

Amazon Elastic Cloud Compute (EC2): Amazon Elastic Block Store (EBS)



Andru Estes

Principal Author

 andru-estes



Amazon EC2 Storage



Two Primary Storage Categories for EC2

Block Storage

File Storage



File Storage

Amazon EFS

Scalable file storage for use with
Amazon EC2

Amazon FSx

High-performance file systems in
the cloud



Block Storage

Amazon EBS

Durable, block-level storage volumes

Instance Store

Temporary block-level storage for instances



File storage services will be discussed later on!



**For now, we will explore the
block storage options.**





Amazon Elastic Block Store (EBS)



Amazon EBS

Scalable, high-performance block storage resources that can be used with Amazon EC2 instances




The volume created and attached at instance launch is referred to as the root data volume.



Amazon EBS Concepts



Network attached drives that act like normal local drives, which allow you to persist your data after an instance is terminated



Normal EBS volumes (*Non multi-attach*) are bound to a single EC2 instance at one time



EBS volumes are deployed and bound to a single Availability Zone



Amazon EBS Information

Capacity

You set a desired amount of space/capacity, and can increase as you need to

Pricing

You get billed for any provisioned capacity, regardless if the space is filled or not



By default, the root EBS volume that is created and attached to an instance at launch is deleted when that instance is terminated!



You can optionally tell AWS to keep the volume by ensuring that “Delete on Termination” is not selected for the volume.



Additional volumes are the opposite, as they are not deleted on termination by default.





Amazon EBS Use Cases

You use these the same way you would use any typical system disk

Designed for production/critical workloads

Example use cases include the following:

- Create a file system
- Run a database
- Run an operating system
- Store data
- Install applications

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



Additional Important Concepts

EBS volumes are automatically replicated within a single Availability Zone to protect against hardware failures

Dynamically increase capacity and change the volume type with no downtime or performance impact to live systems





EBS Volume Types




Amazon EBS Volume Types and Use Cases



General Purpose SSD (gp2, gp3): Web Server Hosting. Workloads that require a good balance of performance and cost-effectiveness.




Provisioned IOPS SSD (io1, io2): EC2-hosted Relational Database with High I/O Demands. Workload that requires consistent and high levels of I/O performance.



Magnetic (sc1 and st1): Data Archiving and Backup. Any large volumes of infrequently accessed data that need to be stored cost-effectively for long-term retention.



Learning the Different Types - General Purpose



gp2:


- Suitable for boot disks and general applications
- Up to 16,000 IOPS per volume
- Up to 99.9% durability

gp3:

- Suitable for high performance applications
- Predictable 3,000 IOPS baseline performance and 125 MiB/s regardless of volume size
- Up to 99.9% durability



Learning the Different Types - Provisioned IOPS



io1:

- Suitable for OLTP and latency-sensitive applications
- 50 IOPS/GiB and up to 64,000 IOPS per volume
- High performance and most expensive
- Up to 99.9% durability

io2:

- Suitable for OLTP and latency-sensitive applications
- 500 IOPS/GiB and up to 64,000 IOPS per volume
- 99.999% durability
- Latest generation Provisioned IOPS volume




EBS Multi-Attach Volumes

Provisioned IOPS EBS volumes that can attach to up to 16 EC2 instances at once

Only work within the same AZ

Instances can read and write to the volumes


You must use cluster-based filesystems (e.g. GlusterFS)




Multi-Attach Volume Use Cases



Shared File Systems



Distributed Databases



High-Performance
Computing



IOPS Versus Throughput

IOPS

Measures the number of read and write operations per second

Important metric for quick transactions, low-latency apps, and transactional workloads

The ability to action reads and writes very quickly

Choose Provisioned IOPS SSD (io1 or io2)

vs.

Throughput

Measures the number of bits read or written per second (MB/s)

Important metric for large datasets, large I/O sizes, and complex queries

The ability to deal with large datasets

Choose Throughput Optimized HDD (st1)



Encryption of EBS volumes





Amazon EBS Encryption

You are able to encrypt both the boot and data volumes of an EC2 instance

Encryption operations happen on the servers that host EC2 instances, meaning there is very little performance impact

The following data types are encrypted:

- Data at rest in the volume
- Data moving between the volume and EC2
- Snapshots created from the volume
- Volumes created from the snapshots

The encryption leverages a service called AWS KMS for AES-256 encryption

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

Exam Pro Tip: If you have an unencrypted EBS volume, you can create a snapshot to create a new, encrypted copy of the volume



This is because you cannot directly encrypt existing unencrypted volumes or snapshots.



**EBS Snapshots are coming
up next...**





Amazon EBS Snapshots

What is an Amazon EBS Snapshot?



EBS Volumes and Snapshots




Point-in-time photographs (copy) of volumes that are incremental in nature




The first snapshot will take time to create because they are a full capture



Snapshots capture all data and configurations at time of creation



You can share snapshots with AWS accounts, as well as between Regions



Snapshots are stored in Amazon Simple Storage Service (S3)



EBS Snapshot Lifecycles



Archive

Amazon EBS Snapshots Archive allows you to store snapshots for lower costs if you do not need fast retrieval




Recycle Bin

Recycle Bin is a data recovery feature to restore accidentally deleted Amazon EBS snapshots and EBS-backed AMIs



Amazon EBS Fast Snapshot Restore



Enables you to create a volume from a snapshot that is fully initialized at creation

Eliminates the latency of I/O operations on a block when it is accessed for the first time

This means that they instantly deliver all of their provisioned performance

Fast snapshot restore can be enabled on snapshots with a size of 16 TiB or less!



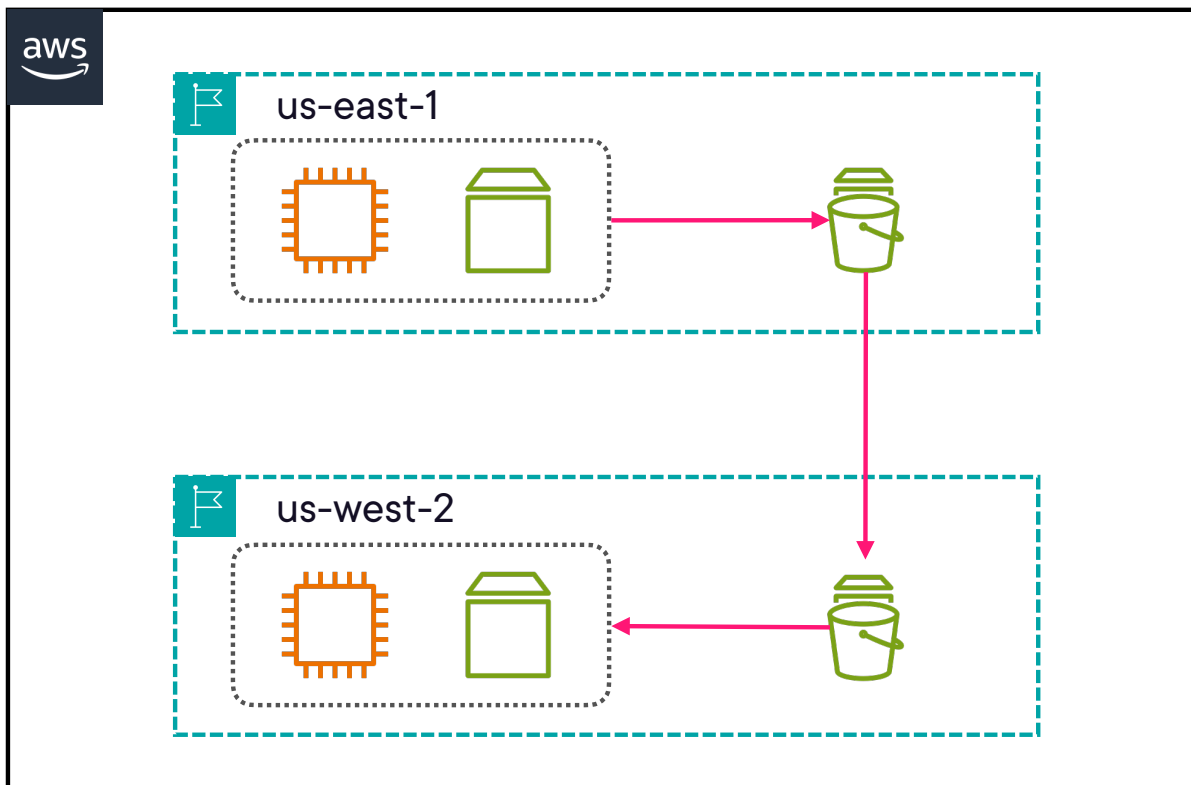
Exam Pro Tip 1: You have no access to the underlying S3 storage that hosts your EBS Snapshots!



Exam Pro Tip 2: Use EBS Snapshots if you need to copy data from an EBS volume to a different AZ or Region.



Demo: Copying Snapshots Between AWS Regions



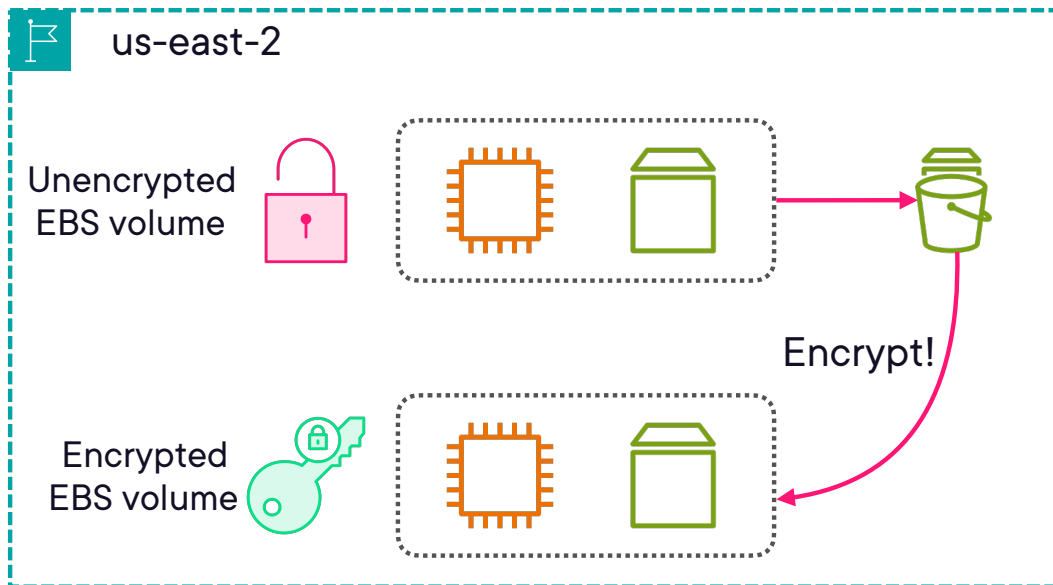
We will take an EBS snapshot of a running instance's EBS volume in us-east-1.

To simulate a disaster recovery effort, we will copy the EBS snapshot to the us-west-2 Region.

After completion, we will restore the snapshot to a new EC2 instance.



Demo: Creating Encrypted Version of Unencrypted EBS Snapshots



Running EC2 instance with an unencrypted EBS volume will have an EBS snapshot taken.

After snapshot creation, we will specify that we want an encrypted version using AWS KMS.

Once the snapshot is ready, we will restore it to a new EC2 instance and see the same data.





EC2 Instance Stores




Temporary block-level storage for your instance




Instance Store Concepts




Located on disks that are physically attached to the host of the EC2




Perfect for temporary storage like buffers, cache, or scratch data



Made up of one or more volumes exposed as block devices on the OS




Sizing depends on the number of available devices and EC2 type/size



Only attached at EC2 instance launch, and not after launch



Instance Store Data Persistence



Reboots

Data persists when the instance is rebooted so that you can use it again after restarting



Stopped, Hibernated, Terminated

When these occur, every block of the instance store volume is cryptographically erased



**Long Story Short: Do not use
Instance Stores to save
critical data!**





Module Summary and Exam Tips



EBS volumes are simply virtual hard disks that get attached to EC2 instances.



You need a minimum of one volume per EC2 instance, which is called the root device volume.



Amazon EBS Volume Types

General Purpose SSD

gp2, gp3

Provisioned IOPS SSD

io1, io2

Magnetic

sc1, st1



EBS SSD Volumes - General Purpose

gp2

Suitable for boot disks and general applications

Up to 16,000 IOPS per volume

Up to 99.9% durability

vs.

gp3

Suitable for high performance applications

Predictable 3,000 IOPS baseline performance and 125 MiB/s regardless of volume size

Up to 99.9% durability



EBS SSD Volumes - Provisioned IOPS

io1

Suitable for OLTP and latency-sensitive applications

50 IOPS/GiB

Up to 64,000 IOPS per volume

High performance and most expensive

Up to 99.9% durability

vs.

io2

Suitable for OLTP and latency-sensitive applications.

500 IOPS/GiB

Up to 64,000 IOPS per volume

Higher performance and less expensive

Up to 99.999% durability

EBS HDD Volumes - Magnetic

st1

Suitable for big data, data warehouses, ETL

Max throughput is 500 MB/s per volume

Cannot be a boot volume.

Up to 99.9% durability

vs.

sc1

Max throughput of 250 MB/s per volume

Less frequently accessed data

Cannot be a boot volume

Lowest cost

Up to 99.9% durability



Multi-attach volumes can attach to up to 16 instances in the same AZ at one time, and require a cluster-based file system!



EBS Volumes and Snapshots




Volumes exist in EBS, whereas snapshots exist in S3 (*No access to this*)




Point-in-time photographs of volumes that are incremental in nature



The first snapshot will take some time to create due to the size



Share snapshots between AWS accounts, as well as between Regions



You can resize EBS volumes on the fly, and change the volume types



EBS Encryption Review

All data at rest is encrypted inside the volume using AWS KMS

All data in flight moving between the instance and the volume is encrypted

All snapshots taken of the volume are encrypted

All volumes created from the snapshots are encrypted



Use snapshots to create an encrypted version of an EBS volume that is not currently encrypted.



Instance Stores

Instance store volumes are sometimes called ephemeral storage (temporary)

These go away when stopped, but they are perfect for temporary storage (buffers, cache, scratch data, etc.)

