

Merge Conflicts Best Practices

Nasuni Corporation – Boston, MA

Introduction

You can synchronize data (“sync” or “merge”) from Nasuni, merging data from your Nasuni Filer with any new or changed data from other Nasuni Filers connected to the same volume. This helps to ensure that everyone in your organization is using the most current data.

Note: *It is impossible to overwrite data in the Nasuni system.*

You can select which days of the week on which to sync data. You can also select at what time of day to start and stop syncing data, and at what frequency. You can configure synchronization using either the Nasuni Filer or Nasuni Management Console user interface.

Note: *If you have directories with tens of thousands of files but few changes during each snapshot, or large files that require multiple snapshots, frequent syncs can increase the system load significantly.*

In addition to scheduled syncs, data is also merged, as necessary, each time a snapshot is performed. Before performing a snapshot on a Nasuni Filer, Nasuni checks to see if any shared data from other Nasuni Filers should be updated on the local Nasuni Filer. Therefore, if snapshots occur frequently, then frequent syncs might not be necessary.

Types of Merge Conflicts

During a synchronization or merge, locally changed data is not overwritten with data from other Nasuni Filers connected to the remote volume. Several different types of merge conflicts are possible.

Data Conflict

A data conflict occurs when two users of two different Nasuni Filers make different changes to their local versions of the same file on the same folder path, during the same time period between syncs or snapshot-related syncs. For example, suppose that the file `Report.doc` is on the same path on both Nasuni Filer X and Nasuni Filer Y. A user on X changes their version of `Report.doc` while a user on Y also changes their version of `Report.doc`. If the user on Y had not made any changes, Nasuni would simply have updated the Y version of `Report.doc` with the changes that the user on X made. However, if the user on Y also made changes to `Report.doc`, then Nasuni recognizes that it cannot simply update the Y version, because that would destroy the changes that the user on Y made.



File Name Conflict

A file name conflict occurs when two users of two different Nasuni Filers give different files the same name on the same folder path. For example, if a user on Nasuni Filer X creates (or renames) a file and names it `Invoice.doc`, while a user on Y also creates (or renames) another file and names it `Invoice.doc` on the same path. If the user on Y had not named their file the same as the file on X, Nasuni would simply have added the X file called `Invoice.doc` to Nasuni Filer Y. However, if the user on Y also named a file `Invoice.doc`, then Nasuni recognizes that it cannot simply add the X file, because that would replace the file that the user on Y named.

Metadata Conflict

A metadata conflict occurs when the metadata of the same file on the same folder path of two different Nasuni Filers is changed, during the same time period between syncs or snapshot-related syncs. For example, the permissions on the file might be changed. A user on X changes the permissions on their version of the file `Presentation.ppt` while a user on Y also changes the permissions on their version of `Presentation.ppt`. If the user on Y had not made any changes, Nasuni would simply have updated the metadata of the Y version of `Presentation.ppt` with the changes that the user on X made. However, if the user on Y also made changes to the metadata of `Presentation.ppt`, then Nasuni recognizes that it cannot simply update the metadata of the Y version, because that would destroy the metadata changes that the user on Y made.

Directory Name Conflict

Directories can only have a name conflict. What would be a data conflict in the case of files is handled by the directory merge process. A directory name conflict occurs when matching cloud and local directory names have different origins at the time of merging. This includes directories that have not been committed to a snapshot yet.

A directory name conflict occurs in these situations, after any necessary snapshots and synchronizations occur:

- At the same path, one Nasuni Filer creates DIR1, then another Nasuni Filer creates DIR1.
- At the same path, one Nasuni Filer renames DIRN to DIR1, then another Nasuni Filer creates DIR1.
- At the same path, one Nasuni Filer renames DIRN to DIR1, then another Nasuni Filer renames DIRM to DIR1.



How Nasuni Handles Merge Conflicts

When Nasuni recognizes that a data conflict, a file name conflict, a metadata conflict, or a directory name conflict has occurred, it performs several actions:

- The local file or directory is always retained unchanged.

Note: *The two Nasuni Filers do not sync at the same time. One performs its sync first, and its local changes are preserved, then the other performs its sync, and its local changes are also preserved.*

- An Error Notification is made on both of the Nasuni Filers to alert users that a merge conflict has occurred. Administrators can receive Notifications by email. The Notification looks like this:

```
Merge process completed for volume remote_ny_files  
with conflicts, view the logs in  
remote_ny_files/.nasuni/sync_logs for details.
```

You can search for messages about merges. Type “Merge” in the Search text box.

- Information about the merge conflict is written to a merge conflict text file. The file name of the merge conflict log file is the date in GMT format. These files are located in the `.nasuni\sync_logs` directory for the volume, and under the name of the Nasuni Filer that encountered the merge conflict.

Note: *By default, only a Filer Administrator has Read permission for the topmost `.nasuni` folder. In addition, with Windows, it is necessary to be able to see hidden files.*

For example, this file:

```
\ny_files\.nasuni\sync_logs\New_York_Office\2013_08_22  
_17.10.10GMT.csv
```

is a summary of the merge conflicts encountered by the `New_York_Office` Nasuni Filer during a merge on August 22, 2013 at 17:10:10 GMT. These merge conflict files are also synchronized, and you can delete them as needed.

The information in the merge conflict file is in CSV format, in this form:

```
<local timestamp>, <GMT timestamp>, <conflict type>,  
<file name>, <local user>, <local filer>,  
<remote user>, <remote filer>, <merge version>
```

such as:

```
2013-12-01 11:58:54EST, 2013-12-01 16:58:54GMT, data conflict,  
"/folder/Report.doc", "pat", "New_York_Office",  
"jan", "Boston_Office", 74680
```



or:

```
2013-12-01 11:58:54EST, 2013-12-01 16:58:54GMT, name conflict,  
"/folder/Report.doc", "pat", "New_York_Office",  
"jan", "Boston_Office", 74680
```

- The conflicting file or directory is saved with a file name or directory name that includes the conflict information, in this format:

```
<name> (data conflict from  
<domain1>_<username1>@<Filer1> with  
<domain2>_<username2>@<Filer2>).<suffix>
```

<name> is the name of the original file or directory and <suffix> is the suffix of the original file or directory. The conflict information in the new name includes the domain names (truncated to 16 characters), user names (truncated to 16 characters), and names of the Nasuni Filers (truncated to 16 characters) for the two conflicting parties.

Note: *If the volume is connected to Active Directory, the user name information (including domain name) appears in the name of the merge conflict file. If the volume is connected to LDAP Directory Services, the user name information appears in the name of the merge conflict file.*

For example:

```
Report (data conflict from NYdom_pat@NewYork with  
BOSdom_jan@Boston).doc
```

or:

```
Invoice (name conflict from NYdom_pat@NewYork with  
BOSdom_jan@Boston).doc
```



Resolving Data Conflicts

In order to resolve a data conflict, the administrator, the users, or someone else knowledgeable about the files must decide which version of the conflicting file to retain, and remove all other versions. There are three possibilities:

- Retain only the version from Nasuni Filer X. (The changes made by the user on Y are lost.)
- Retain only the version from Nasuni Filer Y. (The changes made by the user on X are lost.)
- Create a new version with the changes from both Nasuni Filer X and Nasuni Filer Y.

For example, the administrator could consult with the two users to determine which version to retain, or the users could determine it themselves. The names of the users, the name of the file, and the location of the file are contained in the entry in the merge conflict text file, and, in truncated form, in the name of the conflicting file.

Resolving File Name or Directory Name Conflicts

In order to resolve a name conflict, the administrator must decide which file or directory retains its original name, and then rename the other file or directory. There are three possibilities:

- Retain the name of the file or directory on Nasuni Filer X. (Change the name of the file or directory on Y.)
- Retain the name of the file or directory on Nasuni Filer Y. (Change the name of the file or directory on X.)
- Change both names.

The administrator should consult with the two users to determine which name to change. The names of the users, the name of the file or directory, and the location of the file or directory are contained in the entry in the merge conflict text file, and, in truncated form, in the name of the conflicting file.

Note: *In some situations, it might be necessary to manually merged that the contents of the both directories. If so, it is important to also pay attention to propagating directory metadata like Quota, GL and ACLs as these may need to be either adjusted or re-applied as needed.*



Resolving Metadata Conflicts

In order to resolve a metadata conflict, the administrator, the users, or someone else knowledgeable about the files must decide which version of the conflicting metadata to retain, and remove all other versions. There are two possibilities:

- Retain only the version from Nasuni Filer X. (The changes made by the user on Y are lost.)
- Retain only the version from Nasuni Filer Y. (The changes made by the user on X are lost.)

For example, the administrator could consult with the two users to determine which version to retain, or the users could determine it themselves. The names of the users, the name of the file, and the location of the file are contained in the entry in the merge conflict text file, and, in truncated form, in the name of the conflicting file.

Additional points

If an administrator deletes a conflict file, the file remains in the latest snapshot if it is needed. However, snapshots might only be kept for a certain amount of time, depending on the snapshot retention policy for the volume. Therefore, the snapshot retention policy should be long enough to permit review of merge conflicts.

The longer the time between synchronizations, or snapshots that cause synchronizations, the more likely it is that merge conflicts will occur. However, if you have directories with tens of thousands of files but few changes during each snapshot, or large files that require multiple snapshots, frequent syncs can increase the system load significantly.

The style of how people work in an organization affects the occurrence of merge conflicts. In environments where people often work collaboratively, more merge conflicts can be expected.