Object Storage vs File Storage

Object storage has only been around since the mid-90s. As the relatively new kid to the storage block, there can be some confusion as to how it differs from more traditional storage systems, such as block or file storage.

A Primer on File Storage

File storage has been around for considerably longer than object storage, and is something most people are familiar with. You name your files/data, place them in folders, and can nest them under more folders to form a set path. In this way, files are organized into a hierarchy, with directories and sub-directories. Each file also has a limited set of metadata associated with it, such as the file name, the date it was created, and the date it was
last modified.

This system works (and works well) when you're pulling smaller files, but does experience some issues when retrieving large amounts of data. However, once you start growing, it becomes much harder to find information, especially if there isn't a standard file-naming convention (that everybody adheres to). At the same time, because the files have unique addresses, there is a finite number of files you can store.

When it comes to sharing files on a file storage system, users typically turn to network-attached storage (NAS). This is great for sharing locally, but may result in issues when trying to share files over a wide area network (WAN).

So What is Object Storage?

Object storage essentially bundles the data itself along with metadata tags and a unique identifier. The metadata is customizable, which means you can input a lot more identifying information for each piece of data. These objects are stored in a flat address space, which makes it easier to locate and retrieve your data across regions.

This flat address space also helps with scalability. By simply adding in additional nodes, you can scale to petabytes and beyond.

You can learn more about object storage by watching the video below:

The Difference Between Object and File Storage

Now that you know the basics of both object and file storage, let's look at some of the key differences separating the two.

To start, object storage overcomes many of the limitations that file storage faces. Think of file storage as a warehouse. When you first put a box of files in there, it seems like you have plenty of space. But as your data needs grow, you'll fill up the warehouse to capacity before you know it. Object storage, on the other hand, is like the warehouse, except with no roof. You can keep adding data infinitely – the sky's the limit.

If you're primarily retrieving smaller or individual files, though, then file storage shines with performance, especially with relatively low amounts of data. Once you start scaling, though, you may start wondering, "How am I going to find the file I need?"

In this case, you can think of object storage as valet parking while file storage is more like self-parking (yes, another analogy, but bear with me!). When you pull your car into a small lot, you know exactly where your car is. However, imagine that lot was a thousand times larger— it'd be harder to find your car, right? Because object storage has customizable metadata and all the objects live on a flat address space, it's similar to handing your keys over to a valet. Your car

will be stored somewhere (and in this case, it could even be in another country), and when you need it, the valet will get the car for you. It might take a little longer to retrieve your car, but you don't have to worry about wandering around looking for it.

For a real-life example of why this makes a difference, we can look at X-rays. An X-ray file would have limited metadata associated with it, such as created date, owner, location, and size. An X-ray object, on the other hand, could have any number of metadata tags tied to it, such as patient name, date of birth, injury details, which area of the body was X-rayed — all of this in addition to the same metadata tags that the file had. This makes it incredibly useful for doctors to pull up the relevant information for reference.

This was only a general overview of the differences between object storage and file storage, but it should give you a clearer idea of the advantages of each type. While file storage has its uses, enterprises have started to replace file storage with object storage to meet their unstructured data storage needs.

Cloudian HyperStore offers fully S3-compatible object storage that is infinitely scalable, easy to manage, and cost-efficient.

	OBJECT STORAGE	FILE STORAGE
PERFORMANCE	Performs best for big content and high stream throughput	Performs best for smaller files
GEOGRAPHY	Data can be stored across multiple regions	Data typically needs to be shared locally
SCALABILITY	Can scale infinitely to petabytes and beyond	Potentially scales up to millions of files but can't handle more
ANALYTICS	Customizable metadata allows data to be easily organized and retrieved	Limited number of set metadata tags

Object Storage vs File Storage
Side-by-side comparison

Object Storage vs Block Storage
Find out how object storage differs from block storage.

VIEW PDF >

Cloudian, Inc.
177 Bovet Road, Suite 450, San Mateo, CA 94402
Tel: 1.650.227.2380; Email: info@cloudian.com
www.cloudian.com