AGILISIUM

BIG ON CLOUD. BIG ON DATA.

aws PARTNER Advanced Tier Services

How Log Management and Analytics with **Amazon OpenSearch Serverless** is revolutionizing Operational Analytics?

Learn why log management is essential to the success of cloud-native companies.

Content

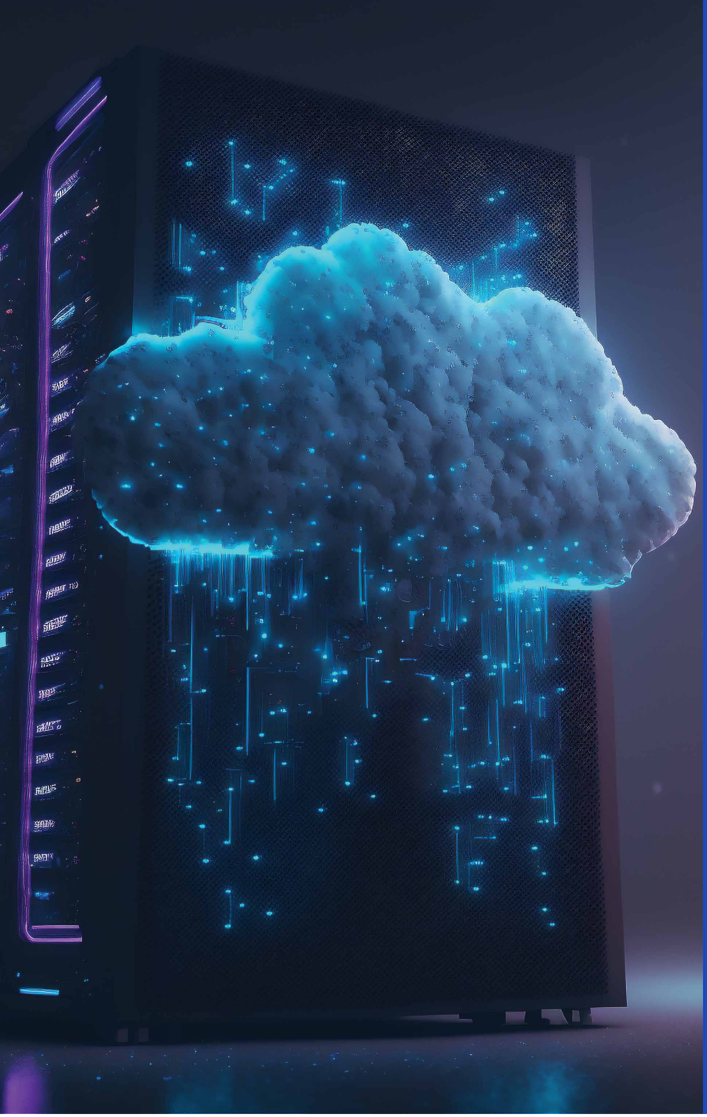

Introduction

Cloud solutions for critical business applications increase scalability, reduce costs, enhance security and compliance, and improve collaboration.

But businesses switching from on-premise to cloud solutions to enjoy these benefits must deal with one challenge: the influx of data from various sources across the cloud environment. Often, these businesses are forced to settle for only deriving partial insights from their data due to the complexity of handling and analyzing it. That is where log management and analytics with OpenSearch Serverless comes in.

Organizations can derive full insights from their data and revolutionize their operational analytics by using Amazon OpenSearch Serverless in log management and analytics. In this blog, we will explore how Amazon OpenSearch Serverless is transforming the way organizations manage and extract insights from their log data.

What is Log Analytics?

Log analytics is the process of analyzing the log of events occurring in an organization's systems and networks. It helps users organize and extract valuable information from the event logs. An event is an occurrence within an IT ecosystem occurrence and includes user activity, system performance, application performance, etc.

What are the four main functions of Log Analytics?

The four main functions of log analytics are normalization, pattern recognition, classification and tagging, and correlation analysis.

Normalization

Normalization is a data management method that involves standardizing log data into a common structure or format. Log analytics help ensure attributes from log entries across an organization's technology stack are represented in the same manner, eliminating inconsistencies.

Pattern recognition

Pattern recognition involves identifying patterns, anomalies, and trends in the data. Log analytics typically use machine learning algorithms to sift through the log data to find patterns, anomalies, or trends.

Classification and tagging

When the patterns, anomalies, and trends have been identified, log analytics classify, and tag data based on its content. Classifying and tagging categorizes data into different groups, such as security events, performance metrics, or error messages.

Correlation analysis

Log analytics also helps with correlation analytics, which means looking for the relationships and dependencies between events. Through correlation analysis, users can quickly identify the root cause of issues across their technology stack.

Challenges customers face with Log Analytics Management

Log analytics management involves collecting, aggregating, storing, analyzing, and visualizing log data. While it offers numerous benefits, it also has some challenges, including the following:

Data overload

As the number of applications, servers, and IoT devices increases, so does the log data they generate. Managing or analyzing this large volume of event data in a big network can be impossible or create a heavy burden on IT resources.

Growing need for real-time search at scale

Customers want to find the right product, service, or information as quickly as possible. To facilitate this, organizations have to provide relevant data in real time. It can be challenging and costly to build and maintain a system that can handle these needs securely at scale.

Log management infrastructure

Users must invest in software and hardware tools and skilled personnel. It can be quite expensive and may divert limited resources from other IT initiatives.

Integration and interoperability: It may be difficult to integrate and correlate log data from various sources if they have different formats. Additionally, log management tools may not be interoperable, creating data silos that prevent a holistic analysis.

Security and compliance

Log data can be sensitive, so it is subject to privacy regulations. Customers must ensure their log analytics management is secure to comply with regulatory requirements.

Despite these challenges, log management and analytics have several benefits, as highlighted below.

What are the real-world benefits of Log management and Analytics?

Some of the most significant real-world benefits of log management and analytics include the following:

Streamlining business operations

Through log management and analytics, businesses can identify opportunities to streamline business operations and optimize system performance. For example, analyzing logs can reveal areas where organizations can automate tasks to improve efficiency.

Improving troubleshooting and resolution

It allows organizations to create and implement a framework for improved root cause analysis, effective troubleshooting, and speedier incident response and resolution. Through log analytics, IT teams can pinpoint the root cause of the problem and resolve it quickly to avoid downtime.

Better resource allocation

By analyzing event logs, organizations can prioritize key items to improve resource allocation and use bandwidth appropriately.

Enhancing cybersecurity

Log management and analytics can enhance cybersecurity measures through proactive and ongoing monitoring. Organizations can identify and address security incidents before they become major problems.

Ensuring compliance

Log management and analytics can help organizations ensure compliance with industry-specific regulatory bodies, such as HIPAA, GDPR, PCI DSS, and internal policies.

Improving sales and marketing campaign effectiveness

Log analytics can evaluate metrics like conversion errors, website traffic, and other metrics to understand and improve the effectiveness of sales and marketing campaigns.

Increasing opportunities for collaboration and timely communication

Log management and analysis ensures everyone has the same information about the state of the infrastructure. It facilitates effective communication and collaboration, improving decision-making and the speed of resolving problems.

Why OpenSearch?

With the explosive growth of log data, organizations need a search and analytics tool that can handle high data volumes in real-time. It is the only way to maximize these real-world benefits of log management and analytics. One such tool is Amazon OpenSearch.

Amazon OpenSearch can help organizations with high log data volume from various sources to ingest, normalize, process, and analyze it in real time and deliver useful insights.

Organizations should consider using OpenSearch as it will help them derive the maximum value from their log data, among other benefits.

What is Amazon OpenSearch?

Amazon OpenSearch is a search and analytics suite used for log analytics, real-time application monitoring, website search, and other use cases. It allows organizations to deploy and scale a search cluster in the cloud without having to manage its underlying infrastructure.

It is the ideal tool for log management and analytics because it can handle large amounts of data and queries and is highly scalable.

Benefits of Amazon OpenSearch

Amazon OpenSearch offers several benefits that make it a powerful tool for search and analytics use cases. They include the following:

Managed OpenSearch

Focus on analysis instead of spending time managing your deployment, and adjusting deployment configurations as requirements change—while using the power of open-source search.

Secure

Meet and maintain high security for authentication, authorization, encryption, audit, and regulatory compliance.

Observability

Deliver log and trace analytics solutions while developing interactive queries and visualizing results with high adaptability and speed.

Cognitive

Use machine learning to detect anomalies in real time, autotune your clusters, and personalize your search results.

Cost-conscious

Eliminate operational overhead and reduce cost with automated provisioning, software installation, patching, storage tiering, and more.

Unique features of Amazon OpenSearch Service functionality

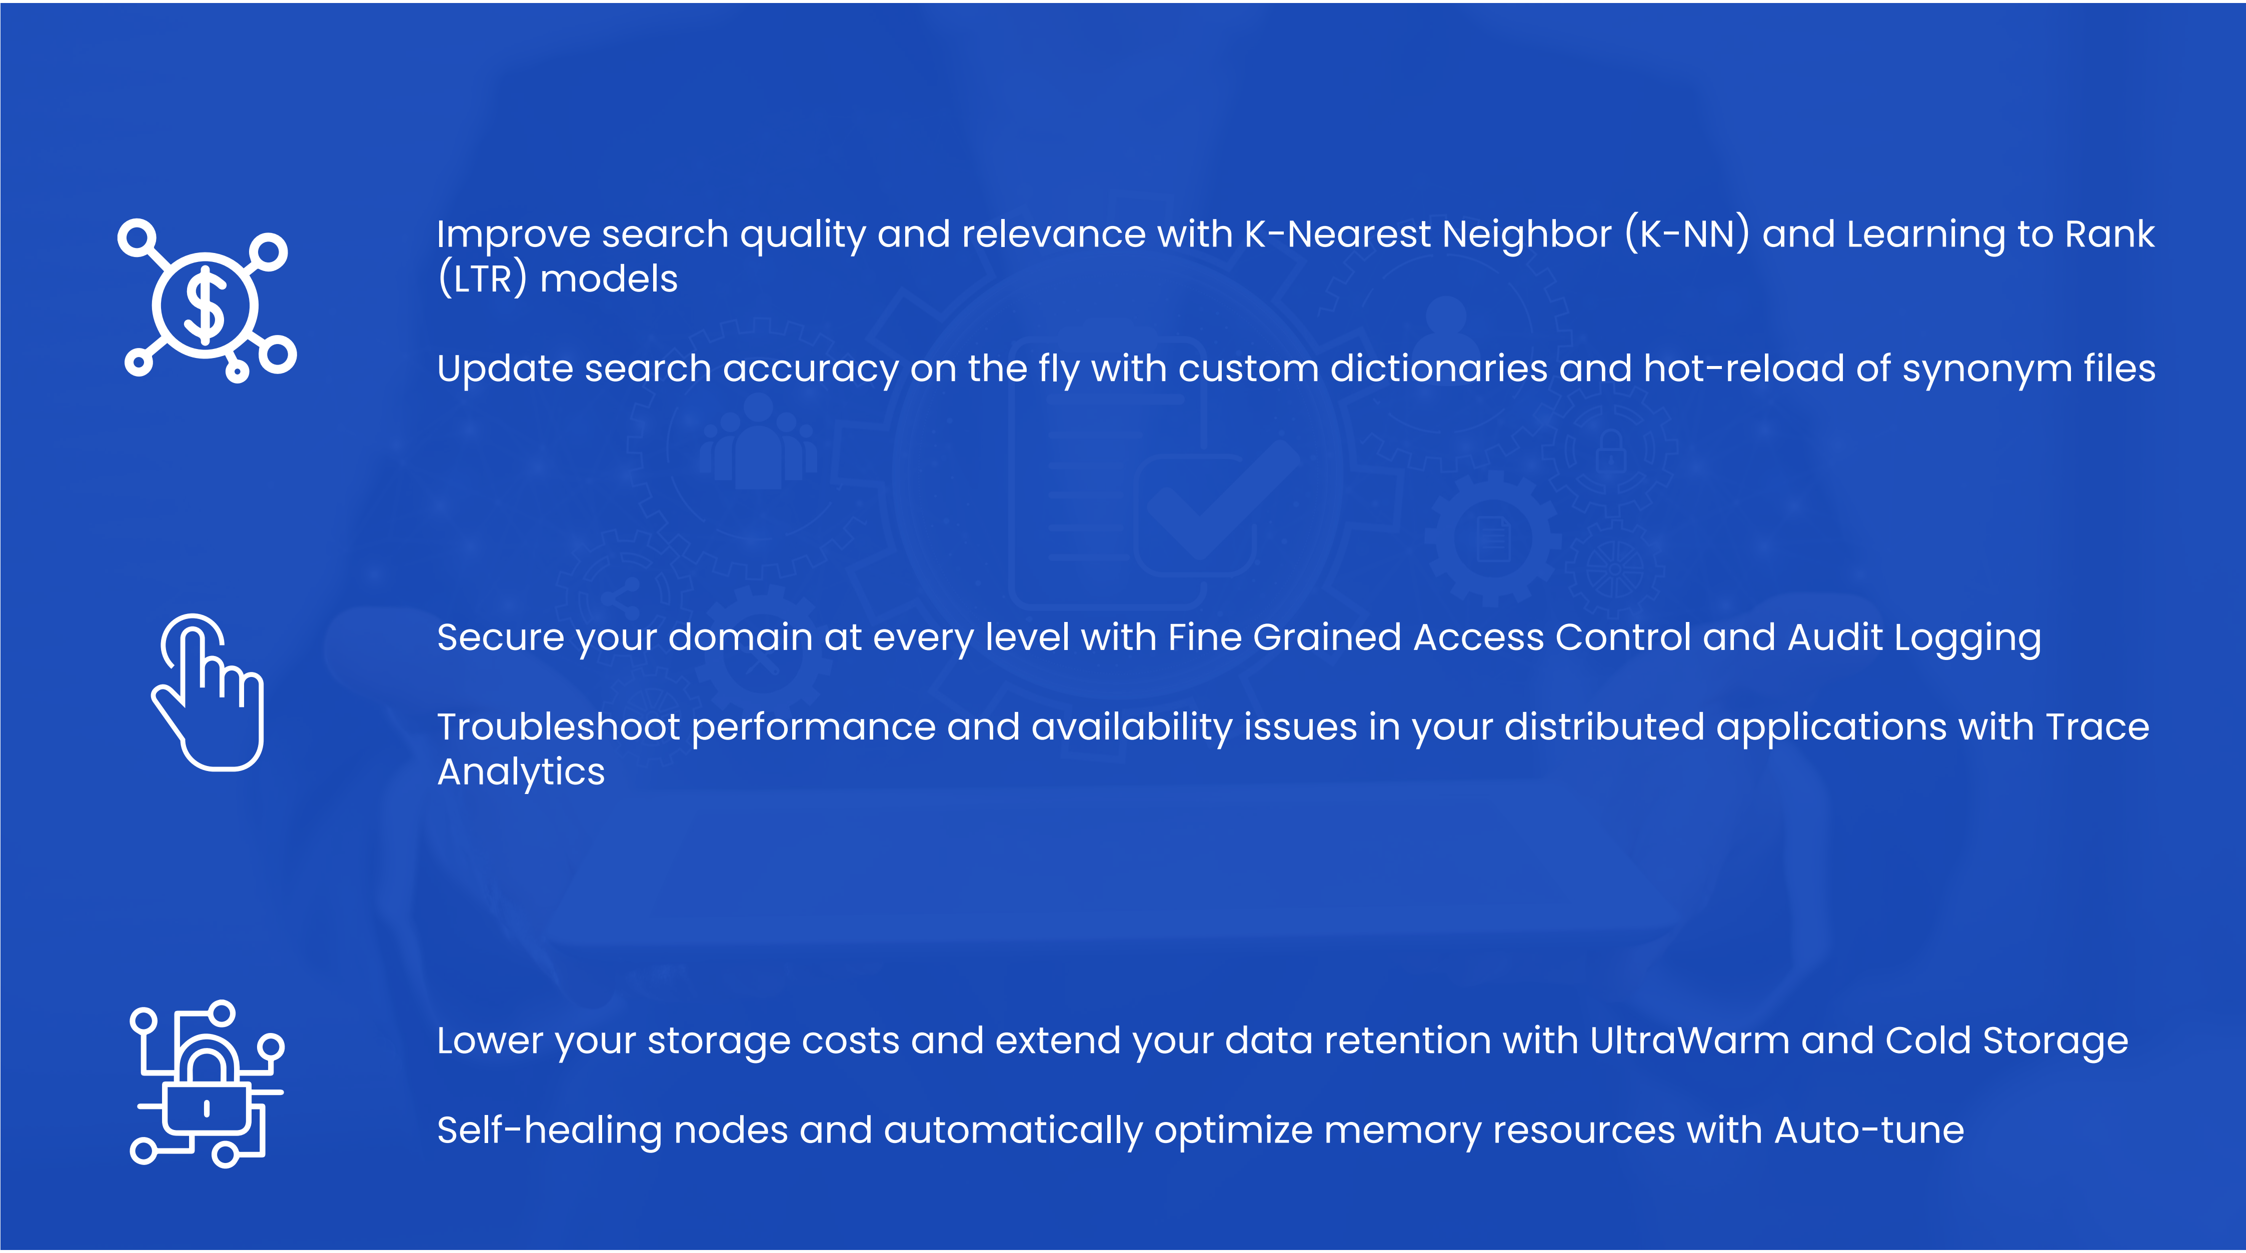

Improve search quality and relevance with K-Nearest Neighbor (K-NN) and Learning to Rank (LTR) models

Update search accuracy on the fly with custom dictionaries and hot-reload of synonym files

Secure your domain at every level with Fine Grained Access Control and Audit Logging

Troubleshoot performance and availability issues in your distributed applications with Trace Analytics

Lower your storage costs and extend your data retention with UltraWarm and Cold Storage

Self-healing nodes and automatically optimize memory resources with Auto-tune

Top use cases of Amazon OpenSearch

The top use cases of Amazon OpenSearch include the following:

1.Observability

Observability is the ability to gain insights into the health and performance of an organization's applications and infrastructure. It involves the analysis of logs, metrics, and traces.

Amazon OpenSearch centralizes log analytics, so organizations can identify or predict performance problems across their operations. It also offers cross-cluster search,which allows users to analyze and query all of their log data via a single intuitive OpenSearch Dashboards interface.

2.Application & Infrastructure Monitoring

Organizations must proactively monitor their applications and infrastructure log data to find issues and solve them before they become major problems.

Organizations can use OpenSearch's real-time search and log analytics capabilities to identify or predict performance problems. It also enables IT teams to do real-time root causes and forensic analysis. This can reduce the Mean Time to Detect (MTTD) and Mean Time to Resolve (MTTR).

3.Search

Search is the ability to quickly find relevant data in applications, websites, and data lake catalogs. Amazon OpenSearch allows organizations to deliver high-quality and personalized search results to customers.

Organizations can also access all Elasticsearch's search APIs, supporting natural language search, auto-completion, faceted search, adjustable ranking, and location-aware search.

4.Security Monitoring

Security monitoring involves keeping data safe and preventing security threats like data breaches.

Amazon OpenSearch Service accelerates security incident detection, forensic analysis, and response. It can quickly analyze logs from disparate applications and systems across an organization's network. This can help detect security threats and mitigate them in real time.

Search Use cases with Amazon OpenSearch Service

Amazon OpenSearch service allows users to add search functionality to an application, site, or document repository. Search functionality will help users quickly find what they need. The search use cases of Amazon OpenSearch include the following:

- **Website search:** OpenSearch service takes in website data, indexes it, ingests it, and returns results at high speeds, no matter the size of the website.
- **Application search:** OpenSearch allows organizations to implement search functionality to applications backed by databases. It delivers AI-powered ranking for search results quickly.
- **Document repository search:** OpenSearch allows organizations to search across a large depository of text in a data lake or internal wiki.

Streaming data Use cases with Amazon OpenSearch Service

Streaming data occurs when data is generated continuously and must be processed and analyzed in real-time as it is collected. Amazon OpenSearch Service can handle streaming data from thousands of sources. Some streaming data use cases include the following:

- **Centralized log analytics:** Consolidate logs from multiple applications and analyze them for various use cases.
- **Observability:** Gain insights into the health and performance of technology and products by analyzing logs, metrics, and traces.
- **Track analytics:** Use and visualize Open Telemetry data for distributed applications - and monitor event flow between applications to identify performance issues.
- **Log analytics:** Analyze logs to interpret and derive useful insights from them.
- **Security analytics:** Take a proactive approach to security through threat monitoring, detection, and alerting features.
- **Anomaly detection:** Leverage ML-based anomaly detection features to automatically detect anomalies during data ingestion.

Through these streaming use cases, organizations can get real- time insights, make faster decisions, and improve operational efficiency.

How does it work?

1

Send data as JSON via REST APIs

2

Data is indexed - all fields searchable, including nested JSON

3

REST APIs, for fielded matching, Boolean expressions, sorting, and analysis

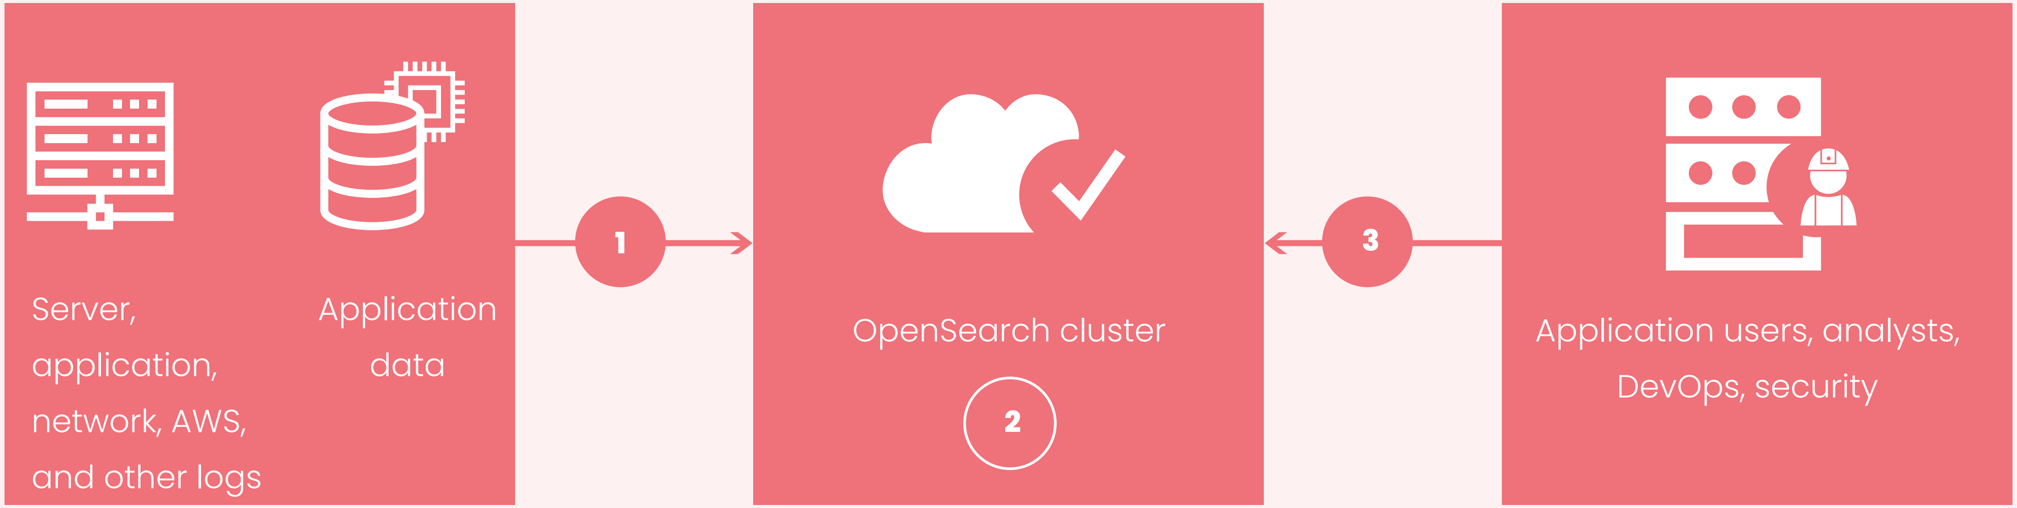

Server, application, network, AWS, and other logs

Application data

1

OpenSearch cluster

2

3

Application users, analysts, DevOps, security

Why Amazon OpenSearch Serverless?

Amazon OpenSearch Serverless is the serverless version of OpenSearch. Being serverless means users can run search and analytics workloads without having to configure, manage, or scale OpenSearch clusters. It has the same features as Amazon OpenSearch but has added benefits like automatic scaling, lower costs, and reduced management overhead.

OpenSearch service offers machine learning capabilities to help users improve search relevance and quality and detect anomalies in real-time streaming data. This allows users to detect and resolve issues before they become critical.

The OpenSearch Service also supports K-Nearest Neighbor (K-NN) and Learning to Rank Models. These features allow it to improve customer experiences by enhancing the relevance of search results. OpenSearch machine learning models are distributed and processed across nodes to ensure high performance at scale. It is also user-friendly, so users don't have to have machine learning expertise to optimize performance and user experience while cutting costs.

Data lifecycle in Amazon OpenSearch Service

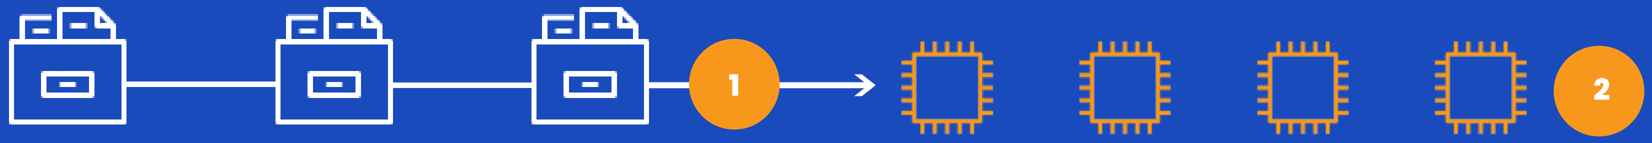

Hot tier
Indexing and fast access

1. Send data to Amazon OpenSearch Service. Index State Management (ISM) automates index migrations or deletions

2. Data is indexed and stored in the hot tier Migrate the index to UltraWarm and Cold

3. Storage for long-term, low cost storage

4. Delete the index at end-of-life

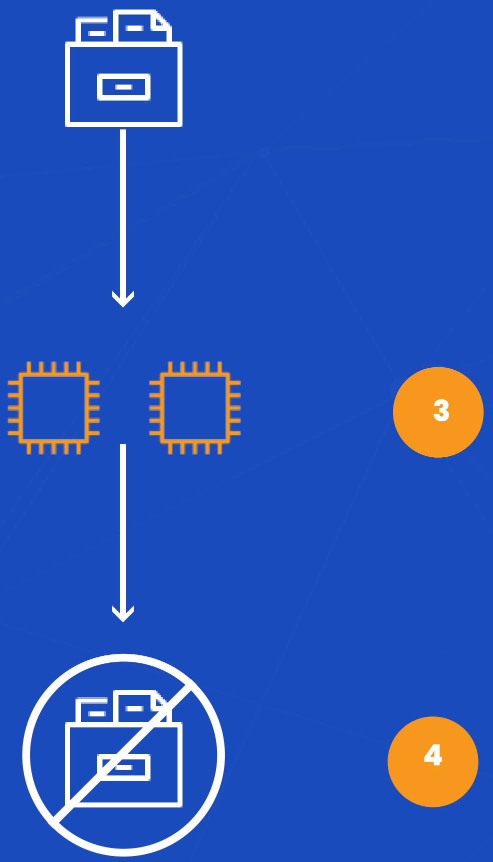

UltraWarm and Cold Storage
Low-cost, longterm retention

Amazon OpenSearch Service data ingestion flow

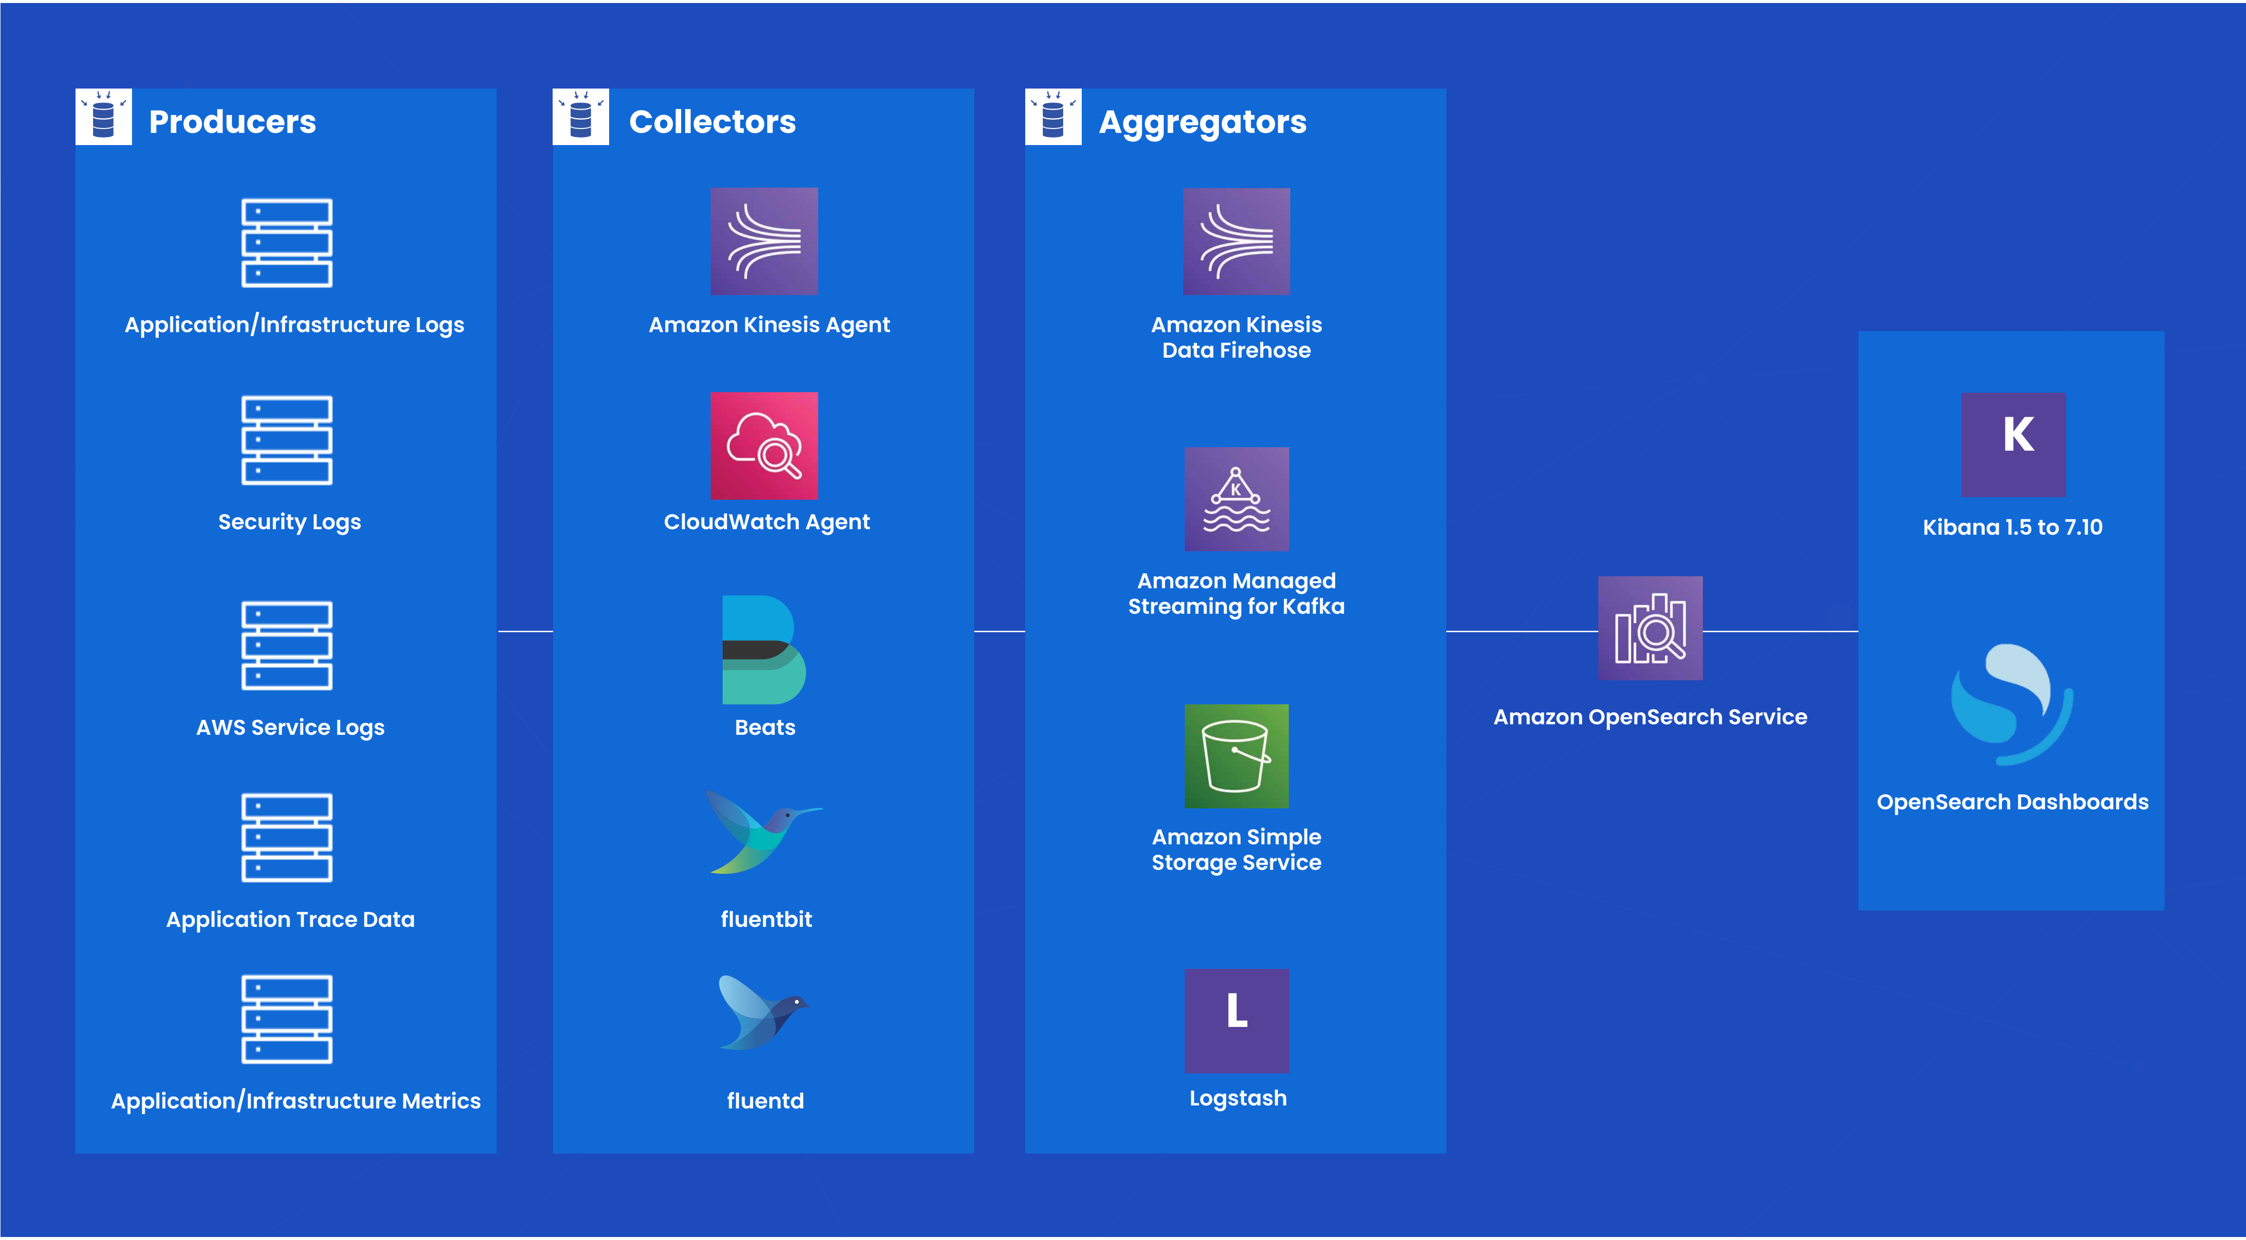

Moving from self-managing open-source solutions to Amazon OpenSearch Service

- Managing and scaling requires dedicated expertise, driving up the total cost of ownership
- Customers need to build or pay for advanced security, alerting, and other features
- Customers need to purchase and manage their own infrastructure

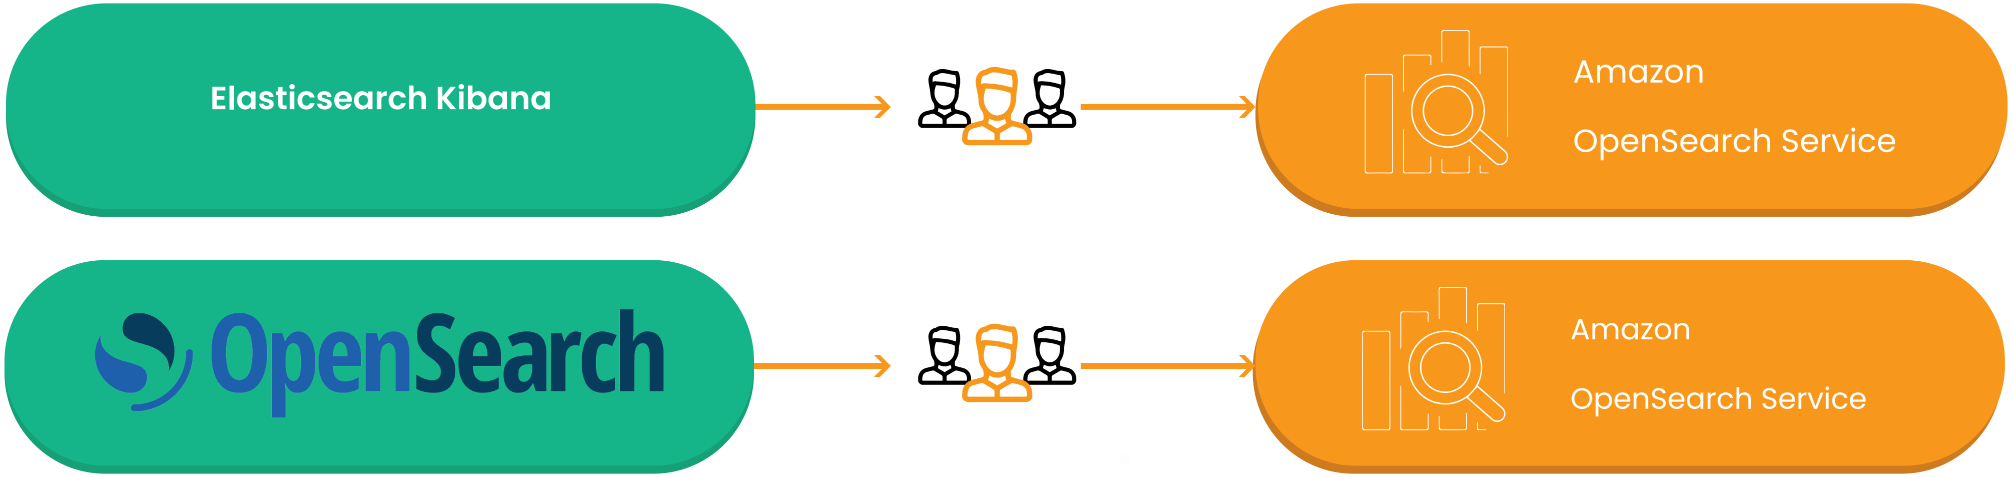

Self-managed vs. Amazon OpenSearch Service

Self-Managed		Managed Service
On-Premises	Amazon EC2	Amazon OpenSearch Service
App dev/optimization	App dev/optimization	App dev/optimization
Hot/warm storage tiers	Hot/warm storage tiers	Hot/warm storage tiers
Plugins (additional cost)*	Plugins (additional cost)*	Plugins
24x7 monitoring & repair	24x7 monitoring & repair	24x7 monitoring & repair
In-place upgrades/patches	In-place upgrades/patches	In-place upgrades/patches
Cluster scaling	Cluster scaling	Cluster scaling
Cross-AZ data transfer cost	Cross-AZ data transfer cost	Cross-AZ data transfer cost
Backups	Backups	Backups
High availability	High availability	High availability
Security (FGAC, Auth)	Security (FGAC, Auth)	Security (FGAC, Auth)
Hardware & OS maintenance	Hardware & OS maintenance	Hardware & OS maintenance
Hardware lifecycle	Hardware lifecycle	Hardware lifecycle
Power/network/HVAC	Power/network/HVAC	Power/network/HVAC

*** SQL querying, Real-time Alerting, Index State Management, Anomaly Detection, Machine Learning**

Moving from licensed solutions to Amazon OpenSearch Service

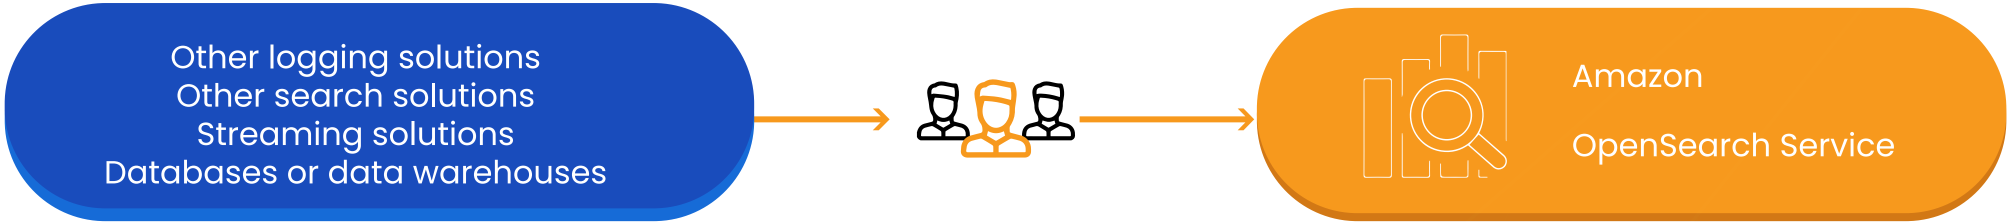

Other logging solutions
Other search solutions
Streaming solutions
Databases or data warehouses

Amazon

OpenSearch Service

- Other, more packaged solutions can drive excessive cost as data volumes grow
- Database solutions and some packaged solutions have lower limits on capacity and higher latency
- Amazon OpenSearch Service is a very flexible tool, supporting search—for application data, but also for logging data. This enables many customers to use Amazon OpenSearch Service for issue debugging and repair

Gain Security analytics with Amazon OpenSearch Service

Amazon OpenSearch Service recently introduced a security analytics feature. The security analytics feature helps organizations with threat monitoring, detection, and alerting features. It allows them to take a proactive approach to threats and neutralize them before they disrupt business operations.

With the OpenSearch Service security features, even users with no prior security experience can investigate security incidents.

Get the most out of Log management and Analytics!

To get the most out of log management and analytics, maintain the log management best practices below:

- Develop a log management policy to ensure everyone follows the same set of protocols for consistency.
- Centralize logs to achieve an end-to-end experience, eliminate data silos, and reduce duplicative efforts.
- Use structured logging to cut down on time spent parsing and simplify analysis.
- Create meaningful log messages that convey information clearly. Include context, timestamps, and unique identifiers.

By following these log management and analytics best practices, organizations can maximize the value of log data as a strategic asset.

Improve Operational analytics with Amazon OpenSearch Serverless

Amazon OpenSearch Service offers a powerful, flexible, and scalable solution for real-time managing and analyzing streaming data. With advanced features, including machine learning capabilities, security analytics, and centralized log management, OpenSearch can help you gain deep insights into your data and detect potential issues faster than ever. By operationalizing OpenSearch with the leading contributor of community-driven, open-source software, you can take full advantage of its capabilities and drive more excellent value for your organization.

If you would like to harness the power of operational analytics with log management and analytics using Amazon OpenSearch to drive your business decisions, get in touch with us today.

Join our Happy Customers Club!

AMGEN

Genmab

Thermo Fisher SCIENTIFIC

neogene THERAPEUTICS

ResMed

ucb

Genentech
A Member of the Roche Group

NBCUniversal

RELIANCE
STEEL & ALUMINUM CO.

IGT

UNIVERSAL
UNIVERSAL MUSIC GROUP

PRIMO WATER

WARNER BROS. PICTURES
WARNER BROS.

About Agilisium

Agilisium is the fastest-growing Cloud Transformation & Data Analytics company with strong expertise in Data lake solutions, Data Warehouse Engineering, Data Migration & Modernization, Data Visualization, and Cloud Optimization services. Agilisium is an AWS Advanced Consulting Partner who helps companies architect, build, migrate, and manage their application workloads to accelerate their journey to the agile cloud, achieve desired business outcomes, and reach new emerging global markets. Learn More at www.agilisium.com.

aws
PARTNER
Advanced Tier
Services

sales@agilisium.com