

Best Practices for Using Cloud Storage

Developing Applications with Google Cloud Platform

CLOUD STORAGE

Version 2.0
Last modified: 2017-09-11



© 2017 Google Inc. All rights reserved. Google and the Google logo are trademarks of Google Inc. All other company and product names may be trademarks of the respective companies with which they are associated.

Performing Operations on Buckets and Objects

Google Cloud Storage Concepts

Resources are entities in Google Cloud Platform including:

- Projects
- Buckets - the basic Cloud Storage container
- Objects - the individual pieces of data that you store in Google Cloud Storage

For more information, see: <https://cloud.google.com/storage/docs/key-terms>

Storage classes

Storage Class	Characteristics	Use Cases	Price (per GB per month)	Name for APIs
Multi-Regional Storage	99.95% availability - Geo-redundant	- Serving website content - Streaming videos - Mobile apps	\$0.026	<code>multi_regional</code>
Regional Storage	99.9% availability - Data stored in a narrow geographic region	Data analytics	\$0.02	<code>regional</code>
Nearline Storage	99.0% availability - Data retrieval costs - Higher per-operation costs 30-day minimum storage duration	- Back-up - Serving long-tail multimedia content	\$0.01	<code>nearline</code>
Coldline Storage	99.0% availability - Data retrieval costs - higher per-operation costs 90-day minimum storage duration	- Disaster recovery - Data archiving	\$0.007	<code>coldline</code>

Google Cloud Storage has four storage classes, with different characteristics, use cases, and prices for your needs.

The Multi-Regional Storage class is designed for 99.95% availability and is geo-redundant. Ideal use cases include serving website content, streaming videos, and storing data for mobile applications.

The Regional Storage class is designed for 99.9% availability and data is stored in a narrow geographic region. Ideal use cases include frequently accessing data in the same region as your Google Cloud DataProc or Google Compute Engine instances that use it, such as for data analytics.

The Nearline Storage class is designed for 99.0% availability, has data retrieval costs and higher per-operation costs, and has a 30-day minimum storage duration. Ideal use cases include back-ups and serving long-tail multimedia content.

The Coldline Storage class is designed for 99.0% availability, has data retrieval costs and higher per-operation costs, and has a 90-day minimum storage duration. Ideal use cases include disaster recovery and data archiving.

For more information, see: <https://cloud.google.com/storage/docs/storage-classes>


Demo: Explore Cloud Storage

1. Create a Cloud Storage bucket in the Console.
2. Upload files to Cloud Storage in the Console.
3. Manipulate Cloud Storage using Cloud Shell.
4. Enable and manage object versioning.

Demo: Exploring Cloud Storage

The following operations are strongly consistent

- Read-after-write
- Read-after-metadata-update
- Read-after-delete
- Bucket listing
- Object listing
- Granting access to resources




Strongly consistent: when you *perform an operation* in Cloud Storage and *receive a success response*, the object is *immediately available* for download and metadata operations.

Strongly consistent operations mean that when you perform an operation in Cloud Storage and you receive a success response, the object is immediately available for download and metadata operations from any location where Google offers service. The following operations in Google Cloud Storage are strongly consistent: read-after-write, read-after-metadata-update, read-after-delete, bucket listing, object listing, and granting access to resources.

For more information, see: <https://cloud.google.com/storage/docs/consistency>

The following operations are eventually consistent

- Revoking access from objects
- Accessing publicly readable cached objects



Eventually consistent: when you perform an operation, it may take some time for the operations to take effect.

The following operations are eventually consistent: revoking access from objects and accessing publicly readable cached objects. Revoking access can take up to a minute for changes to take effect, so operations won't be consistent until the revocation has taken place. Publicly readable cached objects are eventually consistent if the object is in the cache when it is updated or deleted. When the object in the cache is updated or deleted, the operation doesn't take effect until its cache lifetime expires.

For more information, see:

https://cloud.google.com/storage/docs/consistency#eventually_consistent_operations

Use the following request endpoints:

	URIs	HTTP	HTTPS
Typical API Requests	XML: <code>storage.googleapis.com/<bucket>/<object></code> <code><bucket>.storage.googleapis.com/<object></code> JSON: <code>www.googleapis.com/download/storage/v1/b/<bucket>/o/<object-encoded-as-URL-path-segment>?alt=media</code>	✓	✓
CNAME Redirects	Use the following URI in the host name portion of your CNAME record: <code>c.storage.googleapis.com.</code> Example: if you publish <code>travel-maps.example.com</code> CNAME <code>c.storage.googleapis.com</code> , you could access the object at: <code>http://travel-maps.example.com/paris.jpg</code>	✓	
Authenticated Browser Downloads	To download an object using cookie-based authentications: <code>https://storage.cloud.google.com/<bucket>/<object></code>		✓
Content-Based Load Balancing	Create a back end bucket for your Cloud Storage bucket and modify the web-map URL map. To fetch resources: <code>http://[IP_ADDRESS]/static/[OBJECT_NAME]</code> or <code>https://[IP_ADDRESS]/static/[OBJECT_NAME]</code>	✓	✓

Depending on the operation you are performing and your application requirements, you can access Google Cloud Storage through three request endpoints (URIs).

Using the XML API, you can use the two listed URIs, and using the JSON API, you can use the listed URI. These URIs support secure sockets layer (SSL) encryption, so you can use either HTTP or HTTPS. If you authenticate to the Google Cloud Storage API using OAuth 2.0, you should use HTTPS.

A CNAME redirect is a special DNS record that lets you use a URI from your own domain to access a resource (bucket and object) in Google Cloud Storage without revealing the Google Cloud Storage URI. For example, if you published the following, `travel-maps.example.com` CNAME `c.storage.googleapis.com`, you could access an object with the following URI: <http://travel-maps.example.com/paris.jpg>. You can only use CNAME redirects with HTTP.

You can also access resources using authenticated browser downloads, which let you download data through your browser if you are signed in to your Google account and have been granted permission to read the data. Authenticated browser downloads use cookie-based Google account authentication in conjunction with Google account-based ACLs. To download an object using cookie-based authentication, you must use the following URI: <https://storage.cloud.google.com/<bucket>/<object>>. Only HTTPS is supported.

You can modify the content-based load balancing configuration to route requests for

static content to a Cloud Storage bucket. Requests to URI paths that begin with */static* are sent to your storage bucket, and all other requests are sent to your virtual machine instances. You will create a back end bucket for your Cloud Storage bucket, modify the web-map URI map, and fetch resources using either `http://[IP_ADDRESS]/static/[OBJECT_NAME]` or `https://[IP_ADDRESS]/static/[OBJECT_NAME]`.

For more information, see:

Cookie-Based Authentications:

<https://cloud.google.com/storage/docs/authentication.html#cookieauth>

Request Endpoints: <https://cloud.google.com/storage/docs/request-endpoints>

Content-Based Load Balancing:

<https://cloud.google.com/compute/docs/load-balancing/http/adding-a-backend-bucket-to-content-based-load-balancing>

Composite objects and parallel uploads

Combine up to 32 objects into a single new object.

Use cases include:

- Dividing your data and uploading each chunk to a distinct object, composing your final object, and deleting any temporary objects.
- Uploading data to a temporary new object, composing it with the object you want to append it to, and deleting the temporary object.

Compose an object of smaller chunks using gsutil:

```
gsutil compose gs://example-bucket/component-obj-1  
gs://example-bucket/component-obj-2  
gs://example-bucket/composite-object
```

You can compose up to 32 existing objects into a new object without transferring additional object data.

Object composition can be used for uploading an object in parallel: simply divide your data into multiple chunks, upload each chunk to a distinct object in parallel, compose your final object, and delete any temporary objects.

In parallel uploads, you divide an object into multiple pieces, upload the pieces to a temporary location simultaneously, and compose the original object from these temporary pieces.

Some use cases include:

- Fully use your available bandwidth by dividing your data into chunks and uploading them separately.
- Copying many files in parallel with a single command in big data scenarios.

For more information, see:

Composite objects: <https://cloud.google.com/storage/docs/composite-objects>

Working with Objects: <https://cloud.google.com/storage/docs/object-basics>

Design your application to handle network failures with truncated exponential backoff

Truncated exponential backoff:

Is a standard error-handling strategy for network applications.

Periodically retries failed requests with increasing delays between requests.

Should be used for all requests to Google Cloud Storage that return HTTP 5xx and 429 response codes.

```
@retry(wait_exponential_multiplier=1000, wait_exponential_max=10000)
def wait_exponential_1000():
    print "Wait 2^x * 1000 milliseconds between each retry, up to 10 seconds, then 10 seconds afterwards"
```

Truncated exponential backoff is a standard error-handling strategy for network applications in which a client periodically retries a failed request with increasing delays between requests.

Clients should use truncated exponential backoff for all requests to Google Cloud Storage that return HTTP 5xx and 429 response codes, including uploads and downloads of data or metadata.

Understanding how truncated exponential backoff works is important if you are building client applications that use the Google Cloud Storage XML API or JSON API directly, accessing Google Cloud Storage through a client library (some client libraries, such as the Cloud Storage Client Library for Node.js, have built-in exponential backoff), or using the gsutil command line tool (has configurable retry handling).

The Google Cloud Platform Console sends requests to Google Cloud Storage on your behalf and will handle any necessary backoff.

The example shows a backoff implementation for the Python retry library when you retry distributed services and other remote endpoints.

For more information, see: <https://cloud.google.com/storage/docs/exponential-backoff>

Best Practices for Using Cloud Storage

Follow these best practices for naming

Do	Don't
Use globally unique bucket names: us-east1-appdev-bucket01 ✓	Use personally identifiable information (PII): johndoebucket ✗
Use GUIDs or the equivalent if your application needs a lot of buckets: 037763b8-2b55-4887-bbb9-600e2c6c6014 f41b0c0c-9cfc-405d-970e-d4ee46e41279 7f62c421-ce22-4c3f-ad9c-17f342500f62 ✓	Use user IDs, emails , project names/IDs, etc: johndoe ✗ project-gcp-sales ✗
Conform to standard DNS naming conventions: mycompany-unique-bucket01 yourllc-011-eu-central1 ✓	Use IP address notation: 192.168.5.4 ✗
	Use the goog prefix or include any spelling or close misspelling of google: goog_bucket01284 ✗ bucket-for-google-course01 ✗

Google Cloud Storage bucket names are global and publicly visible and must be unique across the entire Google Cloud Storage. If your applications require many buckets, use GUIDs or the equivalent for bucket names. Your application should have retry logic in place to handle name collisions. Consider keeping a list to cross-reference your buckets.

Avoid using any information in bucket names that can be used to probe for the existence of other resources. Use caution if putting personally identifiable information in object or bucket names, because they appear in URLs. Bucket names should conform to standard DNS naming conventions, because the bucket name can appear in a DNS record as part of a CNAME redirect.

If you use part of your company domain name in a bucket name (forward or reverse), Google will try to verify that you own the domain. Avoid using your company domain name in a bucket unless your DNS admin can verify ownership.

For more information, see:

<https://cloud.google.com/storage/docs/bucket-naming#requirements>

Follow these best practices for Cloud Storage traffic

- Consider:
 - Operations per second
 - Bandwidth
 - Cache control
- Design your application to minimize spikes in traffic.
- Use exponential backoff if you get an error.
- For request rates > **1000 write requests/second** or **5000 read requests/second**:
 - Start with a request rate below or near the threshold.
 - Double the request rate no faster than every 20 minutes.

Consider operations per second, bandwidth (how much data will be sent over what time frame), and cache control when designing your application with Cloud Storage. If you specify the Cache-Control metadata on objects, read latency will be lowered on hot or frequently accessed objects. Use exponential backoff if you get an error. For more information on setting object metadata, see:

<https://cloud.google.com/storage/docs/viewing-editing-metadata#edit>

Design your application so that it minimizes spikes in traffic; spread updates out throughout the day. Google Storage has no upper bound on request rate, but for best performance when scaling to high request rates, follow the request rate and access distribution guidelines: If your request rate is less than 1000 write requests per second or 5000 read requests per second, then no ramp-up is needed. If your request rate is expected to go over these thresholds, you should start with a request rate below or near the thresholds and then double the request rate no faster than every 20 minutes. For more information, see:

<https://cloud.google.com/storage/docs/request-rate>

Consider location and availability of your data for storage options

Store data in a region closest to your application's users.



For analytics workloads, store data in regional buckets.



Consider compliance requirements when choosing data location.



FIEC



Consider the location and required availability of your data when choosing your storage options.

Store data in a region closest to your application's users, and consider region-specific compliance requirements when choosing data location. For analytics workloads, store your data in regional buckets to reduce network charges and for better performance (compared to multi-regional).

For more information, see: <https://cloud.google.com/security/compliance>

More location and availability considerations

- Multi-Regional and Regional Storage:
 - Provide the best availability.
 - Are good options for data served at a high rate with high availability.
- Nearline and Coldline Storage are good options for:
 - Infrequently accessed data.
 - Data that tolerates slightly lower availability.

Multi-Regional Storage and Regional Storage provide the best availability, with the trade-off of a higher price, and are good options for data that is served at a high rate with high availability.

Nearline Storage and Coldline Storage are good options for infrequently accessed data and for data that tolerates slightly lower availability.

Secure your buckets using the following options

Use Identity and Access Management (IAM) permissions to grant:



- Access to buckets
- Bulk access to a bucket's objects

Use Access Control Lists (ACLs) to grant:

- Read or write access to users for individual buckets or objects
- Access when fine-grained control over individual objects is required

Signed URLs (query string authentication):

- Provide time-limited read or write access to an object through a generated URL
- Can be created using gsutil or programmatically.

Use Signed Policy Documents to:

- Specify what can be uploaded to a bucket.
- Control size, content type, and other upload characteristics.

Firebase Security Rules provide:



- Granular, attribute-based access control to mobile and web apps using the Firebase SDKs for Cloud Storage

You can control access to your Cloud Storage buckets and objects using these options.

You can use Identity and access Management (IAM) permissions to grant access to buckets and to provide bulk access to a bucket's objects. IAM permissions do not give you fine-grained control over individual objects. For more information, see: <https://cloud.google.com/storage/docs/access-control/using-iam-permissions>

You can use Access Control Lists (ACLs) to grant read or write access to users for individual buckets or objects. It is recommended that you only use ACLs when you need fine-grained control over individual objects. For more information, see: <https://cloud.google.com/storage/docs/access-control/create-manage-lists>

Use signed URLs (query string authentication) to provide time-limited read or write access to an object through a URL you generate. The shared URL provides access to anyone it's shared with for the duration specified. You can create signed URLs using gsutil or programmatically with your application. For more information, see:

Create a signed URL using gsutil:

<https://cloud.google.com/storage/docs/access-control/create-signed-urls-gsutil>

Create a signed URL programmatically:

<https://cloud.google.com/storage/docs/access-control/create-signed-urls-program>

Signed policy documents allow you to specify what can be uploaded to a bucket. They allow greater control over size, content type, and other upload characteristics

than using signed URLs. They are for website owners to allow visitors to upload files to Google Cloud Storage. Signed policy documents only work with form posts. For more information on signed policy documents, see:

<https://cloud.google.com/storage/docs/xml-api/post-object#policydocument>

Firebase security rules provide granular, attribute-based access control to mobile and web applications using the Firebase SDKs for Cloud Storage. For more information, see: <https://firebase.google.com/docs/storage/security>

Consider the following when using gsutil for Cloud Storage

- `gsutil -D` will include OAuth2 refresh and access tokens in the output.
- `gsutil --trace-token` will include OAuth2 tokens and the contents of any files accessed during the trace.
- Customer-supplied encryption key information in `.boto` config is security-sensitive.
- In a production environment, use a service account for `gsutil`.

If you run `gsutil -D` (to generate debugging output) it will include OAuth2 refresh and access tokens in the output. Make sure to redact this information before sending this debug output to anyone during troubleshooting/tech support interactions.

If you run `gsutil --trace-token` (to send a trace directly to Google), sensitive information like OAuth2 tokens and the contents of any files accessed during the trace may be included in the content of the trace.

Customer-supplied encryption key information in the `.boto` configuration is security sensitive. The proxy configuration information in the `.boto` configuration is security-sensitive, especially if your proxy setup requires user and password information. Even if your proxy setup doesn't require user and password information, the host and port number for your proxy is often considered security-sensitive. Protect access to your `.boto` configuration file.

If you are using `gsutil` from a production environment, use service account credentials instead of individual user account credentials. These credentials were designed for such use and protect you from losing access when an employee leaves your company.

For more information, see:

<https://cloud.google.com/storage/docs/gsutil/addlhelp/SecurityandPrivacyConsiderations#recommended-user-precautions>

Consider these additional security best practices

- Use TLS (HTTPS) to transport data.
- Use an HTTPS library that validates server certificates.
- Revoke authentication credentials to applications that no longer need access to data.
- Securely store credentials.
- Use groups instead of large numbers of users.
- Bucket and object ACLs are independent of each other.
- Avoid making buckets:
 - Publicly readable
 - Publicly writable

Always use TLS (HTTPS) to transport your data when you can. This ensures that both your credentials and your data are protected as you transport data over the network.

Make sure that you use an HTTPS library that validates server certificates. A lack of server certificate validation makes your application vulnerable to man-in-the-middle attacks or other attacks. Be aware that HTTPS libraries shipped with certain commonly used implementation languages do not, by default, verify server certificates.

When applications no longer need access to your data, you should revoke their authentication credentials. To do this for Google services and APIs, sign in to your Google Account and click Authorizing applications and sites. On the next page, you can revoke access for applications by clicking Revoke Access next to the application.

Make sure that you securely store your credentials. This can be done differently depending on your environment and where you store your credentials.

Use groups in preference to explicitly listing large numbers of users. Not only does it scale better, it also provides a very efficient way to update the access control for a large number of objects all at once. It's cheaper because you don't need to make a request per-object to change the ACLs.

Before adding objects to a bucket, check that the default object ACLs are set to your requirements first. This could save you a lot of time updating ACLs for individual

objects.

Bucket and object ACLs are independent of each other, which means that the ACLs on a bucket do not affect the ACLs on objects inside that bucket. It is possible for a user without permissions for a bucket to have permissions for an object inside the bucket.

The Google Cloud Storage access control system includes the ability to specify that objects are publicly readable. After a bucket has been made publicly readable, data on the internet can be copied to many places. It's effectively impossible to regain read control over an object written with this permission. The Google Cloud Storage access control system also includes the ability to specify that buckets are publicly writable. Although configuring a bucket this way can be convenient for various purposes, we recommend against using this permission: it can be abused for distributing illegal content, viruses, and other malware, and the bucket owner is legally and financially responsible for the content stored in their buckets.

Consider these best practices for uploading data

- If using XMLHttpRequests:
 - Don't close and re-open the connection
 - Set reasonably long timeouts for upload traffic
- Make the request to create the resumable upload URL from the same region as the bucket and upload location.
- Avoid breaking transfers into smaller chunks
- Avoid uploading content that has both:
 - `content-encoding gzip`
 - `content-type` that is compressed

If you use XMLHttpRequest (XHR) callbacks to get progress updates and detect that progress has stalled, do not close and re-open the connection. Doing so creates a bad positive feedback loop during times of network congestion. When the network is congested, XHR callbacks can get backlogged behind the acknowledgement (ACK/NACK) activity from the upload stream, and closing and reopening the connection when this happens uses more network capacity at exactly the time when you can least afford it.

For upload traffic, set reasonably long timeouts. For a good end-user experience, you can set a client-side timer that updates the client status window with a message (e.g., "network congestion") when your application hasn't received an XHR callback for a long time. Don't just close the connection and try again when this happens.

If you use Google Compute Engine instances with processes that POST to Cloud Storage to initiate a resumable upload, you should use Compute Engine instances in the same locations as your Cloud Storage buckets. You can then use a geo IP service to pick the Compute Engine region to which you route customer requests, which will help keep traffic localized to a geo-region. For resumable uploads, the resumable session should stay in the region in which it was created. Doing so reduces cross-region traffic that arises when reading and writing the session state, which improves resumable upload performance.

Avoid breaking a transfer into smaller chunks if possible, and instead upload the entire content in a single chunk. Avoiding chunking removes fixed latency costs,

improves throughput, and reduces QPS against Google Cloud Storage.

If possible, avoid uploading content that has both content-encoding: gzip and a content-type that is compressed, because this may lead to unexpected behavior. You can also use decompressive transcoding (automatically decompressing a file for the requestor) by storing the file in gzip format in Cloud Storage and setting the associated metadata to Content-Encoding: gzip. For more information, see: https://cloud.google.com/storage/docs/transcoding#decompressive_transcoding

Validate your data

Data can be corrupted during upload or download by:

- Noisy network links
- Memory errors on:
 - Client computer
 - Server computers
 - Routers along the path
- Software bugs

Validate data transferred to/from bucket using:

- CRC32c Hash
 - Is available for all cloud storage objects.
 - Can be computed using these libraries:
 - Boost for C++
 - crcmod for Python
 - digest-crc for Ruby
 - gsutil automatically performs integrity checks on all uploads and downloads.
- MD5 Hash
 - Is supported for non-composite objects.
 - Cannot be used for partial downloads.

Data can be corrupted while it is uploaded to or downloaded from the cloud by noisy network links, memory errors on client or server computers or routers along the path, and software bugs. Validate the data you transfer to/from buckets using either CRC32c or MD5 checksums.

Google Cloud Storage supports server-side validation for uploads, but client-side validation is also a common approach. If your application has already computed the object's MD5 or CRC32c before starting the upload, you can supply it with the upload request, and Google Cloud Storage will create the object only if the hash you provided matches the value Google calculated. Alternatively, users can perform client-side validation by issuing a request for the new object's metadata, comparing the reported hash value, and deleting the object in case of a mismatch.

All Google Cloud Storage objects have a CRC32c hash. Libraries for computing CRC32c include Boost for C++, crcmod for Python, and digest-crc for Ruby. Java users can find an implementation of the algorithm in the [GoogleCloudPlatform crc32 Java project](#). Gsutil automatically performs integrity checks on all uploads and downloads. Additionally, you can use the "gsutil hash" command to calculate a CRC for any local file.

MD5 hashes are supported for non-composite objects. The MD5 hash only applies to a complete object so it cannot be used to integrity check partial downloads caused by performing a range GET.

For more information, see:

Hashes and e-tags: <https://cloud.google.com/storage/docs/bytes-checksums>

CRC32C and Installing crcmod:

<https://cloud.google.com/storage/docs/gsutil/addlhelp/CRC32CandInstallingcrcmod>

You can host static websites.

You can allow scripts hosted on other websites to access static resources stored in a Google Cloud Storage bucket.

You can also allow scripts hosted in Google Cloud Storage to access static resources hosted on a website external to Cloud Storage.

The Cross-Origin Resource Sharing (CORS) topic describes how to allow scripts hosted on other websites to access static resources stored in a Google Cloud Storage bucket.

You can also allow scripts hosted in Google Cloud Storage to access static resources hosted on a website external to Cloud Storage. The website is serving CORS headers so that content on `storage.googleapis.com` is allowed access. It is recommended that you dedicate a specific bucket for this data access.

For more information, see: <https://cloud.google.com/storage/docs/cross-origin>