Virtual Machines in AWS Version 1.1

Step by Step Guide

The information below includes detailed pre-deployment requirements, an in depth step by step guide for the AWS Virtual Machines v1.1 Click to Run deployment, and post deployment steps that will need to be considered.

Page 9: Post deployment activities

✓ Infrastructure Requirements

• **Existing VPC and subnets** - To deploy the click to run solution, a VPC with associated subnets must already exist in the region, with corresponding route tables, gateways etc. The solution does not do any configuration out of the scope of the virtual machine, so to ensure internet access the VPC must be configured properly.

• **Existing key pair:** In order to access the EC2 instance, a key pair must be chosen. The key pair must exist in the selected region at the time of deployment.

Virtual Machines v1.1 Pre-Deployment Steps

1. To look for the IDs of the VPC and Subnet where the solution will be deployed:

   a. Go to the AWS console.

   b. Once in the console, select the region where you will deploy the solution.

   c. Look for the service "VPC".

   d. Go to the "subnets" section.

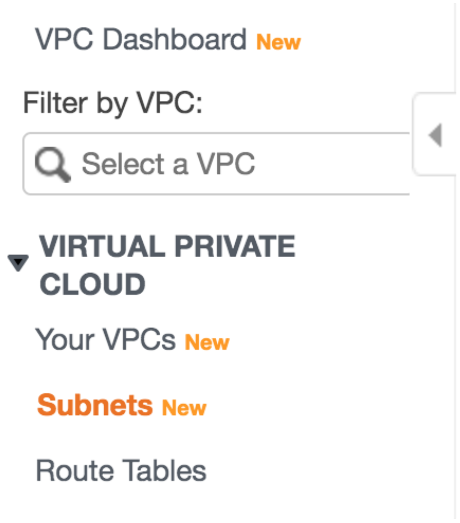

VPC Dashboard New

Filter by VPC:

Select a VPC

VIRTUAL PRIVATE CLOUD

Your VPCs New

Subnets New

Route Tables

e. Then, select from available subnets and VPCs. Take note of corresponding VPC and Subnet names. Please, make sure the subnet belongs in the selected VPC.

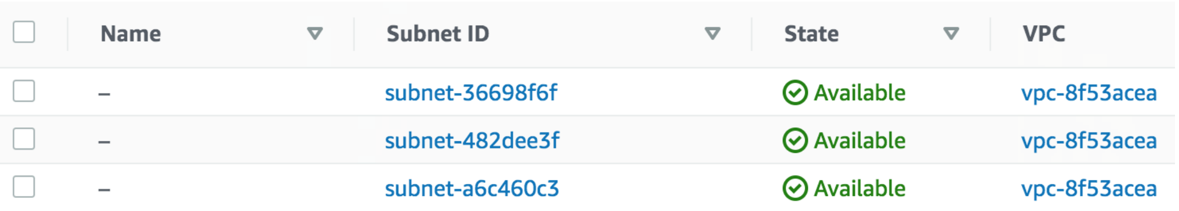

□	Name	Subnet ID	State	VPC
□	–	subnet-36698f6f	✓ Available	vpc-8f53acea
□	–	subnet-482dee3f	✓ Available	vpc-8f53acea
□	–	subnet-a6c460c3	✓ Available	vpc-8f53acea

f. If no VPCs or subnets are listed, proceed to create a VPC and a subnet as per AWS' instructions.

2. To look for the key pair used to access the host:

    a. Go to the AWS console.

    b. Once in the console, select the region where you will deploy the solution.

    c. Look for the service "EC2".

    d. Go to the "Key pairs" section, under "Network & Security".

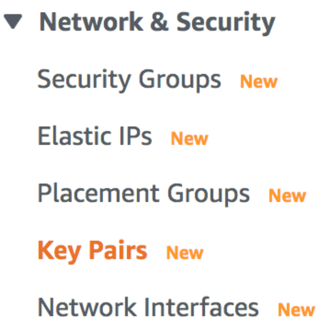

e. Take note of the name of the key pair you want to use to access your EC2 instance.

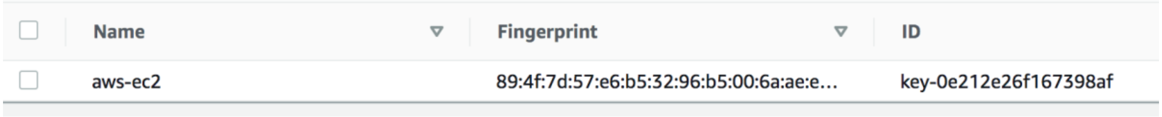

□	Name	Fingerprint	ID
□	aws-ec2	89:4f:7d:57:e6:b5:32:96:b5:00:6a:ae:e...	key-0e212e26f167398af

f. If there are no key pairs listed, create one by clicking on “Create key pair”.

To deploy Centos7 based virtual machines:

Centos 7 based EC2 instances can’t be deployed without first acknowledging the license agreement. This has to be done from the AWS marketplace, in the AWS account where the solution will be deployed. To accept the agreement, subscribe to the following marketplace solution: https://aws.amazon.com/marketplace/pp/B00O7WM7QW

You may now proceed to the solution deployment

Solution Overview:

*“**Elastic Cloud Computing**, or EC2, is a virtual machines solution optimized to run in the AWS cloud. These instances support both Linux and Windows operating systems, and all required drivers and configurations are managed by you as part of the deployment process. There is a variety of optimizations and shirt sizes to choose from, so the instance type that best adapts to your use case can be selected.”*

Parameters & Inputs:

- Input a name for your EC2 instance.
- Select an instance AMI (Different flavors of Linux available).
- Select a host tenancy (default, dedicated).
- Input the name of Key Pair (as noted in the previous section).
- Select the size of the root volume (Minimum 10 Gb for Linux and 40 Gb for Windows, Maximum 1Tb).
- Optionally, enter the URL of a bootstrap script to execute during launch.
- Optionally, enter a execution string to pass to the bootstrap script.
- Enter a CIDR block to allow access from
- Enable or disable HTTP access to the host.
- Enable or disable SSH access to the host.
- Input the ID of the VPC (as noted in the previous section).
- Input the ID of the Subnet (as noted in the previous section).

Deployment Architecture:

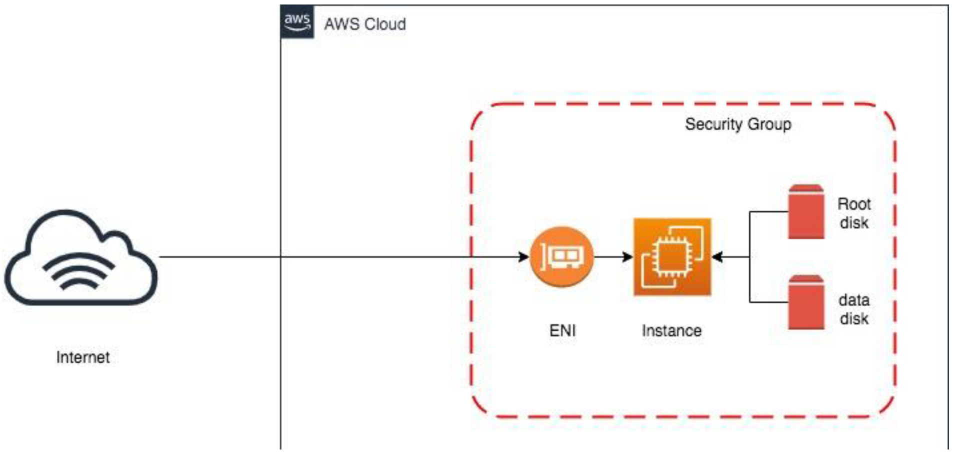

AWS Virtual Machines Deployment and Considerations

Purchase the AWS Virtual Machines V1 Click to Run Solution through StreamOne Marketplace and proceed to the Digital Locker to configure and deploy the solution.

1. Select an available AWS Region: This is the region where the solution will be deployed. Not all regions are available for every solution.

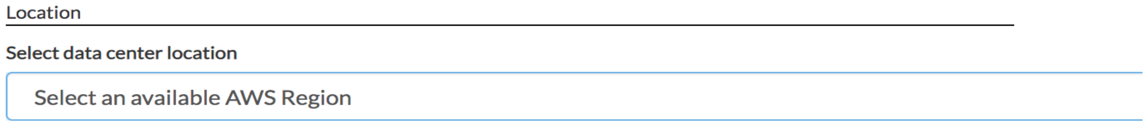

Location

Select data center location

Select an available AWS Region

2. Select a virtual machine name: This will be the name that will be displayed in the EC2 console to identify the EC2 instance. No spaces are allowed.

Basic Information

Virtual machine name

3. Select an instance AMI: This is the image of the operating system that will be deployed with the solution. ** Note: To deploy a Centos Machine, please refer to the instructions above.

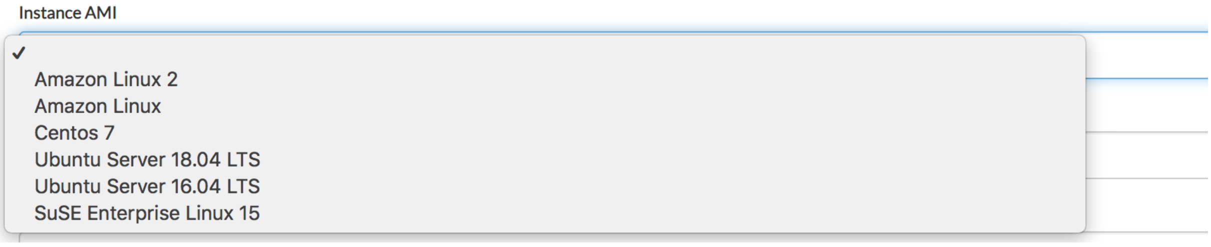

Instance AMI

✓

Amazon Linux 2
Amazon Linux
Centos 7
Ubuntu Server 18.04 LTS
Ubuntu Server 16.04 LTS
SuSE Enterprise Linux 15

4. Select a tenancy mode: The tenancy defines where the EC2 instance runs, in terms of the underlying hypervisor. If "default" is selected, the EC2 instance will be run on shared hardware. If "dedicated" is selected, the EC2 instance will be run on dedicated hardware (additional charges will apply).

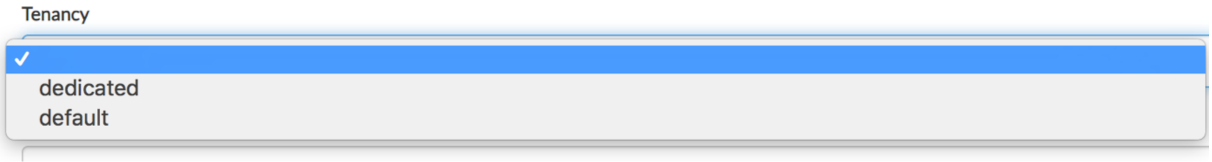

Tenancy

✓

dedicated
default

5. Input an EC2 key pair: The key pair will be used to access the EC2 instance through SSH. To obtain an available key pair or to create one, please follow the steps highlighted in the section above.

Key pair

6. Select disk sizes: Two disks will be provisioned with the EC2 instance, one for the operating system (root volume) and one for data (data volume). Use the slider to select the desired size. (For Windows instances, the minimum root volume must be 40Gb).

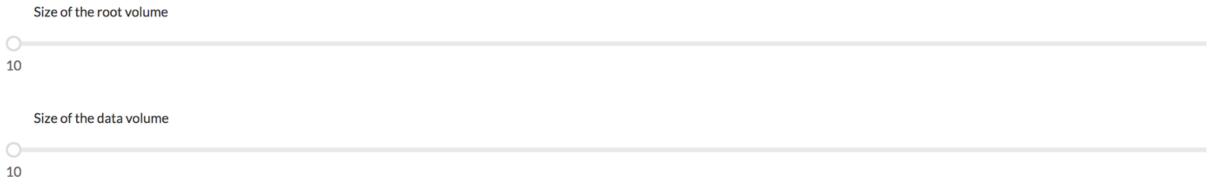

Size of the root volume

10

Size of the data volume

10

7. (Optional) Input the URL of a bootstrap script: The script defined in this field will be executed the first time the EC2 instance is booted. This script has to be a BASH file. Typically, this is used to install additional software, do hardening, or perform additional steps. This is for advanced users only.

(Optional) Bootstrap script

8. (Optional) Input an execution string for the script: If defining a bootstrap script above, it is possible that parameters need to be passed (ie. "--filesystem EFS --language python"). This is for advanced users only.

(Optional) Execution string

9. Input a CIDR block to allow SSH access: SSH access to this EC2 instance will only be allowed from the CIDR block specified in this field. If universal access is required (not recommended), input 0.0.0.0/0 as the CIDR block.

Allowed access CIDR

10. Select whether to allow HTTP access: If HTTP access to the instance is needed, mark the checkbox. Public access on port 80 will be granted in the Security Group. However, please make sure that the subnet and VPC where the EC2 instance is deployed allow access to this host.

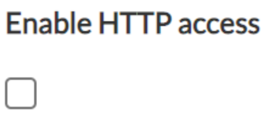
Enable HTTP access

□

11. Select whether to allow SSH access: If SSH access to the instance is needed, mark the checkbox. SSH access will be granted to the Security Group, from the address range specified in the “Allowed access CIDR”. The ssh key specified in the “AWS Keypair” will be required when you log in.

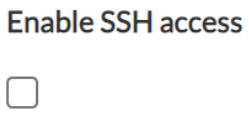
Enable SSH access

□

12. Input the VPCID: The ID of the VPC where the EC2 instance will be provisioned. Input the ID you noted in the prerequisites step.

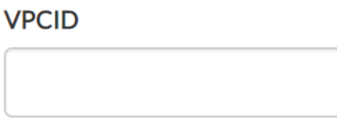
VPCID

13. Input the SubnetID: The ID of the subnet where the EC2 instance will be provisioned. This subnet must be in the same VPC as entered above. Input the ID you noted in the prerequisites step.

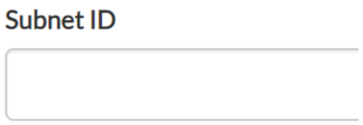
Subnet ID

14. Final check: Please validate your inputs and click “Deploy”. This will change to “Deploying…” and close the window after a few moments. You should then see the solution listed as “In Progress” in the digital locker.

------------------- Proceed to the next page to for post-deployment steps ---------------

Post-Deployment activities for AWS Virtual Machines V1 (EC2 instances)

- Verifying the state of your instance.
- Connecting to your EC2 instance.
- Modifying network access permissions.

Post Deployment Activities

Verifying the state of your instance.

When the instance is first deployed, its state will change to "running" after some minutes. If the instance doesn't start there might have been a problem during startup, so either check the deployment logs or contact our support. To verify the state, go to your AWS console. Once logged in, select the service "EC2". Go to "Instances". Your instance should appear there, with the name you selected as instance name. The state will be displayed in the "instance state" column.

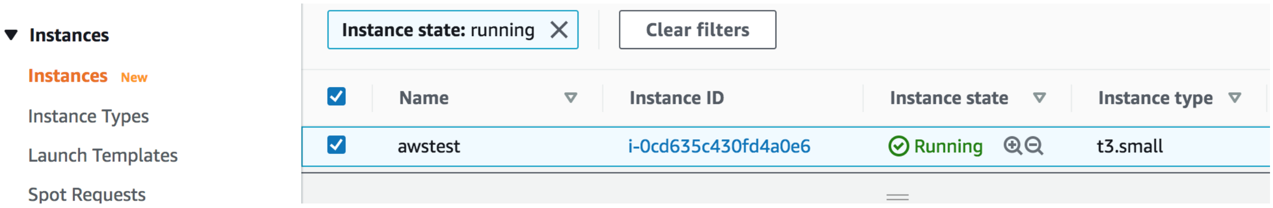

Instance state: running ✕

☑	Name	Instance ID	Instance state	Instance type
☑	awstest	i-0cd635c430fd4a0e6	✓ Running	t3.small

Connecting to your EC2 instance

It might happen that at some point you will have to connect to your EC2 instance,for example to perform administrative tasks or to install software. There are two possible waysto accomplish this: using the AWS console or connecting through SSH.

To connect through the AWS console:

1. Go to your AWS console, and then select the “EC2” service in the top left menu.

2. Go to “instances”.

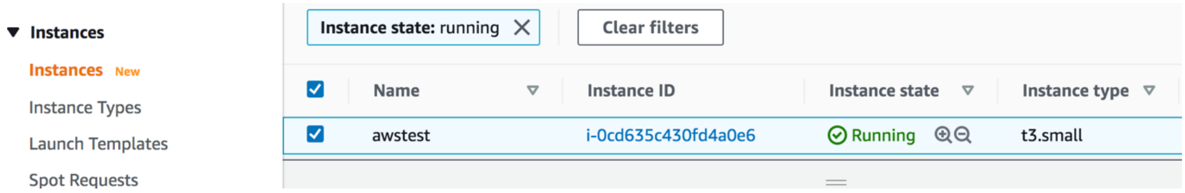

Instance state: running ✕

☑	Name	Instance ID	Instance state	Instance type
☑	awstest	i-0cd635c430fd4a0e6	Running	t3.small

3. In the top right corner, select “connect”:

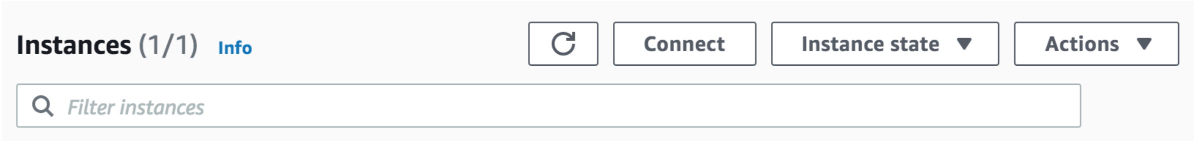

4. Select “Session Manager” as the connection method:

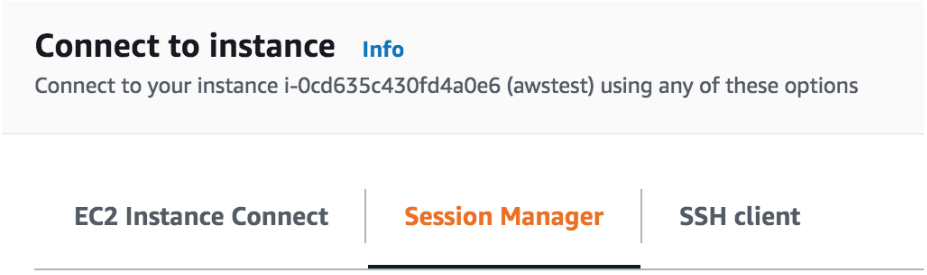
Connect to instance Info

Connect to your instance i-0cd635c430fd4a0e6 (awstest) using any of these options

5. Click on “connect” on the bottom right corner.

6. You will be presented with a web-based ssh session in a new tab.

To connect through the SSH client:

1. Follow the above steps up to step 3.
2. Select “SSH client” as connection method:

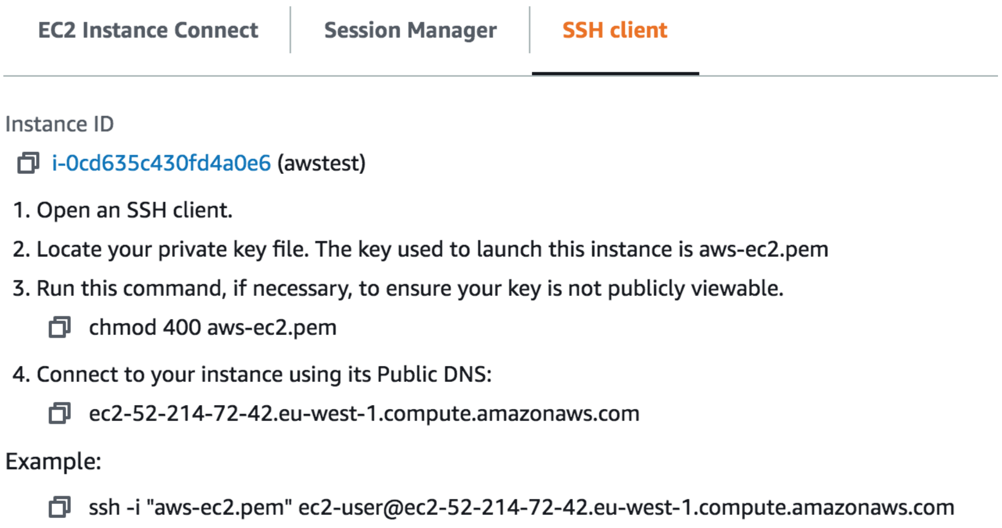

Instance ID

i-0cd635c430fd4a0e6 (awstest)

1. Open an SSH client.
2. Locate your private key file. The key used to launch this instance is aws-ec2.pem
3. Run this command, if necessary, to ensure your key is not publicly viewable.

```
chmod 400 aws-ec2.pem
```

4. Connect to your instance using its Public DNS:

```
ec2-52-214-72-42.eu-west-1.compute.amazonaws.com
```

Example:

```
ssh -i "aws-ec2.pem" ec2-user@ec2-52-214-72-42.eu-west-1.compute.amazonaws.com
```

3. Follow the detailed instructions to connect through SSH. Please note, in order to connect to the machine’s public IP address, you need to have selected the appropriate access CIDR in the field “Remote Access CIDR”, or the connection will be refused. If you have not properly configured your instance, please follow the instructions in the “Modifying network access permissions” in this section.

Modifying network access permissions.

It is entirely possible that the basic permissions issued at instance deployment will need to be modifiedlater on as demands change, or it could be that a mistake was made and the proper CIDR block was not assigned. To modify the network access permissions, follow these steps:

1. Connect to your AWS console. Select the service “EC2”.

2. Go to “instances” in the left navigation bar.

3. Select your instance:

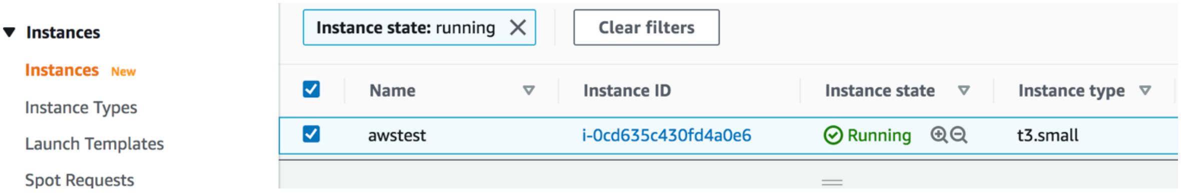

☑	Name	Instance ID	Instance state	Instance type
☑	awstest	i-0cd635c430fd4a0e6	✓ Running	t3.small

4. On the bottom half, select the “Security” tab:

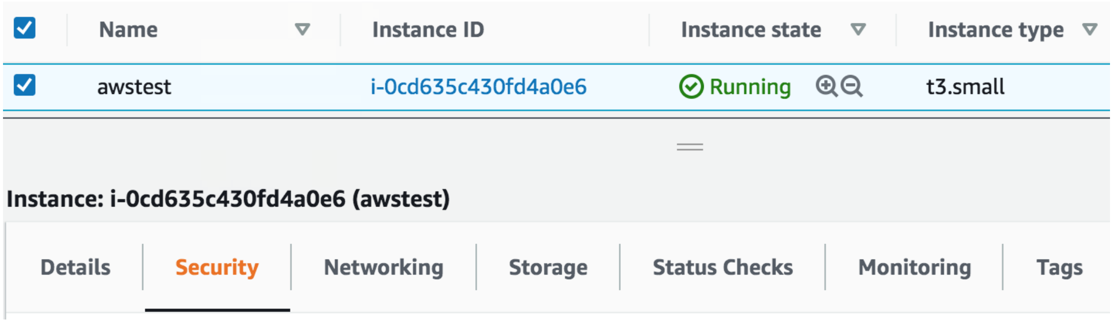

☑	Name	Instance ID	Instance state	Instance type
☑	awstest	i-0cd635c430fd4a0e6	✓ Running	t3.small

Instance: i-0cd635c430fd4a0e6 (awstest)

5. In the “security” section, look for the “Security Groups” and click in the only Security Group that will be displayed therein.

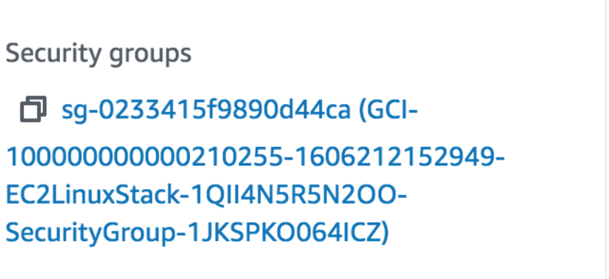

Security groups

sg-0233415f9890d44ca (GCI-1000000000000210255-1606212152949-EC2LinuxStack-1QII4N5R5N2OO-SecurityGroup-1JKSPKO064ICZ)

6. You will be taken to the “Security Groups” interface. There, scroll down until you see the “Inbound rules” section. Your access rules should be displayed there. If you selected “Allow SSH access”, a rule to port 22 will be displayed there, from the CIDR address that you selected:

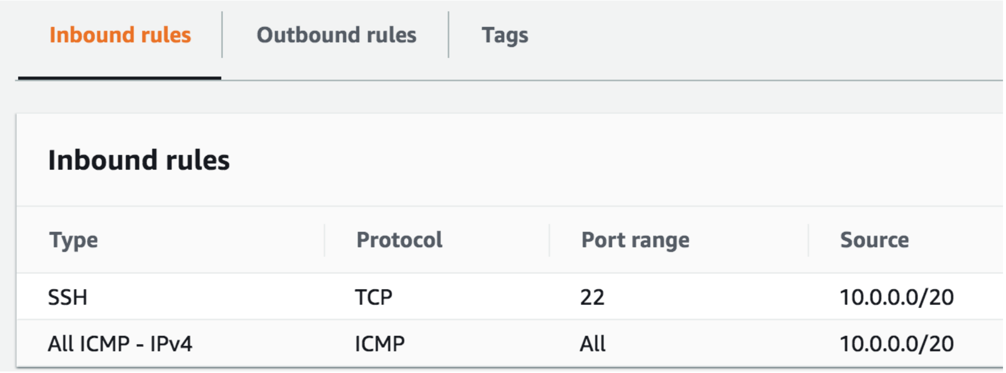

Inbound rules

Type	Protocol	Port range	Source
SSH	TCP	22	10.0.0.0/20
All ICMP - IPv4	ICMP	All	10.0.0.0/20

7. To modify or add access rules, click on “Edit inbound rules”, in the top right corner of the “inbound rules” section.

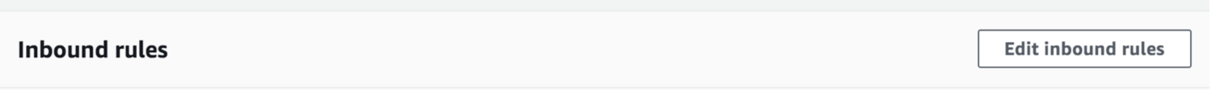

Inbound rules

8. You will be able to modify the “source” and well as the port, or add additional ports. Remember that it is possible that you will have to edit the Security Groups assigned to the subnet to allow access to the host.

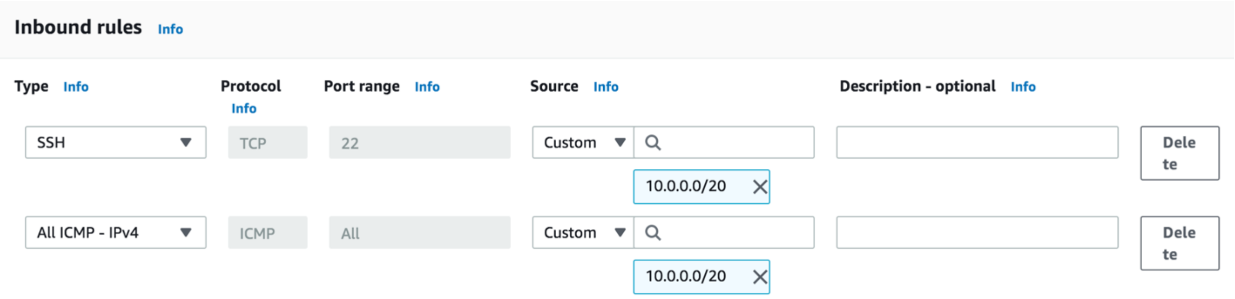

Inbound rules Info

Type Info	Protocol Info	Port range Info	Source Info	Description - optional Info	
SSH ▼	TCP	22	Custom ▼ 10.0.0.0/20 ✕		Delete
All ICMP - IPv4 ▼	ICMP	All	Custom ▼ 10.0.0.0/20 ✕		Delete