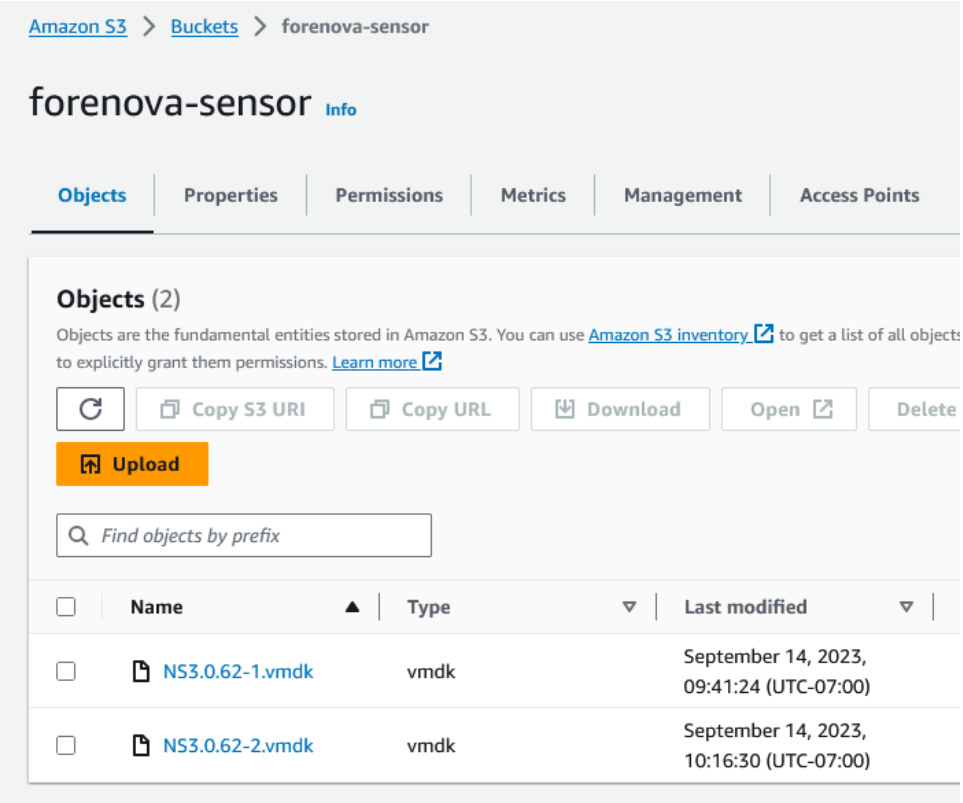
Create a AWS NovaSensor from 1st steps.

Install NovaSensor in VMWare

Just install, after 1st boot, login and shut down the VM

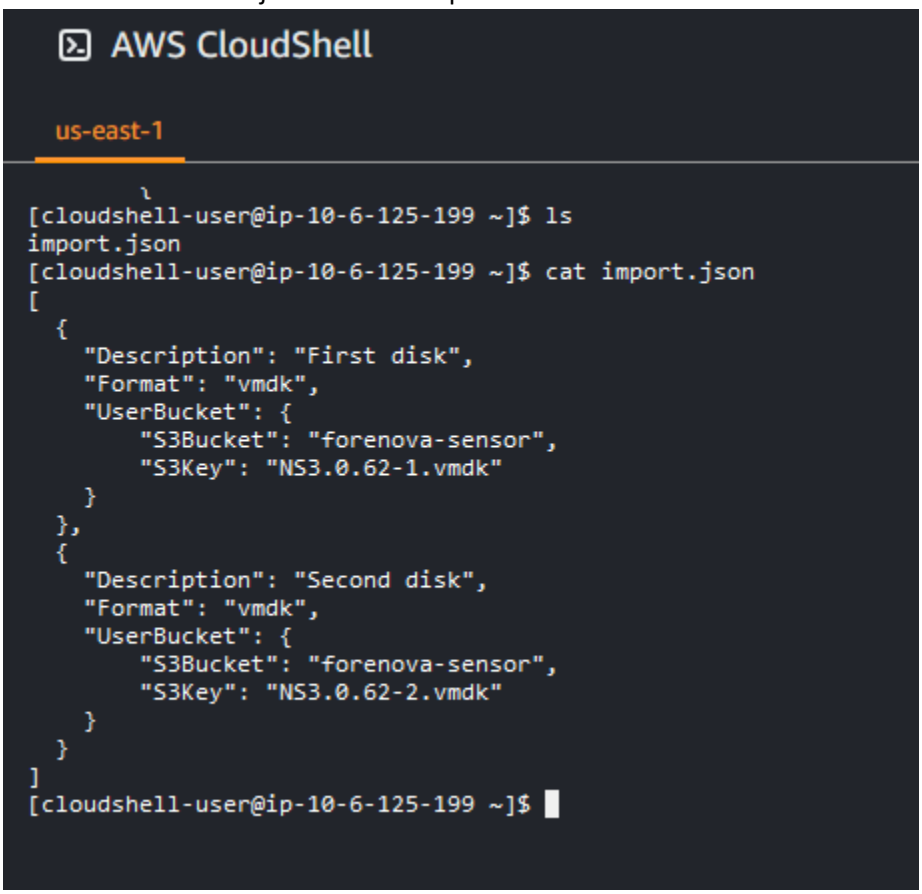
Export the VM, so you have 2 disks (VMDK files)

Upload 2 VMDK's (disk 1 and disk2) to AWS S3



Now that the disks are in AWS, you can import the two disks as an AMI into AWS.
You can do this by the AWS CLI installed local on laptop, or use AWS CloudShell (easy way)
Be sure to start AWS CloudShell in same region as the disks are in S3

In AWS shell create a json file with import details:



Start the import:
aws ec2 import-image --description "NovaSensor" --disk-containers "file://import.json"

Check status with ami-import ID:
aws ec2 describe-import-image-tasks --import-task-ids import-ami-036f29aa5b32b3b02

AWS CloudShell

us-east-1

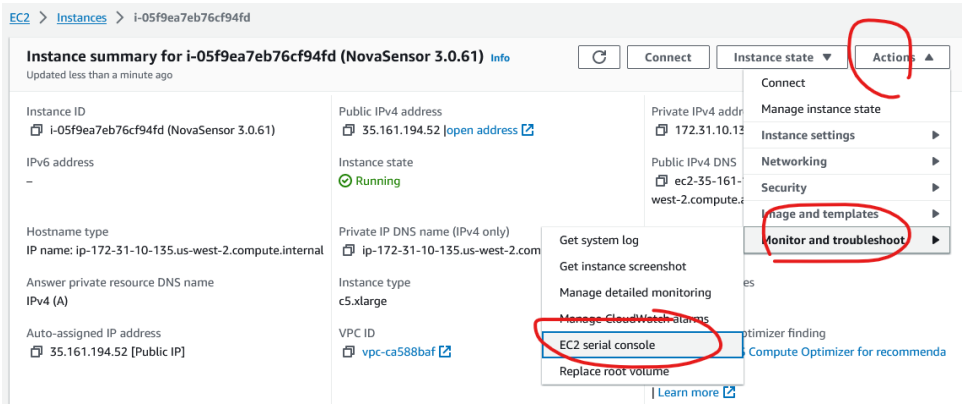
```
}
[cloudshell-user@ip-10-2-66-53 ~]$ aws ec2 describe-import-image-tasks --import-task-ids import-ami-036f29aa5b32b3b02
{
    "ImportImageTasks": [
        {
            "Description": "NovaSensor",
            "ImportTaskId": "import-ami-036f29aa5b32b3b02",
            "Progress": "19",
            "SnapshotDetails": [
                {
                    "Description": "First disk",
                    "DiskImageSize": 6726101504.0,
                    "Format": "VMDK",
                    "Status": "active",
                    "UserBucket": {
                        "S3Bucket": "forenova-sensor",
                        "S3Key": "NS3.0.62-1.vmdk"
                    }
                },
                {
                    "Description": "Second disk",
                    "DiskImageSize": 824320.0,
                    "Format": "VMDK",
                    "Status": "active",
                    "UserBucket": {
                        "S3Bucket": "forenova-sensor",
                        "S3Key": "NS3.0.62-2.vmdk"
                    }
                }
            ],
            "Status": "active",
            "StatusMessage": "converting",
            "Tags": []
        }
    ]
```

When task is finished you have a new AMI that is the basic NovaSensor

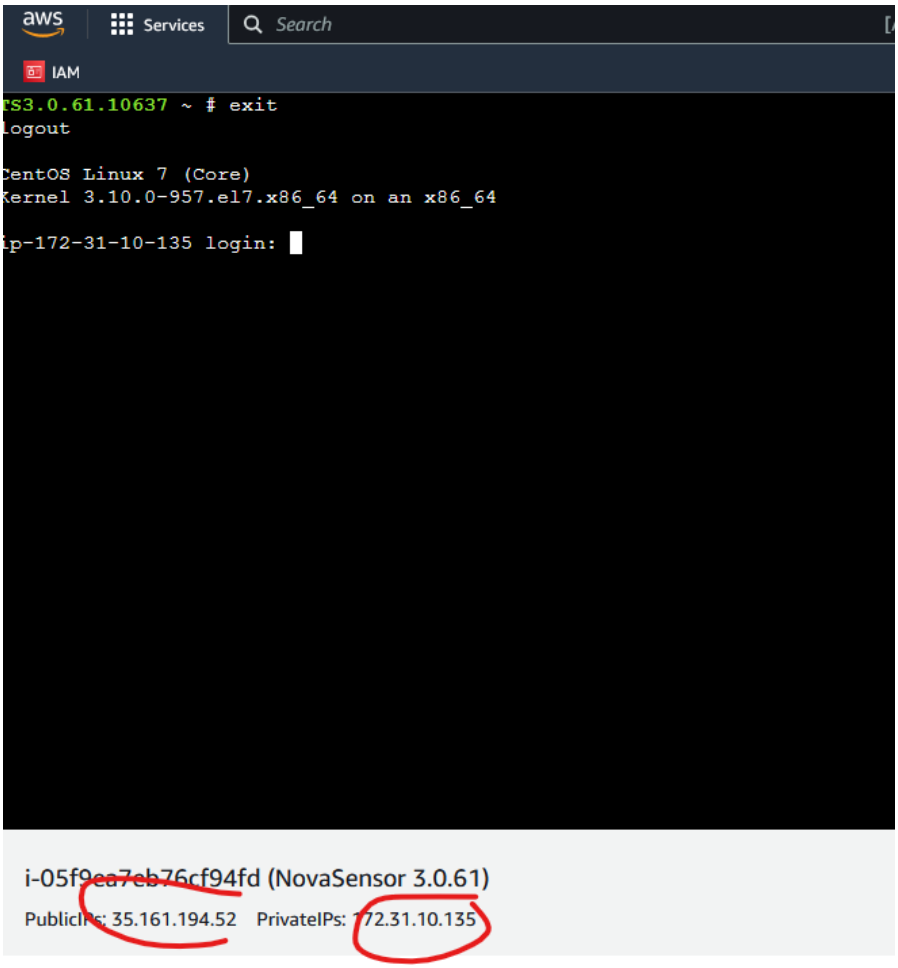
Start a new instance from the AMI:
Important: choose an instance type that has AWS Nitro enabled (so you can use the EC2 serial Console)
In this example I used C5.XLarge (8 GB and 4 cores)

When the instance is created AWS will reserve an IP address (local and if requested also external)
You will see them in the EC2 serial Console.

Connect to instance using AWS serial Concole:



Serial Console will start and you even have access to booting of the machine:



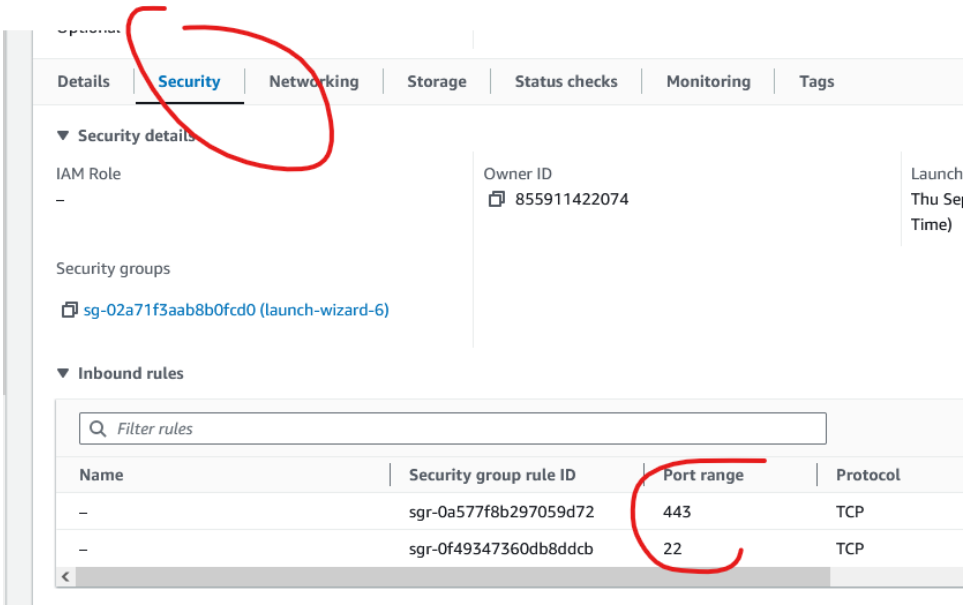
Login with default credentials (admin / adminsangfornetwork)

Public IP and Private IP are in cloudconsole BUT NOT ON NOVASENSOR
Set this ip info on NovaSensor:
ip a add 172.31.10.135/20 dev eth0 (check ip range and subnet in VPC settings)

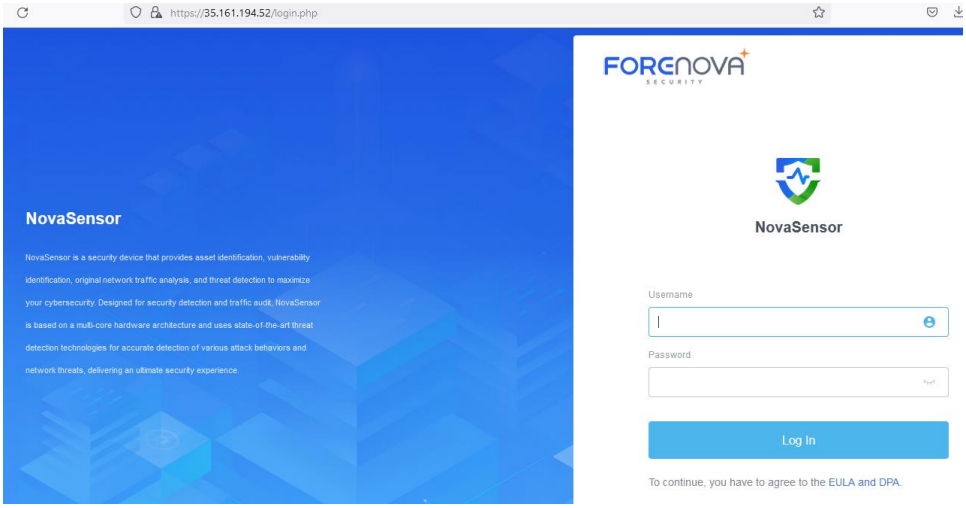
Set the static route:
ip route add 0.0.0.0/0 via 172.31.0.1 dev eth0 (most of the time the first IP in VPC subnet)

Check if you can access outside, by ping 8.8.8.8 for example

Now check if the firewall settings are OK (you need to allow at least 443 for web management)
You can find this in the instance settings at the Security tab.



Try to access the instance from outside:

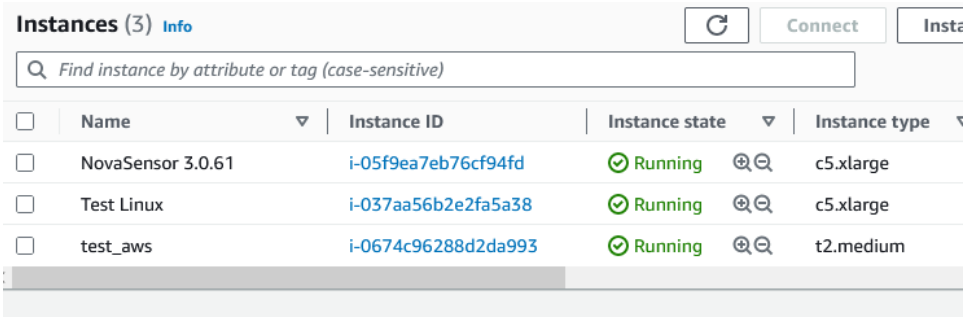


Setting up SPAN data in AWS

To set up the SPAN data in AWS, take the next steps:

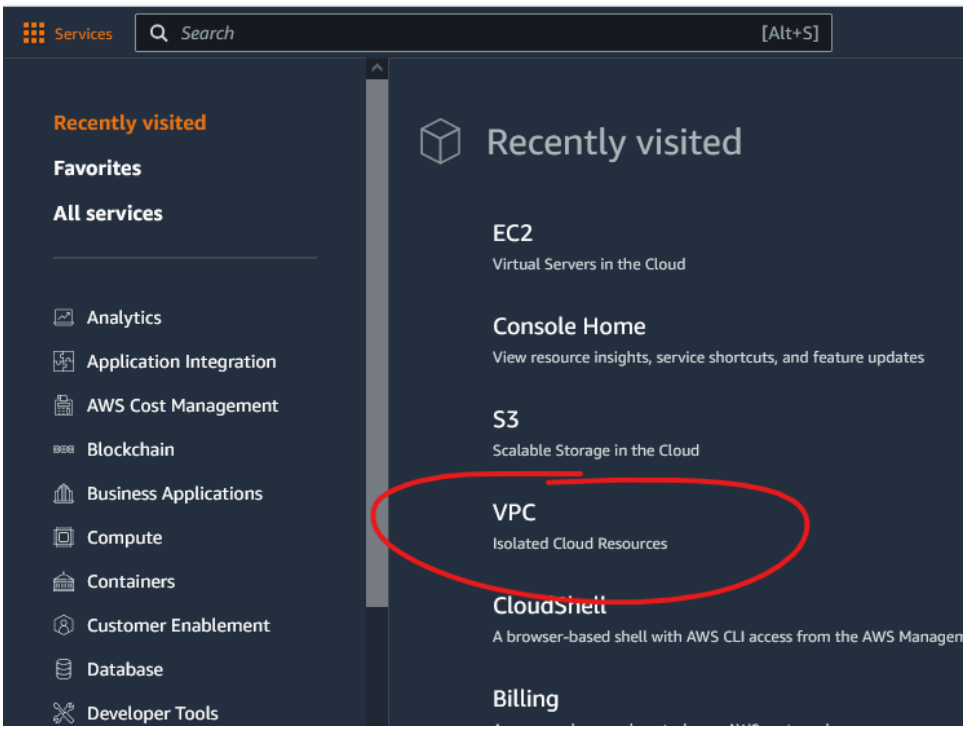
Add an additional interface on the NovaSensor

Create a mirror session in the VPC where the data is that you like to monitor.
In the example I created a new instance (Linux test) that we are going to monitor.
Be aware of that not all instance types are supported for SPAN data in the VPC.
See also: https://docs.aws.amazon.com/vpc/latest/mirroring/what-is-traffic-mirroring.html
For the test I choose a C5.XLarge

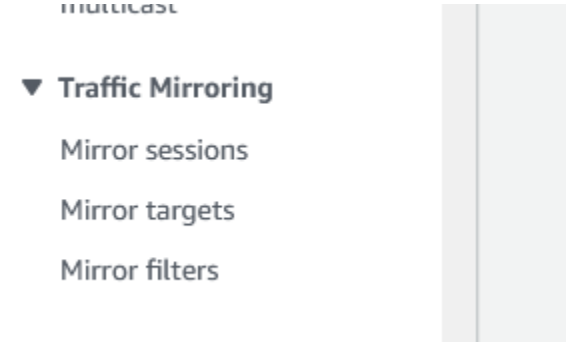


Setting up the mirror session:

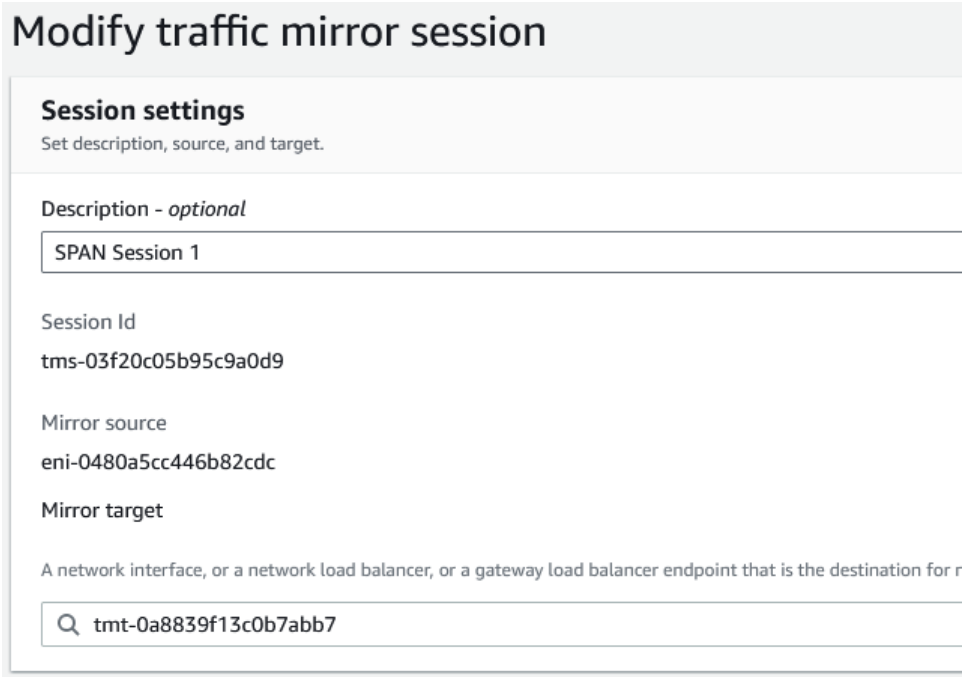
The mirror settings are in the VPC part of AWS, so go there:



Go to Traffic Mirroring in the menu:



Create a mirror session.
As example I created the following session:



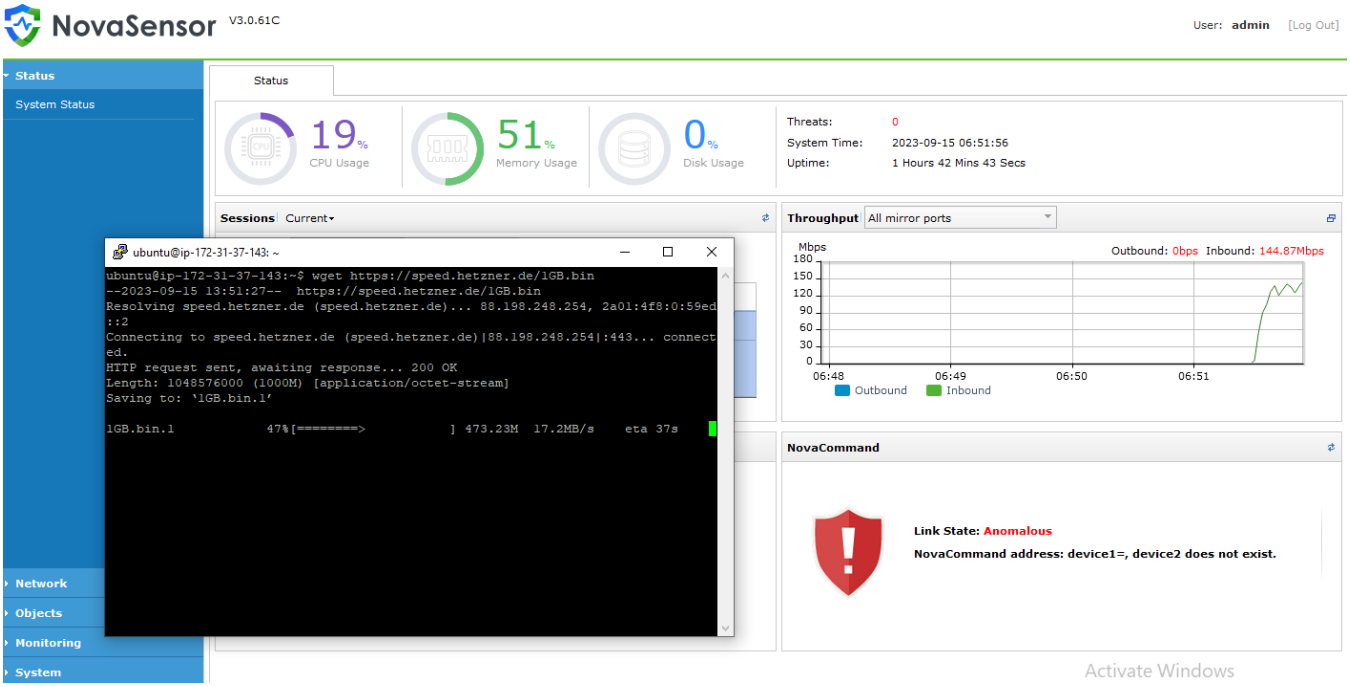
Important, you need to fill in the details like Source, Destination, Target, Session number and filter.
If you don't have these, you can create them on the fly.

Important: adjust the firewall settings on the destination interface (SPAN) and on the source interface as well. For example, if you are not allowed to send out TCP on port 123, you will never see this traffic in your SPAN data.

Example firewall (allow) rules for adapters both in and outbound:

Port range	Protocol	Source
0 - 65535	UDP	0.0.0.0/0
0 - 65535	TCP	0.0.0.0/0

After that, if you test data on the 'source' machine (linux test) you see the details in the NovaSensor:
Test download on Linux showing the throughput in NovaSensor.



The AMI is in 'Oregon' region
The Instance is also running there.
You can delete the instance if you like, username and password are still default.

After editing the NovaSensor instance, you can of course create a new AMI and copy that AMI to other regions or have customers access the AMI for their needs.

I will delete the two instances and the mirror config next week because of costs.
Please let me know if there are any questions on this.