

From Outages to Autonomy: Operating OpenSearch for High-Volume Telemetry

Himanshu Ramchandani, Sonos





Himanshu Ramchandani

Principal Data Engineer, Sonos
OpenSearch Ambassador



Agenda

01

About Sonos & Logging Scale

02

Sonos OpenSearch Use Case

03

OpenSearch Implementation

04

Challenges & Learnings

05

Q&A




About Sonos

Sonos is a leading audio technology company known for its premium wireless home sound systems which offers high-quality audio and seamless multi-room connectivity.



Our Data



Debug product and platform issues quickly



Validate product changes before release



Monitor reliability across connected devices and services



Before OpenSearch...

- Design Complexity
- Service Availability
- Operating Cost
- Customer Support



Sonos OpenSearch Use Case


- Real-Time Log Analysis
- Simplified & Efficient Debugging
- Better Service Availability
- Operating Cost
- Customer Support

aws

 OpenSearch



OpenSearch Service Architecture




This design prevents data loss during traffic spikes and allows independent scaling of ingestion and storage tiers.




Index Templates Management

- Index Patterns, Index Templates & Alias
- Using Alias to write to an Index
- Creating New Index Using ISM policies
- Index Patterns used for searching for relevant logs/events
- Index Templates are created and updated using Sonos Metadata catalog
- Automation run seamlessly through Github Actions



ISM with Rollover Actions

- Use of Hot and Ultra Warm tier for Index State Management
- Each Index gets assigned 1 shard.
- An Index is rolled over to a new write index if
 - Size of the index hits 50 gb **OR**
 - Age of Index is 30 days or more
- Rollover Actions helps in
 - Saving manual efforts of writing & managing scripts for shard management.
 - Data loss prevention.
 - Preserving cluster health.
 - Preventing over-provisioning of shards.





In above example: **firmware-usage-000001** is the initial write index. Once it hits 50gb, **firmware-usage-000002** is created, and all incoming data gets written to this new index. Both indexes roll up to the same alias and are both searchable under the same index pattern.




OpenSearch In Action

- Quick Discovery & Debugging
- Visualizations
- Critical Dashboarding
- Service Readiness



OpenSearch Cluster and Pipeline Observability

- 85+ unique OpenSearch metrics tracked on AWS Cloudwatch
- Real-time visualizations for rapid diagnostics.
- Machine learning-based anomaly detectors and alerts to provide early warnings of any pattern shifts in incoming logs.
- Custom metrics tailored to monitor data-specific values.
- Rapid root cause identification
- Streamlined debugging and triage
- Reduced time to resolution



Challenges & Learning

- OS Cluster and Pipeline Observability
- Lambda Concurrency
- POSTing data on Alias vs Index
- AWS Documentation - Read the Fine Print
- Grok Pipelines



Next in OpenSearch



OpenSearch Dashboard Alerting



Apache Iceberg Onboarding



OpenSearch Serverless



Connect with me



Thank you

Q&A

Scan to leave feedback



From Outages to Autonomy: Operating OpenSearch for High-Volume Telemetry

