



Upgrade Archive or Backup from P4 to P5

A major improvement in P5 is the **new indexing system**. The new, more powerful indexes are based on SQLite and can scale up to accommodate even for the largest of today's systems.

The new indexes come with a few changes that affect updates from P4 when using Backup or Archive. For Synchronize, no special treatment is required.

Backup indexes

After updating to P5, backup jobs will create new **empty** indexes leaving the old indexes intact and browsable in the *Restore* section of the GUI. New and empty indexes imply that the first run of any backup job after the update will always be a full backup. Files saved by previous backups will remain available in the P4 index structure to be restored until they fall out of the retention time. P4 indexes will not be used for further indexing by any P5 job. Furthermore, P4 indexes can only be browsed, whereas P5 can additionally also be searched.

A **major change** in the backup module affects how the backup job determines what to save for the next increment. In P4, the job referred strictly to the latest backup of the backup **plan** that was being executed. In P5, the job refers to the latest files saved by any previous job to the referenced **media-pool**. More specifically, each incremental backup updates the state of the saved files and folders as saved by the latest backup that used that specific pool of volumes.

This implies that:

- when two or more backup plans use the same media pool(s), the jobs will complement each other to actualize the state of that pool. Unlike in P4, incremental backups of different plans are **not independent** of each other. Such setups must be checked after the upgrade. If two plans are to independently run increments, each must have a separate pool. In case a backup plan exclusively uses certain pools and no other plan uses the same, there is no change in the behavior as compared to P4.
- when a backup plan uses alternating media pools in a single **incremental** event, a previous full backup in one pool **will no longer enforce** a full backup in the other pools when they get addressed.
- unlike in P4, it is now possible to create a full backup in one plan and an incremental backup in another plan that both use and refer to the same media pool.

In case backup must be deferred in order to adjust the plans to accommodate the changes, the corresponding plans can be disabled.



Files of a P5 client machine saved by executions of different backup plans are jointly represented in the client's index as **snapshots** of the saved parts of the file system, each taken at the time when the given plan was run. There are two fundamental differences which may come unexpected, if familiarity with P4 indexes exists:

- Snapshots are maintained only for the duration of the retain time as defined in the corresponding backup plan. Snapshot are removed from the index passed the **retain time**.
- Snapshots represent a set of files with a specific version. Some of those files in the same version may have been saved multiple times on different media. Selecting a certain snapshot for **restore** does not necessarily imply the selection of a set of media that will be used for restore. If multiple possibilities exist, the system may choose to restore the files from other media than that particular set that was used to write them to while creating that selected snapshot.

Archive indexes

Same as with the backup indexes, P4 archive indexes remain accessible in the *Restore* section of the GUI and can be used for the immediate restore of data. However, in order to be able to run new archive jobs that add new files to the archives, the indexes must be first converted to the new P5 database format. This conversion is not done automatically, but has to be triggered **manually**.

During conversion the index is read from the old location in the P4 configuration subfolder `conf/index/archive` and written to the new location at `config/index/archive`. In case the index subfolder was moved to some other location and referred to it by a link in the `conf` subfolder, it is recommended to create a similar link in the `config` subfolder.

For instance, if the existing link is

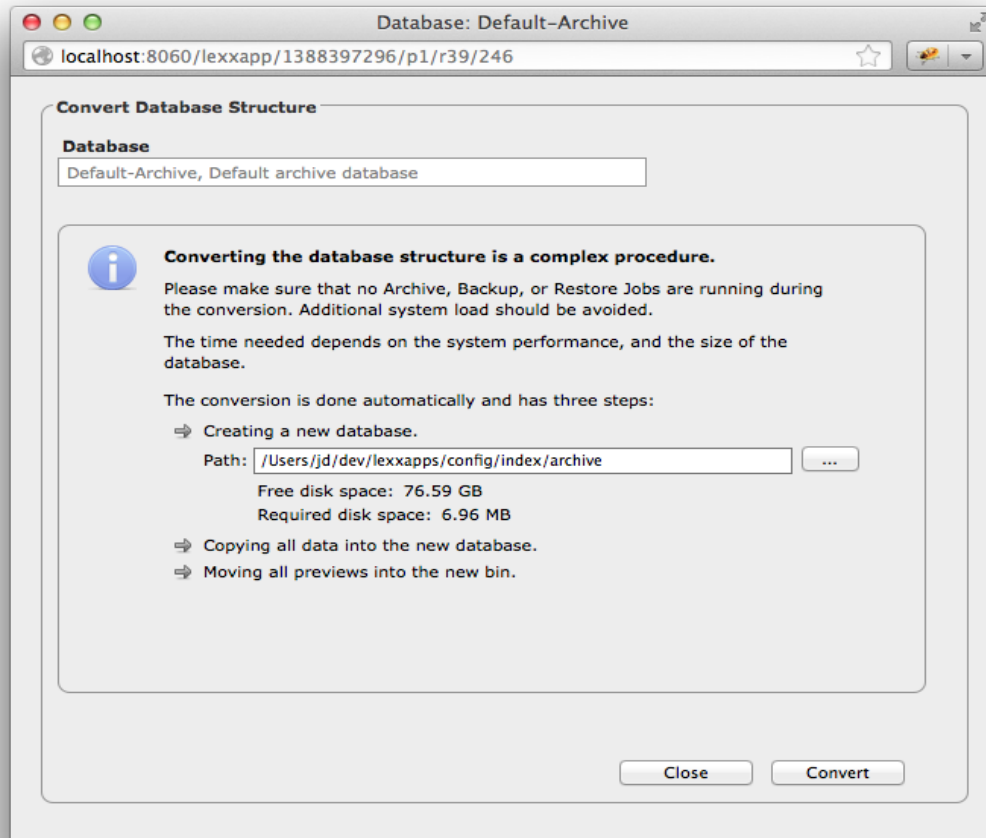
```
conf/index → /Volumes/bigdisk/aw/index
```

then a second link should be created onto the same disk but into another folder, for example:

```
config/index → Volumes/bigdisk/aw/p5index
```

Do not link both of the folders to the same target folder. It would cause the index conversion to fail.

In order to convert the index, navigate to the index databases as described in the manual, select the archive index and click the *Convert* button:



Depending on the index size, the conversion can take between a few minutes and multiple hours. Since the conversion creates high CPU and disk I/O load, please run only one index conversion at a time.

Also make sure that no other jobs access the same index during conversion. The progress of the conversion can be monitored in the P5 job monitor.



Upgrade Backup2Go to Version 5.1

When using the Backup2Go module, an update of the Backup2Go workstations is required, too. Please update the Backup2Go workstations locally, the remote update from the Backup2Go server cannot be used in this case due to the changed folder structure in version 5.1. This affects Backup2Go workstations of all previous versions.

P5 clients for Archive, Backup and Synchronize are not affected. These can be updated using the remote update procedure.