VPC Essentials

RTs = Route Tables

What is a Route Table?

Simplified Definition:
The AWS definition is simple enough, so let's jump right down to it!

AWS Definition:
"A route table contains a ***set of rules***, called ***routes***, that are used to ***determine where network traffic is directed***."

NOTE: Your "default" VPC already has a ***"main"*** route table.

Route Tables:

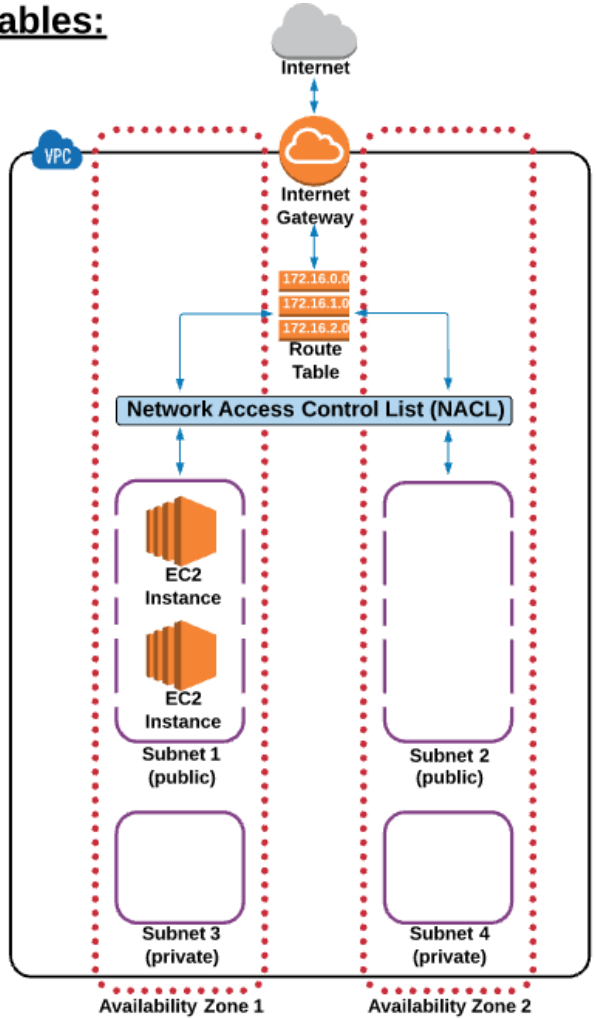

Detach internet gateway

Route Tables:

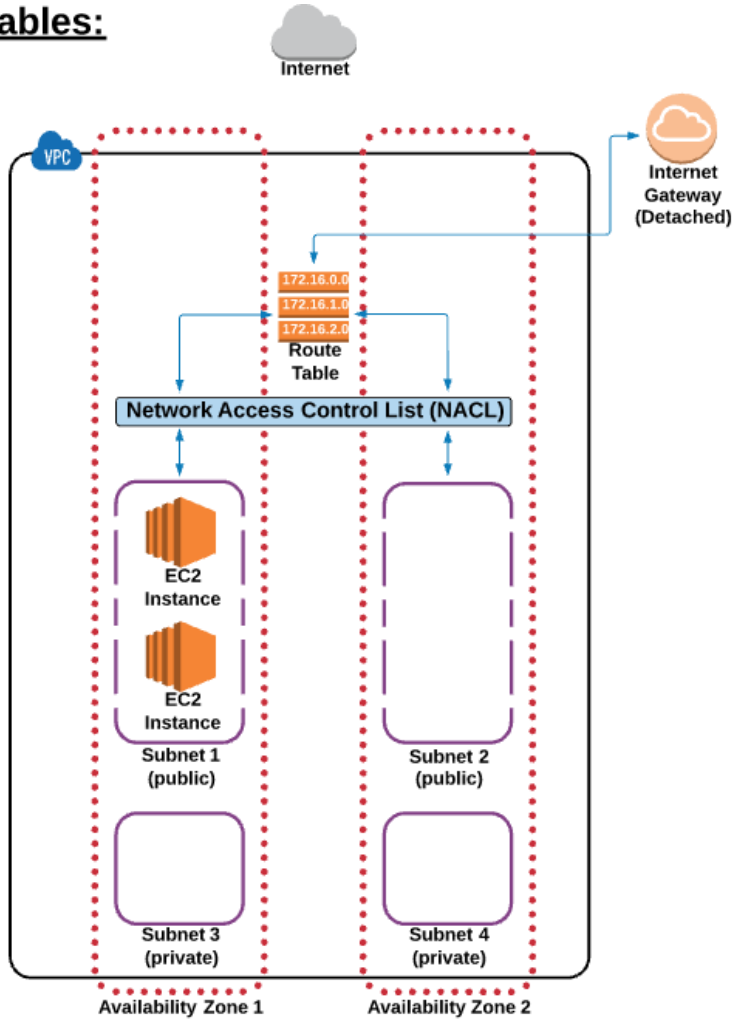

Create a new IGW and attach to VPC

Route Tables:

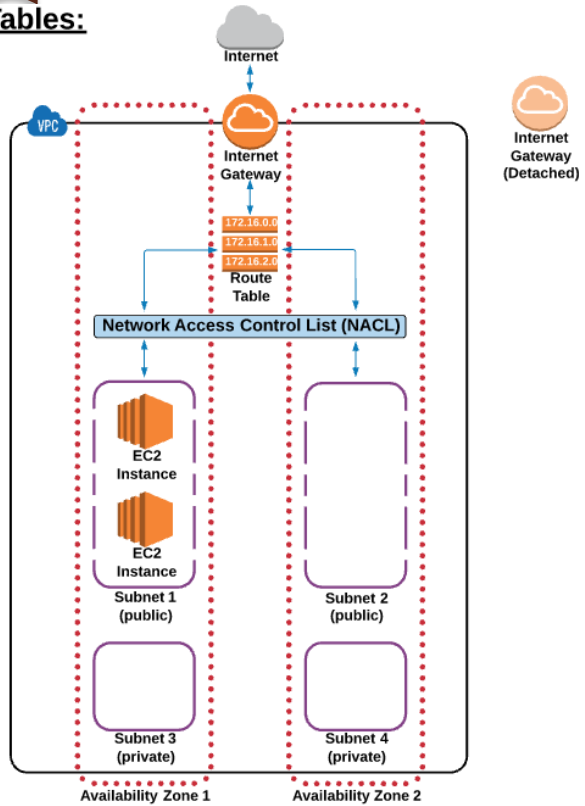

Route table rules and details you need to know:
(1) Unlike an IGW, you can have multiple "active" route tables in a VPC
(2) You cannot delete a route table if it has ***"dependancies"*** (associated subnets)

NACLs = Network Access Control Lists

What is a NACL?

Simplified Definition:
The AWS definition is simple enough, so let's jump right down to it!

AWS Definition:
A network access control list (NACL) 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***.

NOTE: Your "default" VPC already has a NACL in place and associated with the default subnets.

NACLs:

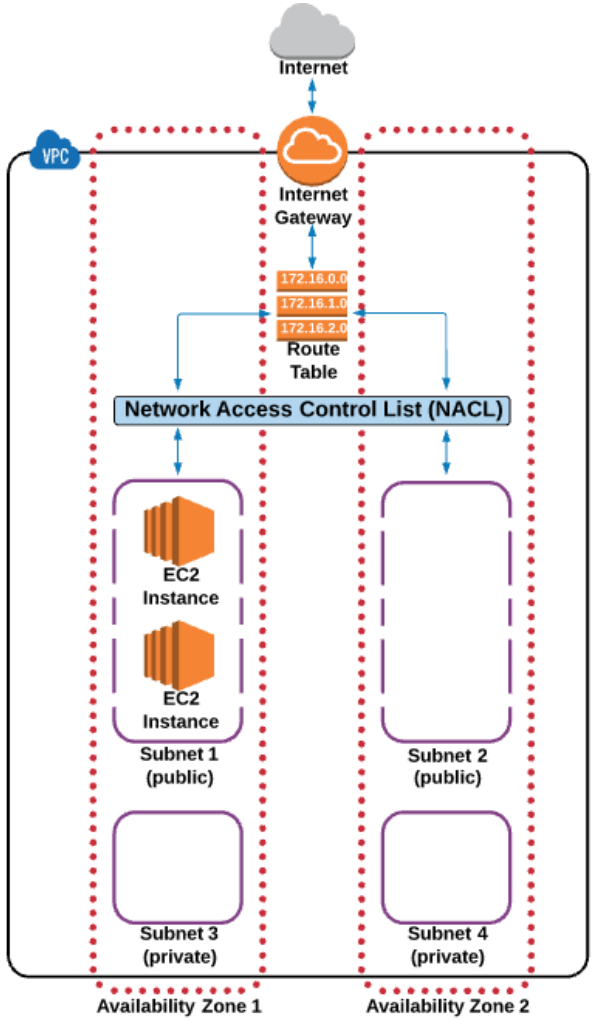

NACLs:

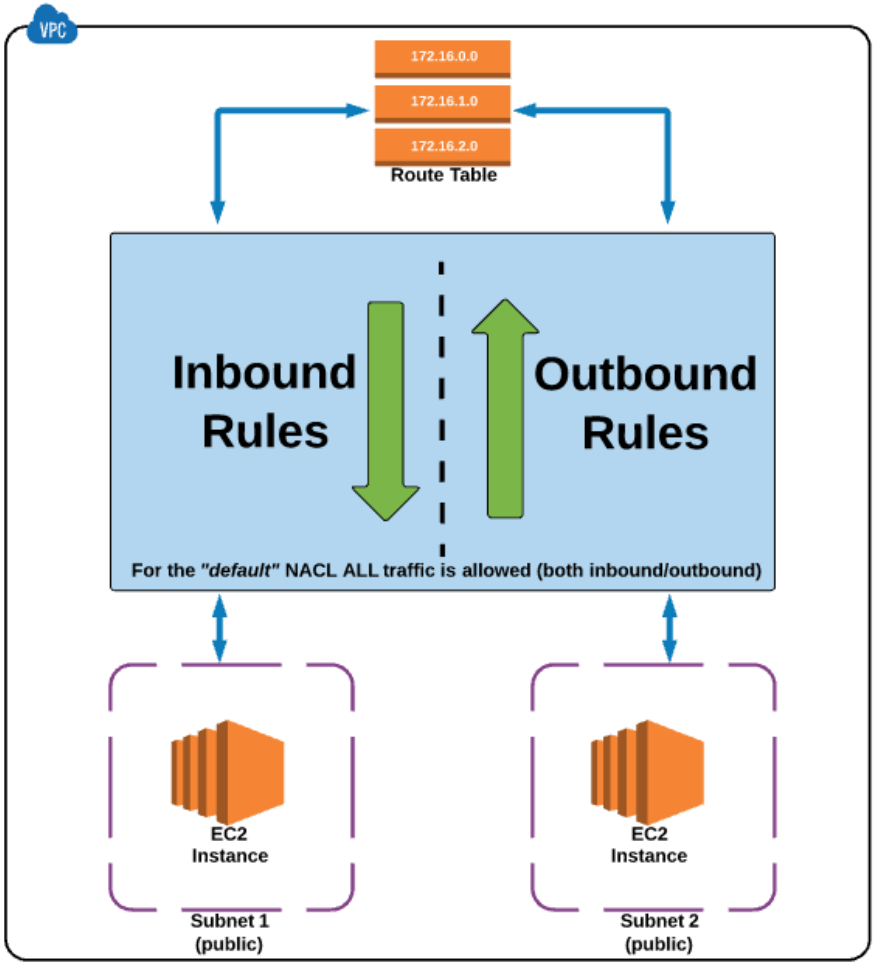

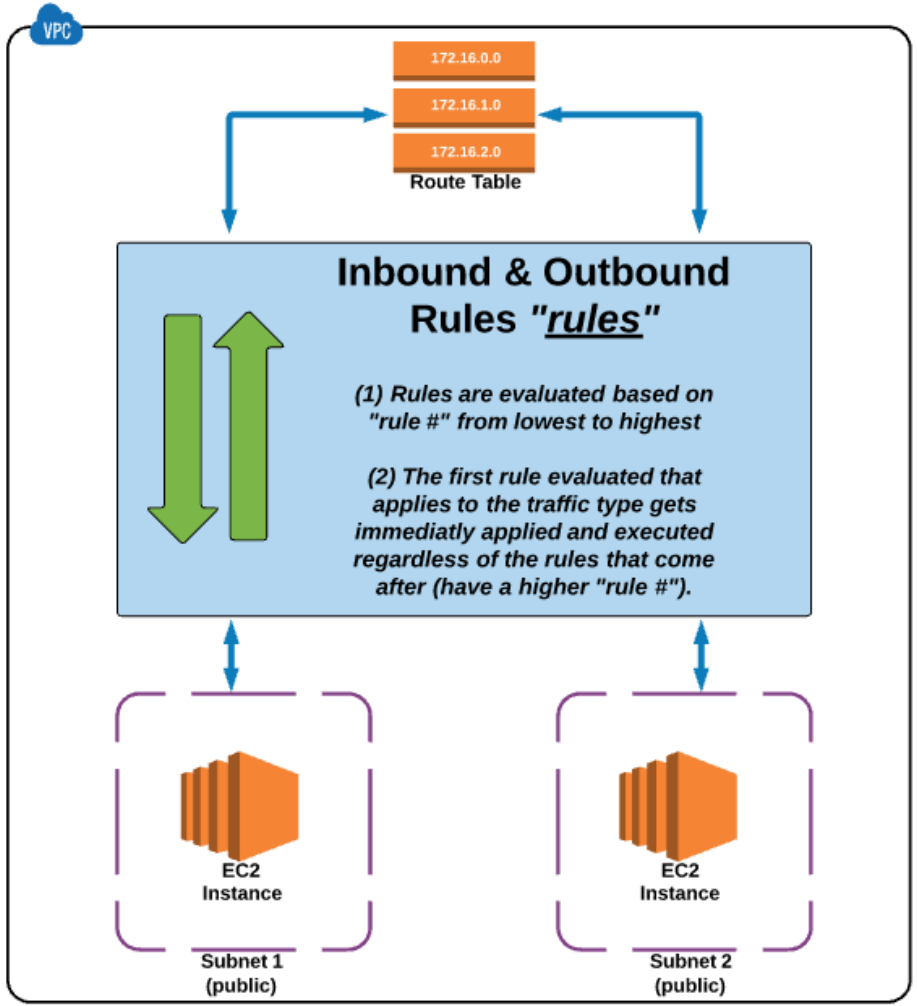

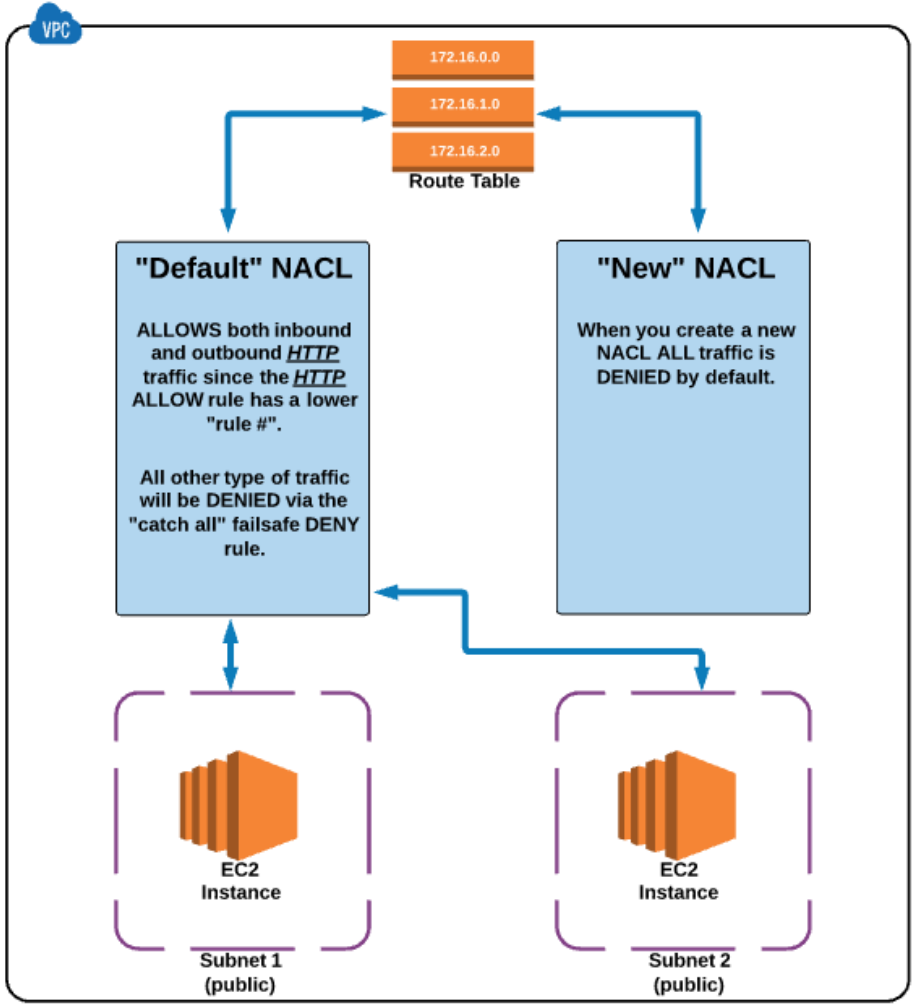
VPC
172.16.0.0
172.16.1.0
172.16.2.0
Route Table
"Default" NACL
ALLOWS both inbound and outbound HTTP traffic since the HTTP ALLOW rule has a lower "rule #".
All other type of traffic will be DENIED via the "catch all" failsafe DENY rule.
"New" NACL
When you create a new NACL ALL traffic is DENIED by default.
EC2 Instance
EC2 Instance
Subnet 1 (public)
Subnet 2 (public)

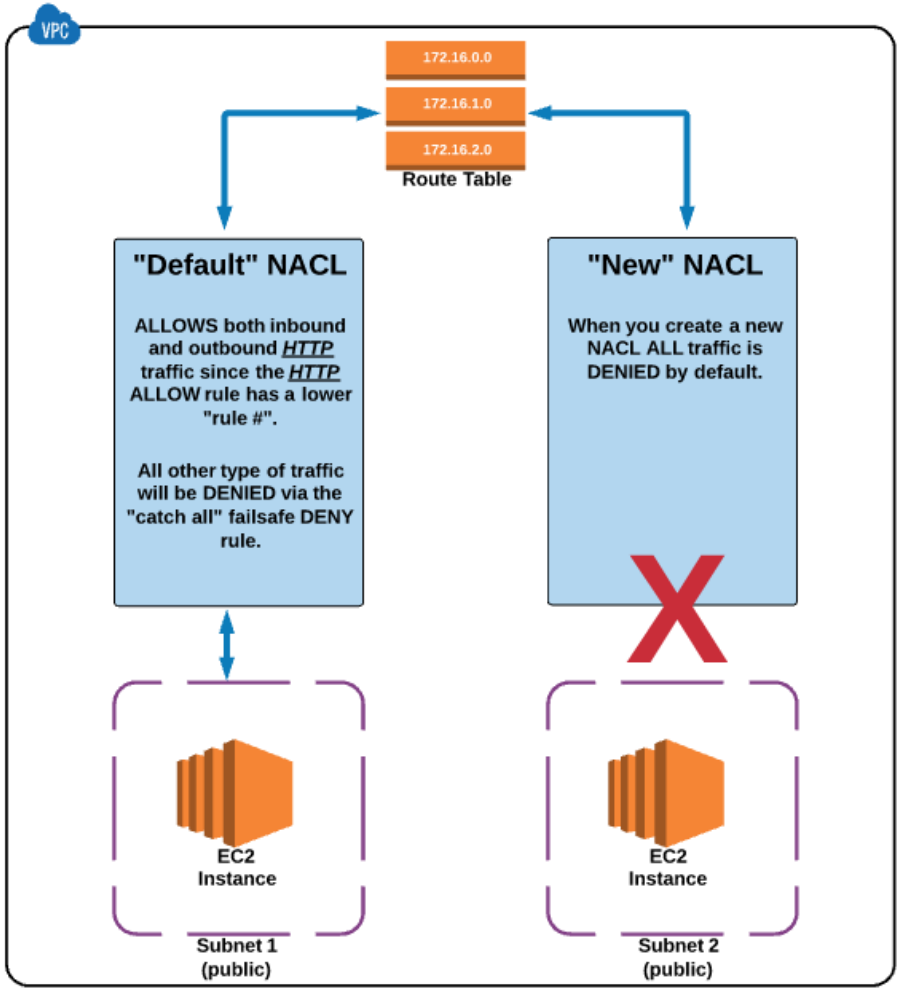
VPC
172.16.0.0
172.16.1.0
172.16.2.0
Route Table
"Default" NACL
ALLOWS both inbound and outbound HTTP traffic since the HTTP ALLOW rule has a lower "rule #".
All other type of traffic will be DENIED via the "catch all" failsafe DENY rule.
"New" NACL
When you create a new NACL ALL traffic is DENIED by default.
EC2 Instance
EC2 Instance
Subnet 1 (public)
Subnet 2 (public)

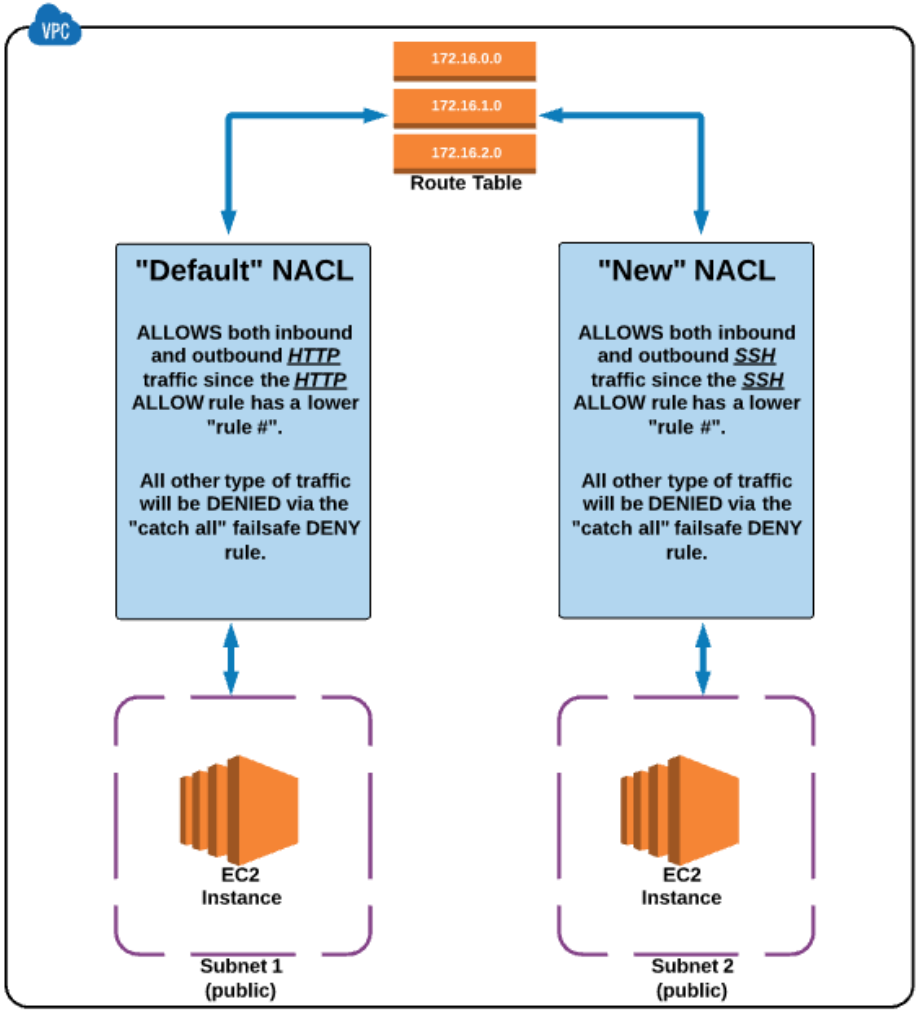

NACL rules and details you need to know (recap):

(1) Rules are evaluated from lowest to highest based on "rule #"

(2) The first rule found that applies to the traffic type is immediatly applied, regardless of any rules that come after it (have a higher "rule #").

(3) The "default" NACL allows all traffic to the default subnets.

(4) Any new NACLs you create DENY all traffic by default.

(4) A subnet can only be associated with ONE NACL as a time.

(6) An NACL allows or denies traffic from entering a subnet. Once inside the subnet, other AWS resources (i.e. EC2 instances) may have an additional layer of security (security groups).

Subnets

What is a Subnet?

Simplified Definition:
A subnet, shorthand for subnetwork, is a sub-section of a network. Generally, a subnet includes all the computers in a specific location. Circling back to the "home network" anology we used in the VPC Baics lesson - if you think about your ISP being a network, then your home network can be considered a subnet of your ISP's network.

AWS Definition:
"When you create a VPC, it spans all of the Availability Zones in the region. After creating a VPC, ***you can add one or more subnets in each Availability Zone***. Each subnet ***must reside entirely*** within one Availability Zone and ***cannot span zones***."

NOTE: Your "default" VPC already has a subnets created by default.

Subnets

Subnets:

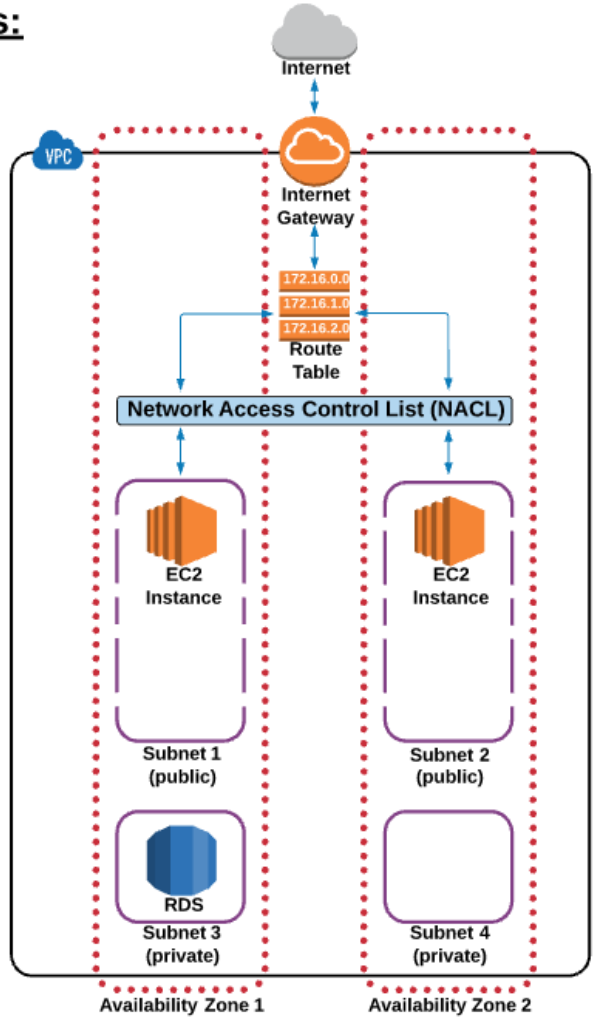

Subnets:

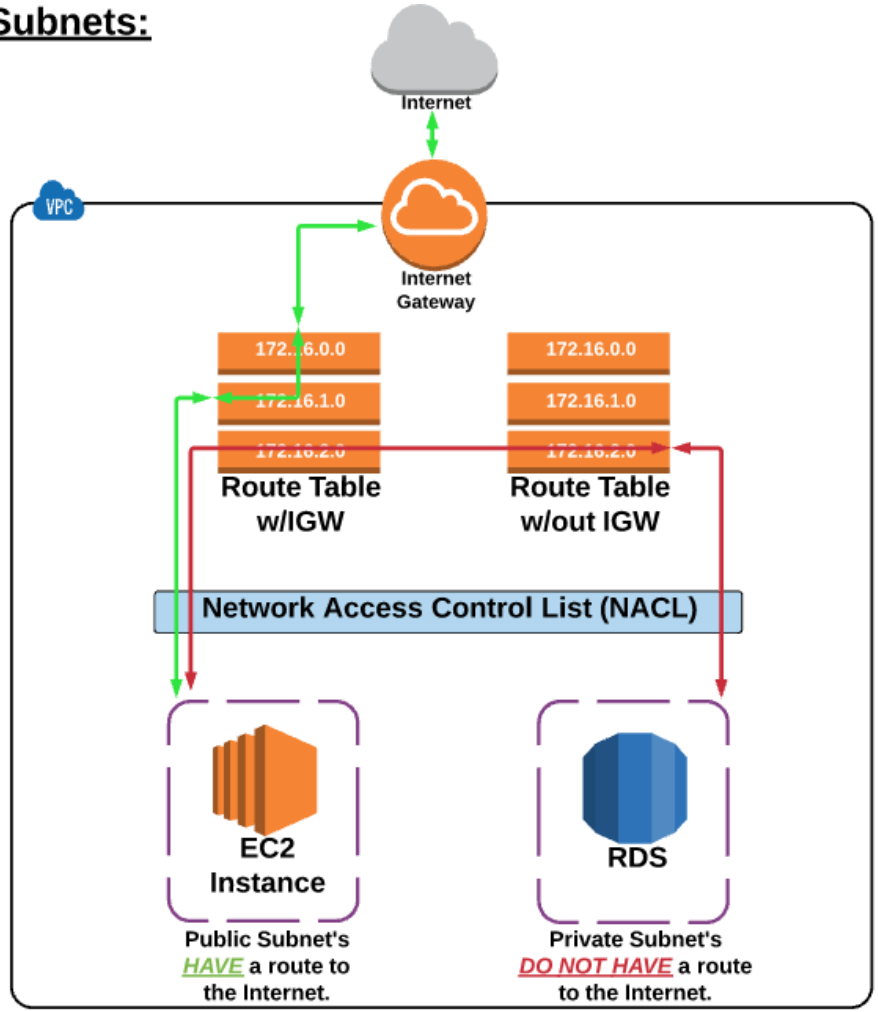

Subnets:

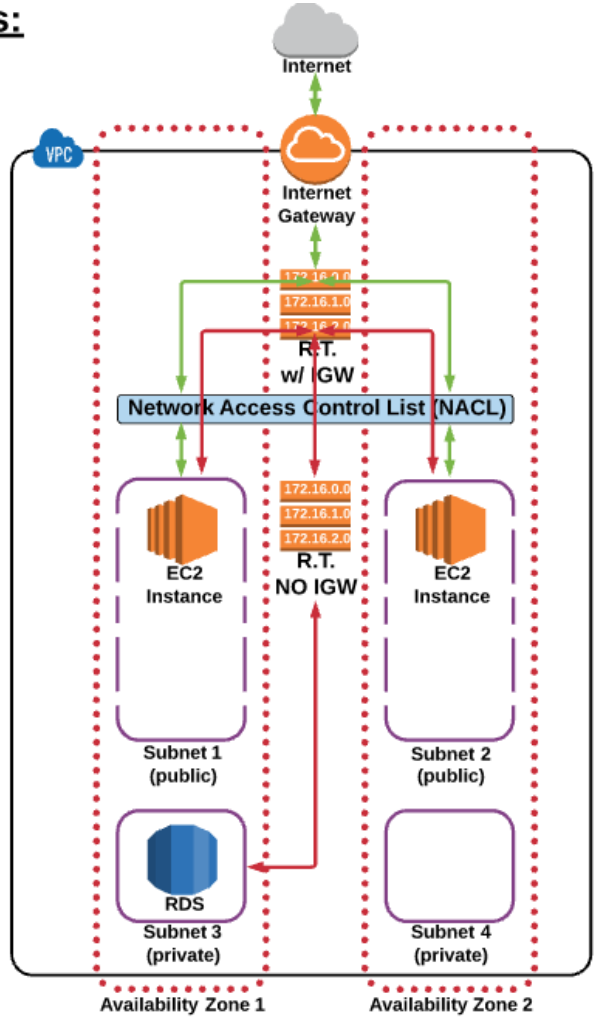

NOTE: Your "default" VPC already has a subnet created by default.

Subnet rules and details you need to know:
(1) Subnets MUST be associated with a route table.
(2) A ***PUBLIC*** subnet ***HAS*** a route to the Internet.
(3) A ***PRIVATE*** subnet ***does NOT have*** a route to the Internet.
(4) A subnet is located in ONE specific Availability Zone.

Subnets

Availability Zones

Availability Zones and VPCs:

Simplified Definition/Explanation:
Any AWS resource that you launch (like EC2/RDS) must be placed in a VPC subnet. Any given subnet must be located in an Availbility Zone. You can (and should) utilize multiple Availabilty Zones to create redundacy in your architecture. This is what allows for ***High Availabilty*** and ***Fault Toleratent*** systems.

AWS Definition/Explanation:
"When you create a ***VPC, it spans all of the Availability Zones in the region***. After creating a VPC, you can add ***one or more subnets in each Availability Zone***. Each subnet must reside entirely within one Availability Zone and cannot span zones.

Availability Zones are distinct locations that are engineered to be isolated from failures in other Availability Zones. By launching instances in separate Availability Zones, you can protect your applications from the failure of a single location."

NOTE: Your "default" VPC already has a subnet created by default.

Availability Zones

Availability Zones:

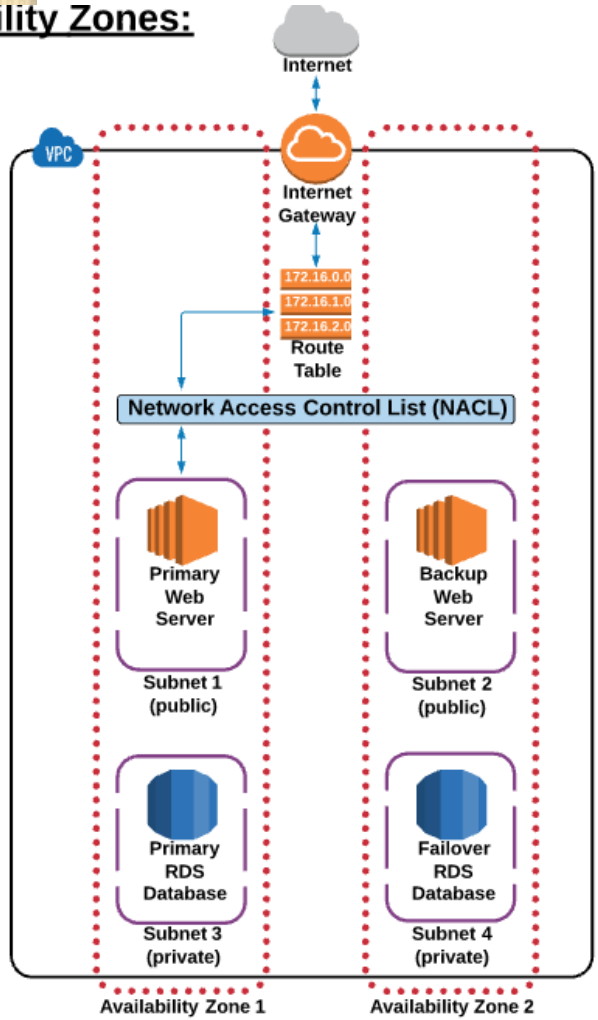

Availability Zones:

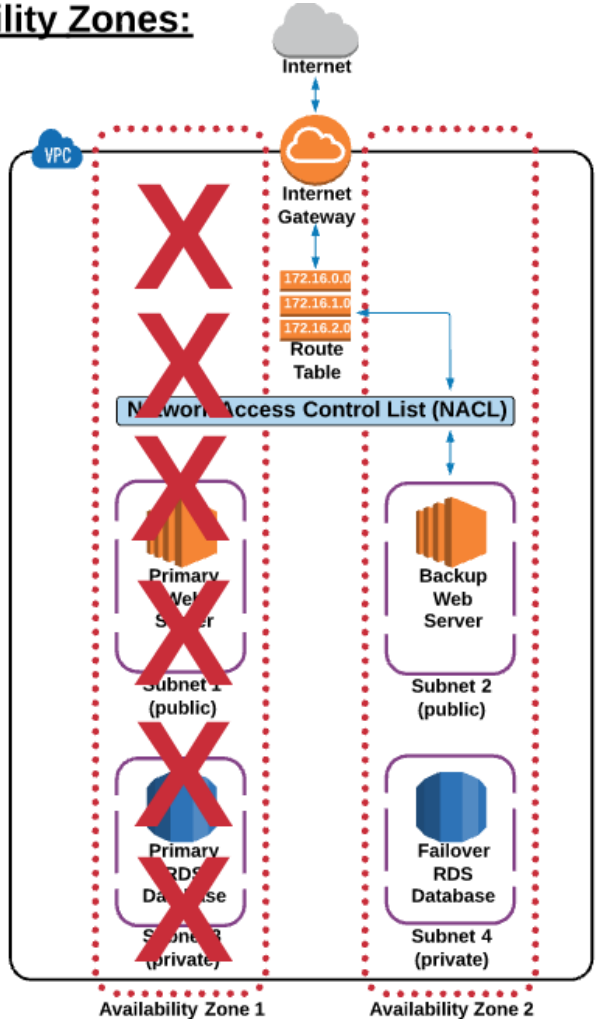

Availability Zones

High Availabilty:

Creating your architecture in such a way that your "system" is always available (or has the least amount of downtime as possible).

What High Availabilty "Sounds" Like:
(1) "I can always access my data in the cloud"
(2) "My website never crashes and is always availaible to my customers"

Fault Tolerant:

The ability of your "system" to withstand failures in one (or more) of its components and still remain availabile.

What Fault Tolerant "Sounds" Like:
(1) "One of my web servers failed, but my backup server immediatly took over"
(2) "If someting in my system fails, it can repair itself."

Availability Zones

Availability Zones and VPCs:

Simplified Definition/Explanation:
Any AWS resource that you launch (like EC2/RDS) must be placed in a VPC subnet. Any given subnet must be located in an Availbility Zone. You can (and should) utilize multiple Availabilty Zones to create redundacy in your architecture. This is what allows for ***High Availabilty*** and ***Fault Tolerateny*** systems.

AWS Definition/Explanation:
"When you create a ***VPC, it spans all of the Availability Zones in the region***. After creating a VPC, you can add ***one or more subnets in each Availability Zone***. Each subnet must reside entirely within one Availability Zone and cannot span zones.

Availability Zones are distinct locations that are engineered to be isolated from failures in other Availability Zones. By launching instances in separate Availability Zones, you can protect your applications from the failure of a single location."

NOTE: Your "default" VPC already has a subnet created by default.

Availability Zones

Proper traffic routing into and out of our AWS Virtual Private Cloud (VPC).

- One IGW attached to the VPC.
- One route table with a route to the Intenet.
- One route table without a route to the Internet.
- Two public subnets - each in a seperate Availabilty Zone.
- Two private subnets - each in a seperate Availabilty Zone.

Lets try this