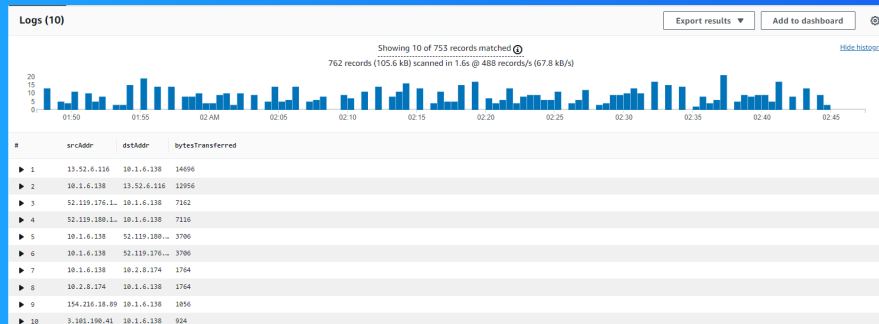




# VPC Monitoring with Flow Logs



mmoorej24@gmail.com





**mmoorej24@gmail.com**  
NextWork Student

[NextWork.org](https://NextWork.org)

# Introducing Today's Project!

## What is Amazon VPC?

AWS VPC is a virtual network dedicated to your AWS account, allowing you to provision and manage a logically isolated section of the AWS cloud where you can launch AWS resources.

## How I used Amazon VPC in this project

I used two Amazon VPCs for peer connecting to monitor logs using the CloudWatch service.

## One thing I didn't expect in this project was...

I didn't expect this project to be so awesome.

## This project took me...

This project took me about two hours and forty minutes to complete.



# In the first part of my project...

## Step 1 - Set up VPCs

In this step we will set up two VPCs from scratch in minutes. Network monitoring can still be done with just a single VPC, but it's great to have the extra challenge and tackle VPC peering in this project too!

## Step 2 - Launch EC2 instances

In this step, we launch two EC2 instances - one in each VPC. Doing this is important to set up the remainder of our project. Our EC2 instances will generate traffic that the VPC Flow Logs will monitor.

## Step 3 - Set up Logs

In this step, we are setting up VPC flow logs to start monitoring network traffic. We are also setting up a storage space for our flow logs.

## Step 4 - Set IAM permissions for Logs

In this step, we provide VPC Flow Logs with the permission to create logs and upload them into our log group in CloudWatch.



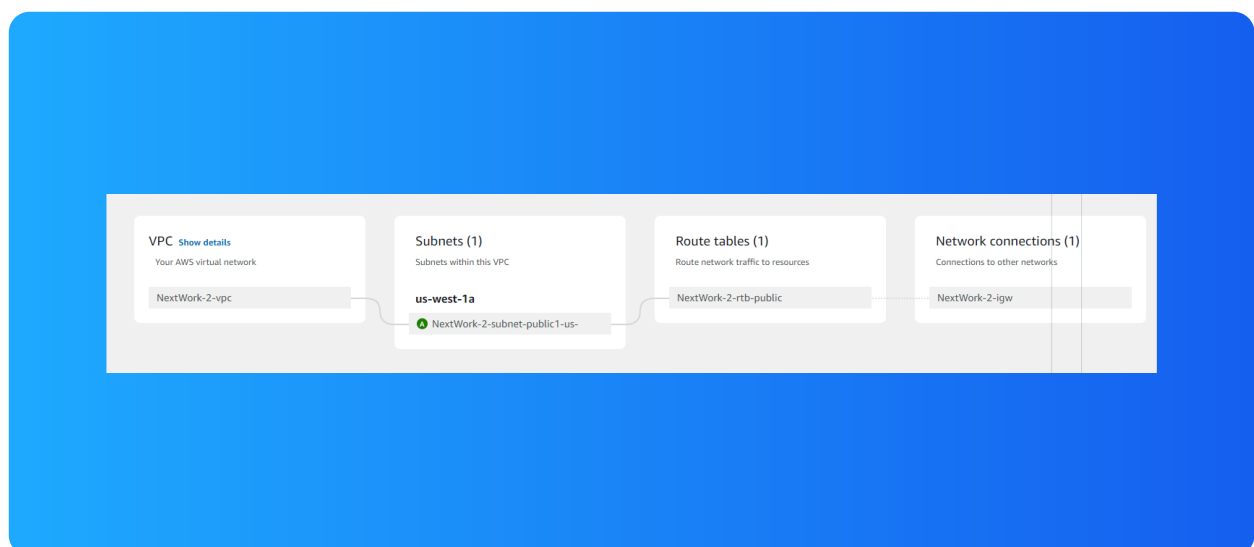
# Multi-VPC Architecture

I started my project by launching two VPCs! We created two public subnets (i.e. one public subnet in each VPC) with no private subnets.

The CIDR blocks for VPCs 1 and 2 are 10.1.0.0/16 and 10.2.0.0/16. They have to be unique because having overlapping CIDR blocks will cause network/routing traffic issues down the line when traffic is needing to go from one VPC to another.

## I also launched EC2 instances in each subnet

My EC2 instances' security groups allow SSH and ICMP type traffic. This is because EC2 Instance Connect will need to access our EC2 instance using SSH-type traffic, and because we need to allow ICMP type traffic for connectivity tests later.





# Logs

Logs are like diary entries for my computer systems - they are detailed records of any kind of activity related to the traffic/resource/AWS service that the log is tracking.

Log groups are groupings of logs i.e. logs that belong to the project/application/source are often in a log group together!

## I also set up a flow log for VPC 1

The screenshot shows the AWS IAM console 'Flow logs (1) info' page. It features a search bar, an 'Actions' dropdown, and a 'Create flow log' button. Below is a table with the following data:

<input type="checkbox"/>	Name	Flow log ID	Filter	Destination type	Destination name	IAM role ARN
<input type="checkbox"/>	NextWorkVPCFlowLog	fl-0286dc57a37abecd9	ALL	cloud-watch-logs	<a href="#">NextWorkVPCFlowLogsGroup</a>	arn:aws:iam:088535984191:role/Next...



# IAM Policy and Roles

I created an IAM policy so that I can define a rule that allows policy holders e.g. our VPC Flow Logs service the ability to create log streams and upload them into CloudWatch.

I also created an IAM role because services like VPC Flow Logs have to be associated with a role instead of JSON! Creating an IAM role will be necessary to give our VPC Logs the access it needs to record and upload logs.

A custom trust policy is specific type of policy! They're different from IAM policies. While IAM policies help you define the actions a user/service can or cannot do, custom trust policies are used to very narrowly define who can use a role.

## Custom trust policy

Create a custom trust policy to enable others to perform actions in this account.

```
1  {
2    "Version": "2012-10-17",
3    "Statement": [
4      {
5        "Sid": "Statement1",
6        "Effect": "Allow",
7        "Principal": {
8          "Service": "vpc-flow-logs.amazonaws.com"
9        },
10       "Action": "sts:AssumeRole"
11     }
12   ]
13 }
```



# In the second part of my project...

## Step 5 - Ping testing and troubleshooting

In this step, I am generating network traffic! This becomes important when we are communicating to cloudworks/cloud networking/or cloud engineering.

## Step 6 - Set up a peering connection

In this step, we are setting up a peering connection so that VPCs 1 and 2 can talk directly with each other.

## Step 7 - Update VPC route tables

In this step, we are updating the route tables for our two VPCs so that traffic bound for the other VPC can be directed to the peering connection instead of the public internet.

## Step 8 - Analyze flow logs

In this step, we are tracking the network data that's been collected on our VPCs, and then analyze that data to extract insights.







# Connectivity troubleshooting

Looking at VPC 1's route table, I identified that the ping test with Instance 2's private address failed because we do not have a route in our VPCs' route tables that directs traffic from one VPC to another.

## To solve this, I set up a peering connection between my VPCs

I also updated both VPCs' route tables so that traffic from one of the VPCs and heading to the other VPC's private IPV4 address can be directed to go through the peer connection instead of the public internet.

Destination	Target	Status	Propagated
0.0.0/0	<a href="#">igw-0b8d9cca3134a39f</a>	Active	No
10.1.0.0/16	<a href="#">pcx-0a8b8b25f1785a0</a>	Active	No
10.2.0.0/16	local	Active	No



# Connectivity troubleshooting

I received ping replies from Instance 2's private IP address! This means setting up the peering connection and then the route table solve the connectivity error of our VPCs' traffic not being able to navigate from one VPC to another.

```
64 bytes from 3.101.190.41: icmp_seq=4 ttl=126 time=0.911 ms
64 bytes from 3.101.190.41: icmp_seq=5 ttl=126 time=0.741 ms
64 bytes from 3.101.190.41: icmp_seq=6 ttl=126 time=1.23 ms
^C
--- 3.101.190.41 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5192ms
rtt min/avg/max/mdev = 0.461/0.821/1.226/0.244 ms
[ec2-user@ip-10-1-6-138 ~]$ ping 10.2.8.174
PING 10.2.8.174 (10.2.8.174) 56(84) bytes of data:
64 bytes from 10.2.8.174: icmp_seq=1 ttl=127 time=1.41 ms
64 bytes from 10.2.8.174: icmp_seq=2 ttl=127 time=1.14 ms
64 bytes from 10.2.8.174: icmp_seq=3 ttl=127 time=0.915 ms
64 bytes from 10.2.8.174: icmp_seq=4 ttl=127 time=0.708 ms
64 bytes from 10.2.8.174: icmp_seq=5 ttl=127 time=0.567 ms
64 bytes from 10.2.8.174: icmp_seq=6 ttl=127 time=1.11 ms
64 bytes from 10.2.8.174: icmp_seq=7 ttl=127 time=0.985 ms
64 bytes from 10.2.8.174: icmp_seq=8 ttl=127 time=0.637 ms
64 bytes from 10.2.8.174: icmp_seq=9 ttl=127 time=1.11 ms
64 bytes from 10.2.8.174: icmp_seq=10 ttl=127 time=0.492 ms
64 bytes from 10.2.8.174: icmp_seq=11 ttl=127 time=0.730 ms
64 bytes from 10.2.8.174: icmp_seq=12 ttl=127 time=0.740 ms
64 bytes from 10.2.8.174: icmp_seq=13 ttl=127 time=0.579 ms
64 bytes from 10.2.8.174: icmp_seq=14 ttl=127 time=0.722 ms
64 bytes from 10.2.8.174: icmp_seq=15 ttl=127 time=0.870 ms
64 bytes from 10.2.8.174: icmp_seq=16 ttl=127 time=0.463 ms
64 bytes from 10.2.8.174: icmp_seq=17 ttl=127 time=0.694 ms
64 bytes from 10.2.8.174: icmp_seq=18 ttl=127 time=0.642 ms
64 bytes from 10.2.8.174: icmp_seq=19 ttl=127 time=0.849 ms
64 bytes from 10.2.8.174: icmp_seq=20 ttl=127 time=0.716 ms
64 bytes from 10.2.8.174: icmp_seq=21 ttl=127 time=0.691 ms
^C
--- 10.2.8.174 ping statistics ---
21 packets transmitted, 21 received, 0% packet loss, time 20567ms
rtt min/avg/max/mdev = 0.463/0.798/1.410/0.234 ms
[ec2-user@ip-10-1-6-138 ~]$ ping 3.101.190.41
PING 3.101.190.41 (3.101.190.41) 56(84) bytes of data:
64 bytes from 3.101.190.41: icmp_seq=1 ttl=126 time=0.472 ms
64 bytes from 3.101.190.41: icmp_seq=2 ttl=126 time=0.672 ms
64 bytes from 3.101.190.41: icmp_seq=3 ttl=126 time=0.738 ms
64 bytes from 3.101.190.41: icmp_seq=4 ttl=126 time=1.16 ms
```



# Analyzing flow logs

Flow logs tell us about the source and destination of the network traffic, the amount of data being transferred, whether the traffic is being rejected or accepted.

For example, the flow log I've captured tells us that traffic went from one address to another. We can also extract that the traffic was accepted by the security groups and network ACLs of my VPC.

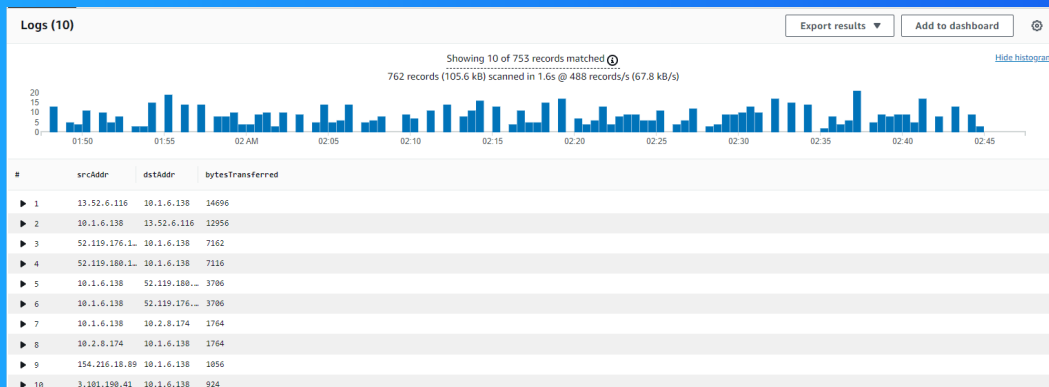
2024-08-10T02:35:30.000Z	2 088535984191 eni-bccc0736cd77e693a 13.52.6.116 10.1.6.138 24862 22 6 2 172 1723257330 1723257374 ACCEPT OK
2 088535984191 eni-bccc0736cd77e693a 13.52.6.116 10.1.6.138 24862 22 6 2 172 1723257330 1723257374 ACCEPT OK	🗑️
2024-08-10T02:35:30.000Z	2 088535984191 eni-bccc0736cd77e693a 10.1.6.138 13.52.6.116 22 24862 6 1 88 1723257330 1723257374 ACCEPT OK
2 088535984191 eni-bccc0736cd77e693a 10.1.6.138 13.52.6.116 22 24862 6 1 88 1723257330 1723257374 ACCEPT OK	🗑️
2024-08-10T02:36:06.000Z	2 088535984191 eni-bccc0736cd77e693a 89.248.163.200 10.1.6.138 54964 5432 6 1 40 1723257366 1723257399 REJECT OK
2 088535984191 eni-bccc0736cd77e693a 89.248.163.200 10.1.6.138 54964 5432 6 1 40 1723257366 1723257399 REJECT OK	🗑️
2024-08-10T02:36:06.000Z	2 088535984191 eni-bccc0736cd77e693a 162.216.150.98 10.1.6.138 54769 52975 6 1 44 1723257366 1723257399 REJECT OK
2 088535984191 eni-bccc0736cd77e693a 162.216.150.98 10.1.6.138 54769 52975 6 1 44 1723257366 1723257399 REJECT OK	🗑️
2024-08-10T02:36:06.000Z	2 088535984191 eni-bccc0736cd77e693a 170.187.163.90 10.1.6.138 52196 25462 6 1 44 1723257366 1723257399 REJECT OK
2 088535984191 eni-bccc0736cd77e693a 170.187.163.90 10.1.6.138 52196 25462 6 1 44 1723257366 1723257399 REJECT OK	🗑️
2024-08-10T02:36:06.000Z	2 088535984191 eni-bccc0736cd77e693a 4.246.247.164 10.1.6.138 48521 1527 6 1 40 1723257366 1723257399 REJECT OK
2 088535984191 eni-bccc0736cd77e693a 4.246.247.164 10.1.6.138 48521 1527 6 1 40 1723257366 1723257399 REJECT OK	🗑️



# Logs Insights

Logs Insights is a CloudWatch feature that analyzes your logs. In Log Insights, you use queries to filter, process and combine data to help you troubleshoot problems or better understand your network traffic!

I ran the query "Top 10 byte transfers by source and destination IP addresses". This query analyzes the top 10 biggest data transfers between IP addresses in your network!





[NextWork.org](https://NextWork.org)

# Everyone should be in a job they love.

Check out [nextwork.org](https://nextwork.org) for  
more projects

