



S3-Compatible Object Storage Use Cases

How to Get **Maximum Value**
from **Object Storage** in the Cloud



Introduction

Modern applications and technologies are creating and consuming more data than ever before. How do you address the need to store, access, and organize this ever-growing amount of data? The answer is increasingly becoming **S3-compatible object storage**.

Object storage is a versatile and cost-effective method for storing static unstructured data that significantly reduces infrastructure costs for developers, businesses, and enterprises. With the ongoing shift to move more infrastructure to the cloud, object storage is becoming more crucial in virtual environments. Use cases of object storage are diverse, ranging from simple multimedia storage to housing large data sets for machine learning.

S3-Compatible Object Storage Use Cases introduces you to object storage, explains the difference between object storage and block storage, and presents examples of how companies such as Coca-Cola, Facebook, and Spotify leverage it. After reading this ebook, you will know if object storage can be a useful tool for your current workload or an upcoming project.

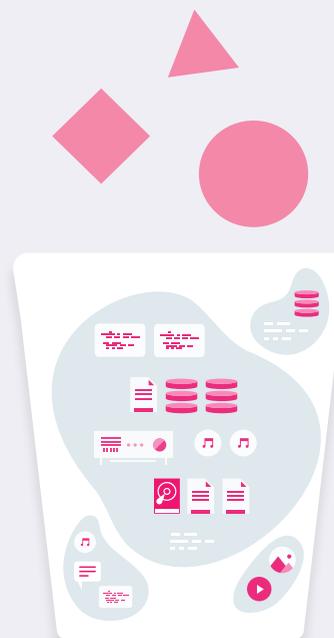
A Brief History of Object Storage

Before the mass adoption of object storage by cloud infrastructure providers, Network Attached Storage (NAS) and block storage were the primary options for storing unstructured data. Both methods succeeded as resilient, scalable solutions. However, as collection and storage rates of unstructured data evolved, users craved new ways to find and consume their data, scale well past terabytes in size, and access data globally.

Researchers at Carnegie Mellon University and the University of California, Berkeley, developed object storage in the late 1990s.

Their goal was to find a data storage solution easier to use and, at the same time, could optimize resources at a larger scale than traditional file and block storage systems.

The researchers separated read and write operations from namespace operations by abstracting many of the lowest layers of storage architecture. In the process, they optimized the performance and scale of both namespace and read and write operations. The result of this combined effort is a system of scalable data easily and directly queried, accessed, and replicated across many physical devices and locations.



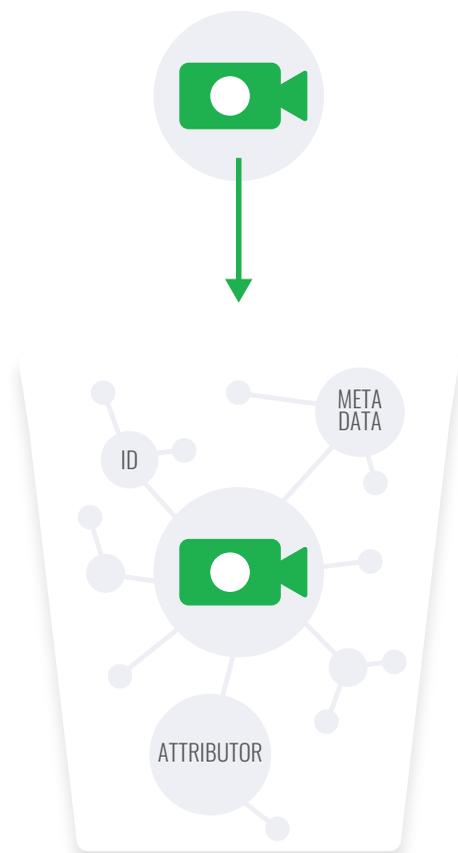
What is Object Storage?

Object storage is a solution that stores data as units, called objects. The objects are housed in containers called buckets.

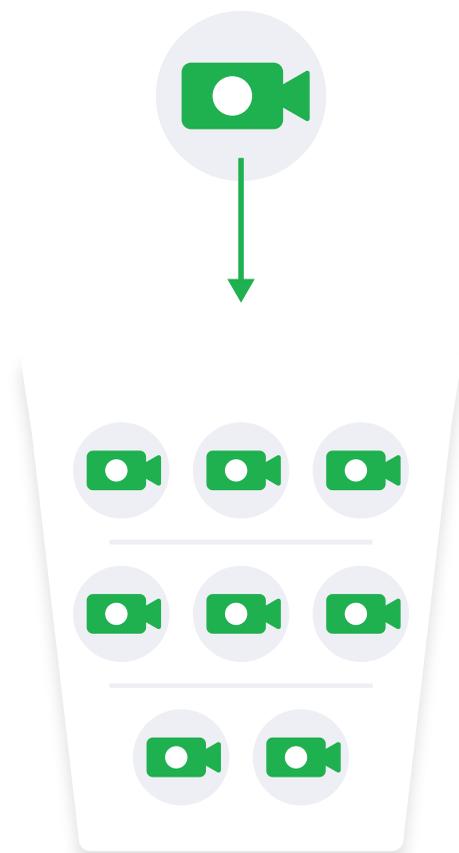
Each object consists of three components: data, metadata, and a global identifier. The global identifier allows you to access an object without knowing its physical location. The rich, configurable metadata can be used to describe any number of arbitrary properties about the object. Each object has its own URL, so accessing the data is often as simple as issuing an HTTP request, either by visiting the object in a browser, retrieving it through the command line, or using service-oriented architecture.

When it comes to object storage, data is stored in its original file type. This speeds up retrieval time because the data doesn't need to be pieced together from smaller units before returning it to the client. The metadata attached to each object provides a rich description of the data, potentially including when it was created, by whom, and how it should be used.

Object Storage



Block Storage



How Object Storage is Different than Block Storage

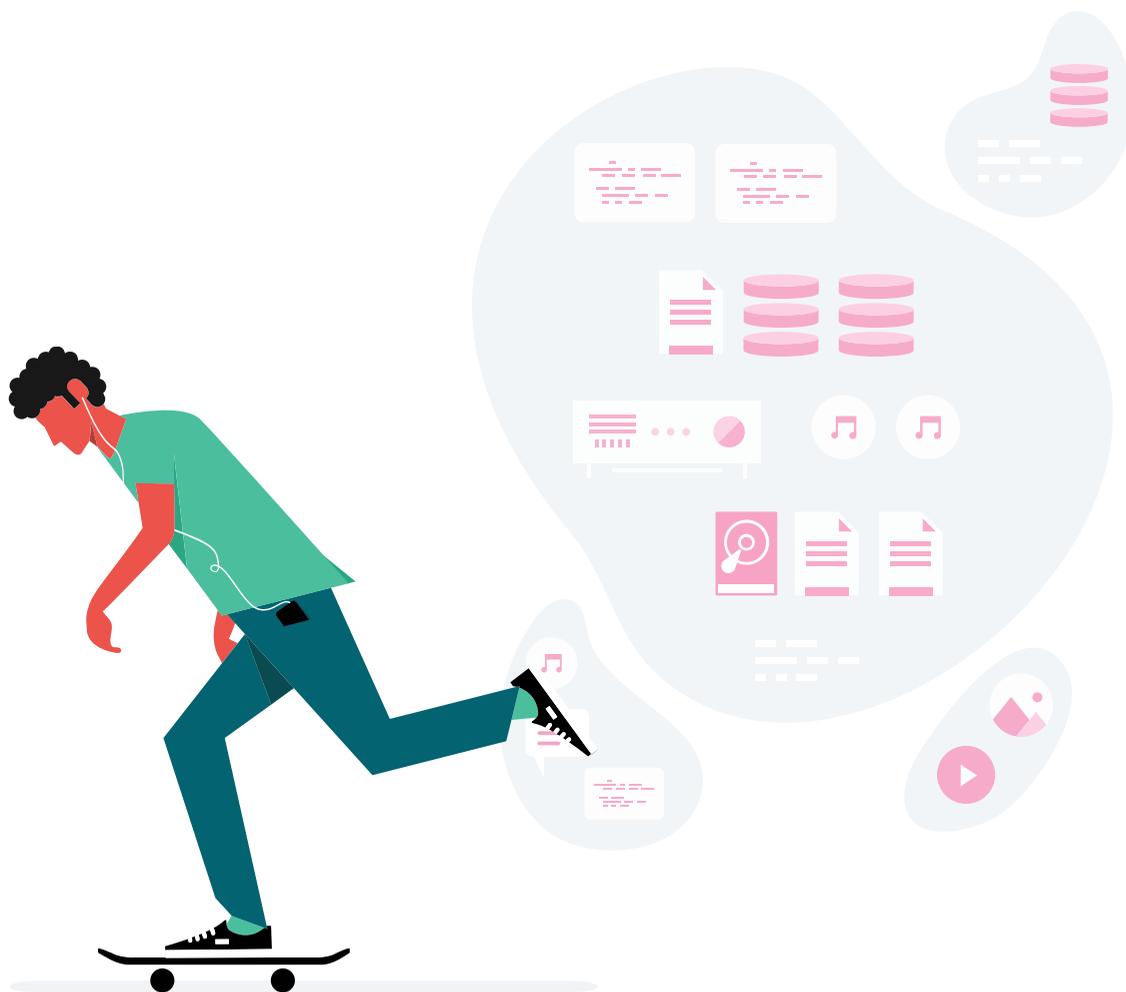
Block storage works by splitting files into small blocks of data. Minimal file metadata is stored alongside these blocks and, in general, descriptive metadata must be stored in a separate file or database.

Block storage stores the location of the blocks in a table. Any time a file is retrieved, the table is queried for the location of the blocks. This is a different table than the one that holds the metadata, which adds extra steps and time to retrieve essential file information.

In order to use a block storage volume, it must be attached to a host server, where it acts as the server's hard drive. In this way, block storage adds additional storage volumes to your cloud infrastructure.

In contrast, object storage does not require a server and instead grants access to the object data via the internet.

While block storage excels with large, frequently updated databases and any other file stores that are consistently changed or updated, object storage is ideal for static data that is rarely altered.



S3-Compatible Object Storage Use Cases



Images, Video, Audio, & Documents

Multimedia assets like images, videos, audio files, and documents are ideally suited for object storage. **Spotify** uses object storage for the millions of songs on their application, while **Facebook** uses object storage for storing photos.

Multimedia files don't change frequently, so there's no need for the write speeds of block storage. And, since object storage gives each file its own URL, it's simple to stream or embed multimedia assets in another program or website without a server. If you're looking for a media server, block storage might be a better fit. But, if you're creating or managing a streaming app, object storage provides a more affordable and easier-to-use solution.



Unstructured Data

The need to store unstructured data was an original catalyst for the creation of object storage. Unstructured data is anything that does not fit into a traditional database, including log files, PDFs, receipts, maps, and multimedia files. What this data has in common—and what makes it unstructured—is that it lacks the internal relational information needed to fall neatly into a database column.

Object storage gives you the ability to configure custom metadata for each piece of unstructured data, making it easy to extrapolate useful information from each object and retrieve objects with similar metadata. Using other storage methods for unstructured data usually requires you to set up an additional database to separately organize the metadata.

Your unstructured data also gains an element of security with object storage. Both block storage and object storage are inherently private, but with one main difference: block storage requires setting up a server to make data public on the internet. With object storage, however, all new objects in a bucket are set to private, and are inaccessible by normal HTTP requests. To make objects public, you can use access controls—no server attached.



Big Data

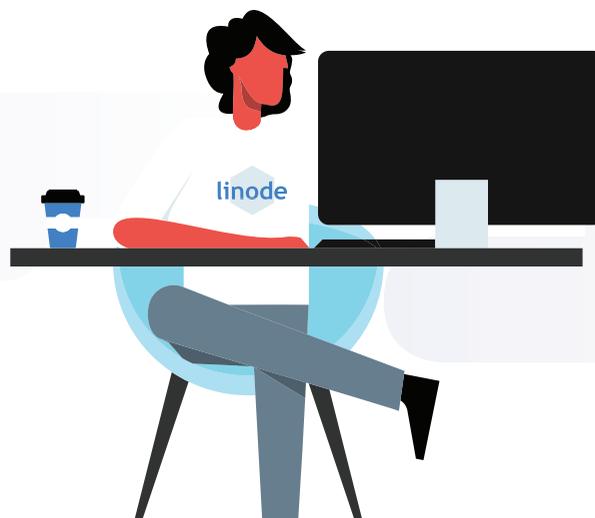
Big data describes data sets that are so large and diverse that it takes specialized tooling to analyze them. Big data falls into the category of unstructured data because it does not fit neatly into databases. Their volume and disparity make big data the ideal candidate for object storage. Neural networks and other machine learning processes create petabytes of data toward training accurate models.

Object storage provides an intuitive system of collating big data from these processes. **Coca-Cola** relies on cloud object storage for training images of mobile proof-of-purchase codes in its rewards program. Not only does object storage add metadata to big data workloads, but its scalability and cost-efficiency make big data a prime use case for object architecture.



Website Files

Object storage can help optimize your website or web application by hosting static assets like images, downloads, or user data in HTTP-accessible buckets. Whether building a dynamic website from scratch, using a content management system like WordPress, or an ecommerce platform like Magento, traditional file or block storage can carry a higher cost for storing website files. Object storage offers an alternative, cost-effective method to store and organize portions of your site or app's assets.





Artifact Storage

As an increasing amount of development and an application's life cycle become automated, more and more artifacts—like bulk collections of logs and version files—are generated from those processes. Discarding artifacts is not a best practice, as those residuals can often be beneficial later in the testing cycle. Developers and stakeholders can use object storage to save artifacts and subsequently share them in the production cycle by distributing their unique URLs.



Cold Storage

Cold storage is the collection of inactive data or data needed at irregular intervals, including backups, regulatory files, or disaster recovery stores. Object storage is a cost-effective and globally available solution for cold storage data.

Significantly more affordable than block storage, object storage also outpaces tape storage for read times and accessibility. Tape storage is a traditional physical storage method of collecting archival data on magnetic tape drives. While tape storage is still frequently used for cold storage, its read times are longer than object storage, and special considerations are necessary when transferring tape drives, like shipping tape safely across long distances.

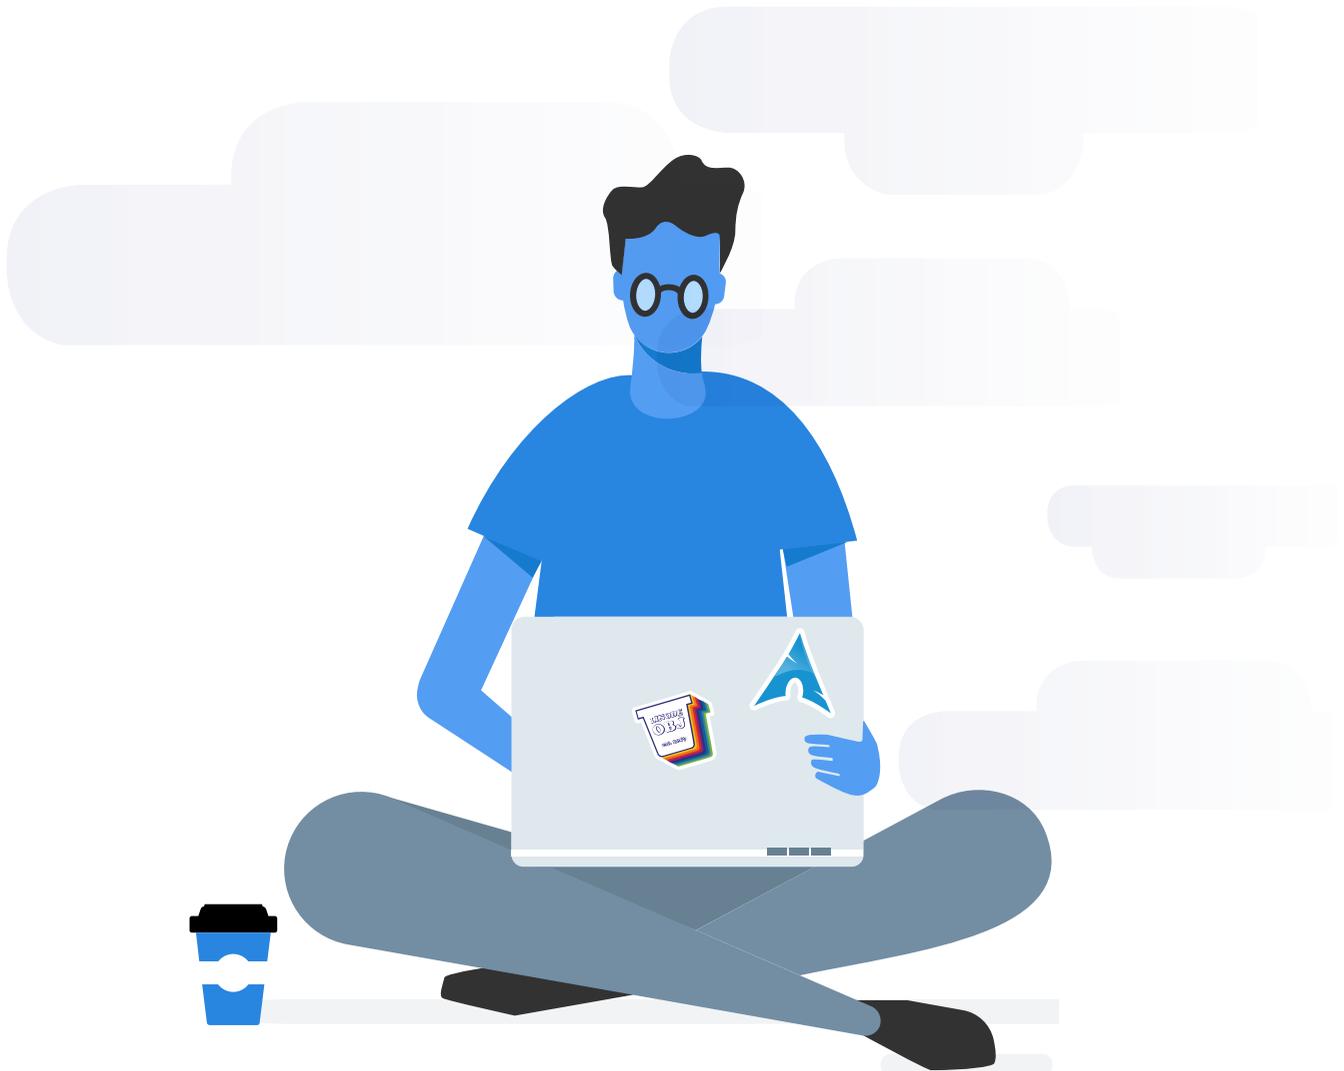




Backups

Most databases offer some method to backup data and prevent critical losses. These backups can be uploaded and stored in object storage. Object storage creates redundancy through object versioning and triple replication. A database dump uploaded to object storage can be scheduled automatically through the command line and ensures resilient file stores.

Buckets can be versioned, so you never lose access to older backups, and objects get replicated three times to ensure they are always available. Object storage's versatility means it can be used as a complementary backup plan for your data if you already rely on other storage tools.



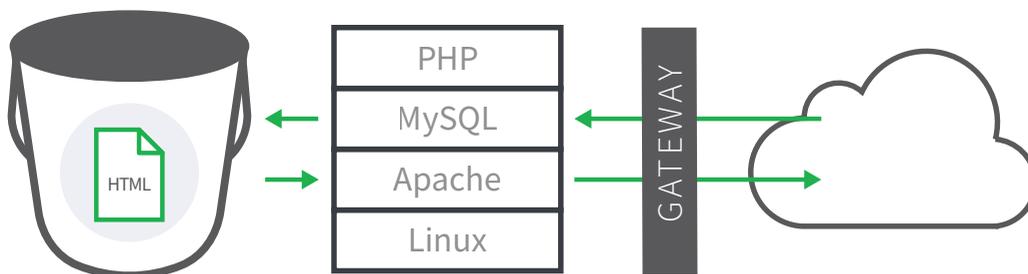


Static Website Hosting

Object storage makes hosting a static website straightforward, with minimal maintenance required. Object storage provides a gateway to objects through a REST API, which means a unique URL over HTTP is available for every object. Traditionally, for static sites, static HTML files would be served by a web server (like NGINX or Apache) running on a virtual machine. This process is unnecessary with object storage because each object has a unique URL. Once you build your static site, you can make it publicly available over the internet by uploading files to an object storage bucket.

Using object storage to host your static site files gives you more time to focus on developing your website instead of routine server maintenance tasks like software upgrades, web server configuration, and security upkeep.

Hosting a Static Site with Object Storage



Benefits and Limitations of Object Storage

Object storage offers clear advantages over other data storage architecture, but it's not right for every project. Developers, small businesses, and enterprises often get drawn to object storage for its scalability, cost efficiency, and encryption, along with other benefits like object versioning and fine-grained access control.

BENEFITS

Object storage is highly available and infinitely scalable because data objects get stored in flat namespaces, which means there is no limit to the amount of stored data.

Ease of access is at the heart of object storage's functionality. Unlike data stored in blocks, data as objects are provisioned with rich metadata descriptions, making for clear organization.

Each object gets a global identifier, so you can retrieve and control object storage data without requiring any infrastructure.

Another benefit of object storage is versioning. Object versioning allows you to store multiple versions of a data object in the same bucket, and access past versions/buckets using a version number.

Versioning helps prevent accidental deletion or overwriting and mitigates the risk of archiving objects since you have quick access to all previous object versions.

In other words, object storage makes mistakes less catastrophic. With a version-enabled bucket, updating an object doesn't overwrite the existing data. Instead, it creates a copy with a new version number. Again, you can retrieve any version of an object by specifying the version number as part of your request.

However, this leads us to one pitfall of this benefit: cost. Object versioning increases the amount of data in your bucket, so your storage costs could increase if your provider offers this feature.

LIMITATIONS

Unlike block storage, stored objects cannot be modified at the block level and must be rewritten in their entirety every time there's a change.

This limitation makes object storage a poor choice for any scenario with many successive read/write operations, such as the needs of databases or transactional data.

Additionally object storage traffic runs over HTTP, so it doesn't benefit from the input/output speeds of a mounted block storage volume. As a rule of thumb, object storage is best in cases where files do not need to be updated frequently.

OUR TAKE

Storage is one of the primary components of modern cloud infrastructure. Object storage is discovering a renaissance of relevance with the need to access and organize large volumes of unstructured data. Object storage, especially, has proven a crucial tool for businesses of all sizes and developers who require a practical and resilient data storage solution.

While block storage remains an important tool for frequently updated files and databases, and tape storage is a cost-effective method for appropriate cold storage purposes, object storage is poised to become the standard choice for unstructured data.

Companies like Linode are making it easier for developers and the enterprises they work with to harness the power of the cloud. Linode Object Storage is a globally-available, S3-compatible method for storing and accessing data that does not require the use of a Linode. To get started setting up object storage, read our guide on [How to Use Linode Object Storage](#) for detailed instructions on creating buckets and uploading objects.

Read our [Host a Static Site using Linode Object Storage](#) to host your own site on object storage.

About Linode

Our mission is to accelerate innovation by making cloud computing simple, affordable, and accessible to all.

Linode accelerates innovation by making cloud computing simple, accessible, and affordable to all. Founded in 2003, Linode helped pioneer the cloud computing industry and is today the largest independent open cloud provider in the world. Headquartered in Philadelphia's Old City, the company empowers more than a million developers, startups, and businesses across its global network of 11 data centers.





The World's Largest Independent Open Cloud

linode.com | Support: 855-4-LINODE | Sales: 844-869-6072
249 Arch St., Philadelphia, PA 19106 Philadelphia, PA 19106