

Libraesva Archiver – Indexing Engine

The Libraesva Email Archiver uses an indexing engine to provide the best performance. This engine is used in almost all the Archiver operations, from the archiving process to the searches.

The index files are stored in the Index disk, which is completely dedicated to this purpose.

These files are fragmented into chunks, which optimal number is determined by the VM resources. The more the resources, the more the chunks number the Archiver can handle.

Over time, the index grows and new chunks are created with the archiving process.

Starting from Libraesva Email Archiver 21.12 (release notes), the Archiver daily checks the indexing engine status and try to optimize it by reducing the chunks number to the optimal value.

×NOTE: the optimization progress proceeds by reducing day by day the chunks number by a little value, so if you have just updated to the 21.12 release you should wait some days to reach the optimal indexing engine status.

During the index optimization process, the indexing engine is online and available, so there's no downtime at all and all the Archiver operations continue as normally.

When you delete emails from the Archiver, the emails are flagged as deleted by the indexing engine but still remain in the index. Purging these messages leads to more fragmentation on the index. The automatic indexing engine optimization will compact again the index to increase overall performance. You can find more information about email deletion at [this page](#).