



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

eServices Administrator's Guide

Sizing for the UCS Database and Index

5/8/2025

Sizing for the UCS Database and Index

The index that is produced by the UCS indexing service can be divided into two parts:

- The inverted index part that enables the full text search
- The data part that enables retrieval of the original data

If all of the original data is kept in the index, it can take as much space as the database itself; the default configuration does not compress it. The inverted index part is considerably smaller, but its size depends on the frequency of words. If the inverted index contains many unique values like IDs and timestamps it will be big. If it contains many common words such as those in the body of e-mails, it will be smaller.

The following table presents an example of sizing for a relatively large database.

Database Name	Database Size	Index Size	Duration
Standard responses	15 MB	6.5 MB	3 seconds
Contacts	5 GB	4.68 GB	2.5 hours
Interactions	20 GB	14 GB	2 hours

In this case the total size of the index files is about 19 GB.

During operations the index can grow to twice the size of the usable data. This is because index operations (such as internal reordering, purge of deleted documents, and concatenation) create new temporary files. To make these operations appear instantaneous, the system creates new files while it is still reading the old files. Then when it is finished creating the new files, it removes the old ones. Therefore, to be safe, free space on the disk hosting the indexes should be three times the size of the index.