AWS re:Invent

DECEMBER 2 – 6, 2024 | LAS VEGAS, NV

SEC235-NEW

Accelerate security analytics across hybrid environments with AWS

Dora Karali

(she/her)
Principal PM
Security Services
AWS

Abhi Khanna

(he/him)
Principal PM
Amazon OpenSearch Service
AWS

Ross Warren

(he/him)
Product Solution Architect
Security Lake
AWS

Agenda

01 Challenges with security data analysis

02 New! Amazon OpenSearch Service zero-ETL integration with Security Lake - benefits

03 Amazon Security Lake and Amazon OpenSearch Service overview

04 Zero-ETL integration features

05 Demo

Challenges with security data analysis

Collecting and managing growing volumes of data from different sources and locations

Logs and alerts in varying formats

Complexity creating and managing data pipelines

Balancing cost-efficient data access with data visibility

Training and usage of multiple analysis tools

Achieving quick mean time to resolution for security issues

Amazon OpenSearch Service zero-ETL integration with Amazon Security Lake

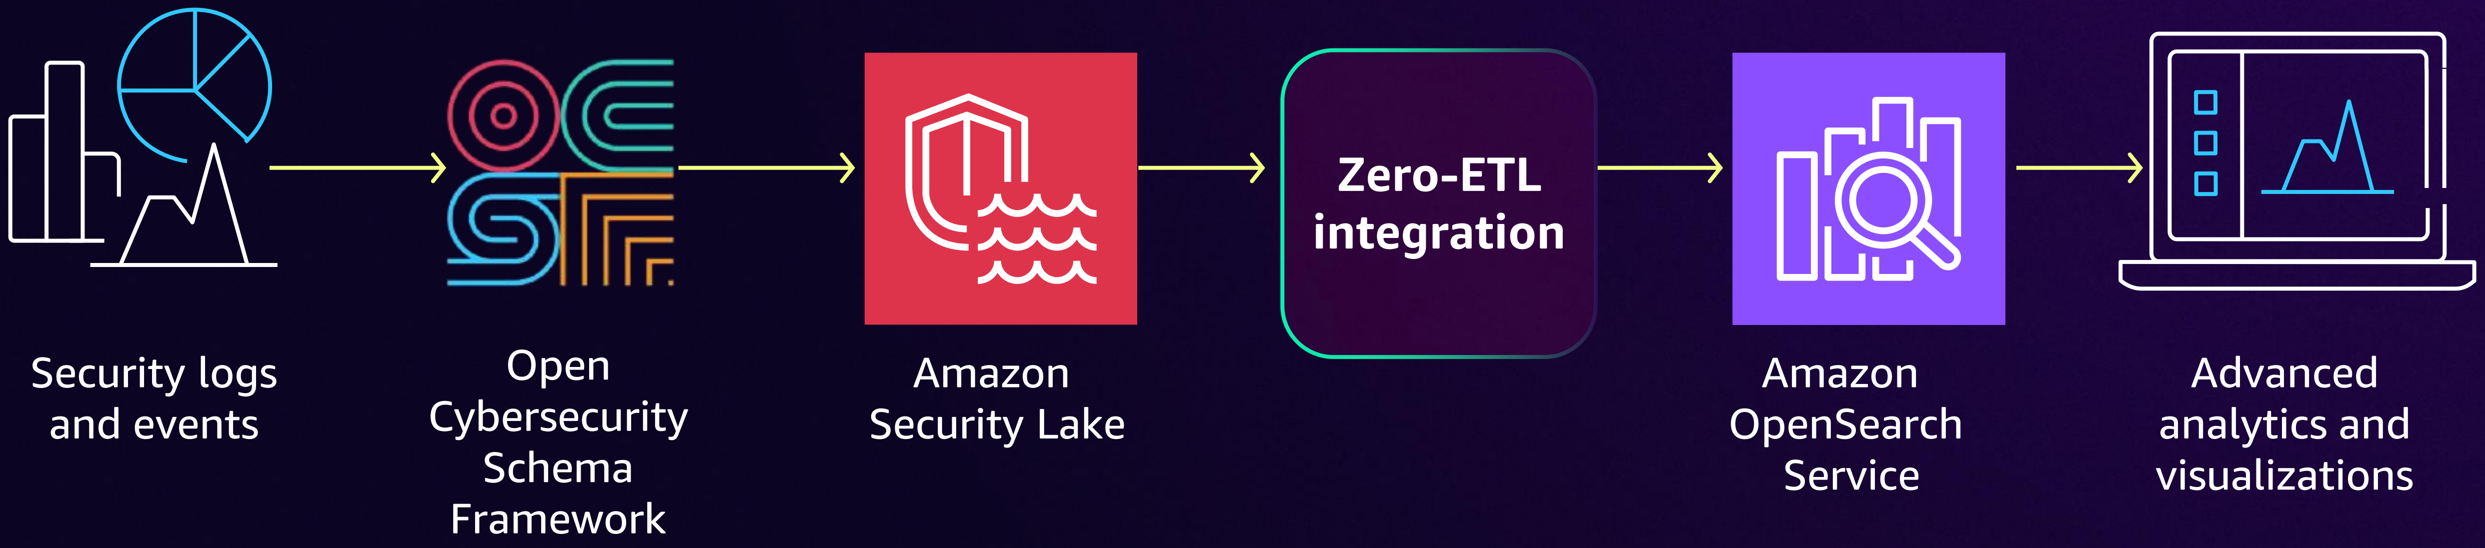

Security logs and events

Open Cybersecurity Schema Framework

Amazon Security Lake

Zero-ETL integration

Amazon OpenSearch Service

Advanced analytics and visualizations

Gain immediate security insights with powerful in-place search, on-demand indexing, and pre-built analytics, eliminating complex data pipelines

Benefits

Security Lake simplifies centralization of data across AWS sources, accounts and regions, and 3rd-party data

Security Lake normalizes data in OCSF schema and prepares data for efficient storage and query access

Directly access your Security Lake data from Amazon OpenSearch with no data pipeline configurations

Full visibility into your Security Lake data with in-place queries in Amazon OpenSearch; reduce ingested data and costs using on-demand indexing

Single tool for real-time and historical data analysis; pre-built queries and dashboards to bring you up and running quickly

Less time on data management so you can focus on resolving security issues

Amazon Security Lake

Automatically centralize security data into a purpose-built data lake

Centralize data automatically from cloud, on-premises, and custom security sources across Regions

Optimize and manage security data for more efficient storage and query performance

Analyze using your preferred analytics tools while retaining control and ownership of your security data

Normalize data to an open standard to easily share and use with multiple analytics tools

Open Cybersecurity Schema Framework (OCSF)

AN OPEN STANDARD THAT CAN BE ADOPTED BY ANYONE TO SIMPLIFY SECURITY DATA NORMALIZATION

NEW! Now part of the Linux Foundation

Open source project to deliver a simplified and vendor-agnostic taxonomy for security data

Speed data ingestion and analysis without the time-consuming, up-front normalization tasks

Combine data from OCSF-compliant sources to break down data silos that slow security teams

Over 200 participating organizations across security ISVs, government, education, and enterprise, with many more using OCSF

Amazon OpenSearch Service

Advanced analytical capabilities to query and analyze security data with powerful visualization and monitoring capabilities

Search: Query at scale to find relevant security events within seconds

Analytics: Securely, easily, and efficiently visualize and analyze your security data

Lower incident response time: Quickly and easily connect all of your data for faster queries and better insights

Alerts: Send security alerts to preconfigured destinations using automated workflows

Zero-ETL with Security Lake features

Quick setup

In-place querying of Security Lake data

On-demand indexing

Pre-built queries and dashboards

Simple setup

1. Create a **subscriber** in Amazon Security Lake
2. Create a **data source for Security Lake** in Amazon OpenSearch Service

Automatically create Amazon OpenSearch Serverless collection and a Dashboards application

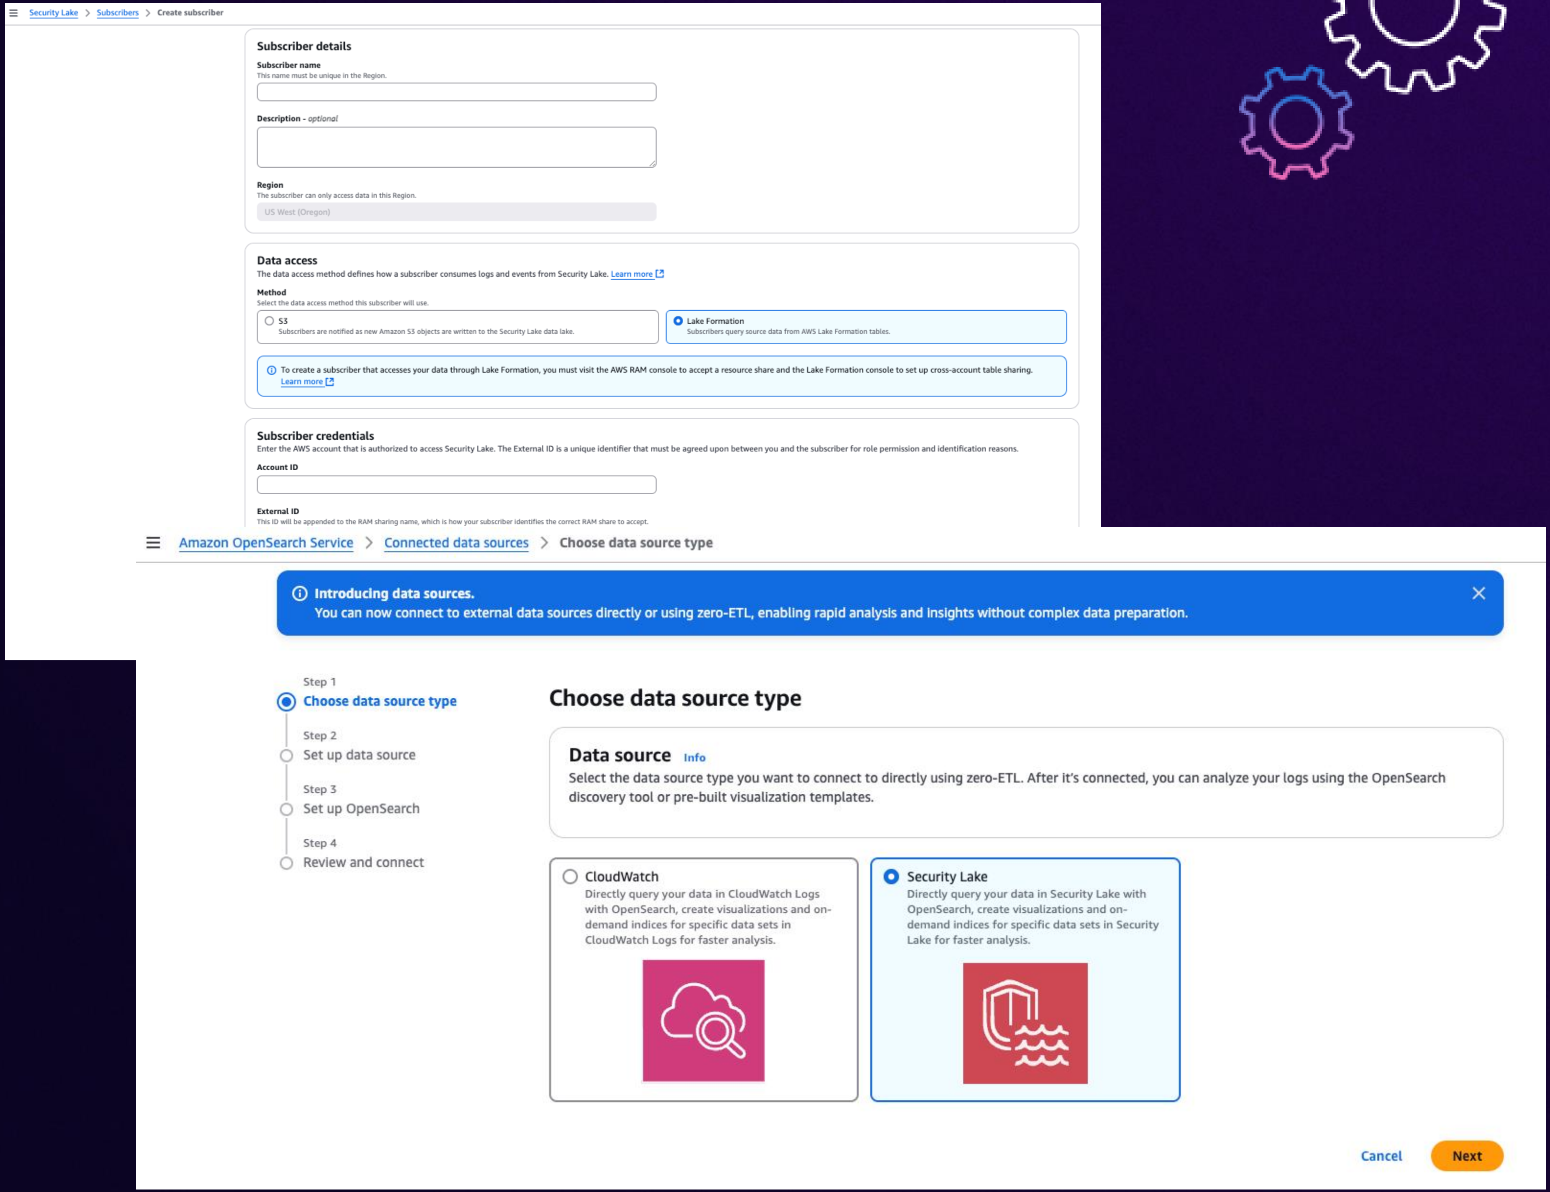

Security Lake > Subscribers > Create subscriber

Subscriber details

Subscriber name
This name must be unique in the Region.

Description - *optional*

Region
The subscriber can only access data in this Region.
US West (Oregon)

Data access
The data access method defines how a subscriber consumes logs and events from Security Lake. Learn more

Method
Select the data access method this subscriber will use.

○ S3
Subscribers are notified as new Amazon S3 objects are written to the Security Lake data lake.

● Lake Formation
Subscribers query source data from AWS Lake Formation tables.

ⓘ To create a subscriber that accesses your data through Lake Formation, you must visit the AWS RAM console to accept a resource share and the Lake Formation console to set up cross-account table sharing. Learn more

Subscriber credentials
Enter the AWS account that is authorized to access Security Lake. The External ID is a unique identifier that must be agreed upon between you and the subscriber for role permission and identification reasons.

Account ID

External ID
This ID will be appended to the RAM sharing name, which is how your subscriber identifies the correct RAM share to accept.

Amazon OpenSearch Service > Connected data sources > Choose data source type

ⓘ **Introducing data sources.**
You can now connect to external data sources directly or using zero-ETL, enabling rapid analysis and insights without complex data preparation.

Step 1
● Choose data source type

Step 2
○ Set up data source

Step 3
○ Set up OpenSearch

Step 4
○ Review and connect

Choose data source type

Data source Info
Select the data source type you want to connect to directly using zero-ETL. After it's connected, you can analyze your logs using the OpenSearch discovery tool or pre-built visualization templates.

○ CloudWatch
Directly query your data in CloudWatch Logs with OpenSearch, create visualizations and on-demand indices for specific data sets in CloudWatch Logs for faster analysis.

● Security Lake
Directly query your data in Security Lake with OpenSearch, create visualizations and on-demand indices for specific data sets in Security Lake for faster analysis.

In-place querying of Security Lake data

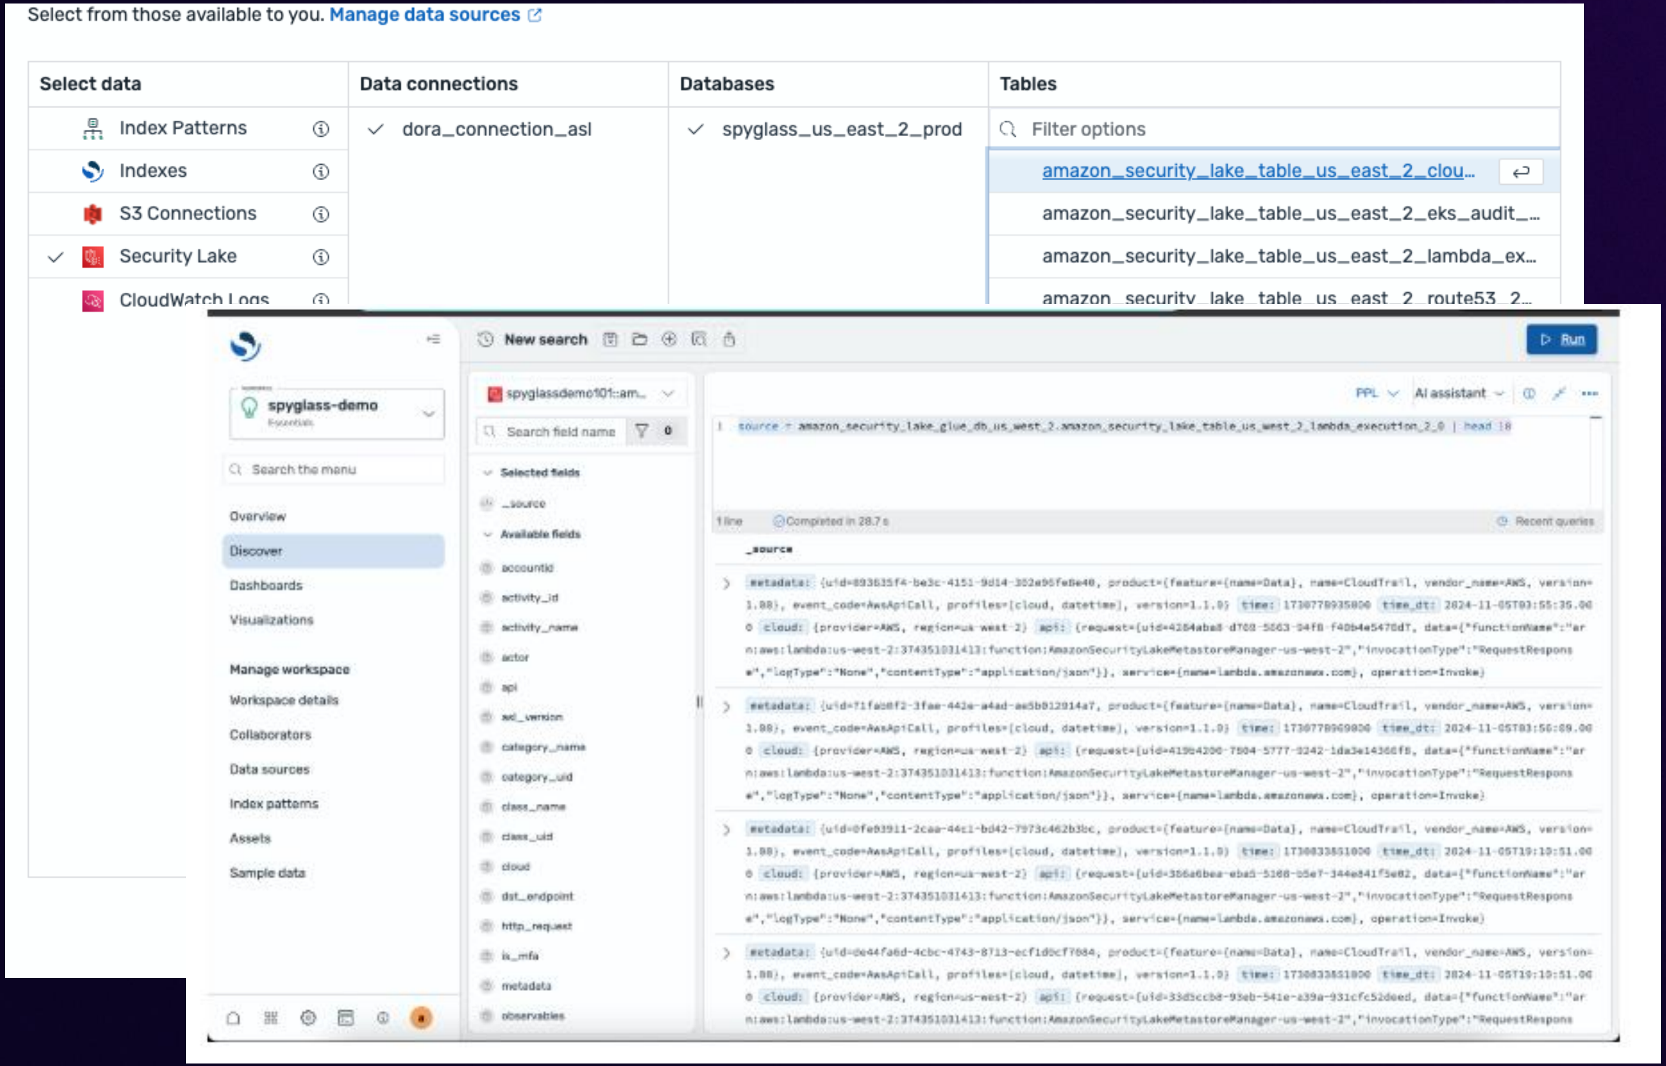

Select from those available to you. Manage data sources

Select data	Data connections	Databases	Tables
Index Patterns	✓ dora_connection_asl	✓ spyglass_us_east_2_prod	Filter options
Indexes			amazon_security_lake_table_us_east_2_clou...
S3 Connections			amazon_security_lake_table_us_east_2_eks_audit_...
✓ Security Lake			amazon_security_lake_table_us_east_2_lambda_ex...
CloudWatch Logs			amazon_security_lake_table_us_east_2_route53_2...

```
source = amazon_security_lake_glue_db_us_west_2.amazon_security_lake_table_us_west_2_lambda_execution_2_0 | head 10
```

```
_source

> metadata: {uid=893635f4-be3c-4151-9d14-302e95fe8e40, product={feature={name=Data}, name=CloudTrail, vendor_name=AWS, version=1.08}, event_code=AwsApiCall, profiles=[cloud, datetime], version=1.1.0} time: 1730778935000 time_dt: 2024-11-05T03:55:35.000 cloud: {provider=AWS, region=us-west-2} api: {request={uid=4284aba8-d708-5063-94f8-f40b4e5470d7, data={"functionName":"arn:aws:lambda:us-west-2:374351031413:function:AmazonSecurityLakeMetastoreManager-us-west-2","invocationType":"RequestResponse","logType":"None","contentType":"application/json"}}, service={name=lambda.amazonaws.com}, operation=Invoke}

> metadata: {uid=71fab8f2-3fae-442e-a4ad-ae5b012914a7, product={feature={name=Data}, name=CloudTrail, vendor_name=AWS, version=1.08}, event_code=AwsApiCall, profiles=[cloud, datetime], version=1.1.0} time: 1730778969000 time_dt: 2024-11-05T03:56:09.000 cloud: {provider=AWS, region=us-west-2} api: {request={uid=41984200-7804-5777-9242-1da3e14366f8, data={"functionName":"arn:aws:lambda:us-west-2:374351031413:function:AmazonSecurityLakeMetastoreManager-us-west-2","invocationType":"RequestResponse","logType":"None","contentType":"application/json"}}, service={name=lambda.amazonaws.com}, operation=Invoke}

> metadata: {uid=0fe03911-2caa-44c1-bd42-7973c462b3bc, product={feature={name=Data}, name=CloudTrail, vendor_name=AWS, version=1.08}, event_code=AwsApiCall, profiles=[cloud, datetime], version=1.1.0} time: 1730833851000 time_dt: 2024-11-05T19:10:51.000 cloud: {provider=AWS, region=us-west-2} api: {request={uid=386a6bea-eba5-5108-b5e7-344e841f5e82, data={"functionName":"arn:aws:lambda:us-west-2:374351031413:function:AmazonSecurityLakeMetastoreManager-us-west-2","invocationType":"RequestResponse","logType":"None","contentType":"application/json"}}, service={name=lambda.amazonaws.com}, operation=Invoke}

> metadata: {uid=de44fa6d-4cbc-4743-8713-ecf1d0cf7084, product={feature={name=Data}, name=CloudTrail, vendor_name=AWS, version=1.08}, event_code=AwsApiCall, profiles=[cloud, datetime], version=1.1.0} time: 1730833851000 time_dt: 2024-11-05T19:10:51.000 cloud: {provider=AWS, region=us-west-2} api: {request={uid=33d5ccb8-93eb-541e-a39a-931cfc52deed, data={"functionName":"arn:aws:lambda:us-west-2:374351031413:function:AmazonSecurityLakeMetastoreManager-us-west-2","invocationType":"RequestRespons
```

Directly query data in Security Lake without ingesting them in OpenSearch

Faster **query times** thanks to Apache Iceberg

Use **SQL or PPL** to query your data across tables using OCSF schema

On-demand indexing

Single-click indexing of your query results from Discover in OpenSearch Dashboards

Indexed views for:

1. Faster querying to support security investigations
2. Visualizations to support security insights

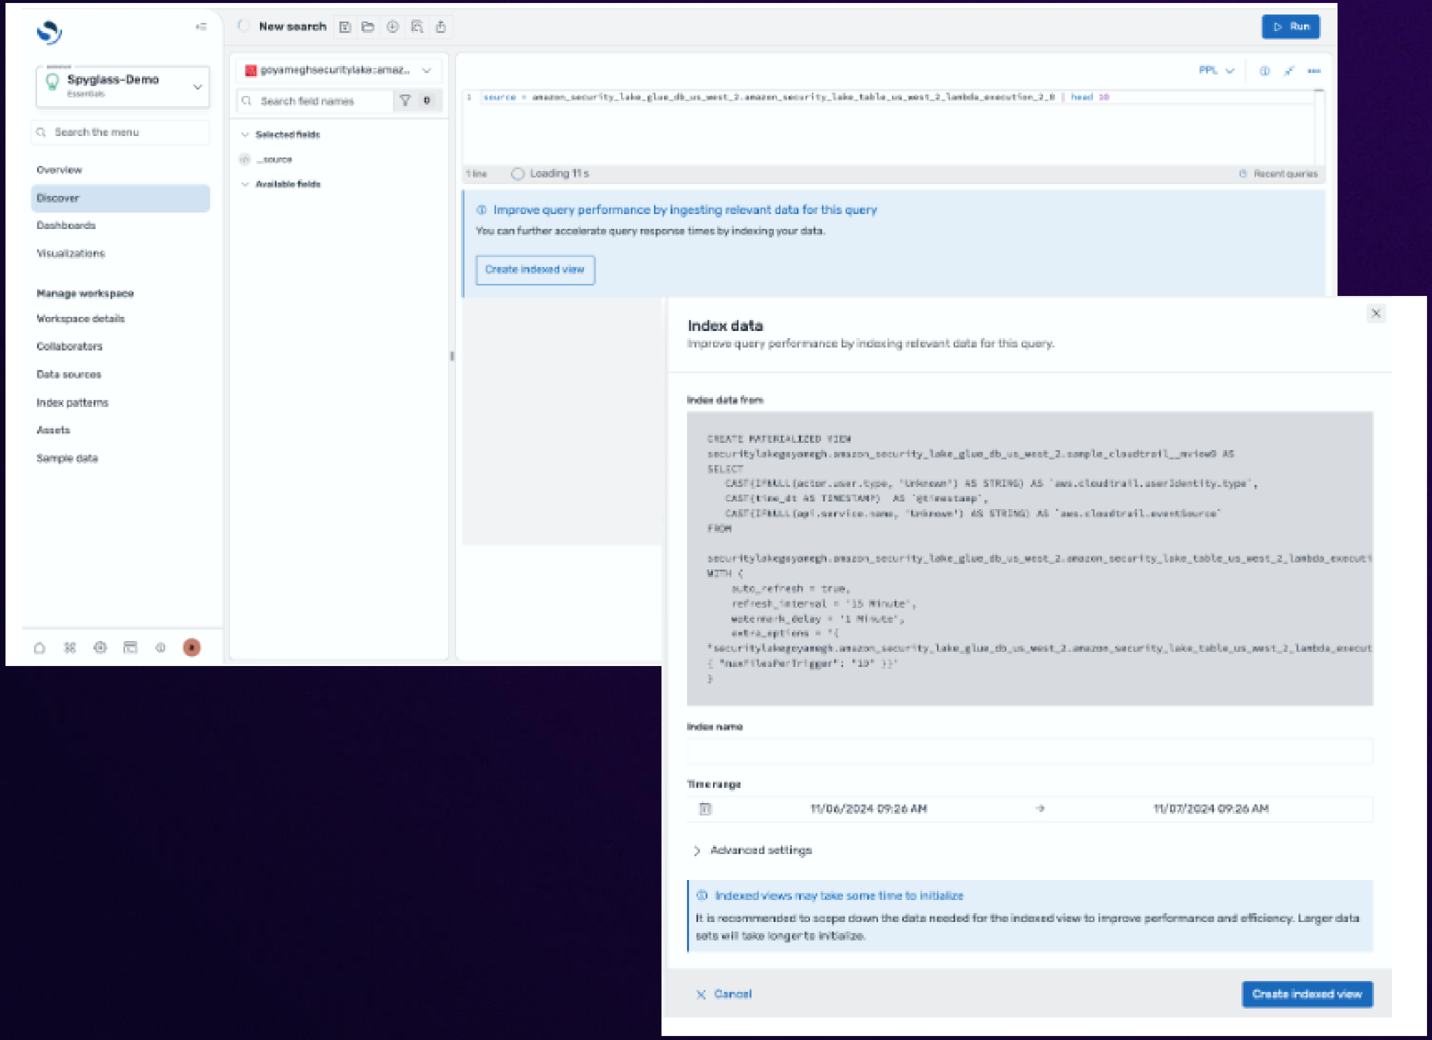

```
source = amazon_security_lake_glue_db_us_west_2.amazon_security_lake_table_us_west_2_lambda_execution_2_0 | head 10
```

Improve query performance by ingesting relevant data for this query

You can further accelerate query response times by indexing your data.

Create indexed view

Index data

Improve query performance by indexing relevant data for this query.

Index data from

```
CREATE MATERIALIZED VIEW
securitylakegoyamegh.amazon_security_lake_glue_db_us_west_2.sample_cloudtrail__mview0 AS
SELECT
   CAST(IFNULL(actor.user.type, 'Unknown') AS STRING) AS `aws.cloudtrail.userIdentity.type`,
   CAST(time_dt AS TIMESTAMP)  AS `@timestamp`,
   CAST(IFNULL(api.service.name, 'Unknown') AS STRING) AS `aws.cloudtrail.eventSource`
FROM

securitylakegoyamegh.amazon_security_lake_glue_db_us_west_2.amazon_security_lake_table_us_west_2_lambda_executi
WITH (
    auto_refresh = true,
    refresh_interval = '15 Minute',
    watermark_delay = '1 Minute',
    extra_options = '{
'securitylakegoyamegh.amazon_security_lake_glue_db_us_west_2.amazon_security_lake_table_us_west_2_lambda_execut
{ "maxFilesPerTrigger": "10" }}'
)
```

Index name

Time range

11/06/2024 09:26 AM → 11/07/2024 09:26 AM

Advanced settings

Indexed views may take some time to initialize

It is recommended to scope down the data needed for the indexed view to improve performance and efficiency. Larger data sets will take longer to initialize.

Pre-built Queries and Dashboards

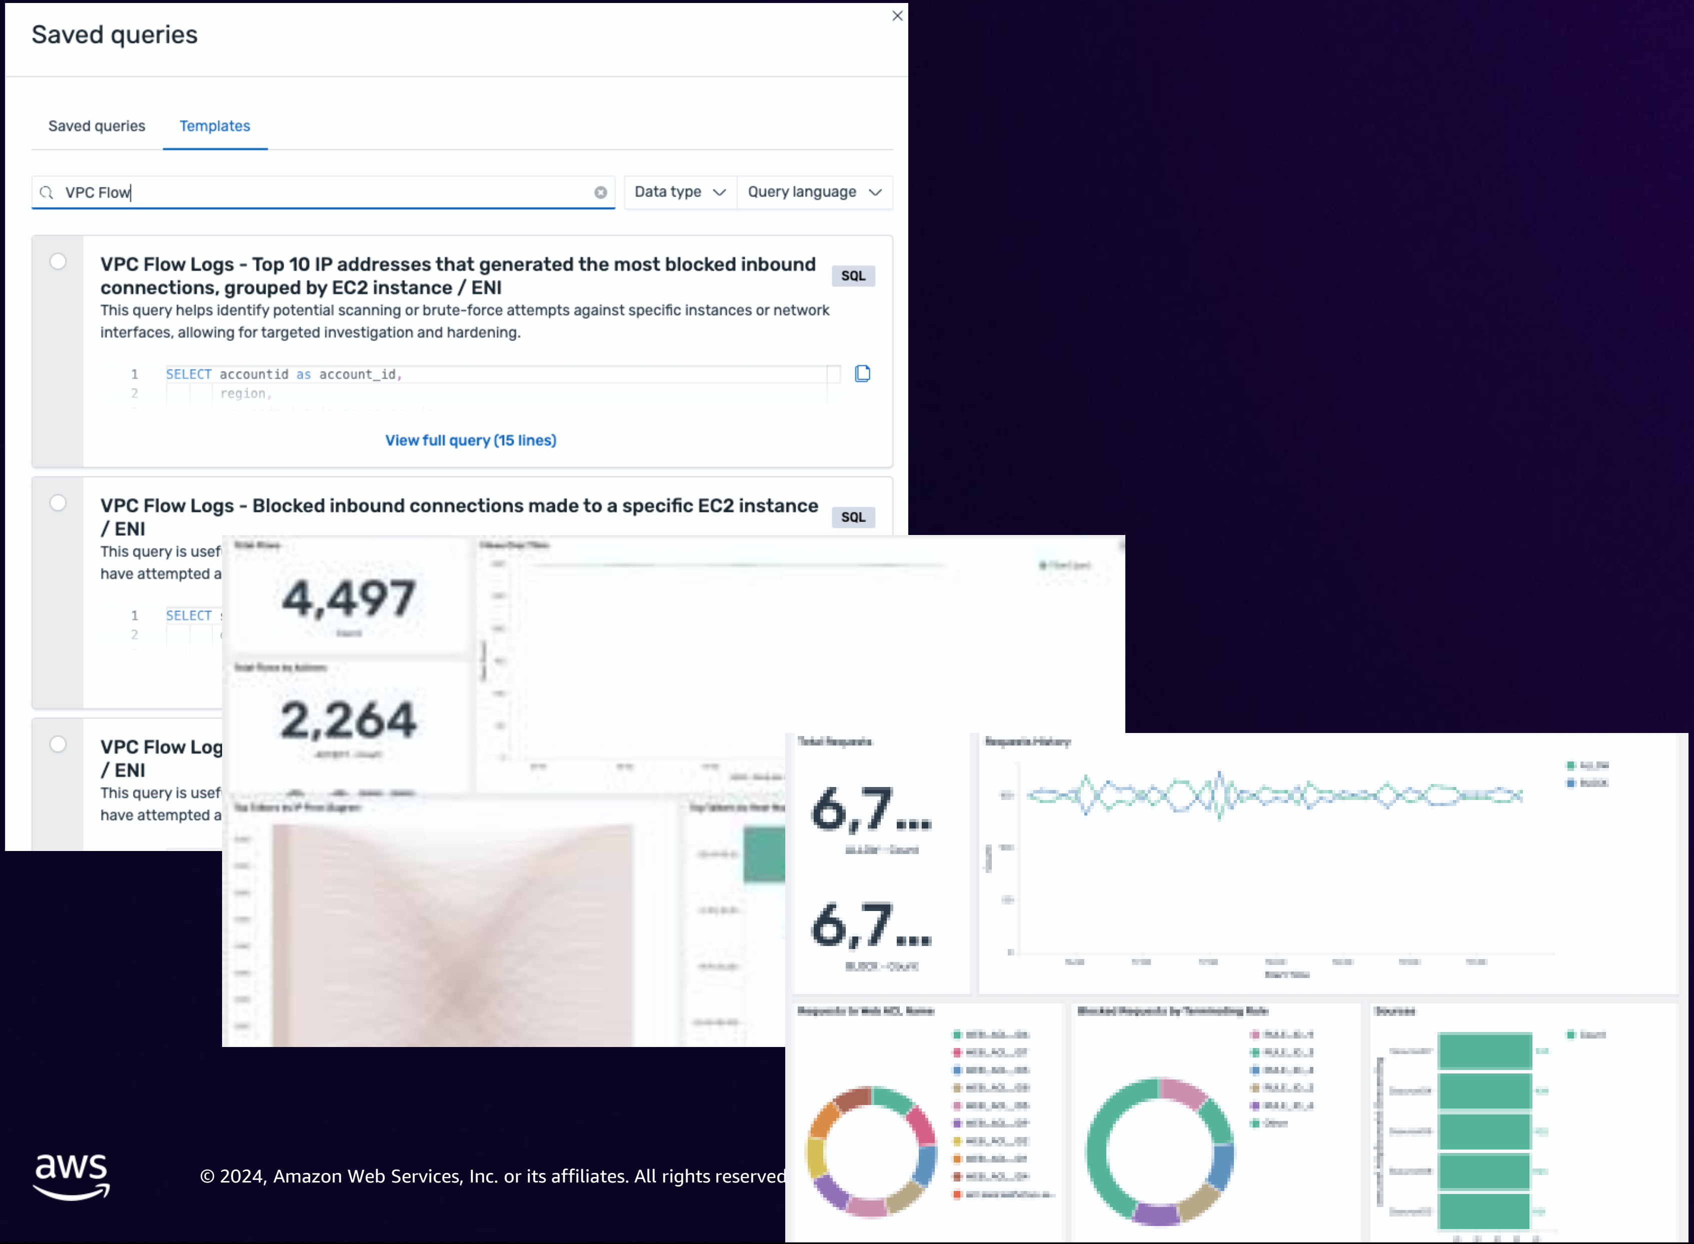

Saved queries

Saved queries | Templates

VPC Flow

Data type | Query language

VPC Flow Logs - Top 10 IP addresses that generated the most blocked inbound connections, grouped by EC2 instance / ENI SQL

This query helps identify potential scanning or brute-force attempts against specific instances or network interfaces, allowing for targeted investigation and hardening.

```
SELECT accountid as account_id,
       region,
```

View full query (15 lines)

VPC Flow Logs - Blocked inbound connections made to a specific EC2 instance / ENI SQL

This query is usef

have attempted a

```
SELECT
```

VPC Flow Log

/ ENI

This query is usef

have attempted a

Use **pre-built** queries and dashboards in OCSF schema for faster insights and on-boarding

200+ pre-built queries

Dashboards for **VPC Flow** Logs, **WAF** logs and **AWS CloudTrail** Management Events

> Managing security data across a complex environment posed challenges. Amazon Security Lake has enabled data sovereignty, improved data visibility, and allowed direct querying without moving data, while ensuring compliance. This solution is expected to expedite incident response and reduce costs, thereby enhancing the protection of clients' data.

Derek Bush

Vice President of Cloud Security, Infor

Demo

Amazon OpenSearch service zero-ETL with Security Lake for:

Simplified Security Data Analytics

Eliminate data duplication and complex ETL processes
Zero-ETL integration allows direct querying of Security Lake data in OpenSearch

Comprehensive Security Visibility

Unified analysis of diverse security data sources
Query and visualize data from AWS using OCSF schema

Accelerated Security Investigations

Faster incident response
Pre-built OCSF-compliant queries, dashboards, and on-demand data acceleration

Optimized Performance and Costs

Balance between query speed and storage efficiency
Flexible options for direct queries, selective indexing, and materialized views

Learn more

What’s New blog

Try Security Lake for 15 days at no cost

Learn more about Amazon OpenSearch Serverless

Learn more about Amazon OpenSearch Service Integrations

Save the date for AWS re:Inforce

JUNE 16 – 18, 2025 | PHILADELPHIA, PA

AWS re:Inforce

Thank you!

Please complete the session survey in the mobile app

Dora Karali

linkedin.com/in/dorakarali/

Abhi Khanna

https://www.linkedin.com/in/abhinavkhanna2016

Ross Warren

linkedin.com/in/rossw7/