About Archiver connector scheduling

Libraesva Email Archiver use a sophisticated engine to handle connector scheduling. This engine has been introduced in v24.1.

This knowledge base will help you to understand how this engine works.

How the engine runs

The Archiver keeps the list of connectors and users to synchronize in a dedicated queue.

Based on the available resources and the users already in the synchronization phase, the Libraesva Email Archiver dynamically and intelligently decides which users to start synchronizing, taking them from the synchronization queue.

There exists only one queue, and this includes both users waiting for synchronization, and those waiting for cleaning. These two operations are handled separately.

For a user to be eligible for cleaning, its first sync must have terminated correctly.

Due to the quite demanding nature of connector cleaning resource-wise, only up to 4 cleaning processes can be running at any given moment.

How connector frequency is handled

The connector may be synchronized less frequently than configured depending on many factors:

- Number of emails
- Size of emails
- Number of connectors
- Number of users

When changing frequency, the engine’s settings are updated in a matter of seconds.

When selecting the “weekly” frequency, all the users relative to the modified connector will be pushed in queue on Friday at 8 PM CET.

What the engine guarantees

Libraesva Email Archiver connector scheduling guarantees that at all times:

- at least 1 user of 2 different connectors (if any) is in the sync phase
- at least 1 PEC user is in the sync phase
- at least 1 non PEC user is in the sync phase
- at least 1 user on its first mail ingestion ever is in the sync phase
- at least 1 user not on its first mail ingestion ever is in the sync phase