Amazon VPC Networking

Amazon Virtual Private Cloud (Amazon VPC) enables you to launch AWS resources into a virtual network that you've defined. This virtual network closely resembles a traditional network that you'd operate in your own data center, with the benefits of using the scalable infrastructure of AWS. So, in simpler words, Amazon Virtual Private Cloud (Amazon VPC) enables the users to define some virtual network and then launch the AWS resources into that virtual network. It gives you full control over various network environments, resources, connectivity, and security. Moreover, it defines how a network should communicate across different Availability Zones or regions. Users have an option of easy customization of the network configuration for their Amazon Virtual Private Cloud(VPC).

Before we go deep into the features or components and vpc cloud connectivity options, let us first take look at basic networking concepts.

IPv4 Address Classes

Class A 1 – 127 255 .00000000 .00000000.00000000 :usable ip Add 16,777,214
Net ID . Host ID. Host ID . Host ID → 8bits Subnet mask

Class B 128 – 191 255 . 255 . 00000000.00000000. usable IP Address:65,534
Net ID . Net ID . Host ID . Host ID → 16bits Subnet mask

Class C 192 – 223 255 . 255 . 255 . 00000000 :usable Ip Address 254
Net ID . Net ID . Net ID .Host ID → 24bits Subnet mask

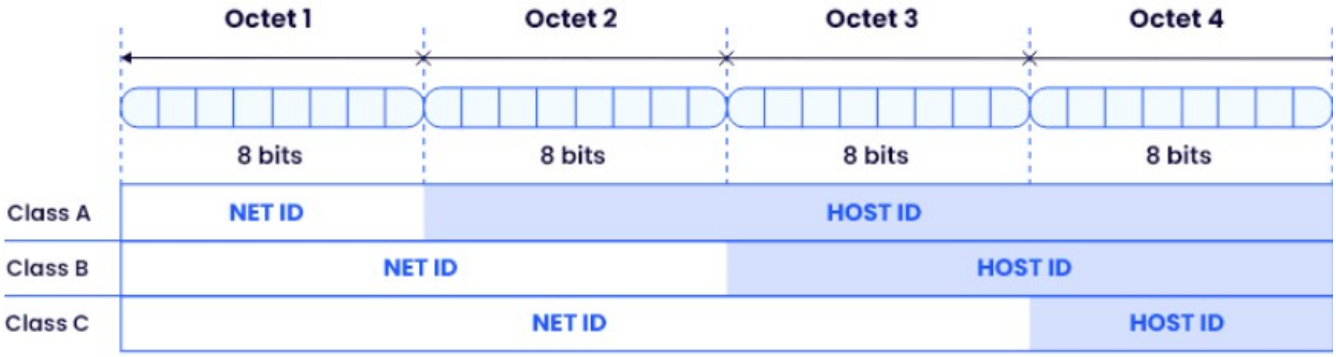

Note: The subnet mask help us to know which portion is the Network ID and Host ID

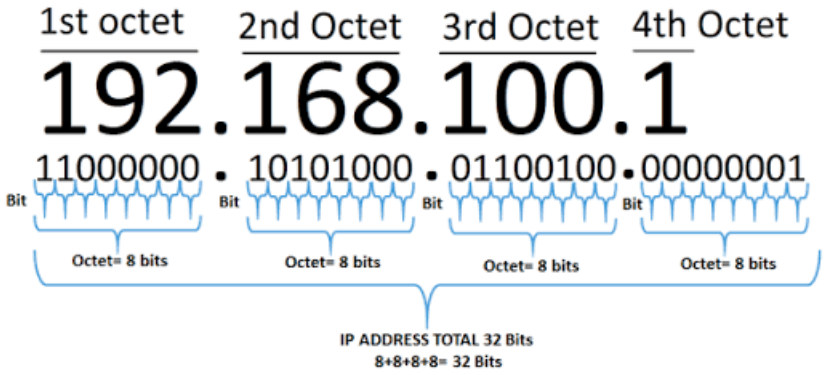

Speciality Address Ranges
Loopback - Only the single 127.0.0.1 address is used, addresses 127.0.0.0 to 127.255.255.255 are reserved. Any address within this block will loop back to the local host.

	Private Address Space
Class A	10.0.0.0 to 10.255.255.255
Class B	172.16.0.0 to 172.31.255.255
Class C	192.168.0.0 to 192.168.255.255

Example: 192.168.0.1 is a class C IP address with 192.168.0 constituting the network ID and 1, the Host ID. → 24bits subnet mask. Written as 192.168.0.1/24.

Classless Inter Domain Routing (CIDR)

Classless Inter-Domain Routing, or CIDR, is a means of allocating internet protocol addresses also known as host addresses, more efficiently compared to the traditional classful network addressing system.

A CIDR IP address may look something like this: **123.45.67.89/12**. This IP address contains two groups of numbers:

- **Network prefix** (123.45.67.89): The binary configuration of a network address
- **Suffix** (/12): The indication of how many bits are in the entire CIDR address

IPv4 addresses are 32-bits long, and while the first 12 bits represent network addresses, the remaining 20 bits represent the available host addresses. It's worth mentioning that every network, by default, has only one subnet containing all host addresses. The 2nd octet which constitute of 8bits is shared between the Net ID and the Host ID

Since CIDR is not restrained by class, it can organize IP addresses into multiple subnets regardless of the IP addresses' value. Compared to traditional subnetting, CIDR enables routers to reach network traffic destinations much quicker.

VPC Overview

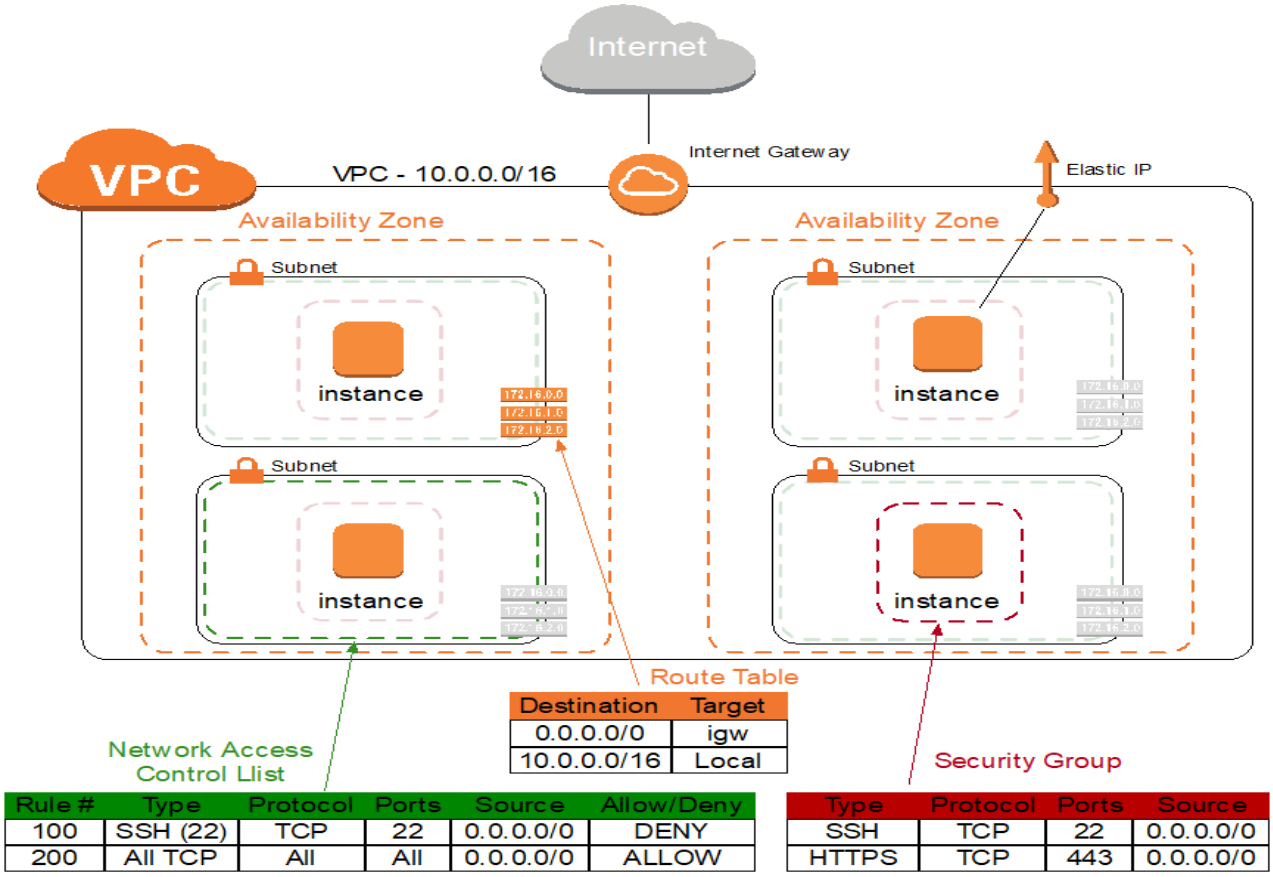

Features
The following features help you configure a VPC and your entire cloud infrastructure to provide the connectivity that your applications need:
1 :Region
An Aws region is a collection of Availability Zones. You can think of a Region as containing multiple data centers within the same geographic area, with no shared fault domains.

2:Virtual Private Cloud (VPC)

A VPC is a set of contained subnets with a common Classless Inter-Domain Routing (CIDR) block (up to a /16 netmask) running in a single geographic area (Region) across multiple data centers (Availability Zones).

A VPC is like a virtual data center, except that it's physically spread out across Availability Zones. VPCs have network connectivity within the Region in which they are created. You can use Internet connectivity, virtual private network (VPN) connectivity, and VPC peering to connect VPCs to other networks

3:Availability Zone

An Availability Zone is a set of buildings, Internet uplinks, and power. You can think of it as a data center, but some Availability Zones contain more than one physical data center.

4: **Subnets**

A subnet is specific to an Availability Zone, and there can be multiple subnets per Availability Zone. Each subnet has a single route table association, and each subnet is associated with a network access control list (network ACL). AWS uses five addresses in each subnet—the first four and the last address.

You may see references to *public* and *private* subnets. This is a shortcut for saying whether a subnet has a direct route to the Internet or not; otherwise, there are no differences. Typically, a public subnet has direct access to the Internet.
A private subnet has no external access—it connects through a virtual private gateway (VGW) to an on-premises data center, or uses Network Address Translation (NAT) for Internet-bound traffic. Subnets also have a setting called **Auto-assign Public IP**, which can give instances public IP addresses by default (and can be overridden per instance).

5:**IP addressing**
You can assign IPv4 addresses and IPv6 addresses to your VPCs and subnets. You can also bring your public IPv4 and IPv6 GUA addresses to AWS and allocate them to resources in your VPC, such as EC2 instances, NAT gateways, and Network Load Balancers.

Public IP address: A public IP address is a publicly routable address that is reachable from the Internet. There are two types of public IP addresses in a VPC: dynamically assigned addresses and Elastic IP addresses. By default, instances connected to the Internet receive a dynamically assigned public address. When the instance is terminated, the address is released. If you require that address to be persistent, you need an Elastic IP address. The public address will not change while the instance is operating.

Elastic IP address (EIP): An Elastic IP address is a static public IP address that is applied to an ENI, and can be associated to another instance after an instance is terminated. This makes it easier for things like ACL rules, DNS entries, and whitelisting in other systems. Every public address (whether it's an Elastic IP address or not) has a private address associated with it. This private address is static unless the Elastic IP address is moved to another subnet. Also, remember that subnets are local to a single Availability Zone.

6:**Routing**
Route Table: They are the set of rules used to decide where the network traffic has to be managed. It specifies the destination i.e IP address and target. The target can be Internet gateway, NAT gateway, Virtual private gateway, etc. With the use of route tables, users can determine where the network traffic will be directed from your subnet or gateway.
Every subnet has a route table, and a single route table can be associated with multiple subnets.
For example, private subnets might have route tables with a default route to a NAT gateway, and public subnets might have route tables with entries to the Internet using an Internet gateway. Private subnets can also have routes back to on-premises data centers using a virtual private gateway.
A default route table is automatically attached to new subnets unless they're manually associated with a different route table. For a higher level of control and flexibility, create a new route table for subnets instead of using the default route table. If a service is local to an Availability Zone (such as NAT Gateway), we recommend defining route tables per Availability Zone.
7:**Security Group: It consists set of firewalls rules that control the traffic for your sample. You can have a single security group associated with multiple instances.**

A security group is like an access control list (ACL) typically applied to routers and firewalls, except that it is applied directly to the instance. Security groups handle the bulk of security in AWS to protect instances, and allow for fine-grained security per instance.

Security groups are stateful and track TCP, UDP, and ICMP connection status. Security groups contain security group rules, which are like ACL entries. Security group rules are whitelist only, and contain an implicit DENY ANY rule. Security groups can reference other security groups, acting similar to a network/object group in firewalls. This capability enables you to permit traffic from certain groups of instances without requiring the use of IP addresses. Overall, security groups are like a Layer 4 distributed firewall.

8:*Network Access Control Lists (NACL):* It is an optional layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets. It adds an additional layer of security to your VPC. Network ACL: Network ACLs are applied per subnet, and provide a broader stroke than security groups.
Network ACLs are ordered and have PERMIT and DENY actions, with an implicit DENY ANY. Network ACLs, unlike security groups, are stateless. Also, network ACLs work only on CIDR ranges and can't reference other objects such as instances. The default network ACL allows all packets.

If you want to block SSH for an entire subnet, you could add a DENY entry for TCP port 22 in the subnet's network ACL. Otherwise, you would need to block SSH on each instance's security group.

9:Elastic network interface (ENI): An elastic network interface (ENI) is like a virtual network interface card (NIC). You can apply multiple ENIs to an instance, and move the ENI to another instance in the same subnet. Multiple Elastic IP addresses can be applied to an ENI. An ENI has a dynamically assigned private address in the assigned subnet, and can optionally have a dynamically assigned public IP address as well. Multiple addresses can be assigned to an ENI.

10:NAT gateways
A NAT gateway is a Network Address Translation (NAT) service. You can use a NAT gateway so that instances in a private subnet can connect to services outside your VPC but external services cannot initiate a connection with those instances.

When you create a NAT gateway, you specify one of the following connectivity types:

- **Public** – (Default) Instances in private subnets can connect to the internet through a public NAT gateway, but cannot receive unsolicited inbound connections from the internet. You create a public NAT gateway in a public subnet and must associate an elastic IP address with the NAT gateway at creation. You route traffic from the NAT gateway to the internet gateway for the VPC. Alternatively, you can use a public NAT gateway to connect to other VPCs or your on-premises network. In this case, you route traffic from the NAT gateway through a transit gateway or a virtual private gateway.
- **Private** – Instances in private subnets can connect to other VPCs or your on-premises network through a private NAT gateway. You can route traffic from the NAT gateway through a transit gateway or a virtual private gateway. You cannot associate an elastic IP address with a private NAT gateway. You can attach an internet gateway to a VPC with a private NAT gateway, but if

you route traffic from the private NAT gateway to the internet gateway, the internet gateway drops the traffic.

The NAT gateway replaces the source IP address of the instances with the IP address of the NAT gateway. For a public NAT gateway, this is the elastic IP address of the NAT gateway. For a private NAT gateway, this is the private IPv4 address of the NAT gateway. When sending response traffic to the instances, the NAT device translates the addresses back to the original source IP address.

External Connectivity

11: Internet gateway

An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. It supports IPv4 and IPv6 traffic. It does not cause availability risks or bandwidth constraints on your network traffic.

An internet gateway enables resources in your public subnets (such as EC2 instances) to connect to the internet if the resource has a public IPv4 address or an IPv6 address. Similarly, resources on the internet can initiate a connection to resources in your subnet using the public IPv4 address or IPv6 address. For example, an internet gateway enables you to connect to an EC2 instance in AWS using your local computer.

An internet gateway provides a target in your VPC route tables for internet-routable traffic. For communication using IPv4, the internet gateway also performs network address translation (NAT). For communication using IPv6, NAT is not needed because IPv6 addresses are public