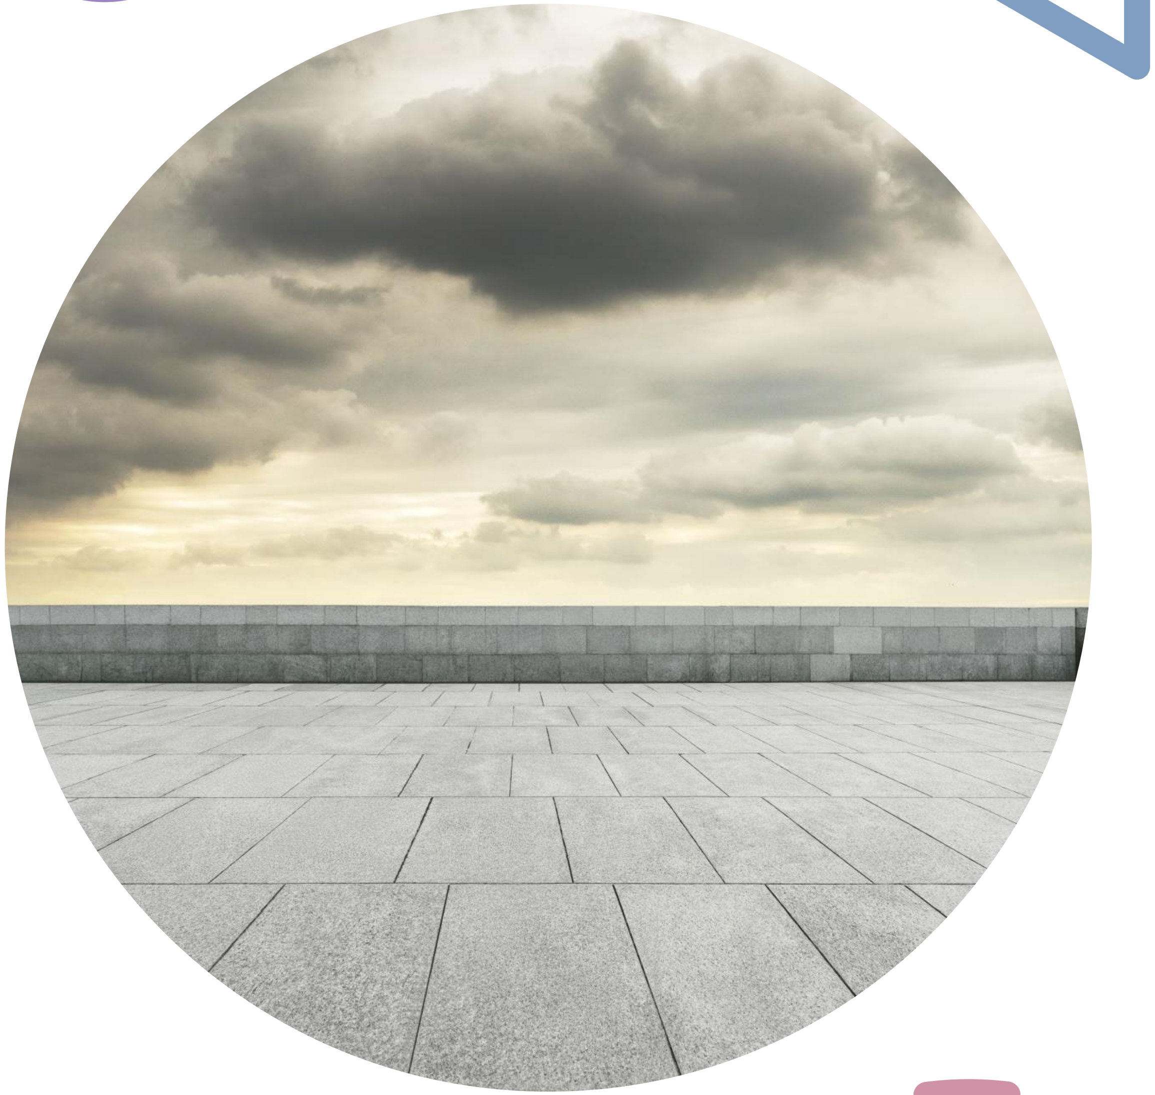

AWS Cloud Practitioner Week-4

Training Course

EC2

From Week-3

EC2 Instance Type Change

- Instance type can be changed only instances that has EBS volume attached.
- Cannot change instance type for **Instance Store** backed EC2
- Steps:
  - Actions => Instance State => Stop
  - Actions => Instance Settings => Change Instance Type
  - Actions => Instance State => Start

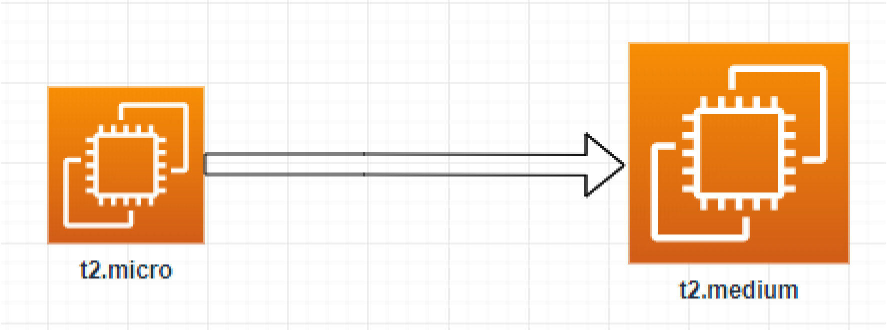

AMI – Amazon Machine Image - Lab

- Launch an EC2 instance and customize it.
- Stop the instance
- Create an AMI form the stopped instance
- Launch an instance from the customized AMI.

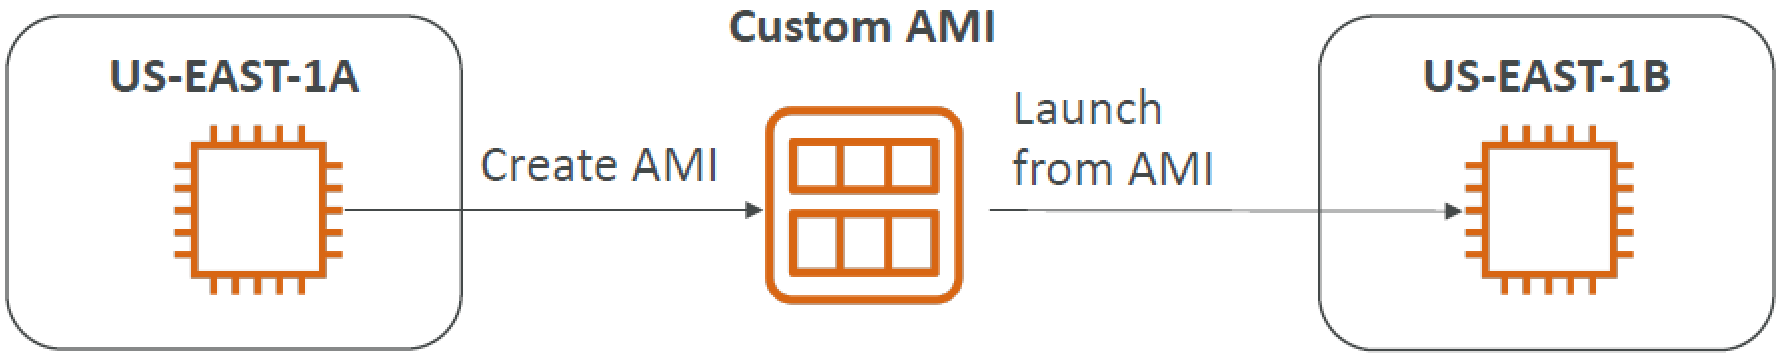

AMI – Amazon Machine Image

AMI = Amazon Machine Image

Faster boot because all software are pre-packaged

AMI Customization:

- You can install your own software
- Build for specific **Region**
- Can be copied across **Regions**
- Launch EC2 instance:
  - Public AMI which is the AWS Managed AMI
  - Your own AMI – Customized by you
  - AWS Marketplace AMI: 3rd party vendor AMI to purchase.

Firewall

What is Firewall?

Generally a firewall is network security device that monitor incoming and outgoing network traffic and permits or block data packets based on the rule set.

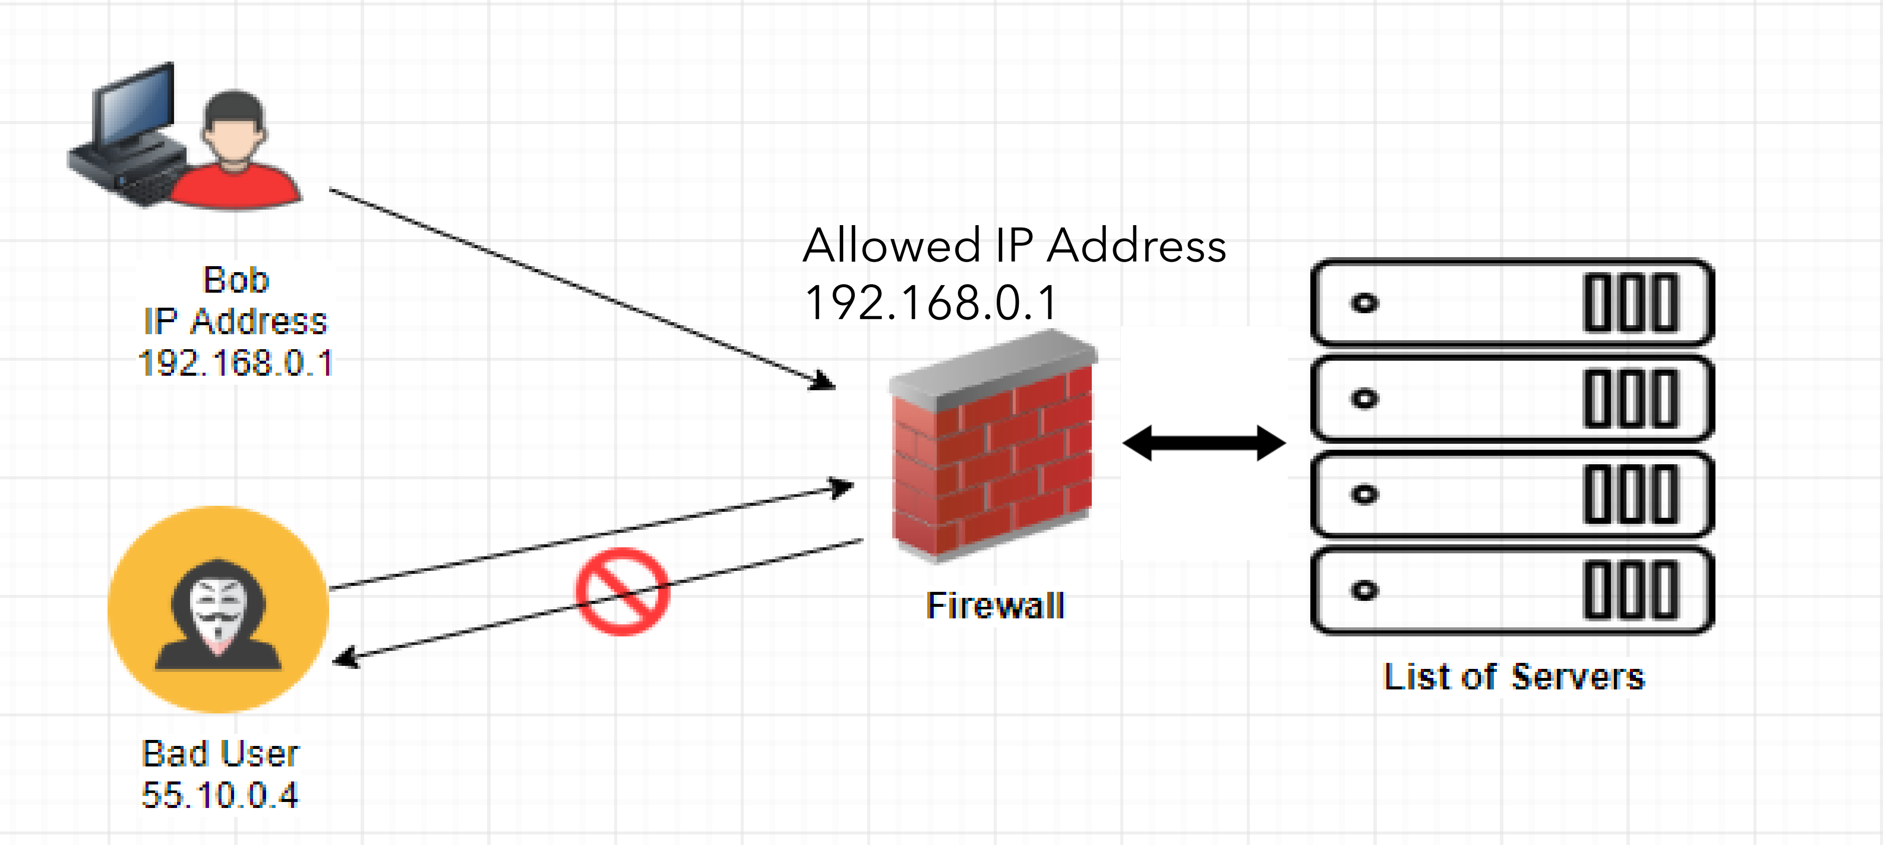

Network Firewall for EC2

Large picture of security, protecting your EC2 Instance.

- Layer-1: **VPC Route Tables**, control the Gateway. You can change the routes to protect access from internet.
- Layer-2: **NACL** is subnet level firewall, can allow/deny for inbound/outbound.
- Layer-3: **Security Group**: virtual firewall EC2 instance level.
- Layer-4: **OS level firewall**, Microsoft firewall, Norton Security installed within the Operating Systems.

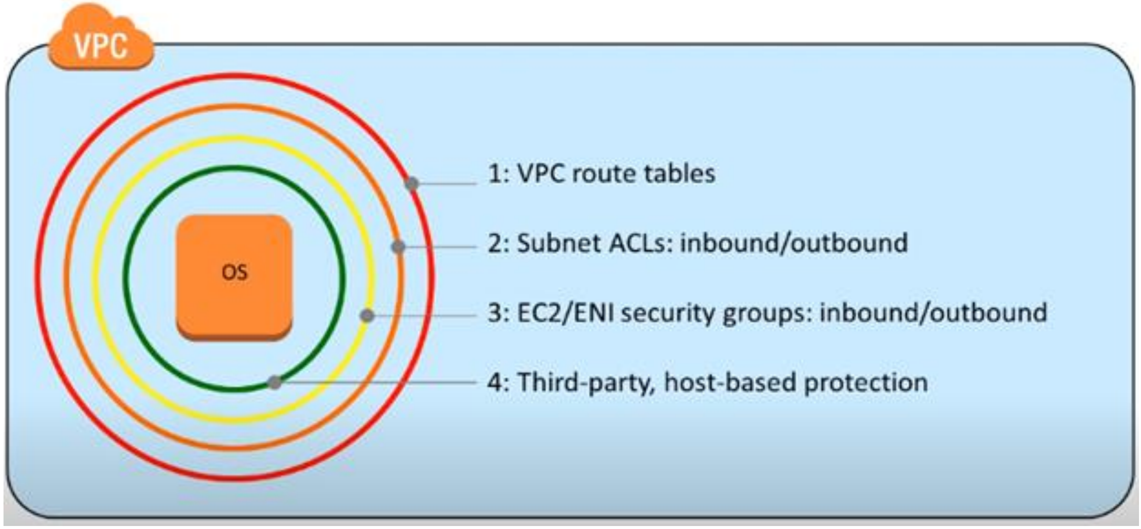
VPC
OS
1: VPC route tables
2: Subnet ACLs: inbound/outbound
3: EC2/ENI security groups: inbound/outbound
4: Third-party, host-based protection

AWS Storages

EBS

Elastic Block Storage

S3

Simple Storage Service

S3 Glacier

Data archiving and backup

EFS

Elastic File Service (Linux)
Network Attached Storage

Storage Gateway

Hybrid Storage

Snowball

To transfer data to AWS

Snowmobile

To migrate large amount of Data to AWS.

FSx

Windows File Server

EBS

Elastic Block Store

Elastic Block Store (EBS) Volume

- **EBS**: (Elastic Block Store)
- **Block** Level Storage attach to EC2
- A network drive you can attach to your instance.
- Data on EBS volume are **persistent**.
- They are bound to a specific Availability Zone, cannot be across multiple Azs.
- To move a volume across, you first need to snapshot it.
- Root EBS volume is mounted to one instance at a time.
- You can detach and attach the volume to another instance.
- Root volume gets terminated along with the instance.

- Free tier allows you up to 30GB of free EBS storage of gp2 (General Purpose) per month.

EBS Volume Diagram for Lab

One or more EBS volume can be attached to one EC2 Instance only.

One EBS volume cannot be shared on multiple EC2 Instances.

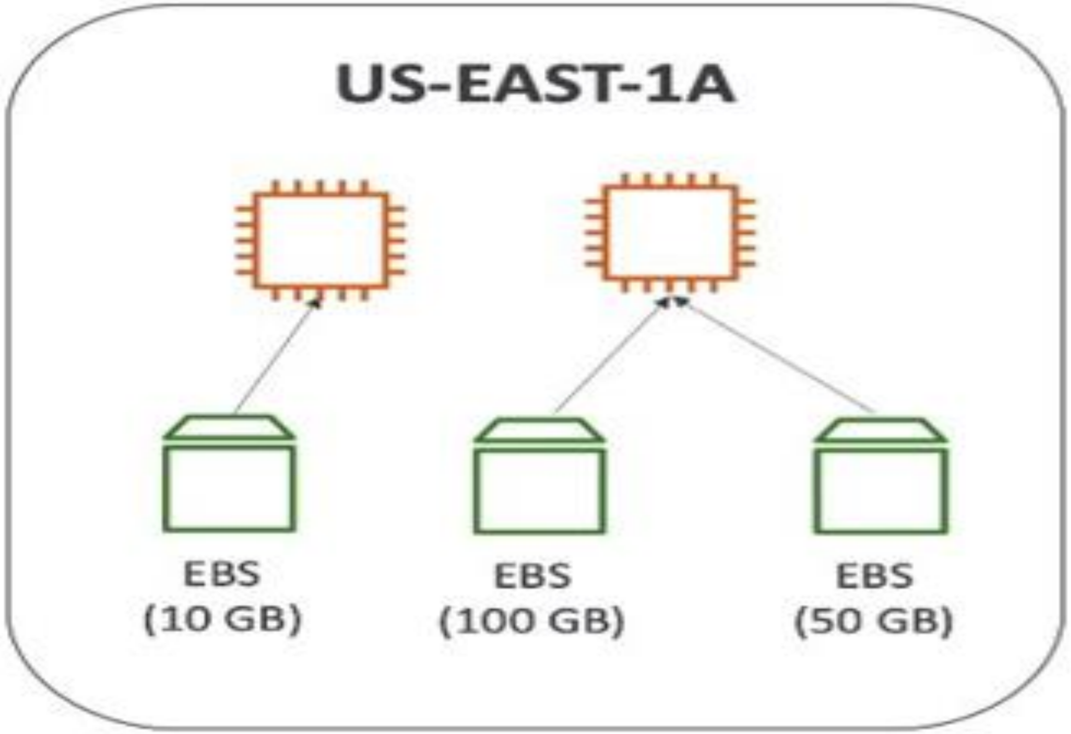

EBS Volume for Windows Server - Lab

Lab

1. Create a second EBS volume in the same AZ as the instance.
2. Choose the default volume Device /dev/xvdf
3. Attach the volume to EC2 instance
4. Partition and format the volume as D: drive.

Re-Sizing Root Volume of EC2 Instance

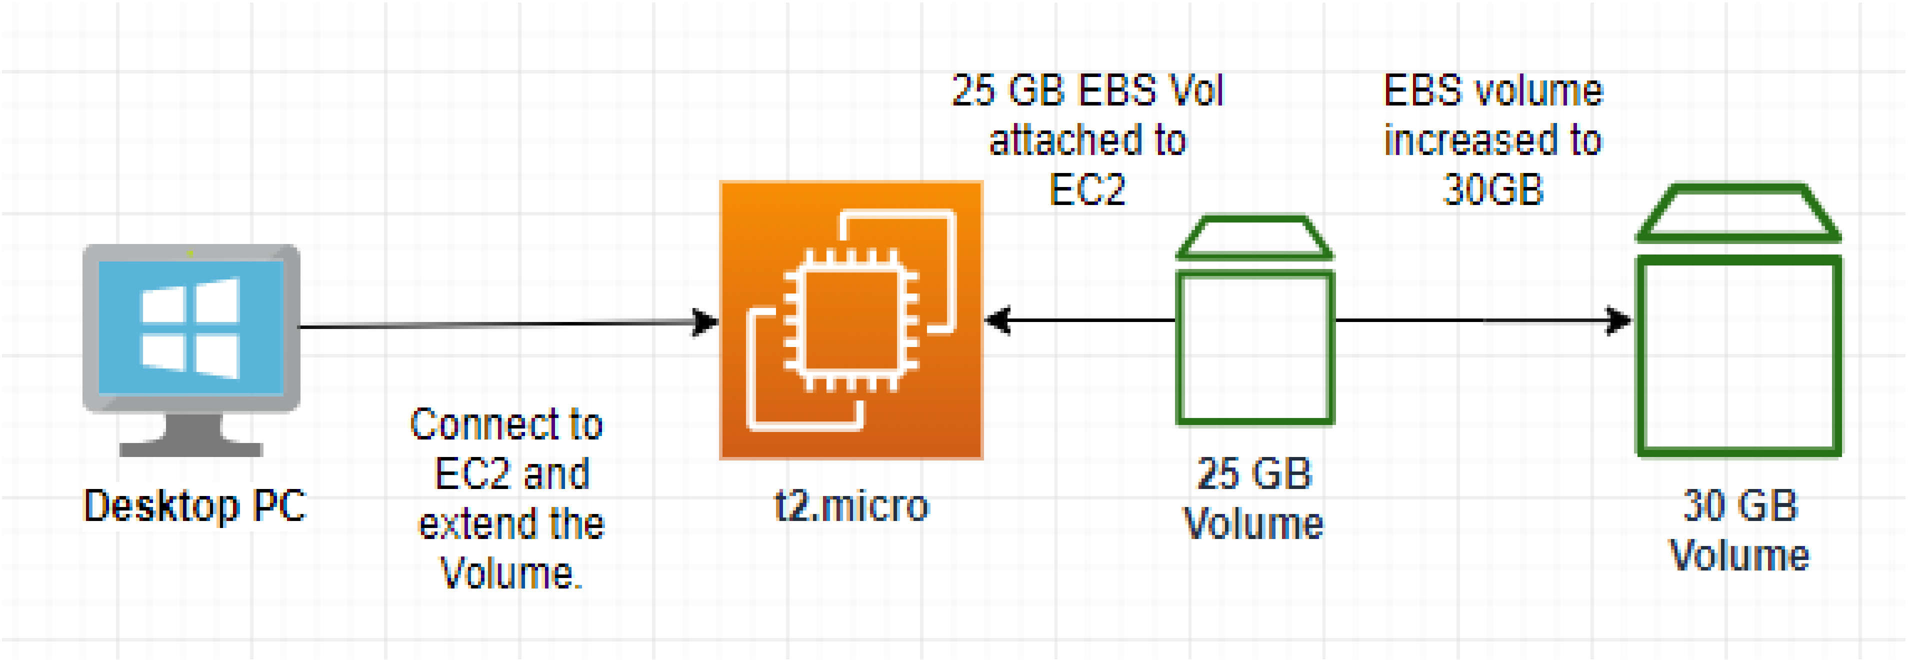
25 GB EBS Vol attached to EC2
EBS volume increased to 30GB
Desktop PC
Connect to EC2 and extend the Volume.
t2.micro
25 GB Volume
30 GB Volume

Taking a Snapshot (backup) of EBS Volume

- A backup of the whole EBS volume attached to a running instance
- No need to detach the volume
- You can copy snapshots across AZ and Region.

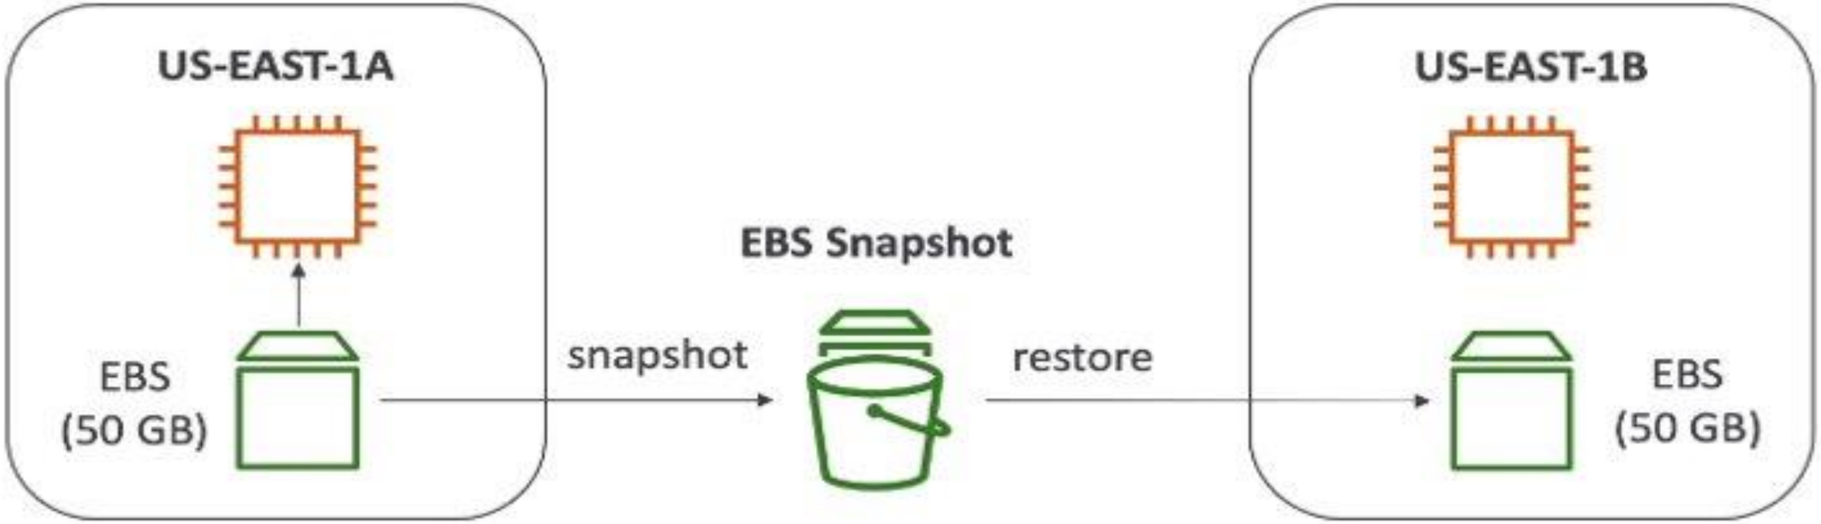

EC2 Instance Store Storage

EC2 Instance Store

- Local to Instance
- Non persistent data store
- Data not replicated
- No Snapshot (volume backup) supported
- Like USB stick, mounted on EC2
- Better I/O performance
- Faster than EBS Storage
- Lose data if instances are stopped.
- Good for buffer, cache, temp data etc.

Ref: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html

EC2 Instance vs EBS Storages

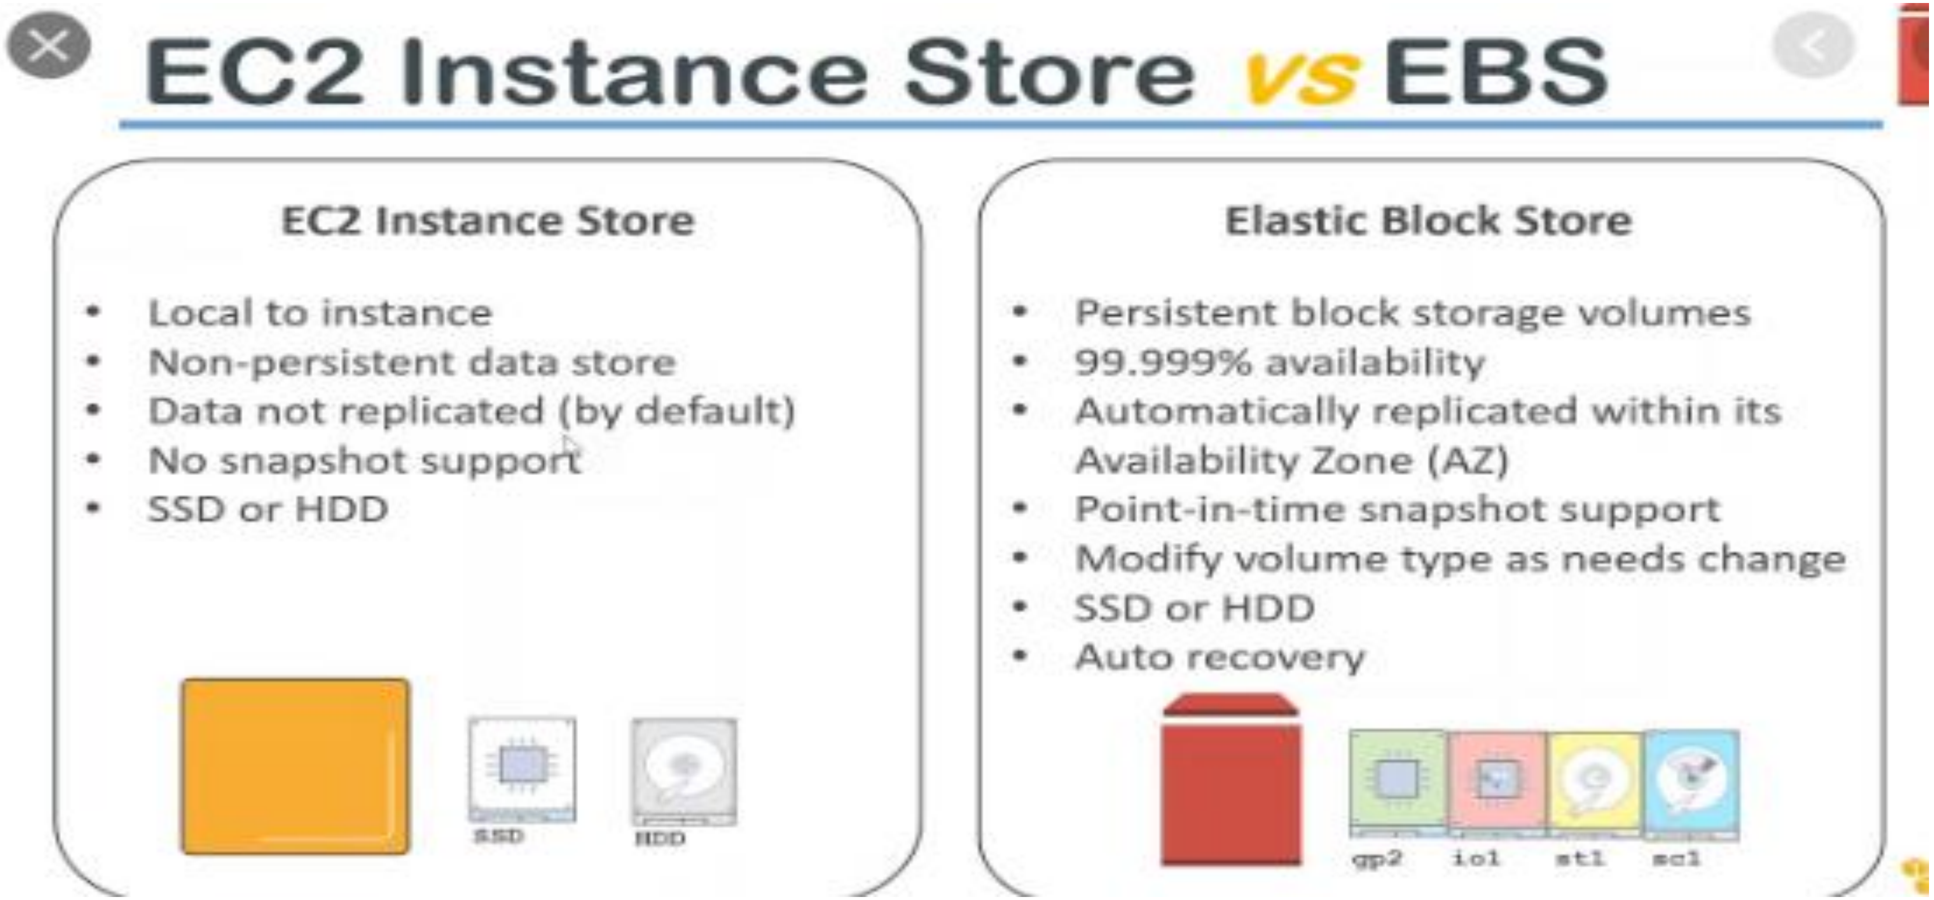

EC2 Instance Store *vs* EBS

EC2 Instance Store

- Local to instance
- Non-persistent data store
- Data not replicated (by default)
- No snapshot support
- SSD or HDD

Elastic Block Store

- Persistent block storage volumes
- 99.999% availability
- Automatically replicated within its Availability Zone (AZ)
- Point-in-time snapshot support
- Modify volume type as needs change
- SSD or HDD
- Auto recovery

S3

Simple Storage System

S3 – Simple Storage Service

- An **object** level storage service.
- Unlimited Storage
- Allows to store **Objects** (files) in **buckets** (directories)
- S3 is global service but specific to a region
- Bucket name must have globally unique name
  - No uppercase / Underscore
- 3-63 characters long
- Fast, highly available, Secure

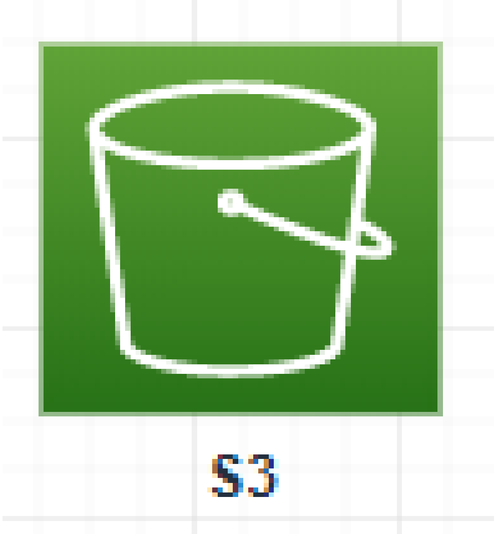

S3 – access via EC2

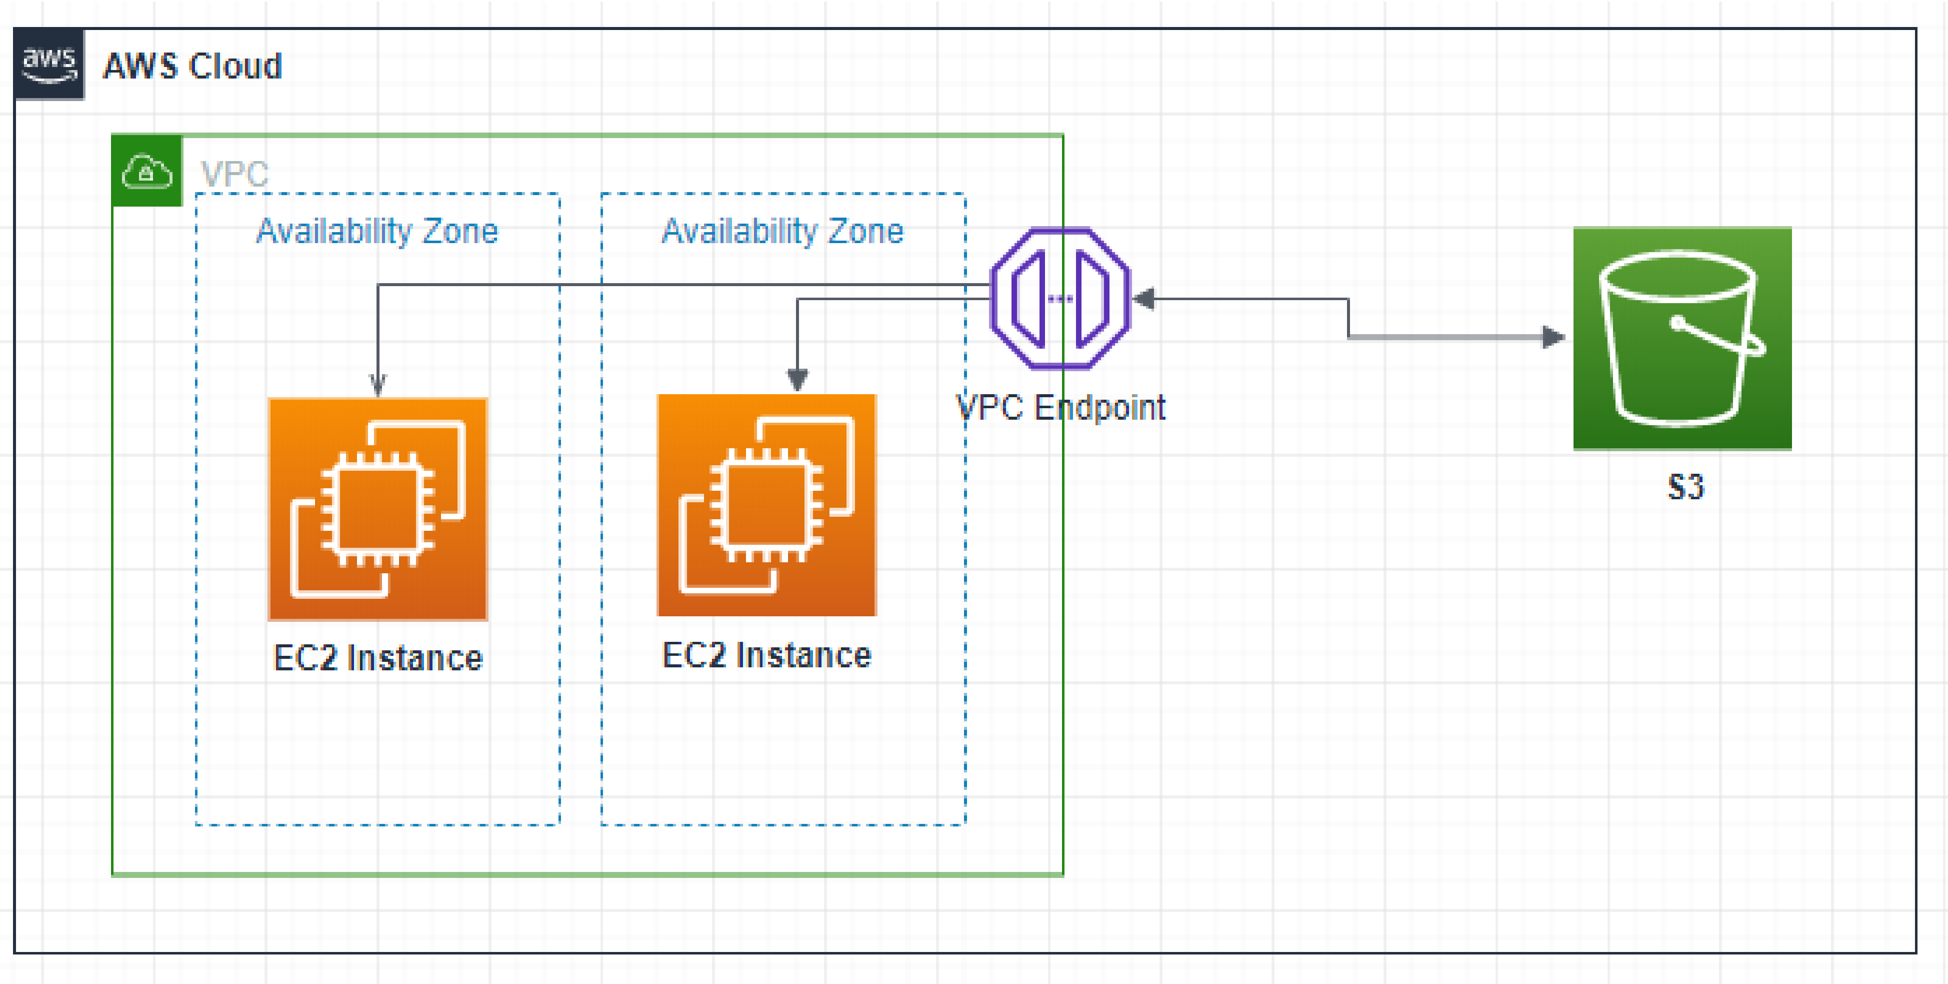

S3 – Use Case

- Static websites from S3
- Backup and storage
- EBS snapshot storage
- Disaster Recovery Data storage
- Media Hosting
- Data Archive
- Data Lakes & big data analytics

S3 Overview

- Object level storage
- Objects = the files
- Key = full path (long name with "/")

  S3://my-bucket/data/my-file.pdf

- Key of composed of:
  - Prefix + Object Name

  S3://my-bucket/data/my-file.pdf

S3 Object (File) Size

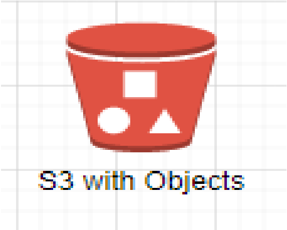

- S3 has unlimited storage
- Max Objects size is 5TB (5000GB)
- 5GB is the limit for an object to upload
- If the object is more than 5GB, use "multi-part upload"

S3 Lab

- Create a bucket
- Upload a file
- Download a file
- Access the via a browser
- Delete a file

Bucket Policies

Bucket Policy

- **JSON based policy**
  - **Resources**: buckets and Objects
  - **Actions**: set of permissions
  - **Effect**: Allow or Deny
  - **Principal**: The account or user to apply the policy to.

- **Use Cases:**
  - Granting Public Access to the bucket
  - Force Objects to be encrypted
  - Grant access to another AWS account (cross account)

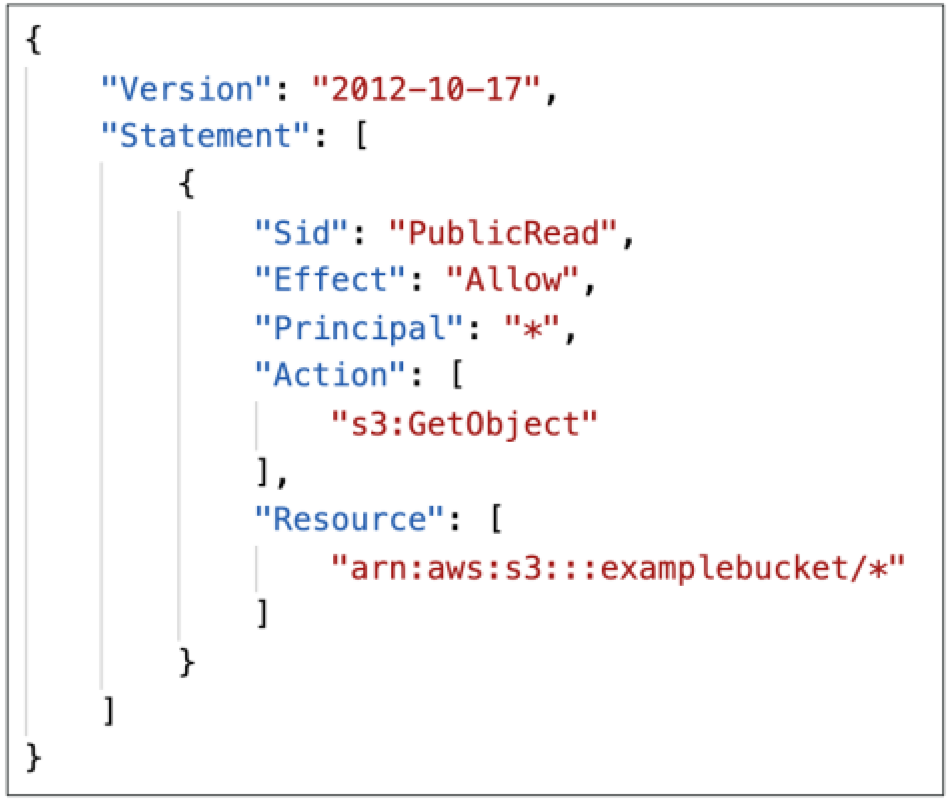

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "PublicRead",
            "Effect": "Allow",
            "Principal": "*",
            "Action": [
                "s3:GetObject"
            ],
            "Resource": [
                "arn:aws:s3:::examplebucket/*"
            ]
        }
    ]
}
```

Bucket Policy – Hands-On

- Uncheck "Block all access"
- Making a file public
- Test the file access from the browser
- Generate a Bucket Policy
- Paste it to S3 Bucket Policy Section

S3 Security

- **S3 Buckets Security access are controlled using:**
  - IAM Policies to IAM users or Roles.
  - Bucket Policies
  - Access Control List (ACL) for Objects and Buckets

AWS Identity and Access Management

- **Object Encryption:**
  - SSE-S3 – (Server-Side Encryption S3)
    - AWS Encrypts the data
  - SSE-KMS - (Server Side Encryption-Key Management Service)
    - AWS Managed
    - Customer Managed - KMS (Key Management Service)

AWS Key Management Service

Public Access – Using Bucket Policy

- Best practice to use Bucket Policy.

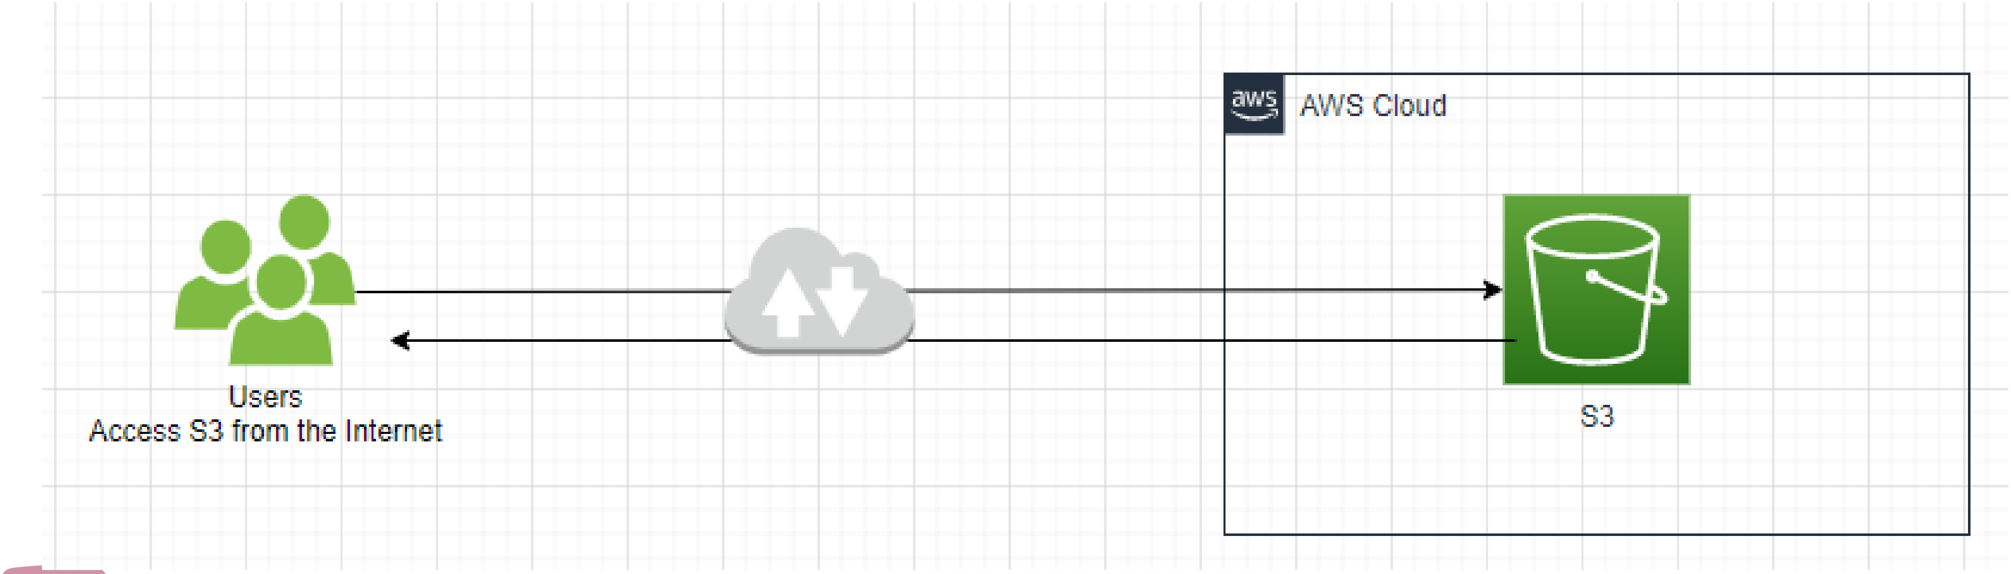

Bucket Access Using IAM Permission

- IAM Policy attached to a user account to access the bucket

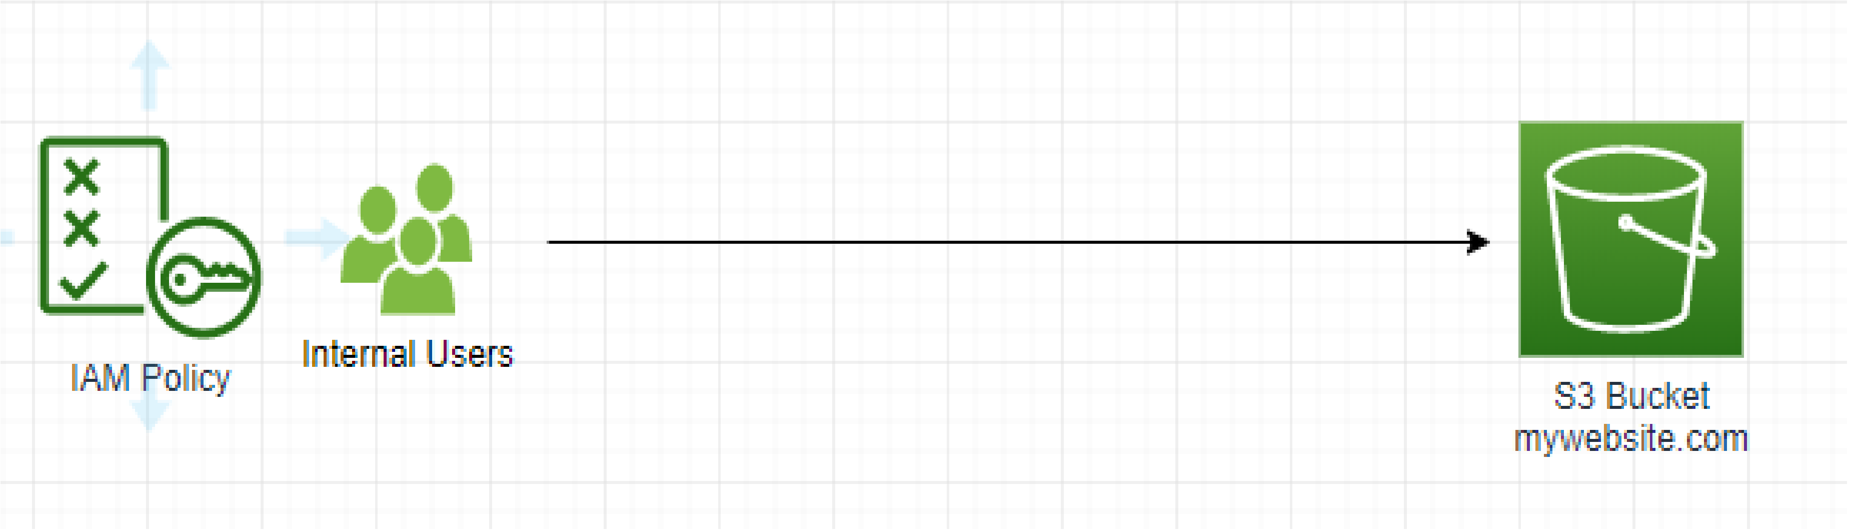

EC2 Instance access – Using IAM Roles

- Use IAM roles for AWS Resource to grant Bucket access.
- Example for EC2 Access:
- Using AWS Cli from the EC2

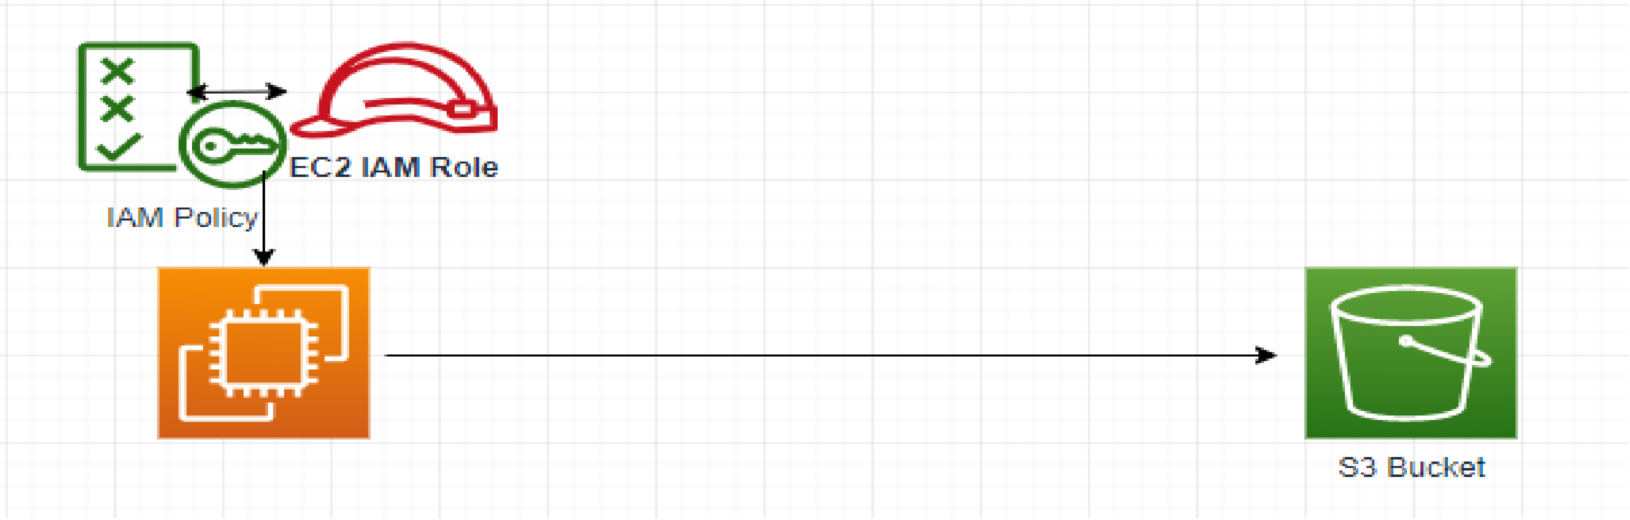

Cross Account Access – Using Bucket Policy

- One bucket can be accessed from one account to another account using bucket policy

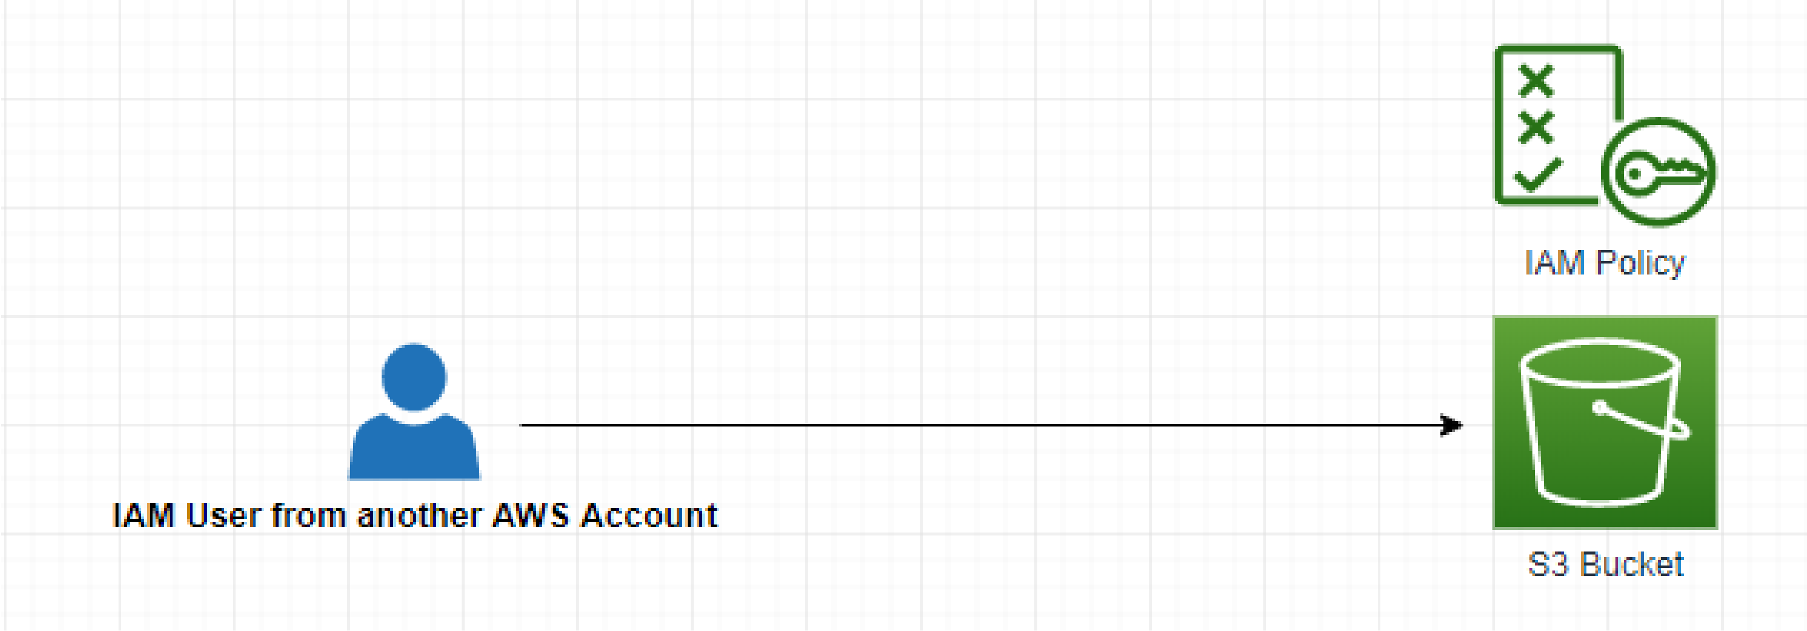

S3 Storage Classes

- Amazon S3 Standard - General Purpose
- Amazon S3 Standard-Infrequent Access (IA)
- Amazon S3 One Zone-Infrequent Access
- Amazon S3 Intelligent Tiering
- Amazon Glacier
- Amazon Glacier Deep Archive

- Ref: https://aws.amazon.com/s3/storage-classes/

S3 Standard – General Purposes

- 99.99% Available
- Used for frequent accessed data

- **Use cases:**
  - Data Lake
  - Big Data Analytic
  - Mobile and Gaming applications, websites

S3 Standard – Infrequent Access (IA)

- 99.99% Availability.
- Lower cost compared to standard.

- **Use cases:**
  - Suitable for data storage that is less frequently accessed but require fast restore when needed.
  - Data store for Disaster recovery, backups.

S3 Intelligent-Tiering

- 99.9% Availability
- Low latency
- High Throughput Performance

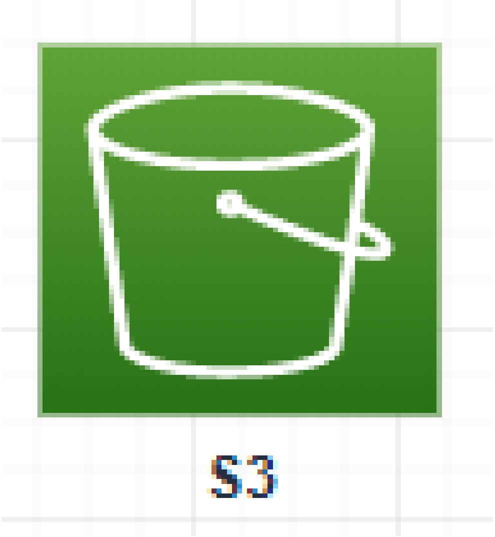

- **Use Case:**
  - Optimize cost by automatically move the data to cheaper storage classes based on these patterns:
    - Frequent access
    - Infrequent access

S3 One Zone - Infrequent Access (IA)

- 99.5% Availability
- Same as IA (Infrequent Access) but data is stored in a single AZ
- Lower cost compared to S3-IA

- **Use cases:**
  - Storing backup copies of on-premises data.

Amazon Glacier & Glacier Deep Archive

- Lowest cost object storage
- Use for Archiving / Backup purpose
- Different options for data retrieval based on the **Time+Fees** for retrieval.
- **Amazon Glacier – cheap:**
  - Fast retrieval: (1 to 5 hours)
  - Standard: (3 to 5 hours)
  - Bulk (5 to 12 hours)
- **Amazon Glacier Deep Archive- cheapest:**
  - Standard retrieval: (12 hours)
  - Bulk retrieval: (48 hours)
- **Use Cases:**

S3 Websites Hosting

- You can host static websites onS3 Bucket accessible in Internet
- Website URL looks like this:
    - <your-bucketname>.s3-website-.amazonaws.com
    - OR
    - <your-bucketname>.s3-website.<AWS-region>.amazonaws.com

Troubleshooting a 403 error:
If the bucket access is not set to public, and it is using bucket policy.
Make sure the bucket policy allows public reads! (**s3:GetObject**)

S3 Website- Public Access

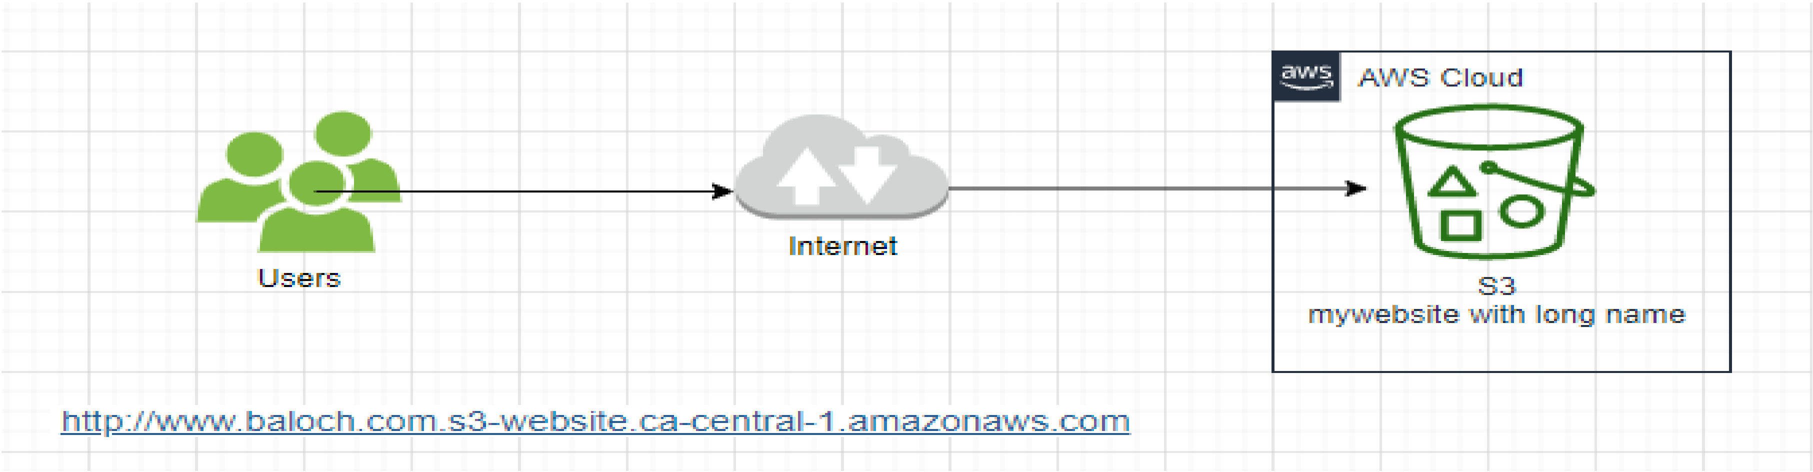

http://www.baloch.com.s3-website.ca-central-1.amazonaws.com

S3 Website – Proper Architecture – Pointing the website to a Domain Name

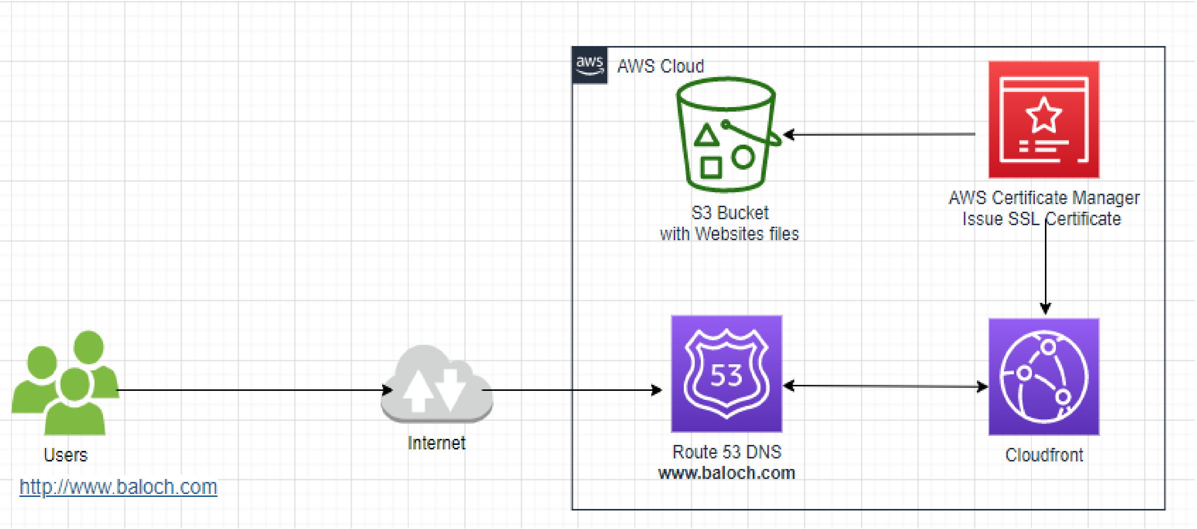

S3 Static Websites Hosting - Lab

1. Register a domain name (baloch.com) using AWS Route 53 Service.
2. Create a bucket with the domain name (Baloch.com)
3. Make the bucket public.
4. Enable Web Hosing on bucket.
5. Upload the website to Baloch.com bucket.
6. Create a public SSL certificate for the domain (Baloch.com) using AWS Certificate Manager service.
7. Create a distribution from AWS CloudFront service.
8. Create a DNS record using AWS Route 53 choosing CloudFront endpoint.
9. Verify the website is working publicly from the internet

AWS S3 Versioning

- Versioning allows to keep multiple versions of objects in one bucket.
- It creates version as 1,2,3
- Best practice to enable versioning against accidently deletes.
- Easy to roll back to a previous object version if deleted.
- Suspending versioning does not delete the previous versions.

- Enabling versioning:
- https://docs.aws.amazon.com/AmazonS3/latest/user-guide/enable-versioning.html

S3 – Server Access Logging

- Best practice to store all logs into a separate S3 Bucket.
- All requests to S3 authorized or denied get logged into another S3 bucket.
- Log data can be viewed and analyzed using Data Analysis tools such as **AWS Athena**. (Athena is a managed AWS service use to query log data.
- **Use Cases:**
  - Log all access to S3 buckets for troubleshooting.
  - Useful for root cause of an issue, audit usage or any data breach.

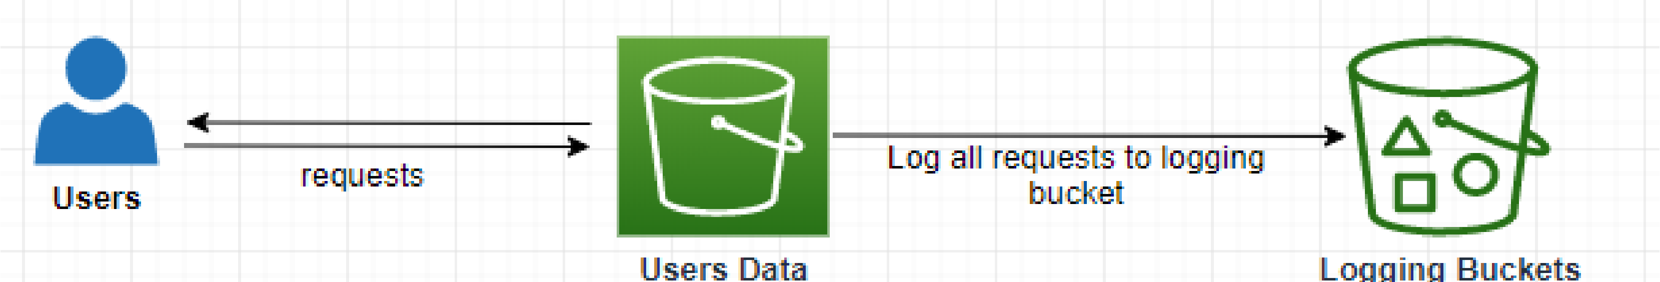

S3 Replication (Management Rule)

- Cross Region Replication (CRR)
- Same Region replication (SRR)
- Copying data from one S3 Bucket to another across accounts.
- Data is copied asynchronously.
- Require IAM permissions to S3.
- Require **S3 Versioning** to be enabled.
- Use Cases:
  - CRR=Compliance, lower latency access
  - SRR=Live replication between accounts (test and production)
-

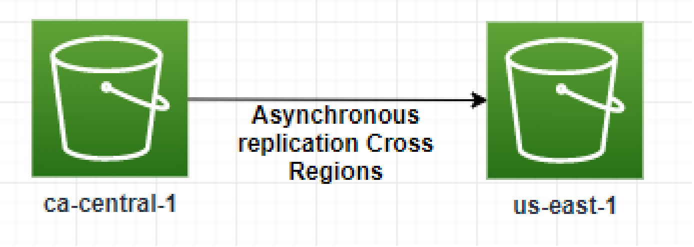

S3 Lifecycle Rule

- Lifecycle allows you to automatically transition objects to Standard -IA or Glacier storage class to save cost.

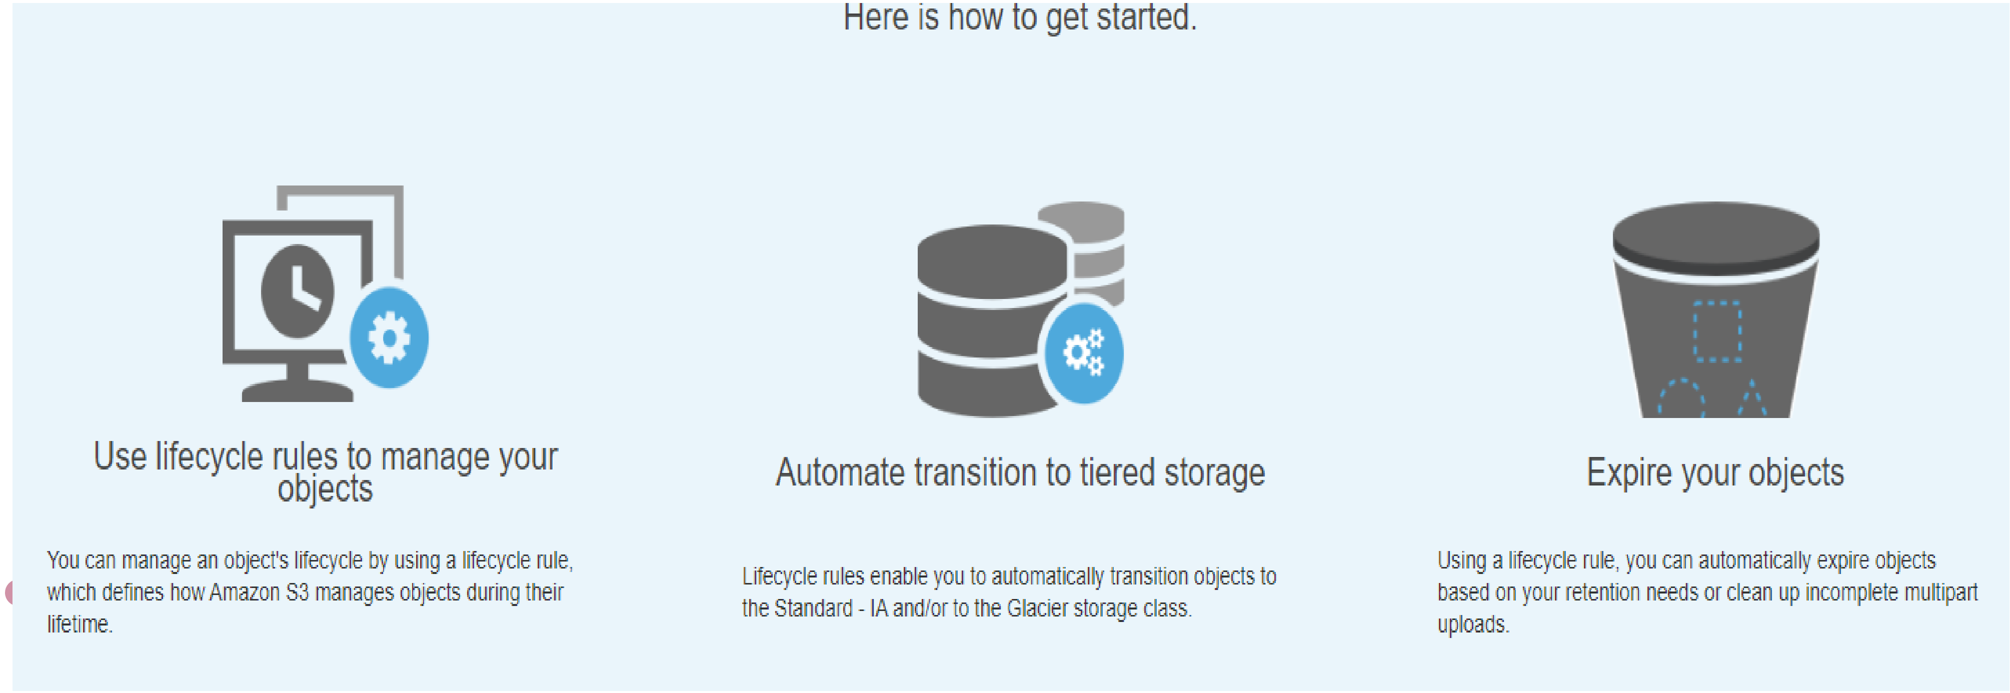

Here is how to get started.

Use lifecycle rules to manage your objects

You can manage an object's lifecycle by using a lifecycle rule, which defines how Amazon S3 manages objects during their lifetime.

Automate transition to tiered storage

Lifecycle rules enable you to automatically transition objects to the Standard - IA and/or to the Glacier storage class.

Expire your objects

Using a lifecycle rule, you can automatically expire objects based on your retention needs or clean up incomplete multipart uploads.

Shared Responsibility Model for S3

- Infrastructure (global security, durability, availability, sustain concurrent loss of data in two facilities)
- Configuration and vulnerability analysis
- Compliance validation

- S3 Versioning
- S3 Bucket Policies
- S3 Replication Setup
- Logging and Monitoring
- S3 Storage Classes
- Data encryption at rest and in transit

AWS Storage Gateway

Hybrid Cloud Storage

AWS Storage Gateway

- Bridge between corporate data and cloud data in S3
- An EC2 virtual server
- Types of Storage Gateway:
    - File Gateway
    - Volume Gateway
    - Tape Gateway

- Use Cases:
    - Disaster Recovery
    - Backup/Restore

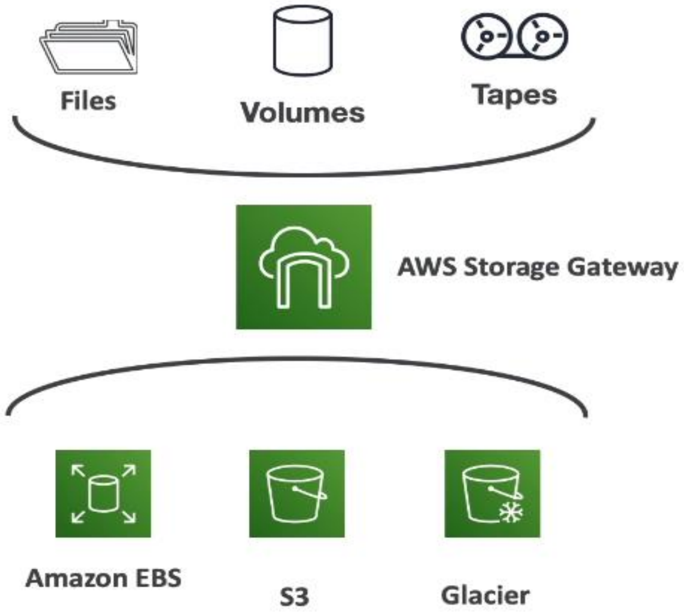

Storage Gateway Architecture

- Understand better

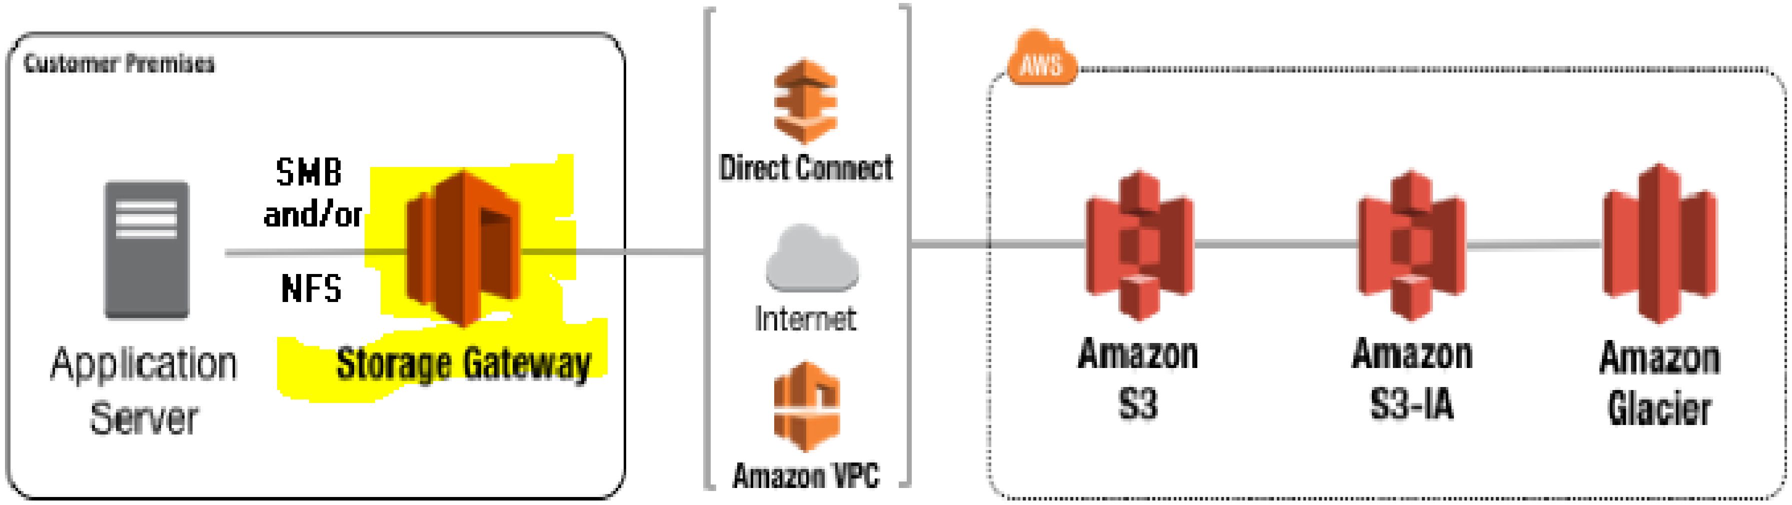

Data Migration

From Corporate to AWS Cloud

AWS Snowball

- A physical data transport solution.
- Helps to move TBs or PBs of data in or out of AWS.

- **Use Case:**
  - Moving large migration from on-premise data center to AWS Cloud.
  - Cloud Migration
  - Disaster recovery
  - Datacenter Decommission

AWS Snowball

- Migrating using upload directly to S3:

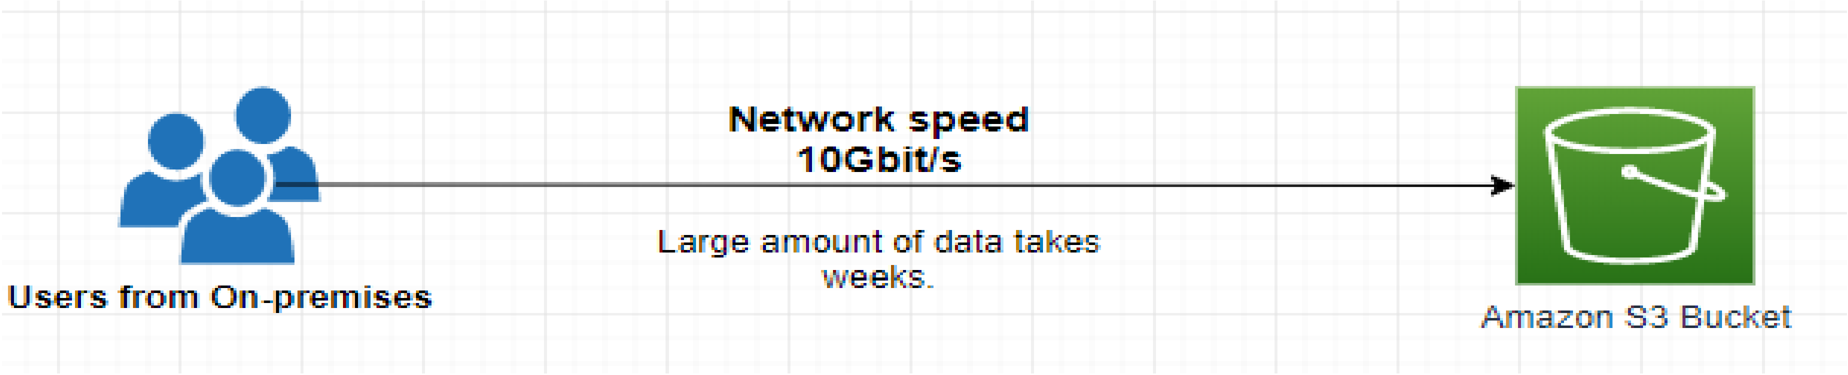

- Migrating data with Snowball:

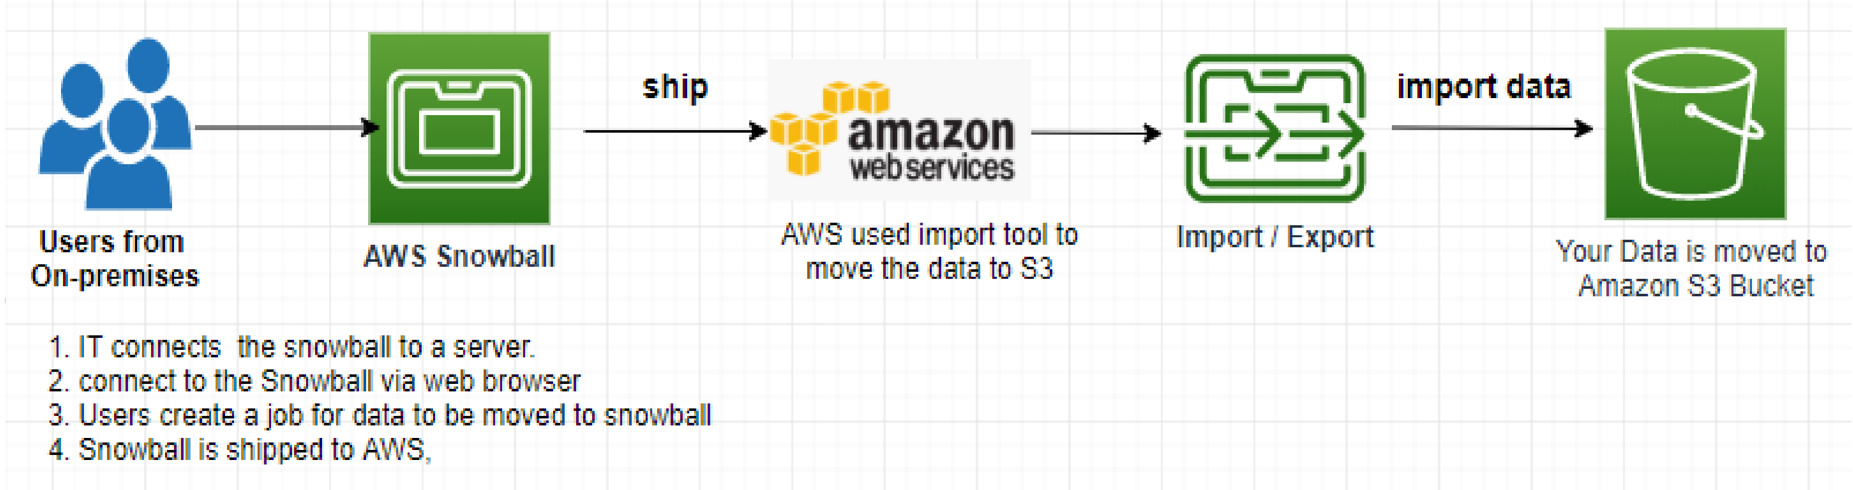

1. IT connects the snowball to a server.
2. connect to the Snowball via web browser
3. Users create a job for data to be moved to snowball
4. Snowball is shipped to AWS,

Ordering a Snowball

- Use **AWS Console** to create a request.
- It will be shipped to your address listed in your account.
- Once you receive it:
  - Connect the Snowball to your on-premise server
  - Copy data using the Snowball software client
- Once done, you ship the device back to AWS.
- AWS will load the data into S3 bucket
- **Data Security:**
  - AWS will completely wipe out the data according to the regulations.
  - Data encryption end to end.

Amazon Snowball Edge

- Up to 100TB capacity
- Built in with a custom EC2 instance

- **Use Cases:**
  - Data Migration, Machine Learning, image data.

AWS Snowmobile

- Transfer **exabytes** of data with multiple Snowmobiles.
- 45-foot long shipping container with computing/networking power inside.
- It is driven close to your corporate data center.
- Connected to your data center network.
- **100 Petabytes** (100,000TBs) Capacity
  (1 Exabytes = 1,000 Petabytes = 1,000,000 Terabytes = 1,024,000,864 Gigabytes)
- Useful for migrating data more **than 10 PB**

- https://aws.amazon.com/snowmobile/

Amazon S3 Summary (Exam Topics)

- **Buckets**: Global unique name, tied to a region.
- **S3 Security**:
  - IAM, Policy
  - S3 Bucket Policy
  - S3 Encryption
- **S3 Websites**: To host a static website.
- **S3 Versioning**: Prevent accidental deletes.
- **S3 Access Logs**: For audit and troubleshooting.
- **S3 Replication**: Same-region or cross region replication (Version).
- **S3 Storage Classes:**
  - Standard, IA, IZ-IA, Intelligent, Glacier, Deep Archive.
- **S3 Lifecycle Rules**: Allows to move objects between S3 Classes.
- **Snowball/Snowmobile:** import corporate data onto S3 through physical device..
- **Storage Gateway**: Hybrid solution to extend CORPORATE storage to Amazon S3.

Question-1

- **Which S3 Storage Class is the most cost-effective for archiving data with no retrieval time requirement?**
  - Amazon Glacier
  - Amazon Glacier Deep Archive
  - Amazon S3 Standard-Infrequent Access
  - Amazon S3 Intelligent Tiering

Question-2

- **What hybrid AWS service is used to allow on-premises servers to seamlessly use the AWS Cloud at the storage layer?**
  - Elastic Block Storage
  - Snowball
  - S3
  - Storage Gateway

Question-3

- **Which of the following services is a petabyte-scale data moving service (as a fleet) in or out of AWS with computing capabilities?**
  - Snowball
  - Snowball Edge
  - Snowmobile

Question-4

- **Which of the following is an exabytes-scale data moving service in or out of AWS?**
  - Snowball
  - Snowball Edge
  - Snowmobile

Question-5

- **Where are objects stored in Amazon S3?**
  - Folders
  - Buckets
  - Files
  - Bin

Question-6

- **What can you use to define actions to move S3 objects between different storage classes?**
  - Scaling Policy
  - Bucket Policy
  - Lifecycle Rules
  - Replication

Question-7

- **Which S3 Storage Class is suitable for less frequently accessed data, but with rapid access when needed, while keeping a high durability and allowing an Availability Zone failure?**
    - S3 Standard
    - Glacier
    - S3 One Zone-Infrequent Access
    - S3 Standard-Infrequent Access

EFS

Elastic File System (Linux)

EFS – Elastic File System

- AWS managed NFS (Network File System)
- Can be mounted and shared on multiple EC2 Instances
- Works only on Linux EC2 Instance.
- Highly Available, More expensive then EBS
- Pay per use

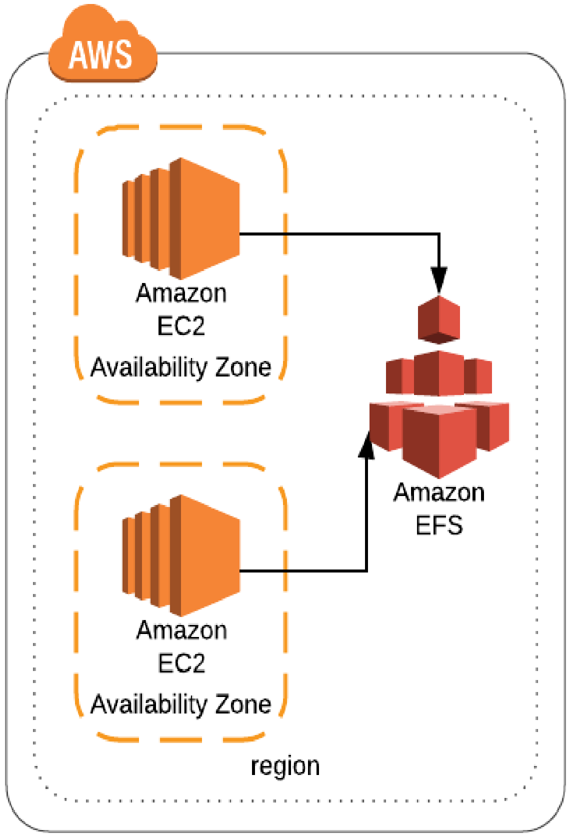

EFS Architecture Example

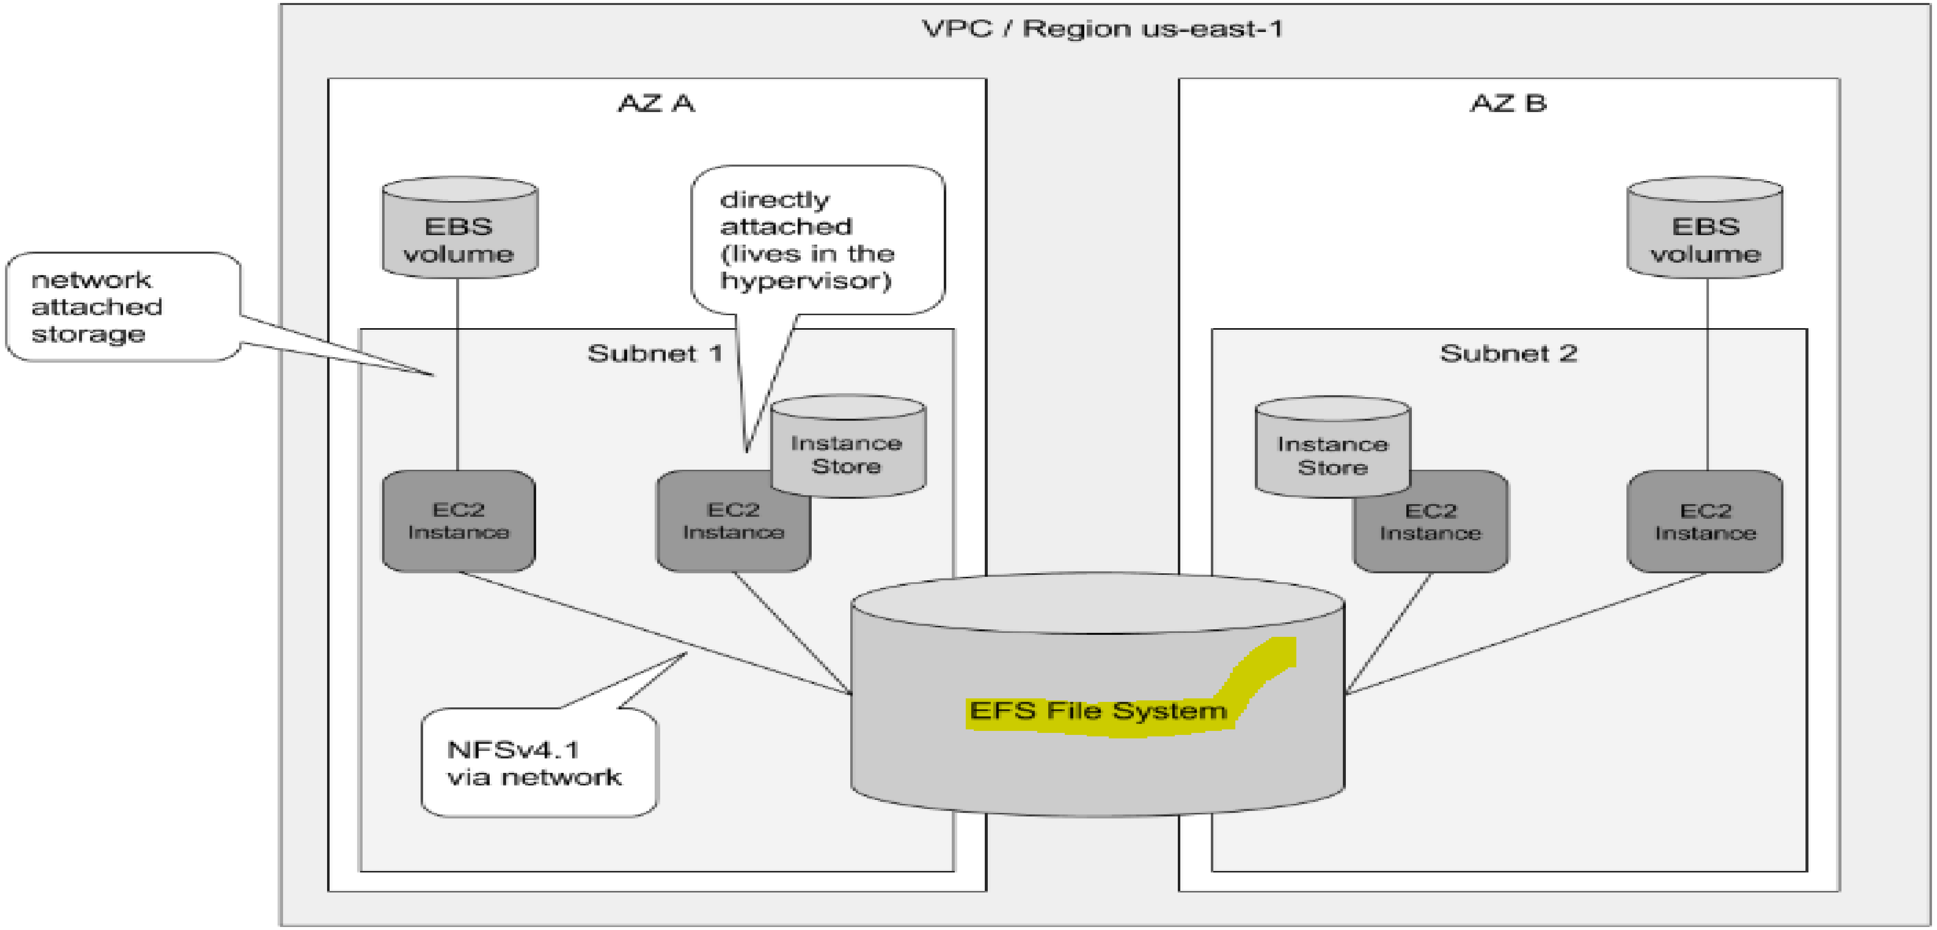

EFx

File Server for Windows

FSx – File Server for Windows

- Managed Service
- Similar to EFS but **FSx** is for Windows Servers.
- Pre-requisite is Microsoft Active Directory.

FSx Architecture Example

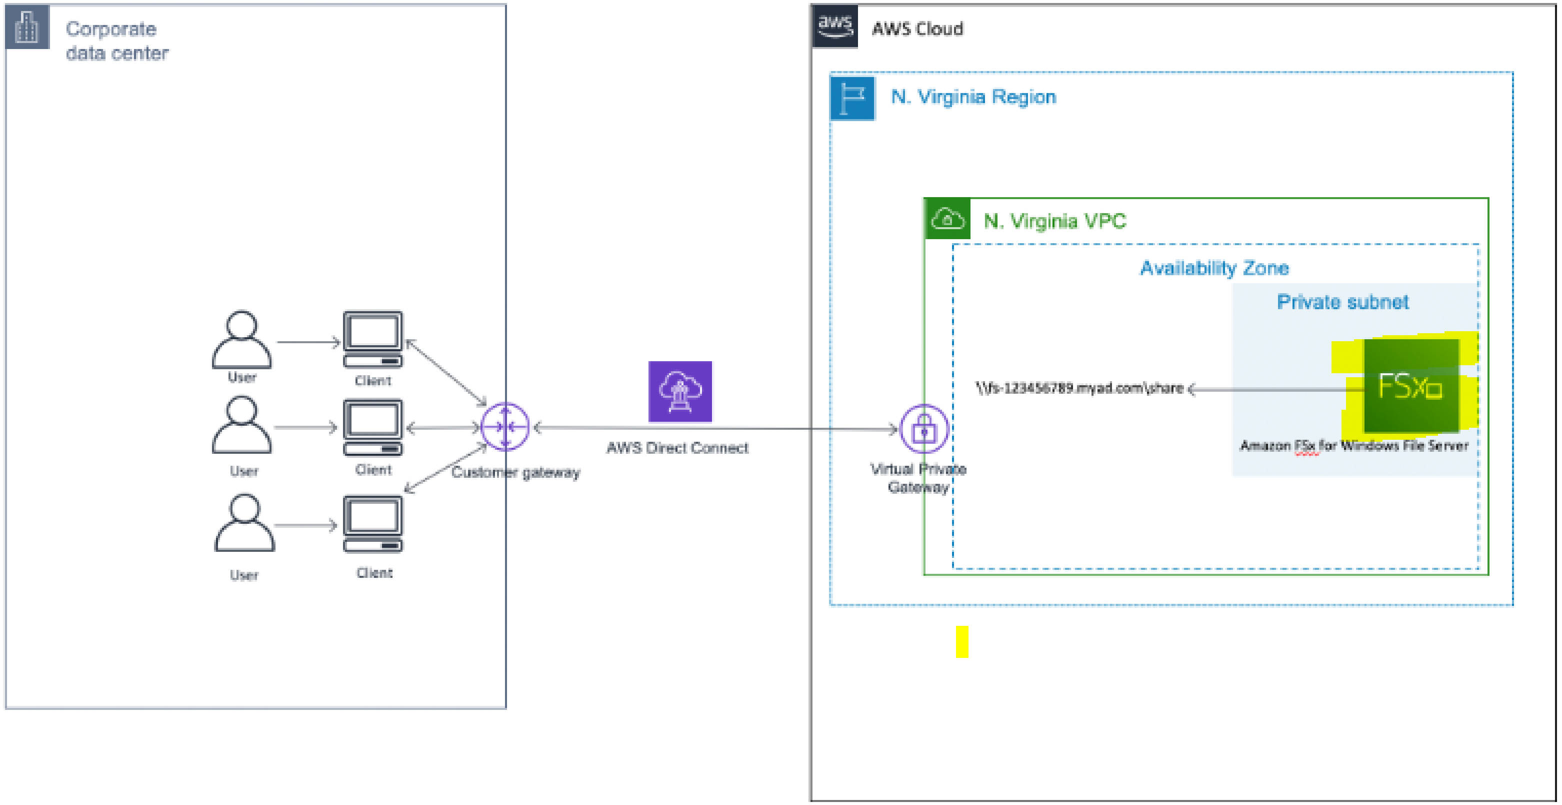

Amazon RDS
(Relational Database Service)

Instruction to Amazon RDS

RDS = Relational Database Service

- AWS managed Service.
- Allows you to create databases in the Cloud managed by AWS.
- RDS offers these types of Database **Engine** (type of databases).
  - Microsoft SQL Server
  - Oracle
  - MySQL
  - MariaDB
  - Aurora (AWS owned Database)
  - Postgres

Relational = Like Excel spreadsheets, with links between them.

Advantage of RDS vs deploying DB on EC2

- Automatic provisioning
- Automated patching
- Automated backup
- Multi-AZ for Disaster Recovery
- Scaling features

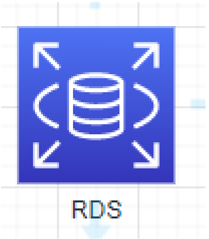

- You cannot connect to RDS instance since back-end is managed by AWS.
- You just deploy your database and configure as per your need.

RDS Architecture

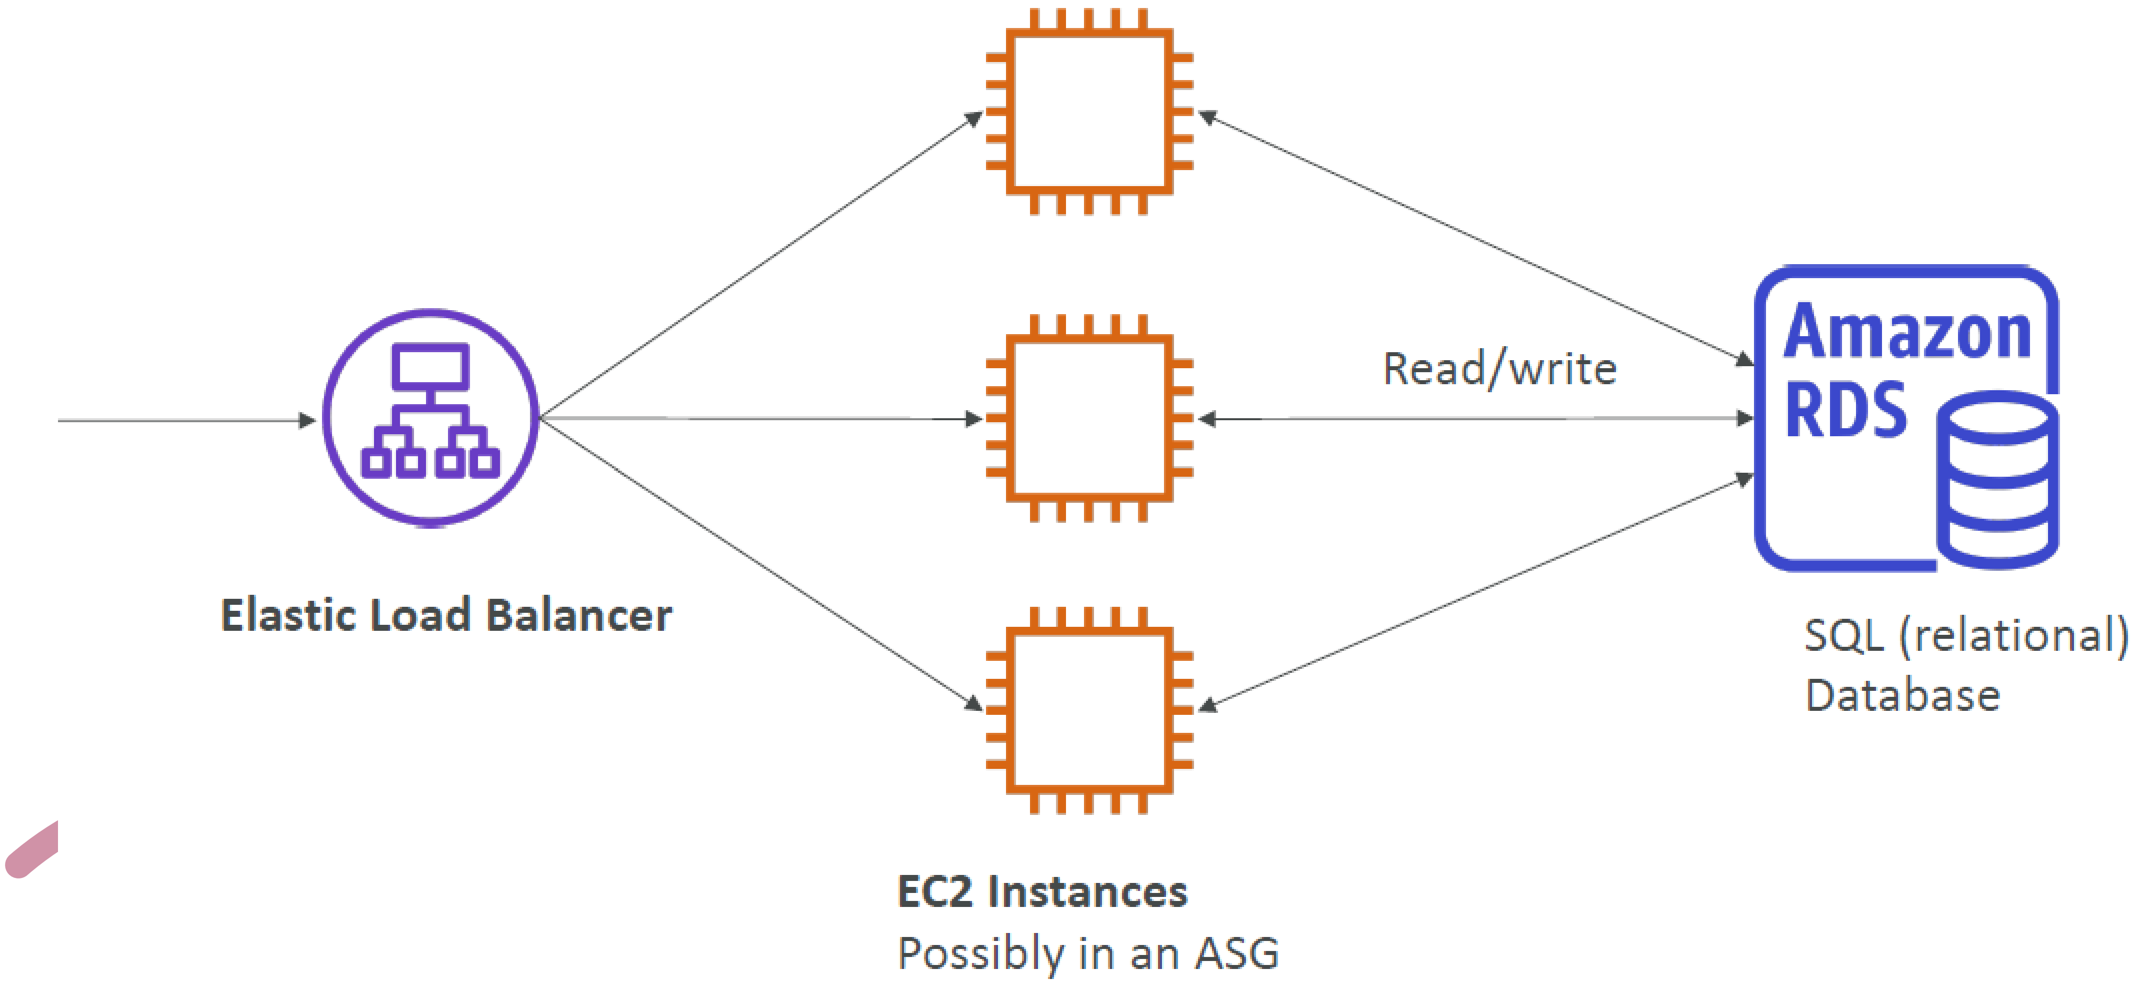

Deploy a test Database

- Create an RDS Instance eligible for Free Tier for SQL Express or MySQL.
- Verify the RDS costs for none-free trier instance.
- Create a DB snapshot
- Using a DB client from your desktop to connect to RDS Endpoint.
- Create a Test SQL database
- Delete the RDS instance