

Cache Configuration

The cache is the local storage of the Nasuni Filer. When running the Nasuni Filer on a virtual platform, you can configure the size of the cache disk and the copy-on-write (COW) disk. On Nasuni hardware appliances, the sizes of the cache disk and the copy-on-write (COW) disk are pre-configured, and there is no need to configure them.

Proper configuration of the cache and related features helps provide the following benefits:

- Enables access to data.
- Improves the performance experience of users and applications.
- Ensures access to data during Internet outages.
- Reduces use of network bandwidth.

This document offers suggestions for configuring the cache of the Nasuni Filer.

About the cache

The cache is the local storage of the Nasuni Filer. All data and metadata that is accessed regularly is kept locally in the cache (“cache resident”). If requested data is not cache resident, the system accesses the requested data from cloud storage, then the data is staged into the cache, and provided for the request.

Note: *You don’t have to keep all data locally. The Nasuni Filer is extremely efficient at keeping only your most-needed data locally, and retrieving data from cloud storage as needed.*

The size and status of the cache are displayed on the Home page.

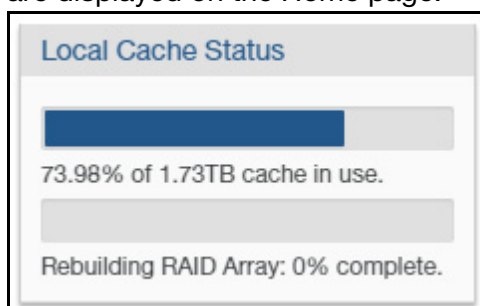


Figure 1-1: Local Storage Status (cache).

Note: *Nasuni’s display of size might differ from other indications of size, such as Windows Explorer and other utilities. Typically, such utilities display only the size of the data currently present in the local cache, while Nasuni displays the full size, regardless of where the data is.*

Any data written to a volume is first staged in the local cache. The first snapshot preserves each data item from the cache in Nasuni's cloud storage. Subsequent snapshots include only the changes between the original version of the data and the most recent version of the data.

The cache continuously tracks "recently changed" data, such as data created or changed since the last snapshot. New data that has been saved to the cache, but that has not yet been preserved in Nasuni's cloud storage, is classified as "new data in cache not yet protected".

A chart of **New Data in Cache (Not yet protected)** appears on the Home page.

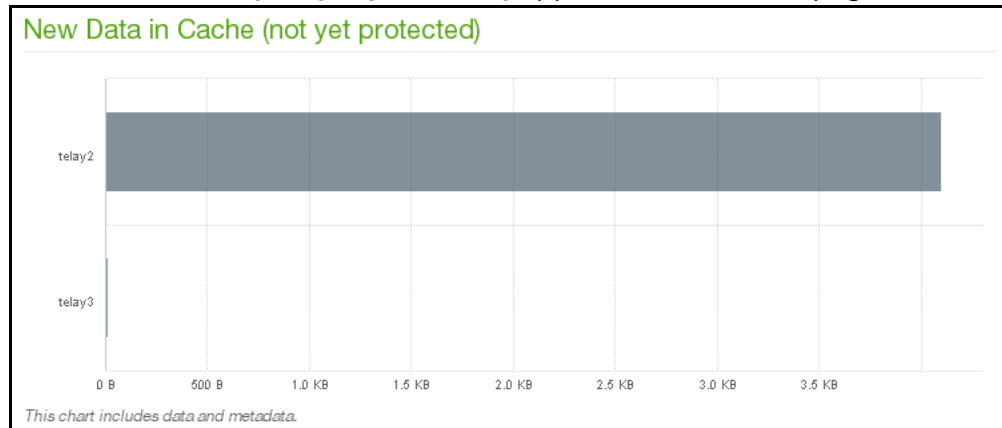


Figure 1-2: New Data in Cache chart.

The most recently used data remains in the cache, including data, metadata, complete files, or parts of files. Saved data that was written once and rarely used afterwards is eventually released (or "evicted") from the cache to free up space for new data. If any of this released data is later requested for reads or writes, the Nasuni Filer retrieves the data from Nasuni's cloud storage and puts it back into the local cache automatically.

Cache structure

This figure illustrates the structure of the cache:

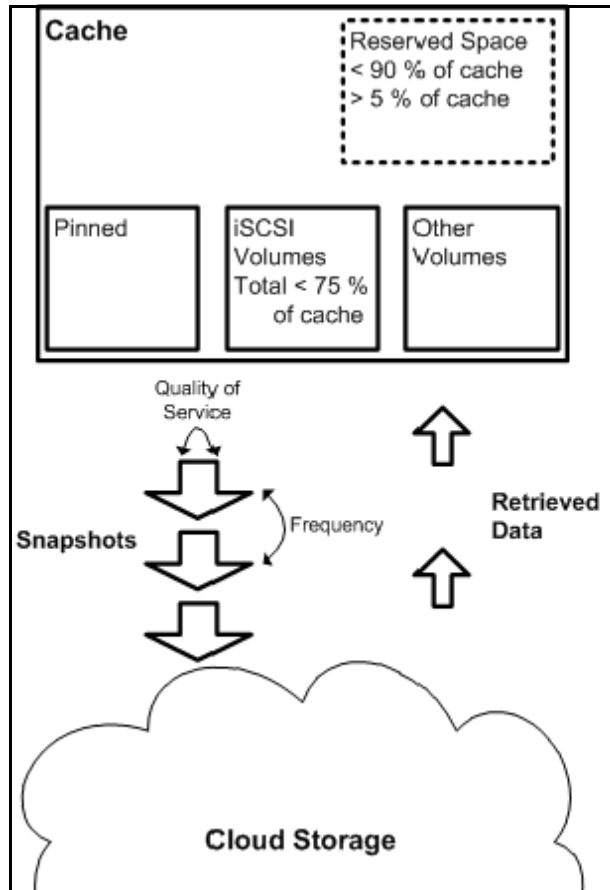


Figure 1-3: Cache structure.

“Pinned” data is not removed from the cache. If too much cache space is taken up by pinned data, an Alert notification is given.

iSCSI volumes are always pinned in the cache. The total size of all iSCSI volumes together can use up to 75 percent of the cache.

Reserved space is space set aside in the cache for incoming data. By default, the Nasuni Filer automatically manages the reserved space. The maximum amount of reserved space is 90 percent of the cache size.

Snapshots preserve data from the cache in Nasuni’s cloud storage. You can configure the frequency of snapshots.

The Quality of Service is the limit on bandwidth for snapshots. The value appears on the Home page.




Figure 1-4: Quality of Service.

Considerations for configuring the cache

This section discusses points to consider when configuring the cache.

Cache size

The size of the cache should be based on the size of the working set (the amount of data being accessed by users or applications on a regular basis), including metadata.

Note: *Cache size has no effect on the cost of Nasuni services. The Nasuni license is based on the storage used, not on cache size.*

Cache disk and copy-on-write (COW) disk (aka snapshot disk)

The snapshot (copy-on-write or COW) disk is used during the snapshot process. If any writes to the Nasuni Filer occur during a snapshot, the previous data from the cache disk is copied to the COW disk, and the new data is written to the cache disk. Hence, the term “copy-on-write”. This allows new writes to take place at any time, even during the snapshot process.

The size of the copy-on-write (COW) disk is determined by how much data can change while a snapshot is running. For example, a heavily used file server, where a lot of data can change while a snapshot is running, might need a larger copy-on-write (COW) disk than a server used mainly for backups, where the snapshot is scheduled for a time when the backup isn't running.

However, if the copy-on-write (COW) disk is too small, then a snapshot might fail, then try again. The snapshot will eventually succeed, but this is very inefficient.

Tip: *The copy-on-write (COW) disk should generally be at least one-quarter the size of the cache disk. Although the Nasuni Filer issues a warning during startup if the COW disk is too small relative to the cache, the Nasuni Filer does not issue this warning if the COW disk is larger than 250 GB.*

Note: *The copy-on-write (COW) disk is the second-largest disk, and the cache disk is the largest disk, by default. However, if you change the relative sizes of the disks, the COW disk might not be the second-largest disk.*

Increasing cache size

For an existing Nasuni Filer, you can only *increase* the size of the cache and copy-on-write (COW) disks (snapshot disks). This means that you cannot decrease a value once you have increased it. Therefore, make sure you understand your resizing requirements before you make changes.

Note: *If necessary, you can create a new Nasuni Filer with a smaller cache size, then perform a Disaster Recovery from the existing Nasuni Filer to the new Nasuni Filer.*

Note: *To add cache space by configuring additional cache disks, contact Nasuni Technical Support.*

Cache size and performance

With a larger cache, more data can be cache resident, and access to that data is faster.

On a hardware appliance, more RAM and CPUs also improve performance. Similarly, on a virtual machine, increased physical resources also improve performance.

Performance issues from cache misses

To determine whether performance issues are caused by cache misses, first find out which specific files are having problems. Then see if those files are cache resident, by clicking **File Browser**, selecting the volume, navigating to each file, and examining its properties in the **Properties** pane. If the **Cache Resident** status is Yes, then the file is cache resident.

*Note: A file might be marked as **Cache Resident** even if only part of the file is present in the cache. To ensure that a given file is cache resident, use **Bring into Cache** on the **File Browser** page.*

If all other investigation fails to solve performance issues, contact Nasuni Technical Support to analyze the logs to see if a lot of cache misses are happening.

Automatically releasing data from cache

If the amount of cache resident data exceeds a specified percentage of the size of the cache, data is released from the cache until the amount of cache resident data no longer exceeds that specified percentage. Before releasing any data from the cache, the Nasuni Filer ensures that the data is stored in the cloud. The Nasuni Filer may perform a snapshot to preserve data in the cloud. Released data is not lost, because the data is stored in the cloud, ready to be retrieved.

The Nasuni Filer releases the least recently used data first, using a sophisticated least recently used (LRU) algorithm.

Loading large amounts of data into the Nasuni Filer

Nasuni's best practices for loading large amounts of data into a Nasuni Filer involve:

- Disable remote access temporarily. This makes the snapshots more efficient since they are less concerned with the response time of updates to the remote sites.
- Decrease snapshot frequency temporarily. This reduces the demand for resources, and allows freer access to the cache.
- Use the Data Migration Service (DMS) to move data. DMS is cache aware and slows down when the disk fills, instead of overfilling.
- Enable as much bandwidth (Quality of Service) as can be spared for when snapshots do occur.

Downloading large amounts of data from cloud

To download data from the cloud to the cache faster, do the following:

- Make sure that you have sufficient download bandwidth from your ISP.
- Adjust the Inbound Limit of the Quality of Service (QoS) rule to the highest rate possible. If the downloading is temporary, you can change or remove the rule when the downloading is complete.
- Ensure that the cache space is large enough to receive the data, so the system doesn't have to spend a lot of time making room. You can increase the size of the cache itself, or increase the reserved space percentage.

Keeping all data locally (cache resident)

To keep all data cache resident, the size of the non-reserved part of the cache must be larger than the total amount of data that you want to keep locally, including metadata. The size of the cache disk must be large enough for the cache. You can then bring all the data into the cache, and pin the data into the cache. If the cache is larger than the total amount of data, all the data remains in the cache, and pinning is not necessary.

Proactively managing the contents of the cache

There are situations where you might need to proactively manage the contents of the cache, because a process might not automatically bring the desired data into the cache. These situations include the following:

- After performing a restore.
- After performing a Disaster Recovery procedure that includes unpinning data.
- When performing data migrations.

You can proactively manage the contents of the cache by performing any of the following actions:

- Pinning data in the cache.
- Accessing files or folders to ensure that they're brought into the cache.
- Bringing data into the cache manually, using **Bring into Cache** on the **File Browser** page.

Different Nasuni Filers for different uses

Different uses of the Nasuni Filer might have different optimum configurations of the cache. For example, the usage pattern for ordinary user access is different from the usage pattern for backups or data mining. To ensure that each use case receives the most suitable configuration, you can set up different Nasuni Filers for different uses. The Nasuni Filer is intelligent and does its best to adapt to any use, but you can help guarantee the best behavior for each use case by isolating the use cases on different Nasuni Filers.

Here are some examples of different uses for Nasuni Filers:

- Performing migrations: Minimize pinned data, infrequent snapshots.
- Primary file server: Pin folders of commonly accessed files, moderate Quality of Service, regular snapshots.
- Hot standby for Disaster Recovery: Minimize pinned data, enable Auto Cache, frequent syncs.
- High-performance application: Pin data of commonly accessed files, high Quality of Service, regular snapshots.

Different volumes for different uses

It is also possible to employ different volumes on the same Nasuni Filer for different purposes. The Nasuni Filer cache adapts to any use (if Reserved Space Percentage is Automatic), but you can also adjust configurations to accommodate volumes with different uses on the same Nasuni Filer.

For example, on a single Nasuni Filer, you might want to use one volume for ordinary user access to data and another for backups. To accommodate these two different uses, you might pin the most commonly used user data in the cache. You might also dedicate the Nasuni Filer to performing as a file server M-F 8-7, but then do backups at 3 AM.

Cache Specifications

The default cache size on disk is 250 GiB for a VM platform.

The current maximum cache size for the Nasuni Filer is 500 TiB. However, each platform also has its own maximum cache size for a single cache disk. (Multiple cache disks are supported.) The following table lists the maximum cache size per single cache disk (by platform) supported by the Nasuni Filer. Visit [VMware features](#) and [Hyper-V features](#) for the most complete information.

Virtual Platform	Maximum Cache Size for a Single Cache Disk
VMware ESXi 5.5 and above NOT using VMFS-5	2 TiB
VMware ESXi 5.5 and above	62 TiB (using VMFS-5) 64 TiB (using physical Raw Disk Mapping (RDM) (aka “pass through disks”))
Microsoft Hyper-V	2 TiB (Server 2008) 500 TiB (using “pass through disks”) (Server 2008) 64 TiB (using VHDX) (Server 2012) 64 TiB (Server 2012 R2)
Microsoft Azure	4,095 GiB (Nasuni Filers support the Microsoft Azure maximum size for a single disk, which is currently 4,095 GiB.)

Virtual Platform	Maximum Cache Size for a Single Cache Disk
Amazon EC2	1 TiB for hard disks, 16 TiB for SSDs

Note: Nasuni supports VMware ESXi 5.5 and above.

Note: To add cache space by configuring additional cache disks, contact Nasuni Technical Support.

Note: Note that you can configure pass-through disks on VMware and Hyper-V.

Quality of cache storage

If you are using a Nasuni Filer hardware appliance, that platform is optimized for the cache.

If you are using virtual machines, ensure that the storage platform meets these specifications:

- Protection, such as 4+1 RAID-5.
- Multiple spindles.
- 10K RPM or faster hard drives.
- Battery-backed write cache ability.

Actions affecting the cache

This section describes the actions you can perform that affect the cache.

Removing data from cache

Deleting a volume removes all the volume's data stored in the cache, as well as data stored in Nasuni's cloud storage.

Disconnecting from a remote volume removes any data in the cache.

Determining if data is in the cache

By examining a folder or file in the **File Browser** display, you can determine if a folder or file is currently in the cache (called "cache resident").

Bringing data into the cache manually

Note: This is optional. Cache management is automatic.

If a file or folder is not in the cache, and you want it in the cache, find it in the **File Browser** display, then click **Bring into Cache**.

Note: If accessed data is not in the cache, it is brought into the cache automatically.

Pinning or unpinning folders in the cache

You can enable or disable pinning folders.

On the Nasuni Filer or the Nasuni Management Console, on the **File Browser** page, select a folder, then click **Edit Cache Settings**.

Note: You cannot unpin iSCSI volumes. iSCSI volumes are always pinned in the cache.

Note: Pinning a folder in the cache does not bring that folder's data into the cache. Pinning a folder prevents that folder's data from being removed from the cache.

Reserving space in cache

By default, the Nasuni Filer automatically manages the percentage of local cache space reserved for new, incoming data, using an advanced algorithm to optimize cache usage. The remainder of the cache retains the data locally that users are most likely to need. However, you can override the percentage reserved for new, incoming data. The percentage that you set applies to all volumes on the Nasuni Filer maximum percentage of the cache reserved for new, incoming data is 90 percent of the cache size. The minimum percentage is 5 percent.

By setting the amount of reserved space, you disable the automatic management of this value.

On the Nasuni Filer, click **Configuration**, then select **Cache Settings**. On the **Cache Settings** page, select **Automatic** or one of the available percentages from the **Reserved Space Percentage** drop-down list.

On the Nasuni Management Console, click **Filers**, then click **Cache Settings** in the left-hand column. Select the Nasuni Filers, then click **Edit Filers**.

Important: The Data Migration Service uses the cache. To ensure that the Data Migration Service has enough space, set Cache Settings to 30 percent or greater.

Snapshots and the cache

Snapshots preserve data from the cache in the cloud, making that data eligible for removal from the cache, if necessary. Thus, scheduling frequent snapshots can help free up room in the cache for incoming writes.

Chunks

Before sending data to the cloud, Nasuni breaks files into optimally-sized pieces for transport between the on-premises cache and cloud storage. This not only disguises the actual sizes of files, but also improves performance. These chunks are then deduplicated, compressed, and encrypted.

If the file is smaller than 1 GiB, the default chunk size is 1 MiB. If the file is 1 GiB or larger, and the appliance has less than 16 GiB of RAM, the default chunk size is 2 MiB. If the file is 1 GiB or larger, and the appliance has 16 GiB of RAM or more, the default chunk size is 10 MiB.

For customer-provided clouds, such as Microsoft Azure, Cleversafe, IBM Cloud Object Storage, EMC ViPR, and EMC ECS, if directed by Nasuni Support, you can adjust the chunk size, or enable or disable Nasuni's compression and deduplication, using the **Cloud I/O** area of the **Volume Overview** page (Nasuni Filer) or the **Volume Cloud I/O** page (NMC). If you do manually change the chunk size, the variable chunk size mentioned above no longer operates. You can restore the variable chunk size mentioned above by leaving the **Chunk Size** field blank and then clicking **Save**.

Scheduling snapshots

You can schedule snapshots on a per-volume basis.

On the Nasuni Filer, on the **Volumes** page, select a volume, then click **Edit Properties**. Click edit next to the **Snapshot Schedule** status. On the **Snapshot Scheduling** page, select the days, start time, stop time, and frequency for snapshots.

On the Nasuni Management Console, on the **Volumes** page, click Snapshot Schedule, select the volumes to change, then click **Edit Volumes**.

Tip: On volumes with Global Locking enabled, we recommend reducing the normal snapshot and synchronization frequency of the volume, because Global Locking provides file synchronization independently of the snapshot and synchronization frequency. This file synchronization occurs when the file gives up the file lock, and when file locks are degraded to read locks (in Advanced mode only).

Manually starting snapshots

On the Nasuni Filer, to take a snapshot sooner than the scheduled time, click "**Take snapshot now**" on the **Volumes** page.

On the Nasuni Management Console, on the **Volumes** page, for the volume that you want to take a snapshot of, click **Take Snapshot**.

Changing Quality of Service rules

The Quality of Service is the limit on inbound and outbound bandwidth for snapshots. The default inbound Quality of Service is unlimited. The default outbound Quality of Service is 10 megabits per second. The larger the outbound Quality of Service, the faster snapshots complete, and the sooner space can become available in the cache.

On the Nasuni Filer, to change the Quality of Service rules, click **Configuration**, then select **Quality of Service** from the list. On the Quality of Service Settings page, select the days, start time, stop time, and bandwidth limit, from 50 kilobits per second to 25 gigabits per second. You can set up to 12 rules.



On the Nasuni Management Console, on the **Filers** page, click **Quality of Service**, select the Nasuni Filers to change, then click **Edit Filers**.

Changing the size of the cache disk or copy-on-write (COW) disk

On the virtual platform, you can change the size of the disk that the Nasuni Filer uses for its cache. You can also change the size of the disk that the Nasuni Filer uses for copy-on-write (COW).

Supported techniques for changing the size of a disk

There are several techniques for achieving cache or snapshot disks of a different size:

Disk	Supported Technique	VMware	Hyper-V	Azure	EC2
Cache	VM user interface	Increase	Increase		
COW	VM user interface	Increase	Increase		
Cache	New Nasuni Filer, then Disaster Recovery	Increase/Decrease	Increase/Decrease	*	Increase/Decrease
COW	New Nasuni Filer, then Disaster Recovery	Increase/Decrease	Increase/Decrease	*	Increase/Decrease
COW	Shut down Nasuni Filer, delete COW, create new COW, restart Nasuni Filer	Increase/Decrease	Increase/Decrease	Increase/Decrease **	Increase/Decrease
Cache	Configuring additional cache disks	Increase ***	Increase ***		

* If you add another disk, the largest disk becomes the cache disk. The second largest disk becomes the COW disk.

** You cannot delete the built-in COW disk, but you can add a new COW disk of a different size. The largest unused disk becomes the new COW disk. Note that this requires assistance from Nasuni Support for configuration.

Note: *** To add cache space by configuring additional cache disks, contact Nasuni Technical Support.

Tip: The copy-on-write (COW) disk should generally be at least one-quarter the size of the cache disk. Although the Nasuni Filer issues a warning during startup if the COW disk is too small relative to the cache, the Nasuni Filer does not issue this warning if the COW disk is larger than 250 GB.

Using the VM user interface to increase the size of the cache disk or copy-on-write (COW) disk

On the VMware ESX/ESXi and Microsoft Hyper-V virtual platforms, you can increase the size of the cache disk and the COW disk.

Note: There must be sufficient disk space in the datastore to increase the size of the disk.

Note: You must shut down the virtual machine to change the size of the disk.

Note: Nasuni supports VMware ESXi 5.5 and above.

Caution: Perform a manual snapshot on the Nasuni Filer before changing the size of the disk.

Tip: Nasuni recommends using Thick provisioning for all virtual disks for performance reasons. On the VMware virtual platform, Nasuni recommends Thick Provision Eager Zeroed.

Warning: Do not attempt to restore the state of the virtual machine from a virtual machine snapshot or backup.

Tip: The copy-on-write (COW) disk should generally be at least one-quarter the size of the cache disk. Although the Nasuni Filer issues a warning during startup if the COW disk is too small relative to the cache, the Nasuni Filer does not issue this warning if the COW disk is larger than 250 GB.

vSphere Client platform

Note: If you have taken a VMware snapshot, you cannot change the size of the disk. You must delete the snapshot first.

Warning: Do not attempt to restore the state of the virtual machine from a virtual machine snapshot or backup.

To change the size of the disk on the vSphere Client platform, follow these steps:

1. Launch the VMware vSphere Client. The VMware vSphere Client login window appears.
2. Log in to the vSphere Client with a valid IP address and your user name and password. The vSphere Client opens.
3. In the pane on the left, expand the host to display the list of virtual machines.
4. Click the name of your Nasuni Filer. This is the name you entered when you installed the Nasuni Filer. Information about your Nasuni Filer appears in the right pane.
5. If the virtual machine is running, then shut down the virtual machine using these steps:
 - a. On the **Console** tab, click in the command line area.
 - b. Press **Enter**.
 - c. Enter the login name and password for the service console. The default login name is `service` and the default password is `service`.



- d. Enter **shutdown**.
 - e. Confirm the shutdown by entering **yes**.
The virtual machine shuts down.
6. On the **Getting Started** tab, click **Edit virtual machine settings**.
Alternatively, on the **Summary** tab, click **Edit Settings**.
Alternatively, right-click the name of your Nasuni Filer in the list of virtual machines, then select **Edit Settings** from the drop-down list.
The Virtual Machine Properties dialog box appears.
 7. On the Hardware tab, select the desired disk in the Hardware list on the left side.
The cache disk is disk 2 and is the largest hard disk.
The copy-on-write (COW) disk is the second-largest hard disk.
Information about the selected hard disk appears on the right side.
 8. Enter or select the new **Provisioned Size**. The Provisioned Size must be less than or equal to the displayed **Maximum Size** of that hard disk. The new Provisioned Size must be greater than the previous Provisioned Size.
For details on the size of the cache disk, see [“Cache Specifications” on page 7](#).
Tip: If necessary, it is possible to reduce the size of the copy-on-write disk. To do this, shut down the Nasuni Filer (see the [Nasuni Filer Administration Guide](#) for instructions), delete the existing copy-on-write disk, create a new copy-on-write disk with the desired size, and then restart the Nasuni Filer.
 9. Click **OK**. The size of the selected hard disk changes to the specified size.
 10. If you have configured more than one cache disk, repeat [step 7](#), [step 8](#), and [step 9](#) for each cache disk.
 11. On the **Getting Started** tab, click **Power on the virtual machine**.
Alternatively, on the **Summary** tab, click **Power On**.
Alternatively, right-click the name of your Nasuni Filer in the list of virtual machines, then select **Power** from the drop-down list, then select **Power On** from the drop-down list.
The virtual machine powers on. After a significant cache resize, the operations after powering on can take some time, with some performance impact during that period. The progress of the cache resize is visible on the Nasuni Filer user interface.

Hyper-V Manager platform

To change the size of the disk on the Hyper-V Manager platform, follow these steps:

1. Launch Hyper-V Manager.
2. In the **Virtual Machines** list, click the name of your Nasuni Filer. This is the name you entered when you installed the Nasuni Filer. Information about your Nasuni Filer appears.



3. If the virtual machine is running, then follow these steps:
 - a. In the **Virtual Machines** list, right-click the name of your Nasuni Filer, then select **Shut Down** from the drop-down list.
Alternatively, in the **Virtual Machines** list, select the name of your Nasuni Filer, then click **Shut Down** in the list of actions on the right side.
The **Shut Down Machine** dialog box appears.
 - b. Click **Shut Down**. The virtual machine shuts down.
4. In the **Virtual Machines** list, right-click the name of your Nasuni Filer, then select **Settings** from the drop-down list.
Alternatively, in the **Virtual Machines** list, select the name of your Nasuni Filer, then click **Settings** in the list of actions on the right side.
The **Settings** dialog box appears.
5. In the **Hardware** list on the left side, select the desired hard drive.
The default label for the cache hard drive is “Hard Drive Cache.vhd”. The cache hard drive is the largest hard drive. The default label for the copy-on-write (COW) hard drive is “Hard Drive Cow.vhd”.
Information about the selected hard drive appears on the right side.
6. Click **Edit**. The **Edit Virtual Hard Disk Wizard** dialog box appears.
7. Select **Expand**, then click **Next**.
8. Enter the **New size** of the disk. The “New size” must be less than or equal to the displayed **Maximum** size of that hard disk. The “New size” must be greater than the previous size. Then click **Next**.
When you resize the cache, Nasuni recommends that the cache disk be no more than four times the size of the snapshot (copy-on-write or COW) disk.
For details on the size of the cache disk, see [“Cache Specifications” on page 7](#).
9. Click **Finish**. The size of the disk changes to the specified size.
10. If you have configured more than one cache disk, repeat [step 5](#), [step 6](#), [step 7](#), [step 8](#), and [step 9](#) for each cache disk.
11. Click **OK**. The **Settings** dialog box closes.
12. In the **Virtual Machines** list, right-click the name of your Nasuni Filer, then select **Start** from the drop-down list.
Alternatively, in the **Virtual Machines** list, select the name of your Nasuni Filer, then click **Start** in the list of actions on the right side.
The virtual machine restarts. After a significant cache resize, the operations after powering on can take some time, with some performance impact during that period. The progress of the cache resize is visible on the Nasuni Filer user interface.

Using the Disaster Recovery procedure to change the size of the cache disk or copy-on-write (COW) disk

On all virtual platforms, you can change the size of the cache disk and the COW disk using the Disaster Recovery procedure.

Tip: The copy-on-write (COW) disk should generally be at least one-quarter the size of the cache disk. Although the Nasuni Filer issues a warning during startup if the COW disk is too small relative to the cache, the Nasuni Filer does not issue this warning if the COW disk is larger than 250 GB.

First, create a new Nasuni Filer on the platform. The new Nasuni Filer should have the correct size of cache and COW disks.

Note: To add cache space by configuring additional cache disks, contact Nasuni Technical Support.

Next, perform a manual snapshot on the existing Nasuni Filer, then shut down the existing Nasuni Filer. Finally, perform a Disaster Recovery procedure to the new Nasuni Filer.

Using a new COW disk to change the size of the COW disk

On all virtual platforms, you can change the size of the COW disk by creating a new COW disk.

Tip: The copy-on-write (COW) disk should generally be at least one-quarter the size of the cache disk. Although the Nasuni Filer issues a warning during startup if the COW disk is too small relative to the cache, the Nasuni Filer does not issue this warning if the COW disk is larger than 250 GB.

First, perform a manual snapshot on the existing Nasuni Filer, then shut down the existing Nasuni Filer.

Next, on all platforms except the Microsoft Azure platform, delete the existing COW disk.

Then, create a new COW disk with the correct size.

Finally, restart the Nasuni Filer.

On the Microsoft Azure platform, the largest unused disk becomes the new COW disk. Note that this requires assistance from Nasuni Support for configuration.

Copyright © 2010-2018 Nasuni Corporation. All rights reserved.

Nasuni Corporation | One Marina Park Drive, Boston, MA 02210 | 1.857.444.8500 | www.nasuni.com