

COLD VS HOT DATA STORAGE: WHAT'S THE DIFFERENCE?



The usage of temperature terms hot and cold to distinguish between storage options originates in the physical ways we have stored data for decades. Items closer to the data center were accessed more regularly, and were located, literally, in storage facilities that were hot.

Items further from the data center had slower loading times, so it became the place to store data you needed to access much less frequently. This type of storage was done differently from hot storage—typically either using old drives or drives that were turned off entirely. These storage types did not generate the kind of heat the other storage facilities created.

Let's take a look at data storage, including how the cloud is affecting how we store and compute data.

Cloud vs on-prem storage: hot, warm, cold

Which storage is hot and which storage is cold can vary depending on the kind of storage architecture you use:

- In a distributed system that uses [edge devices](#), hot storage can serve as both **computational memory** and **storage** for each individual edge device.
- Pure cloud services can offer both hot and cold computational memory and storage, with any off-cloud device using cold storage.

Here's a breakdown of the traditional hot vs cold categories, with warm data storage emerging as a

third type:

	 Hot storage	 Warm storage	 Cold storage
Location	Close to the moment of computation	On a remote server or private network, usually a step away from users	In cloud services
Defining Characteristic	Very fast access	Medium speeds	Different pricing tiers for hot and cold options
Examples	Personal hard drives, SSDs, flash drives	Larger, cheaper, spinning drives	Cloud drives, AWS, Google Cloud Storage

Comparison of data storage types

When to use hot storage

All data that you need to be able to access immediately must be placed in hot storage. This can include data that is:

- Known to change
- Used for customer query purposes
- Used in any current projects

Hot storage requires immediate and reliable access. For example, Amazon and Google's services have a 99.95% availability, while Azure offers up to a 99.99%. Data that comes in from a hot storage system can be called "data streams". Many sophisticated systems [process the flow of data](#) as it flows in from your storage.

Data transfer speeds depend on one primary thing: **How many routes does the data pass through to get from its host to its destination?** Data that is processed closest to its source will be fastest. Data that has to travel over a few different networks and arrive on a developer's laptop can take longer to access.

For example, if data is hosted in Google Storage and the user wishes to retrieve and process that storage through another Google server or within a Google Colab notebook, they should find their process speeds to be fairly quick. If the data is fetched from Google Storage and being transferred to a local external hard drive, the data has to pass through many more routes. That data is also dependent on network speeds and read/write speeds to write the data to a new hard drive.

In machine learning projects, data is read multiple times and needs to be provided to the ML model quickly, so it should be located in hot storage. This data can be on a drive on the modeler's laptop or on an external drive. For large companies with notoriously large datasets can require immediate access to many terabytes or petabytes, and a cloud service provider can help manage their hot storage options. Once the data has been used, or replaced and ready to be retired, the data can be stored in cold storage for a team's data versioning.

When to use cold storage

Cold storage is meant for data that is rarely used. This is data that needs to stick around for some reason, such as legal reasons, compliance, or simple record keeping. Data-versioning is becoming more common, so old versions of datasets are a good item to keep on cold storage. It could be data that is no longer updated, but is still queried. This data is also known as "dormant data".

Cold storage data retrieval can take much longer than hot storage. It can take minutes to hours to access cold storage data, so this data is good to use for projects that allow for patience and planning—not tight deadlines. Cold storage might even require a person to physically sift through a physical set of hard drives, like a library of storage devices, then connect it to a computer and retrieve the data. When it is entirely disconnected from a computer like this, the physical storage is quite literally cold.

In this latter scenario, cold storage can be used to refer to any data that is not stored on the cloud.

Storage in the cloud: pricing options

Many services are moving towards the cloud and so are hot and cold storage options. Their terms, hot and cold, are synonymous with what they were before the cloud. Each major provider has its own hot and cold tiers.

Pricing can be complicated because it depends on [several factors](#), like whether the storage is available in one time zone or across multiple time zones. A good rule of thumbs is that cold storage costs half what hot storage does.

General prices for storage are:

- Hot storage can range from ~\$.10/GB to \$.17/GB
- Cold storage ranges at \$.045/GB to \$.08/GB

Cloud options are changing how we look at data computation and data storage. But the terms hot and cold continue to refer primarily to how accessible your storage is. Fast and easy accessibility is hot. Slow and difficult accessibility is cold.

Additional resources

When it comes to data storage and cloud computing, BMC has you covered. Check out these additional BMC Blogs:

- [What is "Data Center Colocation"? Data Center Colocation Explained](#)
- [Data Center Management Tools: Features, Functions, and How To Choose](#)
- [Data Center Tiers: What Are They and Why Are They Important?](#)
- [What Is A Software Defined Data Center?](#)
- [How Google is Using AI for Data Center Cooling](#)
- [What is a Hyperscale Data Center?](#)