes of logs that pipelines need to process, or widen the pipeline to process more logs.

Option 1: Reduce log volumes

CloudOps teams can reduce the volume of logs they process at several points in the pipeline. They can reduce the volume of logs they **ingest**, for example by selecting only certain sources, or configuring filters that discard "noisy" logs. They also can reduce the volume of logs they **store**, for example by retaining performance-related logs for three months rather than a year, or eliminating failover systems. Another method is to reduce the **searching and querying** of logs. You can achieve this by giving CloudOps engineers quotas on compute cycles, automatically throttling complex queries, or killing queries that take more than a few minutes. ELK stack users reduce log volumes in these and other ways.

Option 2: Widen and extend the pipeline

Alternatively, CloudOps teams can widen and extend their pipelines with new architectural approaches. They aim to widen the pipeline by processing more logs, and extend it by helping you retain logs for longer periods of time. For example, ChaosSearch replaces Elasticsearch, supports Elasticsearch APIs, and features a built-in Kibana interface. ChaosSearch replaces Elasticsearch's Lucene database with a patented, compressed index that creates less processing overhead. It creates one comprehensive, highly compressed physical index of all the log data. Based on that single physical index, users can then configure their own lightweight index views. They use those views to transform the necessary schema and data, effectively creating virtual schema, with much less processing overhead. They automatically absorb changes to source schema without needing to re-index everything.

This low-footprint indexing approach enables faster queries and higher-volume searches than typically is possible with Lucene. You can still integrate it with popular ELK components such as Logstash for ingestion and Elasticsearch for analytics. They can process more real-time and historical logs than before, with lower latency, higher throughput, and fewer expensive compute cycles. Kibana or other popular tools layer on top of this new stack to provide familiar charts, templates, and threshold-based alerts.

Figure 6 illustrates how low-footprint indexing by ChaosSearch streamlines the data pipeline for log analytics.

Figure 6. Low-Footprint Indexing for CloudOps Log Analytics

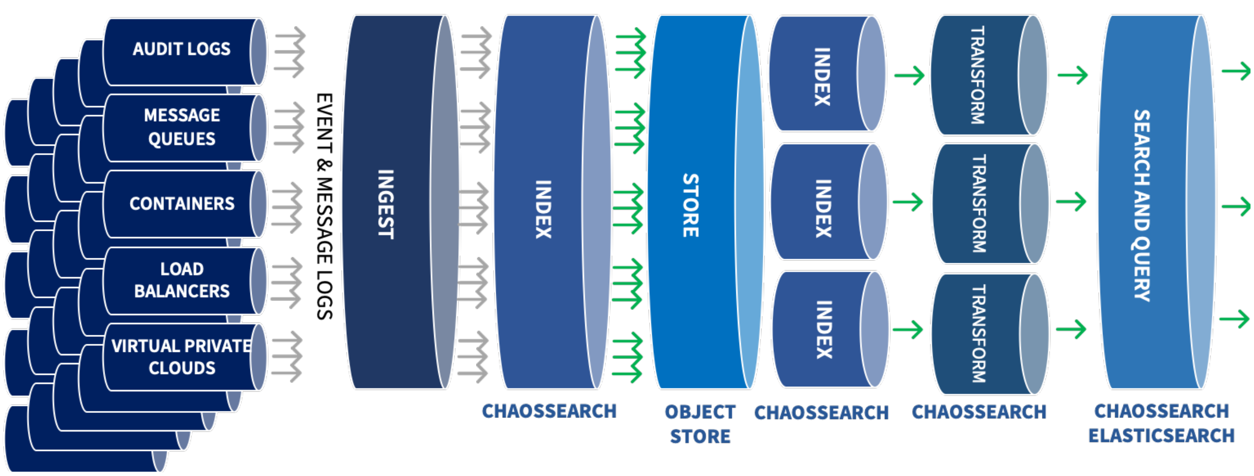

Case Studies

Let's explore two examples of enterprises that streamlined their log analytics pipelines by rearchitecting how they index and therefore analyze logs for CloudOps. Both are customers of ChaosSearch.

Case Study 1: FinTech provider

A European FinTech provider, which processes millions of transactions each day for hundreds of thousands of retailers, struggled to maintain performance standards during the COVID-19 ecommerce boom. A wave of new merchants, customers, and purchases forced them to rapidly scale their cloud resources, generating a volume of logs that overwhelmed their pipelines. This FinTech provider could not ingest and store all those logs, especially logs from AWS CloudTrail and AWS Elastic Load Balancing, within their budget. Their existing indexing methods choked on the log volumes, creating delays in log analytics and therefore slowing their ability to respond to performance issues. Growth was hurting the value of this company's service.

The FinTech revamped its log analytics pipeline by adopting a new, compressed index that dramatically reduced their processing overhead. This freed up compute cycles to ingest, index, store, and analyze more logs faster—which in turn helped them respond to performance issues and maintain SLAs even as transaction volumes boomed. In the meantime, they reduced the compute cost of their log analytics pipeline, and maintained the necessary retention periods to satisfy compliance requirements.

Case Study 2: Educational service platform

Transeo offers a cloud platform that enables students, educators, and various organizations to engage in community service projects and track volunteer hours served. Transeo generates reports that help educators understand student work trends across clubs, teams, and schools. But they struggled to analyze the necessary logs to generate these reports because the Lucene database within their Elasticsearch tool did not efficiently reconcile different file formats. The team at Transeo spent a lot of time reformatting logs from load balancers, containers, and load balancers, which made the reports too expensive to build.

They removed this logjam by changing how they index. They created one physical index of all the logs in their native format, then derived lightweight index views that handle the necessary reformatting with fewer compute cycles and less manual work. As a result, Transeo now can cost-effectively generate community service reports for their educator clients. In addition, they can retain data and logs for ten years to meet clients' compliance policies.

Get Rolling

Log analytics provides enterprises with powerful insights that, if managed well, can improve both the stability and agility of their cloud computing environments. CloudOps teams should follow three guiding principles.

- **Start by easing pain**. Most enterprises already run some applications on cloud platforms, and already analyze the logs to keep things running. And they probably have hit issues already, for example by failing to mitigate performance shortfalls, compute cost overruns, or compliance gaps. Identify the top one or two issues that caused the most business pain—in terms of angry customers, budget owners, compliance officers, etc.—and adapt your data pipelines to prevent such issues in the future. Once this fix is in place, extend that pipeline approach to address other issues.

- **Choose your architecture**. Traditional log pipelines struggle to transform, index, and analyze the high volumes of logs that modern cloud environments generate. If you have the right knowledge and controls, you can dislodge the bottlenecks by reducing log volumes. If you don't have the right knowledge or controls, you run the risk that you miss critical insights. For example, if you retain logs for 3 months rather than 12 months, you won't have last year's performance logs to help you plan your ecommerce workloads for the next Cyber Monday.

- **Collaborate**. Like all technology disciplines, CloudOps requires cross-functional teamwork. CloudOps engineers should teach and learn from colleagues such as DevOps engineers, developers, site reliability engineers (SREs), and IT managers. Your CloudOps, DevOps, and IT roles might already share responsibilities or titles given their common objectives and methods. Consider creating a center of excellence that fosters best practices related to log analytics. Wherever possible, consolidate your logs and activities on common pipeline components. In all likelihood, on-premises IT managers and DevOps engineers can use the same S3 data lake, indexing tool, and visualization tool to address similar use cases.

Cloud computing poses much thornier problems than electricity generation. But a well-crafted CloudOps discipline, armed with smart log analytics, can yield similar outsourcing benefits as Thomas Edison's first electric utilities. Get rolling on your CloudOps initiatives, armed with log analytics, to make this happen.

About Eckerson Group

Wayne Eckerson, a globally-known author, speaker, and consultant, formed Eckerson Group to help organizations get more value from data and analytics. His goal is to provide organizations with expert guidance during every step of their data and analytics journey.

Eckerson Group helps organizations in three ways:

- **Our thought leaders** publish practical, compelling content that keeps data analytics leaders abreast of the latest trends, techniques, and tools in the field.
- **Our consultants** listen carefully, think deeply, and craft tailored solutions that translate business requirements into compelling strategies and solutions.
- **Our advisors** provide one-on-one coaching and mentoring to data leaders and help software vendors develop go-to-market strategies.

RESEARCH CONSULTING ADVISORY

GET • MORE • VALUE • FROM • YOUR • DATA

Eckerson Group is a global research and consulting firm that focuses solely on data and analytics. Our experts specialize in data governance, self-service analytics, data architecture, data science, data management, and business intelligence.

Our clients say we are hard-working, insightful, and humble. It all stems from our love of data and our desire to help organizations turn insights into action. We are a family of continuous learners, interpreting the world of data and analytics for you.

Get more value from your data. Put an expert on your side.
Learn what Eckerson Group can do for you!

About the Sponsor

ChaosSearch empowers data-driven businesses like Blackboard, Equifax, and Klarna to Know Better™, delivering data insights at scale while fulfilling the true promise of data lake economics. The ChaosSearch Data Lake Platform indexes a customer's cloud data, rendering it fully searchable and enabling data analytics at scale with massive reductions of time, cost, and complexity. The Boston-based company raised $40M Series B in December 2020 and is hiring to support its hyper growth. For more information, visit ChaosSearch.io or follow us on Twitter @ChaosSearch and LinkedIn.