

## Snapshot Processing

### Nasuni Corporation – Natick, MA

#### Overview

The purpose of this document is to explain the processing that occurs during snapshots by the Nasuni Filer.

#### Snapshots

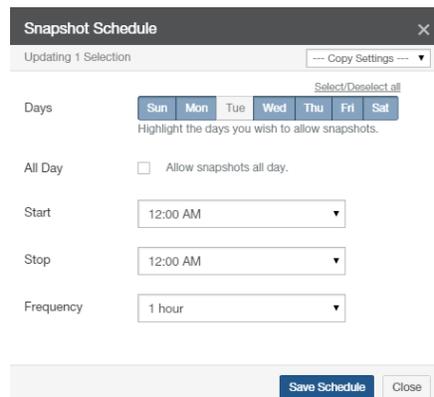
The snapshot process is the act of creating a new version in the UniFS file system.

Snapshots can be triggered under the following circumstances:

- Scheduled snapshot interval: Configurable interval.
- Filer-forced snapshot: When low on available space for writes (cache is > 85% full).
- Customer-requested snapshot via User Interface: **Take Snapshot Now** button.
- Customer-requested snapshot via VSS request.
- Customer-requested when shutting down a Nasuni Filer.

#### Snapshot scheduling

Snapshots are scheduled in the Nasuni Filer UI or in the Nasuni Management Console (NMC) UI.



The screenshot shows a 'Snapshot Schedule' dialog box with the following fields and options:

- Days: A row of buttons for Sun, Mon, Tue, Wed, Thu, Fri, and Sat. The 'Wed' button is highlighted.
- All Day: A checkbox labeled 'Allow snapshots all day.' which is currently unchecked.
- Start: A dropdown menu set to '12:00 AM'.
- Stop: A dropdown menu set to '12:00 AM'.
- Frequency: A dropdown menu set to '1 hour'.
- Buttons: 'Save Schedule' and 'Close' at the bottom right.

Figure 1: Snapshot Schedule in NMC.

Snapshots can be scheduled for any or all days of the week. Snapshots can be scheduled to occur either throughout the day (“**Allow snapshots all day**”) or between specified start and stop times on each selected day.



During each specified snapshot window, you can set the time between snapshots (Frequency). The available Frequency values are different for local volumes (minimum: 1 **hour** between snapshots) and for shared volumes (minimum: 1 **minute** between snapshots).

During the allowed snapshot window, a snapshot is scheduled based on the selected Frequency. If a snapshot is still running when the next snapshot interval occurs, the next snapshot is queued to run immediately after the running one completes. Only one additional snapshot is ever queued to run.

For example, suppose that a snapshot schedule has a Frequency of every 15 minutes, but that each snapshot takes more than 30 minutes to complete:

- At 1:05, a snapshot (#1) is scheduled. Since no other snapshot is running, snapshot 1 begins.
- At ~1:20, snapshot #1 is still running, but another snapshot (#2) is scheduled. Snapshot #2 begins running after snapshot #1 completes.
- At ~1:35, if snapshot #1 is still running, snapshot #2 is already queued, and so no additional snapshots are scheduled.

If a snapshot has been queued to run, or has already started running, it continues to be queued to run (if queued) or to run (if running), **even if the snapshot window has ended**. However, no further snapshots are scheduled outside of the snapshot window.

### **Local volume (remote access disabled)**

For a purely local volume, which does not have Remote Access enabled, the snapshot processing proceeds as follows:

1. At the snapshot frequency interval, a snapshot is scheduled.
2. When the snapshot begins, the Nasuni Filer checks for all unprotected data (namely new or changed data) on the volume.
3. If Antivirus is enabled, an antivirus scan is performed.
4. The Nasuni Filer sends all unprotected data (phase 1) and metadata (phase 2) to the cloud. If there is a lot of data or a low Quality of Service or a low bandwidth from the Nasuni Filer to the cloud, this could possibly take a long time.
5. If the next snapshot interval comes up while the current snapshot is in progress, another snapshot is scheduled, but only one. This new scheduled snapshot does not start until the previous one completes, as well as any other scheduled snapshots for other volumes that might be in the queue also complete.



## **Shared or Remote volume (Remote Access enabled – even if volume is not connected to any other Nasuni Filers)**

For a volume that has Remote Access enabled, even if that volume is not connected to any other Nasuni Filers, the snapshot processing proceeds as follows:

1. At the snapshot frequency interval, a snapshot is scheduled. Note that the snapshot frequency has some randomness built in, so that not all Nasuni Filers are asking for the lock at exactly the same time for the same volume.
2. When the snapshot begins phase 1 (unprotected data), the Nasuni Filer tries to get the snapshot lock for this volume from the NOC.
3. If it succeeds, the snapshot lock is held very briefly to obtain a new snapshot version number, and then the Nasuni Filer releases the snapshot lock.
4. Check with the NOC to see if the latest version of the data is the same as ours. Merge if necessary.
5. The Nasuni Filer then proceeds to queue up unprotected data and send it to the cloud.
6. At the end of 10 minutes (default, but is configurable), the Nasuni Filer stops queuing up data, and then finishes sending the rest of the queued data to the cloud. This 10-minute time limit means that not all unprotected data may be protected in a single snapshot.
7. Now, the Nasuni Filer tries to get the snapshot lock for this volume from the NOC for phase 2 (metadata).
8. If the Nasuni Filer succeeds in getting the snapshot lock, the Nasuni Filer sends all of the metadata that corresponds to the data processed during phase 1 to the cloud.
9. After all of the metadata is protected, the Nasuni Filer releases the snapshot lock for this volume. This generally takes a few minutes, but, depending on the depth or bushiness of the data protected, it can take significantly longer (hours).

If a Nasuni Filer cannot get the snapshot lock from the NOC for either phase, the number of retries is based on the specific snapshot interval set:

- a. If the snapshot interval is less than 5 minutes: skip the snapshot and wait until the next scheduled snapshot.
- b. 5-minute interval: 3 retries
- c. 10-minute interval: 5 retries
- d. 15-minute interval: 7 retries
- e. 25-minute interval or longer: 11 retries

Because there is only one snapshot lock per volume in the NOC, the more Nasuni Filers that are connected to the volume in Read/Write mode, the more possible contention there might be for this lock. Nasuni Filers in Read-Only mode make no changes, so do not add to any possible contention.

For example, suppose that 10 Nasuni Filers share a volume, and all of them are attempting snapshots every 15 minutes. If the metadata phase takes 1 minute for each snapshot, then most or all of the Nasuni Filers should complete a snapshot during that 15 minutes. However, this is not guaranteed. If one Nasuni Filer has a

lot of changes, and holds the lock for 5 minutes, some of the other Nasuni Filers might not complete their snapshots during that interval.

As a second example, suppose that 40 Nasuni Filers are each attempting snapshots every 15 minutes. Again, if the metadata phase takes 1 minute each, there is really no possibility of having all the Nasuni Filers complete their snapshots within that interval. In fact, with that many Nasuni Filers seeking the snapshot lock, between initial tries and retries, there might be a lot of lock contention. In this case, many of the Nasuni Filers take a lot longer to complete a snapshot, in both phases.

### **Folders on volumes with Global File Locking enabled**

For a volume that has Global File Locking enabled, the snapshot processing proceeds as follows:

1. When a file that has Global File Locking enabled is opened, the Nasuni Filer goes to the lockserver to obtain a lock on the file. (These are different locks from the snapshot locks.)
2. When the file is saved or closes, the files are protected in the cloud outside of the regular snapshots. This allows a user on another Nasuni Filer to open the file, go to the lockserver to get the lock for the file, and obtain (sync) the latest version of the file from the cloud. This guarantees that the user is always working on the latest version of the file.
3. However, metadata for the file is not wired into the cloud version of the file until the regular, phase 2, snapshot occurs for the volume. Because of this arrangement:
  - a. New files aren't seen on remote Nasuni Filers until phase 2 of the regular snapshot completes, and the remote Nasuni Filer runs a sync.
  - b. Timestamps of the files under Global Lock are not updated until phase 2 of the regular snapshot completes, and the remote Nasuni Filer runs a sync. Because a Global Lock-enabled file always syncs on lock and open of the file, it opens the latest version, even if the timestamp doesn't reflect this.
  - c. The notes above, related to retries for the snapshot lock in the NOC and possible contention, still apply for these phase 2 snapshots.
4. Any other files or folders on the volume, in folders that do not have Global Lock enabled, snap under the regular snapshot schedule (see above).