

# HOW TO CHOOSE AN OBJECT STORAGE TARGET REPOSITORY



For many mainframers, the concept of [writing to object storage from zSeries mainframes over TCP/IP](#) is a new concept.

The ease of use and the added value of implementing this solution is clear, but there is another question: What to use as a target repository? How do customers decide on a vendor for object storage and whether a private cloud, hybrid cloud, or public cloud should be used? Target repositories can be either an on-premises object storage system, like Hitachi HCP, Cohesity, or a public cloud, such as AWS, Azure or GCP.

The best option for you depends on your individual needs. There are pros and cons in each case. In this post, we break down the factors you need to consider as you choose a target cloud repository that will meet your needs.

## Network bandwidth and external connections

Consider the bandwidth of the OSA cards and external bandwidth to remote cloud, if cloud is an option. Is the external connection shared with other platforms? Is a cloud connection already established for the corporation?

For on-premises storage, network connectivity is required, yet it is an internal network with no external access.

## Amount of data recalled, restored, or read back from repository

There are added costs for reading data back from the public cloud, so an understanding of

expected read throughput is important when comparing costs. If the read rate is high, then consider an on-premises solution.

## **DR testing and recovery plans**

Cloud-based recovery allows recovery from anywhere, and public clouds can replicate data across multiple sites automatically. The disaster recovery or recovery site must have network connectivity to the cloud.

On-premises solutions require a defined disaster recovery setup, a second copy of the object storage off-site that is replicated from the primary site. Recovery at the DR site will access this replicated object storage.

## **Corporate strategies, such as “Mainframe Modernization” or “Cloud First”**

You should be able to quickly move mainframe data to cloud platforms by modernizing backup and archive functions. Cloud also offers either policy-driven and/or automatic tiering of data to lower the cost of cold storage.

If there is no cloud initiative, the on-premises solution may be preferred. Many object storage providers have options to push the data from on-premises to public cloud. So, hot data can be close and cold data can be placed on clouds.

## **Cloud acceptance or preferred cloud vendor already defined**

Many corporations already have a defined cloud strategy and a cloud vendor of choice. You'll want a vendor-agnostic solution.

The knowledge of defining the repository and maintaining it could be delegated to other groups within the organization familiar with and responsible for the corporate cloud.

## **Cyber resilience requirements**

On-premises solutions can generate immutable snapshots to protect against cyberthreats. An air-gapped solution can be architected to place copies of data on a separate environment that can be detached from networks.

Cloud options also include features like versioning, multiple copies of data, and multi-authentication to protect data and allow recovery.

## **Floor or rack space availability**

With an on-premises solution, floor space, rack space, power, etc. are required. With a cloud solution, no on-premises hardware is required.

## **Performance**

There is no clear-cut performance benefit for either solution. It depends upon the hardware and network resources and the amount of data to be moved and contention from other activity in the shop using the same resources.

Cloud customers with performance concerns may choose to establish a direct connection to cloud providers in local regions to prevent latency issues. These concerns are less relevant when a corporate cloud strategy is already in place.

## **Costs**

Cloud storage is priced by repository size and type. There are many add-on costs for features and costs for reading back. There are mechanisms to reduce costs, such as tiering data. Understanding these costs upfront is important.

On-premises object storage requires a minimum of two systems for redundancy, installation, and maintenance.