

Use Tiering With a Microsoft Azure Blob Storage Node

Each object in Azure Blob Storage has an access tier associated with it:

Hot – Optimized for storing data that is accessed frequently.

Cool – Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive – Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements (on the order of hours).

CTERA Portal Azure Blob Storage can use both hot and cool tiers. Data in the cool access tier can tolerate slightly lower availability, but still requires high durability, retrieval latency, and throughput characteristics similar to hot data. For cool data, a slightly lower availability service-level agreement (SLA) and higher access costs compared to hot data are acceptable trade-offs for lower storage costs.

Step-by-step guide

To configure tiering for an Azure Blob Storage node:

- You configure your Azure Blob Storage bucket to tier data that is not frequently accessed in the cool tier with Azure Blob Storage lifecycle management to create a rule-based policy to transition your data to the best access tier. For more details, see <https://azure.microsoft.com/en-us/blog/azure-blob-storage-lifecycle-management-now-generally-available/>.

For example, the following sample policy manages the lifecycle for such data. It applies to block blobs in container portalsn and tiers tier blobs to cool storage 7 days after the last modification.

```
{
  "rules": [
    {
      "name": "rulePortalsn",
      "enabled": true,
      "type": "Lifecycle",
      "definition": {
        "filters": {
          "blobTypes": [ "blockBlob" ],
          "prefixMatch": [ "portalsn" ]
        },
        "actions": {
          "baseBlob": {
            "tierToCool": { "daysAfterModificationGreaterThan": 7 },
          }
        }
      }
    }
  ]
}
```