A reference for the key terms and concepts you should be familiar with when using Cloud Storage.
(/storage/docs/key-terms)

A description of the relationship between Google Cloud projects and Cloud Storage resources, including Cloud Storage service accounts.
(/storage/docs/projects)

An overview and comparison of the storage classes you can choose from when using Cloud Storage.
(/storage/docs/storage-classes)

An explanation of bucket location, including a list of locations currently available and considerations when choosing a location.
(/storage/docs/locations)

The requirements and guidelines to follow when creating buckets and uploading objects.
(/storage/docs/naming)

The requirements for creating buckets with domain names, as well as methods for providing the necessary verification to do so.

(/storage/docs/domain-name-verification)

Objects in Cloud Storage are stored with additional metadata, including custom metadata that you define.

(/storage/docs/metadata)

Pub/Sub Notifications record changes to objects in your Cloud Storage buckets.

(/storage/docs/pubsub-notifications)

The Object Versioning feature of Cloud Storage allows you to retain older versions of your data when you overwrite them with newer versions.

(/storage/docs/object-versioning)

The Object Lifecycle Management feature of Cloud Storage allows you to automatically manage and delete older versions of objects when they meet criteria that you define.

(/storage/docs/lifecycle)

The Requester Pays feature enables costs for accessing buckets and objects to be billed to a requester-specified project instead of the project of the bucket being accessed.

(/storage/docs/requester-pays)

Retention policies prevent objects within your buckets from being modified or deleted for a set period of time. Bucket locking irreversibly prevents your retention policies from being reduced or deleted.

(/storage/docs/bucket-lock)

The Object Transcoding feature of Cloud Storage can automatically decompress files before they are downloaded from your buckets.

(/storage/docs/transcoding)

The composite object feature of Cloud Storage allows you to combine multiple source objects into one new object.

(/storage/docs/composite-objects)

Cloud Audit Logs allows you to generate logs for API operations performed in Cloud Storage.

(/storage/docs/audit-logs)

Learn about how cross-origin resource sharing (CORS) works on Cloud Storage.

(/storage/docs/cross-origin)

An overview page presenting the different methods you can use to control access to your buckets and objects.
(/storage/docs/access-control)

Identity and Access Managament allows you to control who has access to your buckets and objects.
(/storage/docs/access-control/iam)

Uniform bucket-level access prevents permissions from being granted at the object level.
(/storage/docs/bucket-policy-only)

Access control lists allow you to specify user access to individual objects within a bucket.
(/storage/docs/access-control/lists)

Signed URLs allow you to give time-limited access to objects with a URL that may be used by anyone with whom you share the URL.
(/storage/docs/access-control/signed-urls)

Encryption is automatically performed on your data by Cloud Storage. This page also provides an overview of other encryption options you can use to encrypt your object data.
(/storage/docs/encryption)

Encryption is automatically performed on your stored data by Cloud Storage.
(/storage/docs/encryption/default-keys)

Customer-Supplied Encryption Keys allow you to create and manage the keys Cloud Storage uses to encrypt your data.
(/storage/docs/encryption/customer-supplied-keys)

Customer-Managed Encryption Keys allow you to manage encryption keys created by Cloud Key Management Service that Cloud Storage uses to encrypt your data.
(/storage/docs/encryption/customer-managed-keys)

You can encrypt your data before sending it to Cloud Storage.
(/storage/docs/encryption/client-side-keys)

There are several endpoints by which you can access Cloud Storage.
(/storage/docs/request-endpoints)

Hashes and ETags can be used to validate the integrity of transferred data.
(/storage/docs/hashes-etags)

Truncated exponential backoff is a strategy for handling failed requests. This page includes an example of what it looks like when used in Cloud Storage.
(/storage/docs/exponential-backoff)

Resumable uploads are useful when uploading large files to Cloud Storage.
(/storage/docs/resumable-uploads)

Generation numbers allow users to uniquely identify data resources.
(/storage/docs/generations-preconditions)

How consistency is applied in Cloud Storage.
(/storage/docs/consistency)

Learn the best ways to use Cloud Storage's scalability for achieving high request rates.
(/storage/docs/request-rate)

HMAC keys allow you to make signed requests to the XML API.
(/storage/docs/authentication/hmackeys)

Canonical requests form the basis for accessing Cloud Storage resources with signed requests and signed URLs.

(/storage/docs/authentication/canonical-requests)

The Google Cloud Console allows you to interact with Cloud Storage via your browser.

(/storage/docs/cloud-console)

Examples and tips for hosting a static website.

(/storage/docs/static-website)

A summary of the best practices to employ when using Cloud Storage, such as naming guidelines, security considerations, and uploading tips.

(/storage/docs/best-practices)