New Features in ibaHD-Server v2.0.0

Author: Jan Asselman

Date: 9th Feb 2017

Table of contents

1 Dongle license

ibaHD-Server-V2 requires a license update. When upgrading the software, the installer will generate a warning in case the license update is missing:

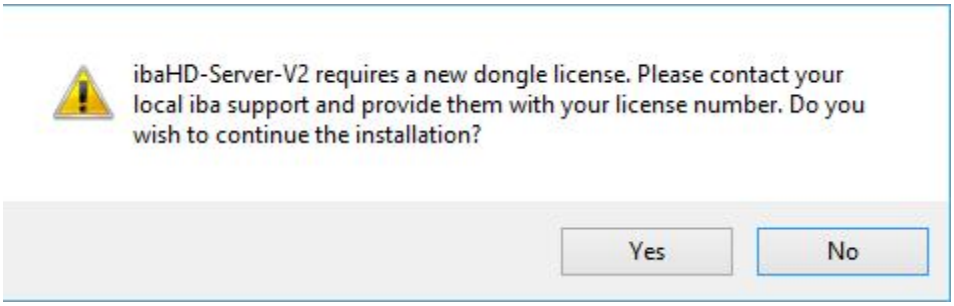

- In case the ibaHD service is started without this dongle license, all configured HD stores will be disabled.

ibaHD-Server-V2 also supports the Extended Update Package (EUP) licensing strategy (see separate document). In a sales region that supports EUP, the installer will display the required information on a license agreement page:

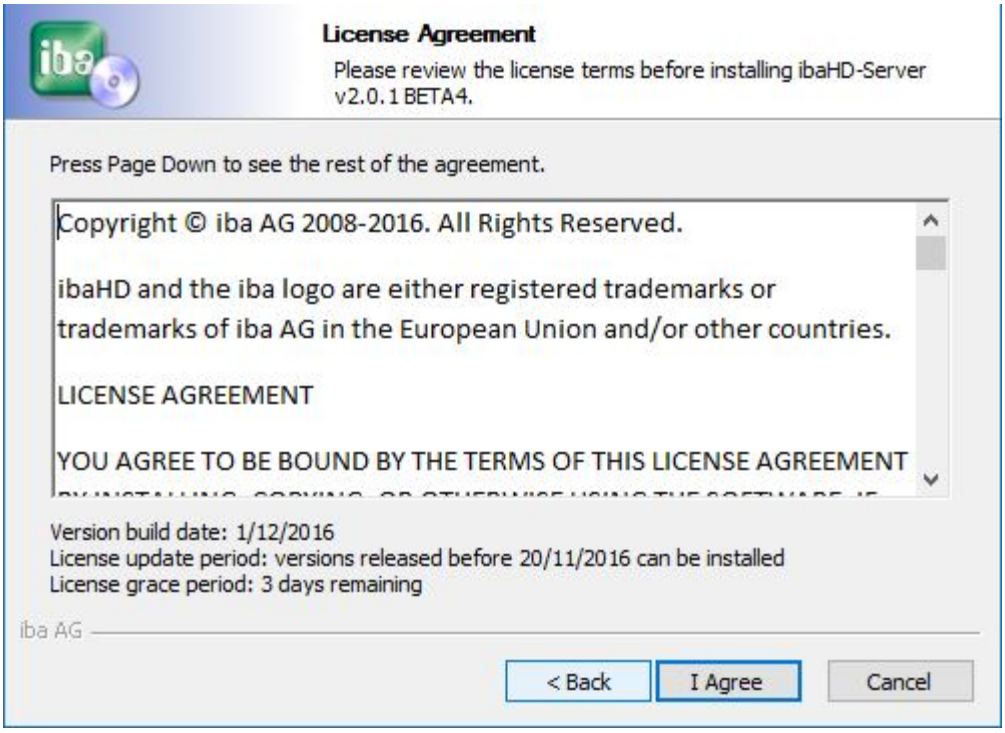

- In case the license update period has expired the installer will generate a warning:

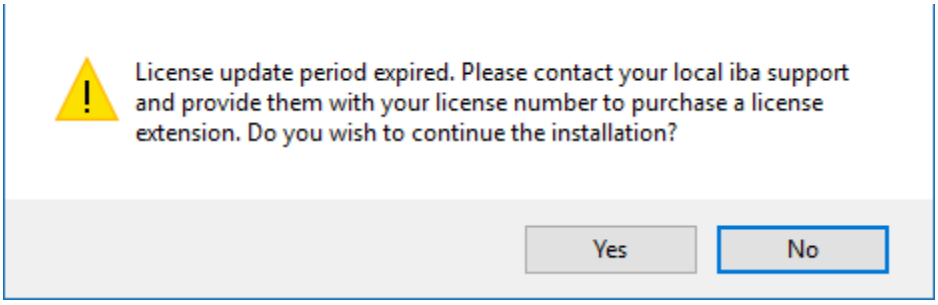

- In case a license grace period is detected the installer will ask to enable it:

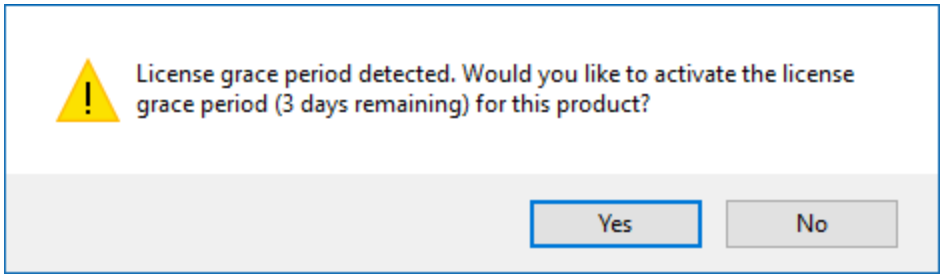

2 Enable/disable store

It is now possible to enable or disable stores using the store configuration dialog.

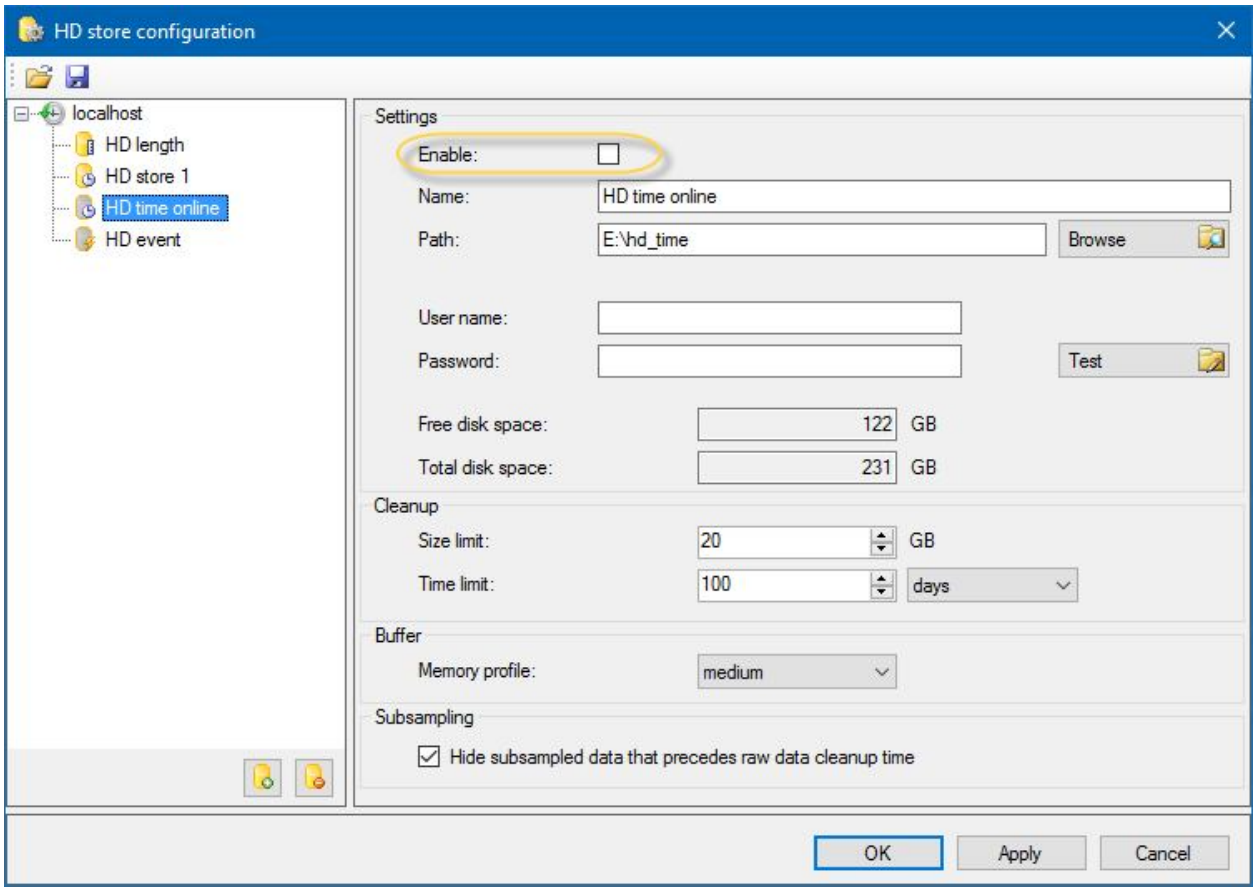

Enabled stores are displayed using a yellow icon. Disabled stores are displayed using a gray icon.

If a store is disabled, read and write access to the store is disabled. And it is not possible to attach or restore backups to the store.

In case there are insufficient licenses available on the iba dongle, stores will now be disabled instead of being removed from the configuration. In case the iba dongle is removed, the disabled stores will automatically be re-enabled when the dongle is inserted.

If PDA client is used to open the store configuration dialog, version 6.37.0 is required to disable stores.

3 Project file

The « General » tab in the HD status application contains two new buttons. One to load and one to save a HD project file.

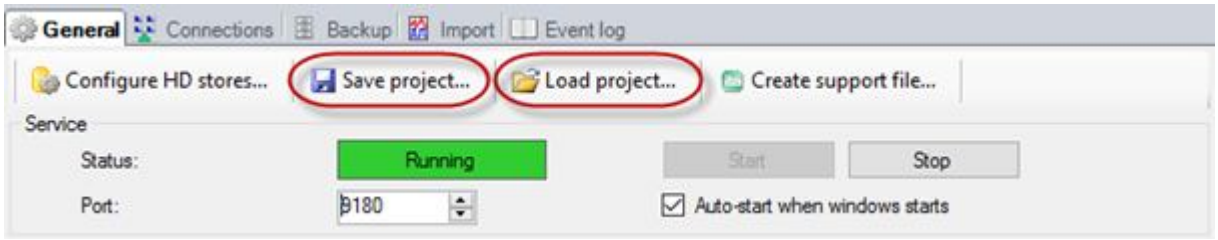

When saving a project you have a couple of options:

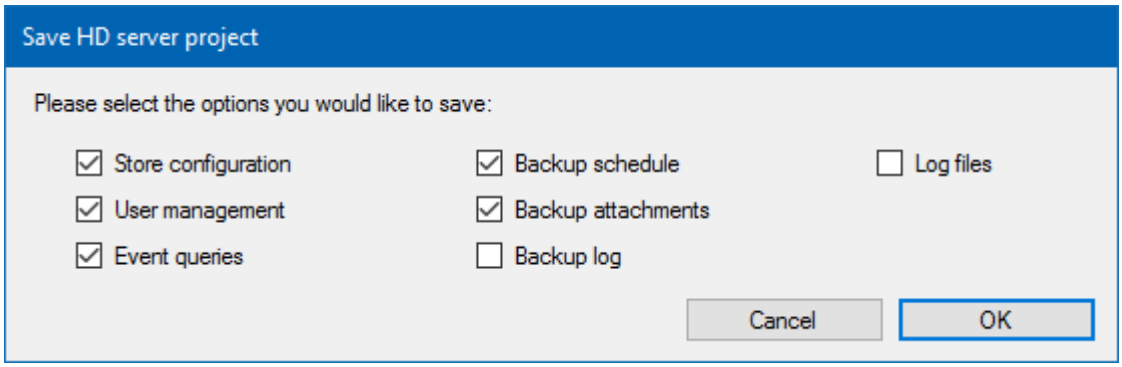

Options that are not included in the project file will be disabled when loading:

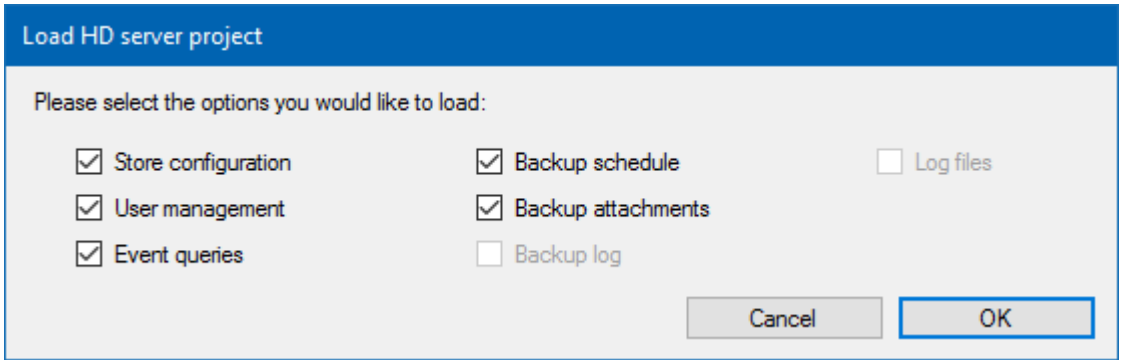

This feature can be useful when moving an HD server to another system.

4 Data file import

An « Import » tab has been added to the HD status application. Click the toolbar button to display the data import wizard.

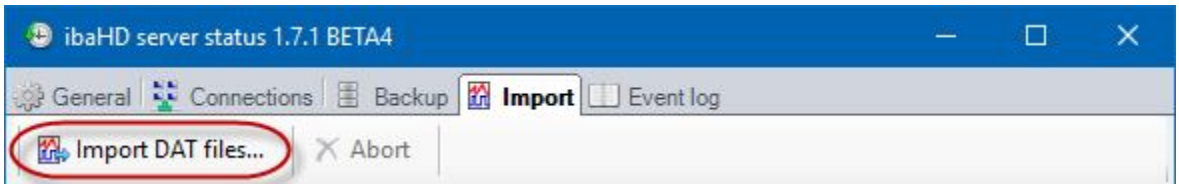

4.1 Source & destination

In the first step, data files and a destination store can be specified.

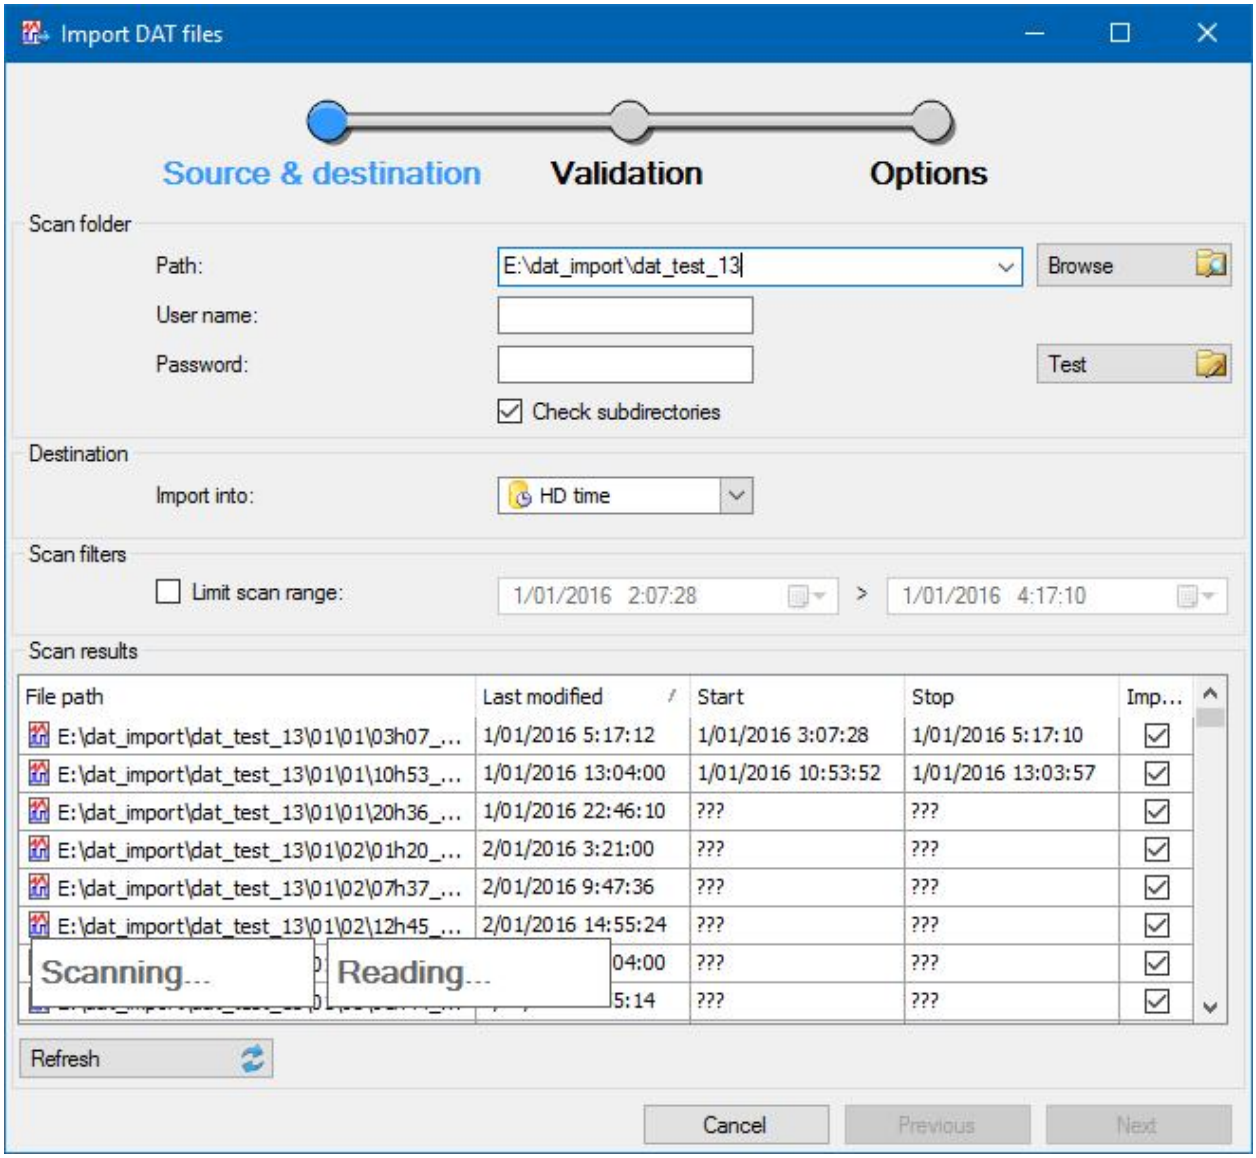

As soon as a scan folder has been entered and a destination store has been selected, scanning starts automatically and scan results will be added to the grid. The « Next » button is disabled as long as scanning is in progress.

When files have been scanned they will automatically be opened to read their data range. Initially the start and stop columns contain three question marks for each of the scanned data files. When a file is read, the start and stop column will receive a proper time stamp.

In general, scanning for data files will be fast, even when the specified scan directory is located in the network. Reading the data range might take more time for files located on a network share.

If you enable the option to limit the scan range, only files inside the specified time range will be scanned.

Subdirectories can be scanned and it is possible to deselect certain files using the checkboxes displayed in the « Import » column.

When the « Next » button is clicked, a dialog might appear reporting the selection of outdated files. A file is outdated if store cleanup would immediately delete the data after importing. Files whose data range has not yet been read, are not considered.

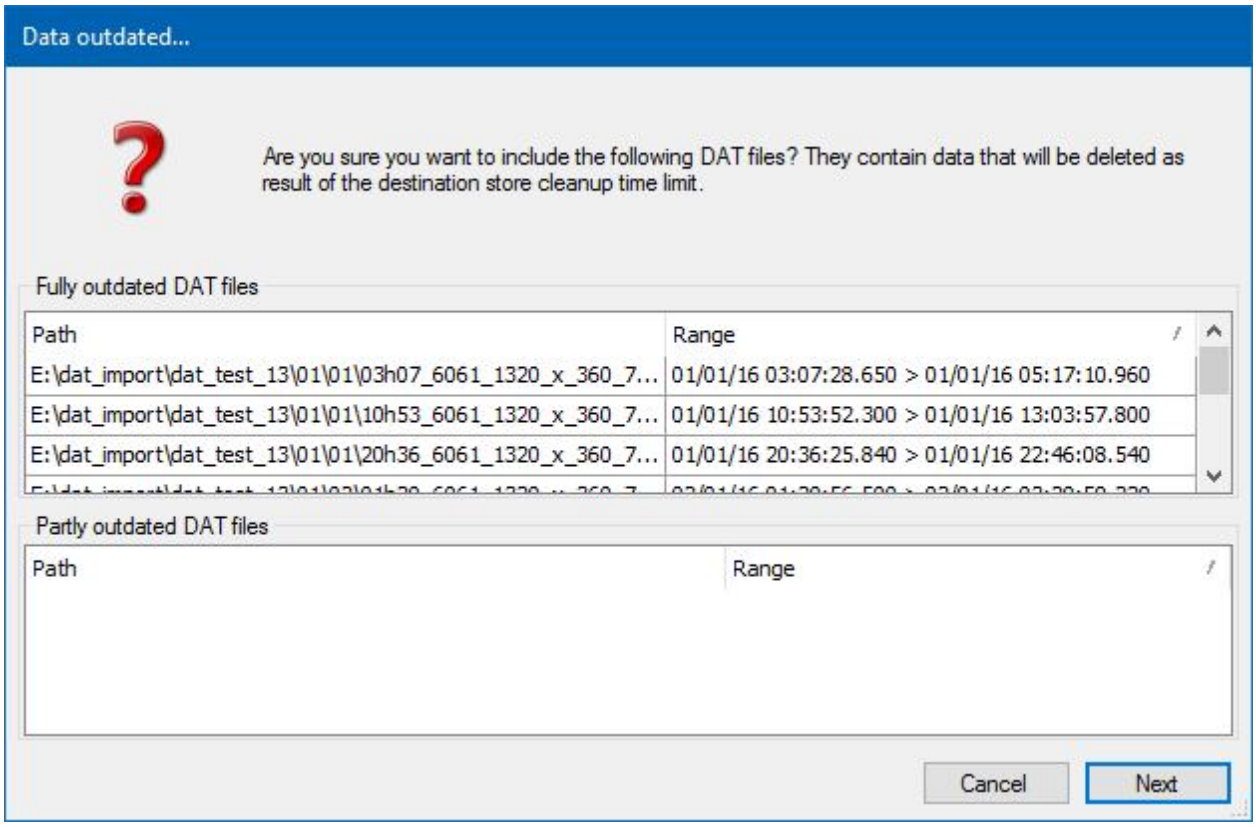

4.2 Validation

In the second step, the user can validate what has been configured in the first step; to make sure that the correct data is imported into the correct store.

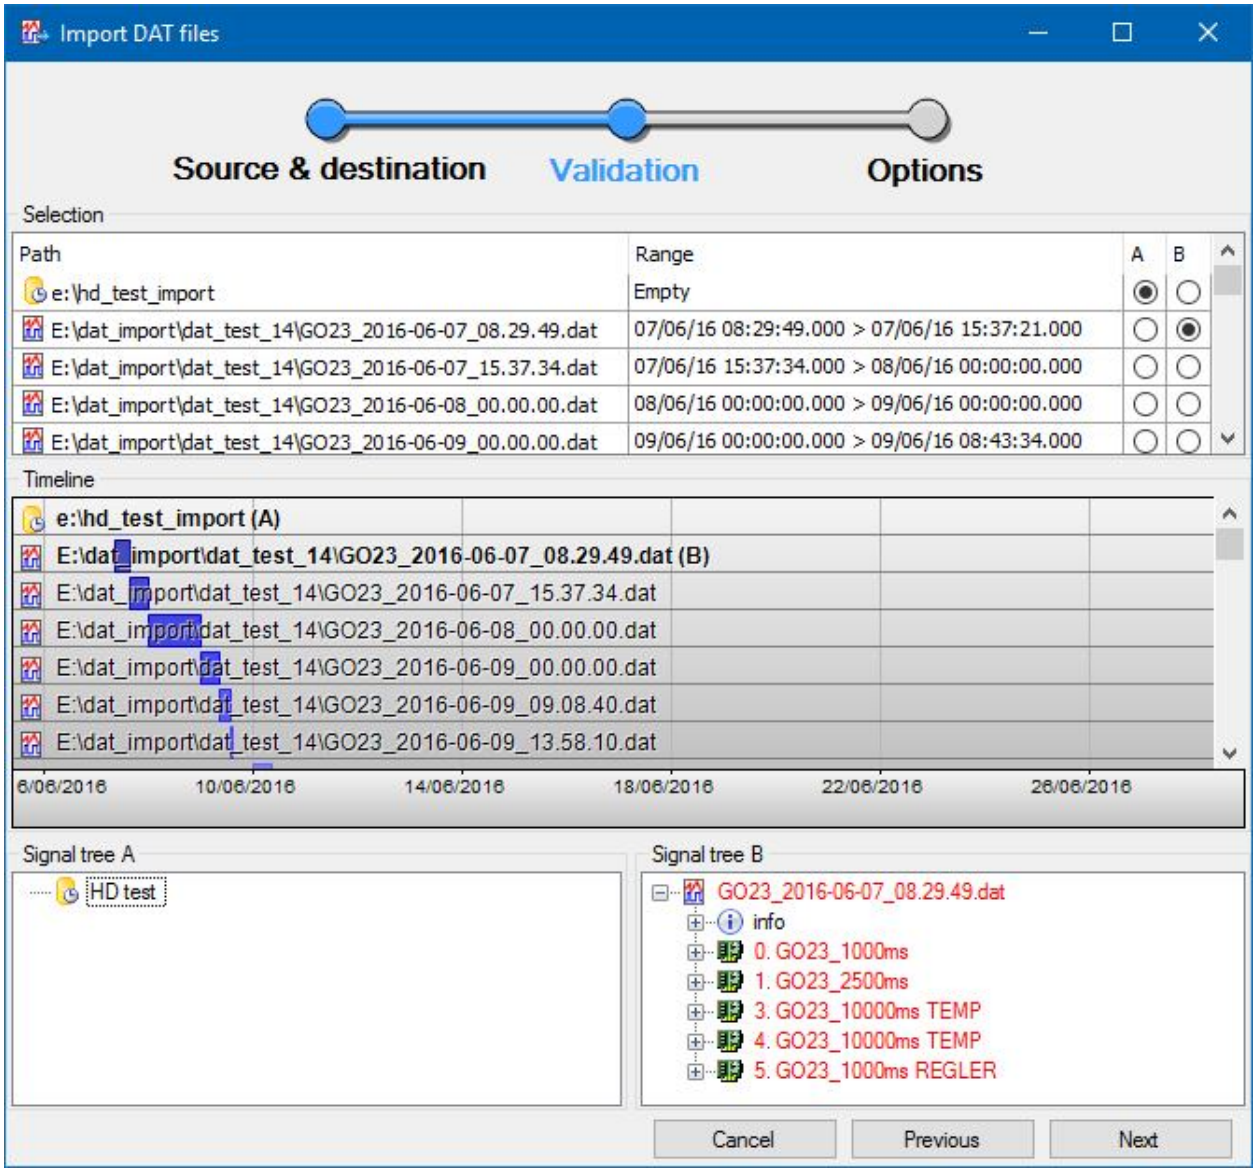

The top grid displays the data files that were selected in the previous wizard step. Files whose data range has not yet been read, are not displayed. The signal trees that are displayed at the bottom can be changed using the radio buttons in column A and B of the top grid. Differences between signal tree A and B are highlighted in red.

The time line visualizes data in the destination store and in the data files. In the image above, we are importing data into an empty store.

4.3 Options

In the third and final step, a couple of options can be enabled.

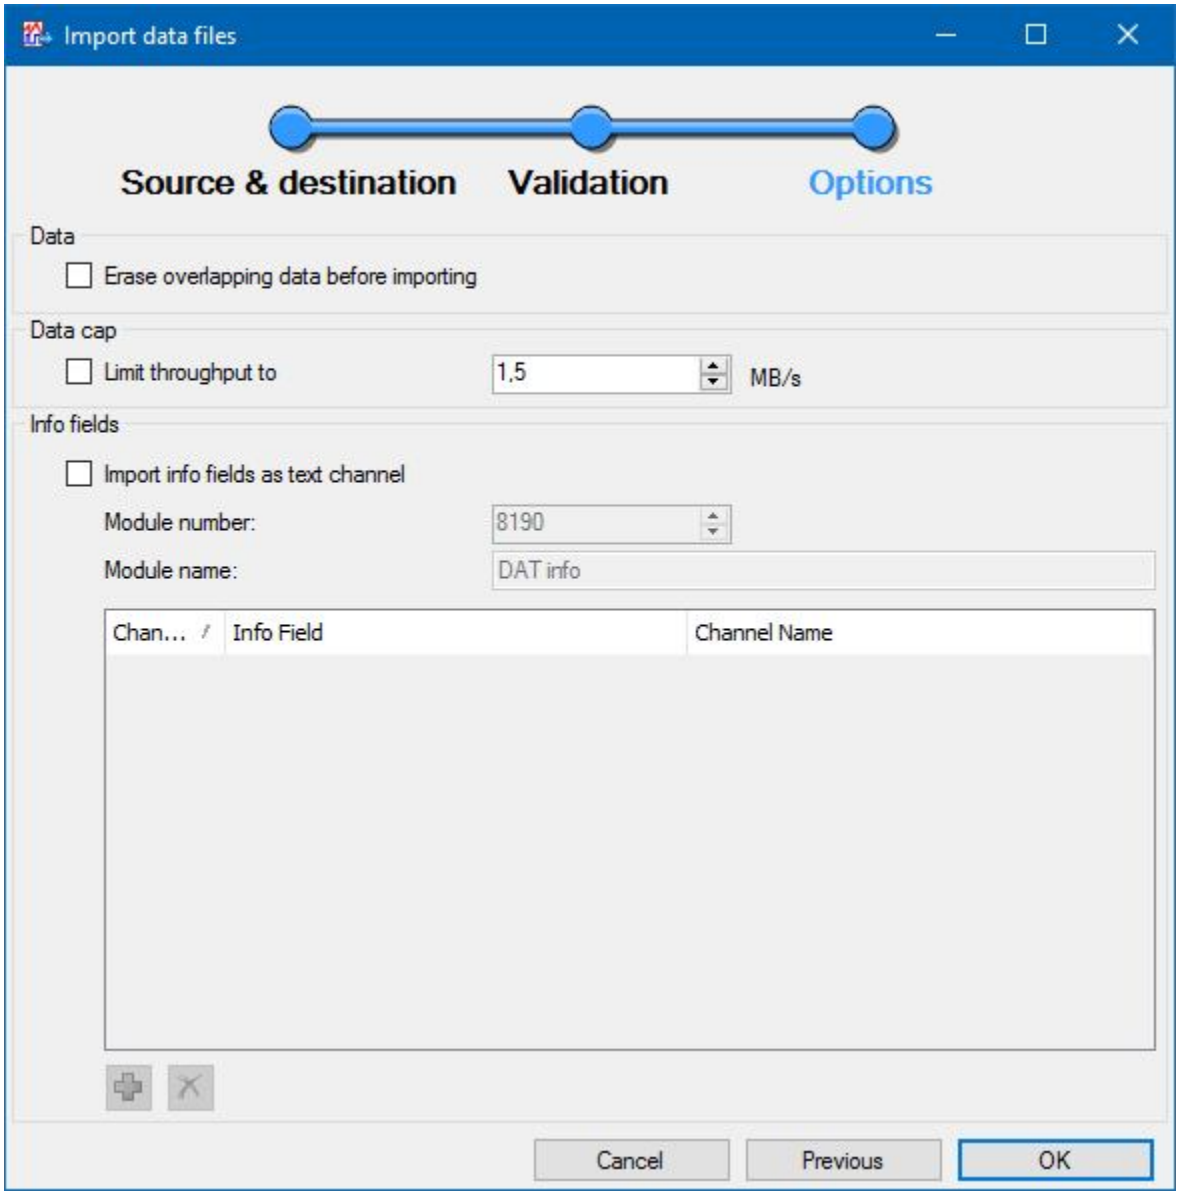

Erase overlapping data before importing
Erase all signal data from the destination store that overlaps with the file data ranges, before importing the data.

Limit throughput
The transfer rate can be limited. This might be important in case the destination store has a concurrent online acquisition.

Import info fields as text channel
Specify info fields that will be read from the data files and will be converted into text channels of the module with specified number and name.

4.4 Progress

Unless the user canceled, the import operation will start when the wizard dialog closes.

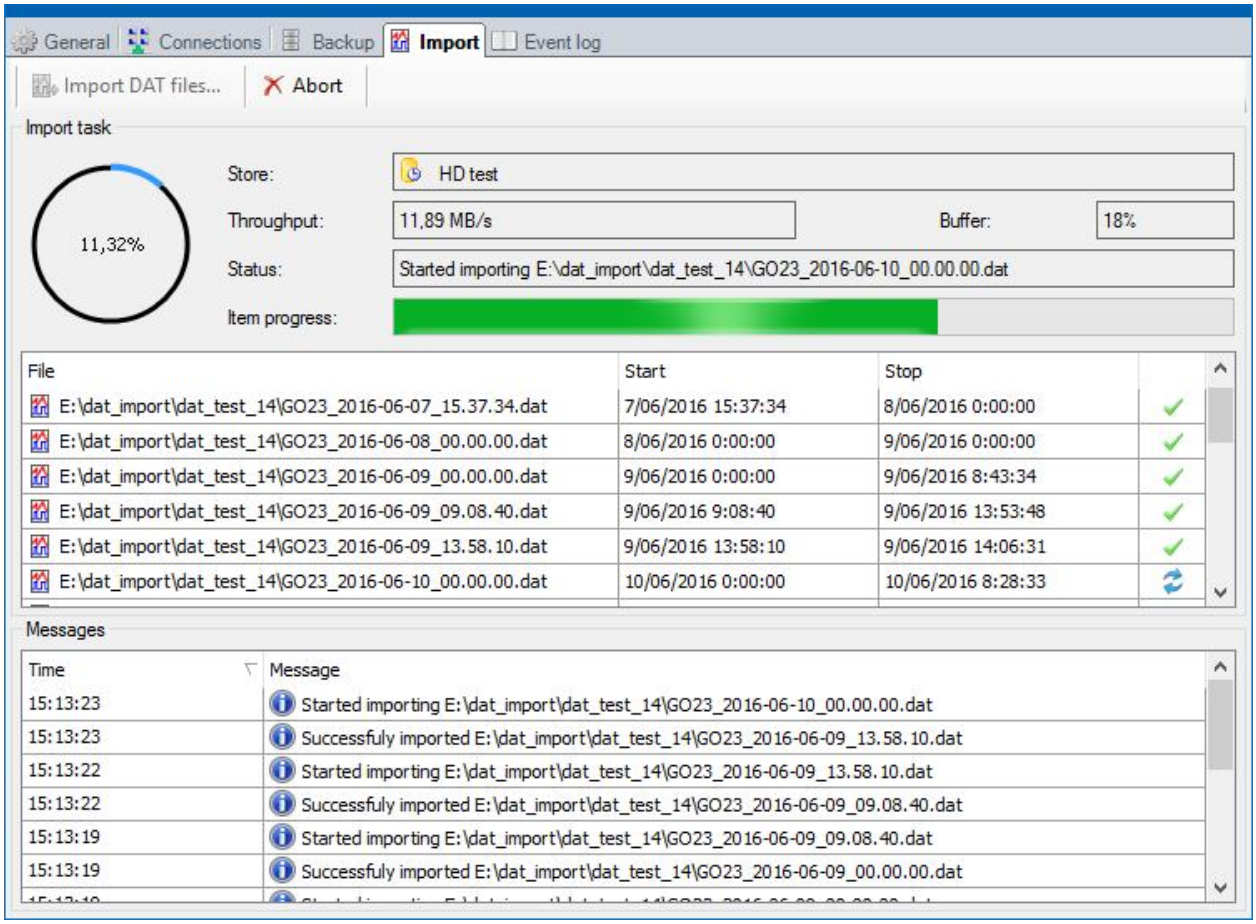

Progress information is displayed at the top and messages are logged in the bottom grid. When the buffer size reaches 100% data transfer will pause until the buffer size has decreased to 50%. The operation can be aborted at all time using the « Abort » button. No rollback occurs when aborting.

Throughput can be influenced by

- configured limit
- internal structure of the files
  - small block sizes will decrease throughput
  - certain data structures are converted on the fly
- time gaps or overlap between consecutive files
  - the HD server will flush its buffer before receiving the data in the next file
- the « Erase overlapping data before importing » option
- overall HD server system load

Imported data is immediately available in all clients. Each HD store allows one online acquisition and one import operation simultaneously. It is planned to allow automated data import using ibaDatCoordinator.

5 Hide subsampled data

A new option has been added to the time based store configuration that hides all subsampled data that precedes the raw cleanup time stamp of the store. This data is hidden for all clients.

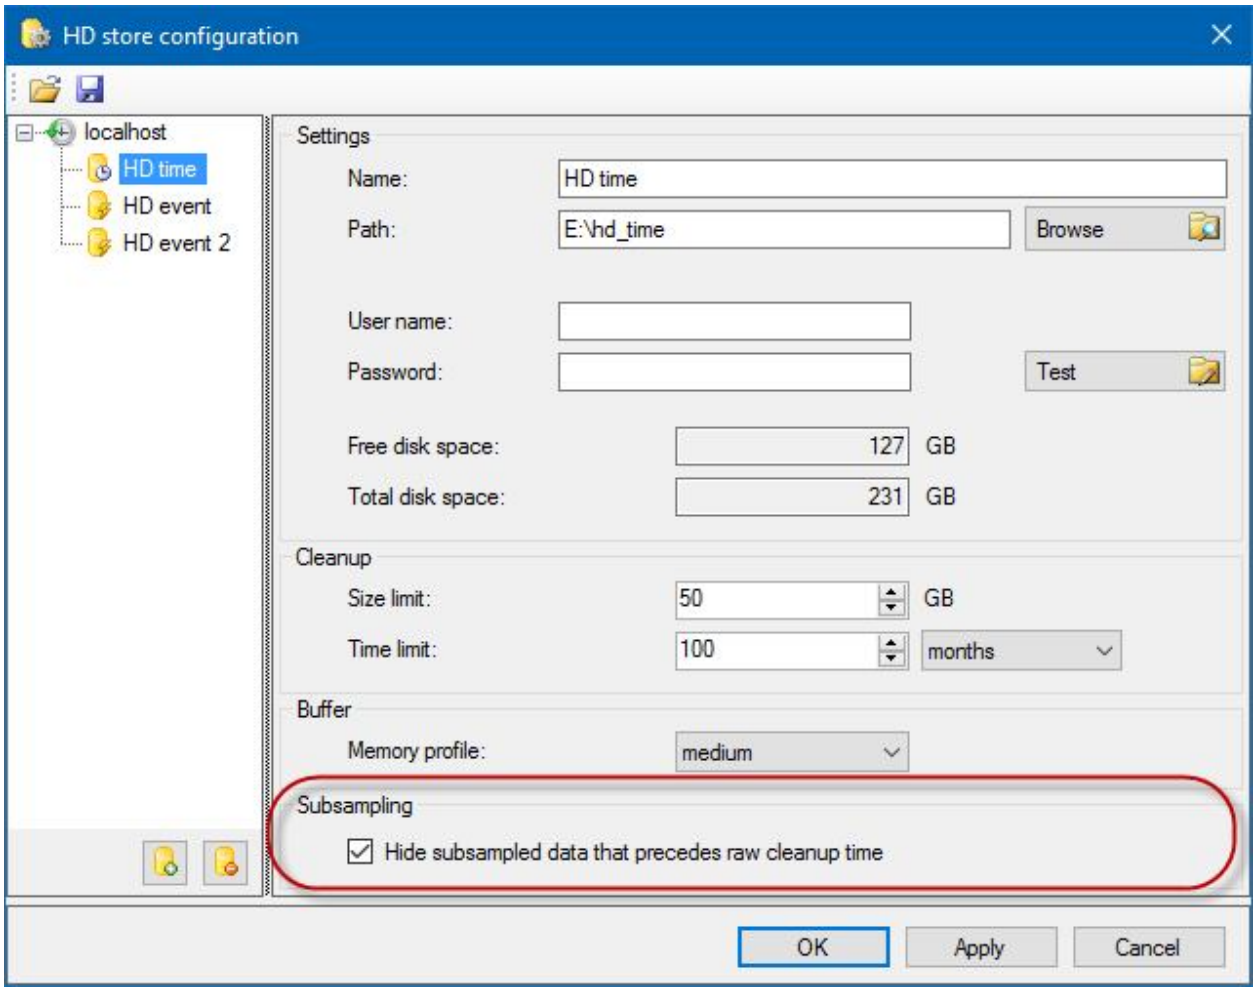

When the time or size limit of a store is reached, the cleanup of the store removes the oldest data in the store. For time based stores, in general, this means that raw data is deleted sooner than sub-sampled data. And 'finer' subsampled data is deleted sooner than 'coarser' subsampled data.

This causes signal data to disappear when zooming into old data in ibaPDA client or when drilling into old data in ibaAnalyzer.

If you enable this new option, then subsampled data will be available only where raw data is available, and thus signal data should never disappear when zooming in or drilling down.

This feature requires the following software versions:

- ibaPDA v6.37.0 or newer
- ibaAnalyzer v6.8.1 or newer

6 Backup

6.1 Mount backups

It is now possible to mount backups.

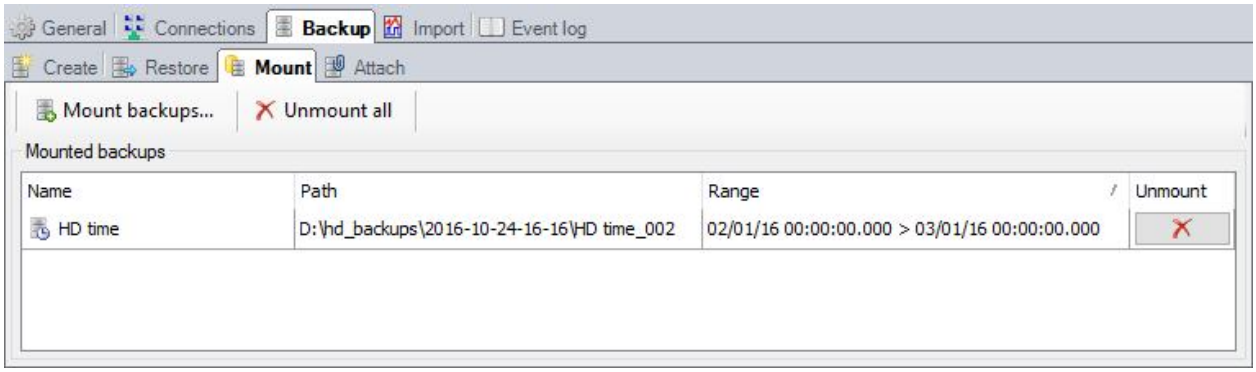

If you mount a backup, it will behave as a HD store. From performance perspective this is the same as configuring a new, empty HD store and attaching a backup to it. But there are two major differences:

- No data can be added to the HD store; it is read-only
- The license counter for the number of configured HD stores is not increased

This enables scenarios where a backup is transferred to a PC without iba dongle for offline data analysis.

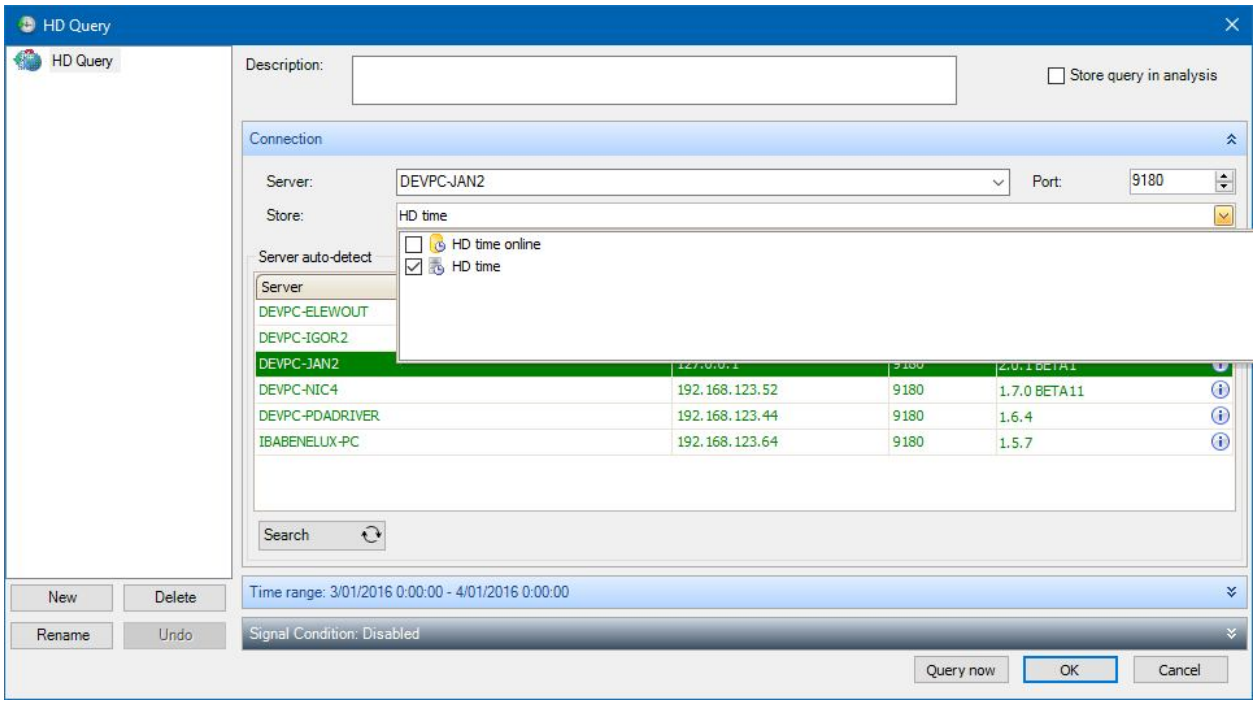

As for normal HD stores it is possible to attach other backups to a mounted backup.

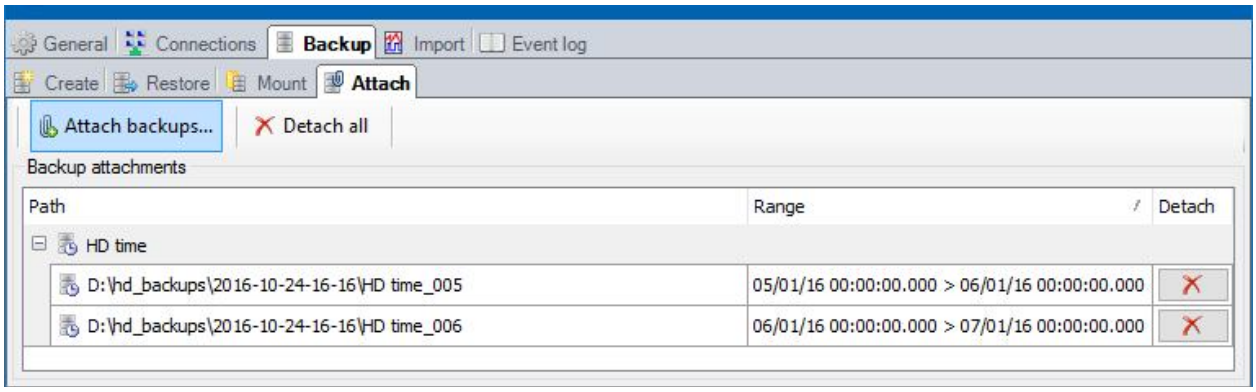

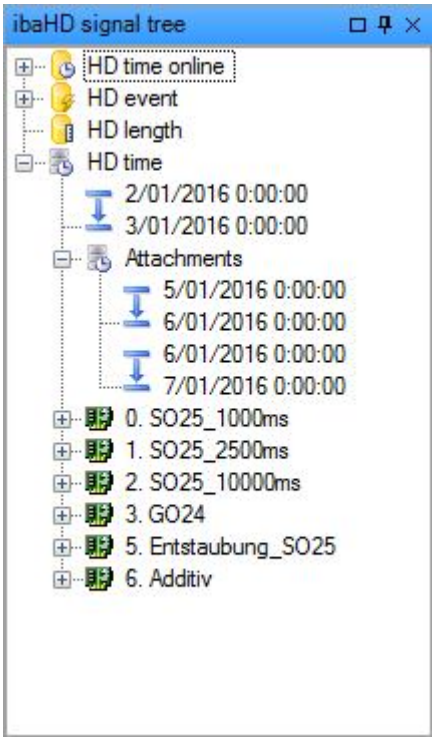

6.2 Split full backup

When creating backups - manually or using the backup scheduler - it is now possible to split a full backup into multiple, smaller backups.

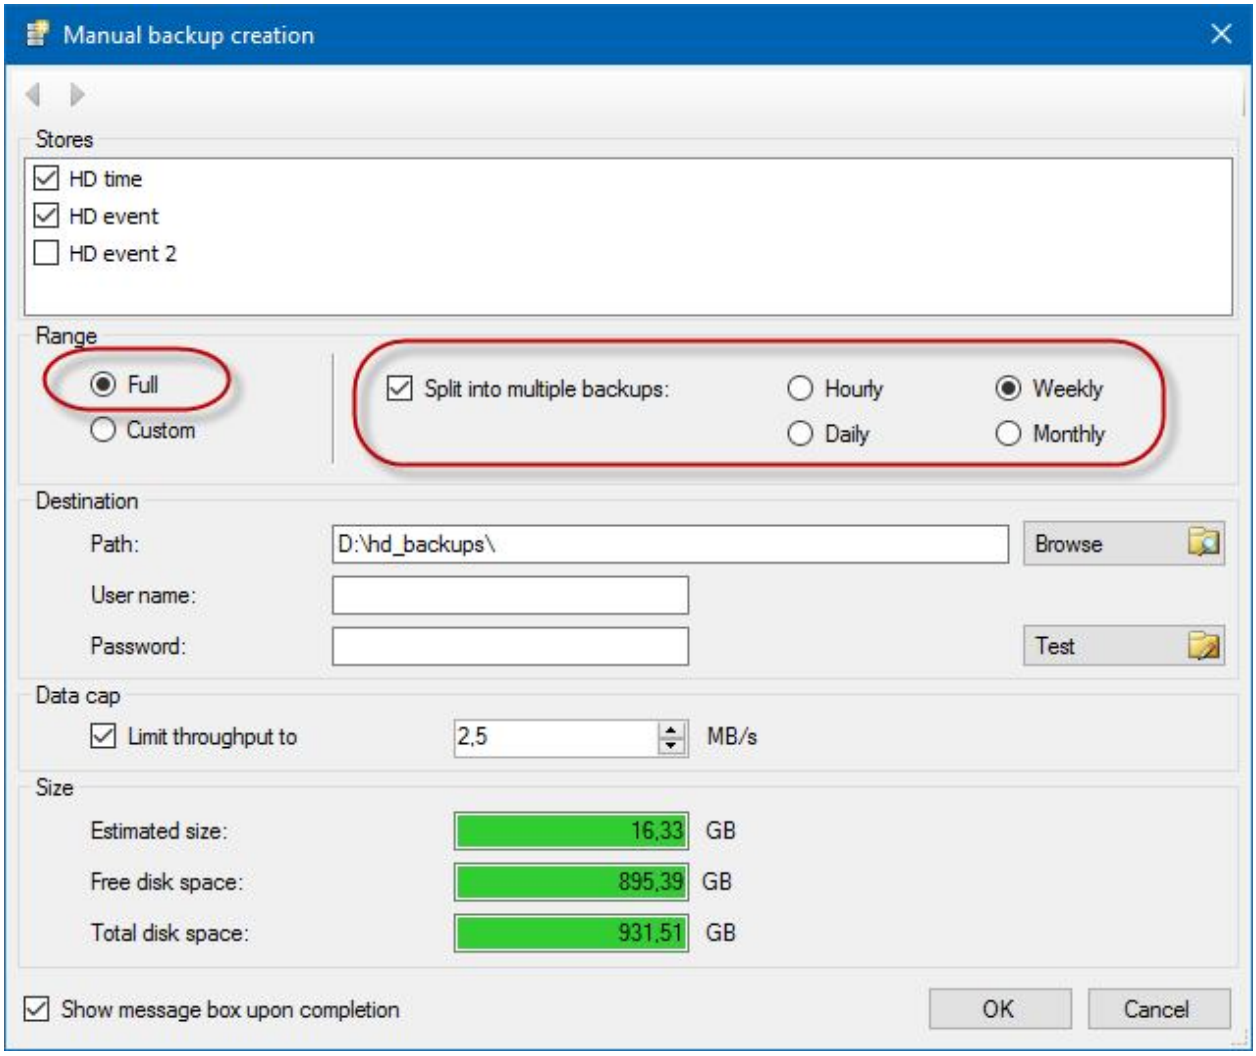

This eases the process of restoring a full backup because you don’t have to execute one single, time-consuming restore task.

Example of « weekly »:

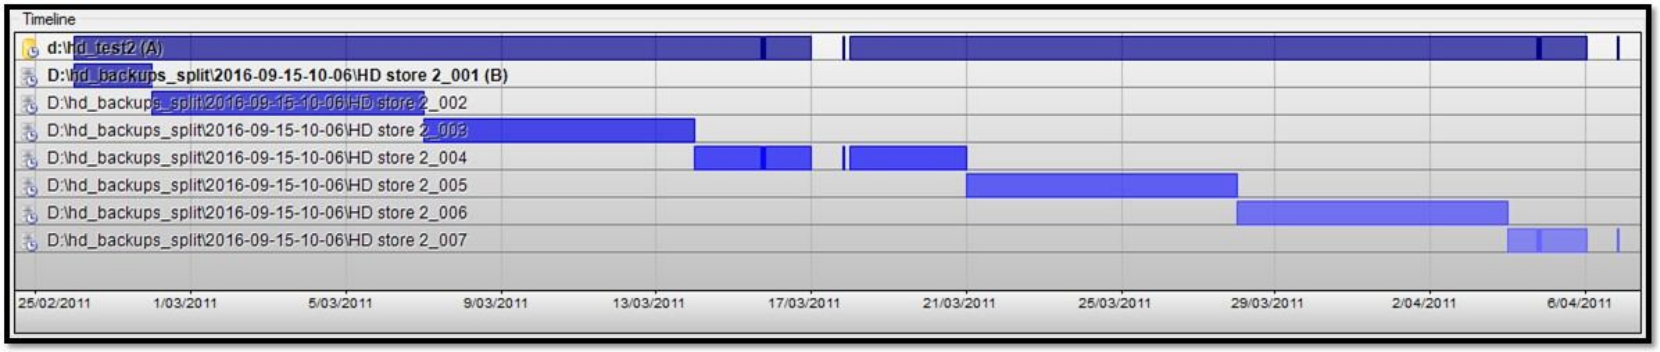

Example of « monthly »:

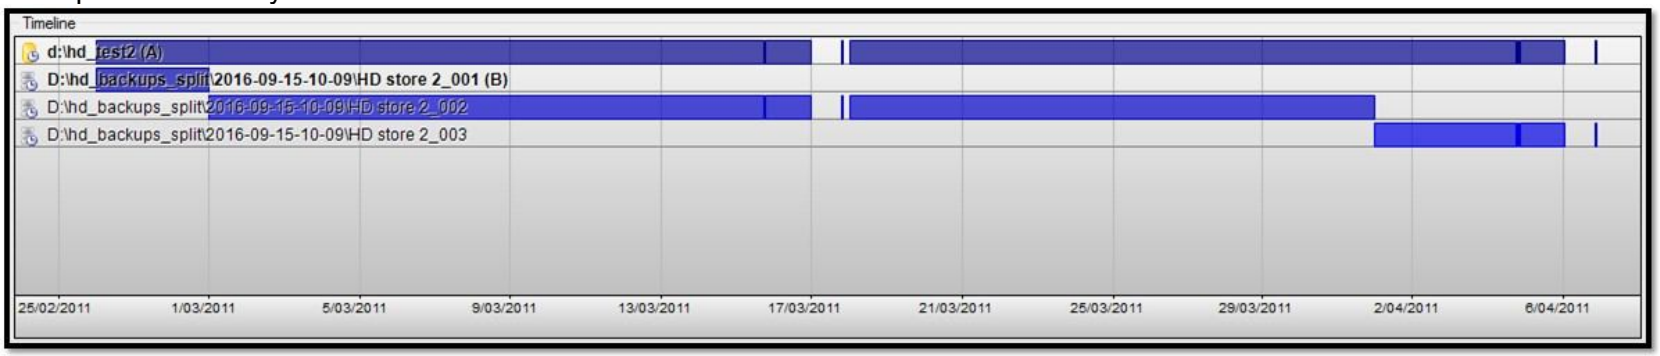

And the system calendar as reference:

	februari 2011						
	MA	DI	WO	DO	VR	ZA	ZO
5	31	1	2	3	4	5	6
6	7	8	9	10	11	12	13
7	14	15	16	17	18	19	20
8	21	22	23	24	25	26	27
9	28	1	2	3	4	5	6
10	7	8	9	10	11	12	13

	maart 2011						
	MA	DI	WO	DO	VR	ZA	ZO
9	28	1	2	3	4	5	6
10	7	8	9	10	11	12	13
11	14	15	16	17	18	19	20
12	21	22	23	24	25	26	27
13	28	29	30	31	1	2	3
14	4	5	6	7	8	9	10

6.3 Data cap

When creating backups - manually or using the backup scheduler - it is now possible to specify a throughput limit.

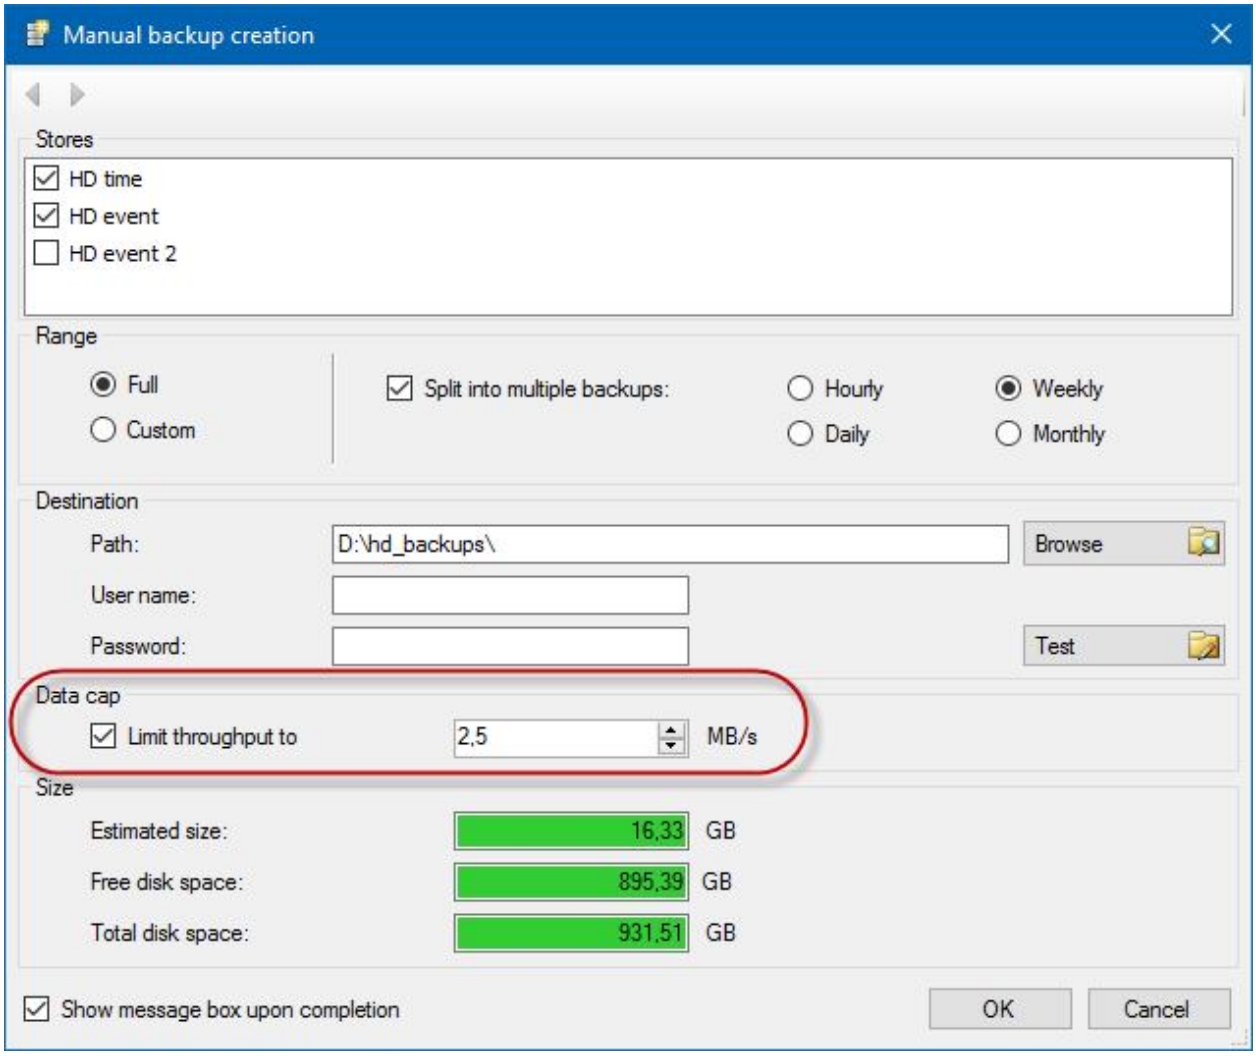

In case the backup destination disk drive also contains one or more HD stores, it can be useful to apply a data cap to the backup task to make sure that the online acquisitions to the HD stores are not influenced.

6.4 Manual backup size estimation

When creating a backup manually, the size of the backup that is about to be created, is now estimated.

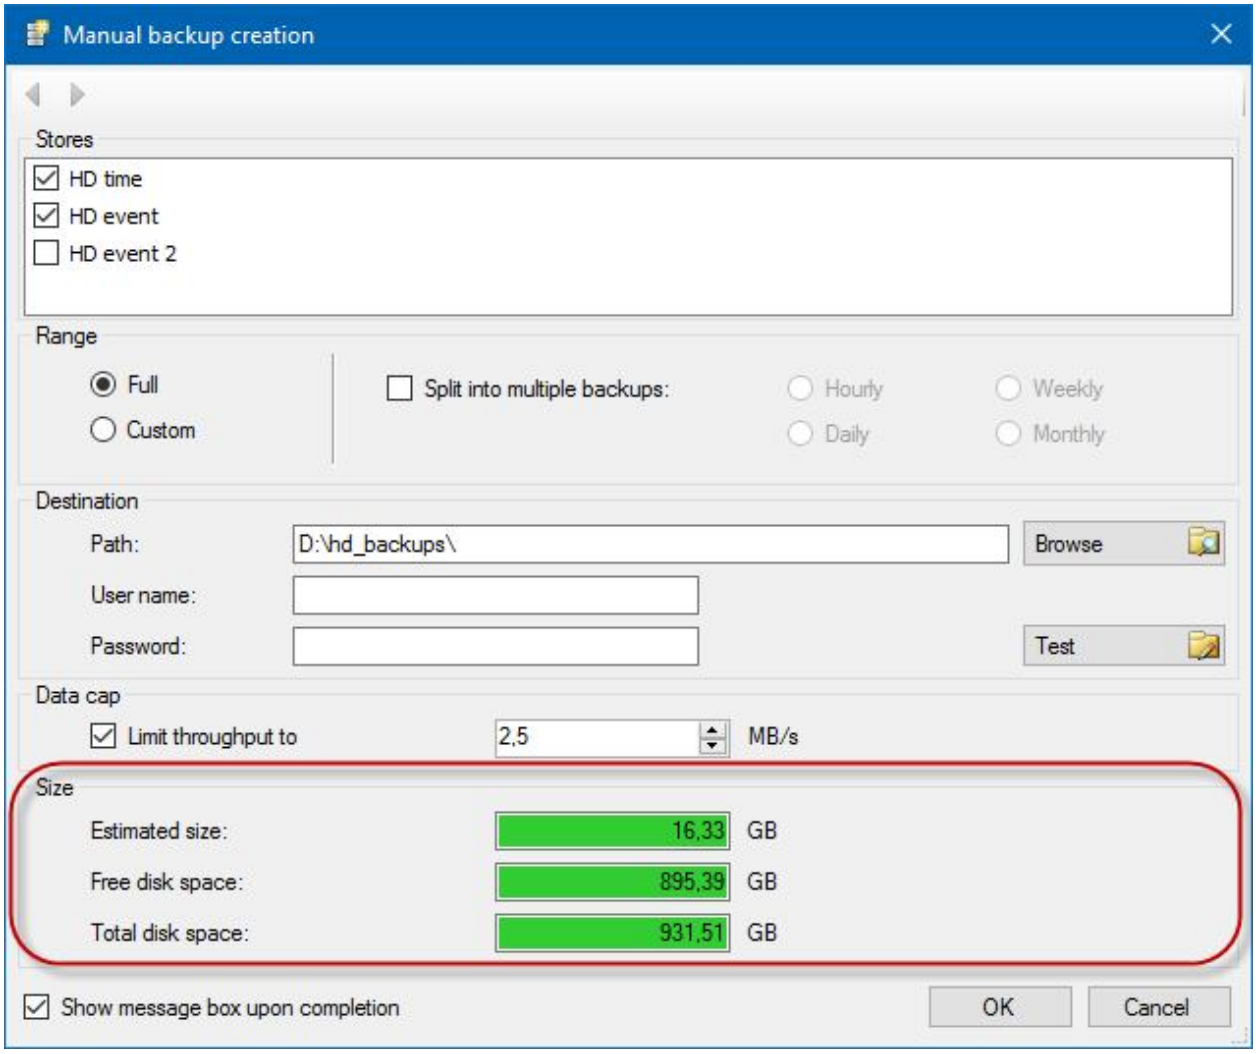

The estimated value is a slight worst case overestimation. The value is automatically adjusted when the selected stores or the backup range are changed. The labels turn red in case insufficient space is available. In case a full backup is split, the overestimated part will grow as the number of splits increase.

6.5 Backup scheduler cleanup

It is now possible to enable cleanup for each task in the backup schedule.

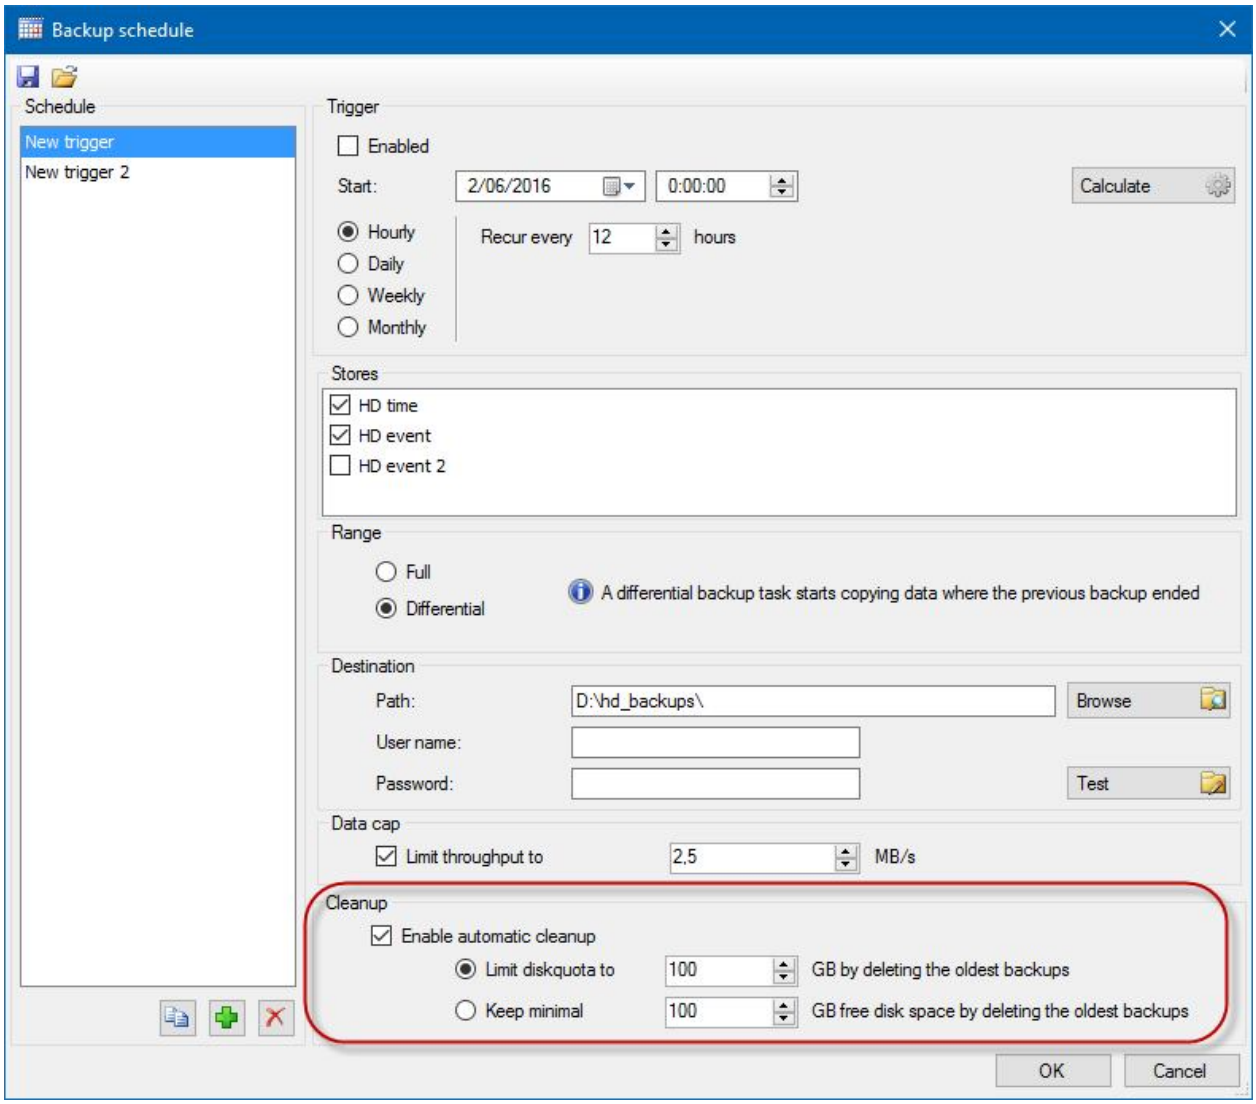

Cleanup will activate if the configured limit is reached while a new backup is being created. It will delete the oldest backup it can find in the destination directory. Sort order is determined by interpreting the subdirectory names. If the interpretation fails (user renamed?) the subdirectory is scanned and the modification date of the backup index file is used for sorting.

Cleanup does not check if backups in the destination directory were created by the task that is currently executing. It deletes any backup in this directory that is marked as oldest.

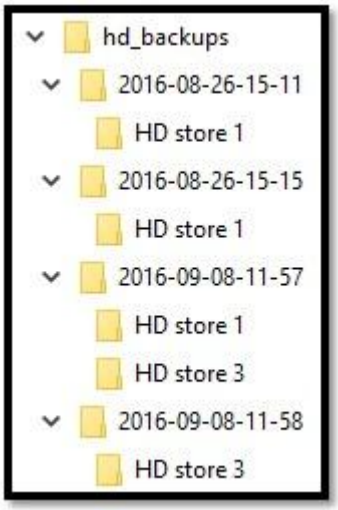

Cleanup info is added to the backup log (orange). In case the size of a single backup exceeds the cleanup limit, the backup task will fail right before the cleanup limit is reached (red).