International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 03 | March - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Logging Aggregation Platform using Cloud Watch

Monalisa kaur, Labony Bhunia, Sruthi Sethumadhavan

Abstract

The Logging Aggregation Platform using Cloud Watch is a scalable solution that leverages multiple AWS services to implement a scalable centralized logging solution.

Executive summary

The Logging Aggregation Platform using Cloud Watch enables organizations to collect, store, monitor, and analyze logs from multiple AWS services and applications in a unified system. By aggregating logs across infrastructure components such as EC2 instances, serverless workloads, containers, and managed services, CloudWatch provides unified observability and operational visibility.

Challenges

AWS Security Lake centralizes security data collection from multiple AWS services including CloudTrail, EKS audit logs, Route 53 resolver logs, Security Hub findings, VPC Flow Logs, and WAF logs, while also supporting custom data sources. To handle sensitive data, it can be integrated with Amazon Comprehend and Macie for automatic PHI/PII identification and masking. This causes below challenges :

Data Protection Limitations AWS Security Lake lacks built-in data protection features, requiring integration with additional AWS services such as Amazon Comprehend or Amazon Macie to identify and mask sensitive data. This dependency on multiple services significantly increases operational costs and adds complexity to the overall security architecture.

Storage Format Constraints Security Lake exclusively stores logs in Parquet format, which limits flexibility in data storage options. This format restriction may not align with existing data processing workflows or downstream analytics tools that require different file formats, potentially creating integration challenges.

Encryption-Related Access Complexity When Security Lake is enabled, it automatically creates encrypted S3 buckets for log storage. While encryption enhances security, it introduces additional complexity when accessing the stored data, requiring proper key management and potentially impacting performance for data retrieval and analysis operations.

Proposed solution

A centralized logging aggregation platform built on Amazon CloudWatch enables organizations to collect, monitor, analyze, and retain logs from multiple applications, servers, and AWS services in a single unified system.

It improves observability, accelerates troubleshooting, enhances security monitoring, and ensures compliance. A key capability of this approach is the integration of data protection policies within CloudWatch Logs. These policies automatically identify and mask sensitive data—such as personally identifiable information (PII), financial information, or credentials—before logs are stored or accessed. This helps organizations enforce security and compliance requirements while still allowing teams to leverage log analytics for troubleshooting, auditing, and operational monitoring.

Pre-requisites

Telemetry config for AWS Organization needs to be turned on.

To turn on telemetry config for the AWS Organization either a management account or a delegated administrator account must be used. CloudWatch will use this account to discover the Organization's AWS resources and their telemetry configurations.

Before turning on telemetry auditing for the Organization, trusted access between AWS Organizations and CloudWatch needs to be enabled.

To use a member account as delegated administrator account, register a member account as Delegated Administrator account in MPA account under **Organizational settings management.**

Quick Setup Overview

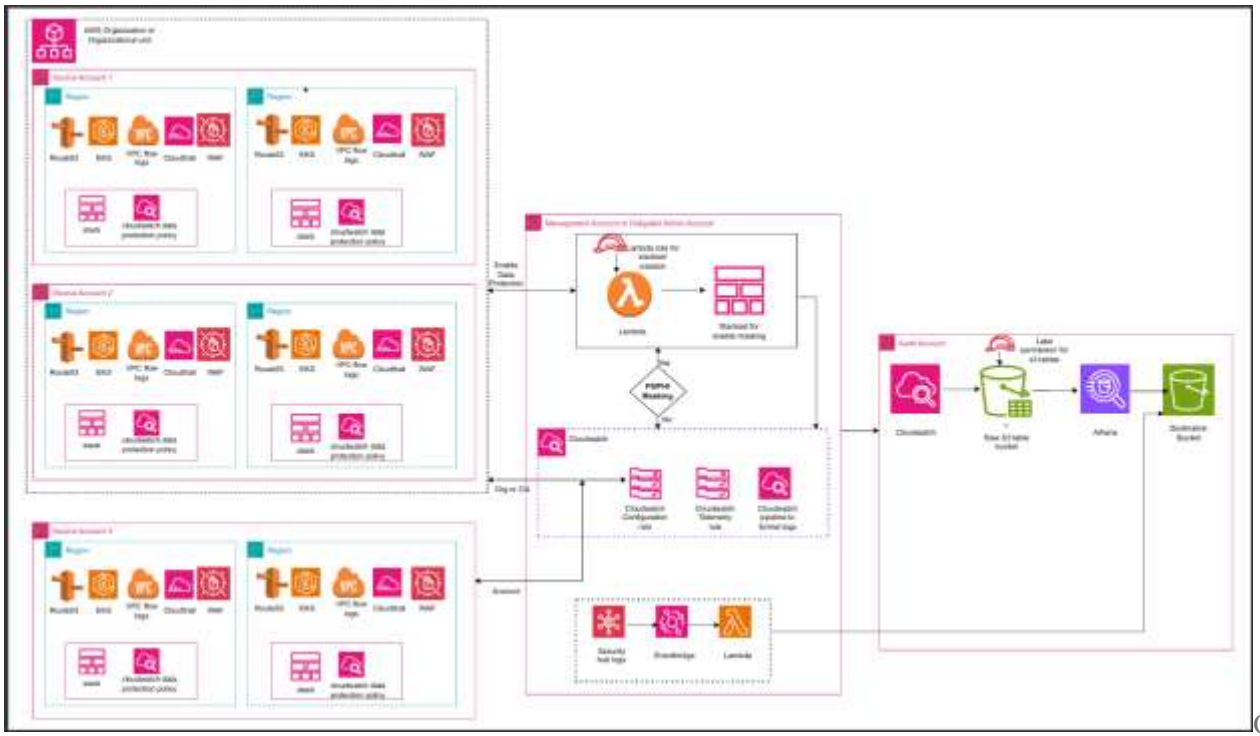

Overall Architecture Diagram

1. Deployment of code will create cloud formation stack in MPA or Delegated admin account to enable cloud watch configuration rule.
2. Enabling Cloud watch telemetry rule is optional hence if selected then stack will create the telemetry rule in MPA/Delegated Admin account.
3. Logs can be aggregated for the resources below from the source accounts :
   - Route53
   - EKS Audit Logs
   - VPC Flow Logs
   - Cloud Trail Logs
   - WAF
4. As part of the centralization rule creation source and destinations can be configured.
5. Source can be Organization or OU or Account and regions.
6. Destination can be Audit account and specific region where the centralized logs will be stored.
7. In the destination account, S3 Table Integration can be enabled under Settings → Account → Integration details.
8. This integration automatically creates the S3 Table bucket under the S3 service.
9. Once the centralization rule is enabled, the configured data sources begin sending logs to the destination (audit) account.
10. Each data source is attached to S3 Tables, and a dedicated S3 bucket is created for each attached resource. After attachment, logs from each data source become visible in S3 Tables.
11. **Integration with AWS analytics services** enables querying S3 Table data through Athena in Amazon S3.
12. A Lake Formation role is created to enable Athena to query the logs table.

13. Logs are available viewable in CSV format and are queried through Athena.
14. A destination bucket is created to store the logs queried by Athena in CSV format, organized under partitioned folders.
15. The Cloud Watch data protection policy can be enabled in source accounts for each of the resources to protect from exposing any sensitive data.
16. Security Hub is not directly supported by CloudWatch, so findings are collected from the Security Hub account (delegated admin/MPA) where all member findings are aggregated.
17. After findings are received, Event Bridge triggers a Lambda which processes events and stores them in the central S3 destination bucket.

Use Cases

1. Multi-Account Observability

- **Unified monitoring** across multiple AWS accounts within a single Region
- **Cross-account troubleshooting** of applications that span different accounts
- **Centralized visibility** without needing to switch between accounts or assume roles
- Support for up to 100,000 source accounts per monitoring account

2. Cross-Region Monitoring

- **Global infrastructure oversight** with dashboards containing metrics and alarms from multiple regions
- **Cross-region log centralization** for comprehensive data collection
- **Unified console experience** allowing you to view resources across different regions from a single interface

3. Enterprise-Scale Log Management

- **Centralized log consolidation** from multiple member accounts into one data repository
- **Automated log replication** using centralization rules
- **Compliance and audit support** with a single source of truth for all log data
- **Cost optimization** through centralized log storage and management

4. Comprehensive Telemetry Aggregation

CloudWatch centralization supports multiple data types:

- **Metrics** from all AWS services and custom applications
- **Log groups** with filtering capabilities for specific subsets
- **Traces** from AWS X-Ray for distributed application monitoring
- **Application Signals** services and Service Level Objectives (SLOs)
- **Application Insights** applications
- **Internet Monitor** data

5. Organizational Governance

- **AWS Organizations integration** for automatic onboarding of new accounts
- **Consistent monitoring standards** across all environments (Dev, QA, Prod)
- **Centralized alerting and incident response**

- **Security and compliance monitoring** across the entire infrastructure

6. Operational Efficiency

- **Single pane of glass** for all observability data
- **Faster incident detection and resolution**
- **Reduced operational overhead** by eliminating custom log aggregation solutions
- **Streamlined troubleshooting** without account boundaries

7. Cost Management

- **No additional charges** for shared logs and metrics in cross-account observability as compared to Security Lake which includes below cost structure :
    - S3-based storage pricing for the data lake
    - Data ingestion and processing costs
    - Analytics query costs (Athena, OpenSearch)
    - Subscriber access charges may apply
- Standard cloud watch pricing for storage and queries
- **Free first trace copy** for distributed tracing
- **Centralized log storage** at $0.05/GB for additional copies
- **Elimination of custom aggregation infrastructure costs**

8. Security and Compliance

- **Control plane observability** through centralized CloudTrail logging
- **Network observability** with VPC Flow Logs and Traffic Mirroring
- **Workload observability** for distributed applications
- **Encrypted log storage** with proper access controls

Conclusion

In this post, we showed you how to set up cloud watch solutions to centralize logging from multiple source accounts and regions.