5 Steps: How to monitor network activity using Amazon Athena & AWS VPC Flow logs?

Are you seeking for a complete step by step guide to building your network monitoring system in AWS VPC? We have got you covered.

Breaking down the complex technical terms:

If you have a complete idea regards these terminologies, you are more than welcome to skip this section, but if you are looking for some easy to understand description relevant to these terms, here we go:

What is VPC?

VPC stands for Virtual Private Cloud. It allows you to have your isolated computing resources within the premises of the public cloud. It permits independent workflow between distinctive organizations.

Ummm, the flow logs?

Flow Logs is a peculiarity that allows you to collect and monitor the information of all the IP addresses travelling from or to your network. Having flow logs help you study the traffic metrics such as the direction of traffic, restrictive security measures, and the traffic that is reaching your cloud.

Amazon Athena:

It is an interactive service that allows you to analyse your data efficiently in Amazon S3. You can pay for solely the executed queries.

But, what about Amazon CloudWatch?

It is a compact service that allows you to monitor as well as manage your data. It comes in hand with several features allowing you to perform actions, analyse data and setup flow logs, all under one roof.

Time to build it:

It’s time to build your network monitoring environment in the VPC using flow logs. Are you ready?
Great, let's dive in!

<u>If you want to integrate with S3:</u>

- **Step 01: Enabling the VPC Flow logs:**
    1. Create a bucket in S3 by giving it a name of your choice.

2. Copy the created bucket's Amazon Resource Name.
3. In your VPC, create the bucket in the same region as that of S3.
4. Don't forget to enable encryption. (Why? Well, it's about sensitive data. Security is a must)

- **Step 02: Creating the VPC Flow logs**:
  1. Reach out the Flow Logs tab at the bottom and click on the Create Flow log.
  2. In the placeholder of Filter, select All. (Why? It defines which sort of traffic you want to filter through, selecting All, makes all the traffic visible)
  3. For the Destination, select Send to S3 (As you are integrating it with S3).
  4. In S3 bucket ARN, paste the ARN you copied below. (It was a smart move to copy before, wasn't it?)
  5. Click on Create to finalize the process. (No need to change any other option in between).

- **Step 03: Navigating the Amazon Athena:**
  1. Open the Amazon Athena.
  2. Run your query by using the following code:

     IMAGE
  3. Important: Don't forget to replace the address between < > with the one of your S3 flow log bucket that you created.

- **Step 04: Adding partitions:**
  1. To make data readable, click on Add Partition.
  2. Use the following query.

     IMAGE
  3. Important: Don't forget to replace the fields relevant to your data.

- **Step 05: Reading data from S3 Bucket:**

  If you want to read your data from S3 bucket, execute the following script:

  IMAGE

If you desire to integrate with CloudWatch:

- **Step 1: Creating flow logs:**
  1. Reach out to the Flow Logs tab and select Create Flow log.
  2. For Filter, select All (All sort of traffic is filtered).
  3. Select Destination as Send to CloudWatch Logs.
  4. Go to the CloudWatch section to hit on to the Logs.
  5. By clicking on Actions and then Create a log group, fill the name of the log group (For Example Flow-Logs) and end the step by clicking on Create log group.
  6. Enter this name in the Destination log Group.

- **Step 2: Creating the IAM Role:**
  1. Open a new tab.
  2. Select Roles succeeded by Create role with the name DeliverVPCFlowLogsRole.
  3. Select the service as EC2 and then click on Next Permissions.

4. By clicking on Create Policy and by selecting JSON, paste the following policy there.

IMAGE

5. Select the tab of Trust Relationship and then Edit Trust relationship by pasting the trust relationship code available at vpc-flow-logs.amazonaws.com.
6. Go back to the Flow logs window and select the IAM Role you just created and click on Create.

Once you have created the flow log, it should appear under the flow logs tab.

- **Step 3: Setting up the CloudWatch Metric Filters:**

Metric Filters allow you to set patterns and then alarm notification for informing you when a certain threshold reaches.

1. Select the CloudWatch Log group you created and then Create Metric Filter.
2. If you want to track the failed SSH attempts, paste the following code in Filter Pattern. (It can differ depending upon your choice)

IMAGE

3. Under the Test Patterns, select Custom Log Data. Paste the data as follows.

IMAGE

4. Click on Next.
5. Hit on Assign Metric.
6. Fill the placeholders as Filter Name: ssh-reject, Metric Name: SSH Rejects, and Metric Value: 1. Click on Next and then Save Changes.

- **Step 4: Creating the Metric Filter Alarm:**

1. Once you have created the Metric Filter, click on Create Alarm.
2. Fill the fields as Threshold: Greater/Equal (or of your choice), Threshold Value: 1 and click on Next.
3. Create a new topic with the name SSH_DENIED_ACCESS_IN_VPC.
4. In the Send Notification To, write your email address to receive the notifications.
5. Set Alarm Name as SSH Rejects. And then hit on Create Alarm.

- **Step 5: Executing the queries:**

You can execute all the queries of your wish in the CloudWatch now. Steps are as follows:
Insights > Queries > Sample queries > VPC flow log queries > Top 10 source IP addresses with the highest number of rejected requests

All in all, as promised, we have provided you with a detailed step by step guide of setting up monitoring environment in your VPC using flow logs and Amazon Athena. Also, it's always a good idea to experiment with new queries.

Wish you happy network monitoring! :)

References:

https://codeburst.io/network-monitoring-with-aws-vpc-flow-logs-and-amazon-athena-de94969f4175

https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html

https://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html

https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs-cwl.html

https://aws.amazon.com/athena/?whats-new-cards.sort-by=item.additionalFields.postDateTime&whats-new-cards.sort-order=desc

Videos:

https://www.youtube.com/watch?v=_A5L4jT-K9I

https://www.youtube.com/watch?v=I_VjSvSSoF4

https://www.youtube.com/watch?v=dyWtRCIO09I

https://www.youtube.com/watch?v=SiUDN95sJIo