



Batch and real-time operational analytics on a modern data architecture

Operational Analytics Pipeline on AWS



Operational Analytics Pipeline on AWS: Batch and real-time operational analytics on a modern data architecture

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Operational analytics pipeline i

Operational Analytics Pipeline on AWS 1

Further reading 2

Diagram history 2


Operational Analytics Pipeline on AWS

Batch and real-time operational analytics on a modern data architecture

Publication date: 2021 ([Diagram history](#))

With this architecture, you can perform operational analytics in batch and real time. You use log information from operational data sources such as [AWS CloudTrail](#) and Amazon VPC flow logs. You visualize operational insights in an OpenSearch dashboard and distribute alerts through [Amazon Simple Notification Service](#) (Amazon SNS).

Operational Analytics Pipeline on AWS



The following steps describe the architecture:

1. Logs stream to [Amazon CloudWatch](#) from various data sources. Example sources include AWS CloudTrail and Amazon VPC flow logs.
2. [Amazon Managed Service for Apache Flink](#) receives logs intended for real-time analysis.
3. An [Amazon S3](#) bucket receives logs intended for batch processing. This bucket serves as a backup for storing the raw logs.
4. Amazon CloudWatch subscription filters stream log data to [Amazon Data Firehose](#) for distribution.

5. Amazon Managed Service for Apache Flink derives metrics and sends them to a [AWS Lambda](#) function for processing. The Lambda function transforms the metrics and ingests them into [Amazon OpenSearch Service](#). If a condition requires alerting, the Lambda function publishes a message to Amazon SNS for distribution.
6. [AWS Glue](#) collects the logs from Amazon S3 and performs bulk processing, transformation, and analysis. AWS Glue uses the Amazon OpenSearch Service connector to ingest the results into Amazon OpenSearch Service.
7. An OpenSearch dashboard visualizes the metrics from Amazon OpenSearch Service.
8. Amazon SNS notifies consumers of alerts by publishing to email and SMS notification topics.

Further reading

For additional information, see the following resources:

- [AWS Architecture Icons](#)
- [AWS Architecture Center](#)
- [AWS Well-Architected](#)

Diagram history

To be notified about updates to this reference architecture diagram, subscribe to the RSS feed.

Change	Description	Date
Initial publication	Reference architecture diagram first published.	January 1, 2021

RSS subscription

To subscribe to RSS updates, you must have an RSS plugin enabled for the browser that you are using.