

## Data Migration Service Best Practices

### Nasuni Corporation – Boston, MA

#### Data Migration Best Practice Recommendations

In order to use Nasuni, it is necessary to first migrate data to the Nasuni Filer, and then to transport data from the Nasuni Filer to cloud storage.

Nasuni recommends migrating data into Nasuni using the Data Migration Service built into the Nasuni Filer. The best practices for setting up and managing the migrations are covered in the **Data Migration Setup** section. Because each Nasuni client has a different quantity of data and a different available bandwidth, Nasuni has the following guidelines to ensure a fast and efficient data migration.

For data sets of less than 1 TB, and with an available bandwidth of 10 to 20 Mbps, Nasuni typically recommends migrating data through the Nasuni Filer and the existing WAN connection. This is on a per-site basis. As a rule of thumb, 1 TB of data migrated through a 20 Mbps WAN connection takes a minimum of 7 days, and typically longer due to Quality of Service limitations during working hours.

For data sets larger than 1 TB, with restricted bandwidth (20 Mbps or less), Nasuni recommends using the Bulk Data Load Service to transport data to cloud storage. With the Bulk Data Load Service, you attach a Nasuni-provided external storage device to a standard USB port, load data onto the external storage device, and then send the external storage device back to Nasuni, where your data is loaded directly into cloud storage.

In most cases, if the bandwidth is limited and the data migration would exceed either 7 days or the requirements to go live, the Bulk Data Load Service is the preferred method for transporting data to cloud storage. If time is not critical, then estimates for the time to migrate data can be made and, if acceptable, migration through the WAN can be accomplished.

Nasuni also recommends using a Nasuni Filer Hardware Appliance with the Bulk Data Load Service, because it provides a larger cache and a dedicated appliance for the process.

For migrations through a VM-based Nasuni Filer, Nasuni recommends setting the cache in the Nasuni Filer performing the migration to the maximum size possible, compared to the data set, during the migration. This allows the migration to proceed at the fastest possible rate. For example, in the case of a 1.5 TB data set to be migrated through a VM Nasuni Filer, a 2 TB cache would be optimal.



The time of migrations also typically depends on the composition of the data set. Migrating millions of very small files usually requires more time than migrating fewer, larger files. This is similar to a normal backup process with many small files.

Consult your Nasuni representative for suggestions for your specific migration, because all data sets and client requirements are different.

## Data Migration Setup

For Nasuni Filer v3.1.5 and later, the recommended best practice for migrating files to the Nasuni Filer while cloning permissions is to use the Data Migration Service (DMS).

Recommended practices:

- The volume must have at least one share created.
- The user connecting to the migration source should be a member of the Backup Operator group on the domain and, ideally, on the local source server.
- This user also should be added as a Filer Administrative User. To add a user as a Filer Administrative User, follow this procedure:
  1. Log in to the Nasuni Filer user interface.
  2. Select **Configuration** → **General CIFS Settings**.
  3. Add the user to the '**Administrative User(s)**' field
  4. Enter the appropriate Active Directory **User Name** and **Password** under the **Active Directory Authentication** section.
  5. Click '**Save CIFS Settings**'.

Ideal Configurations for Migration performance:

- Disable Remote Access until data migration is done, to avoid impacting performance.
- Use a Nasuni Filer that is on the same local network as the source.
- Use a Nasuni Filer in a location with the largest upload speeds to the cloud.
- Ensure that the size of the migration can fit into the cache completely.
- Migrate to a local volume (no remote access enabled).
- Temporarily set snapshots for all possible volumes to run once a day.
  - When running a snapshot, data is not written directly to the cache.
  - Snapshots automatically trigger when the cache is sufficiently filled.
- For large, multi-TB datasets, ask about our Bulk Load Service.



## Setting up exclusions with the Data Migration Service

When creating a migration, you have the option to add exclusions. You add each exclusion on its own line. You can exclude specific files and folders by adding the location starting from the targeted folder on the source.

For example, specifying folders like the following:

```
windows  
data/junk  
data/My Pics
```

matches all of the folders from the connected source that points to 'C'. These entire folders do not migrate.

However, specifying folders like the following:

```
/data/videos  
/data/My Docs  
"/data/My Pics"
```

does not match any folders and those folders are migrated.

If you want to exclude files of a certain type, you can use a wildcard pattern such as the following:

```
documents/*.mov
```

This excludes all .mov files in every subfolder of the documents folder.

You can also exclude similar file names using the \* wildcard:

```
documents/taxes*.pdf  
documents/*2010.pdf  
documents/taxes/*.*
```

You can also exclude a single file, similar to the way you would exclude a folder:

```
documents/taxes2009.pdf
```

## Scheduled migrations fail to run

If you configure multiple migrations, they do not run at the same time. A migration runs at its scheduled time only if no other migration is running. If two migrations are scheduled for exactly the same time, one migration starts and the other starts when the first one finishes. If the start time for one migration occurs while another migration is running, the migration that hasn't started is skipped until its next scheduled start time. To help prevent one migration from interfering with another, schedule them farther apart in time.



## Known issues with the Data Migration Service

There are three known types of files that cannot be read or copied from the source, and that cause the Data Migration Service (DMS) to fail to copy the files. After a certain number of errors failing to copy files, the migration itself goes into a Failed state.

These file types are:

- Windows offline files.
- Archive stub files that require an ancillary program for retrieving the file data. There is no problem with compressed files such as .tar and .zip.
- Shortcuts or symbolic links (UNIX).