How to keep AWS Security Groups Safe using Splunk

Nishant Mehta

Lead DevOps Engineer | Cognizant Technology Solutions

John Ray

Consulting Software Engineer | LexisNexis

Forward-Looking Statements

During the course of this presentation, we may make forward-looking statements regarding future events or plans of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results may differ materially. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, it may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements made herein.

In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only, and shall not be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionalities described or to include any such feature or functionality in a future release.

Agenda

1) **LexisNexis Splunk Journey**

2) **What Problem We are trying to Solve !**

3) **AWS Data Ingestion**
   AWS VPC Flow Logs, AWS Description and Lambda Logs

4) **Data Correlation**
   Correlate all AWS Logs together

5) **Security Group Investigation and Remediation**
   Using Splunk Dashboards and Alerts

6) **Splunk License Cost Savings**
   OnDemand Data Ingestion Model

John Ray

Consulting Software Engineer | LexisNexis

Splunking since 2017 with LexisNexis
Expert in Enterprise Monitoring and ITSM Applications
Focusing on AWS-Splunk Integrations recently!

Nishant Mehta

Lead DevOps Engineer | Cognizant Technology Solutions

Splunk Admin
Working at LexisNexis as Contractor since beginning of LexisNexis Splunk Journey
Loves all things Splunk!

splunk> .conf20

LexisNexis

- LexisNexis Legal & Professional is an information analytics company.
- We combine information, analytics, and technology to help our customers reach essential insights, make more informed decisions, and achieve better outcomes.
- Our mission and purpose is to advance the rule of law around the world.

LexisNexis Data in Numbers

2.50 PB Total Legal Data	119 B Total Documents available for Customers	200 Years of Legal & Judicial data	150 x Size of Wikipedia(doubling every 3 years)
28 M Daily Documents	69,000 Daily Data Sources	70 B Daily Connections	

- What we do:
  - We help lawyers win cases, manage their work more efficiently, serve their clients better and grow their practices.
  - We assist corporations in better understanding their markets, monitoring their brands and competition, and in mitigating business risk.
  - We collaborate with universities to educate students, and we support nation-building with governments and courts by making laws accessible and strengthening legal infrastructures.
- LexisNexis Legal & Professional serves customers in more than 130 countries with 10,000 employees worldwide.
- We are part of RELX, a global provider of information and analytics for professional and business customers across industries.

LexisNexis Splunk Journey

2017

- Splunk Enterprise Instance
- 200 GB/Day Ingestion
- Primarily Collection of Proprietary Application Logs
  - Linux and Windows Universal Forwarders in AWS

Today

- Splunk Cloud Subscription
- Locally Managed Heavy Forwarders and Deployment Server
- 3 TB/Day Ingestion
- Expansion of Data Sources
  - Host/OS, Network and Firewall, Non-Universal Forwarder Sourced Logs
  - AWS Logs for Multiple Accounts
  - CloudTrail, Load Balancers, VPC Flow Logs
  - On Prem (Data Center) Hosts

LexisNexis Splunk Architecture

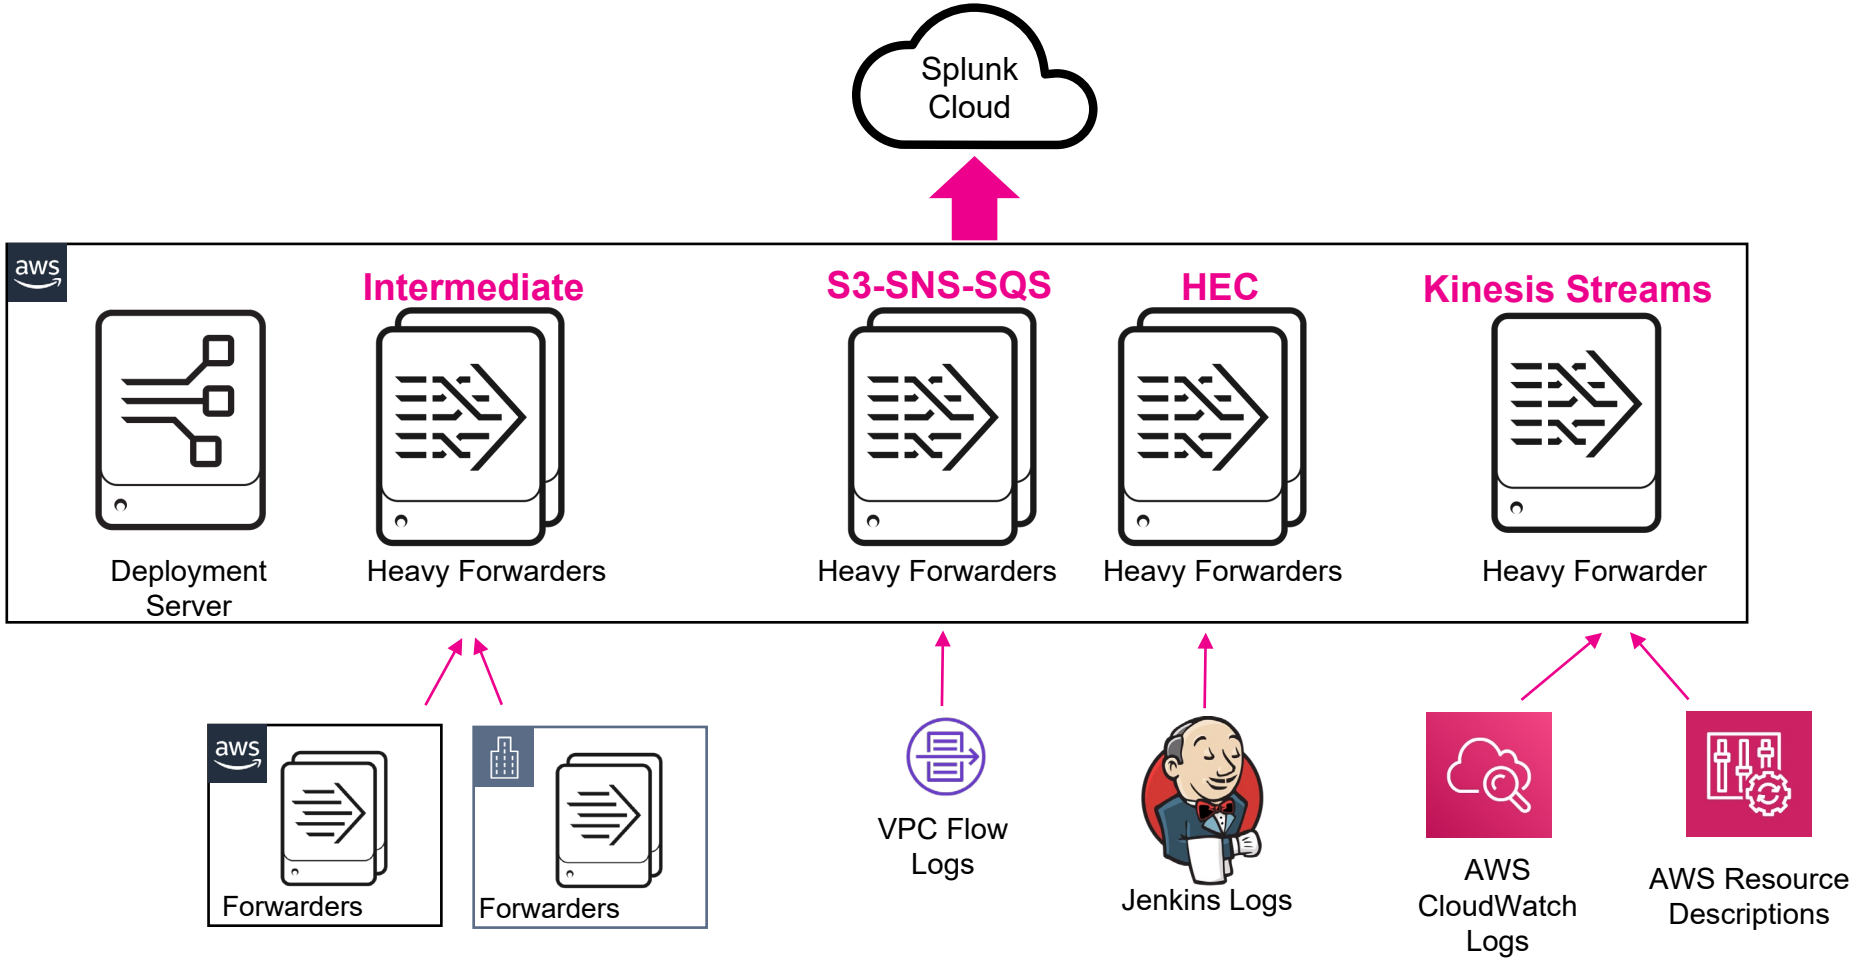

Problem Statement

AWS Security groups are the most important building blocks in AWS cloud deployment. Security Groups act as cloud-based firewalls to protect applications and data for instances within VPCs in both Public and Private subnets. Security groups are a whitelist service which allows customers to expose resources to only whitelisted IP addresses or resources, with access controlled by IP CIDR ranges and ports defined in the security group.

Let's use Splunk to improve AWS Security posture !!!

Our Solution

- Framework comprising of Self-Service Splunk Dashboards and Alert with correlated data from AWS VPC Flow Logs, AWS Description and AWS Lambda Logs so developers can:
  - Audit Security Groups
  - Understand Vulnerabilities
  - Take Corrective Action
  - Monitor Security Group Changes
- Focusing on Monitoring Elastic Load Balancers Security Group Inbound Rules
- Extensible Framework to validate Traffic at Instance Level and Security Group Outbound Rules

AWS Data Ingestion

Different AWS Data Sources Ingestion

Data Correlation

Correlate AWS VPC Flow Logs with SG

SG Investigation and Remediation

Splunk Dashboards and Alerts

License Cost Savings

On-Demand Data Ingestion Model

splunk> .conf20

Data Ingestion

How We are Ingesting different AWS Data Sources

Data Sources

AWS VPC Flow Logs

S3-SNS-SQS Heavy Forwarder Input

AWS Lambda Logs

CloudWatch – Kinesis Input

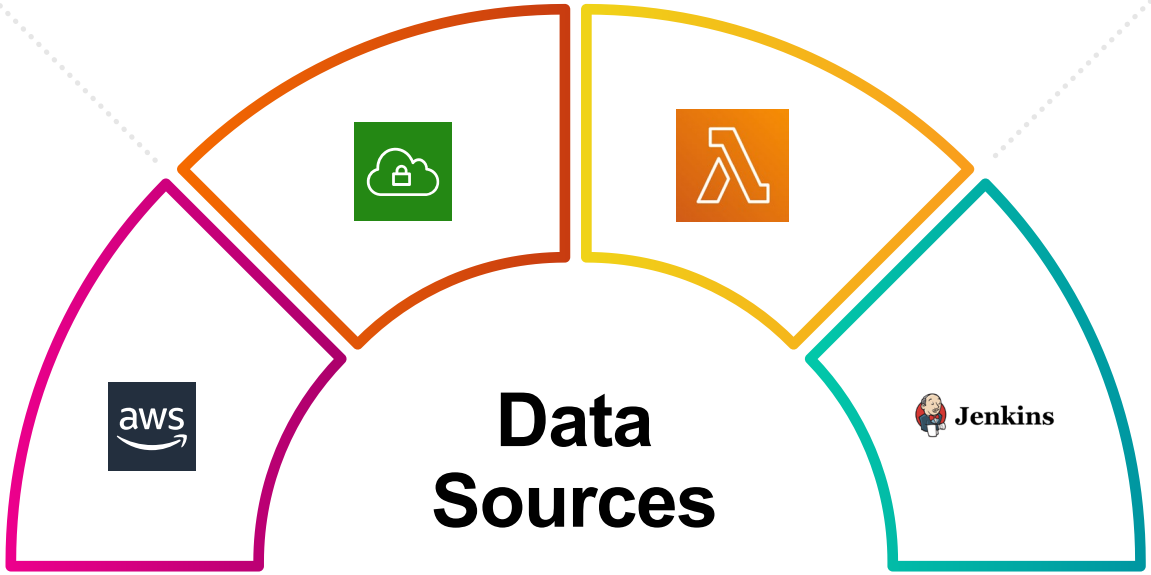

AWS Description

AWS Add-on

Jenkins Pipeline Logs

Splunk Jenkins Plug-in HEC Input

AWS VPC Flow Logs – Data Ingestion

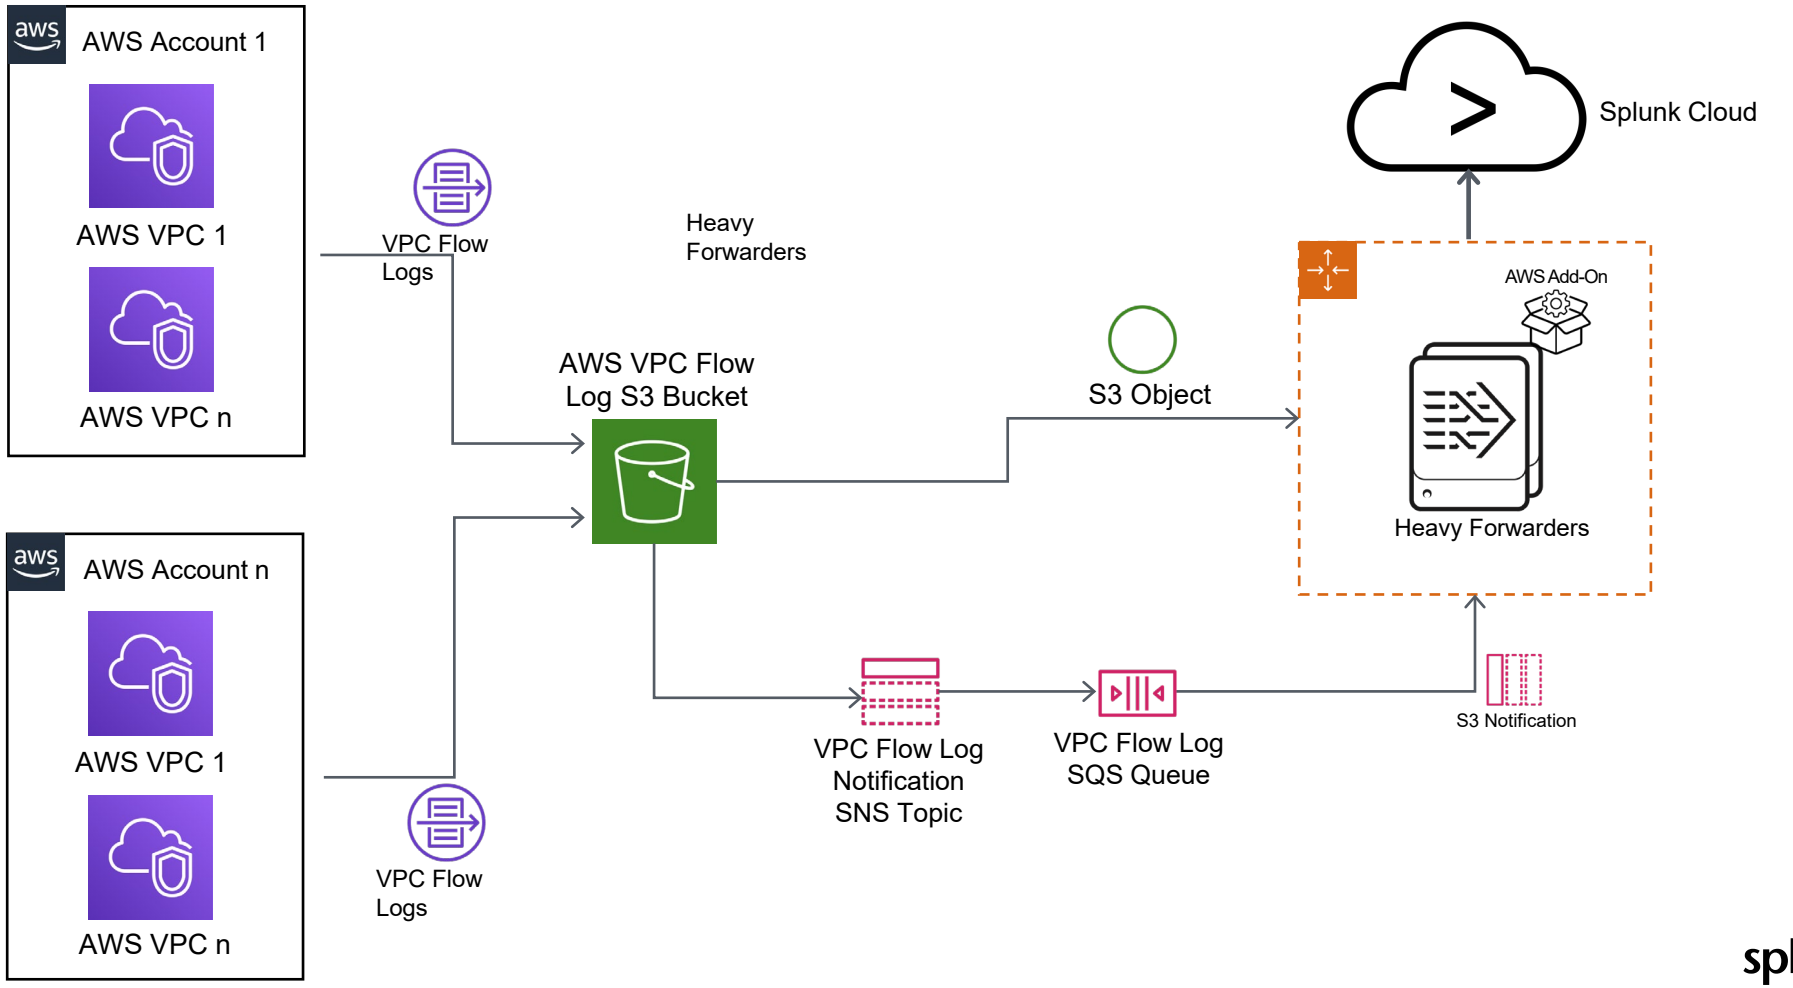

Data Correlation

Building Lookup Tables to correlate SG and VPC Flow Logs

Elastic Network Interface (ENI) – Security Group Mapping

AWS Lambda function invokes describe_network_interfaces operation for all AWS accounts to get ENI and Security Group Mapping Information

https://docs.aws.amazon.com/cli/latest/reference/ec2/describe-network-interfaces.html

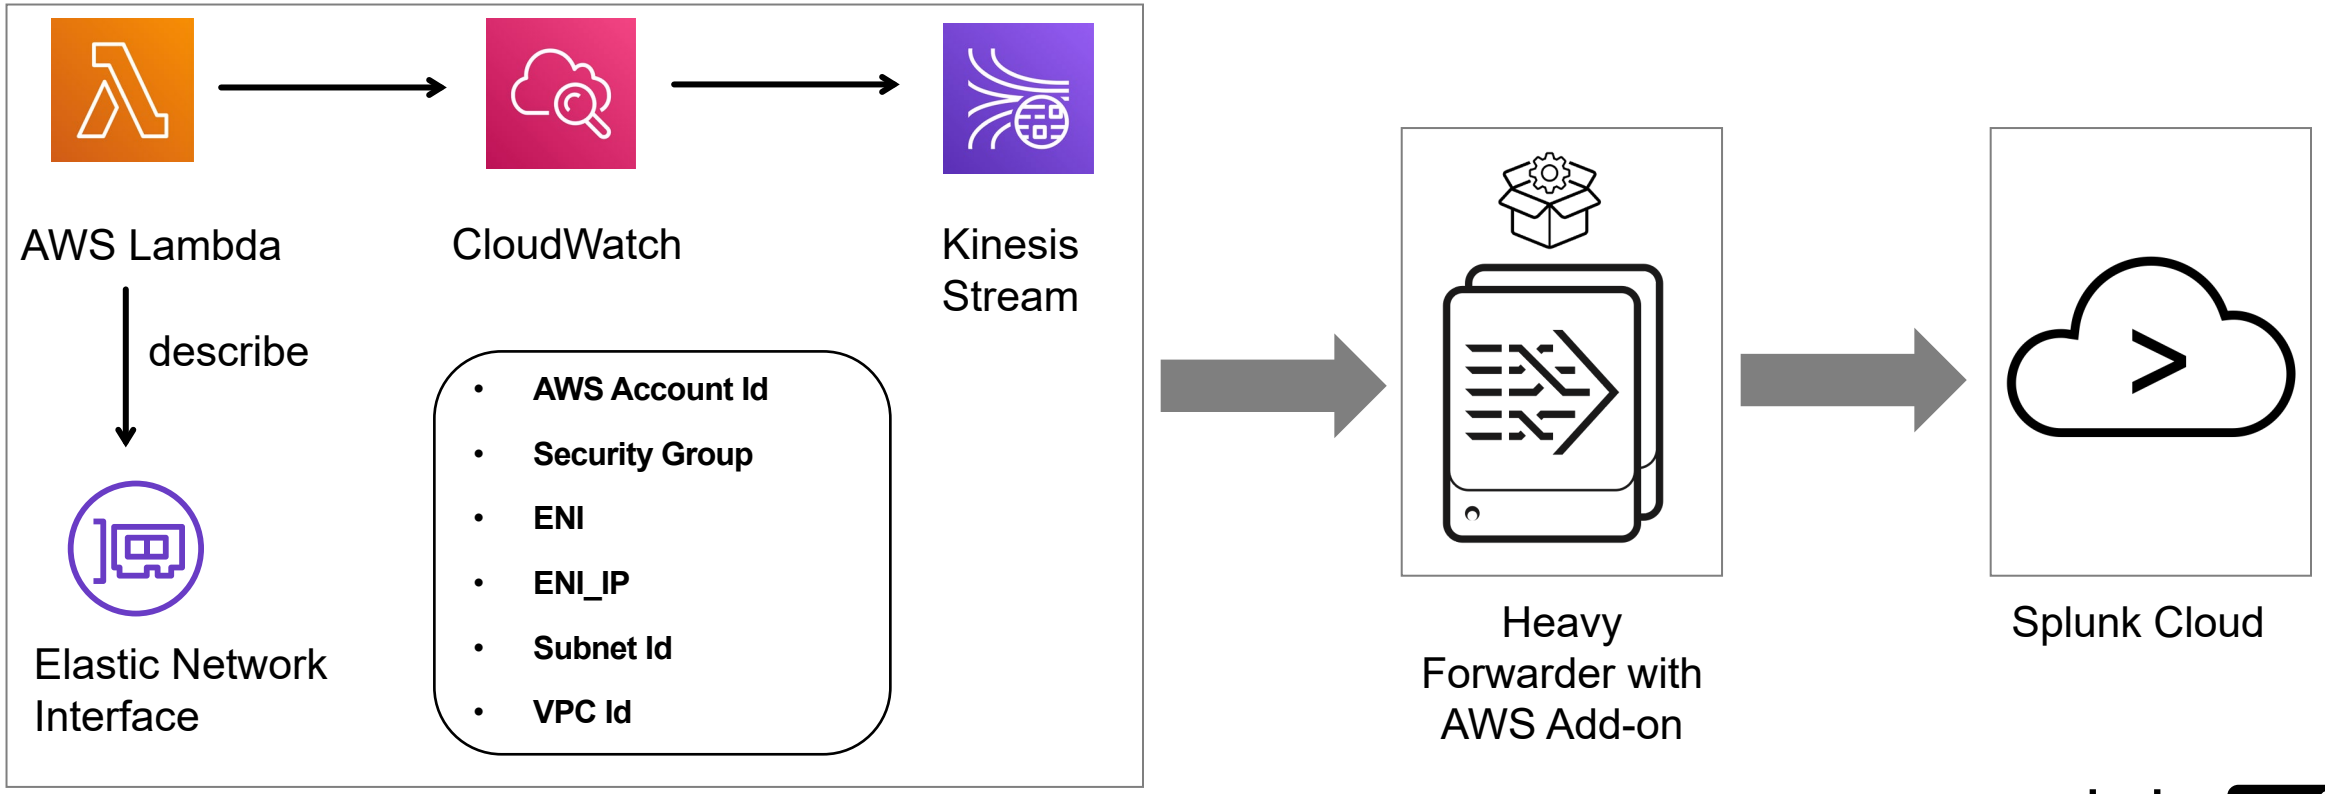

Data Correlation – Lookup Table 1

Create ENI (Elastic Network Interface) – Security Group mapping lookup table based on AWS lambda function logs

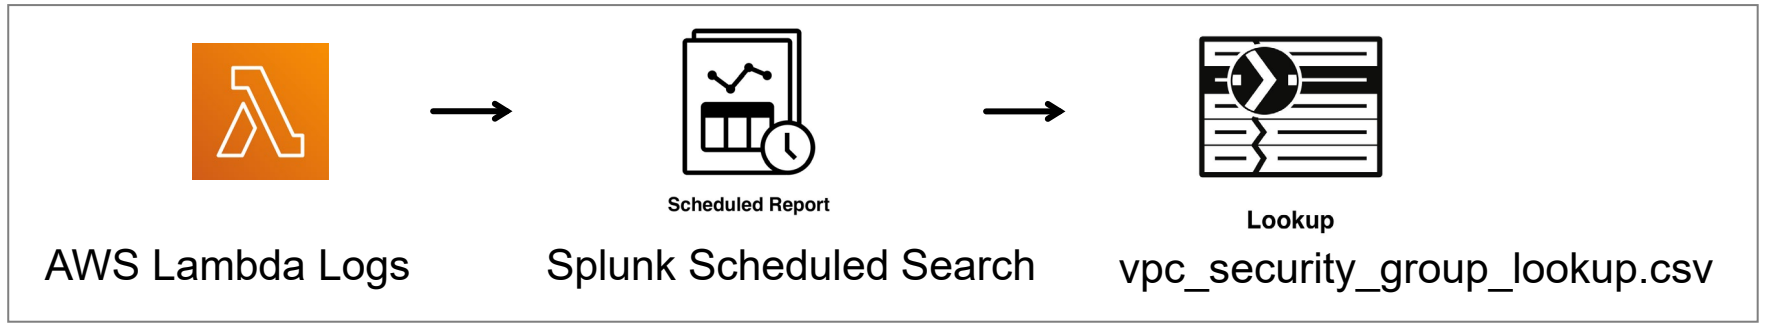

id	account_id	ENI	ENI_IP	subnet_id	vpc_id	InterfaceType
sg-xxxxxxxx	123456789012	eni-12345abcde	10.x.x.x	subnet-1234abcd	vpc-xxxxxxxx	interface
sg-yyyyyyyy	123456789012	eni-67890abcde	10.x.x.x	subnet-2345xyza	vpc-yyyyyyyy	interface
sg-zzzzzzzz	123456789012	eni-22234abcde	10.x.x.x	subnet-5678abcd	vpc-zzzzzzzz	interface

Data Correlation – Lookup Table 2

Asset and Environment Classification for Security Groups based on Security Group AWS Tags

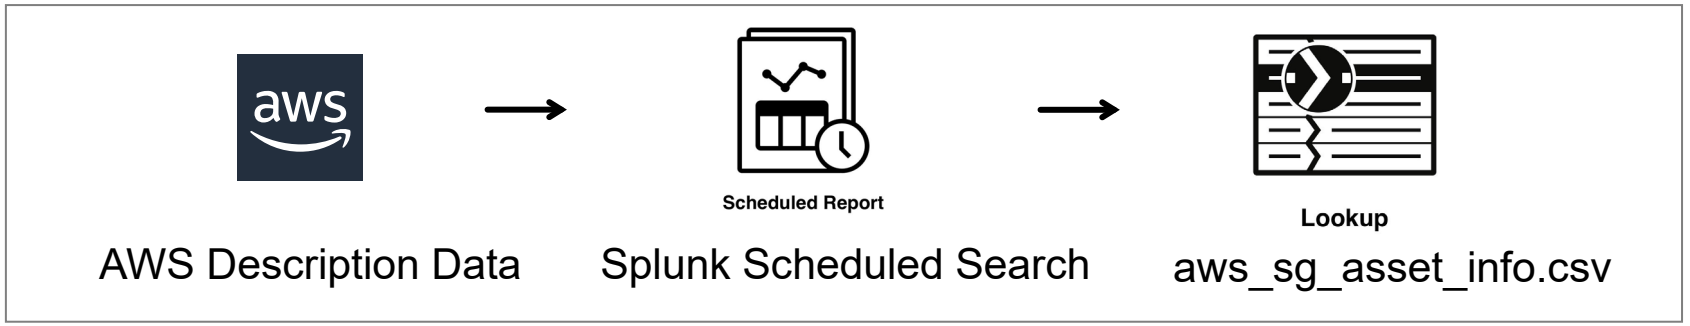

id	name	account_id	AssetGroup	AssetID	AssetName	vpc_id
sg-xxxxxxxx	SG for Search Service	123456789012	Dev	123	Search Service	vpc-xxxxxxxx
sg-yyyyyyyy	SG for Auth Service	123456789012	Cert	456	Auth Service	vpc-yyyyyyyy
sg-zzzzzzzz	SG for Validation Service	123456789012	Prod	789	Validation Service	vpc-zzzzzzzz

Final Data Correlation

With these 2 lookup tables vpc_security_group_lookup.csv and aws_sg_asset_info.csv, we can find ENIs for respective Security Group from non-production and production accounts belonging to different assets (services/applications)

version	account-id	interface-id	srcaddr	dstaddr	srcport	dstport	protocol	packets	bytes	start	end	action	status
2	123456789012	**eni-12345abcde**	10.x.x.x	10.x.x.x	20641	80	6	20	4249	1418530010	1418530070	ACCEPT	OK

SG Investigation and Remediation

Security Group Investigation and Remediation

1. **Insecure Security Groups Discovery**
   Splunk Dashboard – Use AWS Lambda Logs to detect Security Groups with 10.0.0.0/8
2. **VPC Flow Logs Analysis**
   Splunk Dashboards – Use VPC Flow Logs and Data Correlation with Security Groups
3. **Update Security Group Rules**
   Cloud Formation Template Updates and Release Management
4. **Deployment Monitoring**
   Splunk Configuration Rule Driven Alert & Dashboard - Detect Unusual Reject Traffic

1

Insecure Security Group Detection

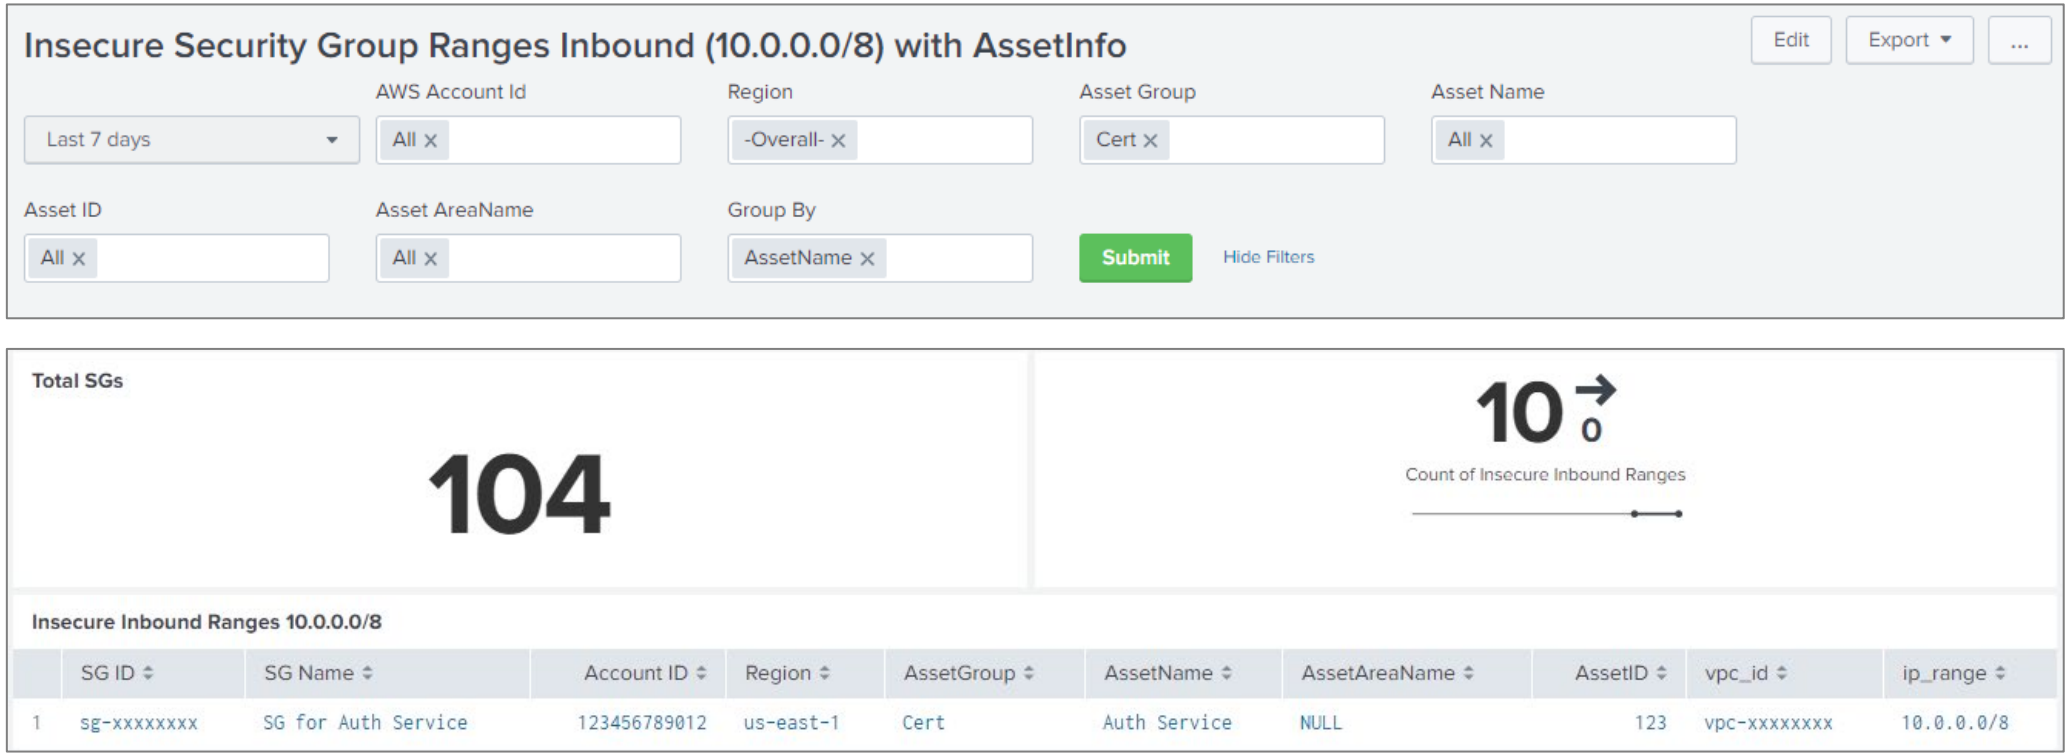

← Filter Criteria for production/non-production Assets SGs

← SG Statistics

← SG Details

VPC Flow Logs Analysis – Dashboard 1

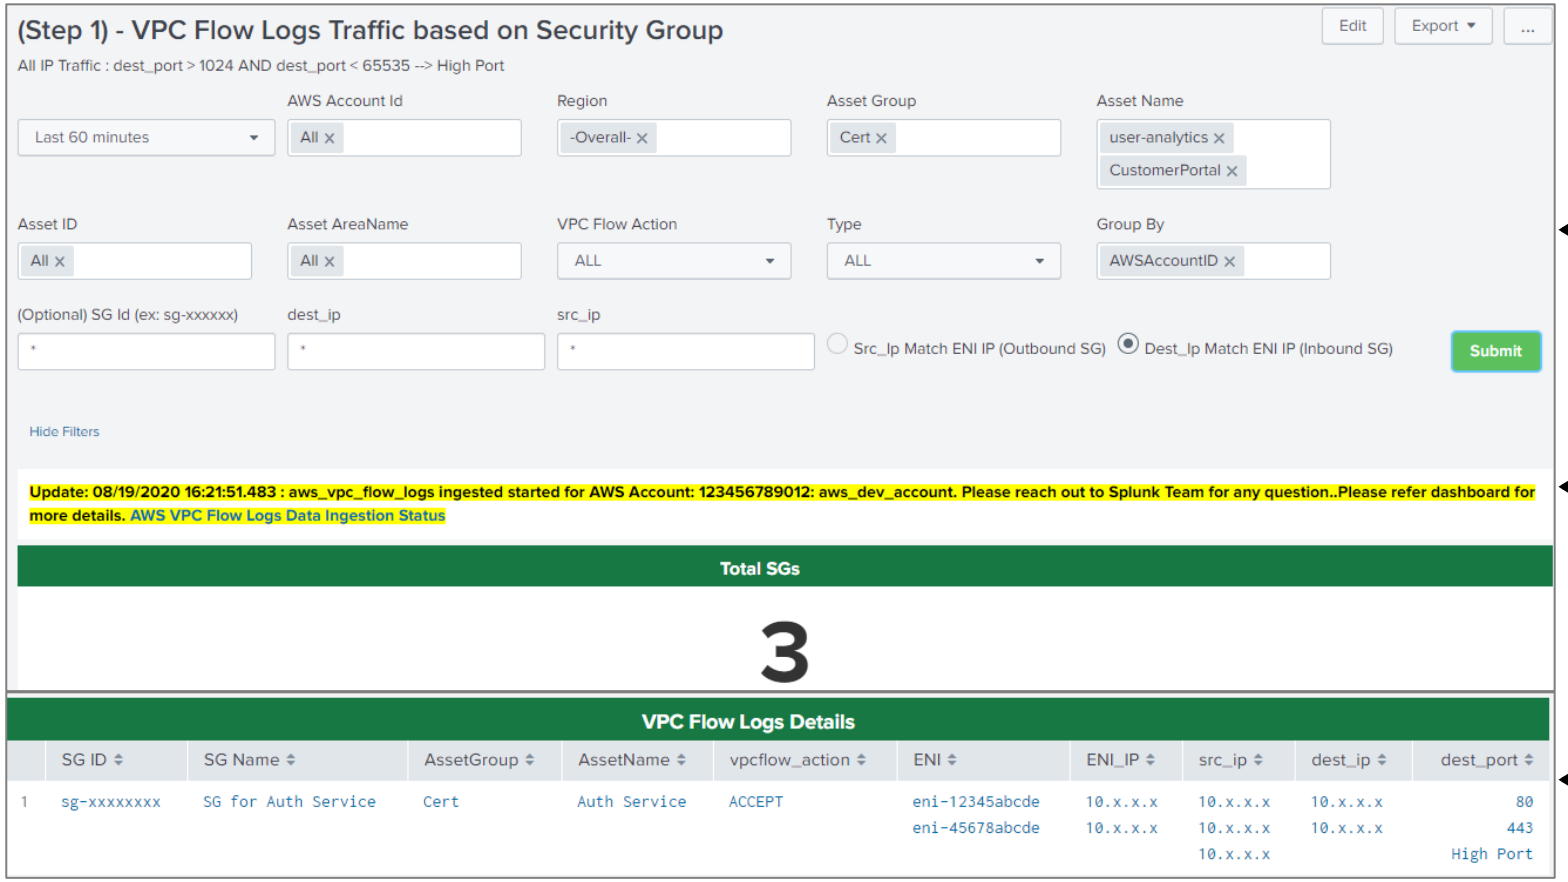

Filter Criteria for production/non-production Assets SGs

OnDemand Jenkins Jobs Logs to show VPC Flow Logs Ingestion Status

VPC Flow Logs Analysis with SG

splunk> .conf20

2

VPC Flow Logs Analysis – Dashboard 2

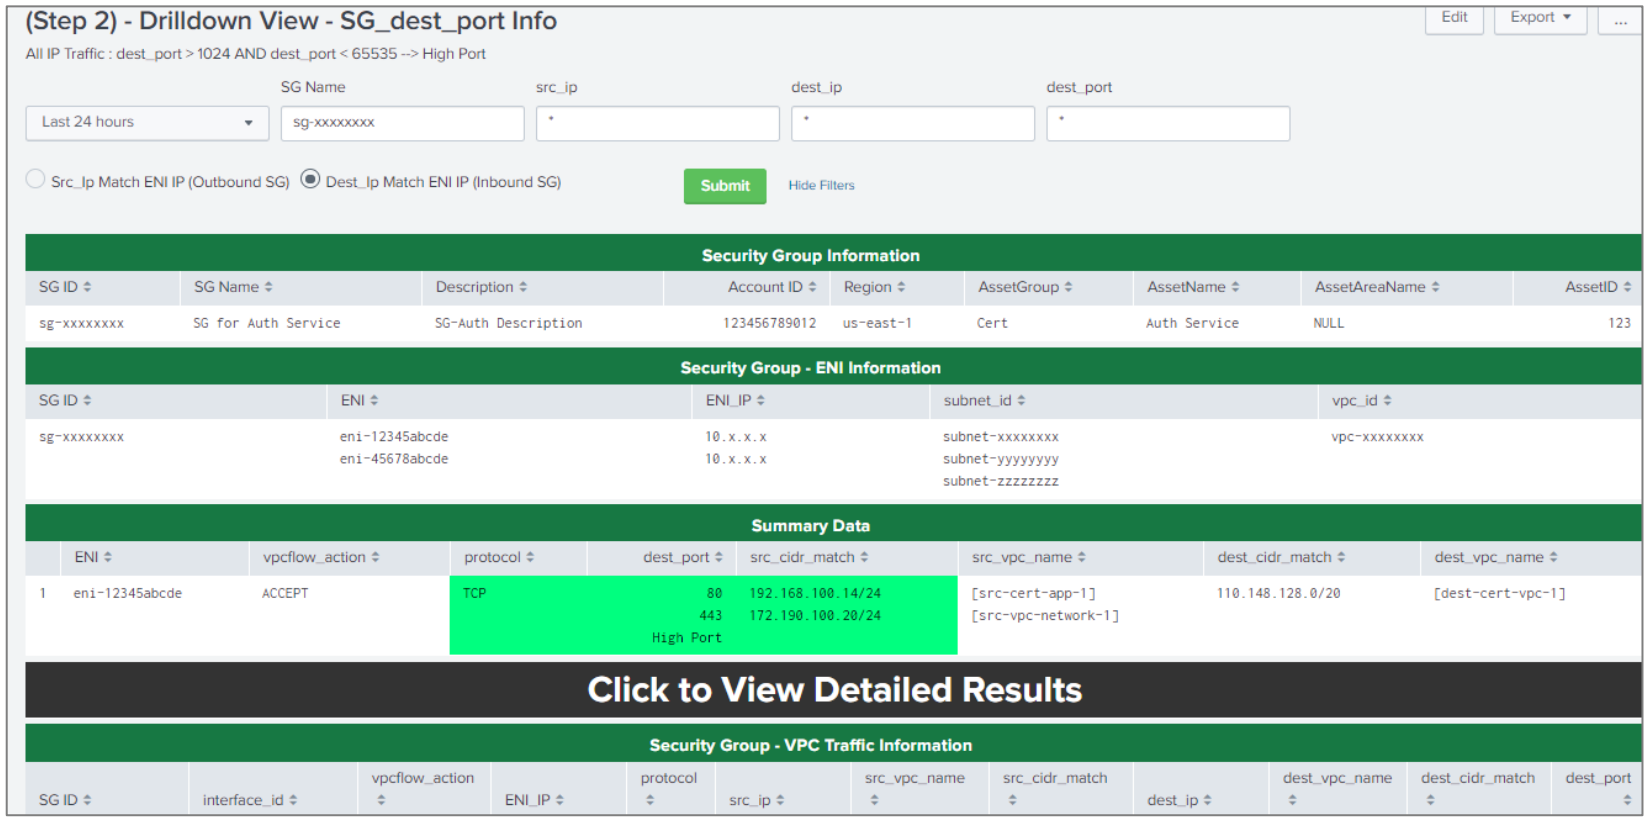

← SG Information

← SG – ENI Mapping

← VPC Traffic Information with CIDR Blocks

← Detailed Results broken down by IP

splunk> .conf20

Update Security Group Rules

- Developers review Security Group findings
- Audit findings with Network Engineer and DevOps Cloud Engineering Team
- Developers update Cloud Formation Templates and Stage Changes
- Work with Release Management to deploy changes every week

Deployment Monitoring

Are we rejecting valid traffic because of Security Group rule change?
...We have a way to detect unusual reject traffic by Configuration Rule driven Splunk Alert

AssetID	alert_title	assetArea	dest_port	is_active	protocol
123	Auth Service	auth	80,443	Y	*
456	Search Service	search	80,443	Y	*
789	Validation Service	validation	80,443	Y	*

Security Groups Updates…

1. sg-xxxxxxxx needs to be modified to accept traffic only from certain CIDR blocks
2. sg-xxxxxxxx needs to be modified to accept traffic only from certain ports
3. RDP/SSH traffic should Not be configured on ELB SG
4. sg-yyyyyyyy is a public facing and they do not need all the networks. Only 80/443 for default route (0.0.0.0/0)
5. sg-yyyyyyyy this is an internal SG and we need to replace the 10.0.0.0/8 with the list provided
6. Detect why certain valid traffic is getting REJECTED because SG rules are not sufficient
7. Detect why certain traffic is getting REJECTED because source is not making connection on right port
8. sg-zzzzzzzz has ICMP is locked down to only LexisNexis CIDR blocks

Splunk License Cost Savings

OnDemand Data Ingestion Model for AWS VPC Flow Logs

VPC Flow Log Ingestion Filtering

Filter S3 Event Notifications

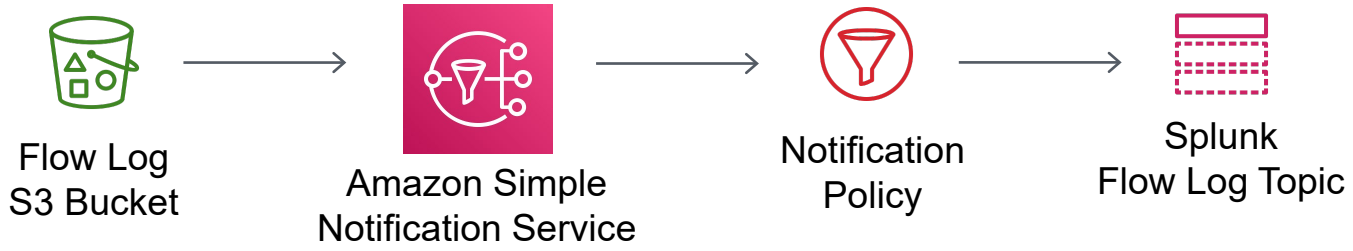

Filter With Splunk Heavy Forwarder

VPC Flow Log Ingestion Filtering

Example: Filter By AWS Account Id Using S3 Event Notification Prefix Filters

AWS Reference: https://docs.aws.amazon.com/AmazonS3/latest/dev/NotificationHowTo.html

Use S3 Notification Policy

- Filter Notifications Sent to SNS Topic for Splunk to Process

Example Object Name:

- s3://aws-logs-to-splunk/AWSLogs/123456789012/vpcflowlogs/us-east-1/2020/08/04/123456789012 _vpcflowlogs_us-east-1_fl-017ded307f939a89e_20200804T0000Z_2814abba.log.gz
- CLI To Apply Filter:
  - aws s3api put-bucket-notification-configuration --bucket aws-logs-to-splunk --notification-configuration file://./newPolicy.out

S3 Notification Policy Document

```
{
    "TopicConfigurations": [
        {
            "Filter": {
                "Key": {
                    "FilterRules": [
                        {
                            "Name": "Prefix",
                            "Value": "AWSLogs/123456789101"
                        }
                    ]
                }
            },
            "Id": "SendFlowLogNotificationsToSplunk0",
            "TopicArn": "arn:aws:sns:us-east-2: 123456789101:vpcflow-splunk-topic",
            "Events": [
                "s3:ObjectCreated:*"
            ]
        }
    ]
}
```

Example: Filter for REJECTs Only Using Props/Transforms on Heavy Forwarder

Default VPC Flow Log Fields

version	account-id	interface-id	srcaddr	dstaddr	srcport	dstport	protocol	packets	bytes	start	end	action	status
2	123456789010	eni-12345abcde	10.x.x.x	10.x.x.x	20641	22	6	20	4249	1418530010	1418530070	**REJECT**	OK

.conf File	Stanza
props.conf	[source::s3://aws-vpcflowlogs-to-splunk...] TRANSFORMS-all_vpcflowlogs = DropAllVPC,KeepVPCRejects
transforms.conf	[DropAllVPC] REGEX = .* DEST_KEY = queue FORMAT = nullQueue [KeepVPCRejects] REGEX = REJECT DEST_KEY = queue FORMAT = indexQueue

Summary

- With this framework, we have improved our AWS security posture by updating Security Group Rules for non-production and production Services and Applications
- We have targeted 1500 Security Group Updates and with this framework, we have cut down total estimates of ~7000 FTE hours to ~3500 FTE hours for Security Group Remediation Project

Learnings

- Correlation of AWS VPC Flow Logs, AWS lambda logs and AWS description logs
- Self-Service Splunk Dashboards for Developers and Network teams to audit Security Groups
- On-Demand Data Ingestion Model to help with Splunk License Savings

What Next…

- Update OnDemand Data Ingestion Model to Ingest data by specific VPC
- Self-Service Jenkins Job for Developers to ingest ACCEPT, REJECT or ALL traffic by specific AWS Account

Thank You

Please provide feedback via the

SESSION SURVEY