

Side Load Feature

Nasuni Corporation – Boston, MA

Overview

When Nasuni first supported the Disaster Recovery (DR) process, it was intended to be used to recover from true disasters such as hardware failure or buildings falling down. However, the reality is that, much of the time, the recovery process is actually done to upgrade from one piece of hardware to another or, in the less usual case, upgrade between very significant changes in Nasuni's underlying operating system.

For recoveries in similar situations, the challenge is that the old Nasuni Filer contains a populated cache while the new Nasuni Filer will have an empty cache. The user then has degraded performance (more cache misses) and needs to use a lot of inbound bandwidth to re-populate the cache on the new Nasuni Filer over the following hours, days, or weeks.

The Side Load feature was created to help address this issue by “side loading” the data that is resident in the cache of the old Nasuni Filer to the new Nasuni Filer's cache.

Note: *You can only perform the Side Load procedure as part of the entire Recovery procedure. See the [Recovery](#) document for complete details. In particular, before performing the Side Load procedure, you must perform the procedure “Preparing the Original Source Nasuni Filer (if available)”, which includes running the `preparedr` command on the console of the original source Nasuni Filer, as described in the Recovery document.*

Also, before performing the Side Load procedure, you must perform the procedure “Recovering the Nasuni Filer”, which includes the original source Nasuni Filer becoming decommissioned, as described in the Recovery document

Tip: *While the Side Load process typically reduces the time necessary to populate the cache of a Nasuni Filer, this is not always the case. For example, situations that involve many small files might not experience significant time savings, due to the overhead of handling the metadata for each file.*

Planned DR

The first part of the Side Load process is for the user to perform a planned recovery. This uses the `preparedr` command on the console. The old Nasuni Filer is set to read-only. A final push of data to the cloud ensures that all the data in the cache has been protected. After the push has completed, the old Nasuni



Filer remains powered on during the Side Load process. The user then powers on the new Nasuni Filer and starts a normal recovery procedure.

Side Load Configuration

After the new Nasuni Filer is online, the user goes to the Side Load page on either the Nasuni Filer UI or the NMC UI. The user provides the IP address of the old Nasuni Filer, and the username and password of an administrative user on the old Nasuni Filer. This username/password is used to authenticate with the old Nasuni Filer's UI.

After authentication, an API key is generated and used for all future side load communication (via HTTPS) between the old and new Nasuni Filers.

The Side Load process checks if the old Nasuni Filer is accessible. The Side Load process also checks whether the volumes on the new Nasuni Filer are present on the old Nasuni Filer.

If there is data on the old Nasuni Filer that is not yet protected, the Side Load process still proceeds. However, data that is not yet protected is not loaded to the new Nasuni Filer, and a warning is displayed.

Also, if the new Nasuni Filer's cache is smaller than the cache of the old Nasuni filer, the Side Load process still proceeds. However, after the cache of the new Nasuni Filer fills, then the Side Load process stops.

Side Load Process

Assuming that all of the above checks are successful, the Side Load process copies the database of the old Nasuni Filer to the new Nasuni Filer.

After that copying process has completed, for every file on the old Nasuni Filer, the Side Load process retrieves the directory metadata from the cloud for the file's parent directory. If the file manifest is still present in the directory listing, the Side Load process copies the file by connecting to the old Nasuni Filer via HTTPS and requesting the data for the file.

If only some of the chunks of the file are located in the cache on the old Nasuni Filer, the Side Load process copies only those chunks.

The Side Load process copies only file data from the old Nasuni Filer. The Side Load process retrieves all additional metadata on files, such as ownership and ACLs, from the cloud. If the file data has changed in any way, so that there is a new file manifest handle, the Side Load process copies the data from the cloud and not from the old Nasuni Filer.

Older data from the old Nasuni Filer does not overwrite newer data in the new Nasuni Filer. The progress of the Side Load process is displayed. If the Side Load process is interrupted, when it restarts, it continues where it left off. If the Side Load process fails, the user receives a notification. The user can cancel a running Side Load process by clicking the cancel button on the Side Load page.



Also, while Side Load is primarily a background process that iterates through a list of files from the database and copies the data to the new Nasuni Filer, Side Loading data on demand is also possible, if a request for a file comes in before it is brought into the cache. In this case, the process first checks the old Nasuni Filer for the data, and then copies data chunks from the cloud that are not present on the old Nasuni Filer.

Example Use Case

A client would like to upgrade from a NF200 to a NF400 to increase the performance of their Nasuni Filer. After swapping in the new hardware and performing a planned disaster recovery, the Side Load process (once properly configured) proactively populates the new Nasuni Filer's cache from the old Nasuni Filer rather than copying data from the cloud.

Operations Details

Nasuni supports a Recovery process (see the [Nasuni Filer Administration Guide](#)) that enables you to recover the Nasuni Filer after a true disaster, such as the loss of a data center. However, most of the time, clients perform the Recovery process in order to upgrade from one piece of hardware to another.

In such a situation, there is a working Nasuni Filer in their data center that contains active data in the cache. Performing the Recovery process results in a new Nasuni Filer that has an empty cache. The client often then re-populates the new cache with data, which can require considerable inbound bandwidth from the cloud, and which can take days, weeks, or even months to complete.

The Side Load feature enables you to transfer cache data directly from the source Nasuni Filer to the new Nasuni Filer. The source Nasuni Filer must already be decommissioned.

Tip: Only one Side Load process is permitted at a time for each Nasuni Filer.

Tip: Only the Admin user can perform the Side Load process.

Tip: The source Nasuni Filer must be:

- *Running;*
- *Decommissioned;*
- *Using release 7.0 or above.*

Warning: If the total data available on the source Nasuni Filer exceeds the available cache space on the new Nasuni Filer, you receive a warning message and notification.

You can configure Notifications to notify you by email when the Side Load process completes.



Starting the Side Load process

To start the Side Load process from the source Nasuni Filer to the current Nasuni Filer, follow these steps:

1. Click **Services**, then select **Side Load** from the list. The **Side Load** page appears.

Side Load

Perform a "Side Load" to populate Filer cache from a healthy/decommissioned Filer (post DR). Decommissioned Filer must be of version 7.0 or greater.

Decommissioned Filer

Host Address

Authenticate using the decommissioned Filer's credentials.

Admin Username

Admin Password

Connect and Start

Figure 10-34: SideLoad page.

Note: If this Nasuni Filer is under Nasuni Management Console control, this page is not available on the Nasuni Filer. Instead, use the Nasuni Management Console to view information or perform actions.

2. In the **Host Address** text box, enter the host address of the source Nasuni Filer. The source Nasuni Filer must already be decommissioned.
3. In the **Username** text box, enter the username for the specified source Nasuni Filer. The source Nasuni Filer must already be decommissioned.
4. In the **Password** text box, enter the password for the specified Username for the specified source Nasuni Filer. The source Nasuni Filer must already be decommissioned.
5. Click **Connect and Start**.
A connection is established with the data of the source Nasuni Filer. Data begins moving to the current Nasuni Filer. After the data transfer starts, you can view the progress of the Side Load process.



- When the Side Load process completes, the **Complete** label appears on the bar graph.

Tip: Record any information you want to retain from the screen before clicking Done.

Click **Done**.

Viewing and controlling the Side Load process

To view or control the progress of the Side Load process, follow these steps:

- Click **Services**, then select **Side Load** from the list. If the Side Load process has not yet completed, the **Side Load - Status** page appears.

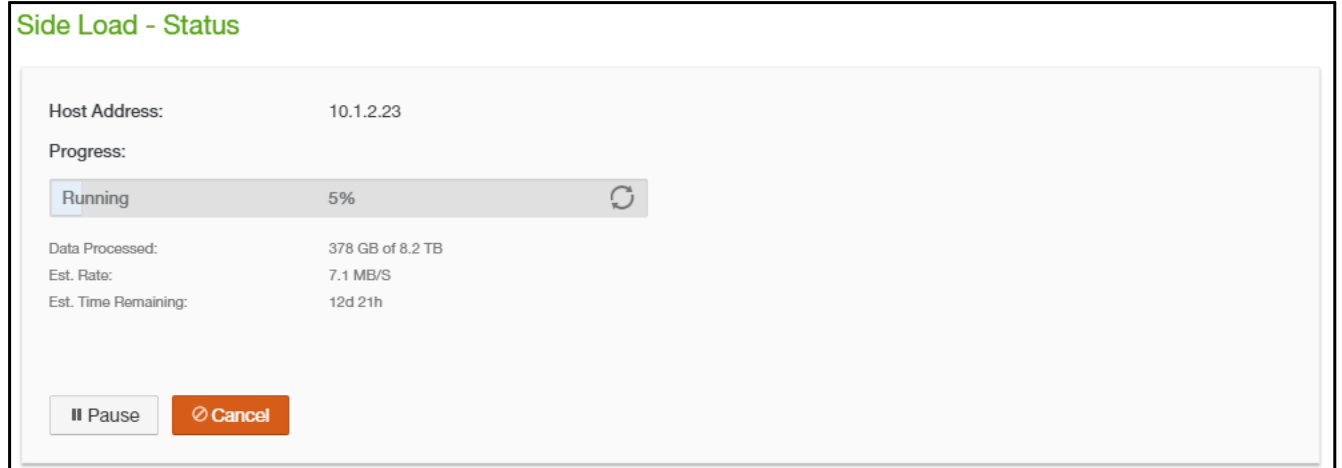


Figure 10-35: Side Load - Status page.

Note: If this Nasuni Filer is under Nasuni Management Console control, this page is not available on the Nasuni Filer. Instead, use the Nasuni Management Console to view information or perform actions.

The following information appears:

- **Host Address:** The host address of the source Nasuni Filer. The source Nasuni Filer must already be decommissioned.
- **Progress:** A bar graph indicating the progress of the Side Load process. The percentage of the Side Load process that is complete appears. If the Side Load process is running, the **Running** label appears. If the Side Load process is paused, the **Paused** label appears.
- **Data Processed:** The amount of data processed (in KB, MB, GB, or TB) and the total amount of data to process (in KB, MB, GB, or TB).



- **Est. Rate:** The estimated rate of data transfer (in KB/S, MB/S, GB/S, or TB/S).
 - **Est. Time Remaining:** The estimated time until the Side Load process is complete.
2. To pause a running Side Load process, click **Pause**. The Side Load process pauses indefinitely.
The bar graph label changes to **Paused**.

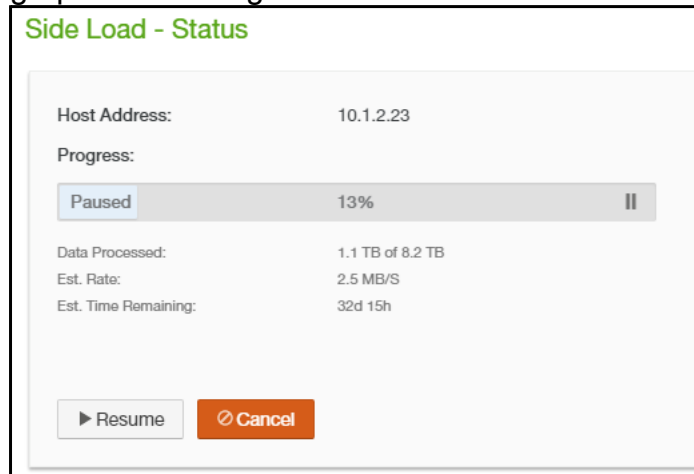


Figure 10-36: Side Load - Status page, paused.

To continue with the Side Load process after a pause, click **Resume**. The Side Load process continues. The bar graph label changes to **Running**.

3. To cancel the Side Load process, click **Cancel**. The **Cancel Side Load** dialog box appears.




Figure 10-37: Cancel Side Load dialog box.

To cancel the Side Load process, click **Continue**.



If the Side Load process is canceled, the bar graph label changes to **Canceled**.

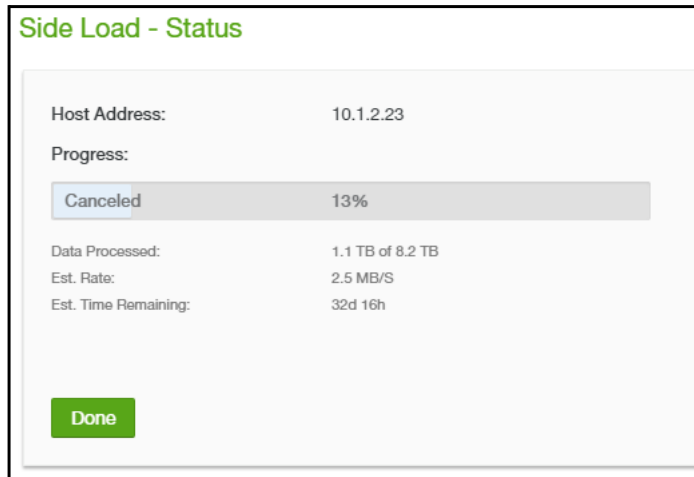


Figure 10-38: Side Load status page, canceled.

Tip: Record any information you want to retain from the screen before clicking Done.

4. When the Side Load process completes, the **Complete** label appears on the bar graph.

Tip: Record any information you want to retain from the screen before clicking Done.