

# Local versus distributed storage

This page explains the benefits and drawbacks of local versus distributed storage in MicroCloud and helps you understand when you should use each.

In MicroCloud, you can set up both types of storage, or only one, depending on your needs. For production environments, we recommend that you use both. For detailed hardware requirements for both local and distributed storage, see [Storage requirements](#).

## Setup

Local storage is simpler: each cluster member has its own local disk, and workloads on that member read/write directly to it. It does not provide data redundancy.

Distributed storage (using [Ceph RBD](#)) replicates data across multiple disks. Workloads access the shared storage over a Ceph network, set up by MicroCloud using MicroCeph. Due to this more complex setup and replication, distributed storage uses more CPU, RAM, and network bandwidth.

## Access

With local storage, a workload on one cluster member cannot access data on other members unless it's copied over, shared over the network, or the workload is migrated over.

With distributed storage, workloads can access shared data cluster-wide. Access occurs over the Ceph network; RBD volumes are mapped only on the member where the workload runs.

## Speed

Local storage is often faster because each read/write operation stays on the local disk.

Distributed storage can be slower due to network latency and data replication overhead, though the speed difference can be small depending on hardware and network configuration. This overhead is most noticeable for write-heavy or latency-sensitive workloads.

## Instance migration

[Migrated between cluster members](#). For instances using local storage, data is transferred over the network along with those instances. MicroCloud (LXD) creates a temporary snapshot file, then sets up an NBD (Network Block Device) listener to transfer disk data over the network.

Skip to content

 latest ▼

Data stored in a distributed setup is accessible to all cluster members and does not need to be transferred during instance migration. This means that instances using distributed storage can be migrated more quickly.

## High availability and data redundancy

With local storage, you risk a higher potential of data loss and workload disruption if a local disk or its host machine fails. Even if you set up a backup and recovery procedure, you'll still lose any data since the last backup, it will take longer to restore data from backups than with distributed storage.

Distributed storage avoids a single point of failure. By default, MicroCloud replicates data three times across three disks. This protects against data loss and reduces recovery time; if one disk fails, data is still available from another disk. Since the data remains available, workloads can restart on other members without needing to restore from backups. For more information about HA in MicroCloud, see [High availability](#).

## Recommendations

For production deployments where you need both performance and resilience, use both local and distributed storage. Choose which to use depending on the workload type:

- Use local storage for temporary workloads where you need fast disk access, and data loss is less of a concern.
- Use distributed storage when you have long-running workloads with uptime requirements that need data redundancy.

## Related topics

### How-to guides

- [How to configure Ceph networking](#)

### Explanation

- [About MicroCloud](#)

### Requirements

- [Storage requirements](#)

[Skip to content](#)

  [latest](#) ▼