Amazon S3 (Simple Storage Service):

Amazon S3 is a scalable object storage service provided by AWS. It allows users to store and retrieve any amount of data from anywhere on the internet. S3 is widely used for its durability, scalability, and ease of integration with other AWS services.

How S3 Works:

1.**Object-based Storage**: S3 stores data as objects, which include the data itself, metadata, and a unique identifier. These objects are stored in "buckets."

2.Bucket Structure: A bucket is a container for objects. Users can upload objects to buckets, and each object is identified by a unique key.

3.Global Access: S3 allows data access over the internet using API calls, SDKs, or the AWS Management Console.

4.High Durability & Availability: S3 is designed for 99.999999999% (11 nines) durability, with data automatically replicated across multiple availability zones.

Use Cases of S3:

- **Data Storage & Backup**: Regular backups for applications, systerns, and databases. Big Data & Analytics: Storing large datasets for analysis and processing.
- Website Hosting: Hosting static websites (HTML, CSS, JavaScript).
- **Content Distribution**: Storing media files for distribution (e.g., videos, software).
- **Disaster Recovery:** Storing critical data for quick recovery

Difference between EBS and S3:

EBS (Elastic Block Store):

Type: Block storage for EC2 instances.

Use Case: Low-latency storage for databases, file • •

Attachment: Only attached to one EC2 instance at a time.

• **Type**: Object storage for unstructured data (images, backups).

• **Use Case**: Highly scalable and accessible from anywhere.

• **Attachment**: Not tied to any single instance, can be accessed globally

Cold Data vs. Hot Data:

Hot Data: Frequently accessed data requiring lowlatenc storage (e.g., S3 Standard or EBS).

Cold Data: Rarely accessed data stored in lower-cost solutions (e.g., S3 Glacier) for long-term retention.

Object Storage Class vs. Block Storage Class:

Object Storage (S3): Stores data as objects with unique identifiers, suited for unstructured data (documents, media).

Block Storage (EBS): Stores data in blocks, similar to a hard drive, designed for structured data (databases, file systems).

S3 Storage Classes:

1.S3 Standard: General-purpose storage for frequently accessed data. All Bookmarks

2.S3 Intelligent-Tiering: Automatically moves data between frequent and infrequent access tiers to save costs.

3.S3 Standard-IA (Infrequent Access): Lower cost for infrequently accessed data.

4.S3 One Zone-IA: Similar to Standard-IA but stores data in a single availability zone.

5.S3 Glacier: Low-cost archival storage for data rarely accessed.

6.S3 Glacier Deep Archive: Lowest-cost storage for long-term, rarely accessed data.

Advantages of S3:

- **Scalability**: Automatically scales to accommodate any amount of data.
- **Durability**: 99.999999999% durability with automatic replication across multiple locations.
- **Security**: Supports encryption, access control policies, and integrates with AWS Identity and Access Management (IAM).
- **Cost-Effective**: Pay for what you use, with multiple storage classes to optimize cost based on data access frequency.

Disadvantages of S3:

- **Access Latency**: Not suitable for low-latency, high-
- Performance applications like databases (use EBS for that).
- **Complexity in Management**: Managing large
- Datasets or setting up lifecycle policies may become complex for certain use cases.

Features of S3:

- **Versioning**: Tracks changes to objects, allowing you Of revert to previous versions.
- **Lifecycle Management**: Automates the movement of Data to lower-cost storage classes.
- **Cross-Region Replication**: Automatically replicates Data across AWS regions for better durability.
- **Encryption**: Supports server-side and client-side Encryption for securing data.
- **Access Control**: Fine-grained control over who can

Access or modify data using IAM policies, bucket Policies, and access control lists (ACLs).

Go to the **AWS Management Console** and sing in with your credentials.

In the search bar, type S3 and click on the Amazon S3 service.

Step-by-Step Guide for Uploading a File to Amazon S3:

Step 1: Log in to the AWS Management console

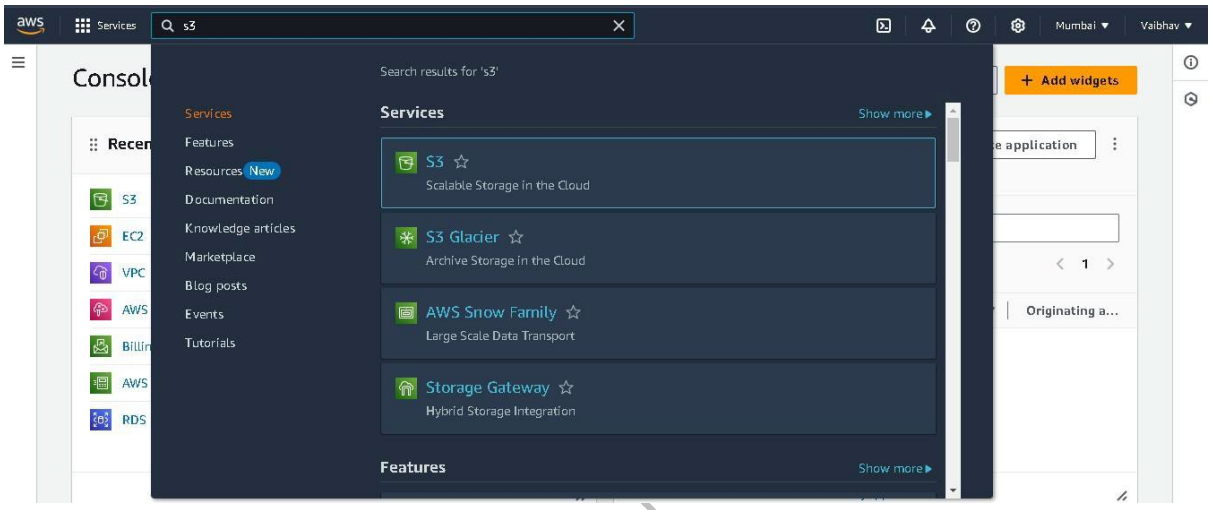

Step 2: Create a New S3 Bucket

- Once in the S3 dashboard, click Create Bucket.
- Enter a Bucket Name (it must be unique globally across all AWS users)

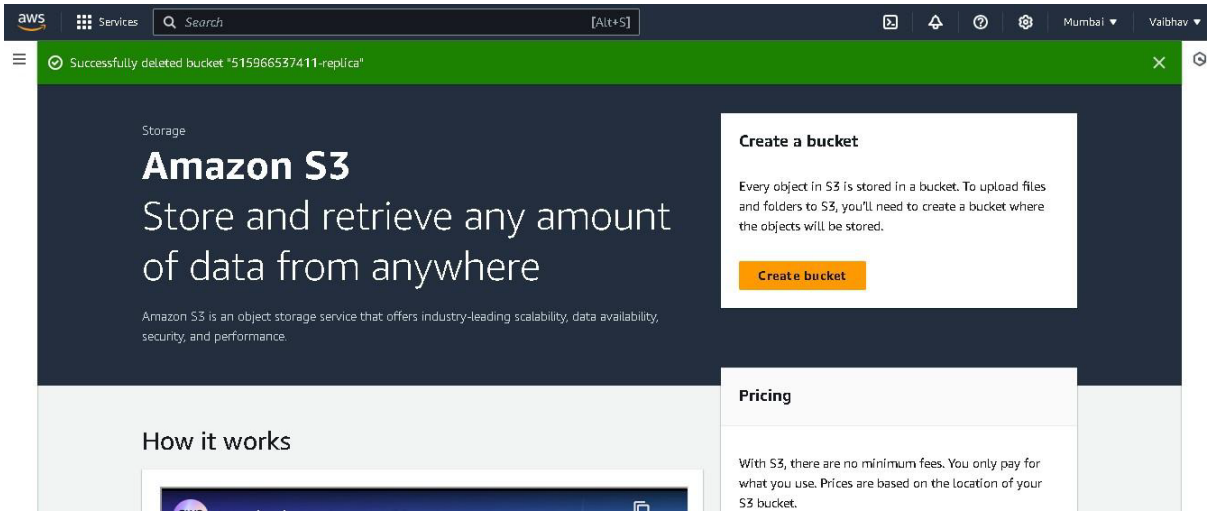

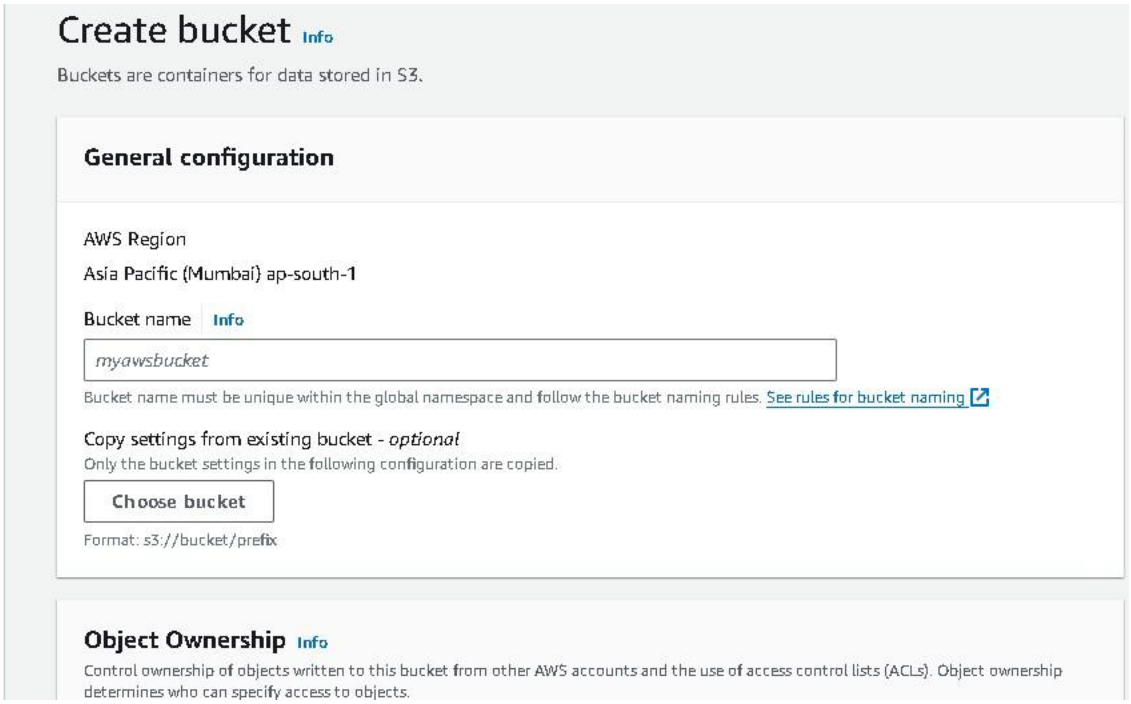

- Select the AWS Region where you want to store your data (choose a region closest to your location for better performance.)
- Configure bucket settings.

Block Public Access: You can keep public access blocked unless you want to make your files accessible to the public.

• **Bucket Versioning**: You can enable or disable versioning. This feature keeps multiple version of an object.

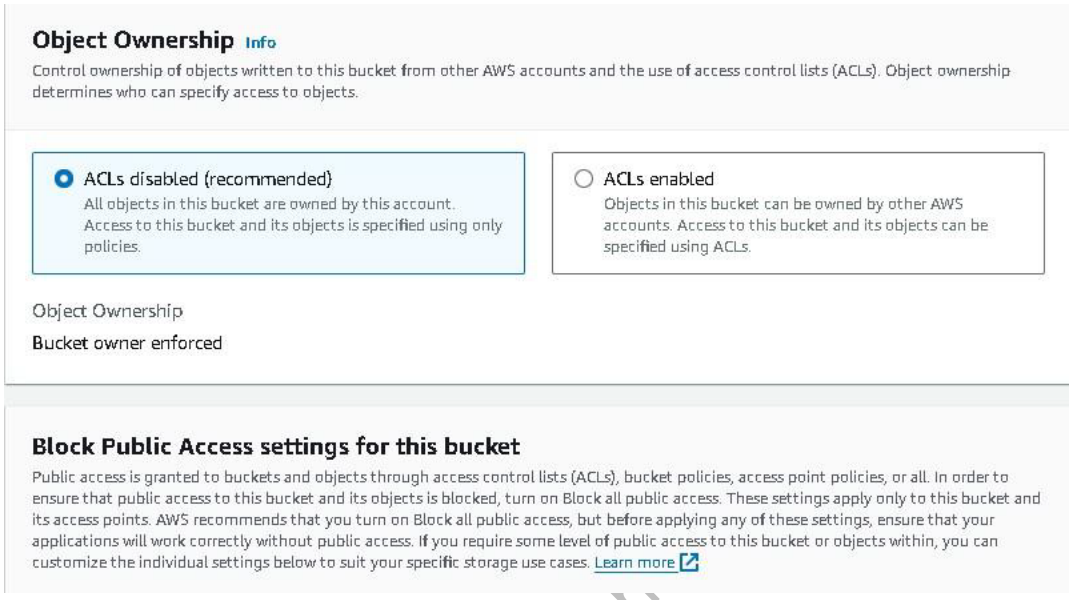

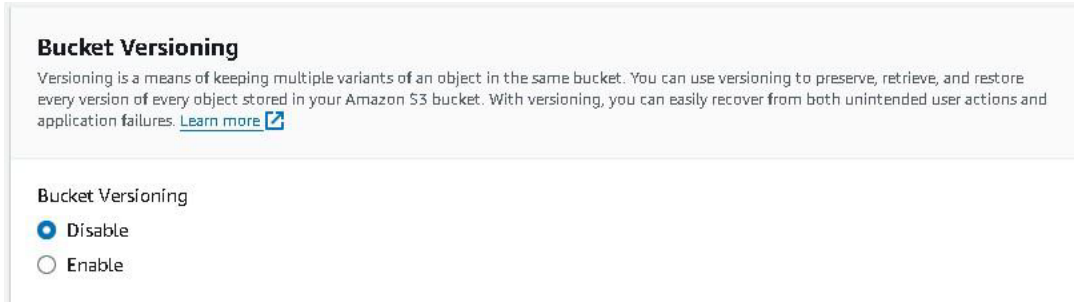

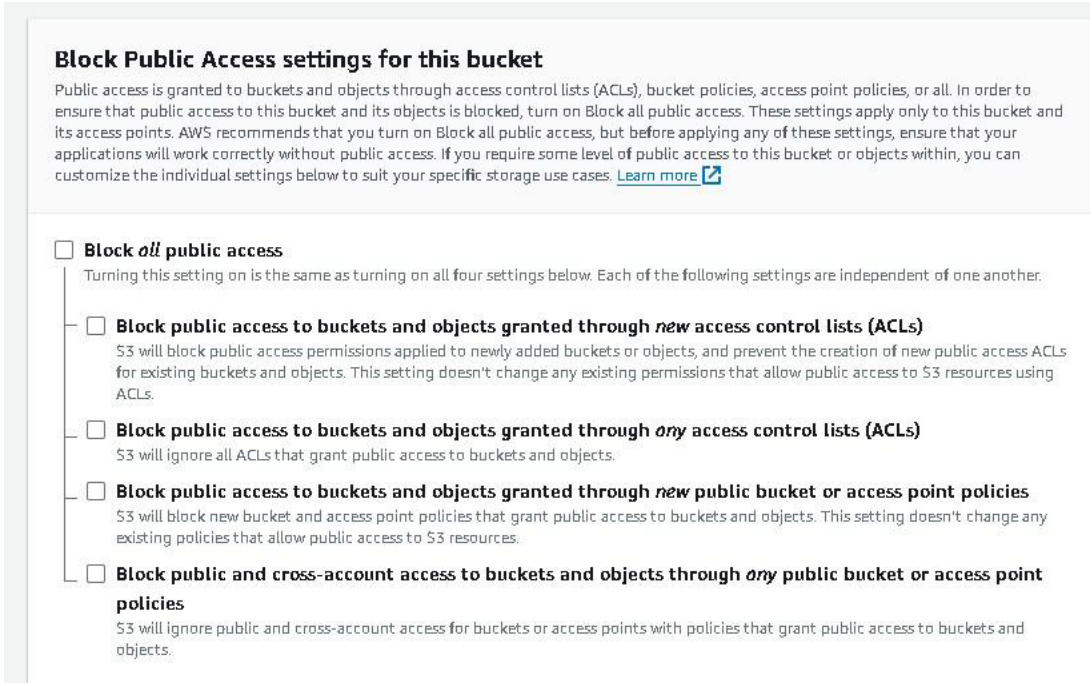

Scroll down and click Create Bucket.

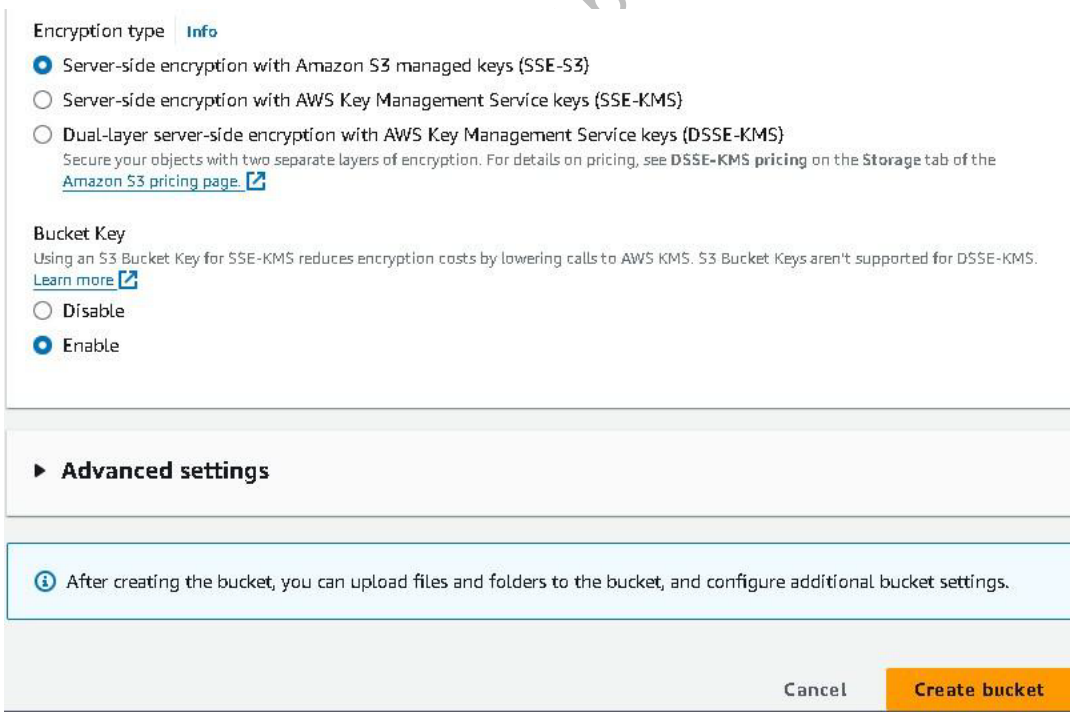

Upload a File to the S3 bucket

- Once the bucket is created, click on your bucket's name to open it.

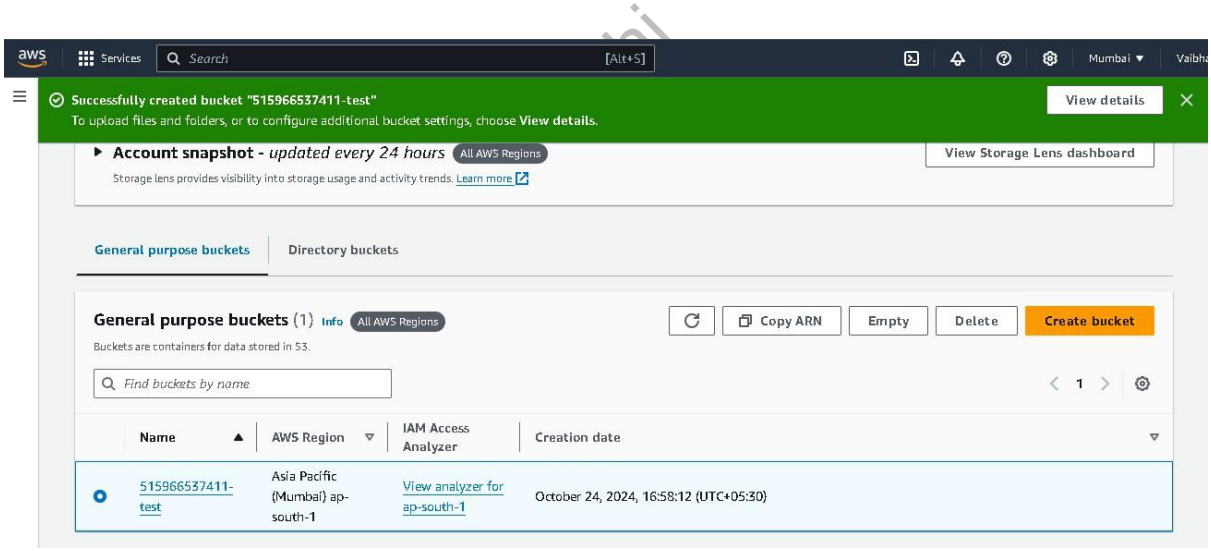

Click upload on the bucket's dashboard.

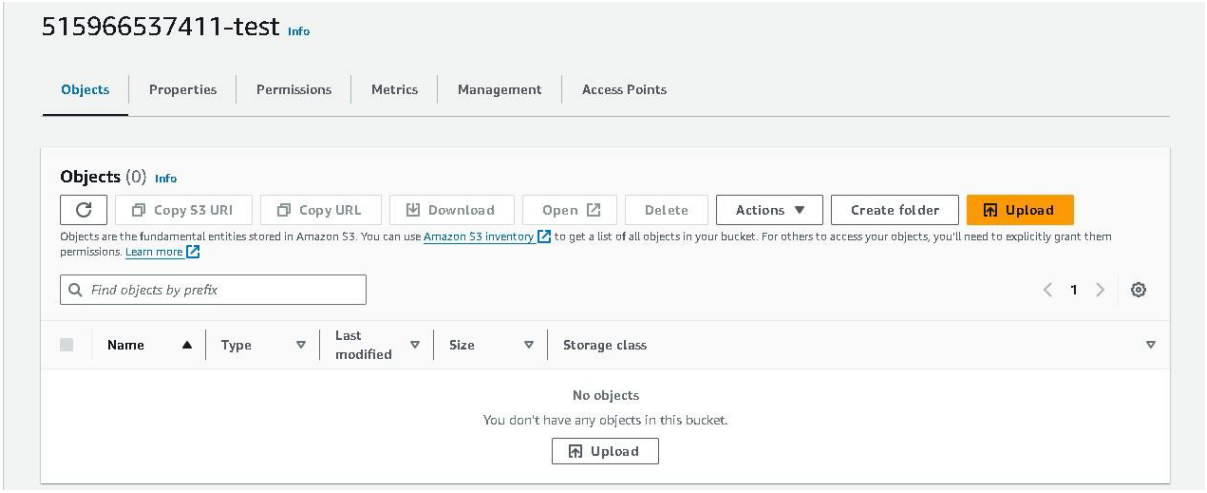

In the Upload window, either drag and drop a file from your local machine or click add files to browse for the file you want to upload.

- After adding the file, you can configure permission:
- Manage public permissions: Decide if the file Should be public or private.

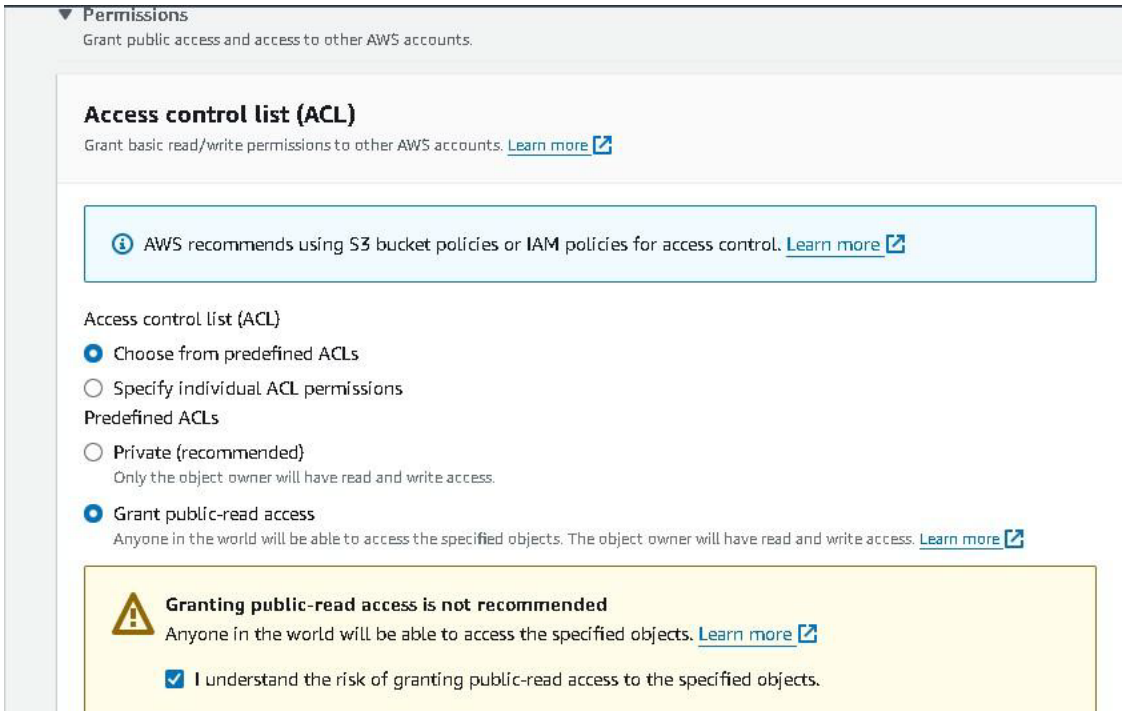

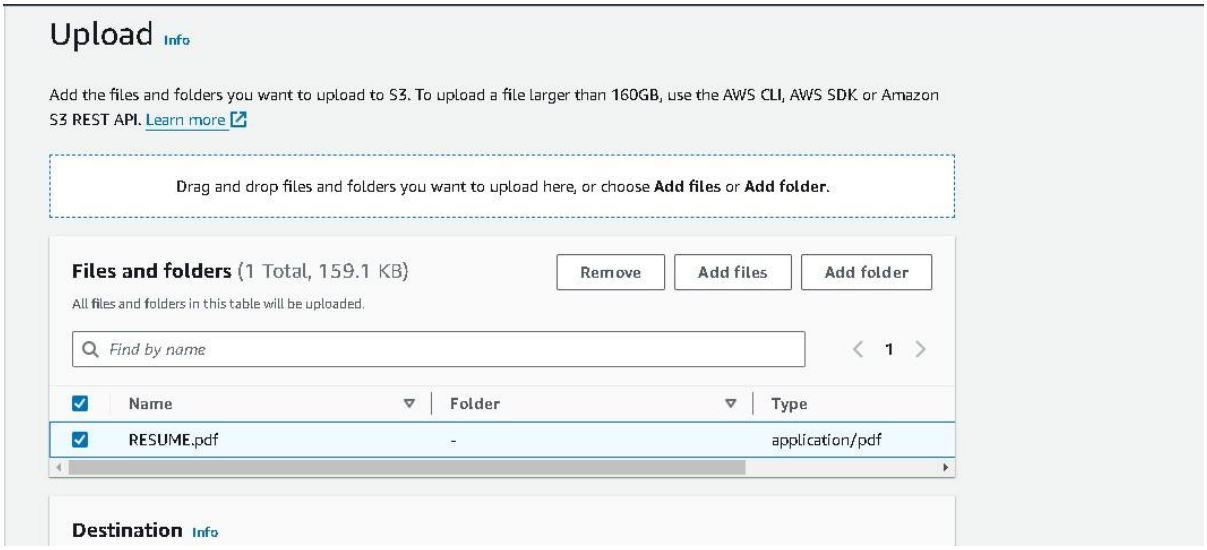

- You can leave storage class settings as default (S3 Standard), or choose other storage classes like Glacier for archival.
- Once you've set the file and permissions, click Upload.

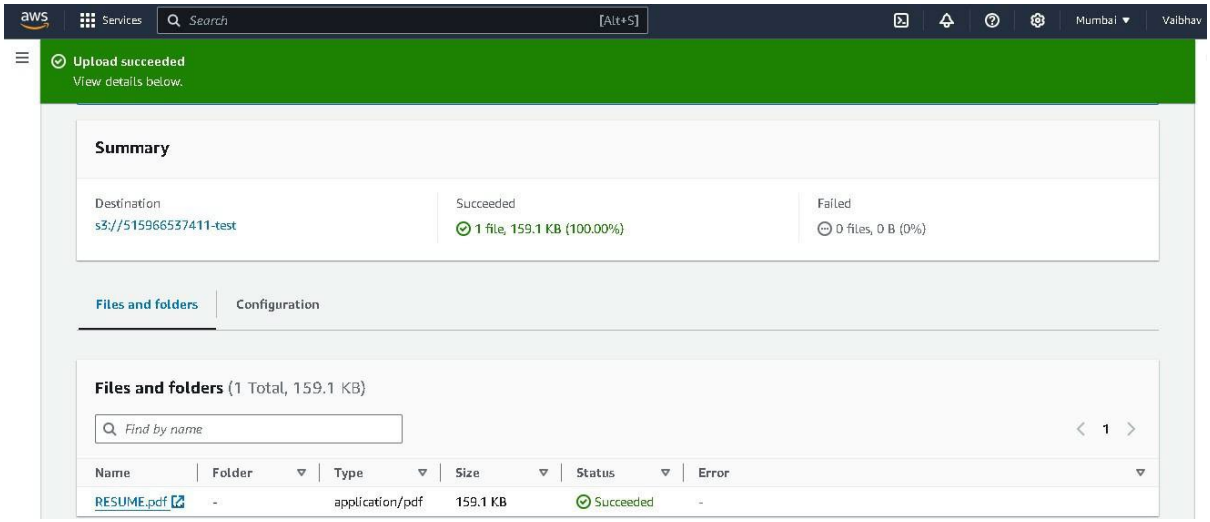

Verify the Upload

Once the file is uploaded, you will see it in the bucket. You can click on the file name to view its properties.

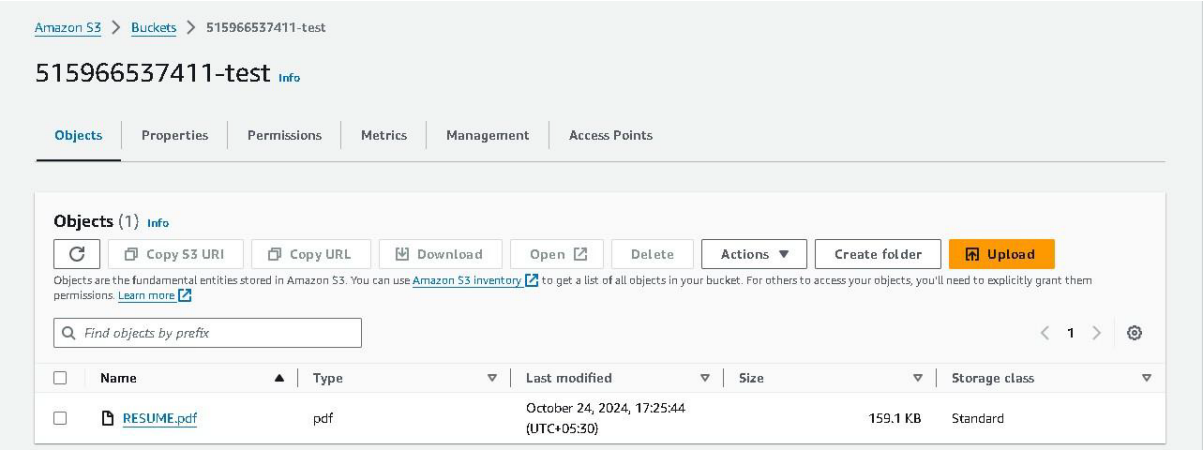

The file has a URL that you can use to access it

(based on permissions). The URL format will look like this:

https://<bucket-name>.s3.<region>.amazo naws.com/<file-name>

- After uploading, you can manage your file's permissions, move it to different storage classes (like Glacier), or delete it.
- If you want to download the file from the bucket, You can either use the AWS CLI or the S3 console.

Additional Point:

Public Access: If you want the file to be publicly accessible, make sure to set the object's permissions accordingly.
Storage Classes: By default, objects are uploaded to the S3 Standard storage class. You can later change it to Glacier or Intelligent-Tiering, depending on your use case.

Reminder: Delete S3 Bucket After Completing Your Work
Once you've completed your tasks and no longer need the S3 bucket, make sure to delete it to avoid unnecessary storage costs. Deleting the bucket will permanently remove all data, so ensure you have saved any important files. Always clean up unused resources after finishing your work.