

GENESYS

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

Searching the UCS Database

Contents

- 1 Searching the UCS Database
 - 1.1 Overview
 - 1.2 Collocated and Distributed Deployment
 - 1.3 More About Indexing
 - 1.4 Sizing
 - 1.5 Recommendations

Searching the UCS Database

Overview

UCS supports full-text search of contacts, interactions, and standard responses using the Platform SDK (PSDK) Contact API.

To enable full text search, the stored data must be prepared, an operation known as indexing. Indexing converts, abbreviates, and sorts data to allow fast searching in large data sets.

There are various ways of preparing the data for the index; one way is to divide text data into words. Another way of preparing text data is *stemming*, or storing only the stem of a word rather than all its forms (for example, *run, running, runs,* and *ran* all stored as simply *run*). UCS uses *analyzers* to perform word division and stemming; you can configure which analyzer applies to which database field. In addition,

- Examples of the syntax used for searching are available.
- Many attributes are not directly searchable because of how they are stored; however you can make such attributes searchable.

Collocated and Distributed Deployment

This refers to the location of the indexes of the primary and backup UCS instances. The two types of deployment require different settings of the shared option (index section).

Collocated

In this type both primary and backup instances use the same storage path for the index. This is recommended in the following cases:

- For smaller deployments.
- Either both UCS instances are on the same machine or two instances on separate machines share the index files over a network.
- Read/write performance is not expected to be an issue.
- Network latency between UCS hosts and the shared file system is negligible

With collocated deployment, sharing must be enabled (shared = true) so that only the primary instance updates the index. It can also be enabled when sharing the index over a network.

Distributed

In this type each instance his its own set of index files. Note that the backup instance periodically

reads the database to replicate contact, interaction, and standard response updates to its local index. Here is a list of cases in which distributed deployment is recommended:

- For larger deployments.
- When the two UCS instances are on separate networks.
- For optimum index/search performance.
- In disaster recovery scenarios.

In this type of deployment, both UCS applications must have sharing disabled (shared = false).

More About Indexing

It is important to note that the UCS indexing service enables searching for text data, but does not guarantee its retrieval.

When an object is found it must be retrieved using standard PSDK methods. These methods are documented in the API References of the Platform SDK.

When designing an application, data coming from search results should be considered as cached data, like the preview that can be seen in web search engines. This data is optimized for search and cannot be used to update UCS database objects. For example, the StructuredText of an interaction is rendered from HTML to text. All formatting is lost if this data is re-inserted into the interaction.

Documents stored in the index will at least have the following retrievable fields:

- id—Unique identifier in the UCS database. This is what enables retrieval of the object from UCS.
- IndexationDate—The date this index was created. The format is UTC (yyyy-MM-dd'T'HH:mm:ss.SSS'Z').
- DocumentType—Type of the document, selected from the following:
 - EmailIn
 - EmailOut
 - PhoneCall
 - CallBack
 - CoBrowse
 - Chat
 - Interaction (Open Media)
 - Contact
 - StandardResponse

Example:

id = 00001a4EM4TD007K

IndexationDate = 2013-10-02T13:54:54.318Z

Type = chat

In the index, even dates and numbers are text. This format enables alphabetic sorting, so that, for example, 21 December 2008 comes before 1 June 2009.

Default configuration enables the indexing of Contacts, Interactions and Standard Responses.

Sizing

Read about sizing for the database and its index.

Recommendations

- Distributed is recommended (index/shared=false).
- Store the index on a local dedicated SSD (Solid State Drive).
- If needed, run index rebuild while UCS is in backup mode, not primary mode. It is always better to rebuild off production hours.
- It is possible to copy a healthy index directory from one UCS to another while both instances are stopped.
- UCS Memory (XMX setting) may need to be increased if more than 3 months of data is kept in main database.