Amazon EFS

❖ AMAZON EFS (ELASTIC FILE SYSTEM):

- Amazon **Elastic File System (EFS)** Provides scalable file storage for EC2 instances and on-premises resources.
- EFS is a regional and zone-based service where the data is replicated across multiple availability zones in your AWS region.
- Can attach multiple EC2 instances to EFS as a centralized file storage, unlike an EBS volume which can only be attached to one EC2 instance at a time.
- EFS is growing and shrinking automatically as you add and remove files.
- You pay only for the storage used by your file system and there is no minimum fee or setup cost.
- EFS provides file system access semantics, such as strong data consistency and file locking.
- EFS supports the **Network File System (NFSV4)** protocol.

NOTE: Microsoft Windows based EC2 instances are not supported.

EFS BENEFITS:

- Dynamic Elasticity
- Shared File Storage
- Seamless Integration
- Automatically Scales
- Scalable Performance
- Low Cost
- Secure

EFS FEATURES:

- Elastic and Scalable
- Shared, concurrent file access
- Performance modes
- Throughput modes
- Storage Classes and Lifecycle Management
- Highly available and durable
- Security and Compliance
- Encryption
- Data transfer and backup
- Low latency file operations

➢ EFS STORAGE CLASSES:

- Amazon EFS offers a range of storage classes that are designed for different use cases.
- EFS storage classes are:
  - EFS Standard & Standard–Infrequent Access (Standard-IA)
  - EFS One Zone & One Zone–Infrequent Access (EFS One Zone-IA).

EFS STANDARD AND STANDARD-IA:

- These storage classes are regional storage classes that are designed to provide continuous availability to data, even when one or more Availability Zones in an AWS Region are unavailable.
- **EFS Standard** storage class is used for frequently accessed files. It is the storage class to which customer data is initially written for Standard storage classes.
- **Standard–IA** storage class reduces storage costs for files that are not accessed every day.

EFS ONE ZONE AND ONE ZONE–IA:

- storage classes are designed to provide continuous availability to data within a single Availability Zone.
- The EFS One Zone storage classes store file system data and metadata redundantly within a single Availability Zone in an AWS Region.
- **EFS One Zone**–Standard is used for frequently accessed files.
- **EFS One Zone–IA** storage class reduces storage costs for files that are not accessed every day.

➢ EFS BACKUP & RESTORE:

- There are two options for backing up EFS file systems.
  - EFS-to-EFS backup solution.
  - AWS backup service

EFS-TO-EFS:

- This backup method is suitable for all amazon EFS file systems in all AWS regions.
- AWS DataSync is recommended for the backup, basically copying data files between two EFS file systems in the same or different regions, and in the same or different accounts.

AWS BACKUP:

- EFS integrates with the AWS backup service which can be used to provide a comprehensive backup solution for EFS file systems.
- AWS Backup is designed to simplify the creation, migration, restoration, and deletion of backups, while providing improved reporting and auditing.
- AWS Backup provides incremental backups of the EFS file systems.

RESTORE:

- We can restore from a backup to a new EFS file system or to the same EFS File system.
- When you perform a full restore, the entire file system is restored.

➢ EFS LIFE CYCLE MANAGEMENT:

- EFS lifecycle management automatically manages cost-effective file storage for your file systems.
- Lifecycle management migrates files that have not been accessed for a set period of time to the EFS Standard–IA or One Zone–IA storage class, depending on your file system.
- Amazon EFS lifecycle management uses an internal timer to track when a file was last accessed, and not the POSIX file system attributes that are publicly viewable.
- You define that period of time using the Transition into IA lifecycle policy.
- Amazon EFS supports two lifecycle policies.

TRANSITION INTO IA:

- It instructs lifecycle management when to transition files into the file systems' Infrequent Access storage class.
- The Transition into IA lifecycle policy has the following values:
  - None
  - 7, 14, 30, 60, 90 days since last access

TRANSITION OUT OF IA:

- It instructs intelligent tiering when to transition files out of IA storage.
- The Transition out of IA lifecycle policy can have the following values
  - None
  - On first access

➢ PERFORMANCE MODES:

- File system performance is typically measured by using the dimensions of latency, throughput, and Input/Output operations per second (IOPS).
- Amazon EFS performance across these dimensions depends on your file system's configuration.

THROUGHPUT MODE:

- There are two throughput modes to choose from for your file system.
  - Enhanced
  - Bursting

ENHANCED:

- Provides more flexibility and higher throughput levels for workloads with a range of performance requirements.

BURSTING:

- Provides throughput that scales with the amount of storage for workloads with basic performance requirements.
- File-based workloads are typically spiky, driving high levels of throughput for short periods of time, and low levels of throughput the rest of the time.

- To support a wide variety of cloud storage workloads, Amazon EFS offers two performance modes.
  - General Purpose mode
  - Max I/O mode

ELASTIC:

- Use this mode for workloads with unpredictable I/O. With Elastic Throughput, performance automatically scales with your workload activity and you only pay for the throughput you use (data transferred for your file systems per month).

PROVISIONED:

- Use this mode if you can estimate your workload's throughput requirements. With Provisioned mode, you configure your file system's throughput and pay for throughput provisioned.

➢ EFS DATA ENCRYPTION:

- EFS supports two forms of encryption for file systems, encryption of data in transit and encryption at rest.
- It uses AWS KMS for key management.
- EFS automatically manages the keys for encryption in transit, client will be encrypted at rest.

➢ NETWORK ACCESS:

MOUNT TARGETS:

- EFS system has mount targets. Mount targets are ENIs that are deployed in subnet we specify.
- Mount targets have security groups that should allow the communication from the instance's security groups.
- NFSV4 protocol runs on TCP port 2049.
- We can configure one mount target per AZ to support all EC2 instances in any of the subnets in that AZ. Mount targets are Highly Available.