

# Amazon Elastic Cloud Compute (EC2): EC2 Networking and Performance Scenarios



**Andru Estes**


Principal Author

 andru-estes





# **Configuring an Elastic Network Interface (ENI)**



## Elastic Network Interface (ENI)


Logical networking component that is part of a VPC and is meant to represent a virtual network card for compute resources



# Elastic Network Interfaces Concepts



Created and bound  
within a single  
Availability Zone



Attach, detach, and  
reattach as needed to  
any EC2 in the same AZ



ENIs offer many  
attributes that can be  
customized



# Customizable ENI Attributes

Primary private IPv4 address from the subnet CIDR

Primary private IPv6 address from the subnet CIDR

Secondary private IPv4 addresses from your subnet CIDR

One Elastic IP address (IPv4) for each private IPv4 address

Security groups with your customized rules

One ephemeral, public IPv4 address





## Ephemeral Public IPv4 Addresses

When you assign a public IP to your instances in a public subnet, the IP address is from Amazon's pool of public IPv4 addresses

When it is disassociated from the instance, it goes back into that pool of IP addresses

Your IP is released when the following occurs:

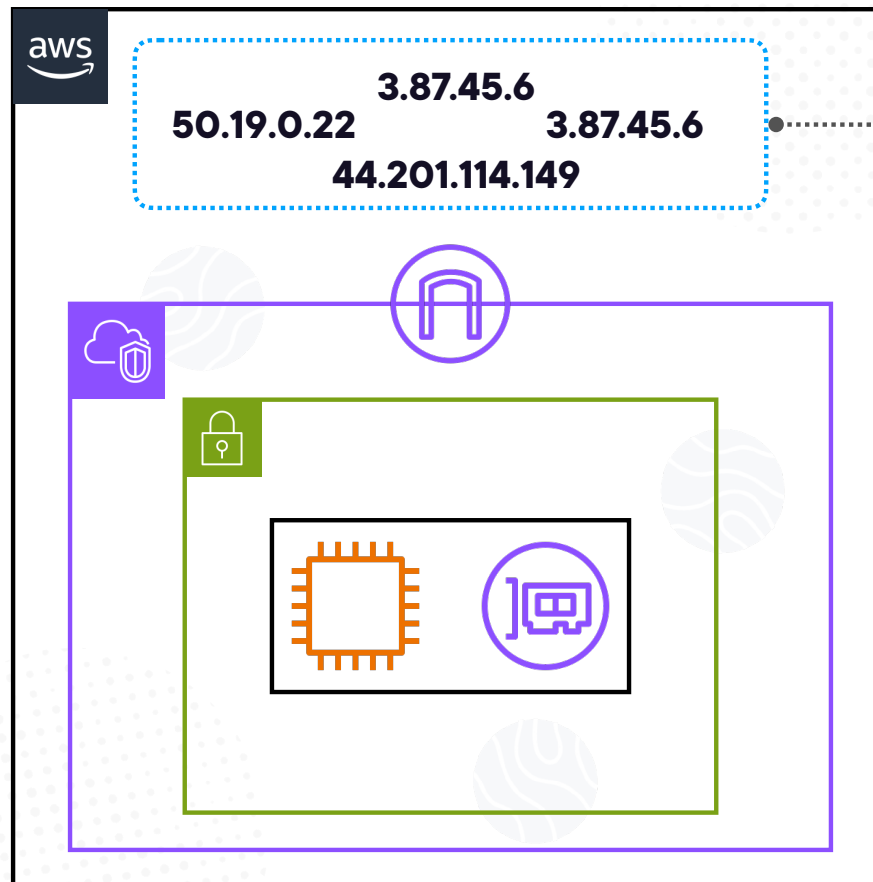
- Instance is stopped, hibernated, or terminated
- An Elastic IP address is associated

In other words, they are temporary, also known as **ephemeral!**

Image Source: <http://unsplash.com>



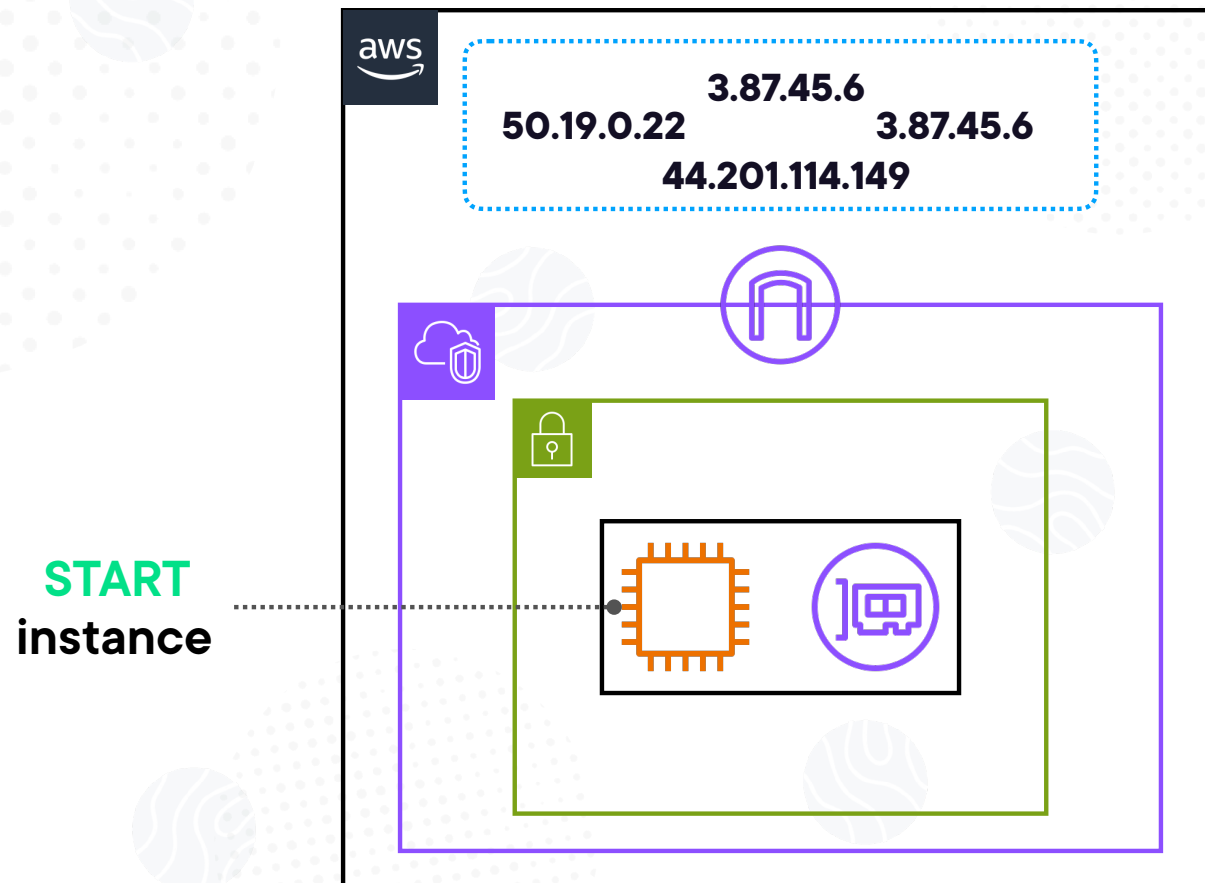
# The Ephemeral Public IPv4 Process



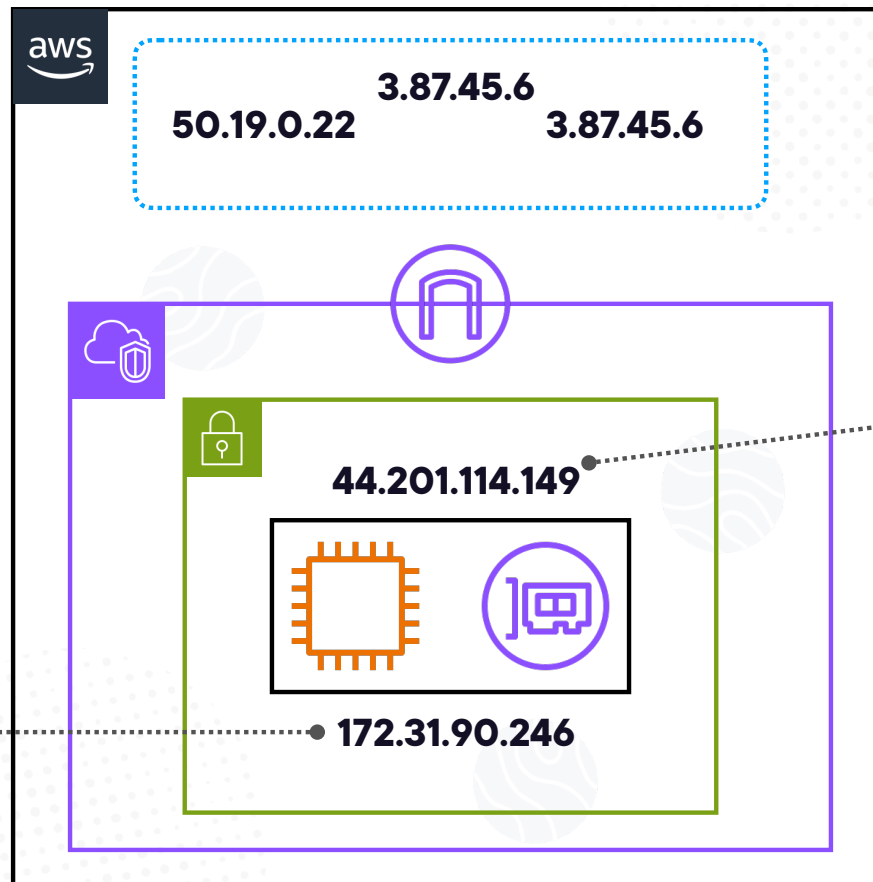
AWS has a large pool of available ephemeral public IPv4 addresses



# The Ephemeral Public IPv4 Process



# The Ephemeral Public IPv4 Process



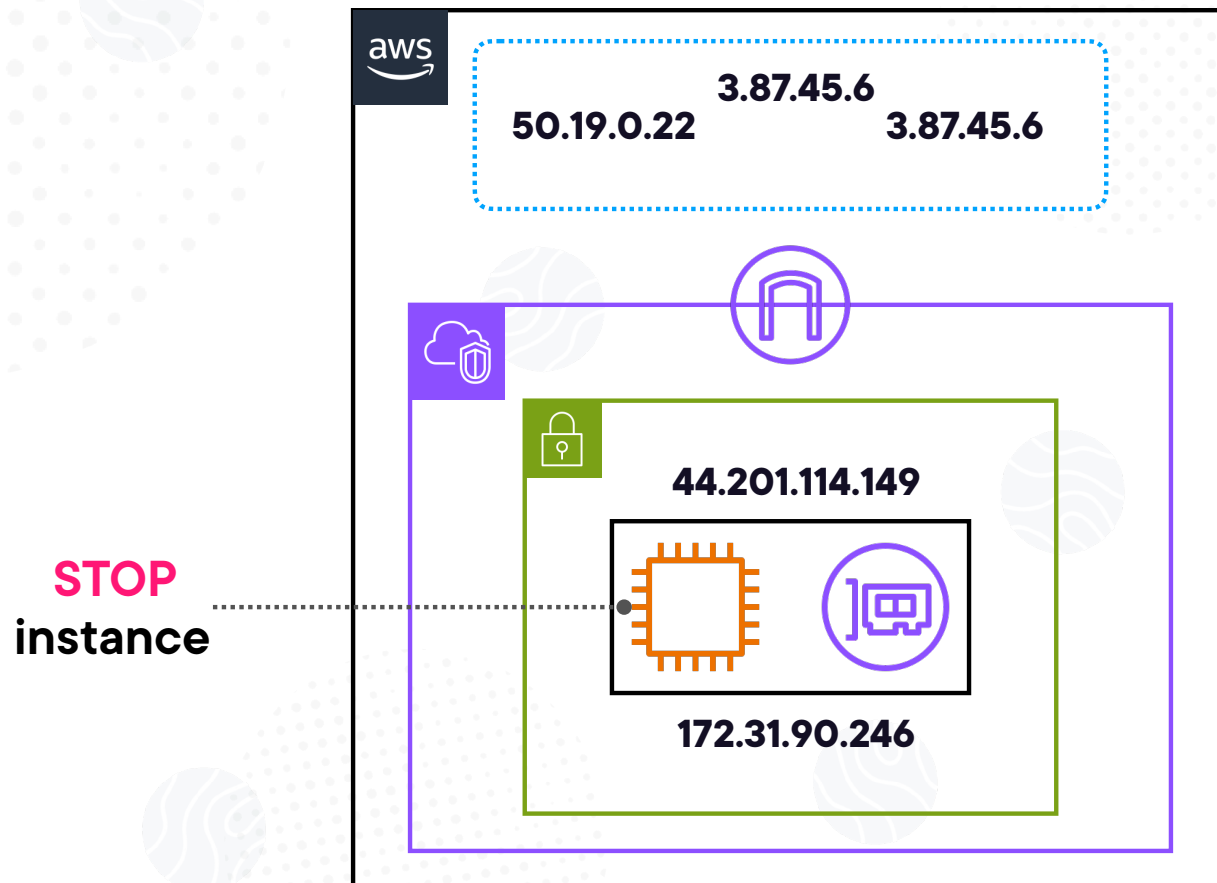
The attached ENI also receives an internal, private IPv4 address from the subnet CIDR

Assigning a public IP address pulls from the Amazon pool of available IPv4 addresses

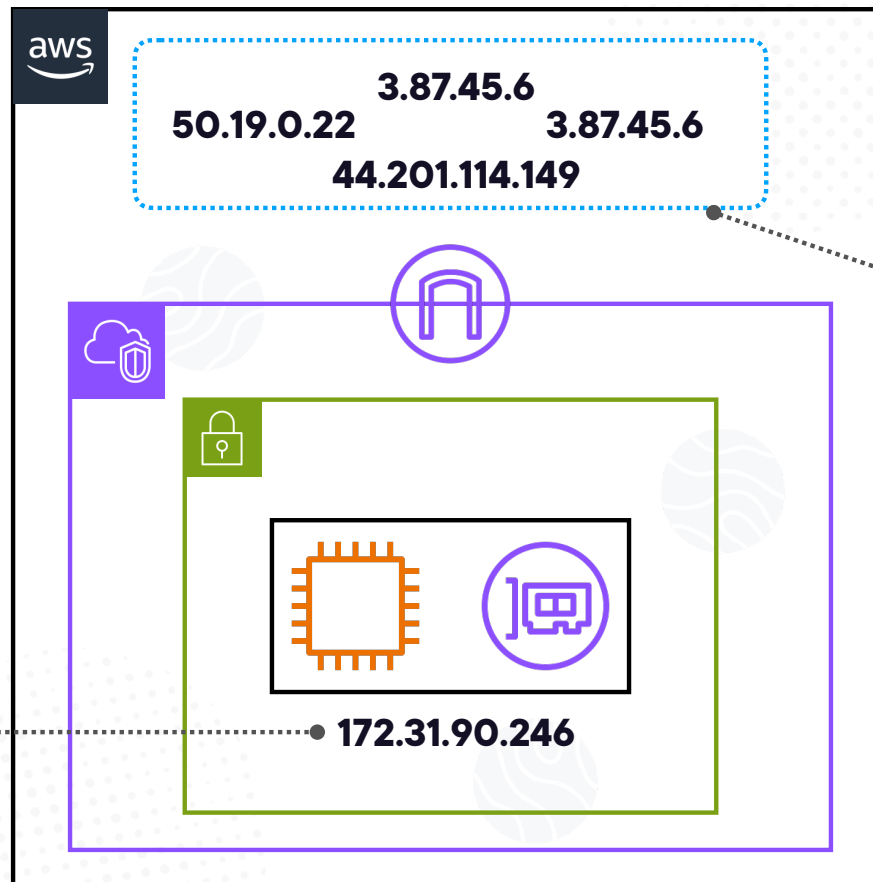
The IP is associated with your ENI

It also receives a public DNS like: [ec2-44-201-114-149.compute-1.amazonaws.com](https://ec2-44-201-114-149.compute-1.amazonaws.com)

# The Ephemeral Public IPv4 Process



# The Ephemeral Public IPv4 Process

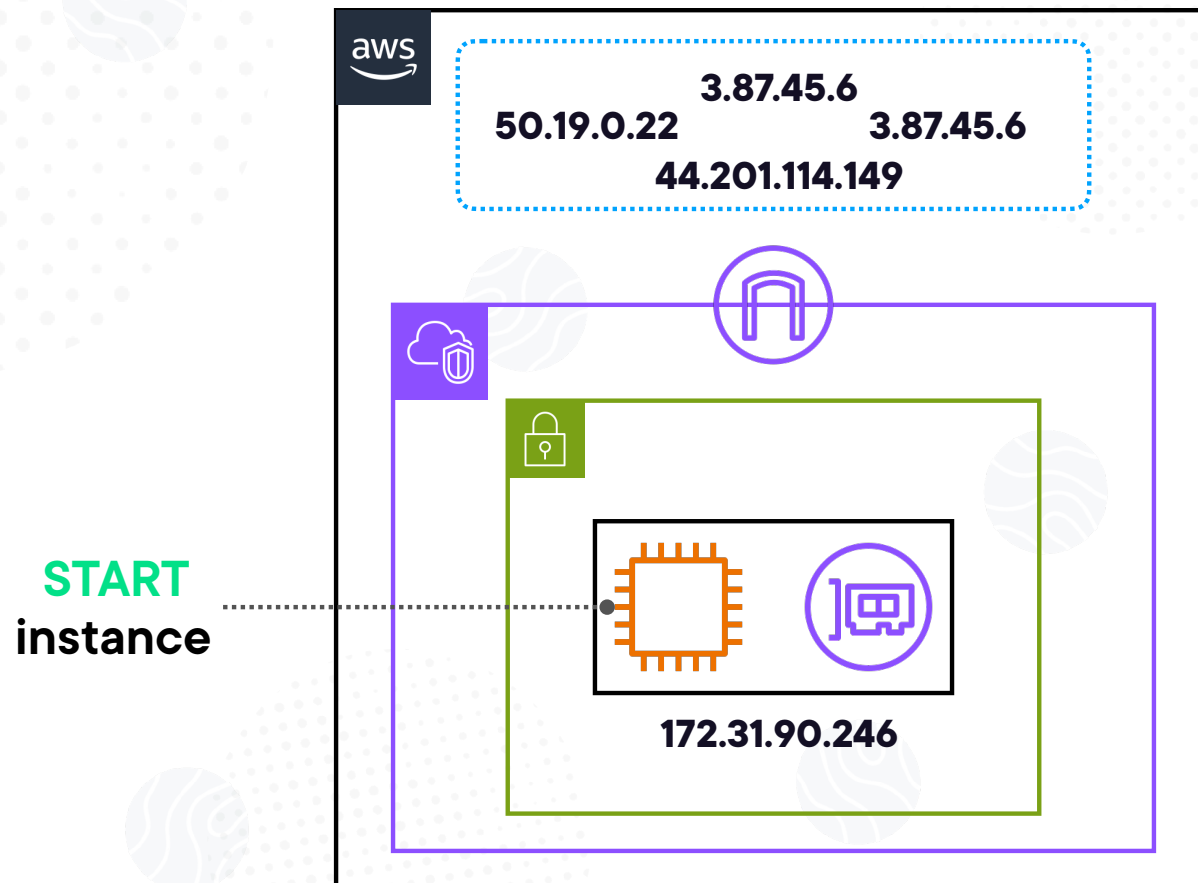


Stopping the instance does **NOT** affect the private IPv4 assigned

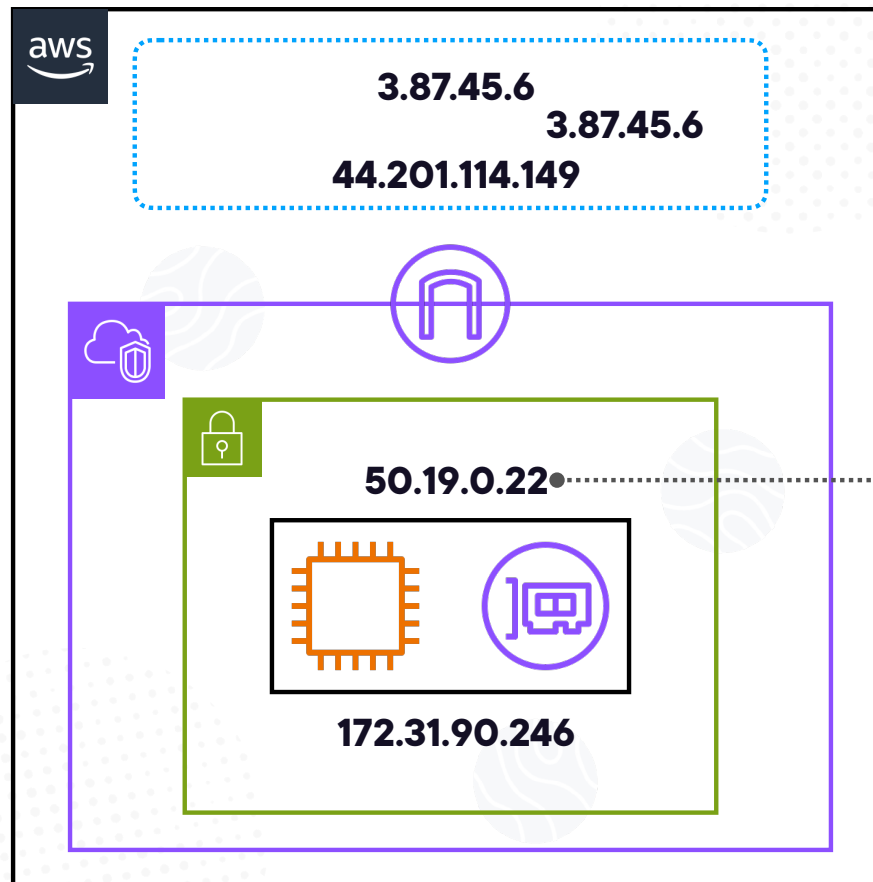
Stopping the instance releases this assigned public IPv4 address back into the pool of addresses



# The Ephemeral Public IPv4 Process



# The Ephemeral Public IPv4 Process



Starting the instance results in a reassignment of an available IPv4 address

You should assume that this will more than likely not be the same IP address as before

**Next up, let's look at using  
Elastic IP Addresses (EIPs)  
to solve for static IP  
mapping!**





# Assigning Static IPv4 with Elastic IP Addresses (EIPs)



**Elastic IP Addresses are public, static IPv4 addresses that you can map to different compute.**



# Elastic IP Address (EIP) Concepts



An EIP is a static IP address that will never change



They can only be used within the specific Region they are created in



You first allocate an address, and then you associate the address



Associating an EIP means the public DNS of the instance will change



Soft limit of 5 EIPs per Region, per Account (*you can request an increase*)



**While these are usable for  
masking failures of  
compute, etc., within  
architectures...**



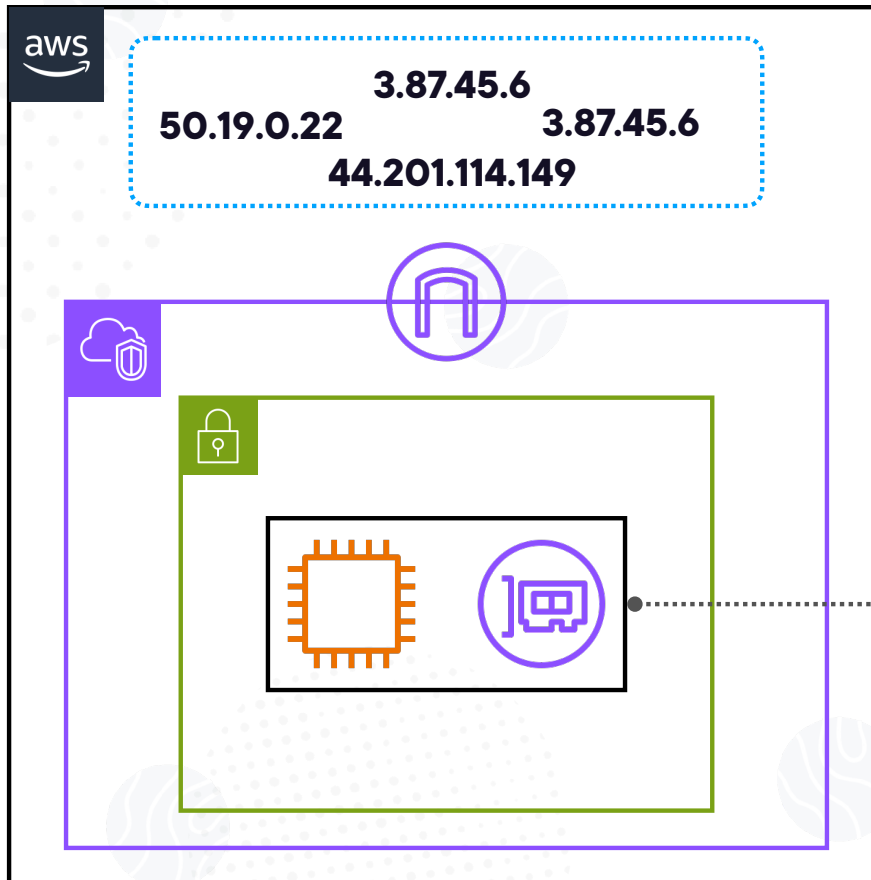
**It is generally a best practice to avoid using these unless absolutely necessary!**



**You should, instead,  
leverage DNS names that  
you can map to ephemeral  
IP addresses.**



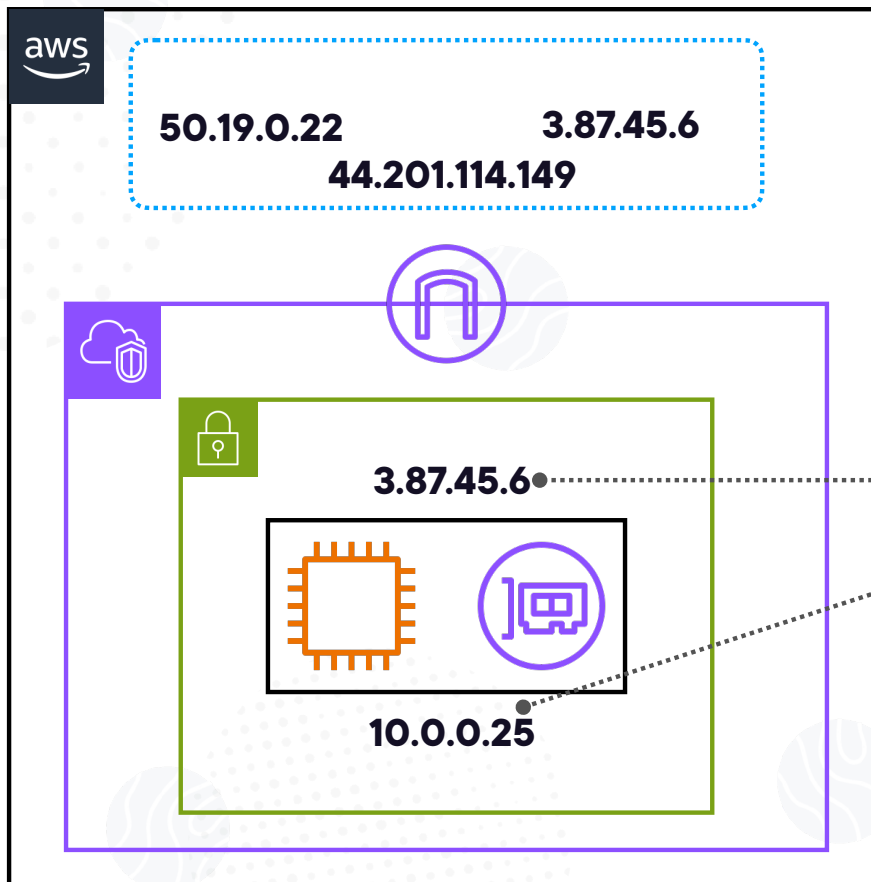
# Demo: Associating an EIP



We will have a publicly reachable EC2 instance and leverage an Amazon provided public IPv4 address

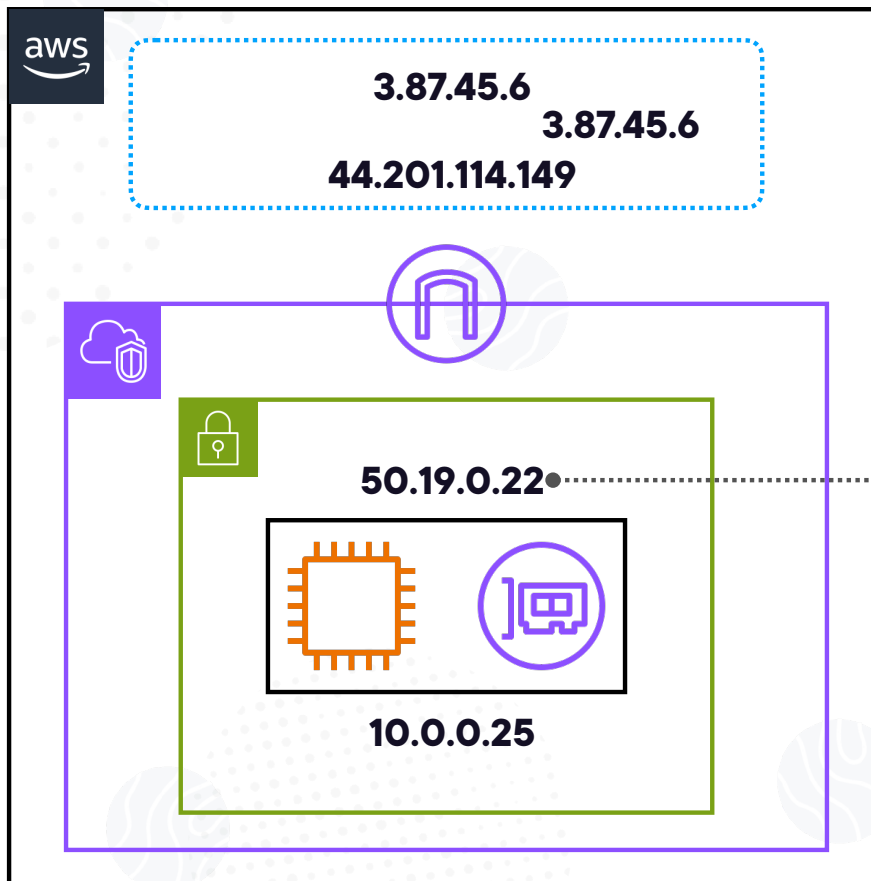


# Demo: Associating an EIP



Launching the instance will result in both an ephemeral public IP address being assigned, as well as a private subnet CIDR IP address being assigned

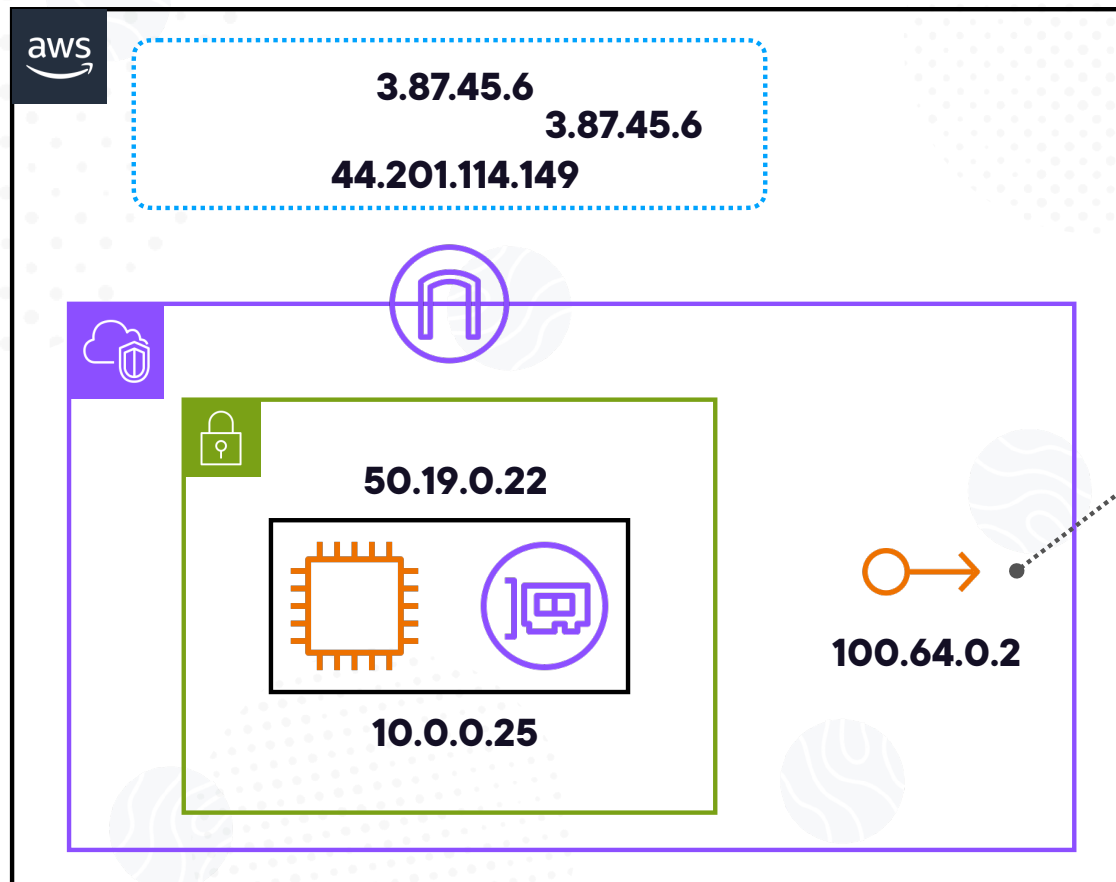
# Demo: Associating an EIP



To show how ephemeral IP addresses work, we will then **stop and start** the instance and watch the public IP address get reassigned

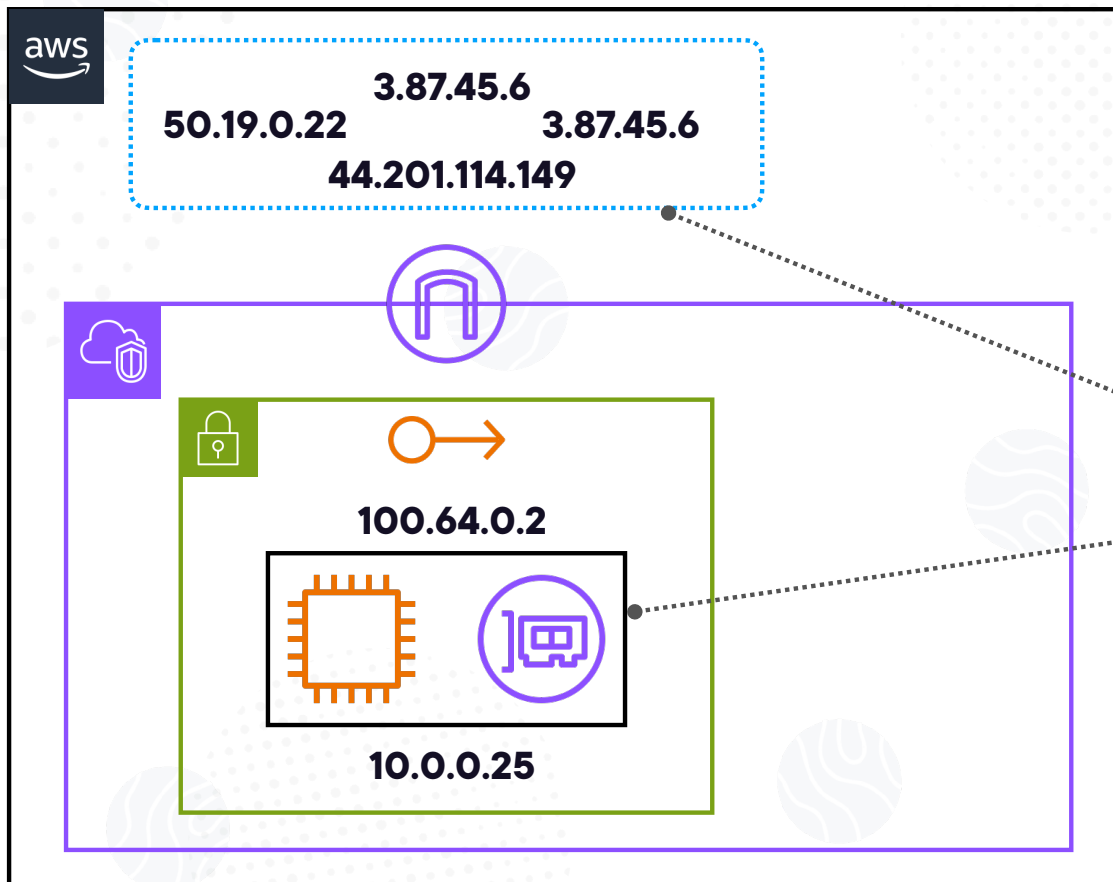


# Demo: Associating an EIP



To solve for this changing of ephemeral IP addresses, we will create an EIP within our VPC for static mapping

# Demo: Associating an EIP



Associating the EIP will result in a release of the ephemeral public IP address

Now, anytime we reboot or stop the instance, the EIP will remain associated



# **Dual-home EC2 Instances**

<https://t.me/learningnets>





## Multiple Network Interfaces

Dual-home instances typically means that there are multiple ENIs attached which live in different network bands

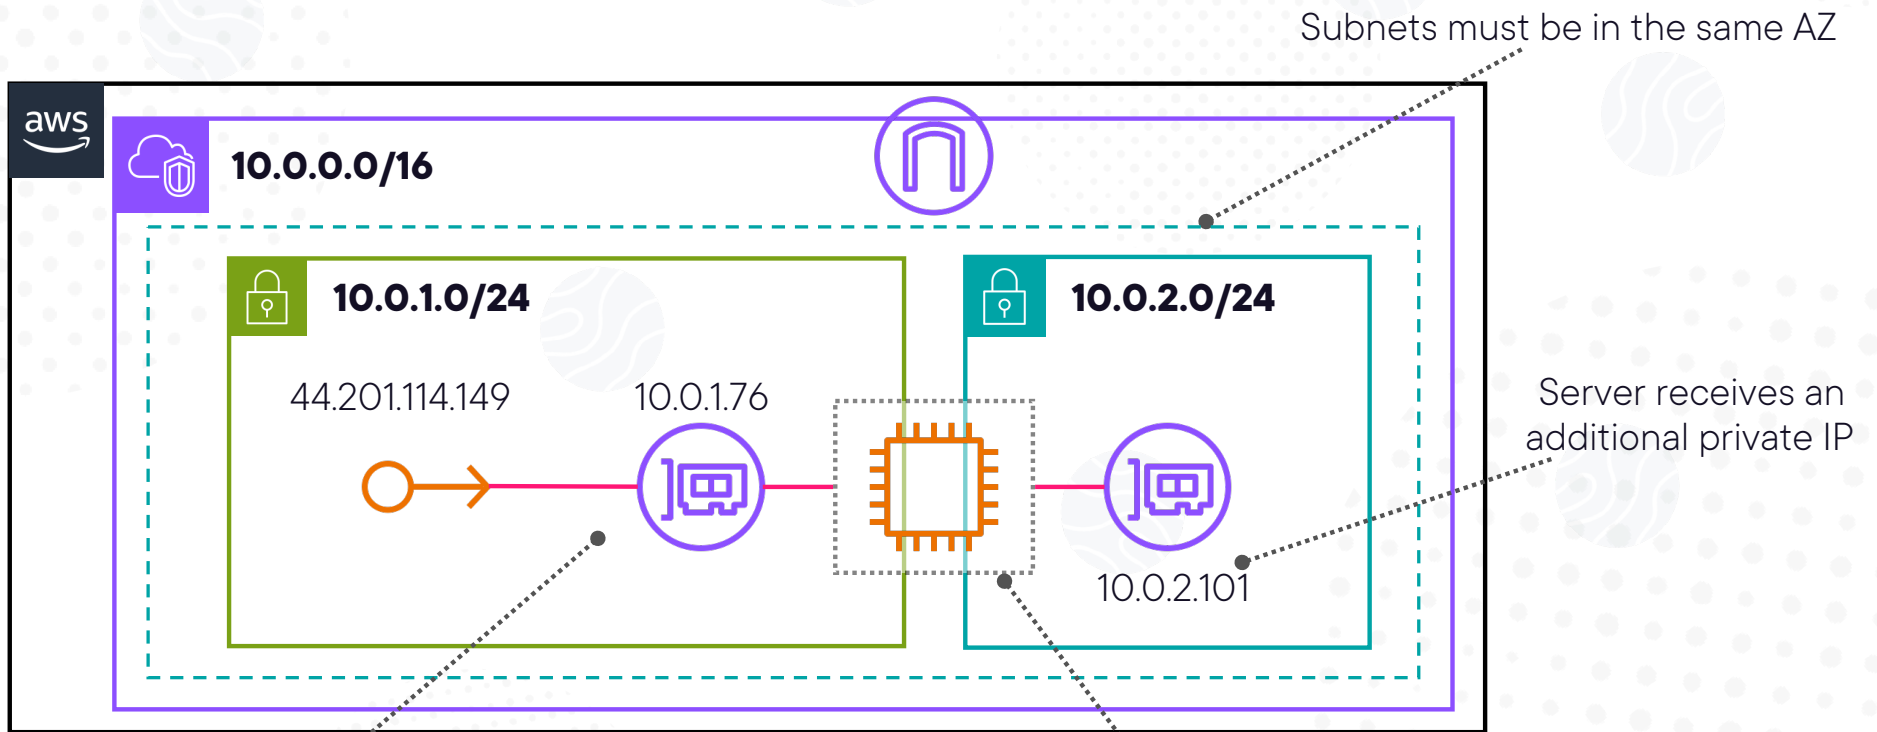
EC2 networking is split between two different subnets/tiers

**Example:** Management subnet for restrictive access, public subnet for public web access

Image Source: <https://unsplash.com/>



# Dual-home Architecture Diagram



Public IPv4 address/Elastic IP address associated and a private IPv4 assigned

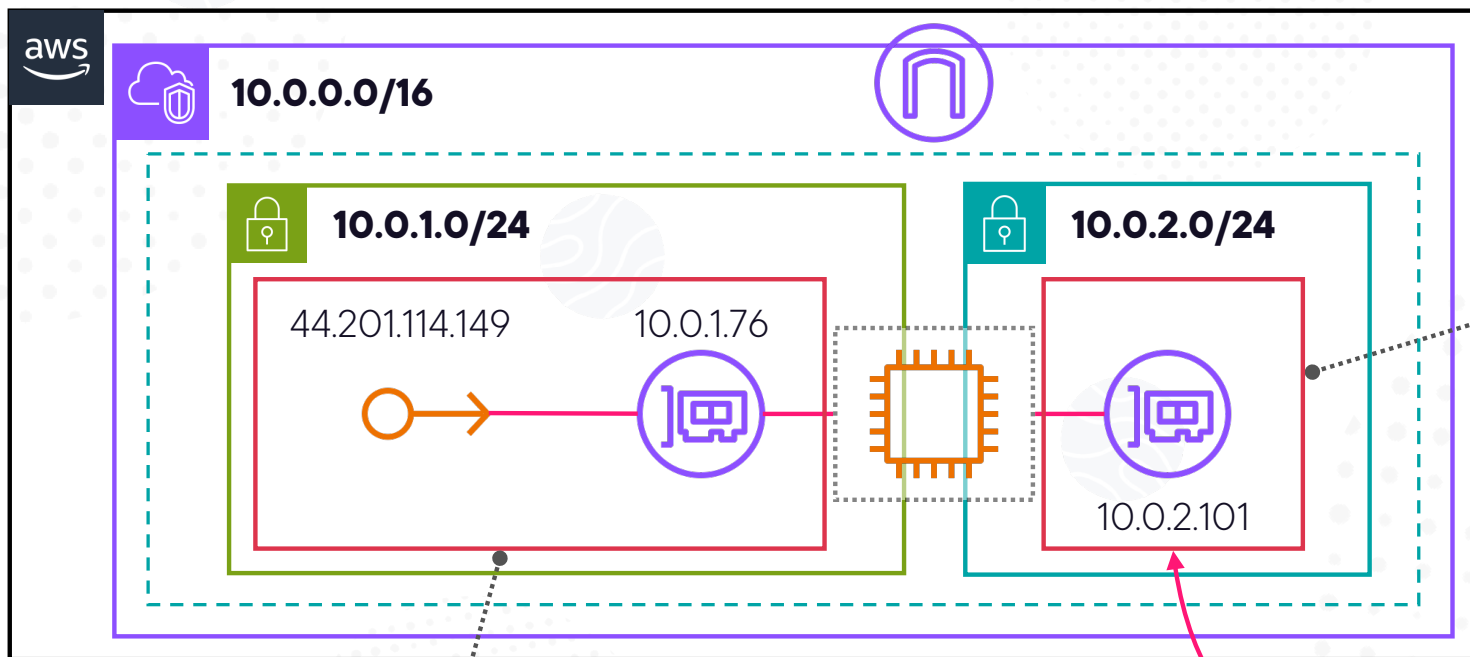
Server has two network bands, so it can talk to public resources, and then it has a separate secured ENI for internal networks

Server receives an additional private IP

Subnets must be in the same AZ



# Dual-home Architecture Diagram

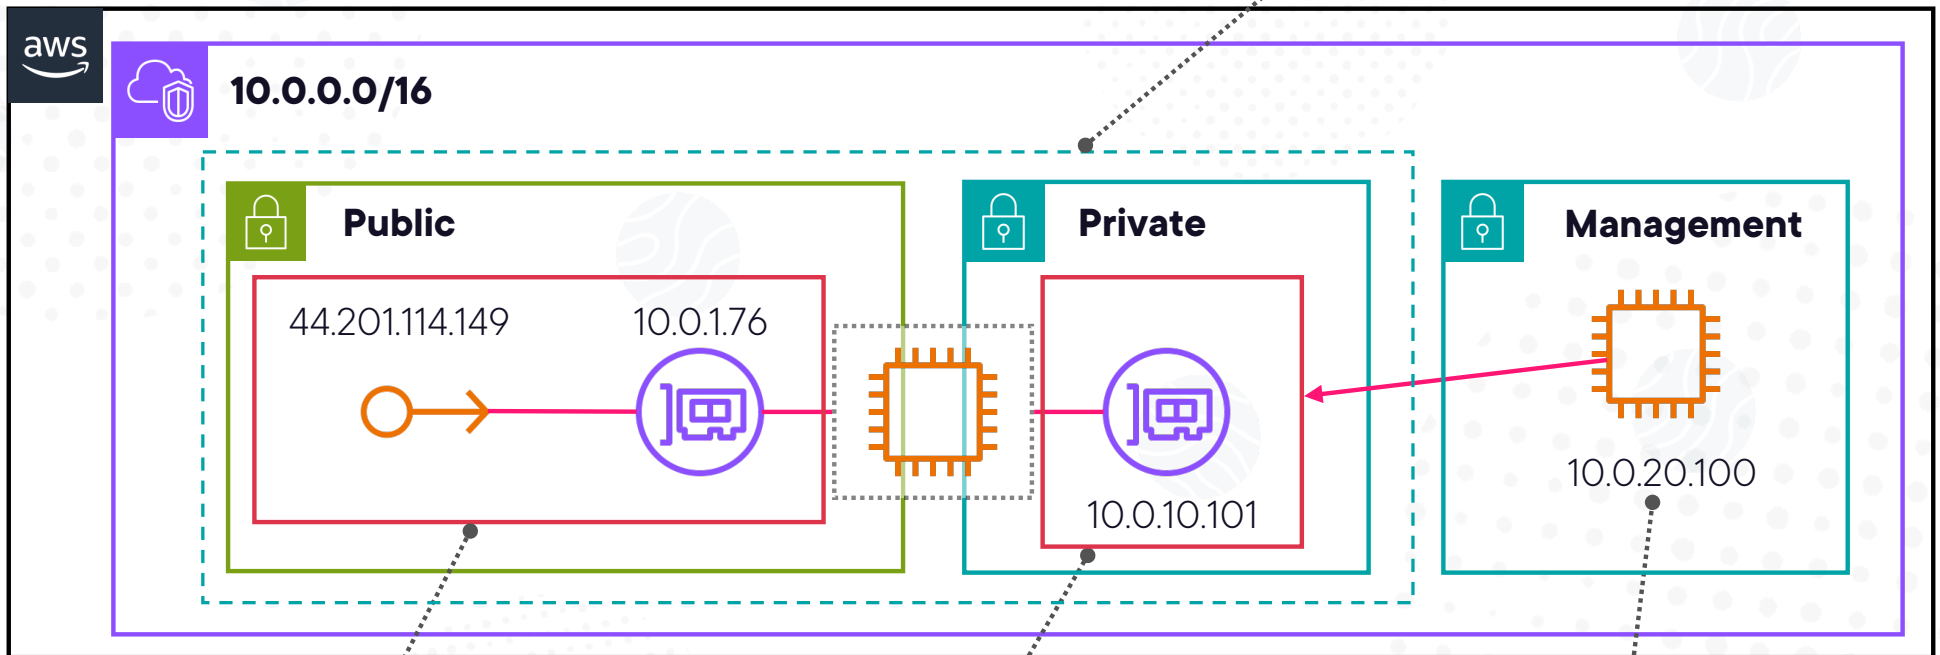


**External security group** can allow application traffic (*HTTP, HTTPS, etc.*) from anywhere

**Internal security group** can be restricted to only allow internal traffic (*SSH, RDP, etc.*) for secure management of the resource

# Demo: Creating a Dual-home EC2 Instance

ENIs must be in the same Availability Zone



**Primary security group** will allow HTTP from anywhere

**Secondary security group** is restricted to only allow internal SSH traffic from the management subnet

Test connection from a **management subnet**



# **Increasing Performance with EC2 Placement Groups**


<https://t.me/learningnets>




**Placement groups are optional groupings of EC2 instances to influence their placement.**




# Three Placement Group Types



**Cluster**



**Spread**

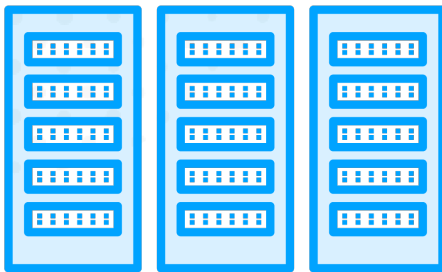


**Partition**



# Three Placement Group Types

Key characteristic



Cluster


- Puts instances closely together in the same AZ
- Offers extremely low-latency network performance
- Typically used for high-performance computing (HPC)
- Not all instance types support this placement group!




# Three Placement Group Types



**Cluster**




**Spread**



**Partition**



# Three Placement Group Types




**Spread**


- Places small groups of instance across different underlying hardware
- Meant to reduce common hardware failures
- Recommended for critical applications for uptime
- Typically used for small number of instances



# Three Placement Group Types



**Cluster**



**Spread**



**Partition**



# Three Placement Group Types



Partition


- Spreads instances across logical partitions
- Ensures groups do not share underlying hardware
- Each partition has its own set of racks
- Used for distributed/replicated workloads
- Hadoop, Cassandra, Kafka



# Exam Tips to Remember



You can add a stopped instance to an existing placement group!



AWS recommends homogenous instances in cluster placement groups.

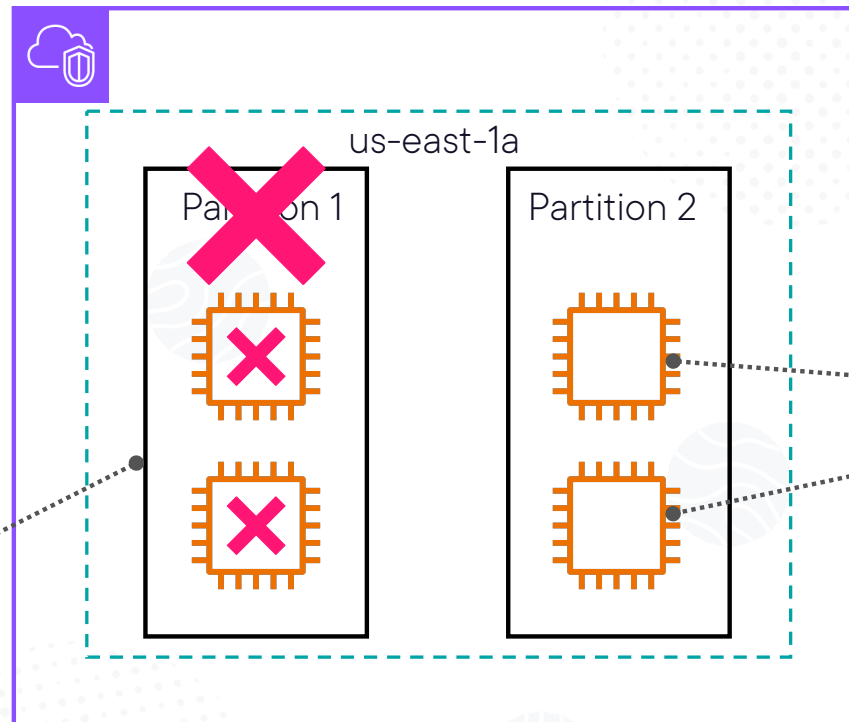


You cannot merge placement groups!



# Demo: Launching EC2 Instances in a Placement Group

## Partition




If a partition's underlying hardware has failures, then the other partition will continue to run

EC2 does its "best effort" to evenly distribute instances across the partitions



# Demo: Launching EC2 Instances in a Placement Group

## Spread



If one EC2's underlying hardware has failures, then the rest will continue to run

Each EC2 instance is placed on a completely separate/distinct server rack

This means each instance has its own networking and power source





# **AWS Outposts for Localized Compute**

<https://t.me/learningnets>





## AWS Outposts

Outposts brings the AWS data center directly to you, on-premises

Allows you to have a large variety of AWS services in your data center

You can have Outposts in sizes such as:

- 1U or 2U servers (Outpost Servers)
- A 42U rack (Outpost Racks)

*The above are just industry standards for measuring servers and rack space*

Image Source: <https://unsplash.com/>



# Benefits of AWS Outposts

Create a hybrid cloud where you can leverage AWS services inside your own data center

Bring the AWS Management Console, APIs, and SDKs into your data center, allowing uniform consistency in your hybrid environment

AWS can manage the infrastructure for you. You do not need a dedicated team to look after your Outposts infrastructure



# Outposts Family Members

## Outposts Racks

Available starting at a single 42U rack

Scale up to 96 racks

AWS compute, storage, database, and other services locally

Gives the same AWS infrastructure, services, and APIs in your own data center

vs.

## Outposts Servers

Individual servers in 1U or 2U form factor

Useful for small space requirements, such as retail stores, branch offices, healthcare provider locations, or factory floors

Provides local compute and networking services

**Scenarios involving extending AWS to your data center will likely include AWS Outposts.**






# Enhanced Networking for EC2


<https://t.me/learningnets>




# EC2 Networking Card Types



**Elastic Network Interface (ENI)**



**Enhanced Networking (EN)**



**Elastic Fabric Adapter (EFA)**



# EC2 Networking Card Types

## ENI

Best for basic, day-to-day networking tasks

## EN

Uses single root I/O virtualization (SR-IOV) for high-performance

## EFA

Accelerates HPC and machine learning applications



**SR-IOV allows a single physical network device to present itself as multiple virtual devices to the operating system and hypervisor.**




**This allows much higher I/O performance (*measured in packets per second and latency*) while minimizing CPU utilization!**




# Enhanced Networking

Your instance type dictates the supported Enhanced Networking options...



## Elastic Network Adapter (ENA)

Supports network speeds of up to 100 Gbps for supported instance types.  
Choose this if possible.




## Intel 82599 VF Interface

Supports network speeds of up to 10 Gbps for supported instance types.  
Typically used on older instances.



# Elastic Fabric Adapter (EFA)



A network device you attach to your EC2 instance to accelerate High Performance Computing (HPC) and machine learning (ML) applications



EFA offers OS-Bypass, which allows HPC and ML applications to communicate directly with the network interface hardware



Provides lower, more consistent latency, and higher throughput than the TCP transport traditionally used in cloud-based HPC systems




Only linux! This is not supported for Windows instances at this time!





# **Module Summary and Exam Tips**






## Elastic Network Interface (ENI)

Logical networking component that is part of a VPC and is meant to represent a virtual network card for compute resources




# Ephemeral Public IP Addresses and Static EIPs



Ephemeral public IPs are from Amazon's pool of public IPv4 addresses



You lose your IP when instances are **stopped**, **hibernated**, or **terminated**



You also lose the IP whenever an Elastic IP address is associated



Use an EIP if you need a static, never changing public IPv4 address




**Dual-home EC2 instances  
have networking split  
between two different  
subnet layers.**



**Dual-home subnets must be  
in the same Availability  
Zone!**




# Remember the Three Placement Groups




## Cluster

Low network latency,  
high network  
throughput, HPC



## Spread

Useful for individual,  
critical EC2 instances  
and applications



## Partition

Multiple EC2 instances;  
Think HDFS, HBase, and  
Cassandra





## **AWS Outposts**

**Scenario about extending AWS to your data center?**


**Think AWS Outposts!**

**AWS Outposts rack for large deployments**

**AWS Outposts servers for smaller deployments**




# Remember the EC2 Networking Card Types




**ENI**

**Best for basic, day-to-day networking tasks**



**EN**

**Uses SR-IOV for high-performance**



**EFA**

**Used to accelerate HPC and ML applications**

